IBM Systems Director 6.3 Best Practices

Provides additional guidance beyond standard instructions

Covers basic and advanced features

Learn from the experts in the field

Rufus Credle
Srikanth Aithal
Nicolas Bour
Stephane Bourdeaud
Tomi Mannikainen
Olaf Menke
Juan P Fernandez Sabate

ibm.com/redbooks

Redbooks

International Technical Support Organization

IBM Systems Director 6.3 Best Practices

November 2013

SG24-8141-00

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

First Edition (November 2013)

This edition applies to IBM Systems Director V6.3 Express, Standard, and Enterprise Edition.

Contents

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 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:

Active Memory™
AIX 5L™
AIX®
BladeCenter®
DB2®
Domino®
DS4000®
DS6000™
DS8000®
Electronic Service Agent™
EnergyScale™
Global Technology Services®
i5/OS™
IBM Flex System™
IBM Flex System Manager™
IBM SmartCloud®
IBM Systems Director Active Energy Manager™
IBM®
Lotus®
Power Systems™
POWER6®
POWER7®
PowerPC®
PowerSC™
PowerVM®
POWER®
PureFlex™
PureSystems™
Redbooks®
Redpapers™
Redbooks (logo) ®
Storwize®
System i®
System p®
System Storage®
System x®
System z®
Systems Director VMControl™
Tivoli®
WebSphere®
XIV®
z/VM®
zEnterprise®

The following terms are trademarks of other companies:

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

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

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

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

This IBM® Redbooks® publication describes the positioning of the IBM Systems Director in the complete management range. It also compares the IBM Systems Director with the IBM Flex Systems Manager (FSM) and describes the environments for which each tool is best suited.

This publication helps you plan, install, tailor, and configure the IBM Systems Director on different platforms. It contains information about required system resources and which network ports are used. It shows how to use the Workload Estimator to select the appropriate hardware for IBM Systems Director server and provides information about the IBM Systems Director Editions.

Best practices are covered for the basic management tasks that are available in IBM Systems Director, including how to perform discovery; how to collect inventory on discovered resources; how to deploy agent, driver, and firmware updates; how to manage hardware events; and other miscellaneous tasks.

An overview of best practices is provided for using IBM Systems Director VMControl™. Systems Director VMControl is a cross-platform product that assists you in rapidly deploying virtual appliances to create virtual servers that are configured with the operating system and software applications that you want. It also enables you to group resources into system pools, which enable you to centrally manage and control the different workloads in your environment.

The following plug-in offerings are described:

- Energy monitoring and management features offered by IBM Systems Director Active Energy Manager™ along with the best practice, which needs to be followed in using the IBM Systems Director Active Energy Manager.
- The IBM AIX® Profile Manager is a tool that can help implement and monitor the security of all AIX servers in a production environment but also implement and monitor the system compliance of those AIX servers.
- Best practices and the most important questions to ask before creating Workload Partition Manager (WPAR) and WPAR Manager infrastructure. In addition, how you can manage and relocate WPARs using WPAR Manager graphical interface and the command-line interface.
- Network Control basic functionalities and how to plan for Network Control deployments and also a number of common scenarios with best practices.
- The IBM Systems Director Service and Support Manager describes how to set up and how to handle serviceable events.
- Best practices for the Storage Monitoring and Management capabilities offered by IBM Systems Director server.

This book is for IBM IT specialists and IT architects, IBM Business Partners, and clients, who are utilizing or considering implementing IBM Systems Director.

Authors

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

Rufus Credle is a Certified Consulting IT Specialist at the ITSO, Raleigh Center. In his role as Project Leader, he conducts residencies and develops IBM Redbooks and IBM Redpapers™ publications. Subjects include network operating systems, enterprise resource planning (ERP) solutions, voice technology, high availability, clustering solutions, web application servers, pervasive computing, IBM and OEM e-business applications, IBM WebSphere® Commerce, IBM industry technology, IBM System x®, and IBM BladeCenter®. Rufus' various positions during his IBM career include assignments in administration and asset management, systems engineering, sales and marketing, and IT services. He has a BS degree in Business Management from Saint Augustine's College. Rufus has been employed at IBM for 33 years.

Follow Rufus on Twitter: http://twitter.com/rcredle1906

Join Rufus' network on LinkedIn:

http://www.linkedin.com/pub/rufus-p-credle-jr/1/b/926

Srikanth Aithal is a Senior Systems Management Test Engineer in the IBM Systems and Technology Group. He is based in Bengaluru, India. He holds a degree in Computer Science from the Visvesvaraya Technological University in Karnataka. Srikanth has worked at IBM since 2007 and is currently working for Cloud Systems Software group. He focuses primarily on virtualization management on IBM Power servers and IBM Systems management stack. He also has expertise on IBM System x and IBM BladeCenter systems and has delivered many client trainings around IBM BladeCenter and IBM systems management.

Join Srikanth's network on LinkedIn:

http://in.linkedin.com/pub/srikanth-aithal/13/922/6b8

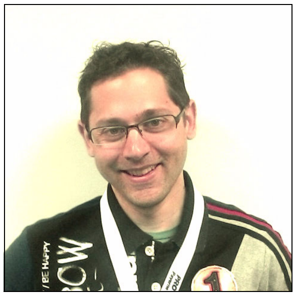

Nicolas Bour is a UNIX Systems Administrator in Switzerland. He has 11 years of experience in the AIX field. His areas of expertise include IBM AIX, IBM Power Systems™, IBM Systems Director, Security, and Linux Red Hat. He is an IBM Certified Advanced System Expert (Power5) and Red Hat Certified Systems Administrator.

Join Nicolas' network on LinkedIn:

http://ch.linkedin.com/pub/nicolas-bour/11/378/729

Stephane Bourdeaud is a Certified Infrastructure Architect for the IBM Global Technology Services® Delivery organization. He is based in Aubière, France. Stephane has over 15 years of experience in the IT industry and focuses primarily on virtualization technologies on Intel platforms. He works for the Global Technology Services Delivery Technology and Engineering organization, where he is a member of the core team for the Virtualization and Distributed Server Management Specialty Area. He is a VMware Certified Professional and teaches a class on virtualization at the Institut Supérieur d'Informatique de Modélisation et de leurs Applications (ISIMA), an engineering school in Clermont Ferrand, France. Stephane is also a contributor to the IBM Expert Integrated System Blog (http://expertintegratedsystemsblog.com).

Follow Stephane on Twitter: https://twitter.com/BourdeaudS

Join Stephane's network on LinkedIn:

http://fr.linkedin.com/pub/stephane-bourdeaud/57/4b8/168

Tomi Mannikainen is an IBM PureSystems™ Client Technical Specialist for IBM Systems and Technology Group. He is based in Helsinki, Finland. He has 17 years of experience in the IT industry with a focus on service provider environments, networking, virtualization, and cloud provisioning technologies. Tomi is a contributor in the PureHeat community and he actively participates in beta and iiCap programs. He is a Flex Systems Manager black belt. He holds a Master's degree in Business Administration.

Join Tomi's network on LinkedIn:

http://fi.linkedin.com/pub/tomi-mannikainen/5/a20/449

Olaf Menke is a Consultant and Subject Matter Expert for systems management. He has worked in the IBM Technical Support Services (TSS) Software Service in IBM Germany for the past two years. Prior to this position, he was a Systems Engineer and IBM System x and BladeCenter specialist in the System x Pre-Sales team in Germany. He has over 16 years of experience in support of computer systems and software. He holds a degree in Information Technology from the Technische Universitaet in Dresden. His areas of expertise include System x, BladeCenter, IBM PureFlex™, Systems Director, and management hardware. He is an IBM Certified Specialist for PureFlex and IBM Certified Expert for IBM System x and BladeCenter.

Join Olaf’s network on LinkedIn:

`http://www.linkedin.com/pub/olaf-menke/5/48a/9aa`

Juan P Fernandez Sabate is a Delivery Technology and Engineering Architect for the IBM Delivery Center that is located in Argentina. He holds a degree in Computer Engineering from the National University of Tucuman in Argentina. He has over eight years of experience in IT and has been working for IBM for last six years. His areas of expertise include IBM System i®, IBM Tivoli® products, and IBM Systems Director. He focuses primarily on providing Automation and Monitoring solutions.

Join Juan Pablo’s network on LinkedIn:

`http://www.linkedin.com/pub/juan-pablo-fernandez-sabate/4/527/b36`

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

Tamikia Lee
IBM International Technical Support Organization, Raleigh Center

Venkata s Amulothu (Kishore), Larry Bailey, Elodie Bourdeaud, Stephen Boyte, Deanna Brown, Kimberly K Button, Steven M Hancock, David Kasberg, Jay Kidambi, Uday K Kumbhari, Achuth Maniyedath, John McGinn, Ranjith Rajagopalan Nair, Brian Noordyke, Christian Ovsenik, Pushkar Patil, Lee Prissel, Rahul B Rege, Hector L. Sanchez, Chris Seyer, Jin Ming Shang, Barb Smith, and Liyan Zhang.
IBM

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:

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:

  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

1

Positioning IBM Systems Director

This chapter describes the positioning of the IBM Systems Director in the complete management range. It also compares the IBM Systems Director with the IBM Flex Systems Manager (FSM) and describes the environments for which each tool is best suited.

This chapter contains the following topics:

- 1.1, “Description of IBM Systems Director” on page 2
- 1.2, “Functionality of IBM Systems Director” on page 2
- 1.3, “Positioning of IBM Systems Director” on page 3
- 1.4, “IBM Systems Director and IBM Flex Systems Manager” on page 4
- 1.5, “Best practices” on page 5

1.1 Description of IBM Systems Director

IBM Systems Director is a platform management tool with enhanced functionality for managing both physical and virtual resources in a heterogeneous environment. IBM Systems Director uses industry standards to support multiple operating systems and virtualization technologies across IBM platforms and non-IBM x86 platforms.

IBM Systems Director provides tools for managing, monitoring, and handling your heterogeneous environment including System x, Power Systems, IBM System z®, Storage, and networking components as physical or virtual systems.

IBM Systems Director provides administrators with a single vantage point through its web interface or command line tool (SMCLI) and reduces IT management complexity and costs.

1.2 Functionality of IBM Systems Director

IBM Systems Director consists of a basic IBM Systems Director server. It includes all necessary functions to manage the different platforms and plug-ins or Advanced Managers, which enhance the functionality of the IBM Systems Director server.

The following basic functions are provided at no charge on IBM hardware with IBM Systems Director:

- *Discovery of systems* that brings the systems into the IBM Systems Director server and allows access to the systems.
- *Inventory of systems* for hardware and software inventory that is stored in the IBM Systems Director Database.
- *Event management* that handles incoming events, Event Action Plans, and monitoring with thresholds.
- *Update management* that downloads the latest update information; runs a compliance check; updates the firmware, drivers, and IBM Systems Director server and agents, and rollout of agents.
- *Platform management* that includes specific management tasks for the different IBM platforms such as Power Systems, System z, System x, and BladeCenter.

Detailed information about these basic functions can be found in this IBM Redbooks publication in Chapter 4, “Basic management tasks” on page 77.

To expand the functionality, there are plug-ins or the Advanced Manager available. These are:

- *VMControl* is a tool for handling, monitoring, and management of virtual environments across brands (VMware vSphere, Kernel-based Virtual Machine (KVM), IBM PowerVM®). See Chapter 5, “VMControl” on page 247 for detailed information.
- *Active Energy Manager* is a tool that includes monitoring of power usage (at no charge) and management of power usage (fee based) that includes power safe and power capping functionality for servers and infrastructure components. See Chapter 6, “Active Energy Manager” on page 345 for detailed information.
- *AIX Profile Manager* is used to implement and monitor security standards (IBM LLS, MLS, and HLS models; and PCI, SOX, Dod, Hippa standard models). It is also used to implement and control system settings on all AIX and VIO servers against a client predefined standard. See Chapter 7, “AIX Profile Manager” on page 375 for detailed information.

- *AIX PowerVM Workload Partition Manager* is used to manage all the workload partition (WPAR) infrastructure, which includes all kinds of WPARs and all global environment. See Chapter 8, “Workload Partition Manager” on page 403 for detailed information.
- *Network Control* is used for management of top of rack (ToR) switches and switches in BladeCenter and Flex systems. See Chapter 9, “Network Control” on page 419 for detailed information.
- *Service and Support Manager* is a *call home* tool to automatically create hardware defect calls to IBM. See Chapter 10, “Service and Support Manager” on page 433 for detailed information.
- *Storage Control* is a tool for the management of different storage devices such as the IBM Storwize® V7000, SAN Volume Controller (SVC), DS3/4/5000, IBM DS8000®, and third-party devices from EMC and Hitachi. See Chapter 11, “Storage Management solutions and Storage Control” on page 451 for detailed information.

1.3 Positioning of IBM Systems Director

IBM Systems Director is a platform management tool. The Upward Integration modules from IBM Systems Director provide an integration into enterprise management tools, such as Tivoli, HP OpenView, CA Unicenter, and the Microsoft Systems Management Server family.

As shown in Figure 1-1, IBM Systems Director is positioned below the enterprise tools. Additionally, IBM Systems Director can manage all of the platforms, service processors, and platform-specific tools in one common program and console. This helps administrators focus on their work without the need to connect to different consoles.

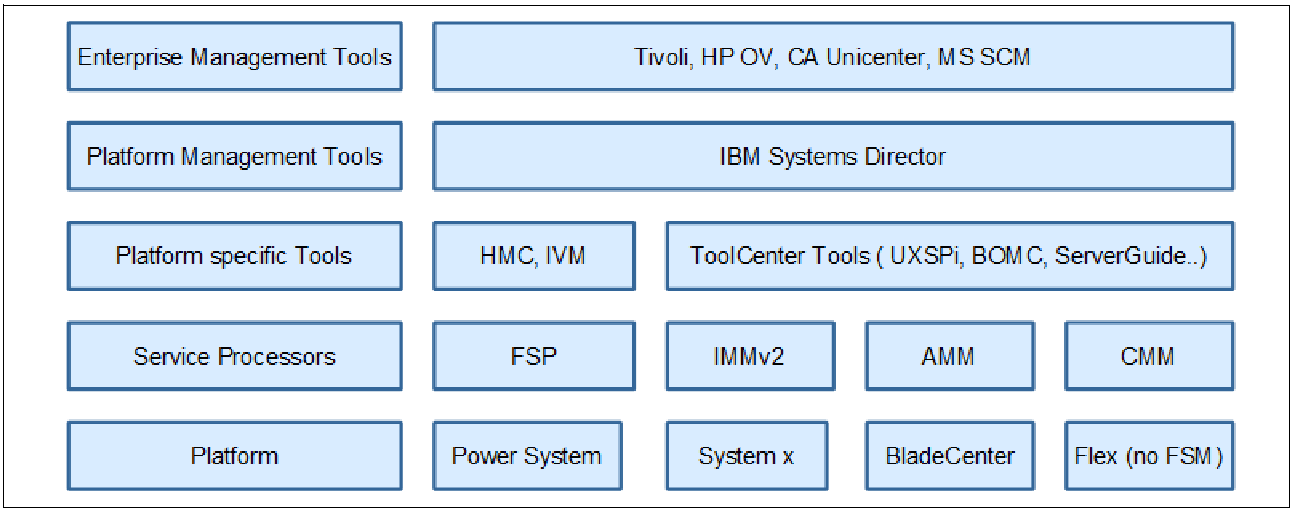

Figure 1-1 Positioning IBM Systems Director

As a platform management tool, IBM Systems Director supports the management of IBM and non-IBM hardware and drives common tasks through the following platform-specific manager:

- IBM Power Systems management
  - HMC, IVM, and VIOS appliances
  - Power Systems server, Power Blade server
  - AIX, IBM i, and Linux on Power operating systems

- IBM BladeCenter chassis management
  - IBM BladeCenter chassis components such as switches and blades
  - VMware virtual server inside a BladeCenter
- IBM System x Management
  - System x and BladeCenter server systems
  - Windows and Linux operating systems
- System z management
  - IBM z/VM® hypervisor
  - Linux on z installed on z/VM virtual server and also on partitions without z/VM
- IBM System Storage® Management
  - Integrated RAID Adapter (such as LSI) and ServeRaid Adapter family
  - Network Storage, such as IBM DS3000, IBM DS4000®, IBM DS5000, IBM DS6000™
  - Storage switches (SAS, Fibre Channel, and Converged Network (CN) switches) such as IBM BladeCenter SAS, IBM Systems Networking CN switches, and switches from Brocade, QLogic, Nortel, and Cisco.
- IBM Flex System™ hardware
  - No FSM installed
  - x86 compute nodes only

1.4 IBM Systems Director and IBM Flex Systems Manager

IBM Systems Director and Flex Systems Manager (FSM) are based on the same basic code modules (the FSM reuses the IBM Systems Director code and some of its advanced managers) but they are enhanced with specific elements for each of them. They are both platform management tools. However, they have different specifications and support different types of hardware:

- IBM Systems Director is for *general, universal* use.

  This means that IBM Systems Director can manage rack, tower, and blade servers across brands. IBM Systems Director can also manage Flex Systems, but only when there is no FSM installed in the Flex Chassis and only for x86 compute nodes. IBM Systems Director supports third-party hardware. IBM Systems Director must be manually installed. There is no pre-installation option.
- Flex Systems Manager (FSM) is for *specialized* use.

  FSM is optimized for managing the Flex and PureFlex systems. FSM brings a new user interface (UI), the Flex Explorer. It contains additional functionality through the graphical view to the components in the Flex Chassis such as an operating system provisioning engine and configuration patterns for Flex System components. The Flex Systems Manager is a completely pre-installed appliance on a specialized compute node for the Flex and PureFlex systems. The FSM brings a deeper integration with hypervisors and system patterns for the configuration of the environment on the Flex Chassis.

Figure 1-2 on page 5 shows a comparison of the two tools.

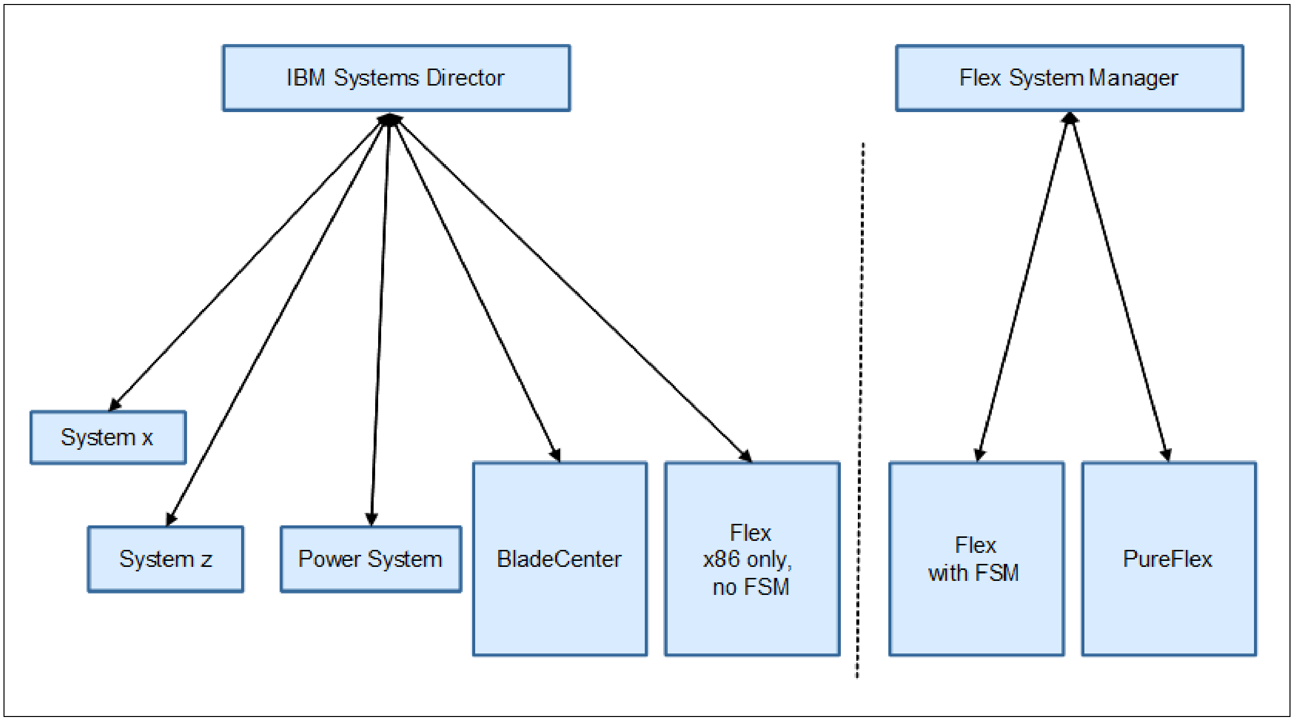

Figure 1-2 IBM Systems Director versus Flex Systems Manager

> **Note:** If you have both Flex System and System x or BladeCenter hardware in your infrastructure, you can have an FSM and IBM Systems Director server running side by side. You cannot manage Flex System components from both the FSM and IBM Systems Director. The components can be managed only on the FSM. There is currently no available solution to manage both IBM Systems Director and FSM from the same console.

1.5 Best practices

The following list describes the best practices for positioning and usage of IBM Systems Director:

- Use IBM Systems Director for all of your cross brand systems, including Flex Enterprise chassis without FSM installed.
- Prefer Flex Systems Manager for managing all systems inside a Flex or PureFlex system because it has additional functionalities.
- Use IBM Systems Director *only* as a platform management tool.
- Start using IBM Systems Director with the basic tasks. When you get experienced with the basic tasks and the handling of IBM Systems Director, you can expand the functionality by installing the plug-ins or Advanced Managers.
- Plan the installation and usage of IBM Systems Director carefully, using the requirements and architectural decisions in the following chapters.

2

Planning

This chapter describes the planning requirements for IBM Systems Director. It contains information about required system resources and which network ports are used. It shows how to use the IBM Systems Workload Estimator to select the right hardware for IBM Systems Director server and provides information about the IBM Systems Director Editions.

This chapter also addresses some of the most common architectural decisions you will be faced with when planning for an IBM Systems Director deployment.

This chapter contains the following topics:

2.1 System resources

Table 2-1 is a guide for the installation of IBM Systems Director. Use Table 2-1 to estimate the system resources to allocate to the logical partition (LPAR), virtual, or physical machine on which the IBM Systems Director server will be installed.

Table 2-1 IBM Systems Director hardware requirements for medium to large environments

Operating system	Processor	Memory	Disk storage
AIX/Linux	Four processors, POWER5, IBM POWER6®, or POWER7®: ► Entitlement = 4 ► Uncapped ► Virtual processor = 8 ► Weight = default	16 GB	30 GB
Microsoft Windows	Four processor cores (two dual-core processors or one quad-core processor)	16 GB	30 GB plus space for Update Manager files
Linux on x86	Four processor cores (two dual-core processors or one quad-core processor)	16 GB	30 GB plus space for Update Manager files
Guest OS on virtualized environment on x86	Four vCPUs	16 GB	30 GB plus space for Update Manager files
► Recommendations are based on 64-bit Java virtual machine (JVM). ► Recommendations are based on POWER6, Intel Xeon processor numbers for x86. ► I/O requirements: SCSI/serial-attached SCSI (SAS) adapters and multiple 10 K - 15 K rpm disks. ► Suggested: Two processors minimum and 8-GB memory minimum. ► Disk storage depends on advanced managers, the used database (local/remote), and the number of systems models for the Update Manager repository size. ► Advanced manager might require more memory (for performance).			

Table 2-1 references a medium-to-large environment. In a small environment, the amount for memory and CPU usage can be smaller. But for performance reasons, you should also use the size that is described for small environments

Installation sizes are summarized in Table 2-2.

Table 2-2 Definitions for small, medium, and large installations

Configuration size	Managed systems
Small	< 500 managed systems
Medium	500 > managed systems < 1000
Large	> 1000 managed systems

The installation media for the IBM Systems Director server includes an integrated IBM DB2® database. Use the integrated DB2 database as the default database to simplify the installation and reduce the need for a database administrator. If you plan to use Storage Control, a DB/2 database (local or remote) is required.

The components in Figure 2-1 show points to consider when you design and implement the IBM Systems Director server from a loading viewpoint.

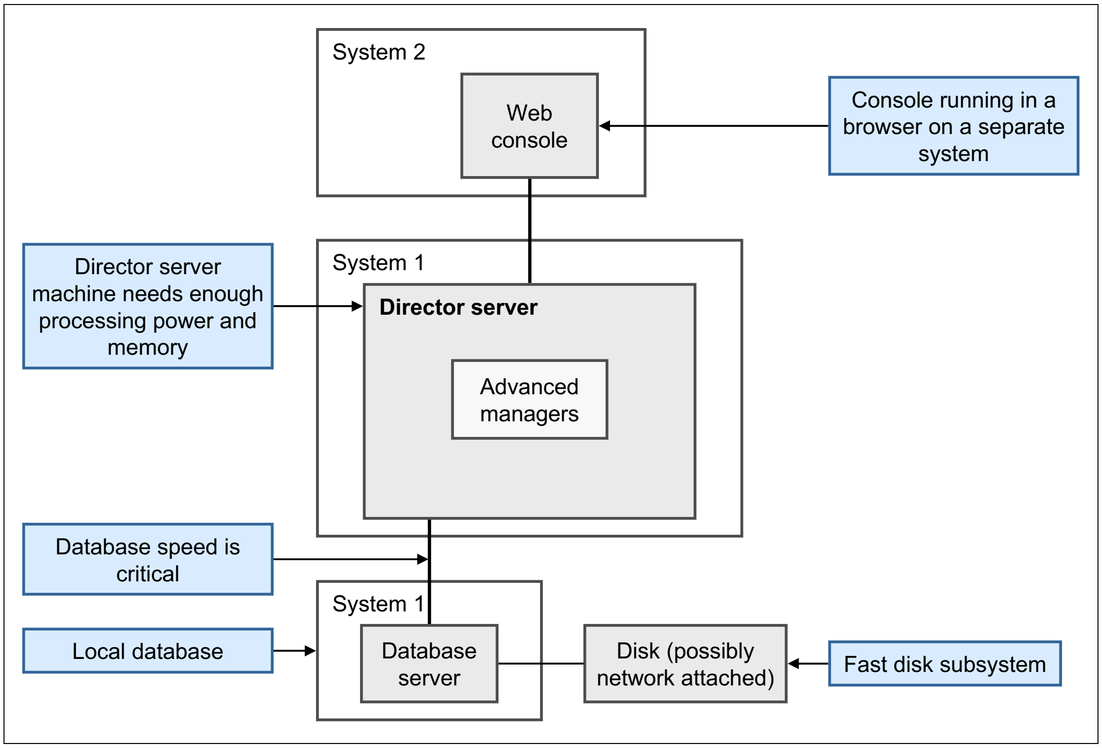
System 2
Web console
Console running in a browser on a separate system
System 1
Director server
Director server machine needs enough processing power and memory
Advanced managers
Database speed is critical
System 1
Local database
Database server
Disk (possibly network attached)
Fast disk subsystem

Figure 2-1 Systems Director components

2.2 IBM Systems Workload Estimator

IBM Systems Workload Estimator for Systems Director 6.3 is a web-based tool that can size hardware for systems that run the Systems Director server. The tool is presented in a Q&A format and requests user input.

The tool provides information about physical systems but it can also be used to get figures for a virtual environment (for example, running IBM Systems Director in a VM). In this case, use the information to get the physical system and transform this into a virtual environment. For example, if the physical server should have four cores you should use four vCPUs instead for a virtual server running the IBM Systems Director server.

Use the following steps to launch and use the IBM Systems Workload Estimator.

1. Launch the IBM Systems Workload Estimator for Systems Director 6.3 from this URL:

   http://pic.dhe.ibm.com/infocenter/director/pubs/index.jsp?topic=%2Fcom.ibm.director.plan.helps.doc%2Fwle.html

2. When the web page displays, click **WLE for IBM Systems Director 6.3** to start Workload Estimator.

3. You might be asked to enter user demographic information, as shown in Figure 2-2. Enter the appropriate information about this page and then click **Continue**.

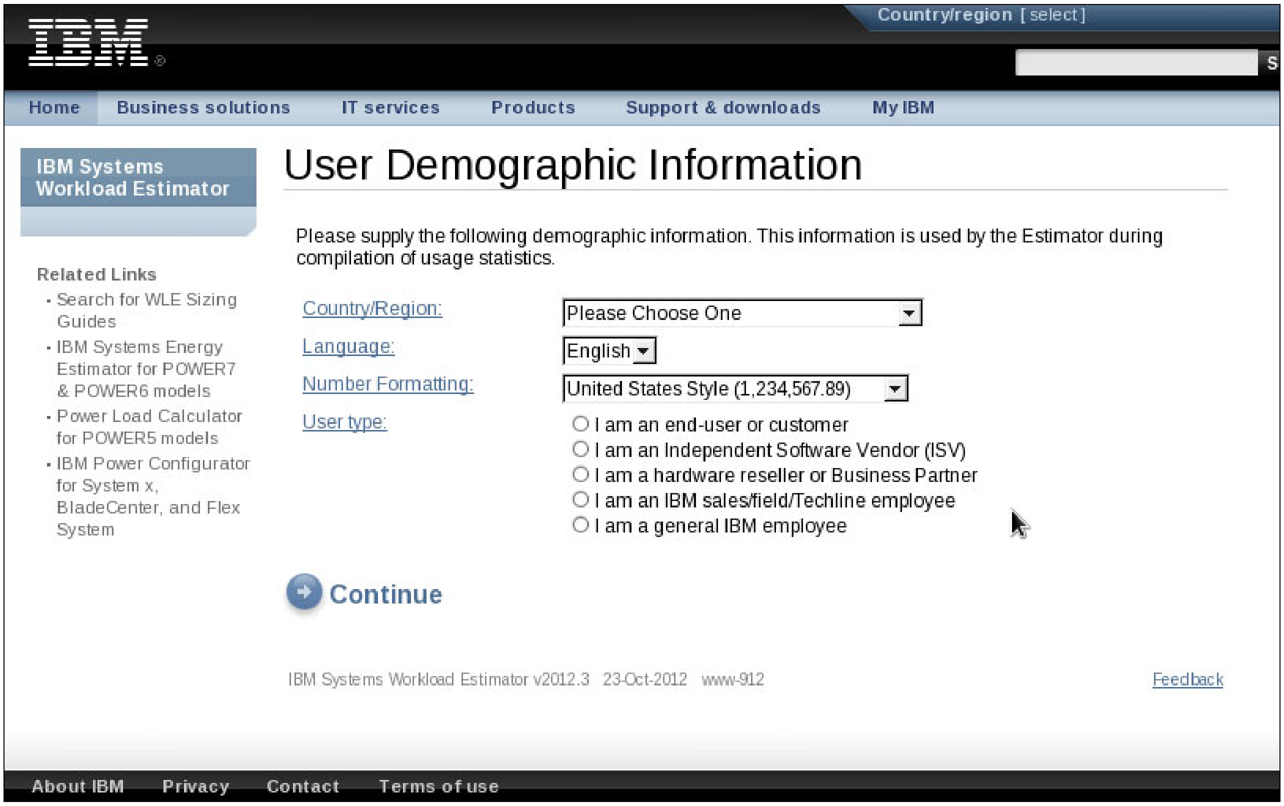
IBM
Country/region [select]
Home
Business solutions
IT services
Products
Support & downloads
My IBM
IBM Systems Workload Estimator
Related Links
Search for WLE Sizing Guides
IBM Systems Energy Estimator for POWER7 & POWER6 models
Power Load Calculator for POWER5 models
IBM Power Configurator for System x, BladeCenter, and Flex System
Continue
Feedback
About IBM
Privacy
Contact
Terms of use

User Demographic Information

Please supply the following demographic information. This information is used by the Estimator during compilation of usage statistics.

Country/Region: Please Choose One

Language: English

Number Formatting: United States Style (1,234,567.89)

User type:
○ I am an end-user or customer
○ I am an Independent Software Vendor (ISV)
○ I am a hardware reseller or Business Partner
○ I am an IBM sales/field/Techline employee
○ I am a general IBM employee

IBM Systems Workload Estimator v2012.3 23-Oct-2012 www-912

Figure 2-2 User Demographic Information window

4. If you are not asked for demographic information, the window shown in Figure 2-3 is displayed, which shows the platform, database, and plug-in type. The Active Energy Manager plug-in is listed due to the additional I/O, network traffic, and processor utilization activity that results from collecting data from your energy consumption. Click **Continue**.

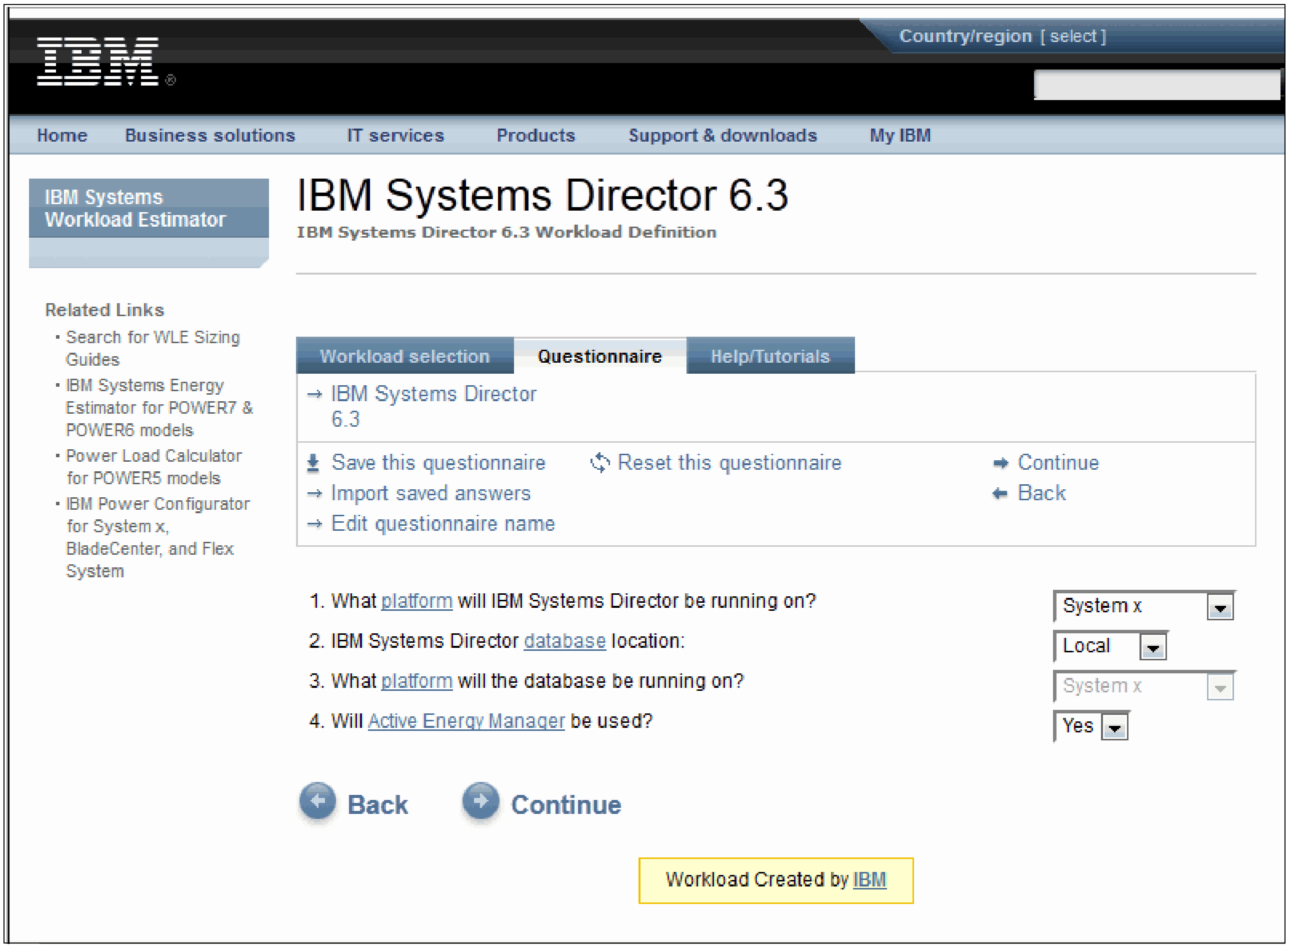
IBM
Country/region [select]
Home
Business solutions
IT services
Products
Support & downloads
My IBM
IBM Systems Workload Estimator
Related Links
Search for WLE Sizing Guides
IBM Systems Energy Estimator for POWER7 & POWER6 models
Power Load Calculator for POWER5 models
IBM Power Configurator for System x, BladeCenter, and Flex System
Workload selection
Questionnaire
Help/Tutorials
→ IBM Systems Director 6.3
Save this questionnaire
Reset this questionnaire
Continue
→ Import saved answers
Back
→ Edit questionnaire name
Back
Continue
Workload Created by IBM

IBM Systems Director 6.3

IBM Systems Director 6.3 Workload Definition

1. What platform will IBM Systems Director be running on? System x
2. IBM Systems Director database location: Local
3. What platform will the database be running on? System x
4. Will Active Energy Manager be used? Yes

Figure 2-3 Platform choice

5. Figure 2-4 requests information about the operating system on which you chose to install the Systems Director server, the estimated number of physical systems and operating systems with the managed environment, and the number of concurrent console users. Enter the appropriate information and click **Continue**.

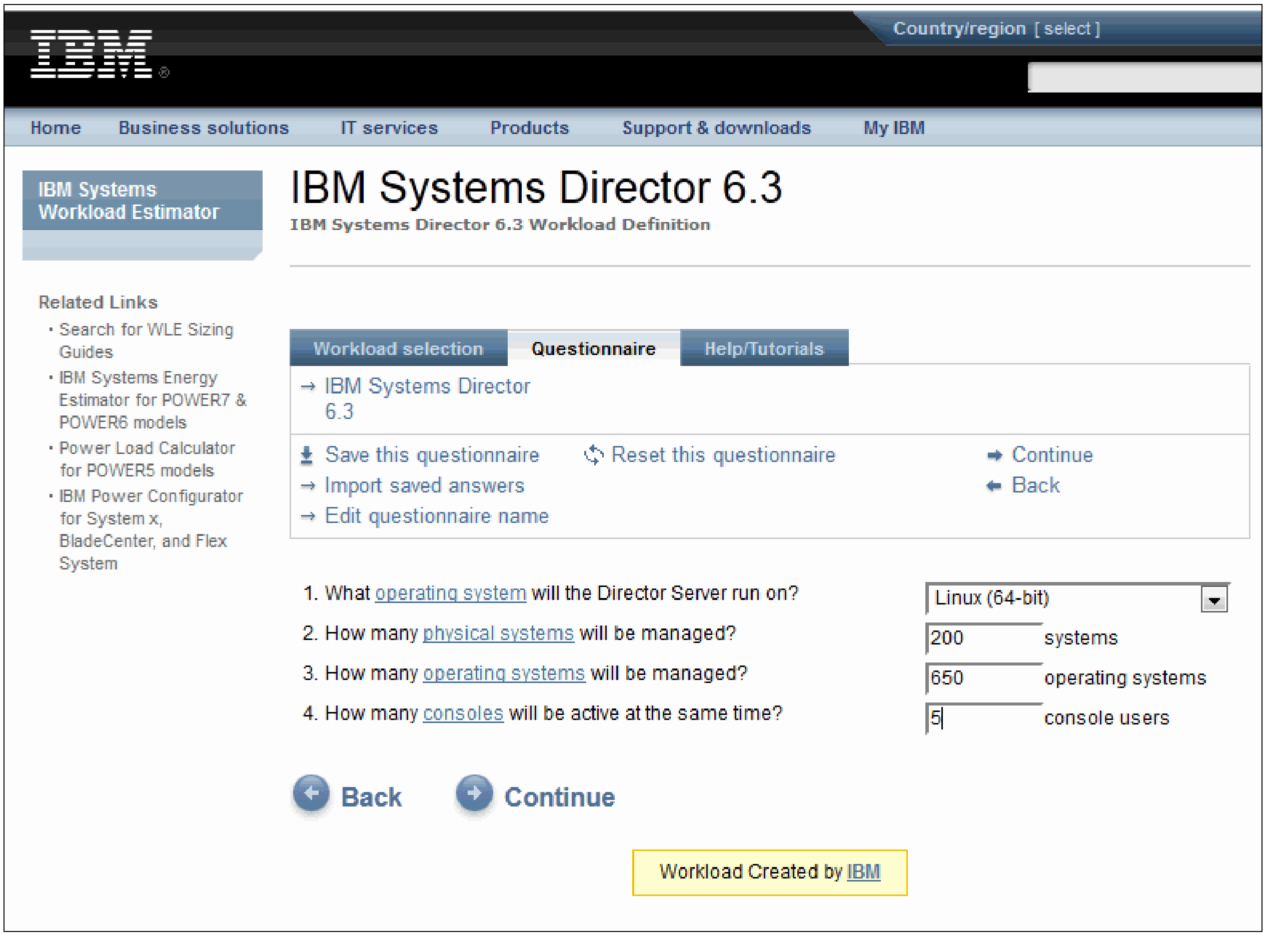

1. What operating system will the Director Server run on? Linux (64-bit)
2. How many physical systems will be managed? 200 systems
3. How many operating systems will be managed? 650 operating systems
4. How many consoles will be active at the same time? 5 console users

Figure 2-4 Operating system and physical system information usage

6. The estimator provides guidance for the number of required disk drives for internal or external storage (Figure 2-5). Enter the appropriate information and click **Continue**.

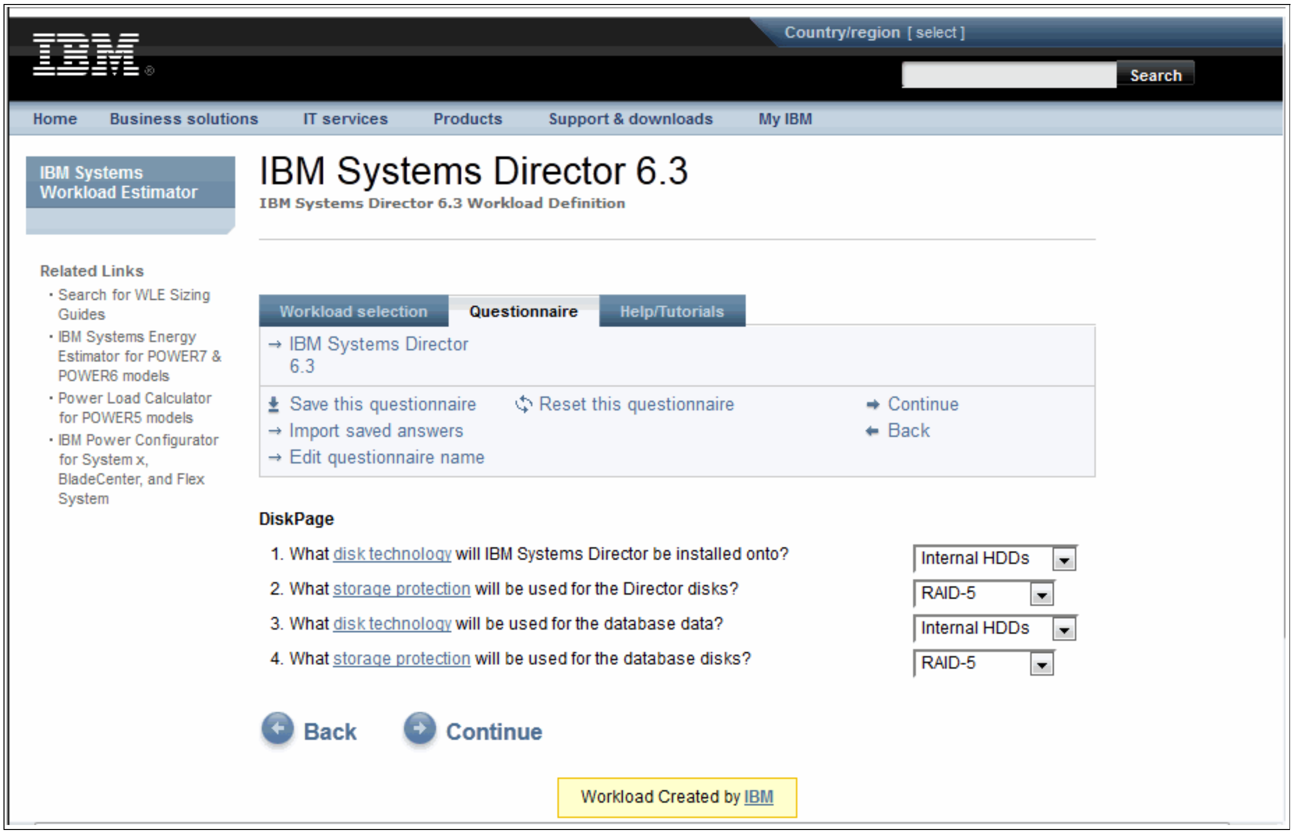
Country/region [select]
Search
Home
Business solutions
IT services
Products
Support & downloads
My IBM
IBM Systems Workload Estimator
IBM Systems Director 6.3
IBM Systems Director 6.3 Workload Definition
Related Links
Search for WLE Sizing Guides
IBM Systems Energy Estimator for POWER7 & POWER6 models
Power Load Calculator for POWER5 models
IBM Power Configurator for System x, BladeCenter, and Flex System
Workload selection
Questionnaire
Help/Tutorials
→ IBM Systems Director 6.3
Save this questionnaire
Reset this questionnaire
Continue
→ Import saved answers
Back
→ Edit questionnaire name
Back
Continue
Workload Created by IBM

DiskPage

1. What disk technology will IBM Systems Director be installed onto? Internal HDDs
2. What storage protection will be used for the Director disks? RAID-5
3. What disk technology will be used for the database data? Internal HDDs
4. What storage protection will be used for the database disks? RAID-5

Figure 2-5 Disk and storage information

After entering all of your system information, the Workload Estimator provides two *estimated* outputs. One estimate is for an immediate solution and the other estimate is for a growth solution (Figure 2-6).

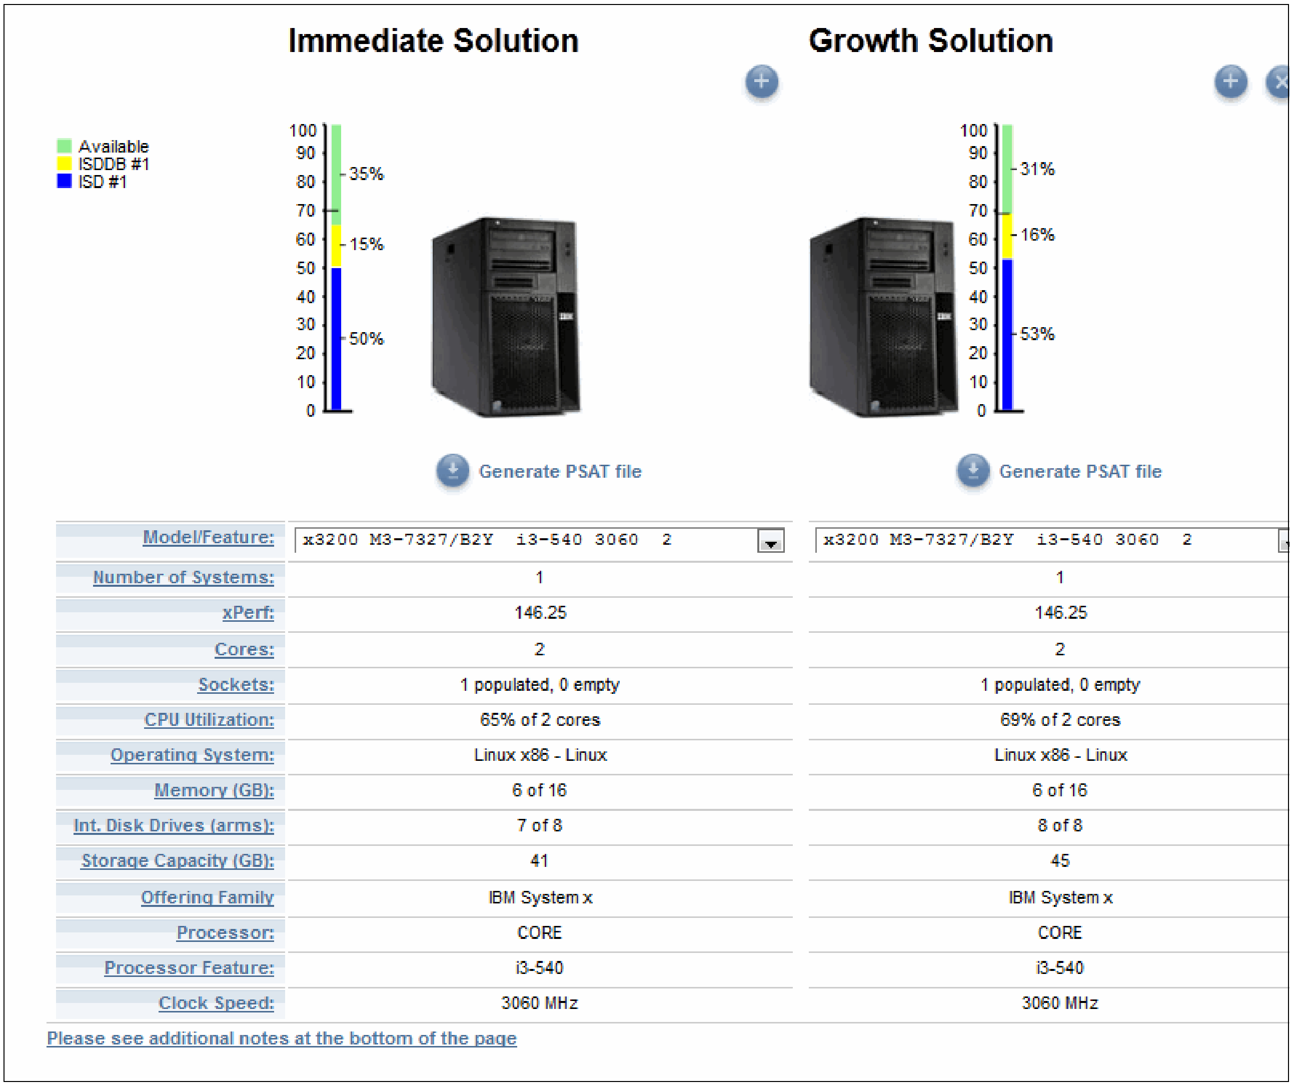

	Immediate Solution	Growth Solution
Model/Feature:	x3200 M3-7327/B2Y i3-540 3060 2	x3200 M3-7327/B2Y i3-540 3060 2
Number of Systems:	1	1
xPerf:	146.25	146.25
Cores:	2	2
Sockets:	1 populated, 0 empty	1 populated, 0 empty
CPU Utilization:	65% of 2 cores	69% of 2 cores
Operating System:	Linux x86 - Linux	Linux x86 - Linux
Memory (GB):	6 of 16	6 of 16
Int. Disk Drives (arms):	7 of 8	8 of 8
Storage Capacity (GB):	41	45
Offering Family	IBM System x	IBM System x
Processor:	CORE	CORE
Processor Feature:	i3-540	i3-540
Clock Speed:	3060 MHz	3060 MHz

Please see additional notes at the bottom of the page

Figure 2-6 Workload Estimator proposed solution

You can further modify the configuration by reviewing the selected system and by using the modify section to change the configuration, as shown in Figure 2-7.

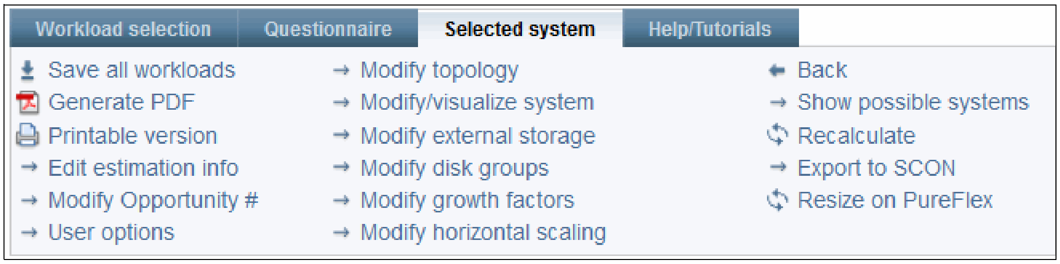

Figure 2-7 Modify the Workload Estimator selection

This output does not imply that you acquire new hardware. However, this output can be used as a guide to place a system in an environment that has available resources.

2.3 Before you begin

Before the installation, review the requirements that are applicable to the operating system that you use for the installation and the current hardware environment.

Management server: Carefully plan the hardware and virtualization environment to be managed by the management server.

- Hardware requirements are listed in the Information Center at the following site:

  http://pic.dhe.ibm.com/infocenter/director/pubs/topic/com.ibm.director.plan.helps.doc/fqm0_r_hardware_requirements.html

- The supported operating systems are listed in the Information Center at the following site:

  http://pic.dhe.ibm.com/infocenter/director/pubs/index.jsp?topic=%2Fcom.ibm.director.plan.helps.doc%2Ffqm0_r_supported_operating_systems.html

- Security features and considerations are documented in the Information Center at the following site:

  http://pic.dhe.ibm.com/infocenter/director/pubs/index.jsp?topic=%2Fcom.ibm.director.security.helps.doc%2Ffqm0_c_security.html

Perform the following primary tasks:

- Install the Systems Director server. (Chapter 3, “Installation” on page 35)
- Start the Systems Director server. (3.1.5, “Starting IBM Systems Director” on page 55)
- Update the Systems Director server. (4.3.9, “Updating the Systems Director server” on page 134)

File system requirements that are needed for the installation are documented at the following website:

http://pic.dhe.ibm.com/infocenter/director/pubs/index.jsp?topic=%2Fcom.ibm.director.plan.helps.doc%2Ffqm0_r_hardware_requirements_servers_running_aix.html

The installation of the Systems Director server is not the only step during the setup phase. Figure 2-8 on page 16 illustrates the interaction within Systems Director and the interactions among the components of the server:

- Command-line interface (CLI) interaction
- Operating system and hardware
- Network speed
- Disk subsystem
- Database activity
- Concurrent users
- Managed systems

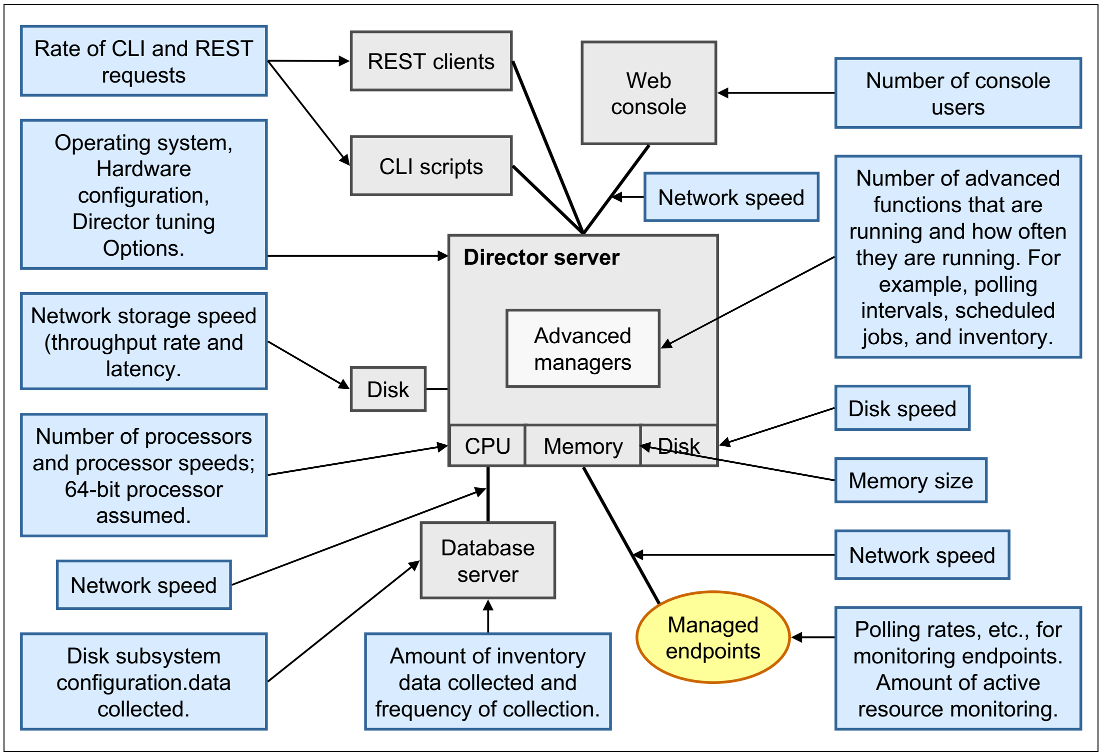
Rate of CLI and REST requests
REST clients
Web console
Number of console users
Operating system, Hardware configuration, Director tuning Options.
CLI scripts
Network speed
Number of advanced functions that are running and how often they are running. For example, polling intervals, scheduled jobs, and inventory.
Director server
Advanced managers
Network storage speed (throughput rate and latency.
Disk
Disk speed
CPU
Memory
Disk
Memory size
Number of processors and processor speeds; 64-bit processor assumed.
Database server
Network speed
Network speed
Managed endpoints
Disk subsystem configuration.data collected.
Amount of inventory data collected and frequency of collection.
Polling rates, etc., for monitoring endpoints. Amount of active resource monitoring.

Figure 2-8 Systems Director interaction

When it comes to placement of the IBM Systems Director server, network connectivity is critical, including DMZ and network firewalls. If firewalls are placed between the management server and the systems to be managed, changes must be made to allow for the required information flow.

A list of all TCP/IP ports that are used by IBM Systems Director are listed at the following site:

http://pic.dhe.ibm.com/infocenter/director/pubs/index.jsp?topic=%2Fcom.ibm.director.plan.helps.doc%2Ffqm0_r_all_available_ports.html

Figure 2-9 displays a flow from the IBM Systems Director server to a discovered Hardware Management Console (HMC) and AIX operating systems.

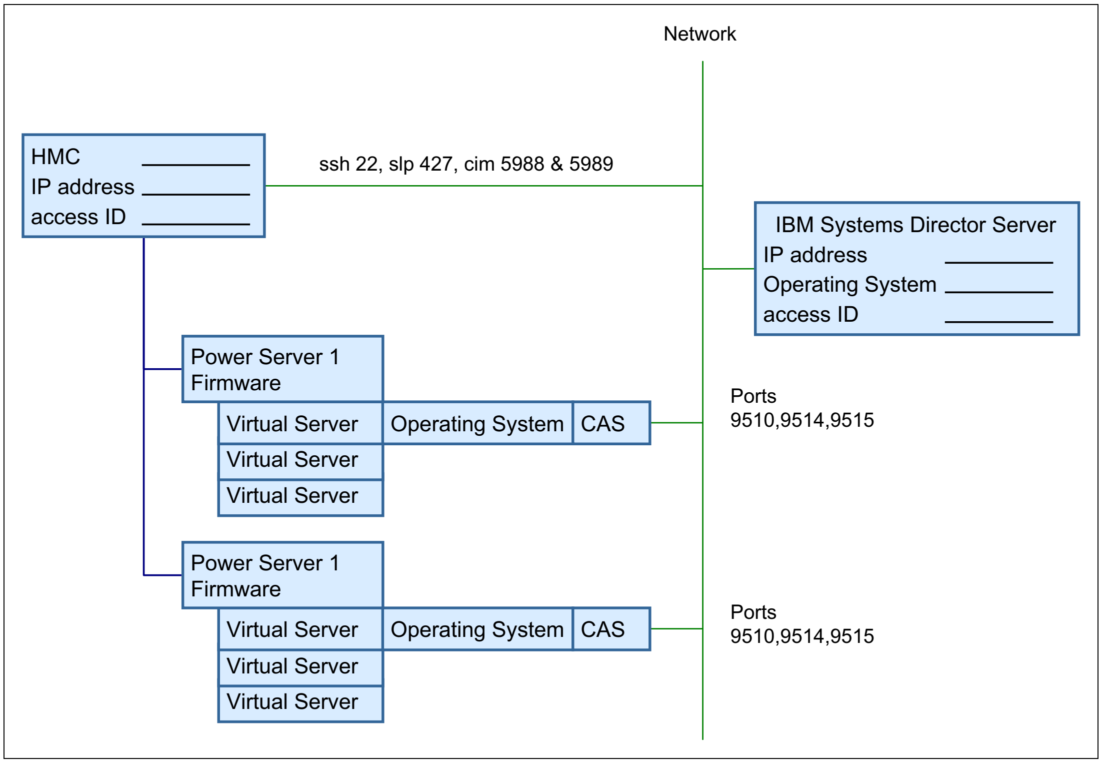

Figure 2-9 Sample connectivity for Power Systems

Before installing IBM Systems Director server, ensure that no server ports are in use by using the **netstat** and **rmsock** commands. Figure 2-10 lists examples of the **netstat** command for active ports.

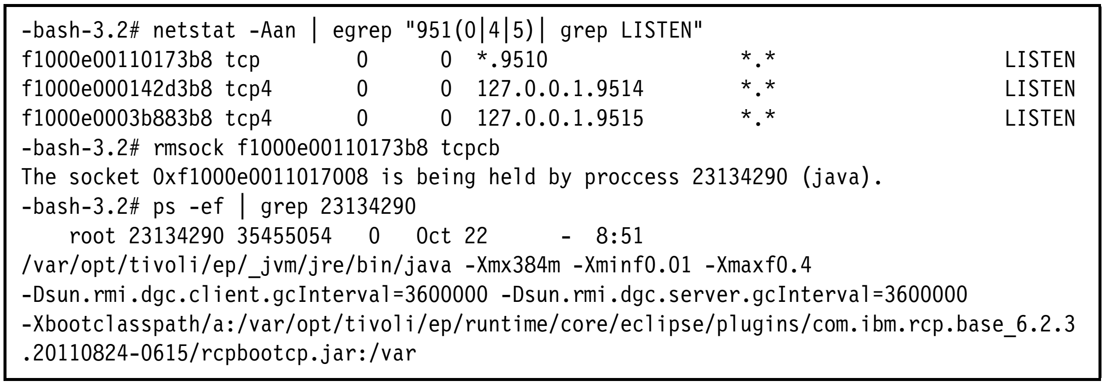

```
-bash-3.2# netstat -Aan | egrep "951(0|4|5)| grep LISTEN"
f1000e00110173b8 tcp        0      0  *.9510           *.*                   LISTEN
f1000e000142d3b8 tcp4       0      0  127.0.0.1.9514   *.*                   LISTEN
f1000e0003b883b8 tcp4       0      0  127.0.0.1.9515   *.*                   LISTEN
-bash-3.2# rmsock f1000e00110173b8 tcpcb
The socket 0xf1000e0011017008 is being held by proccess 23134290 (java).
-bash-3.2# ps -ef | grep 23134290
    root 23134290 35455054   0   Oct 22      -  8:51
/var/opt/tivoli/ep/_jvm/jre/bin/java -Xmx384m -Xminf0.01 -Xmaxf0.4
-Dsun.rmi.dgc.client.gcInterval=3600000 -Dsun.rmi.dgc.server.gcInterval=3600000
-Xbootclasspath/a:/var/opt/tivoli/ep/runtime/core/eclipse/plugins/com.ibm.rcp.base_6.2.3
.20110824-0615/rcpbootcp.jar:/var
```

Figure 2-10 The netstat -Aan and rmsock commands

By using the **netstat** and **rmsock** commands, you can see which process is holding the port and take corrective action to free the port before the IBM Systems Director server installation.

An alternate to the **netstat** and **rmsock** commands is to use the **lsof** command to list open files. You can download **lsof** from this link as part of the AIX Expansion Pack:

http://www-03.ibm.com/systems/power/software/aix/expansionpack/index.html

The **netcat** command is an option for both Transmission Control Protocol (TCP) and User Datagram Protocol (UDP) to check connectivity between machines, as shown in Figure 2-11.

```
-bash-3.2# netcat -zv 9.42.171.23 389
xs-2120rhelppc.itso.ral.ibm.com [9.42.171.23] 389 (ldap) open
-bash-3.2# netcat -zv 9.42.171.23 9510
xs-2120rhelppc.itso.ral.ibm.com [9.42.171.23] 9510 (?) open
-bash-3.2#
```

Figure 2-11 The netcat command

2.4 Best Practice for planning endpoint management

IBM Systems Director has multiple methods to manage different endpoints. The method depends on the type of equipment that you plan to manage. Servers are the most common endpoints.

Complete these prerequisites to manage endpoints by using Systems Director:

- *Verify that the Domain Name System (DNS) functions correctly for both forward and reverse lookup.*

  Systems Director uses standard networking technologies, such as DNS, to identify and communicate with the endpoints.

- *Open the necessary firewall ports.*

  Systems Director uses several ports to communicate with various endpoints that need to be open. Each type of device requires a group of ports. Determining which ports to open depends on what you plan to manage. See section 2.5, “Firewall ports” on page 19 for a list by function.

  For more information about ports that are used by the Systems Director server, managed systems, and important port considerations, go to the following site:

  http://pic.dhe.ibm.com/infocenter/director/pubs/index.jsp?topic=%2Fcom.ibm.director.plan.helps.doc%2Ffqm0_r_all_available_ports.html

- *Determine how you want to manage your systems.*

  You can choose agentless, Platform Agent, or Common Agent. For more information, see section 2.7.1, “Determine which agent to use to manage your systems” on page 26.

2.5 Firewall ports

Table 2-3 through Table 2-20 on page 24 list the TCP and UDP ports that need to be open for specific Systems Director functions to work correctly.

Table 2-3 BladeCenter advanced management module (AMM) (out-of-band)

Protocol	Description	IBM Systems Director server port	Managed endpoint port
Service Location Protocol (SLP)	Discovery	427, TCP, UDP Inbound, Outbound	427, TCP, UDP Inbound, Outbound
SLP	Communication and discovery	14252,[a] TCP, UDP Inbound, Outbound	14252, TCP, UDP Inbound, Outbound
TCP Command Mode	Ongoing communication and management	6090, TCP Outbound	6090, TCP Inbound
UDP	Native events	13991, UDP Inbound	13991, UDP Outbound
Simple Network Management Protocol (SNMP)	SNMP communication/traps	162, TCP, UDP Inbound, Outbound	162, TCP, UDP Inbound, Outbound
Trivial File Transfer Protocol (TFTP)/SNMPv3	Updating AMM firmware	69, UDP Inbound; 121, UDP Inbound, Outbound	69, UDP Outbound; 121, UDP Outbound, Inbound

[a] On both the server and the MEP, the source port plus the next 25 ports (or 75 when more systems are discovered concurrently) must be open. The source port can be changed in the `slp.prop` files.

Table 2-4 Integrated management module 1 (IMM1) and IMM2 (out-of-band rack servers)

Protocol	Description	IBM Systems Director server port	Managed endpoint port
SLP	Discovery	427, TCP, UDP Inbound, Outbound	427, TCP, UDP Inbound, Outbound
UDP	Native events	13991, UDP Inbound	13991, UDP Outbound
SLP	Communication and discovery	14252,[a] TCP, UDP Inbound, Outbound	14252, TCP, UDP Inbound, Outbound

[a] On both the server and the MEP, the source port plus the next 25 ports (or 75 when more systems are discovered concurrently) must be open. The source port can be changed in the `slp.prop` files.

Table 2-5 Management module (MM), Remote Supervisor Adapter (RSA 1/2)

Protocol	Description	IBM Systems Director server port	Managed endpoint port
SLP	Discovery	427, TCP, UDP Inbound, Outbound	427, TCP, UDP Inbound, Outbound
SLP	Communication and discovery	14252,[a] TCP, UDP Inbound, Outbound	14252, TCP, UDP Inbound, Outbound
CIM	Ongoing communication and management	5988 or 15988[b], TCP Inbound, Outbound (unsecure); 5989 or 15988[b], TCP Inbound, Outbound (secure)	5988 or15988[b], TCP Inbound, Outbound (unsecure); 5989 or 15988[b]TCP Inbound, Outbound (secure)
TFTP	Updates for System x/Flex servers that run ESXi	69, UDP Outbound	69, UDP Inbound

a. On both the server and the MEP, the source port plus the next 25 ports (or 75 when more systems are discovered and managed concurrently) must be open. The source port can be changed in the `slp.prop` files.

b. CIM Server for SUSE Linux Enterprise Server.

Table 2-6 Flex Chassis Management Module (CMM)

Protocol	Description	IBM Systems Director server port	Managed endpoint port
SLP	Discovery	427, TCP, UDP Inbound, Outbound	427, TCP, UDP Inbound, Outbound
SLP	Communication and discovery	14252,[a] TCP, UDP Inbound, Outbound	14252, TCP, UDP Inbound, Outbound
CIM	Ongoing communication and management	5988 or 15988[b], TCP Inbound, Outbound (unsecure); 5989 or 15988[b], TCP Inbound, Outbound (secure)	5988 or 15988[b], TCP Inbound, Outbound (unsecure); 5989 or 15988[b], TCP Inbound, Outbound (secure)
SFTP	Update CMM firmware	9520, TCP Outbound	9520, TCP Inbound

a. On both the server and the MEP, the source port plus the next 25 ports (or 75 when more systems are discovered and managed concurrently) must be open. The source port can be changed in the `slp.prop` files.

b. CIM Server for SUSE Linux Enterprise Server.

Table 2-7 Hardware Management Console (HMC)

Protocol	Description	IBM Systems Director server port	Managed endpoint port
SSH	Ongoing communication with limited management	22, TCP Outbound	22, TCP Inbound
CIM	Ongoing communication and management	5989, TCP Inbound, Outbound (secure)	5989, TCP Inbound, Outbound (secure)

Table 2-8 Windows agentless

Protocol	Description	IBM Systems Director server port	Managed endpoint port
DCOM	Ongoing communication with limited management	135, TCP, UDP Outbound (software installation); 137 - 139, TCP, UDP Outbound; 445, TCP, UDP Outbound	135, TCP, UDP Inbound (software installation); 137 - 139, TCP, UDP Inbound; 445, TCP, UDP Inbound

Table 2-9 Linux agentless

Protocol	Description	IBM Systems Director server port	Managed endpoint port
SSH	Ongoing communication with limited management	22, TCP Outbound	22, TCP Inbound

Table 2-10 AIX agentless

Protocol	Description	IBM Systems Director server port	Managed endpoint port
SSH	Ongoing communication with limited management	22, TCP Outbound	22, TCP Inbound

Table 2-11 VMWare ESXi

Protocol	Description	IBM Systems Director server port	Managed endpoint port
SLP	Discovery	427, TCP, UDP Inbound, Outbound	427, TCP, UDP Inbound, Outbound
CIM	Ongoing communication and management	5988, TCP Inbound, Outbound (unsecure); 5989, TCP Inbound, Outbound (secure)	5988, TCP Inbound, Outbound (unsecure); 5989, TCP Inbound, Outbound (secure)

Table 2-12 IBM Systems Director Platform Agent

Protocol	Description	IBM Systems Director server port	Managed endpoint port
SLP	Discovery	427, TCP, UDP Inbound, Outbound	427, TCP, UDP Inbound, Outbound
CIM	Ongoing communication and management	5988 or 15988[a], TCP Inbound, Outbound (unsecure); 5989 or 15988[a], TCP Inbound, Outbound (secure)	5988, 15988[a] TCP Inbound, Outbound (unsecure); 5989, 15988[a] TCP Inbound, Outbound (secure)

[a] CIM Server for SUSE Linux Enterprise Server

Table 2-13 IBM Systems Director Common Agent (CAS)

Protocol	Description	IBM Systems Director server port	Managed endpoint port
SLP	Discovery	14252, TCP, UDP Inbound, Outbound	14252, TCP, UDP Inbound, Outbound
CAS	All ongoing communication and management	9510, TCP Inbound, Outbound; 9511 - 9513, TCP Inbound; 20000, TCP Inbound	9510, TCP Inbound; 9511 - 9513, TCP Outbound; 20000, TCP Outbound

Table 2-14 I/O modules

Protocol	Description	IBM Systems Director server port	Managed endpoint port
SSH	Ongoing communication with limited management	22, TCP Outbound	22, TCP Inbound
SNMP	Monitoring	162, TCP, UDP Inbound, Outbound	162, TCP, UDP Inbound, Outbound
TFTP/SFTP/FTP	Firmware updates	FTP 20 - 21 TCP Inbound; SFTP 9520 TCP Inbound; TFTP 69 UDP Inbound	FTP 20 - 21 TCP Outbound; SFTP 9520 TCP Outbound; TFTP 69 UDP Outbound

Table 2-15 SNMP Devices

Protocol	Description	IBM Systems Director server port	Managed endpoint port
SNMP	SNMP communication/traps	162, TCP, UDP Inbound, Outbound	162, TCP, UDP Inbound, Outbound

Table 2-16 IBM Systems Director server Service and Support Manager

Protocol	Description	IBM Systems Director server port	Managed endpoint port
HTTPS	Communication with IBM	443, TCP Outbound	N/A
FTP	Service log upload	21, TCP Outbound	N/A

Table 2-17 IBM Systems Director server Update Manager

Protocol	Description	IBM Systems Director server port	Managed endpoint port
HTTP	Check for updates, Download updates	80, TCP, Outbound	N/A
HTTPS	Check for updates, Download updates	443, TCP, Outbound	N/A

Table 2-18 IBM Systems Director server web interface

Protocol	Description	IBM Systems Director server port	Managed endpoint port
HTTP	HTTP communication with IBM Systems Director web interface (auto redirects to HTTPS)	8421, TCP, Inbound	N/A
HTTPS	HTTPS communication with IBM Systems Director web interface	8422, TCP, Inbound	N/A

Table 2-19 Default Managed DB2 database on IBM Systems Director server

Protocol	Description	IBM Systems Director server port	Managed endpoint port
FCM	Database communication	50010, TCP Inbound, Outbound	N/A

Table 2-20 IBM Systems Director CLI (smcli)

Protocol	Description	IBM Systems Director server port	Managed endpoint port
TCP	Command-line interface (CLI)	2044, TCP Inbound, Outbound	N/A

2.6 IBM Systems Director Editions

When you want to install IBM Systems Director server, you can select either to download the Director installation package from the download site or you can select an IBM Systems Director Edition.

The IBM Systems Director Editions provide a bundle of IBM Systems Director, Advanced Managers, and Service and Support.

There are different versions of editions available, depending on your requirements and the operating system on which the IBM Systems Director will be installed. The following versions are available:

- *Express* (available for System x and Power Systems)

  This version includes management tools that can reduce operational complexity and provide the ability to remotely discover, monitor, configure, and update your systems.

- *Standard* (Available for System x, Power Systems, and System z)

  This version includes all of the capabilities of the Express Edition and adds advanced deployment, monitoring, and control features under the same console. It also includes features for optimized energy usage and capacity.

- *Enterprise* (available for Power Systems)

  This version includes all of the capabilities of the Express and Standard Editions plus features to enable automated resource provisioning and balancing that are required for a dynamic infrastructure. It also provides enterprise management tools for advanced troubleshooting, capacity analysis, and reporting.

Table 2-21 provides an overview of the features of the IBM Systems Director Editions by platform. This helps you to find the right IBM Systems Director Edition for your requirements.

Table 2-21 IBM Systems Director Editions v6.3.2

Products	IBM System Director Editions					
	System x[a]		Power Systems			System z
	Express	Standard	Express	Standard	Enterprise	Standard
IBM Systems Director v6.3.2	x	x	x	x	x	x
IBM Systems Director Service and Support Manager v6.3.2	x	x	x	x	x	x
IBM Systems Director Active Energy Manager - Monitoring v4.4.2	x	x	x	x	x	x

Products	IBM System Director Editions					
	System x[a]		Power Systems			System z
	Express	Standard	Express	Standard	Enterprise	Standard
IBM Systems Director VMControl Express v 2.4.2	x	x	x	x	x	x
IBM Systems Director Active Energy Manager - Monitoring v4.4.2		x		x	x	x
IBM Systems Director Network Control v1.4.0 (with license restrictions[b])		x		x	x	x
IBM Systems Director VMControl Standard Edition - Image Management v2.4.2		x[c]		x	x	x
IBM Systems Director VMControl Standard Edition - System Pools v2.4.2					x	
IBM Systems Director Storage Control v 4.2.1.1					x	
Tivoli content (AIX only)						
IBM Tivoli Monitoring (ITM) v6.3					x	
IBM Tivoli Common Reporting v3.1					x	
Platform-unique functions						
IBM Upward Integration for Microsoft System Center 4.0		x				
IBM Upward Integration Module for VMware (ISV) v 1.3		x				
IBM Tivoli Provisioning Manager for OS Deployment v7.1.1[d]		x				
IBM Systems Director Agent for z/BX Power and x86 Blades v6.3.2						x
IBM Virtual Media Key (IMMv1) or IBM Integrated Management Module (IMMv2) Advanced Upgrade		x[e]				

[a] Also supports Flex Systems (x86 compute node, without FSM only).

[b] License includes the ability to perform network device management (discovery, inventory, monitoring status) and to use network diagnostics. All other features of network control (including topology view, interface to vendor configuration tools, and configuration of host VLANs) require a full-use license of network control.

[c] Kernel-based Virtual Machine (KVM) support only.

[d] TPMfOSD is a separate installation. It is not integrated into the IBM Systems Director console or managed through the IBM Systems Director.

[e] Edition can be ordered with or without this feature.

The Advanced Manager, which is part of the IBM Systems Director editions can be licensed separately for installations of the IBM Systems Director. The editions that are listed above are not used.

Service and Support is included in the IBM Systems Director Editions. If you need service and support for your IBM Systems Director installation and have no use of an IBM Systems Director Edition, contact your IBM representative and ask for service offerings for the IBM Systems Director.

Tivoli Provisioning Manager for OS Deployment IBM Systems Director Edition (TPMfOSD ISD Edition) is a separate installation, which is not integrated into IBM Systems Director Console. The TPMfOSD is a tool for bare metal installations for IBM System x servers (for Windows, Linux, and VMWare OS) and IBM PowerPC® Systems (for AIX and Linux). There is no support for the IBM Power Systems server in this product version.

For more information about this tool, the installation process, and usage, see the Tivoli information center at the following link:

http://pic.dhe.ibm.com/infocenter/tivihelp/v3r1/index.jsp?topic=/com.ibm.tivoli.tpm.osd.doc/welcome/osdhome.html

2.7 Architectural decisions and recommendations

This section contains recommendations on some of the architectural decisions that you will have to make when planning for an IBM Systems Director server deployment. For each decision, all possible alternatives are described, as well as the benefits and disadvantages for each decision. Some of the recommendations might not apply to your environment, based on your requirements and constraints.

2.7.1 Determine which agent to use to manage your systems

IBM Systems Director performance and scalability are intrinsically linked to the allocated resources, the number of discovered systems, and the associated management type. One of the key decisions that you will have to make is which agent, if any, you will use to manage your endpoints.

The following server platforms are available:

- Agentless
- Platform Agent
- Common Agent

Agentless

Agentless management provides *hardware alerting* out-of-band either through the IMM or AMM. With agentless management, you can inventory your systems by using the distributed component object model (DCOM) for Windows or Secure Shell (SSH) for Linux.

For more information about agentless systems, see the following site:

http://pic.dhe.ibm.com/infocenter/director/pubs/index.jsp?topic=%2Fcom.ibm.director.install.helps.doc%2Ffqm0_t_preparing_agentless_managed_systems.html

If you think agentless monitoring is right for you, consider the following factors:

- To get accurate hardware event information, your system requires a management module. You should also discover this management module *out-of-band* (meaning that you should discover it directly by using the management module IP address) because functionality, such as launch-in-context remote control or power operations, are only available when the management module has been discovered out of band.

- Agentless hardware monitoring might not include certain event types from internal devices such as Redundant Array of Independent Disks (RAID) adapters or Peripheral Component Interconnect (PCI) controllers. These require additional CIM modules, which are included only with the Platform Agent (the Common Agent includes the Platform Agent).
- Unless you have a requirement to monitor operating system resources and processes with IBM Systems Director, and you do not have a requirement to get extended events from certain internal devices such as RAID controllers and PCI cards, and if you are not planning on using VMControl with your Microsoft Hyper-V or Linux KVM hosts, agentless is probably the best option to start with for Windows and Linux servers.
- Agentless is your only option for managing VMware vSphere ESXi hosts. You should not attempt to discover ESXi hosts directly. Instead, you need to discover your vCenter Server endpoint as described in section 5.4, “Managing VMware vSphere with VMControl” on page 282.
- On Linux, agentless access can be configured with the **sudo** utility, as shown at the following site:

  http://pic.dhe.ibm.com/infocenter/director/pubs/topic/com.ibm.director.install.helps.doc/fqm0_t_setting_up_access_to_agentless_systems.html
- If the credentials that you used to unlock access for your agentless endpoint change (for example, when your password expires or is changed), you have to manually update those credentials in IBM Systems Director to recover access to your agentless endpoints. Go to **Actions** → **Security** → **Configure Access** to do this action.
- Inventory after an agentless discovery creates a server MEP object. This should not be confused with the management module, which requires out-of-band discovery.

Platform Agent

Platform Agent is the lightweight agent that is installed on Windows or Linux systems that provides everything that agentless management provides. Platform Agent also provides additional operating system level monitoring.

For more information about Platform Agent systems, go to the following site:

http://pic.dhe.ibm.com/infocenter/director/pubs/index.jsp?topic=%2Fcom.ibm.director.main.helps.doc%2Ffqm0_c_platform_agent.html

The following scenarios should be considered when you think the Platform Agent is right for you:

- If you have Microsoft Hyper-V or Linux KVM hosts that you are planning to manage with IBM Systems Director VMControl, the Platform Agent is required.
- If you need to monitor operating system processes or resource utilization using operating system performance library objects, the Platform Agent is also required. See Table 2-22 on page 29 for a list of supported operating system monitors for the Platform Agent.
- The Platform Agent includes additional CIM providers that include additional hardware events from components such as internal RAID adapters or PCI cards.

Note: Previous versions of IBM Systems Director require the installation of a separate RAID management software. CIM providers for IBM hardware RAID controllers are now included with the Platform Agent.

- If the credentials that are used to unlock access to the Platform Agent change, you have to request access again to unlock the CIM protocol. Until you do that, managed endpoints show only partial access.
- The Platform Agent is required if you are planning on using Storage Management with your managed endpoints.

Common Agent

The *Common Agent* provides everything that the Platform Agent offers (it includes the Platform Agent code) and adds functionality. It can monitor the operating system performance, services, and processes.

For more information about Common Agent systems, go to the following site:

http://pic.dhe.ibm.com/infocenter/director/pubs/index.jsp?topic=%2Fcom.ibm.director.main.helps.doc%2Ffqm0_c_common_agent.html

Consider the following factors when you think the Common Agent is right for you:

- The Common Agent uses more hardware resources than the Platform Agent and has a larger footprint. This might not be an issue for most environments where physical servers are under-utilized, but it is nevertheless something to keep in mind. If you are concerned with resource utilization, ensure that you plan for benchmarking to determine the exact impact that the Common Agent will have in your environment.
- Systems running the Common Agent can be managed only by a single Common Agent server. This means that an endpoint running the Common Agent can be managed only by a single IBM Systems Director server.
- You can install the Common Agent on top of the platform agent by using IBM Systems Director release management capabilities.
- You should manage AIX systems using the Common Agent.
- You should avoid using the Common Agent when you are already using the IBM Tivoli Monitoring agent because a Common Agent can be managed only by one Common Agent server. If you must use both Common Agents (Tivoli and IBM Systems Director), it might require special configuration, as documented at the following site:

  https://www.ibm.com/developerworks/community/wikis/home?lang=en#!/wiki/W3e8d1c956c32_416f_a604_4633cd375569/page/Coexistence+of+Director+V6+CAS+Agent+with+other+Tivoli+CAS+Agents
- Using the Common Agent on non-IBM hardware requires additional licensing (IBM Systems Director server comes with 20 licenses only for non IBM hardware). For more information, go to the following site:

  http://pic.dhe.ibm.com/infocenter/director/pubs/topic/com.ibm.director.install.helps.doc/fqm0_t_obtaining_licenses.html
- If you need to monitor operating system processes or resource utilization using operating system performance library objects, the Platform Agent is also required. See Table 2-22 on page 29 for a list of supported operating system monitors for the Common Agent.
- The *Manage Processes* and *Process Monitors* actions are only available with the Common Agent.
- The Common Agent is not affected if the credentials used to unlock the endpoint change. This is because the Common Agent uses its own credentials after the endpoint has been initially unlocked.

Agent capabilities

The IBM Systems Director Information Center helps you choose the level of agent capabilities to deploy on managed systems:

http://pic.dhe.ibm.com/infocenter/director/pubs/index.jsp?topic=%2Fcom.ibm.director.plan.helps.doc%2Ffqm0_t_agent_tiers.html

The preceding link also contains additional links that show, per platform, what type of functionality each agent provides.

Table 2-22 shows the agent capabilities for the Windows platform.

Table 2-22 List of available Windows Server 2008 R2 monitors by agent type

Monitor	Agentless	Platform Agent	Common Agent
CPU% Utilization	No	Yes	Yes
Disk Space Remaining	No	Yes	Yes
Memory Usage	No	Yes	Yes
Primary File System Percent Space Available	No	Yes	Yes
Process Count	No	Yes	Yes
Disk Workload	No	No	Yes
Disk Space Used	No	No	Yes
IP Packets Received with Errors/sec	No	No	Yes
IP Packets Received/sec	No	No	Yes
IP Packets Sent/sec	No	No	Yes
Locked Memory	No	No	Yes
Primary File System Percent Space Used	No	No	Yes
TCP Connections	No	No	Yes
UDP Datagrams Received/sec	No	No	Yes
UDP Datagrams Sent/sec	No	No	Yes

If you are still unsure which agent you need, list your requirements and then plan to test functionality starting from agentless and working your way up to the Platform Agent and then the Common Agent.

2.7.2 Determine how to deploy your agents

After you have decided which type of agent to use on your systems, you need to select the best way to install the agents.

The following options are available to install the agents:

- Install the platform or Common Agent manually on your systems using binary packages, then register them on your IBM Systems Director server. This is desirable when you have multiple distributed locations with limited wide area network (WAN) capabilities.

  The main downside of deploying agents in this manner is that it requires additional labor because agent installation has to be performed endpoint by endpoint. In addition, installation packages have to be manually copied to each endpoint.

- Include the Platform or Common Agent that is installed in your gold image and register them later on your IBM Systems Director server.

  Again, this may be the best choice in distributed environments. Cloning pre-installed IBM Systems Director agent instances requires additional steps as documented in the Information Center:

  `http://pic.dhe.ibm.com/infocenter/director/pubs/index.jsp?topic=%2Fcom.ibm.director.discovery.helps.doc%2Ffqm0_t_discovering_systems_mirrored_image.html`

  The main disadvantage of doing this is that it requires additional configuration and therefore additional labor. Also, the agent version may quickly become outdated, requiring an update of the gold image.

- Discover all of your systems that are agentless, then use the Release Management and Agent Installation wizard from your IBM Systems Director console.

  This is the preferred method because it helps ensure that up-to-date agents are properly deployed, configured, and registered. You can also do it on multiple systems concurrently, if your local or wide area network can handle the load.

  This method can also be used to upgrade a Platform Agent to a Common Agent. You cannot use central management to uninstall agents.

2.7.3 Determine if you should run Systems Director server in a virtual machine

Keep in mind the following factors when selecting a server platform for running your IBM Systems Director server instance:

- The best way to run IBM Systems Director server is in a virtual machine. If your virtual machine is correctly sized (see 2.1, “System resources” on page 8), running IBM Systems Director in a virtual machine will not affect overall performance or functionality.
- There is no real high-availability solution for IBM Systems Director server. If you run it on a physical server platform, it might be more difficult to restore service quickly, whereas virtual infrastructure usually includes automated protection against hardware failure.
- If you run IBM Systems Director in a virtual machine, keep in mind that the license agreement requires that the underlying physical server that runs the hypervisor must be IBM server hardware. For example, if you have a multivendor cluster, this might require that you implement host affinity rules.

- Always plan to install the IBM Systems Director server instance into its own file system (as opposed to installing it in the same file system partition as the operating system). It is difficult to predict file growth with IBM Systems Director, especially when problems occur and multiple Java core dumps or log files are generated. Although this happens less frequently than with previous releases, it is still a potential risk and thus an IBM Systems Director instance should be isolated from other applications' file systems.

2.7.4 Determine what type of database to use with Systems Director server

IBM Systems Director server 6.3.2 comes with an embedded full unlimited version of IBM DB2 for exclusive use by IBM Systems Director. This means that there is no reason to consider using an external database. For more information about database selection, go to the following site:

http://pic.dhe.ibm.com/infocenter/director/pubs/topic/com.ibm.director.plan.helps.doc/fqm0_t_selecting_the_ibm_director_database_application.html

> **Note:** If you are planning on using Storage Control, you cannot use Microsoft SQL Server or Oracle as your database management system (DBMS) for IBM Systems Director. The embedded DB2 or an external DB2 server are your only options.

2.7.5 Determine the networks that your server should be connected to

Depending on the type of traffic, some environments might require network separation. For example, your company policy may require separate networks for out-of-band management modules (such IMM or AMM), system management traffic (such as remote connections or monitoring), and production data (where applications and user traffic are located).

In these types of environments, where should you place the IBM Systems Director server primary network interface?

The answer depends on how you manage your infrastructure. For example, if you are planning on managing your hardware using only out-of-band management modules, it might be better to install your IBM Systems Director server in the management module network segment.

In most cases, your Systems Director has its primary interface in the administration network segment and also requires connectivity to the management module network segment (either directly via a secondary interface, which is the preferred method, or through a firewall).

> **Warning:** Service Location Protocol (SLP) discovery through firewalls can be tricky and requires special configuration, as documented at the following site:
>
> http://pic.dhe.ibm.com/infocenter/director/pubs/topic/com.ibm.director.plan.helps.doc/fqm0_r_ports_considerations.html

In general, your IBM Systems Director should not require network connectivity to the production network because it should be considered as a system management platform. However, if you need to restrict IBM Systems Director traffic to a single network interface, this is only supported on Windows and requires specific configuration, as documented at the following site:

http://pic.dhe.ibm.com/infocenter/director/pubs/topic/com.ibm.director.agent.helps.doc/fqm0_t_binding_pa_specific_ip_addresses.html

Finally, if you are planning on using the Update Manager functionality of IBM Systems Director (which enables firmware updates for your managed endpoints), Internet connectivity (either direct or through a proxy server) is highly desirable because you would otherwise have to manually import all applicable updates.

2.7.6 Determine how to make your solution highly available

There is no specific solution for making IBM Systems Director highly available. The best strategy to enhance availability is to take advantage of your infrastructure capabilities (such as HA and DRS clusters if you run your IBM Systems Director server in a VMware virtual machine) to protect the IBM Systems Director instance from hardware failure.

In addition, you should have a backup strategy for your IBM Systems Director server as highlighted in 4.7, “Backup” on page 233.

2.7.7 Determine how to deal with distributed environments

Previous versions of IBM Systems Director had a hierarchical management server (HMS) functionality built-in because it was not reliable enough and did not scale well. Support for HMS has been dropped from the product.

This means that if you have a large distributed environment (for example, you are a retailer with hundreds of remote stores that have servers you want to be able to monitor and update), consider the following factors:

- If you have a few managed endpoints (less than 10) in each location and you are not planning on using Update Manager (for firmware or agent updates) for your remote locations, even a modest wide area network bandwidth may be enough to support basic hardware monitoring functions from a central IBM Systems Director server instance in your data center.
- If you are planning on using Update Manager, plan to have at least 10 Mbps of network bandwidth because some of the update packages can be fairly large (several hundred MB).
- If you have many managed endpoints and are planning to use Update Manager and do not have 10 Mbps or more of network bandwidth, have a local IBM Systems Director server instance, which will have to be managed on its own.

Note: Most firmware updates and agent update packages can also be distributed because any other software packages using an alternative software distribution solution might be more suitable to distributed environments.

2.8 Best practices

To summarize, consider the following best practices when planning for IBM Systems Director:

- If it covers your requirements, *prefer agentless systems* because they prevent you from having to keep agents up to date. This recommendation is only applicable if you have a System x server without any internal storage with configured out-of-band management modules.
- Plan to use *privileged credentials with a non-expiring password* for agentless endpoints. Otherwise, you have to manually update all existing credentials endpoint by endpoint every time that the password is changed or when it expires.

- *If you cannot use non-expiring privileged credentials, use domain credentials*. When the domain credentials change, revoke access to all managed endpoints by using those credentials, then request access and enter the new credentials.
- Whenever possible, *deploy agents centrally from the IBM Systems Director server console.* This ensures that you have registered properly and have up-to-date endpoints.
- *Run IBM Systems Director server in a virtual machine.* This helps enhance availability of the service.
- Use *+16 GB RAM and four CPU cores* or four *vCPUs* as best practice or for performance of the IBM Systems Director server.
- *Use the embedded DB2 database* because it has no limitations (as opposed to previous versions of IBM Systems Director). A local database helps with overall performance and reduces complexity.
- *Always discover management modules out of band*, even if they have already been discovered in band because this unlocks key functionality.
- *Plan for Internet connectivity for your IBM Systems Director server* so that it can retrieve updates directly from the web.
- *For distributed environments, plan to use IBM Systems Director only for hardware monitoring* and leverage your existing software distribution solution to push agent and firmware updates.
- *Use the Workload Estimator* to get basic figures for the system running Systems Director.
- *When possible, use one of the more advanced IBM Systems Director Editions.* This provides, in addition to the IBM Systems Director and the Advanced Manager software, a Service and Support offering (SWMA) for one or three years. Choose the edition that best suits your requirements and your environment.
- *Open the necessary firewall ports* in your environment for IBM Systems Director server and managed endpoints.
- *Verify that the Domain Name System (DNS) functions correctly* for both forward and reverse lookup.

3

Installation

This chapter provides information and best practices for installing IBM Systems Director server on different platforms.

The following topics are covered:

- 3.1, “Installation of IBM Systems Director server on an x86 platform” on page 36
- 3.2, “Installing IBM Systems Director server on an AIX platform” on page 57
- 3.3, “Installation of IBM Systems Director server for a Linux on Power platform” on page 68
- 3.4, “Installation of the IBM Systems Director agent on Linux x86” on page 73
- 3.5, “Best practices” on page 76

3.1 Installation of IBM Systems Director server on an x86 platform

The IBM Systems Director server runs on a Windows or Linux platform on x86 systems, but only on hardware that is branded IBM (a license requirement). The IBM Systems Director Information Center has a complete list of all supported hardware at the following site:

http://pic.dhe.ibm.com/infocenter/director/pubs/index.jsp?topic=%2Fcom.ibm.director.plan.helps.doc%2Ffqm0_r_hardware_compatibility.html

One way to use IBM Systems Director on x86 is to install the IBM Systems Director server on a virtual machine under the control of a hypervisor, such as VMware ESX or Linux Kernel-based Virtual Machine (KVM). A virtual machine offers advantages that make it a useful installation method:

- Hardware independence of certain components and drivers.
- Simple extensibility of resources (memory and processor allocation).
- The ability to perform snapshots before the installation of plug-ins and advanced managers.

When you use a virtual machine, configure four vCPUs and 16 GB or more of memory. The disk drive size depends on the number of systems and the database that is used for the IBM Systems Director installation.

3.1.1 Supported operating systems

The IBM Systems Director server is supported on the following operating system versions that run on IBM x86 servers (the values with an * indicate enhancements that are coming with IBM Systems Director Version 6.3.3):

- Linux 32-bit:
  - Red Hat Enterprise Linux Advanced Platform, Version 5.0 (supports Updates 1, 2, 3, 4, 5, 6, 7, 8, and 9*)
  - Red Hat Enterprise Linux, Version 5.0 (supports Updates 1, 2, 3, 4, 5, 6, 7, 8, and 9*)
  - Red Hat Enterprise Linux Advanced Platform, Version 6.0 (supports Updates 1, 2, 3, and 4*)
  - Red Hat Enterprise Linux, Version 6.0 (supports Updates 1, 2, 3, and 4*)
  - SUSE Linux Enterprise Server 10 for x86 (supports Service Packs 1, 2, 3, and 4)
  - SUSE Linux Enterprise Server 11 for x86 (supports Service Packs 1 and 2)
- Linux 64-bit:
  - Red Hat Enterprise Linux Advanced Platform, Version 5.0, for AMD64 and EM64T (supports Updates1, 2, 3, 4, 5, 6, 7, 8, and 9*)
  - Red Hat Enterprise Linux, Version 5.0, for AMD64 and EM64T (supports Updates 1, 2, 3, 4, 5, 6, 7, 8, and 9*)
  - Red Hat Enterprise Linux Advanced Platform, Version 6.0, for AMD64 and EM64T (supports Updates 1, 2, 3, and 4*)
  - Red Hat Enterprise Linux, Version 6.0, for AMD64 and EM64T (supports Updates 1, 2, 3, and 4*)
  - SUSE Linux Enterprise Server 10 for AMD64 and EM64T (supports Service Pack 1, 2, 3, and 4)

  - SUSE Linux Enterprise Server 11 for AMD64 and EM64T (supports Service Packs 1 and 2)
- Windows 32-bit:
  - Windows Server 2003, Enterprise, and Standard Editions, Release 2 (supports SP 2)
  - Windows Server 2003, Enterprise, and Standard Editions (supports SP 2)
  - Windows Server 2008, Enterprise, and Standard Editions (supports SP 1 and SP 2)
- Windows 64-bit:
  - Windows Server 2003, Enterprise, and Standard x64 Editions, Release 2 (supports SP 2)
  - Windows Server 2003, Enterprise, and Standard x64 Editions (supports SP 2)
  - Windows Server 2008, Enterprise, and Standard x64 Editions (supports SP 1 and SP 2)
  - Windows Server 2008, Enterprise, and Standard x64 Editions, Release 2 (with or without SP 1)

A detailed list of supported operating systems for the IBM Systems Director server and agents is available at the following site:

http://pic.dhe.ibm.com/infocenter/director/pubs/index.jsp?topic=%2Fcom.ibm.director.plan.helps.doc%2Ffqm0_r_os_supported_by_ibm_director_631.html

You can install IBM Systems Director on a guest OS that runs on a virtualized environment. The supported guest operating systems are supported by both IBM Systems Director and the hypervisor. The following conditions must be true:

- The OS platform is supported by IBM Systems Director.
- The OS platform is supported as a guest OS by a hypervisor.
- The hypervisor is supported by IBM Systems Director.

With these three conditions, IBM Systems Director's support of the OS platform extends to running it as a guest OS on that hypervisor. See the hypervisor product documentation for a list of supported operating systems.

The following hypervisors for the x86 environment are supported:

- VMware ESX 4.0.*x* and 4.1.*x*
- VMware ESXi 4.0.*x* and 4.1.*x* (under the control of VMware vCenter)
- VMware vSphere 5.0.*x* and 5.1.*x* (under the control of VMware vCenter)
- Linux KVM
- Windows Server 2012 and Windows Server 2008 and 2008R2, Enterprise, Standard, and Datacenter x64 Editions with Hyper-V role-enabled

Required resources: The IBM Systems Director server is only supported on hardware that is branded IBM. Therefore, the hypervisor must run on IBM hardware to meet the license requirements.

Resources that are required for running the IBM Systems Director server are referenced in section 2.1, “System resources” on page 8; 2.2, “IBM Systems Workload Estimator” on page 9; and section 2.3, “Before you begin” on page 15.

If you run firewalls in your environment, ensure that the necessary ports for the IBM Systems Director server are open. A list of the TCP/IP ports that are used by the IBM Systems Director server can be found at the following site:

http://pic.dhe.ibm.com/infocenter/director/pubs/index.jsp?topic=%2Fcom.ibm.director.plan.helps.doc%2Ffqm0_r_ports_for_the_management_server.html

3.1.2 Installation of Linux on x86 systems

For the installation, complete the following checks before the installation of IBM Systems Director. These checks can be completed in any order:

- Sizing

  The IBM Systems Workload Estimator for IBM Systems Director 6.3 is a web-based tool. This tool provides hardware sizing suggestions for systems that run the IBM Systems Director server. Launch the Workload Estimator from the following site:

  http://pic.dhe.ibm.com/infocenter/director/pubs/index.jsp?topic=%2Fcom.ibm.director.plan.helps.doc%2Fwle.html

- Required libraries and packages for Linux on System x

  Table 3-1 lists the required libraries to install IBM Systems Director 6.3.2 on Linux.

  Tip: SLES11 installs all required packages and libraries, by default.

Table 3-1 Required libraries for Linux on x86

Linux distribution	Management server	Common Agent	Platform Agent
RHEL	openssh libstdc++.so.5 libm.so.6 libgcc_s.so.1 libc.so.6 libdl.so.2 libpthread.so.0 unzip libaio libcrypt.so.1 libnsl.so.1 libpam.so.0 librt.so.1 bind-utils net-tools libstdc++.so.6 libuuid.so.1 libexpat.so.0	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 libexpat.so.0	libstdc++.so.5 bind-utils net-tools libpam.so.0 libstdc++.so.6 libuuid.so.1 libcrypt.so.1 unzip libexpat.so.0

Linux distribution	Management server	Common Agent	Platform Agent
SLES10 SLES11	openssh libstdc++.so.5 libm.so.6 libgcc_s.so.1 libc.so.6 libdl.so.2 libpthread.so.0 unzip libaio libcrypt.so.1 libnsl.so.1 libpam.so.0 librt.so.1 bind-utils net-tools libstdc++.so.6 libuuid.so.1 libexpat.so.1	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 libexpat.so.1	libstdc++.so.5 bind-utils net-tools libpam.so.0 libstdc++.so.6 libuuid.so.1 libcrypt.so.1 unzip libexpat.so.1

The libraries that are listed in Table 3-1 on page 38 can be installed in RedHat Linux by installing the following packages:

- `compat-libstdc++-33.i686`
- `libstdc++-4.4.6-4.el6.i686`
- `zlib-1.2.3-27.el6.i686`
- `pam-1.1.1-10.el6_2.1.i686`
- `compat-expat1-1.95.8-8.el6.i686`

In addition, fulfill the following steps on your RedHat installation:

1. Disable SE Linux:

   ```
   # vi /etc/selinux/config
   SELINUX=DISABLED
   ```

2. Disable IP tables:

   ```
   # service iptables stop
   # chkconfig --level 0123456 iptables off
   ```

As a best practice, disable the local firewall (iptables) to avoid heavy configuration with iptables. But if your security standards require that you have iptables up and running, configure it with all of the ports that are listed in section 2.5, “Firewall ports” on page 19.

Detailed steps for installation on Linux for x86

Use the following steps to install the IBM Systems Director server on Linux on x86:

1. Download the installation package from the IBM Systems Director downloads, which are available at the following website:

   http://ibm.com/systems/software/director/downloads/mgmtservers.html

2. Extract the contents of the installation package with the following command:

   `tar -zxvf` *`package_name`*

   The *`package_name`* is the file name of the download packages. Alternatively, you can mount the DVD image to your system.

3. Change to the directory of the installation script. Use the following command to run the pre-installation check:

   ```
   ../checkds/./checkds.sh
   ```

   Reports are generated and results are displayed in the command window or the default browser. For more information, see the `/checkds/readme.txt` file.

   When the result shows no errors (return code= 0) or you can explain the error message that is displayed, continue with the installation.

   In this example, an error code/return code of 34 was returned. The IBM Systems Director server is installed on a virtual system. No baseboard management controller (BMC) or Intelligent Peripheral Management Interface (IPMI) driver is installed. Figure 3-1 on page 41 shows the output from the pre-installation check in the browser.

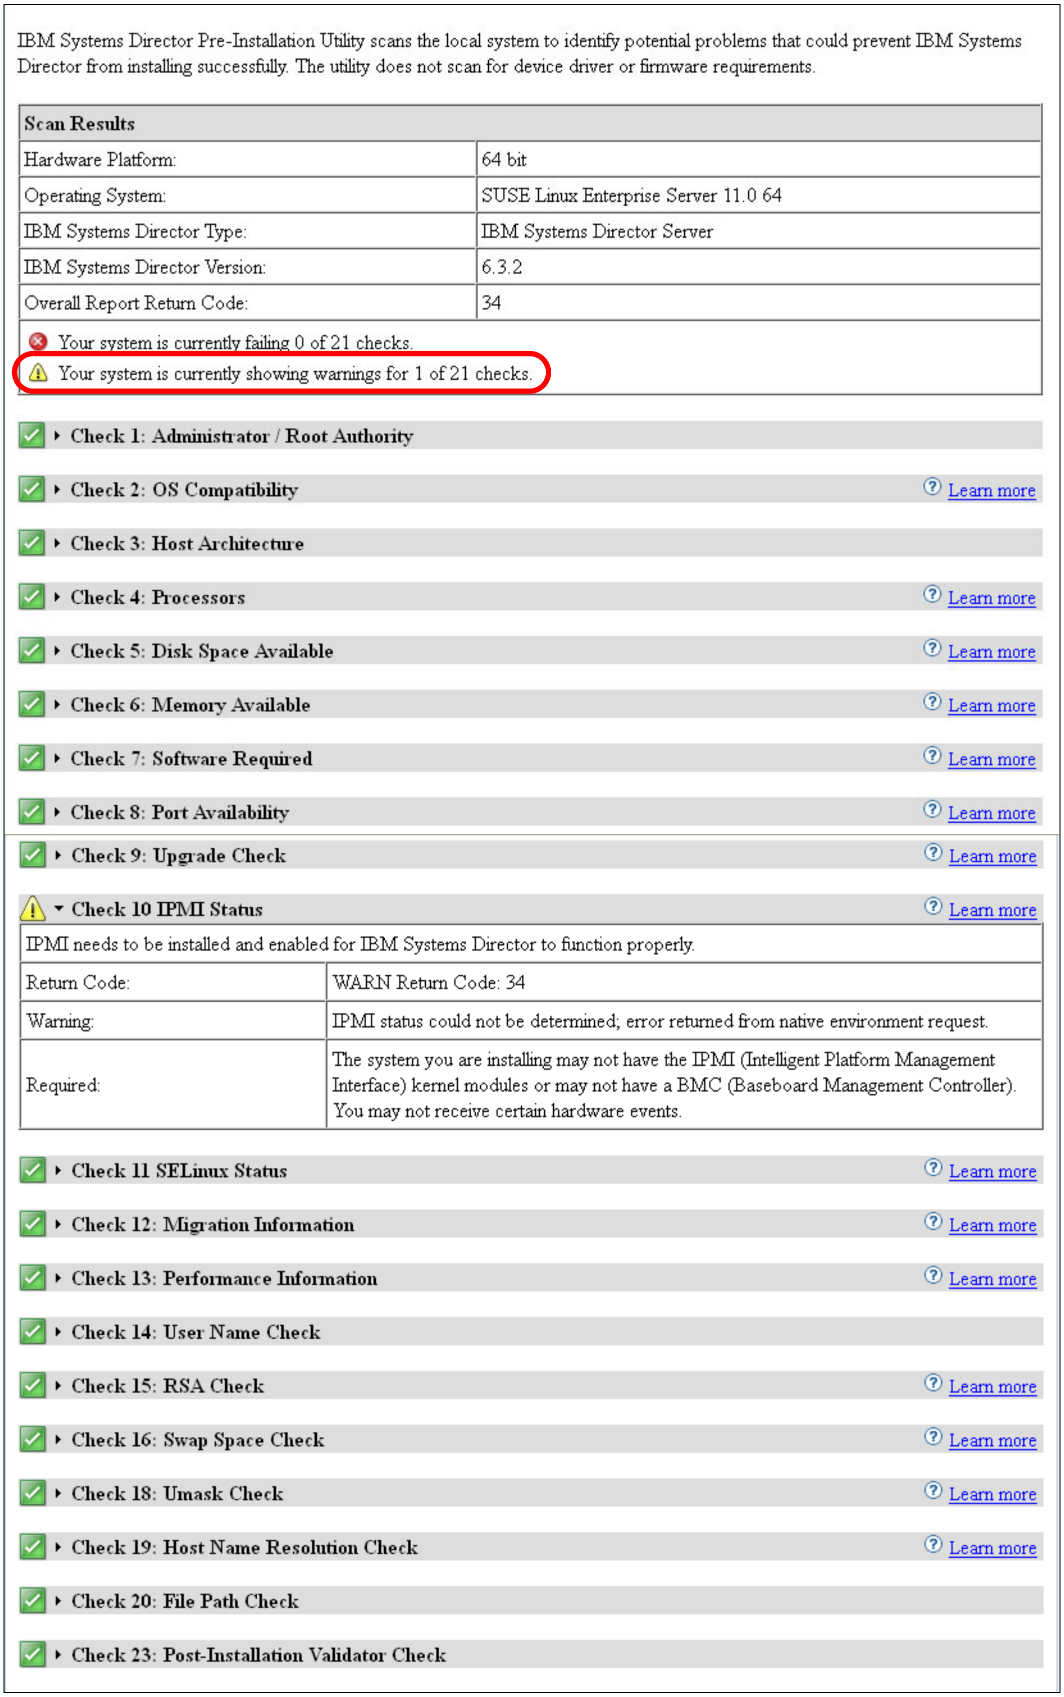

IBM Systems Director Pre-Installation Utility scans the local system to identify potential problems that could prevent IBM Systems Director from installing successfully. The utility does not scan for device driver or firmware requirements.

Scan Results	
Hardware Platform:	64 bit
Operating System:	SUSE Linux Enterprise Server 11.0 64
IBM Systems Director Type:	IBM Systems Director Server
IBM Systems Director Version:	6.3.2
Overall Report Return Code:	34

⊗ Your system is currently failing 0 of 21 checks.

⚠ Your system is currently showing warnings for 1 of 21 checks.

✓ ▸ Check 1: Administrator / Root Authority

✓ ▸ Check 2: OS Compatibility ? Learn more

✓ ▸ Check 3: Host Architecture

✓ ▸ Check 4: Processors ? Learn more

✓ ▸ Check 5: Disk Space Available ? Learn more

✓ ▸ Check 6: Memory Available ? Learn more

✓ ▸ Check 7: Software Required ? Learn more

✓ ▸ Check 8: Port Availability ? Learn more

✓ ▸ Check 9: Upgrade Check ? Learn more

⚠ ▾ Check 10 IPMI Status ? Learn more

IPMI needs to be installed and enabled for IBM Systems Director to function properly.

Return Code:	WARN Return Code: 34
Warning:	IPMI status could not be determined; error returned from native environment request.
Required:	The system you are installing may not have the IPMI (Intelligent Platform Management Interface) kernel modules or may not have a BMC (Baseboard Management Controller). You may not receive certain hardware events.

✓ ▸ Check 11 SELinux Status ? Learn more

✓ ▸ Check 12: Migration Information ? Learn more

✓ ▸ Check 13: Performance Information ? Learn more

✓ ▸ Check 14: User Name Check

✓ ▸ Check 15: RSA Check ? Learn more

✓ ▸ Check 16: Swap Space Check ? Learn more

✓ ▸ Check 18: Umask Check ? Learn more

✓ ▸ Check 19: Host Name Resolution Check ? Learn more

✓ ▸ Check 20: File Path Check

✓ ▸ Check 23: Post-Installation Validator Check

Figure 3-1 Pre-installation check result

If you run your system from the command line, the report is in text format and looks similar to the output in Figure 3-2.

```
Java:
/isd632/standard_linux_x86_Director_base/server/checkds/jvm/xlinux/bin/java

Starting IBM Systems Director Pre-Installation Utility...
Finished analysing system
Creating reports...

Install Readiness Text report being written to
        /tmp/checkds/reports/checkDS_Text_20121214_174936.txt
Install Readiness Error Text report being written to
        /tmp/checkds/reports/checkDS_Error.txt
Install Readiness Detailed HTML report being written to
        /tmp/checkds/reports/checkDS_Detailed_20121214_174936.html
Install Readiness Summary HTML report being written to
        /tmp/checkds/reports/checkDS_Summary_20121214_174937.html


Your system is currently showing warnings for 1 of 21 checks.

WARN Check 10 IPMI Status

IPMI status could not be determined; error returned from native environment request.

The system you are installing may not have the IPMI (Intelligent Platform Management
Interface) kernel modules or maynot have a BMC (Baseboard Management Controller).

You may not receive certain hardware events.

Overall Report Return Code: 34
```

Figure 3-2 Text output from the pre-installation utility

If recommendations or errors exist, you must address them before you can continue the installation. For information about the problems, see the report. After the problems are fixed, run the pre-installation check again.

4. To install the IBM Systems Director server, from within the directory of the installation script, use one of the following commands:

   – To accept the default settings, enter this command:

     `./dirinstall.server`

   – To use the response file, enter this command:

     `./dirinstall.server -r` */directory/response.rsp*

     The *directory* is the local directory to which you copied the response file and *response.rsp* is the name of the response file.

   – To force a clean installation, regardless of the existing data, enter this command:

     `./dirinstall.server -g`

Tip: If you previously installed IBM Systems Director on this system, data is saved in the `/var/tmp/director_save_630` directory, by default. The data is not removed even if you uninstall the previous installation. If you want a clean installation, use **`./dirinstall.server -g`** to ensure that you do not inadvertently migrate this data. IBM Systems Director 6.3.*x* installs cleanly and the data from the previous installation is preserved.

5. The installation runs now with the default setting or with the settings from the response file. If you use the default integrated DB2 database, the installation automatically creates the settings to use DB2. If you use another supported database, you must configure the database for use with IBM Systems Director.

   For information about how to configure these databases, go to the following site:

   http://pic.dhe.ibm.com/infocenter/director/pubs/index.jsp?topic=%2Fcom.ibm.director.configdir.helps.doc%2Ffqm0_t_config_database_application_cfgdbcmd.html

6. When the license agreement displays, confirm that you accept this agreement by entering `1`. The Pre-Installation Utility runs. In Figure 3-3, return code 34 is returned, indicating that no IPMI driver is installed. However, because the installation occurs in a virtual environment, this is not a problem. Enter `1` to continue as shown in Figure 3-3.

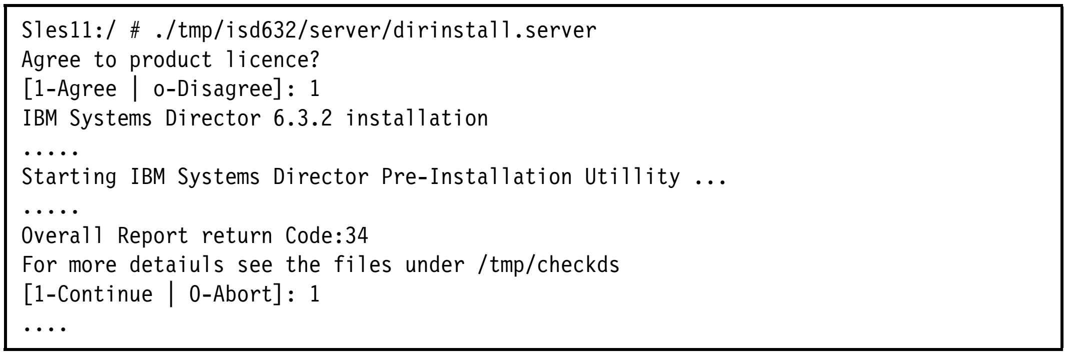

```
Sles11:/ # ./tmp/isd632/server/dirinstall.server
Agree to product licence?
[1-Agree | o-Disagree]: 1
IBM Systems Director 6.3.2 installation
.....
Starting IBM Systems Director Pre-Installation Utillity ...
.....
Overall Report return Code:34
For more detaiuls see the files under /tmp/checkds
[1-Continue | 0-Abort]: 1
....
```

Figure 3-3 Running dirinstall.server script

7. The Director server, components, features, and embedded DB2 database are installed. When the installation completes, a message appears similar to the message shown in Figure 3-4.

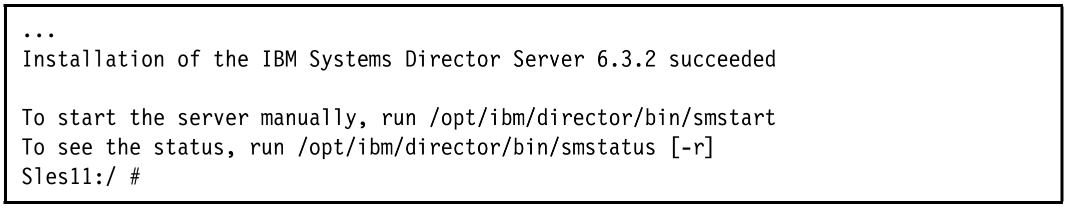

```
...
Installation of the IBM Systems Director Server 6.3.2 succeeded

To start the server manually, run /opt/ibm/director/bin/smstart
To see the status, run /opt/ibm/director/bin/smstatus [-r]
Sles11:/ #
```

Figure 3-4 Completing the IBM Systems Director installation

8. After the installation completes, configure the Agent Manager and then start the IBM Systems Director.

Important: Do not start the IBM Systems Director before you configure the Agent Manager.

9. To configure the Agent Manager (if you did not configure it during the installation process), run the following command:

   ```
   install_root/bin/configAgtMgr.sh
   ```

10. Respond to the `configAgtMgr.sh` script prompts:

    – Agent Manager

      Enter `1` to use the Agent Manager that is installed with this server (suggested), or enter `0` to use an existing Agent Manager (advanced).

    – Resource Manager

      Enter the Resource Manager user ID that you want to set for the Agent Manager. The user ID does not need to be an operating system user ID. Remember this user ID. If you want to use the same Agent Manager with another system, you need this user ID.

      Enter and verify the Resource Manager password to set for the Agent Manager.

    – Agent Registration password

      Enter and verify the Agent Registration password to set for your Agent Manager. This password can be the same password for the Agent Manager. This password is used to register the Common Agents with Agent Manager.

    – IP address and port for Agent Manager:

      - Enter the IP address for the existing Agent Manager.

        If you selected 0 (use an existing Agent Manager), you must provide the IP address of the existing Agent Manager.

      - Enter the port number for the existing Agent Manager.

        If you selected 0 (use an existing Agent Manager), you must provide the port number of the existing Agent Manager. The port number must be a valid number 0 - 65535.

11. Start IBM Systems Director processes on the management servers by running the **`smstart`** command:

    ```
    install_root/bin/smstart
    ```

12. To check the status of the IBM Systems Director, run the following command:

    ```
    install_root/bin/smstatus -r
    ```

    When this command returns a value of `Active`, the server is started and active.

3.1.3 Installation on Windows

Use the following steps to install the IBM Systems Director server on a Windows server:

1. Download the installation package from the following link and decompress it:

   http://ibm.com/systems/software/director/downloads/mgmtservers.html

2. Double-click the `IBMSystemsDirectorServerSetup64.exe` file to start the installation process.

3. The pre-installation check runs. If the check runs successfully, a green mark and the "`No warnings or errors were found`" message displays (Figure 3-5). Click **Next** to continue. If problem and error messages appear, fix them, return to this window, and run the installation program again.

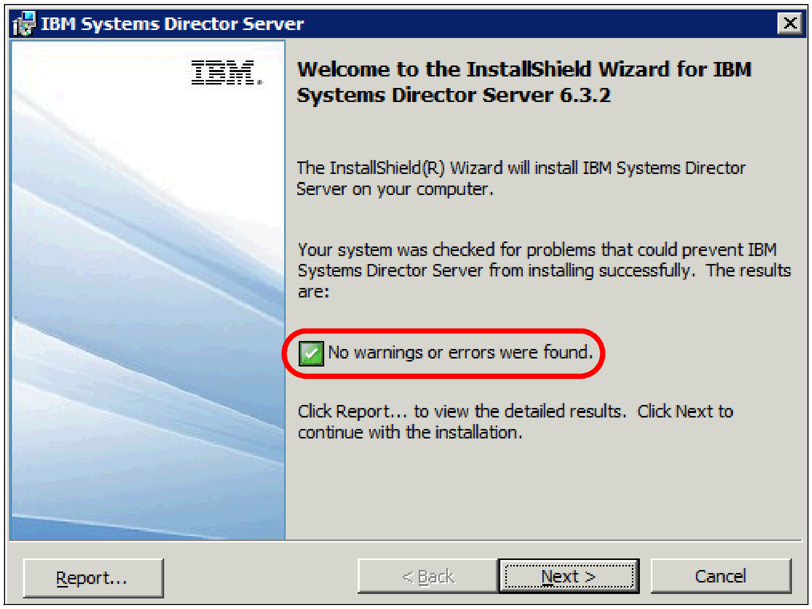

IBM Systems Director Server

Welcome to the InstallShield Wizard for IBM Systems Director Server 6.3.2

The InstallShield(R) Wizard will install IBM Systems Director Server on your computer.

Your system was checked for problems that could prevent IBM Systems Director Server from installing successfully. The results are:

✓ No warnings or errors were found.

Click Report... to view the detailed results. Click Next to continue with the installation.

Figure 3-5 Welcome to the InstallShield wizard for IBM Systems Director server 6.3.2

4. Agree to the license agreement and click **Next** to continue.

5. Specify the folder where you want to install the software (Figure 3-6) and click **Next** to continue.

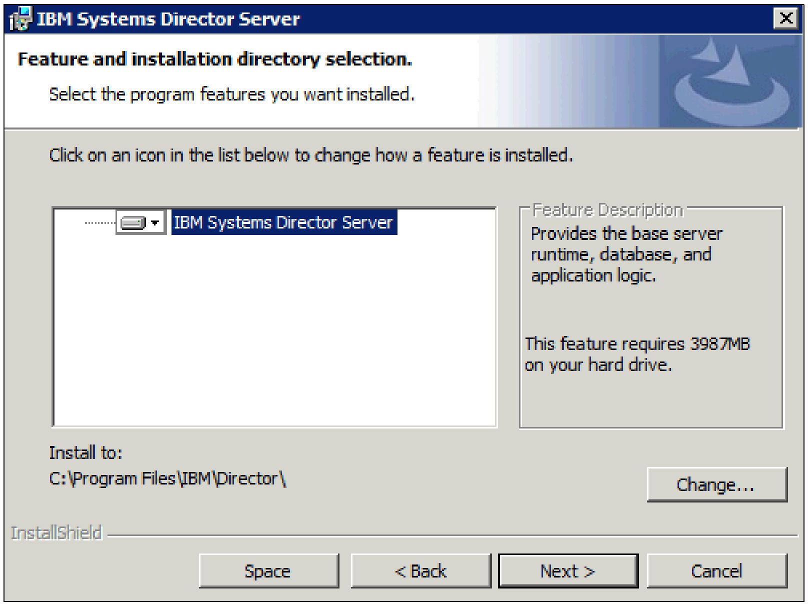

IBM Systems Director Server

Feature and installation directory selection.

Select the program features you want installed.

Click on an icon in the list below to change how a feature is installed.

IBM Systems Director Server

Feature Description

Provides the base server runtime, database, and application logic.

This feature requires 3987MB on your hard drive.

Install to:
C:\Program Files\IBM\Director\

Figure 3-6 Feature and installation directory selection

6. In the next window (Figure 3-7), select the installation type. Two types are available:

   – Click **Basic** to use the embedded database and default ports and install the Common Agent Services (CAS) server with the Director installation. You type the user ID and password only one time. The installation program uses the user ID and password for all settings.

   – Click **Advanced** if you want to use a database other than the embedded, managed DB2. Also, if you install a second IBM Systems Director server in your environment, use the Advanced setup to select the existing CAS server. With the Advanced setup, you can define a different user ID and password for the CAS server. With the Advanced setup, you can use different default ports for the IBM Systems Director server.

   This installation uses the Basic setup type. Click **Next** to continue.

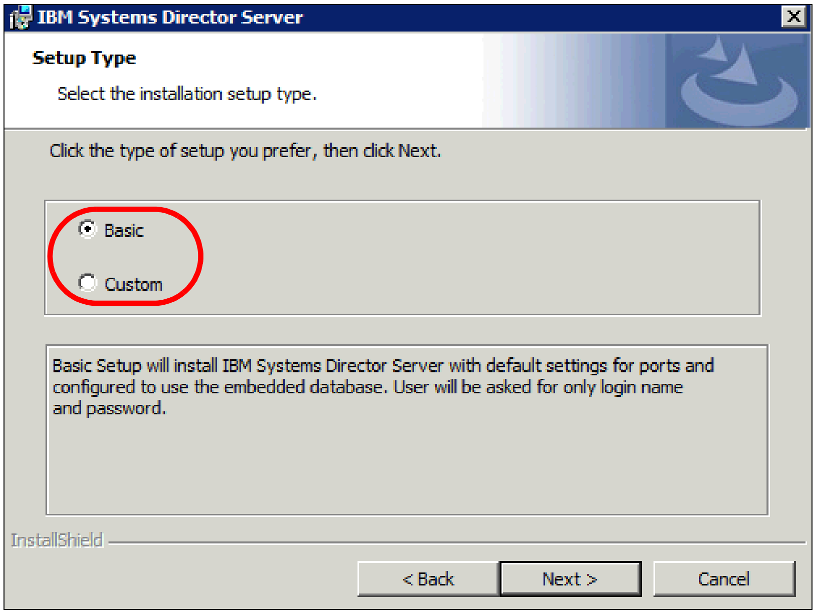

IBM Systems Director Server

Setup Type

Select the installation setup type.

Click the type of setup you prefer, then click Next.

● Basic

○ Custom

Basic Setup will install IBM Systems Director Server with default settings for ports and configured to use the embedded database. User will be asked for only login name and password.

Figure 3-7 Setup type

7. Type in the credentials that you want to use for the IBM Systems Director server: computer name, user name, and password (Figure 3-8). Click **Next** to continue.

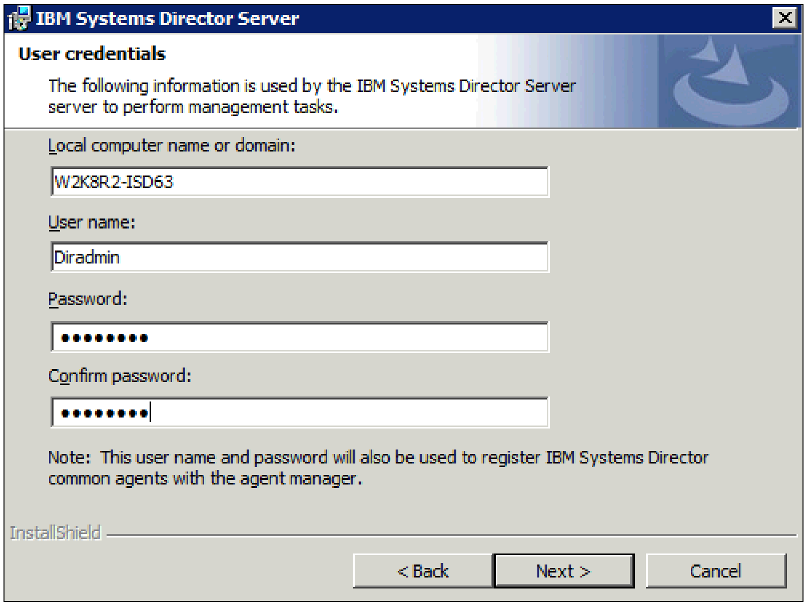

IBM Systems Director Server

User credentials

The following information is used by the IBM Systems Director Server server to perform management tasks.

Local computer name or domain: W2K8R2-ISD63

User name: Diradmin

Password: ••••••••

Confirm password: ••••••••

Note: This user name and password will also be used to register IBM Systems Director common agents with the agent manager.

Figure 3-8 User credentials

8. When you are ready to begin the installation process, click **Install** (Figure 3-9).

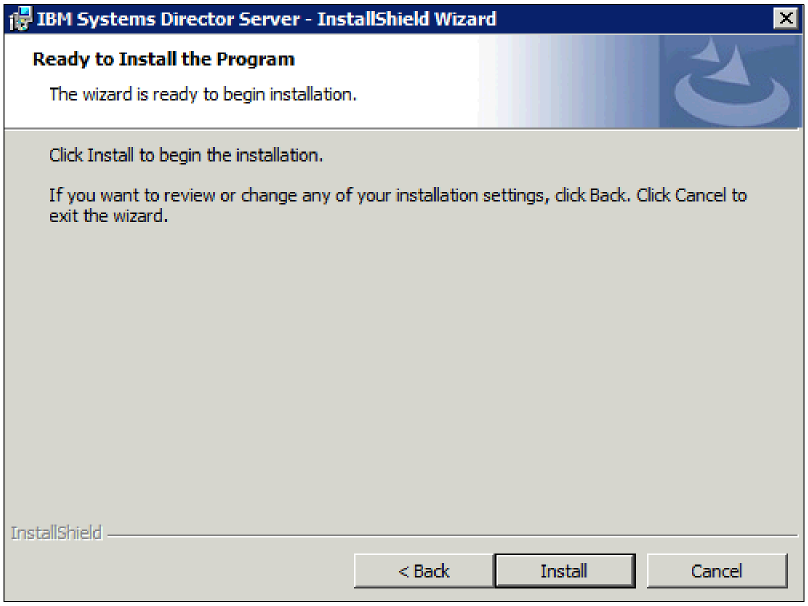

IBM Systems Director Server - InstallShield Wizard

Ready to Install the Program

The wizard is ready to begin installation.

Click Install to begin the installation.

If you want to review or change any of your installation settings, click Back. Click Cancel to exit the wizard.

Figure 3-9 Ready to Install the Program window

9. The IBM Systems Director server database is installed. You can see the progress of the installation process (Figure 3-10).

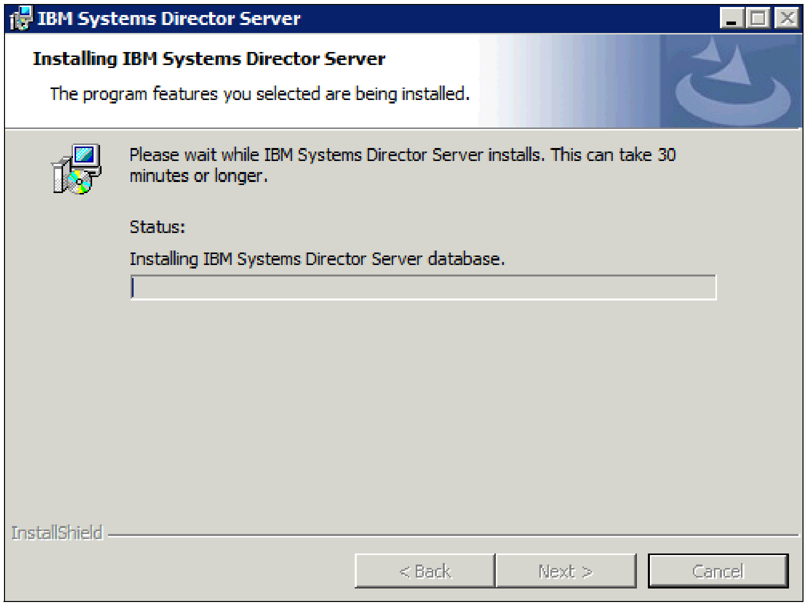
IBM Systems Director Server

Installing IBM Systems Director Server

The program features you selected are being installed.

Please wait while IBM Systems Director Server installs. This can take 30 minutes or longer.

Status:

Installing IBM Systems Director Server database.

Figure 3-10 Installing the IBM Systems Director server

10.The IBM Systems Director Common Agent is installed next (Figure 3-11).

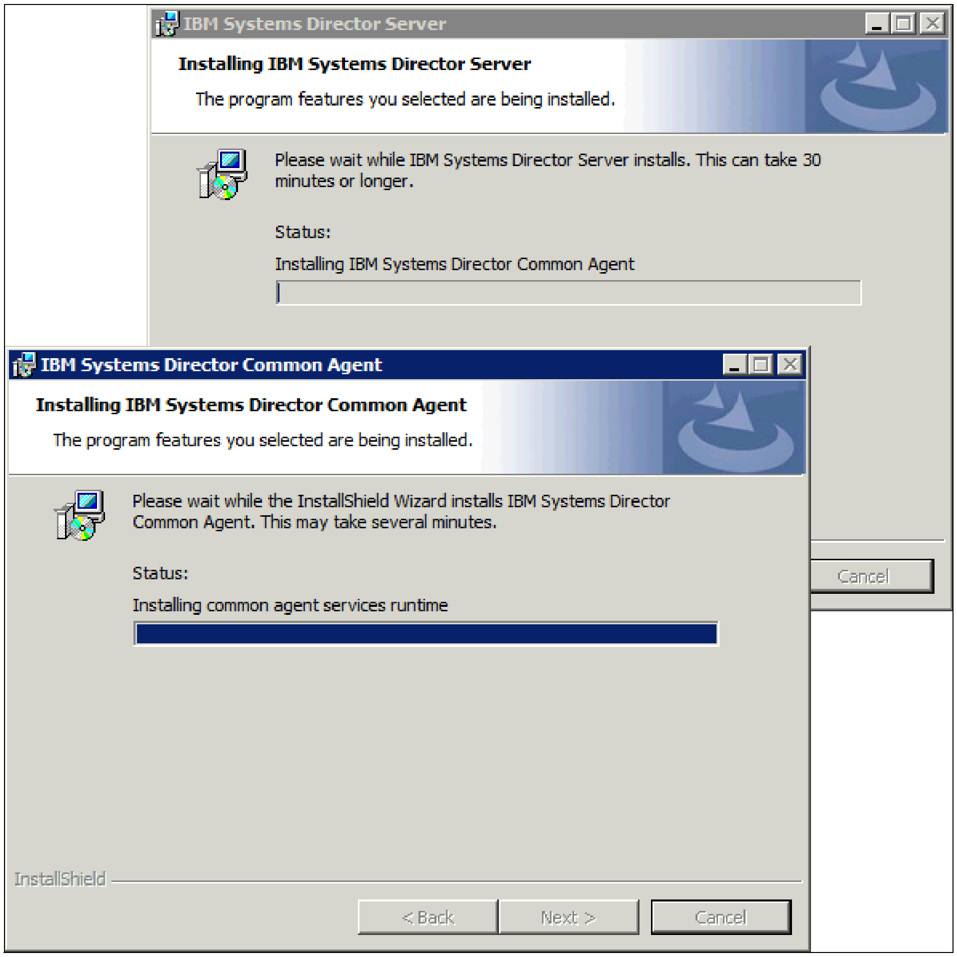

IBM Systems Director Server

Installing IBM Systems Director Server

The program features you selected are being installed.

Please wait while IBM Systems Director Server installs. This can take 30 minutes or longer.

Status:

Installing IBM Systems Director Common Agent

IBM Systems Director Common Agent

Installing IBM Systems Director Common Agent

The program features you selected are being installed.

Please wait while the InstallShield Wizard installs IBM Systems Director Common Agent. This may take several minutes.

Status:

Installing common agent services runtime

Figure 3-11 Install Common Agent and Common Agent services

11. The IBM Systems Director Platform Agent packages are installed (Figure 3-12).

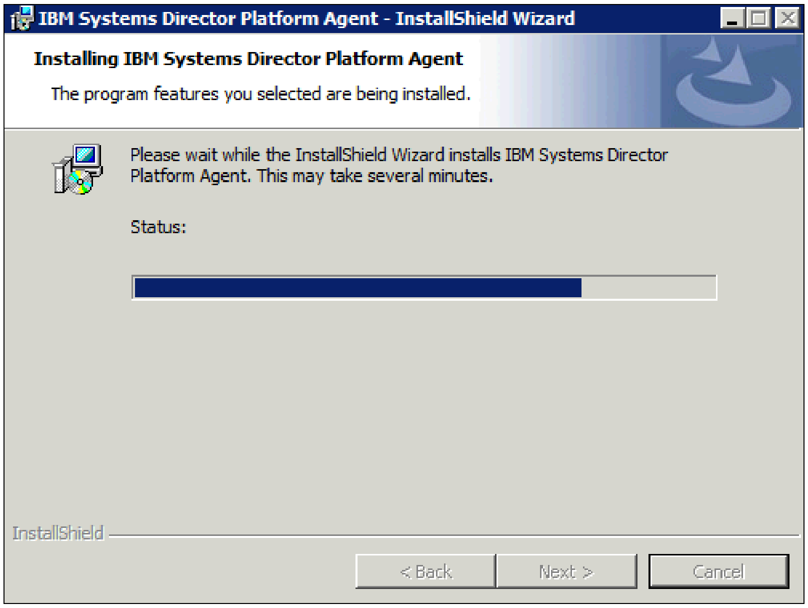

Installing IBM Systems Director Platform Agent

The program features you selected are being installed.

Please wait while the InstallShield Wizard installs IBM Systems Director Platform Agent. This may take several minutes.

Status:

Figure 3-12 Platform Agent installation

12. The files for the IBM Systems Director server are installed (Figure 3-13).

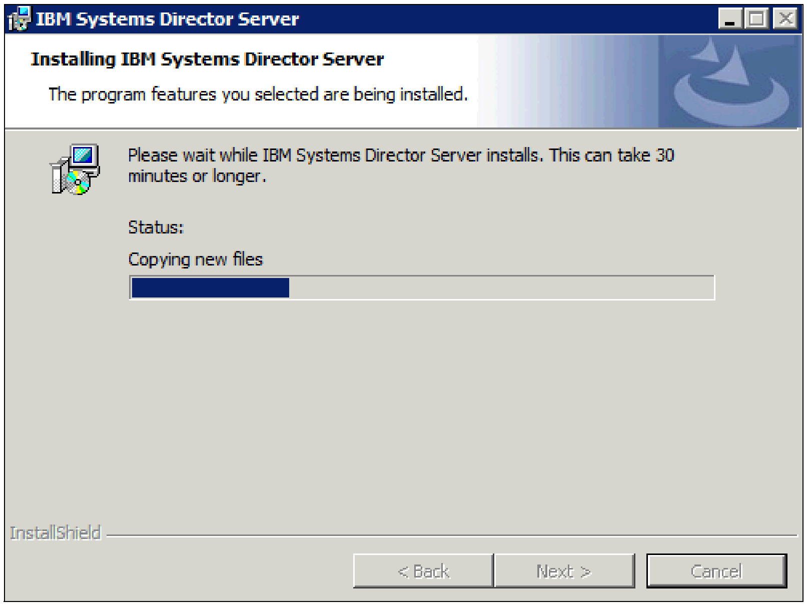

Installing IBM Systems Director Server

The program features you selected are being installed.

Please wait while IBM Systems Director Server installs. This can take 30 minutes or longer.

Status:

Copying new files

Figure 3-13 Installing the IBM Systems Director server

13. Additional features and plug-ins are installed (Figure 3-14).

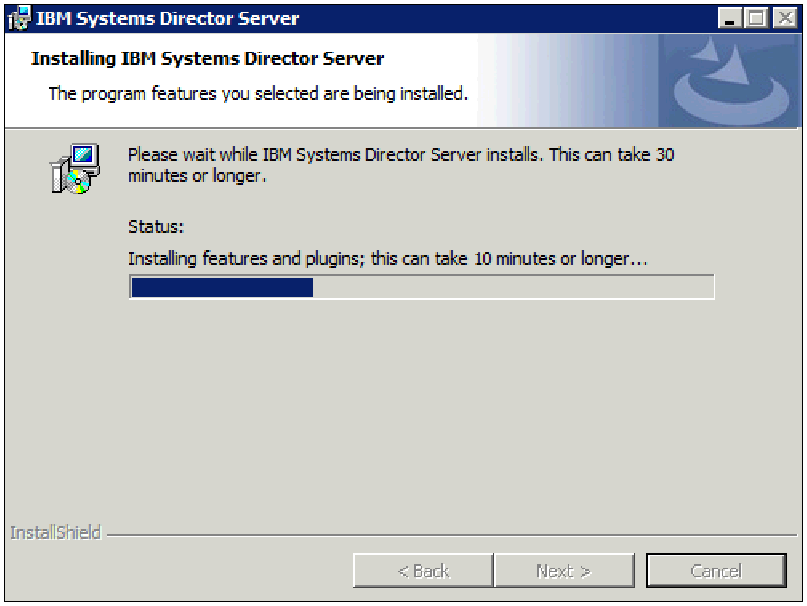

IBM Systems Director Server

Installing IBM Systems Director Server

The program features you selected are being installed.

Please wait while IBM Systems Director Server installs. This can take 30 minutes or longer.

Status:

Installing features and plugins; this can take 10 minutes or longer...

Figure 3-14 Installing features and plug-ins

14. The Agent Manager (CAS server) is installed (Figure 3-15).

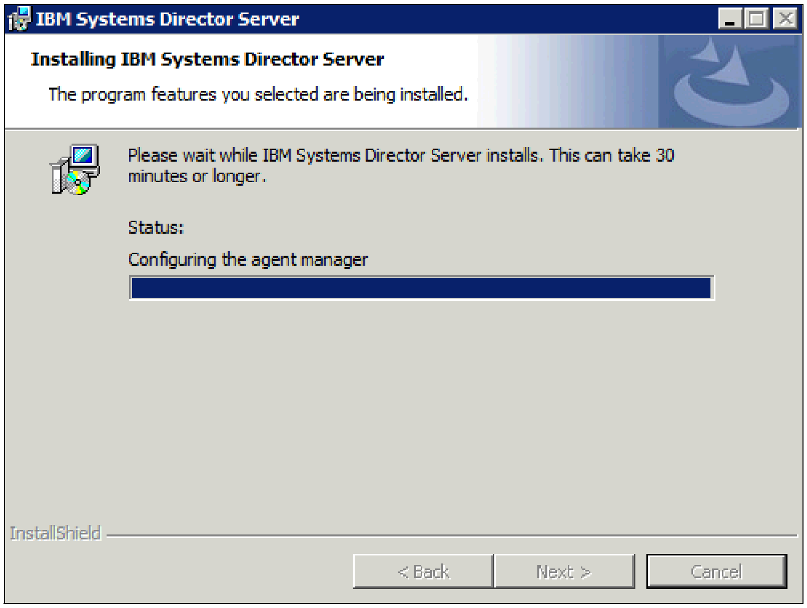

IBM Systems Director Server

Installing IBM Systems Director Server

The program features you selected are being installed.

Please wait while IBM Systems Director Server installs. This can take 30 minutes or longer.

Status:

Configuring the agent manager

Figure 3-15 Installing Agent Manager

15. When the installation completes, the InstallShield Wizard Completed window displays (Figure 3-16). You can view the Windows Installer log. Complete the installation by clicking **Finish**.

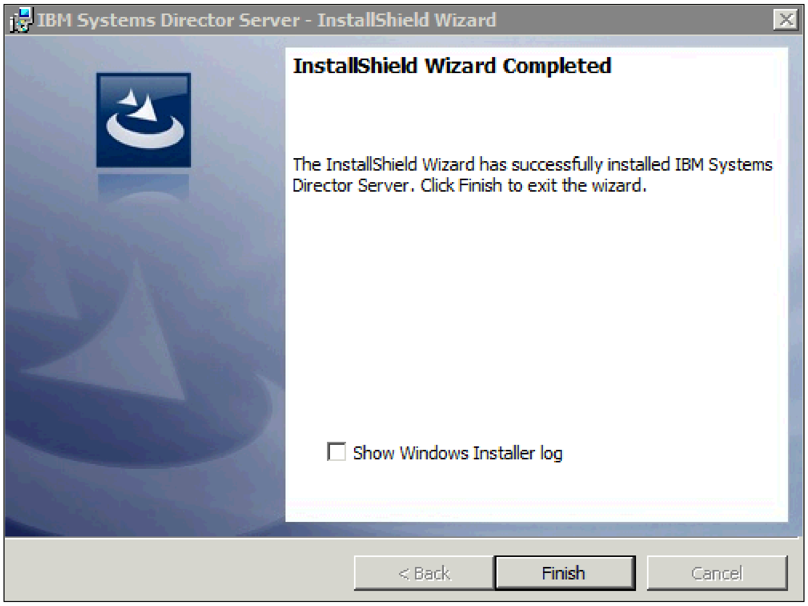
IBM Systems Director Server - InstallShield Wizard
InstallShield Wizard Completed

The InstallShield Wizard has successfully installed IBM Systems Director Server. Click Finish to exit the wizard.

□ Show Windows Installer log

Figure 3-16 InstallShield Wizard Completed window

After the installation is finished, the IBM Systems Director server starts automatically. You can check the status of the IBM Systems Director server through the status icon or by using the **smstatus -r** command. When the status icon shows a green circle or the status shows as active, the IBM Systems Director server is up and running.

3.1.4 Post Installation Validation tool

With IBM Systems Director 6.3.2, IBM provides the new Post Installation Validation (PIV) tool. This tool can be run after the installation process to check the installation for completeness and errors. The tool is in the `/piv` folder in the installation medium or directory.

The PIV analyzes the installation logs for errors and checks for services, ports, Agent Manager configuration, and database configuration. The PIV also checks whether the server is active by using the **smstatus** command.

The PIV tool is small (less than 5 MB). The PIV tool is written in Python and includes a small Python interpreter. (Python must be installed on the system to run PIV.) PIV is not a health checker for a running IBM Systems Director; it is only a tool to verify the installation.

The tool is a command-line tool and is run in the following way:

- Linux on x86:

```
<Install_directory>\bin\piv> .\PostInstallValidator_xLin.sh
```

- Windows:

```
<Install_directory>\bin\piv> .\PostInstallValidator_Win.exe
```

Command options are available. You can see the full list of command-line options by using the **-h** option.

The following command options are the most important options:

-o or **--output**	Specify the location of the installation report
-c or **--config**	Specify the location of the configuration file
-s or **--silent**	Run the tool silently
-r or **--report**	Open the text report on completion (Windows only)
-j or **--nohtmlreport**	Do not create an HTML report
-d or **--detailed**	Include detailed information in the report
-w or **--wait**	Wait to return until the installation is completed

The tool generates a report. If a graphical user interface (GUI) is available, you receive an HTML report (if not, clear by using the **-j** option). If only a command-line environment is available, you receive a text report.

The reports are in the following directories, by default:

- Linux/AIX reports are in the `/tmp` directory.
- Windows reports are in the `%temp%` directory.

The PIV HTML report looks similar to the example of the installation in Windows Server 2008 R2 x64 (Figure 3-17 on page 54).

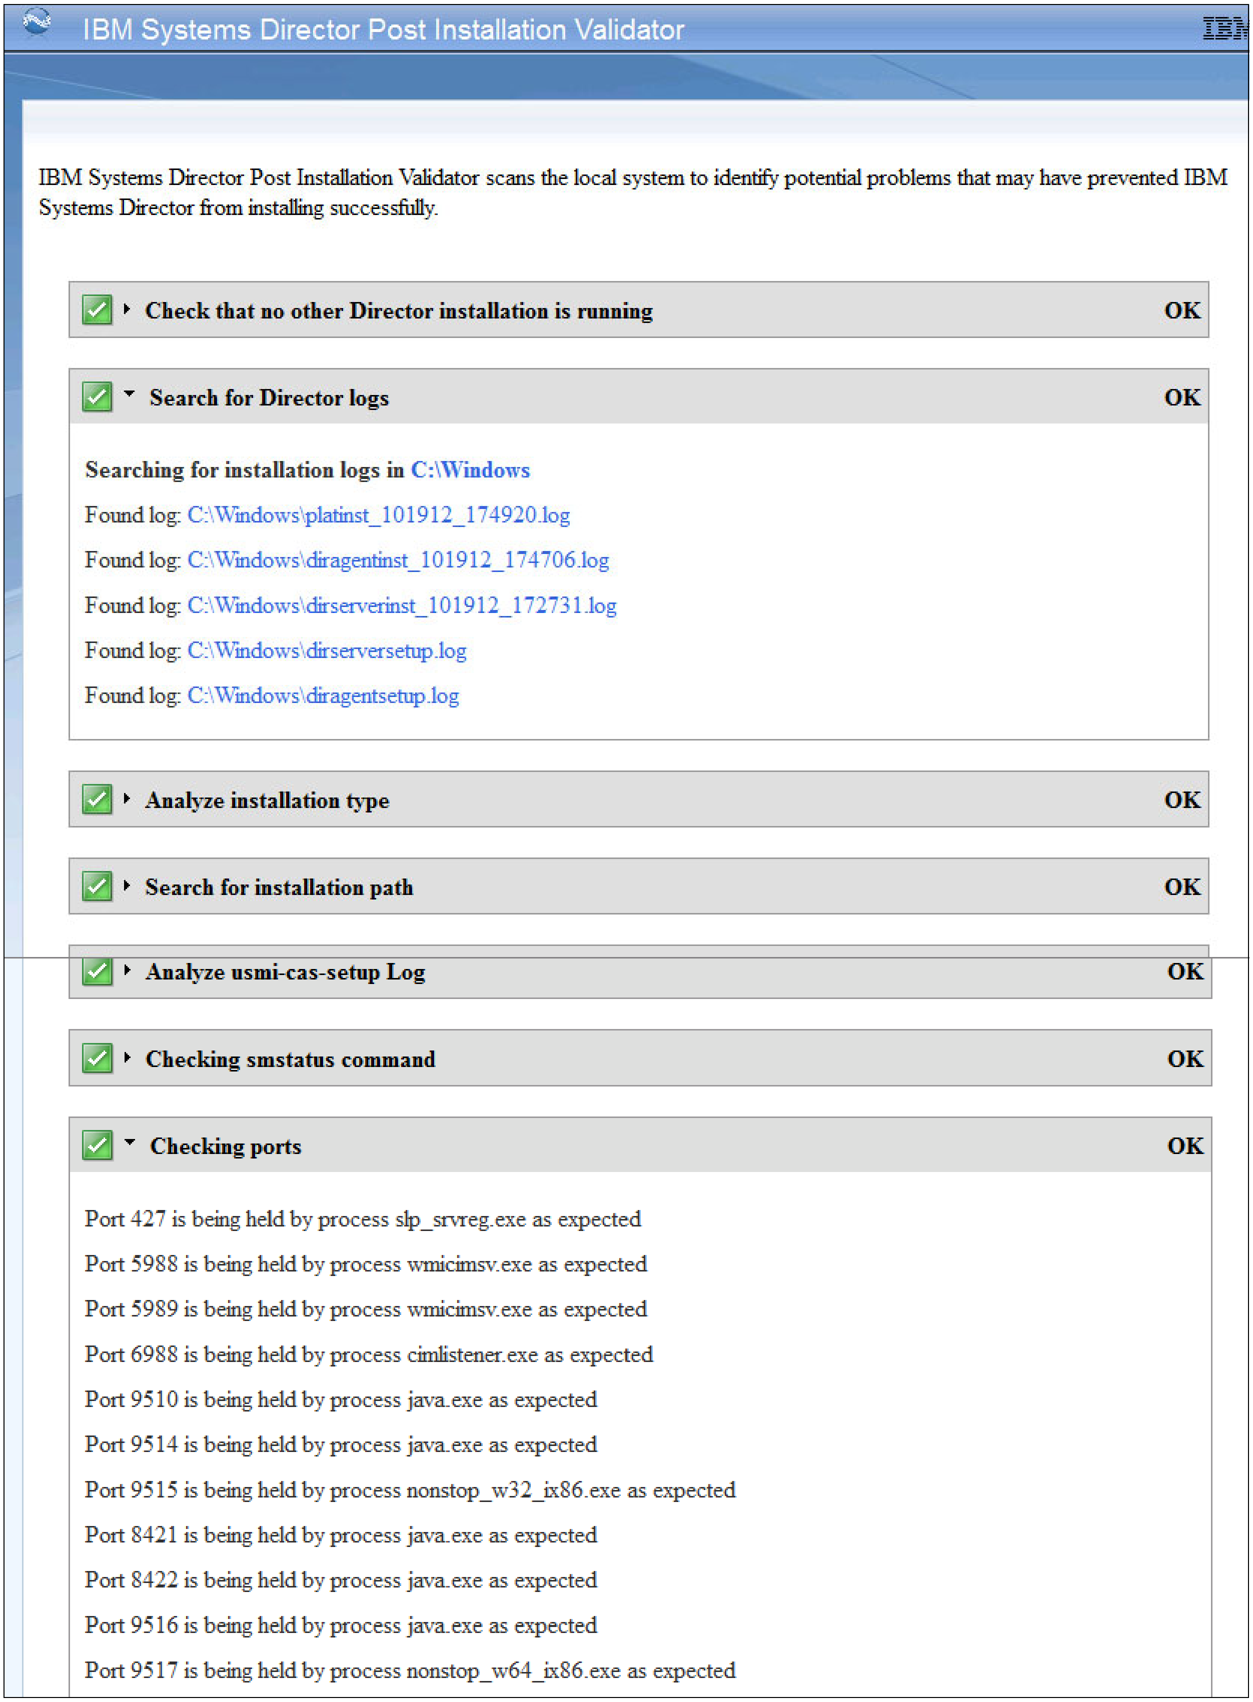

IBM Systems Director Post Installation Validator

IBM Systems Director Post Installation Validator scans the local system to identify potential problems that may have prevented IBM Systems Director from installing successfully.

✓ ▸ **Check that no other Director installation is running** **OK**

✓ ▾ **Search for Director logs** **OK**

Searching for installation logs in C:\Windows

Found log: C:\Windows\platinst_101912_174920.log

Found log: C:\Windows\diragentinst_101912_174706.log

Found log: C:\Windows\dirserverinst_101912_172731.log

Found log: C:\Windows\dirserversetup.log

Found log: C:\Windows\diragentsetup.log

✓ ▸ **Analyze installation type** **OK**

✓ ▸ **Search for installation path** **OK**

✓ ▸ **Analyze usmi-cas-setup Log** **OK**

✓ ▸ **Checking smstatus command** **OK**

✓ ▾ **Checking ports** **OK**

Port 427 is being held by process slp_srvreg.exe as expected

Port 5988 is being held by process wmicimsv.exe as expected

Port 5989 is being held by process wmicimsv.exe as expected

Port 6988 is being held by process cimlistener.exe as expected

Port 9510 is being held by process java.exe as expected

Port 9514 is being held by process java.exe as expected

Port 9515 is being held by process nonstop_w32_ix86.exe as expected

Port 8421 is being held by process java.exe as expected

Port 8422 is being held by process java.exe as expected

Port 9516 is being held by process java.exe as expected

Port 9517 is being held by process nonstop_w64_ix86.exe as expected

Figure 3-17 PIV HTML report

The text report looks similar to the report shown in Figure 3-18.

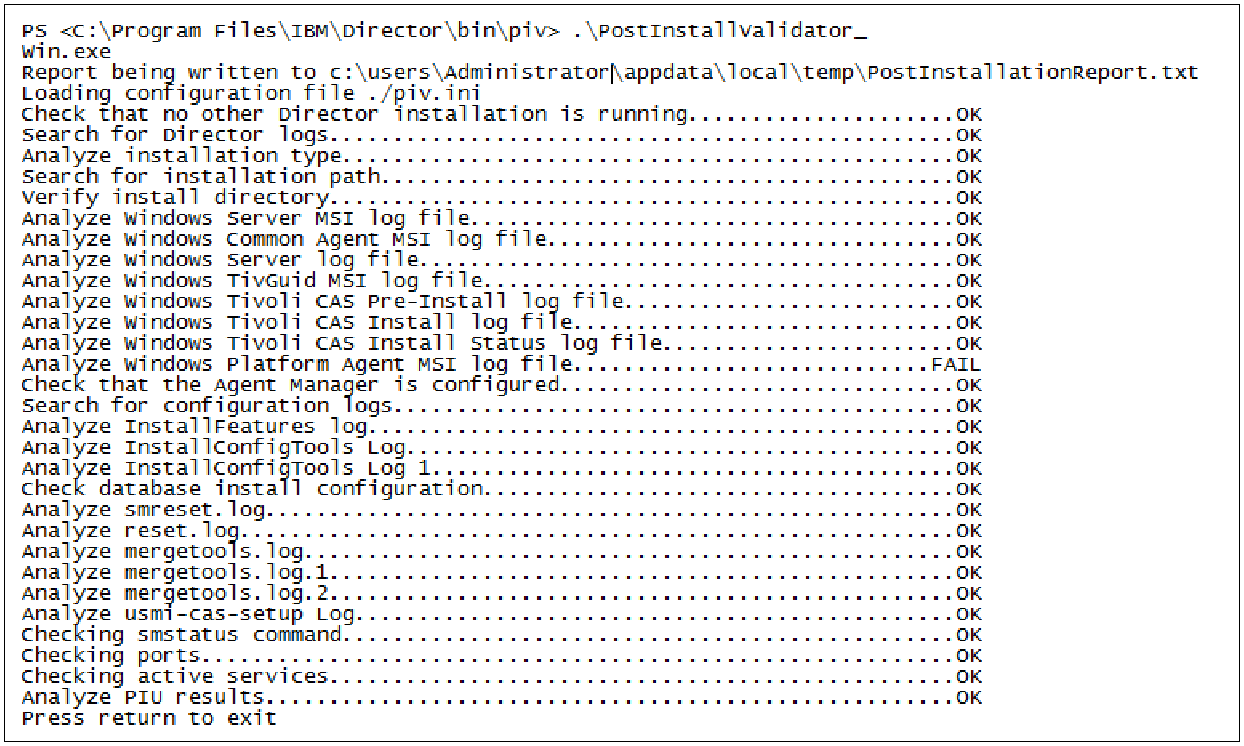

```
PS <C:\Program Files\IBM\Director\bin\piv> .\PostInstallValidator_
win.exe
Report being written to c:\users\Administrator\appdata\local\temp\PostInstallationReport.txt
Loading configuration file ./piv.ini
Check that no other Director installation is running.....................OK
Search for Director logs.................................................OK
Analyze installation type................................................OK
Search for installation path.............................................OK
Verify install directory.................................................OK
Analyze Windows Server MSI log file......................................OK
Analyze Windows Common Agent MSI log file................................OK
Analyze Windows Server log file..........................................OK
Analyze Windows TivGuid MSI log file.....................................OK
Analyze Windows Tivoli CAS Pre-Install log file..........................OK
Analyze Windows Tivoli CAS Install log file..............................OK
Analyze Windows Tivoli CAS Install Status log file.......................OK
Analyze Windows Platform Agent MSI log file............................FAIL
Check that the Agent Manager is configured...............................OK
Search for configuration logs............................................OK
Analyze InstallFeatures log..............................................OK
Analyze InstallConfigTools Log...........................................OK
Analyze InstallConfigTools Log 1.........................................OK
Check database install configuration.....................................OK
Analyze smreset.log......................................................OK
Analyze reset.log........................................................OK
Analyze mergetools.log...................................................OK
Analyze mergetools.log.1.................................................OK
Analyze mergetools.log.2.................................................OK
Analyze usmi-cas-setup Log...............................................OK
Checking smstatus command................................................OK
Checking ports...........................................................OK
Checking active services.................................................OK
Analyze PIU results......................................................OK
Press return to exit
```

Figure 3-18 PIV text report

3.1.5 Starting IBM Systems Director

The IBM Systems Director automatically starts after the installation. No reboot of the system is necessary. You can start and stop the IBM Systems Director server using the command line. The command line is the best method because all necessary services are started and stopped in the correct sequence.

Use the following commands to start or stop IBM Systems Director:

- Linux:
  - **`smstart`** to start the IBM Systems Director server
  - **`smstop`** to stop the IBM Systems Director server
- Windows
  - **`net start dirserver`** to start the IBM Systems Director server
  - **`net stop dirserver`** to stop the IBM Systems Director server

Initial logon

After the IBM Systems Director server is in the active status, use the following steps to log on to the IBM Systems Director web interface.

1. Open your browser and type in the following address:

   `http://hostname_or_IP_address:8421/ibm/console/logon.jsp`

   or

   `https://hostname_or_IP_sddress:8422/ibm/console/logon.jsp`

2. At the first access, a window opens to show you that the connection is untrusted (Figure 3-19). Click **Add Exception**.

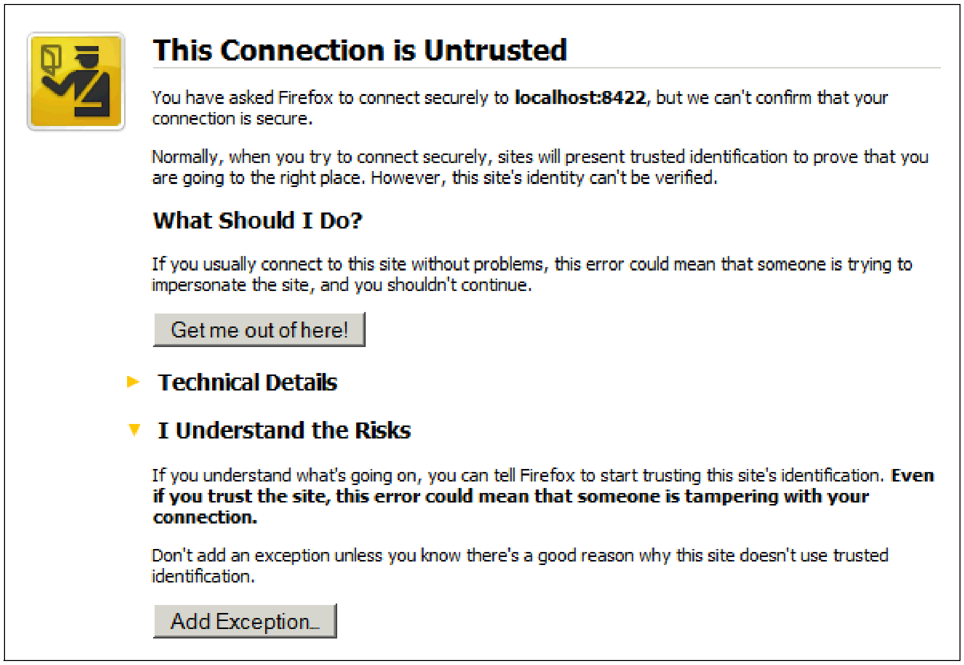

This Connection is Untrusted

You have asked Firefox to connect securely to **localhost:8422**, but we can't confirm that your connection is secure.

Normally, when you try to connect securely, sites will present trusted identification to prove that you are going to the right place. However, this site's identity can't be verified.

What Should I Do?

If you usually connect to this site without problems, this error could mean that someone is trying to impersonate the site, and you shouldn't continue.

If you understand what's going on, you can tell Firefox to start trusting this site's identification. **Even if you trust the site, this error could mean that someone is tampering with your connection.**

Don't add an exception unless you know there's a good reason why this site doesn't use trusted identification.

Figure 3-19 Untrusted connection

3. A new window opens (Figure 3-20). Click **Confirm Security Exception**.

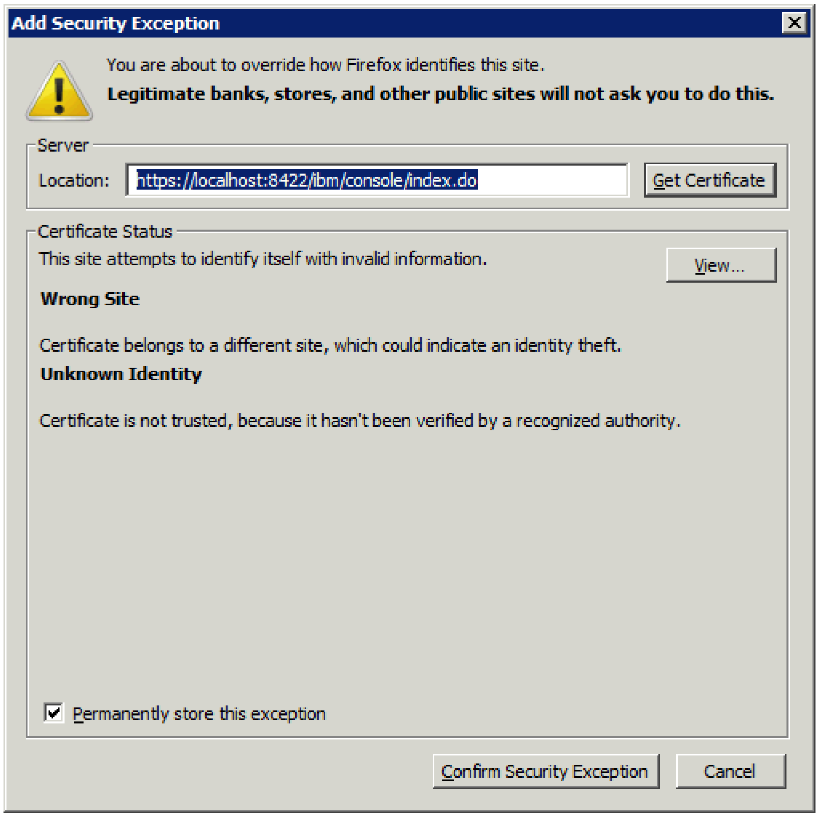

You are about to override how Firefox identifies this site.

Legitimate banks, stores, and other public sites will not ask you to do this.

Server

Location: https://localhost:8422/ibm/console/index.do

Certificate Status

This site attempts to identify itself with invalid information.

Wrong Site

Certificate belongs to a different site, which could indicate an identity theft.

Unknown Identity

Certificate is not trusted, because it hasn't been verified by a recognized authority.

☑ Permanently store this exception

Figure 3-20 Confirm security exception

4. The login window displays (Figure 3-21). Enter your user ID and password and click **Log in**.

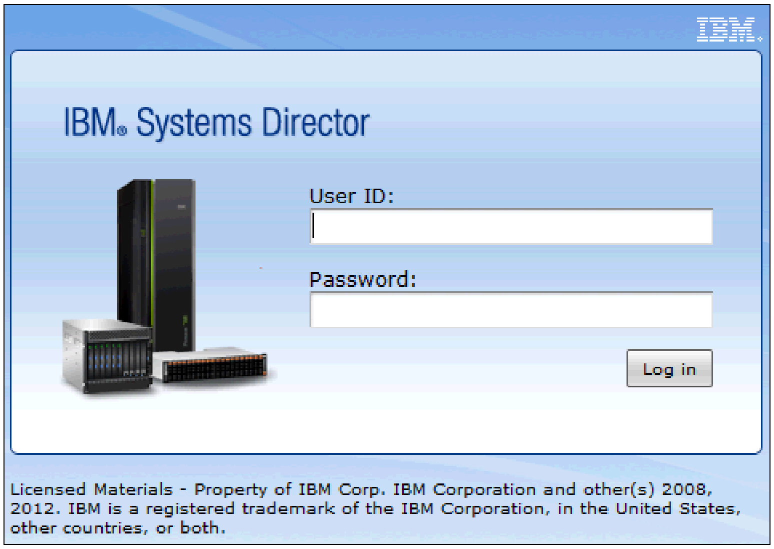

Licensed Materials - Property of IBM Corp. IBM Corporation and other(s) 2008, 2012. IBM is a registered trademark of the IBM Corporation, in the United States, other countries, or both.

Figure 3-21   Logon window

3.2 Installing IBM Systems Director server on an AIX platform

Useful information is provided about the process of installing IBM Systems Director on an AIX platform. Resources that are required for running the IBM Systems Director server are referenced in section 2.1, “System resources” on page 8, and section 2.2, “IBM Systems Workload Estimator” on page 9.

For the installation of IBM Systems Director on AIX, go to the following site:

http://publib.boulder.ibm.com/infocenter/director/pubs/topic/com.ibm.director.install.helps.doc/fqm0_t_installing.html

The IBM Systems Director Management Server code can be sourced from this site:

http://www-03.ibm.com/systems/software/director/downloads/mgmtservers.html

3.2.1 Downloading the software

The IBM Systems Director source can be downloaded in two formats: an .iso file or a GZIP (.gz) file. Download your preferred file type and place it in a temporary file system on the AIX server using one of the following commands:

- ISO:

```
loopmount -i express_aix_Director_base.iso -o “-V cdrfs -o rw” -m /mnt
```

- GZIP:

```
gunzip -c express_aix_Director_base.tar.gz | tar -xvf -
```

The installation of the IBM Systems Director server comes with an embedded pre-installation check utility. This option is enabled, by default, and is referenced in the `dirserv.rsp` file. The suggestion is to leave the pre-installation check enabled, which is the default (Figure 3-22).

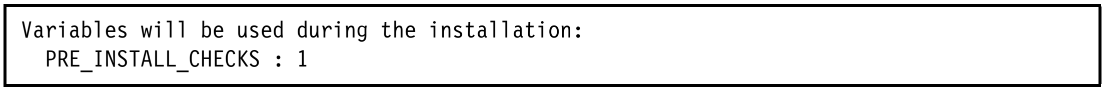

```
Variables will be used during the installation:
  PRE_INSTALL_CHECKS : 1
```

Figure 3-22   Pre-installation checks are enabled by default in the dirserv.rsp file

The preferred practice is to also run the **`checkds`** utility before you run **`dirinstall.server`**. The **`checkds.sh`** script is in the server folder:

```
/mnt/server/checkds/checkds.sh
```

On successful completion of the **`checkds`** script, proceed with the installation.

3.2.2 Prerequisites

For the installation, complete the following checks in any order before the installation of IBM Systems Director.

Yellow pages

Ensure that the yellow pages group is not running and is in an inoperative state (Figure 3-23). Otherwise, the embedded DB2 for the IBM Systems Director server does not install successfully.

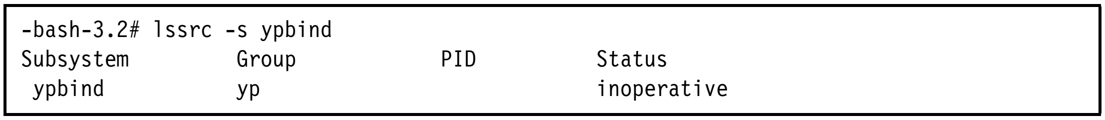

```
-bash-3.2# lssrc -s ypbind
Subsystem         Group            PID          Status
 ypbind           yp                            inoperative
```

Figure 3-23   Yellow pages

OS level

Check that the OS level of the AIX server is supported:

http://publib.boulder.ibm.com/infocenter/director/pubs/topic/com.ibm.director.plan.helps.doc/fqm0_t_planning_to_install_ibm_director_server.html

ulimits

Ensure that the `fsize` setting is set to `unlimited` because `fsize` determines the maximum allowable file size. See Figure 3-24.

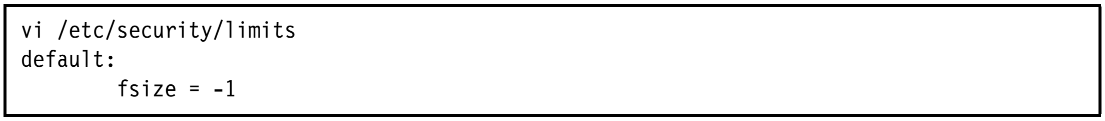

```
vi /etc/security/limits
default:
        fsize = -1
```

Figure 3-24   ulimits

The IBM Systems Director installation file is larger than 2 GB. Log out of the terminal session to activate the `fsize` changes.

Required installation files

Ensure that the filesets for AIX are installed for Secure Shell (ssh) and Secure Sockets Layer (ssl) at the following level or greater. See Figure 3-25.

```
-bash-3.2# lslpp -L | egrep "ssh|ssl"
  openssh.base.client     5.0.0.5302    C     F    Open Secure Shell Commands
  openssh.base.server     5.0.0.5302    C     F    Open Secure Shell Server
  openssh.license         5.0.0.5302    C     F    Open Secure Shell License
  openssl.base             0.9.8.801    C     F    Open Secure Socket Layer
  openssl.man.en_US       0.9.8.1800    C     F    Open Secure Socket Layer
-bash-3.2#
```

Figure 3-25 ssh/ssl filesets

Check that no previous installation of IBM Systems Director exists (Figure 3-26). If a previous installation exists, uninstall it.

```
lslpp -l | egrep -i “directorserver|directorcomm|directorplat|cimserver|cas”
```

Figure 3-26 lslpp check

If file sets from a previous installation are returned or for the file sets that are listed in Figure 3-25, remove the associated files (Figure 3-27).

```
installp -ug DirectorServer DirectorCommonAgent DirectorPlatformAgent cas.rte
sysmgt.cimserver.pegasus.rte
```

Figure 3-27 installp -ug

If a service is locked and cannot be removed, use the **lsof** or **rmsock** command to determine which port or file prevents the removal of the associated files. Then, remove the file systems that are associated to the files (Figure 3-28).

```
rm -rf /opt/ibm/director
rm -rf /opt/ibm/icc
rm -rf /opt/ibm/tivoli
```

Figure 3-28 folder removal

Note: It is not compulsory to remove the file sets that are referenced with the **-ug** option. However, by removing these file sets, you eliminate any issues with previous installations for agents or the server, which leads to a smoother installation.

Volume groups

The suggestion is to leave root volume group (rootvg) primarily for the operating system. Create a separate volume group for the additional storage that is required for IBM Systems Director on an alternate disk. By keeping rootvg lean and clean, you can recover more easily. The IBM Systems Director server recovery is described in 4.7, “Backup” on page 233.

Because the DB2 installation is restricted, the DB2 installation path is also restricted (Figure 3-29).

```
/home/dirinst1
/opt/ibm/director/db2
```

Figure 3-29 DB2 default paths

Changing this path to an alternate path for the system backup and restoration of rootvg is beneficial (Figure 3-30).

```
mklv -y “isddb2” -t jfs2 rootvg 10G
crfs -v jfs2 -d isddb2 -m /isddb2 -A yes
mount /isddb2
```

Figure 3-30 Changing the path

Tip: This installation has only one disk. It is advisable to have n+1 and to mirror the volume groups that are associated to the disks. Repeat the commands in Figure 3-30 for `/opt/ibm/director` on the alternate volume group if you want.

The **`checkds`** script looks for 3 GB or greater of paging space (Figure 3-31).

```
    -bash-3.2# lsps -a
    Page Space Physical Volume   Volume Group    Size %Used Active  Auto  Type Chksum
    hd6 hdisk0           rootvg        1024MB     2   yes   yes    lv     0
    -bash-3.2# chps -s 2 hd6
    -bash-3.2# lsps -a
    Page Space      Physical Volume   Volume Group    Size %Used Active  Auto  Type
    Chksum
    hd6             hdisk0            rootvg        3072MB     1   yes   yes    lv     0
```

Figure 3-31 Paging space

3.2.3 Installation

Because the `iso` file is mounted on `/mnt`, we changed the path to `/mnt/server/`, which is the location of the **`dirinstall.server`** executable script. Before you run the installation script, change the default DB2 path of the database. Because the media is mounted in read-only mode, copy the file to a temporary directory. Edit the `dirserv.rsp` file with a text editor. Figure 3-32 shows the database path that we chose for the installation. It is referenced by the **`DB_DATAPATH`** variable.

```
# Used to specify where the managed DB2 database will be stored when
# managed DB2 database is selected.  If not specified the default path will be
# /home/dirinst1.  If the path does not exist, it will be created.
DB_DATAPATH=/isddb2
```

Figure 3-32 dirserv.rsp

When you start the installation for IBM Systems Director, name the executable the name that is shown in Figure 3-33 to point to the changed response file.

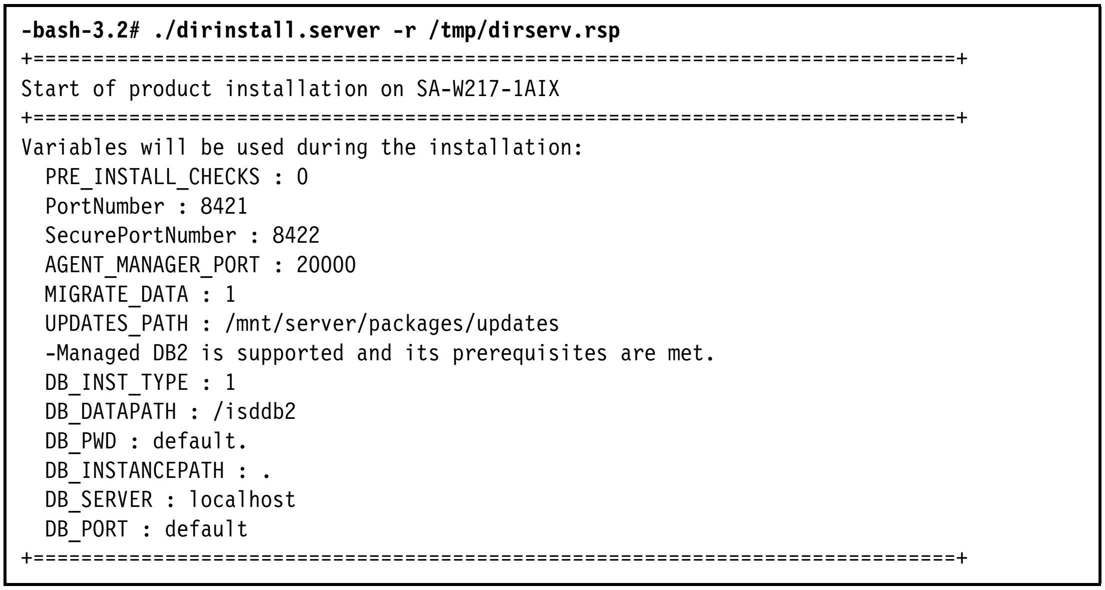

```
-bash-3.2# ./dirinstall.server -r /tmp/dirserv.rsp
+===============================================================================+
Start of product installation on SA-W217-1AIX
+===============================================================================+
Variables will be used during the installation:
  PRE_INSTALL_CHECKS : 0
  PortNumber : 8421
  SecurePortNumber : 8422
  AGENT_MANAGER_PORT : 20000
  MIGRATE_DATA : 1
  UPDATES_PATH : /mnt/server/packages/updates
  -Managed DB2 is supported and its prerequisites are met.
  DB_INST_TYPE : 1
  DB_DATAPATH : /isddb2
  DB_PWD : default.
  DB_INSTANCEPATH : .
  DB_SERVER : localhost
  DB_PORT : default
+===============================================================================+
```

Figure 3-33 Specifying a response file

Successful installation is similar to Figure 3-34.

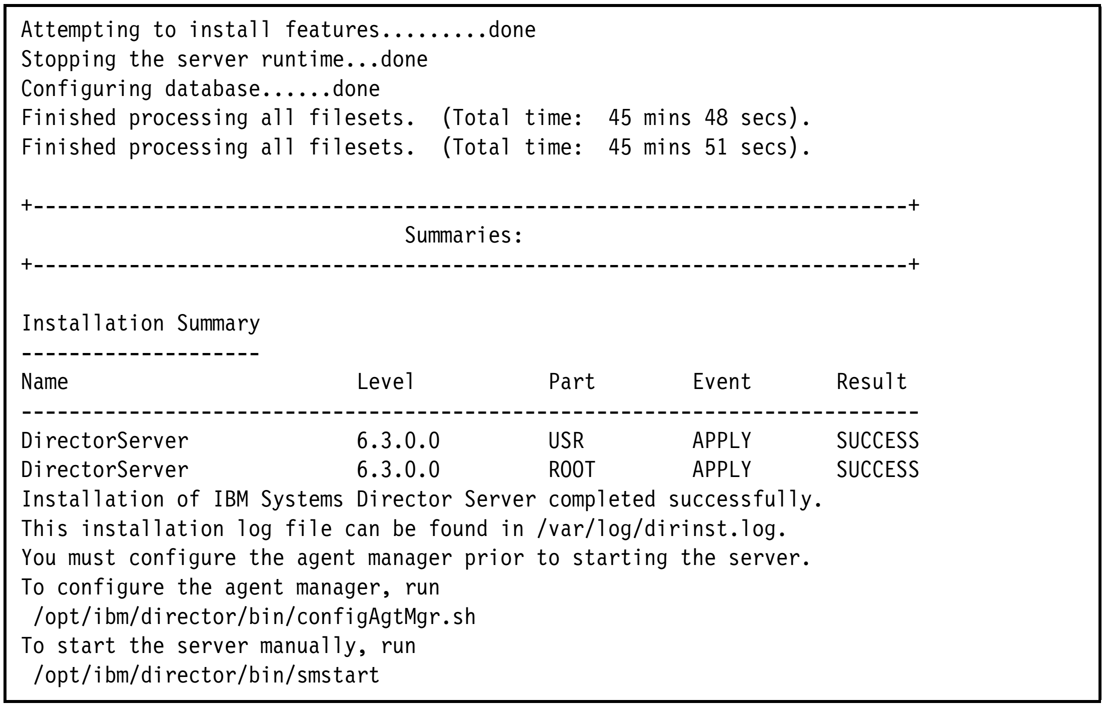

```
Attempting to install features.........done
Stopping the server runtime...done
Configuring database......done
Finished processing all filesets.  (Total time:  45 mins 48 secs).
Finished processing all filesets.  (Total time:  45 mins 51 secs).

+---------------------------------------------------------------------------+
                                Summaries:
+---------------------------------------------------------------------------+

Installation Summary
--------------------
Name                        Level           Part        Event       Result
---------------------------------------------------------------------------
DirectorServer              6.3.0.0         USR         APPLY       SUCCESS
DirectorServer              6.3.0.0         ROOT        APPLY       SUCCESS
Installation of IBM Systems Director Server completed successfully.
This installation log file can be found in /var/log/dirinst.log.
You must configure the agent manager prior to starting the server.
To configure the agent manager, run
 /opt/ibm/director/bin/configAgtMgr.sh
To start the server manually, run
 /opt/ibm/director/bin/smstart
```

Figure 3-34 Successful installation

The Agent Manager provides authentication and authorization services for managed systems that have common installed agents here:

`http://pic.dhe.ibm.com/infocenter/director/pubs/index.jsp?topic=%2Fcom.ibm.director.cli.helps.doc%2Ffqm0_r_cli_configAgtMgr.html`

As shown in Figure 3-34 on page 61, the Agent Manager must be configured *before* you start the IBM Systems Director server. See Figure 3-35.

```
/opt/ibm/director/bin/configAgtMgr.sh
-bash-3.2# /opt/ibm/director/bin/configAgtMgr.sh
Enter 1 to use the Agent Manager installed with this server (recommended)
Enter 0 to use an existing Agent Manager (advanced) : 1
Enter Resource Manager username : itso
Enter Resource Manager password :isd4itso
Re-Enter Resource Manager password :
Enter Agent Registration password :
Re-Enter Agent Registration password :
Re-Enter Agent Registration password :
[Add] [Element]: AgentManagerUserID [Value]: itso
[Add] [Element]: AgentManagerPassword [Value]:
{aes:3C5SnKQL63SjkEy44Gs+vHE6nQzC+Dil1NzNvSiAzk=}fFn7zXZpwvH3wYuP1yCIw==
[Add] [Element]: ManagerRegistrationPassword [Value]:
{aes:3C5SnKQL63SjkEy44Gs+vHE6nQzC+Dil1NzNvSiAzk=}fFn7zXZpwvH3wYuP1yCIw==
DataSourceConfig.sh=0
DataStoreInstall.sh=0
GenerateCertificates.sh=0
EncryptAMProps.sh=0
WebConfig.sh=0
usmi-cas-setup.sh=0
-bash-3.2#
```

Figure 3-35 Agent Manager configuration

The return codes of all called scripts must be `0`. Because all tasks are now successfully completed, we start IBM Systems Director (Figure 3-36).

```
-bash-3.2# export /opt/ibm/director/bin
-bash-3.2# export PATH=$PATH:/opt/ibm/director/bin
-bash-3.2# smstart
Starting IBM Director...
The starting process may take a while. Please use smstatus to check if the server is
active.
-bash-3.2# smstatus -r
Starting
Active
```

Figure 3-36 Starting IBM Systems Director

Figure 3-36 confirms that the server returned an `Active` state. Now, you can change the DB2 parameters that relate to the system setup, if necessary.

3.2.4 DB2 settings

Disabling remote access for the DB2 user is not required. However, it helps to prevent issues with user IDs and in-house AIX security policies:

http://www-01.ibm.com/support/docview.wss?uid=isg1IC83082

Stop IBM Systems Director and make the following changes:

1. Stop the IBM Systems Director server (Figure 3-37).

```
-bash-3.2# smstop
Shutting down IBM Director...
```

Figure 3-37 Stopping IBM Systems Director

2. Edit the user file (Figure 3-38).

```
-bash-3.2#vi /etc/security/user
dirinst1:
        admin = false
        rlogin=false
```

Figure 3-38 Edit the user properties file

3. Edit the `sshd_config` file (Figure 3-39).

```
vi /etc/ssh/sshd_config
# Added to restrict remote access
DenyUsers   dirinst1
```

Figure 3-39 Edit the sshd_config file

4. Restart `sshd` (Figure 3-40).

```
-bash-3.2#  stopsrc -s sshd
0513-044 The sshd Subsystem was requested to stop.
-bash-3.2# startsrc -s sshd
0513-059 The sshd Subsystem has been started. Subsystem PID is 16973974.
-bash-3.2#1
```

Figure 3-40 Restart sshd

5. Restart IBM Systems Director (Figure 3-41).

```
-bash-3.2# smstart
Starting IBM Director...
The starting process may take a while. Please use smstatus to check if the server is
active.
```

Figure 3-41 Restart IBM Systems Director

Because the response file for the DB2 installation is customized, check whether the database path is configured as requested in the response file. Figure 3-42 shows file system usage.

```
df -g
-bash-3.2# df -g /home/dirinst1 /isddb2
Filesystem    GB blocks      Free %Used    Iused %Iused Mounted on
/dev/hd1           2.00      1.95    3%      211     1% /home
/dev/isddb2       10.00      9.31    7%      104     1% /isddb2
```

Figure 3-42 File system usage

By using **db2** commands, query the database parameters to confirm the path within DB2 (Figure 3-43).

```
-bash-3.2# su - dirinst1
$ db2 get dbm config | grep "database pa"
 Default database path                       (DFTDBPATH) = /isddb2
$
```

Figure 3-43 DB2 path

Note: During the installation, you might not edit the `dirserv.rsp` file. You can change the default database path after the installation by using the **db2relocatedb** command. This command does not require a backup and restore. For more information, go to the following site:

http://pic.dhe.ibm.com/infocenter/db2luw/v9r7/index.jsp?topic=%2Fcom.ibm.db2.luw.admin.cmd.doc%2Fdoc%2Fr0004500.html

Because IBM Systems Director is installed and no endpoint discoveries or additional tasks are complete, back up IBM Systems Director in its current state. Before you complete the save, create another lv, fs, and mount. Or, use the **smsave** command.

Optionally, create another mount point to point the **smsave** to an alternate directory (Figure 3-44).

```
mklv -y "isdbkup" -t jfs2 rootvg 10G
crfs -v jfs2 -d isdbkup -m /isdbkup -A yes
mount /isdbkup
```

Figure 3-44 Back up lv and fs

After you create another mount point, run **smsave** with options (Figure 3-45).

```
-bash-3.2# smstop;smsave -targetDir /isdbkup
Shutting down IBM Director...
Command is running. Monitor live status and results in /opt/ibm/director/log/smsave.log

ALR1325I: The lightweight runtime has started.
com.ibm.net.SocketKeepAliveParameters

Command completed successfully
```

Figure 3-45 The smsave command with options

Figure 3-46 shows the **smsave** command with no options.

```
-bash-3.2# smstop;smsave;smstart
Shutting down IBM Director...
Command is running. Monitor live status and results in /opt/ibm/director/log/smsave.log
ALR1325I: The lightweight runtime has started.
com.ibm.net.SocketKeepAliveParameters
Command completed successfully
Starting IBM Director...
The starting process may take a while. Please use smstatus to check if the server is
active.
-bash-3.2# smstatus
Active
```

Figure 3-46 smsave with no options

Figure 3-47 shows the location of both backups.

```
-bash-3.2# ls -al /isdbkup
total 8
drwxr-xr-x    4 root     system          256 Nov 07 11:27 .
drwxr-xr-x   26 root     system         4096 Nov 07 10:40 ..
drwxr-xr-x    8 root     system          256 Nov 07 11:30 2012_11_7_11.27.1
drwxr-xr-x    2 root     system          256 Nov 07 10:40 lost+found
-bash-3.2# ls -al /opt/ibm/director/backup
total 8
drwxr-xr-x    3 root     system          256 Nov 07 11:30 .
drwxr-xr-x   30 root     system         4096 Oct 22 14:30 ..
drwxr-xr-x    8 root     system          256 Oct 22 14:32 2012_10_22_14.29.29
-bash-3.2#
```

Figure 3-47 Location of backups

IBM Systems Director backups are discussed in 4.7, “Backup” on page 233.

3.2.5 Initial login

After you successfully start IBM Systems Director, log in to the server through the user interface (UI), as shown in Figure 3-48 on page 66.

In this example, `SA-W217-1AIX.itso.ral.ibm.com` is the host name of the server where IBM Systems Director is installed. Or, you can use the native IP address of the server. This example uses the following URL:

```
https://SA-W217-1AIX.itso.ral.ibm.com:8422/ibm/console/logon.jsp
```

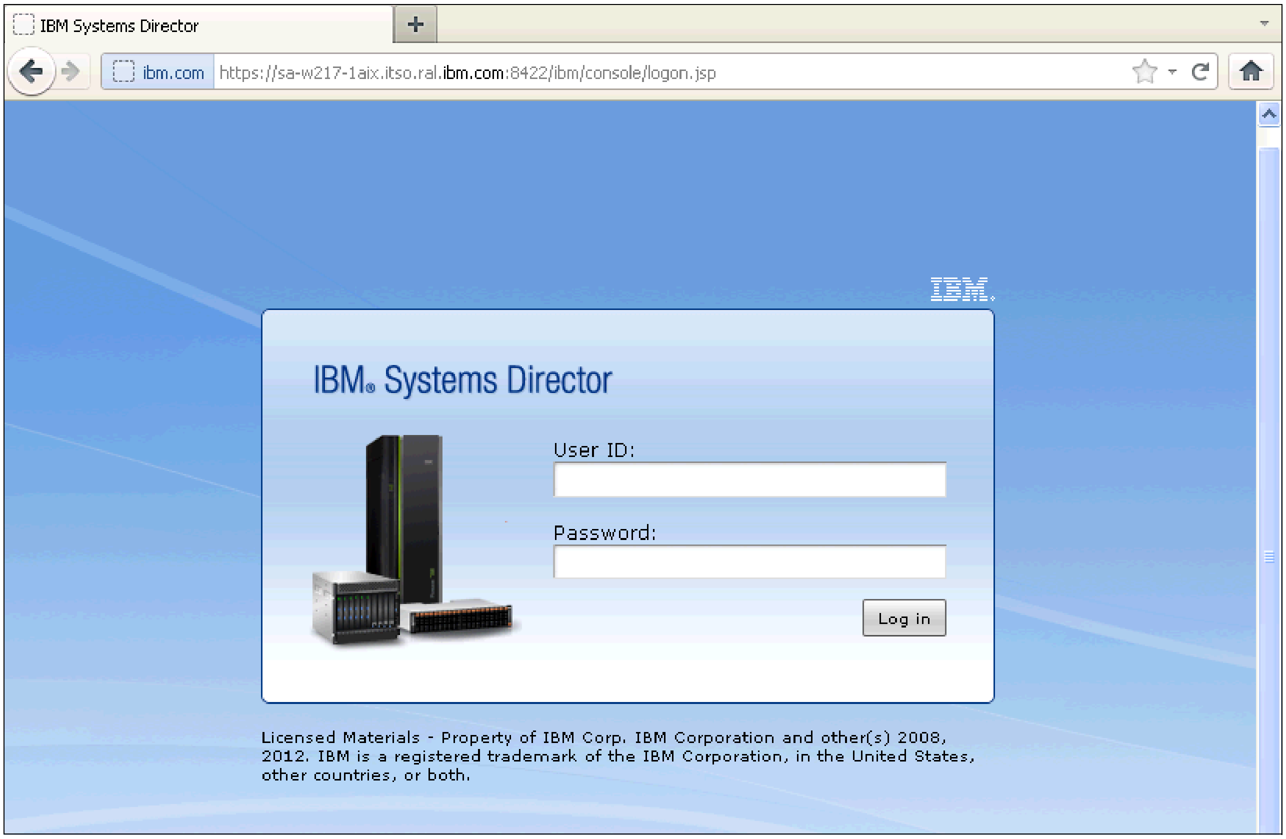

IBM Systems Director
https://sa-w217-1aix.itso.ral.ibm.com:8422/ibm/console/logon.jsp
IBM.

IBM® Systems Director

User ID:

Password:

Log in

Licensed Materials - Property of IBM Corp. IBM Corporation and other(s) 2008, 2012. IBM is a registered trademark of the IBM Corporation, in the United States, other countries, or both.

Figure 3-48 IBM Systems Director login page

After you log in to IBM Systems Director, select the **Plug-ins** tab. Select **IBM Systems Director server** to get an overview of the server and associated properties (Figure 3-49).

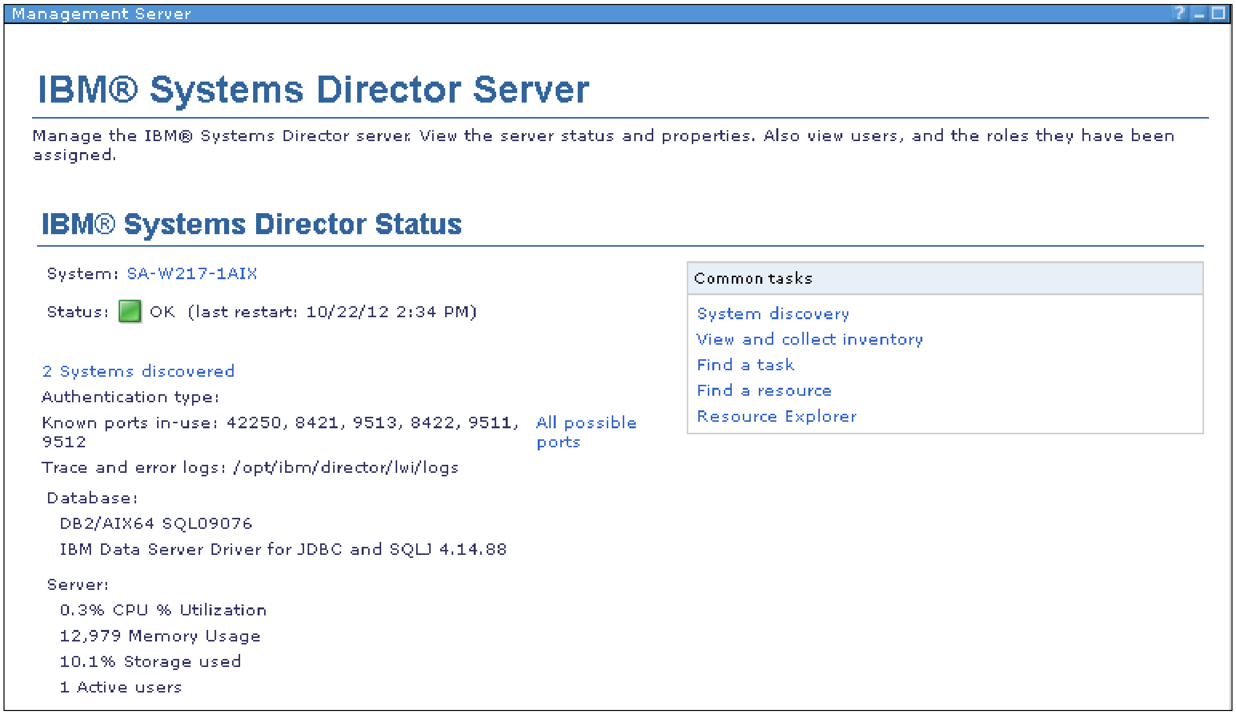

IBM® Systems Director Server

Manage the IBM® Systems Director server. View the server status and properties. Also view users, and the roles they have been assigned.

IBM® Systems Director Status

System: SA-W217-1AIX

Status: 🟩 OK (last restart: 10/22/12 2:34 PM)

2 Systems discovered

Authentication type:

Known ports in-use: 42250, 8421, 9513, 8422, 9511, 9512 All possible ports

Trace and error logs: /opt/ibm/director/lwi/logs

Database:
DB2/AIX64 SQL09076
IBM Data Server Driver for JDBC and SQLJ 4.14.88

Server:
0.3% CPU % Utilization
12,979 Memory Usage
10.1% Storage used
1 Active users

Common tasks
- System discovery
- View and collect inventory
- Find a task
- Find a resource
- Resource Explorer

Figure 3-49 IBM Systems Director server

3.2.6 Installing the IBM Systems Director license

The IBM Systems Director license is on the root path of the ISO or GZIP file that you transferred to the installation server. The license key can be imported by using the command line. Or, use the UI for a license key that is stored locally on the computer that is used to access the UI (Figure 3-50).

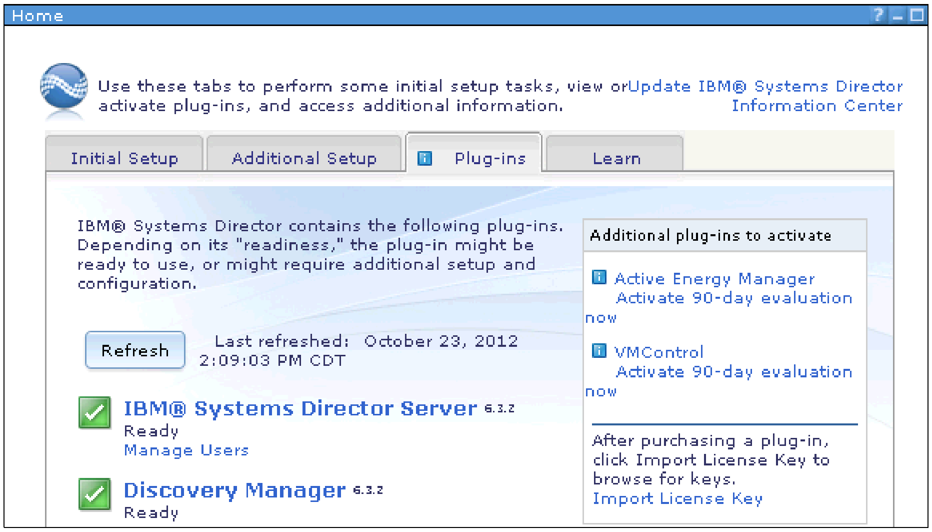

Home

Use these tabs to perform some initial setup tasks, view or activate plug-ins, and access additional information.

Update IBM® Systems Director Information Center

Initial Setup | Additional Setup | Plug-ins | Learn

IBM® Systems Director contains the following plug-ins. Depending on its "readiness," the plug-in might be ready to use, or might require additional setup and configuration.

Refresh

Last refreshed: October 23, 2012 2:09:03 PM CDT

✓ IBM® Systems Director Server 6.3.2
Ready
Manage Users

✓ Discovery Manager 6.3.2
Ready

Additional plug-ins to activate

Active Energy Manager
Activate 90-day evaluation now

VMControl
Activate 90-day evaluation now

After purchasing a plug-in, click Import License Key to browse for keys.
Import License Key

Figure 3-50 UI license import

Figure 3-51 shows an example of importing the license key from the server installation code.

```
-bash-3.2# importkey ISD_express_edition_power.lpsa
International Program License Agreement

Part 1 - General Terms

BY DOWNLOADING, INSTALLING, COPYING, ACCESSING, CLICKING ON
AN "ACCEPT" BUTTON, OR OTHERWISE USING THE PROGRAM,
LICENSEE AGREES TO THE TERMS OF THIS AGREEMENT. IF YOU ARE
ACCEPTING THESE TERMS ON BEHALF OF LICENSEE, YOU REPRESENT
AND WARRANT THAT YOU HAVE FULL AUTHORITY TO BIND LICENSEE
TO THESE TERMS. IF YOU DO NOT AGREE TO THESE TERMS,

- DO NOT DOWNLOAD, INSTALL, COPY, ACCESS, CLICK ON AN
"ACCEPT" BUTTON, OR USE THE PROGRAM; AND

- PROMPTLY RETURN THE UNUSED MEDIA, DOCUMENTATION, AND
PROOF OF ENTITLEMENT TO THE PARTY FROM WHOM IT WAS OBTAINED

Press Enter to continue viewing the license agreement, or
enter "1" to accept the agreement, "2" to decline it, "3"
to print it, or "99" to go back to the previous screen.
1
Importing license keys.
 IBM Systems Director Express Edition
All keys imported successfully.
```

Figure 3-51 CLI license import

For more information about the license, go to the following site:

http://pic.dhe.ibm.com/infocenter/director/pubs/index.jsp?topic=%2Fcom.ibm.director.editions.doc%2Feditions_power_express_license.html

3.3 Installation of IBM Systems Director server for a Linux on Power platform

Useful information is described about the process of installing IBM Systems Director for a Linux on Power platform. For the installation of IBM Systems Director for Linux on Power, go to the following site:

http://publib.boulder.ibm.com/infocenter/director/pubs/topic/com.ibm.director.install.helps.doc/fqm0_t_installing.html

3.3.1 Downloading the software

The IBM Systems Director management server code can be sourced from the following site:

http://www-03.ibm.com/systems/software/director/downloads/mgmtservers.html

Place the code on the Linux on Power system by using a file transfer method of your choice. After the code is on the Linux system, use this command to extract the installation files:

```
tar -zxf express_Power_Linux_Director_base.tar.gz
```

3.3.2 Prerequisites

To determine whether the Linux distribution fulfills the software requirements, change to the `/server/checkds` directory. Each Linux distribution that is supported has additional Red Hat Package Manager (RPM) packages that need to be installed. Look for the **`checkds.sh`** script in the `checkds` folder. This script checks the state of the server and whether the server is supported for the IBM Systems Director server. If all the required RPM packages are installed, look for a return code of `0`.

The **`checkds`** script invokes a `/checkds/checklists/lin-server-chklist.properties` checklist file. This file is unique for each supported OS on which the IBM Systems Director server is installed.

Before the installation of IBM Systems Director for Linux on Power, review Table 3-2 on page 69. Use the links in the footnotes to source the additional required RPM packages.

Table 3-2 Software requirements for Linux on Power

Installation scenario	**Required RPM packages on the Agent**	**Required RPM packages on the server**
Red Hat Enterprise Linux Advanced Platform, version 5.*x* on IBM Power Systems	► compat-libstdc++-<version>.ppc.rpm[a] ► servicelog-0.2.9-0.ppc64.rpm[b] ► openssl097a-0.9.7a-9.<*version*>.ppc.rpm[c] ► librtas-1.3.4-0.ppc64.rpm[b]	► vacpp.rte[d,e] ► Required RPM packages on Agent
Red Hat Enterprise Linux Advanced Platform, version 6.* on IBM Power Systems	► compat-libstdc++-33.ppc[a] ► libstdc++-4.4.4-13.el6.ppc.rpm[a] ► pam-1.1.1-4.el6.ppc.rpm[a] ► servicelog-1.1.7-2.el6.ppc64[a] ► librtas-1.3.4-2.el6.ppc[a] ► libservicelog-1.1.9-4.el6.ppc[a] ► expat-2.0.1-9.1.el6.ppc[a] ► compat-expat1-1.95.8-8.el6.ppc[a]	► vacpp.rte[d,e] ► Required RPM packages on Agent
SUSE Linux Enterprise Server 10 on IBM Power Systems	► compat-libstdc++-<*version*>.ppc.rpm[a] ► libservicelog-1.1.9-1.ppc.rpm[b] ► servicelog-1.1.7-1.ppc.rpm[b] ► lsvpd-0.16.0-1.ppc.rpm[b] ► librtas-1.3.5-1.ppc.rpm[b]	► vacpp.rte[d,e] ► Required RPM packages on Agent
SUSE Linux Enterprise Server 11 on IBM Power Systems	► libstdc++33-3.3.3-11.9.ppc64.rpm[a] ► libservicelog-1.1.9-1.ppc.rpm[b] ► servicelog-1.1.7-1.ppc.rpm[b] ► lsvpd-0.16.0-1.ppc.rpm[b] ► librtas-1.3.5-1.ppc.rpm[b] ► pam-32bit-1.0.2-20.1.ppc64.rpm[a]	► vacpp.rte[d,e] ► Required RPM packages on Agent ► gcc-4.3-62.198.ppc64.rpm ► gcc-c++-4.3-62.198.ppc64.rpm ► libstdc++43-devel-4.3.3_20081022-11.18.ppc64.rpm ► gcc43-c++-4.3.3_20081022-11.18.ppc64.rpm ► glibc-devel-2.11.1-0.17.4.ppc64.rpm ► linux-kernel-headers-2.6.32-1.4.13.noarch.rpm
SUSE Linux Enterprise Server 11 SP2 on IBM Power Systems	► libstdc++33-3.3.3-11.9.ppc64.rpm[a] ► libservicelog-1.1.9-1.ppc.rpm[b] ► servicelog-1.1.7-1.ppc.rpm[b] ► lsvpd-0.16.0-1.ppc.rpm[b] ► librtas-32bit-1.3.6-010.1.ppc64.rpm[a] ► ppc64-diag-2.4.2-0.14.12.ppc64.rpm[a] ► libvpd2-2.1.3-0.9.1.ppc64.rpm[a] ► pam-32bit-1.0.2-20.1.ppc64.rpm[a] ► pam-modules-32bit-11-1.22.1.ppc64.rpm[a]	

a. Obtain this RPM package from the operating system distribution media. There might be minor version variations from the versions that are listed, which are acceptable.

b. Obtain this RPM package from IBM Service and productivity tools for Linux on Power Systems at: `https://www14.software.ibm.com/webapp/set2/sas/f/lopdiags/home.html`. Select your Linux distribution and then select the appropriate tab for your version. Follow any special instructions for each RPM package. For RHEL5, if the listed RPM version is not available on the website, get it from the RHEL4 tab.

c. Obtain this RPM package from the operating system distribution media in addition to openssl 0.9.8, which is installed by default.

d. Server only.

e. Obtain the `tar.gz` package from `https://www-304.ibm.com/support/docview.wss?uid=swg24030460`. Untar and install the three included RPM packages. This action applies for all platforms.

Note: The 64-bit RPM package file names include `ppc64`. The 32-bit RPM package file names include `ppc`. If the listed RPM package shows `ppc`, you need the 32-bit version. The Platform Agent does not install if you show the 64-bit version only.

For software requirements, see the following site:

http://www.ibm.com/developerworks/wikis/display/WikiPtype/Software+requirements+for+Director+6.3+on+Linux+on+Power

Run the **`checkds.sh`** script (Figure 3-52). Check for the return code `0`. If the return code is not `0`, review, fix, and run again.

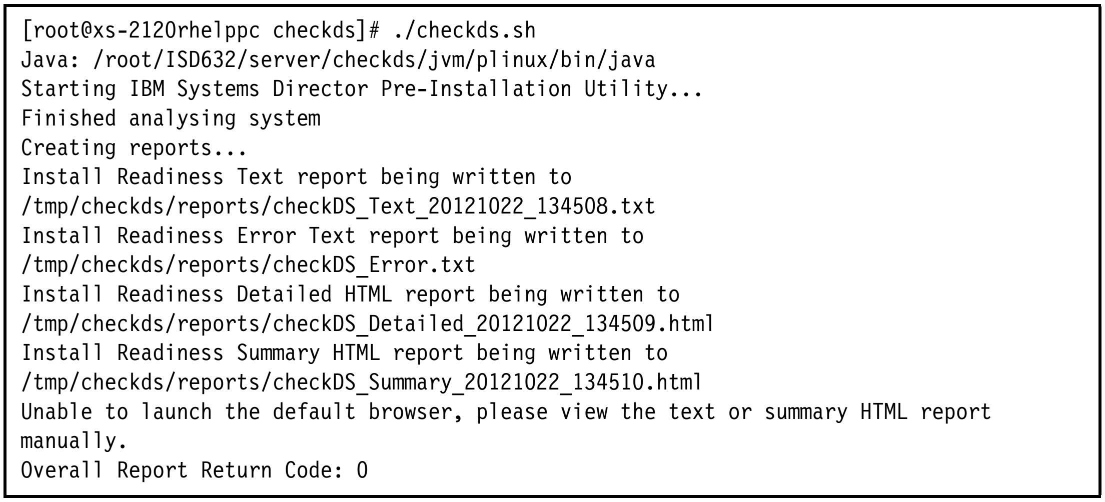

```
[root@xs-2120rhelppc checkds]# ./checkds.sh
Java: /root/ISD632/server/checkds/jvm/plinux/bin/java
Starting IBM Systems Director Pre-Installation Utility...
Finished analysing system
Creating reports...
Install Readiness Text report being written to
/tmp/checkds/reports/checkDS_Text_20121022_134508.txt
Install Readiness Error Text report being written to
/tmp/checkds/reports/checkDS_Error.txt
Install Readiness Detailed HTML report being written to
/tmp/checkds/reports/checkDS_Detailed_20121022_134509.html
Install Readiness Summary HTML report being written to
/tmp/checkds/reports/checkDS_Summary_20121022_134510.html
Unable to launch the default browser, please view the text or summary HTML report
manually.
Overall Report Return Code: 0
```

Figure 3-52 Running the checkds script

3.3.3 Installing the IBM Systems Director server

After you see the return code 0 from the **checkds** script (Figure 3-52 on page 70), proceed with the installation (Figure 3-53).

```
[root@xs-2120rhelppc server]# ./dirinstall.server
Agree to product license?
[1-Agree|0-Disagree]:1
..../....
Enter 1 to use the Agent Manager installed with this server (recommended)
Enter 0 to use an existing Agent Manager (advanced) : 1
Enter Resource Manager username : isd4itso
Enter Resource Manager password :
Re-Enter Resource Manager password :
Enter Agent Registration password :
Re-Enter Agent Registration password :
[Add] [Element]: AgentManagerUserID [Value]: isd4itso
[Add] [Element]: AgentManagerPassword [Value]:
{aes:3C5SnKQL63SjkEy44Gs+vHF6nQzC+Dil1NzNvSiAzzk=}fFn7zXZpwvsH3wYuP1yCIw==
[Add] [Element]: ManagerRegistrationPassword [Value]:
{aes:3C5SnKQL63SjkEy44Gs+vHF6nQzC+Dil1NzNvSiAzzk=}fFn7zXZpwvsH3wYuP1yCIw==
DataSourceConfig.sh=0
DataStoreInstall.sh=0
GenerateCertificates.sh=0
EncryptAMProps.sh=0
WebConfig.sh=0
usmi-cas-setup.sh=0
Installation of the IBM Systems Director Server 6.3.2 succeeded.
To start the server manually, run /opt/ibm/director/bin/smstart.
To see the status, run /opt/ibm/director/bin/smstatus [-r].
```

Figure 3-53 Summary output of dirinstall.server

After the installation completes successfully, start IBM Systems Director for Linux on Power (Figure 3-54).

```
[root@xs-2120rhelppc server]# /opt/ibm/director/bin/smstart
Starting IBM Director...The starting process may take a while. Please use smstatus to
check if the server is active.
[root@xs-2120rhelppc server]# /opt/ibm/director/bin/smstatus -r
Starting
Active
```

Figure 3-54 Starting IBM Systems Director for Linux on Power

After the IBM Systems Director server returns an Active status (Figure 3-54 on page 71), go to the login page (Figure 3-55).

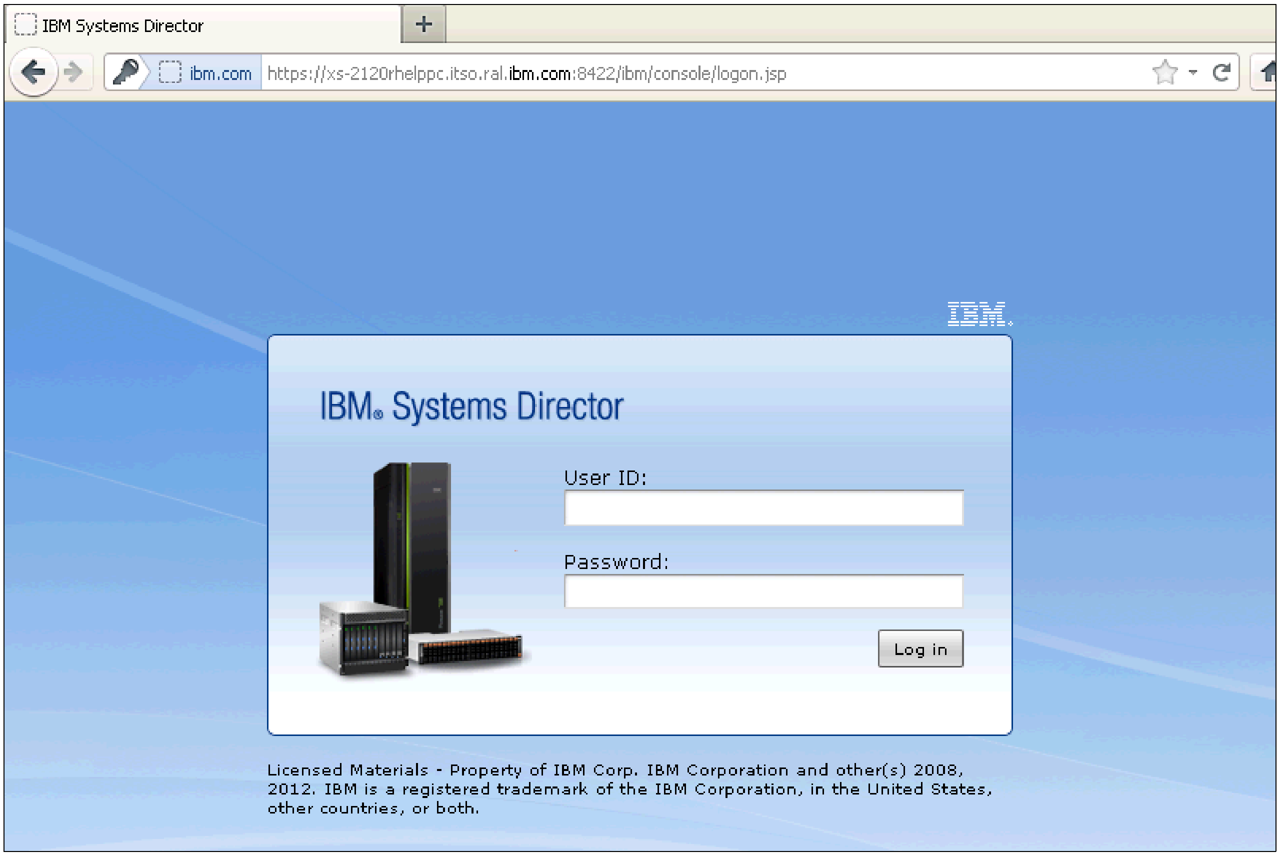

Licensed Materials - Property of IBM Corp. IBM Corporation and other(s) 2008, 2012. IBM is a registered trademark of the IBM Corporation, in the United States, other countries, or both.

Figure 3-55 IBM Systems Director login page

After you log in, select the **Plug-ins** tab and select **IBM Systems Director server** to see an overview of the server and associated properties (Figure 3-56).

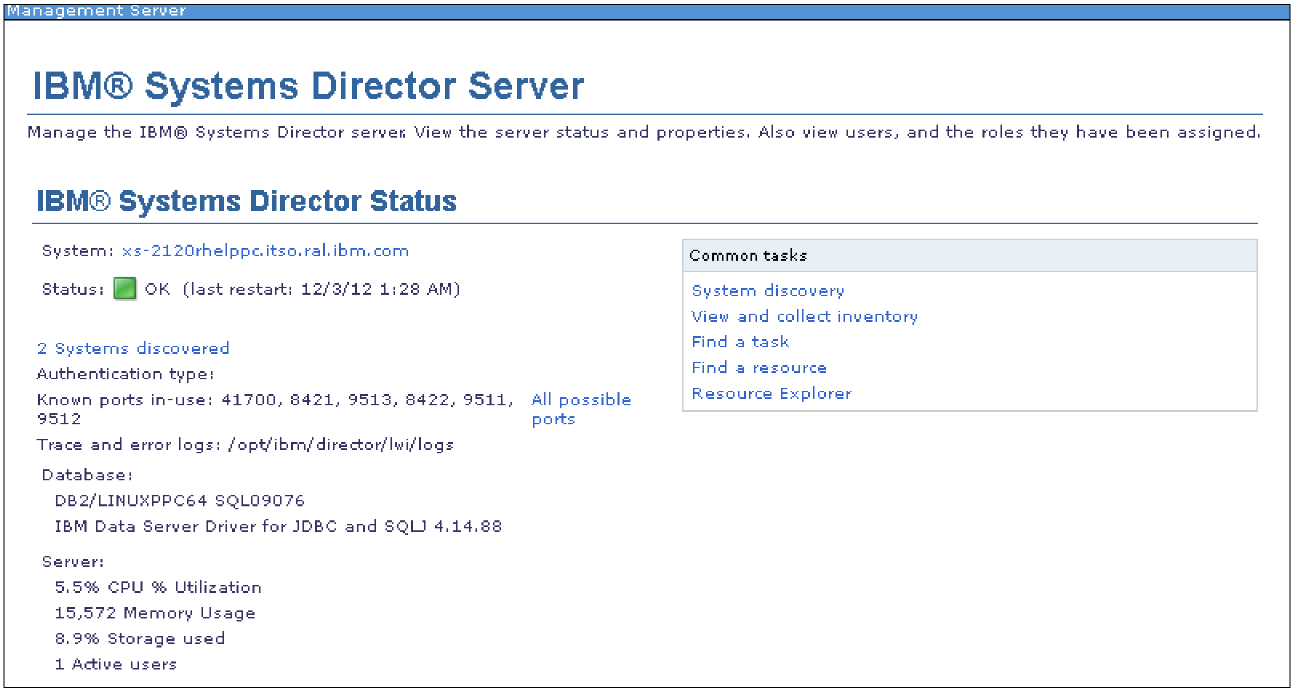

Management Server

IBM® Systems Director Server

Manage the IBM® Systems Director server. View the server status and properties. Also view users, and the roles they have been assigned.

IBM® Systems Director Status

System: xs-2120rhelppc.itso.ral.ibm.com

Status: 🟩 OK (last restart: 12/3/12 1:28 AM)

2 Systems discovered

Authentication type:

Known ports in-use: 41700, 8421, 9513, 8422, 9511, 9512 All possible ports

Trace and error logs: /opt/ibm/director/lwi/logs

Database:
DB2/LINUXPPC64 SQL09076
IBM Data Server Driver for JDBC and SQLJ 4.14.88

Server:
5.5% CPU % Utilization
15,572 Memory Usage
8.9% Storage used
1 Active users

Figure 3-56 IBM Systems Director server

3.4 Installation of the IBM Systems Director agent on Linux x86

Because the Information Center does not explain how you can install the IBM Systems Director Platform or Common Agent without disabling SELinux and iptables, we cover those steps here because this is a common requirement.

Use the following steps to deploy the Platform or Common Agent on a Linux operating system MEP running on x86 without disabling SELinux or the iptables service.

1. Ensure that you have prepared your Linux system as documented at the following site:

   http://pic.dhe.ibm.com/infocenter/director/pubs/topic/com.ibm.director.install.helps.doc/fqm0_t_preparing_to_install_core_on_xseries.html

   > **Tip:** At the time of writing, on a default installation of RHEL 6.3, you need to install the commands shown in Example 3-1 to satisfy dependencies for the Platform Agent. You can find those packages on the RHEL 6.3 installation media, which can be added to your yum repository by running the **vim /etc/yum.repos.d/rhel-dvd.repo** command and following the instructions in the text shown in Example 3-2 on page 74.

Example 3-1 Packages to install on default RHEL 6.3

```
yum install compat-libstdc++-33.i686
yum install compat-expat1.i686
yum install pam-1.1.1.i686
yum install libstdc++.i686
yum install libuuid.i686
```

Example 3-2   Content of /etc/yum.repos.d/rhel-dvd.repo

```
[rhel-dvd]
name=Red Hat Enterprise Linux $releasever - $basearch - DVD
baseurl=file:///<path where you have mounted your RHEL installation media>/Server/
enabled=1
gpgcheck=1
gpgkey=file:///<path where you have mounted your RHEL installation
media>/RPM-GPG-KEY-redhat-release
```

2. Discover the operating system MEP by running a System Discovery and unlock it with root.

> **Note:** You can use sudo to unlock access to your MEPs, as documented at the following link:
>
> http://pic.dhe.ibm.com/infocenter/director/pubs/topic/com.ibm.director.security.helps.doc/fqm0_t_access_agentless_with_sudo_and_keypair_authentification.html?resultof=%22%73%75%64%6f%22%20
>
> However, until the agent is installed, try not to use sudo access because this prevents the agent installation from completing successfully. After the agent is deployed, you can revoke access and reconfigure it with sudo.

3. Run a full inventory on the MEP by selecting it and going to **Actions** → **Inventory** → **Collect Inventory** in the Resource Explorer view.
4. On the Linux system, edit the iptables configuration to open necessary ports:
   a. Use the `iptables -nL --line-numbers` command to list existing rules and note the line number of the `-A INPUT -j REJECT --reject-with icmp-host-prohibited` rule, as shown in *red* in Example 3-3.

   Example 3-3   Showing the active rules in iptables

   ```
   [root@SBvmrhel1 tmp]# iptables -nL --line-numbers
   Chain INPUT (policy ACCEPT)
   num  target     prot opt source               destination
   1    ACCEPT     all  --  0.0.0.0/0            0.0.0.0/0           state
   RELATED,ESTABLISHED
   2    ACCEPT     icmp --  0.0.0.0/0            0.0.0.0/0
   3    ACCEPT     all  --  0.0.0.0/0            0.0.0.0/0
   4    ACCEPT     tcp  --  0.0.0.0/0            0.0.0.0/0           state NEW
   tcp dpt:22
   5 REJECT    all  --  0.0.0.0/0          0.0.0.0/0          reject-with icmp-host-prohibited

   Chain FORWARD (policy ACCEPT)
   num  target     prot opt source               destination
   1    REJECT     all  --  0.0.0.0/0            0.0.0.0/0
   reject-with icmp-host-prohibited

   Chain OUTPUT (policy ACCEPT)
   num  target     prot opt source               destination
   ```

   b. If you see that line (normally applies only to RHEL), you need to delete it before adding new rules. To delete it, use the **`iptables -D INPUT`** ***n*** command where ***n*** is the line number (in Example 3-3, it is line 5).

c. Add the required firewall rules by running the commands that are shown in Example 3-4.

Example 3-4 Adding required rules in iptables

```
iptables -A INPUT -p tcp --dport 427 -j ACCEPT
iptables -A INPUT -p udp --dport 427 -j ACCEPT
iptables -A INPUT -p tcp --dport 15988 -j ACCEPT
iptables -A INPUT -p tcp --dport 15989 -j ACCEPT
iptables -A INPUT -p udp --dport 161 -j ACCEPT
iptables -A INPUT -p udp --dport 162 -j ACCEPT
```

d. *Optionally*, add the Virtual Network Computing (VNC) ports if you are deploying to a virtual server managed by VMControl and want to enable remote console access by running this command:

```
iptables -A INPUT -p tcp --dport 5900:5950 -j ACCEPT
```

e. Add the reject rule back by running this command:

```
iptables -A INPUT -j REJECT --reject-with icmp-host-prohibited
```

f. Save the rules by running this command:

```
/sbin/service iptables save
```

Tip: If you are using a default installation of Novell SLES, you do not need to add firewall rules because there are none that are defined by default, which means that all network traffic is allowed.

5. You can now deploy the platform agent on your Linux operating system MEP by right-clicking the MEP and selecting **Release Management** → **Install Agent**.
6. Collect inventory again as explained in step 4 on page 74 now that the agent has been deployed.

3.5 Best practices

To summarize, keep in mind the following best practices when installing IBM Systems Director:

- *Prepare your system carefully*. Check the prerequisite (see Chapter 2, “Planning” on page 7 for this).
- *Run the pre-check* before starting the installation. Fix all shown problems.
- *Follow the installation steps*, read the instruction for each of the steps, and put in the requested information.
- *Use a technical user* (with local admin rights) as credential for installing the IBM Systems Director.
- To avoid problems, the user name that is used for the IBM Systems Director installation should not contain *ibm* in the name.
- *Select the embedded DB2 database* as the IBM Systems Director database. This makes the installation and configuration easier.
- After finishing the installation, *run the Post Installation Validator* to check that the installation runs and finishes successfully.

4

Basic management tasks

This chapter covers best practices for the basic management tasks that are available in IBM Systems Director. These tasks include how to perform discovery; how to collect inventory on discovered resources; how to deploy agent, driver, and firmware updates; how to manage hardware events and other miscellaneous tasks.

The following topics are covered:

- 4.1, “Discovery” on page 78
- 4.2, “Inventory” on page 86
- 4.3, “Updates” on page 104
- 4.4, “Event management” on page 143
- 4.5, “Hardware Management Console and AIX Launch-in-Context” on page 194
- 4.6, “Security” on page 208
- 4.7, “Backup” on page 233

4.1 Discovery

Systems Director 6.3.2 can discover various types of endpoints. For multiple discovery options and multiple resource types to discover, see the Discovery section of the web interface.

For more information about the discovery and inventory processes of Systems Director, see the Systems Director Discovery Information Center:

http://pic.dhe.ibm.com/infocenter/director/pubs/index.jsp?topic=%2Fcom.ibm.director.discovery.helps.doc%2Ffqm0_t_discovering_and_inventorying_resources.html

Figure 4-1 displays the available options to specify the endpoints that need to be discovered. You can limit the discovery process to a single address or range of sequential addresses. Or, you can use a discovery profile.

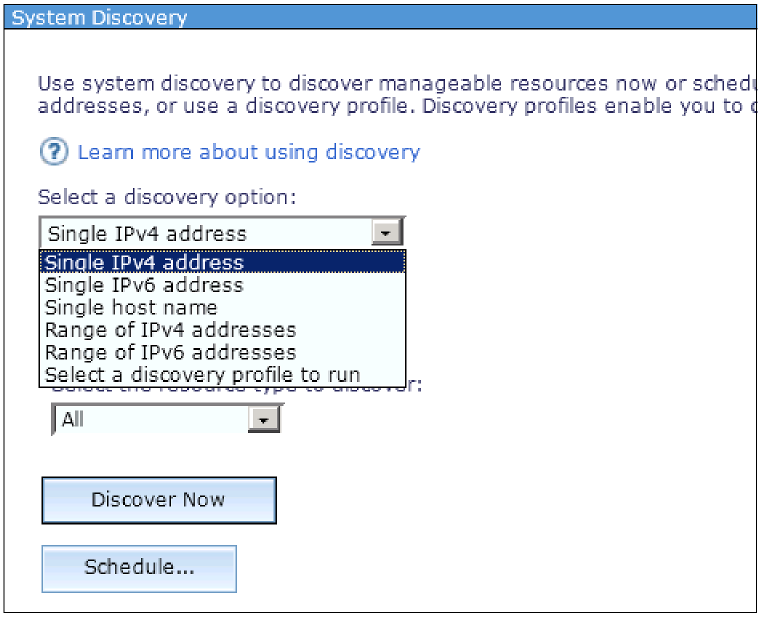

System Discovery

Use system discovery to discover manageable resources now or schedu
addresses, or use a discovery profile. Discovery profiles enable you to d

Select a discovery option:

Figure 4-1 Discovery options

Figure 4-2 shows the resource options that are available to discover. Use resource types to limit the discovery process to protocols that are based on the resource type to discover.

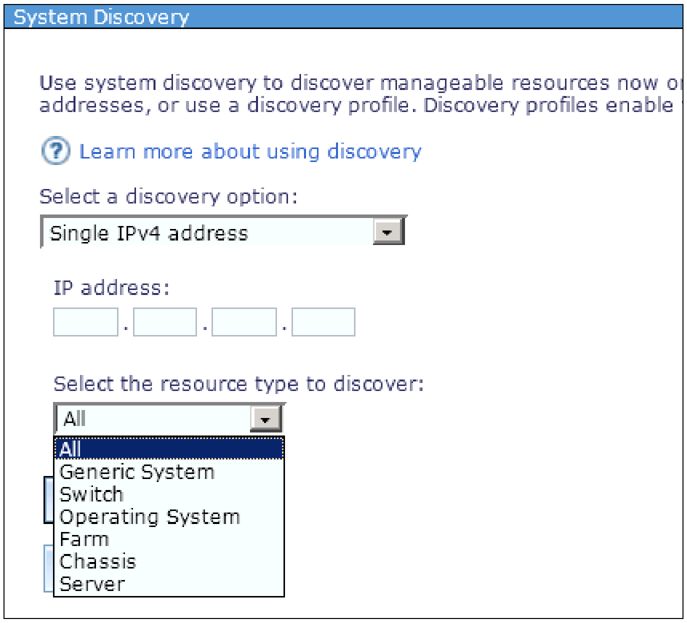
System Discovery

Use system discovery to discover manageable resources now o
addresses, or use a discovery profile. Discovery profiles enable

Select a discovery option: Single IPv4 address

IP address: . . .

Select the resource type to discover: All

Figure 4-2 Resource options

Tip: IBM Systems Director will not discover Windows Server 2012 systems unless you prepare those systems as documented in the information center:

`http://pic.dhe.ibm.com/infocenter/director/pubs/index.jsp?topic=%2Fcom.ibm.director.install.helps.doc%2Ffqm0_t_preparing_windows_server_2012_managed_systems.html`

4.1.1 Discovery profiles

Using discovery profiles is the best way to efficiently and effectively perform a discovery within Systems Director. Figure 4-3 displays the Advanced Tasks section, which gives you the options to manage discovery profiles.

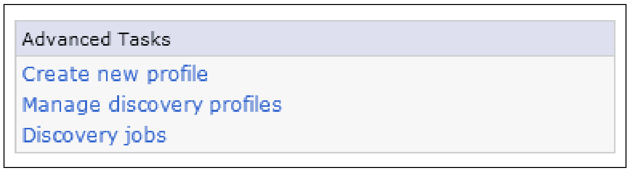

Figure 4-3 Advanced Tasks page

Figure 4-4 shows an example of creating a discovery profile for a BladeCenter chassis.

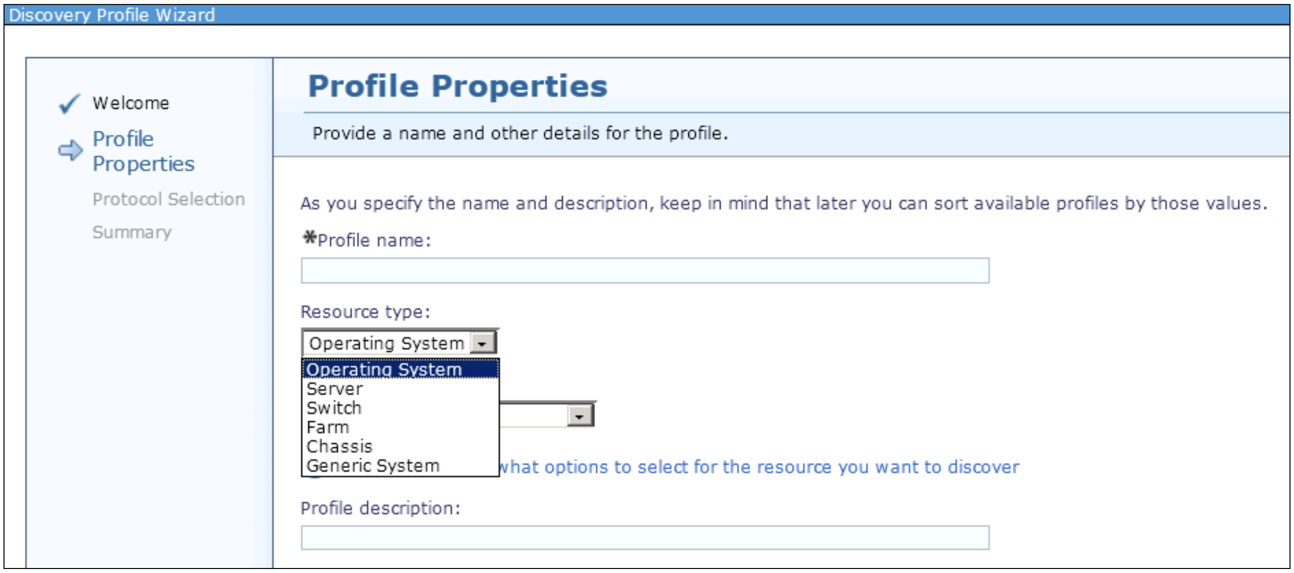

Discovery Profile Wizard

Profile Properties

Provide a name and other details for the profile.

As you specify the name and description, keep in mind that later you can sort available profiles by those values.

*Profile name:

Resource type:

Operating System

what options to select for the resource you want to discover

Profile description:

Figure 4-4 Profile Properties page

Figure 4-5 is the Protocol Selection page. If the chosen resource type supports additional discovery protocols, the protocols are listed.

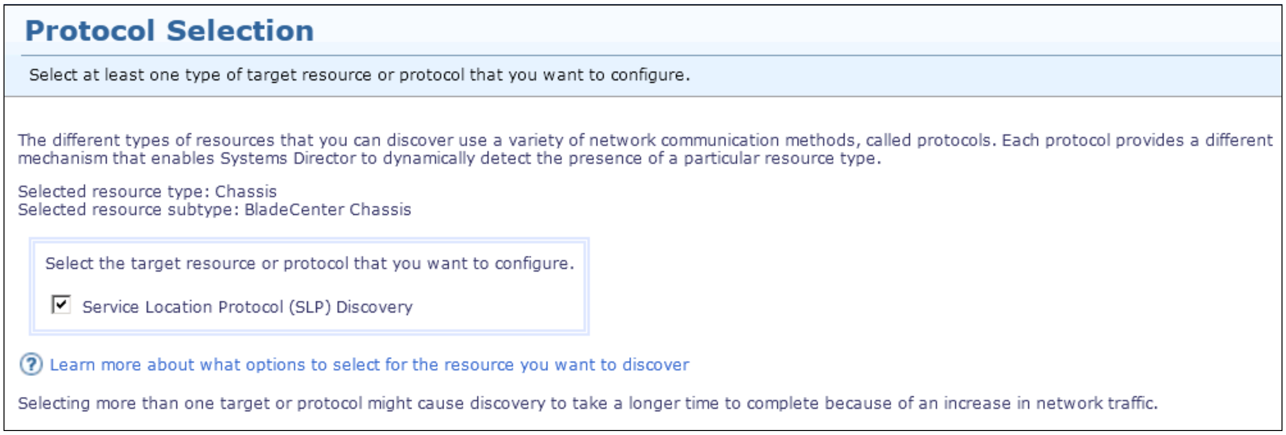

Protocol Selection

Select at least one type of target resource or protocol that you want to configure.

The different types of resources that you can discover use a variety of network communication methods, called protocols. Each protocol provides a different mechanism that enables Systems Director to dynamically detect the presence of a particular resource type.

Selected resource type: Chassis
Selected resource subtype: BladeCenter Chassis

Select the target resource or protocol that you want to configure.

☑ Service Location Protocol (SLP) Discovery

Learn more about what options to select for the resource you want to discover

Selecting more than one target or protocol might cause discovery to take a longer time to complete because of an increase in network traffic.

Figure 4-5 Protocol Selection page

Figure 4-6 shows where you configure how to discover your endpoints. Depending on the network, Unicast, Multicast, and Broadcast can be used. For best results, the use of Unicast is advised. With this option, you can specify an IP address or sequential range of IP addresses of endpoints. You also can import a group of nonsequential IP addresses by importing a text file or a CSV file that contains one IP address per line.

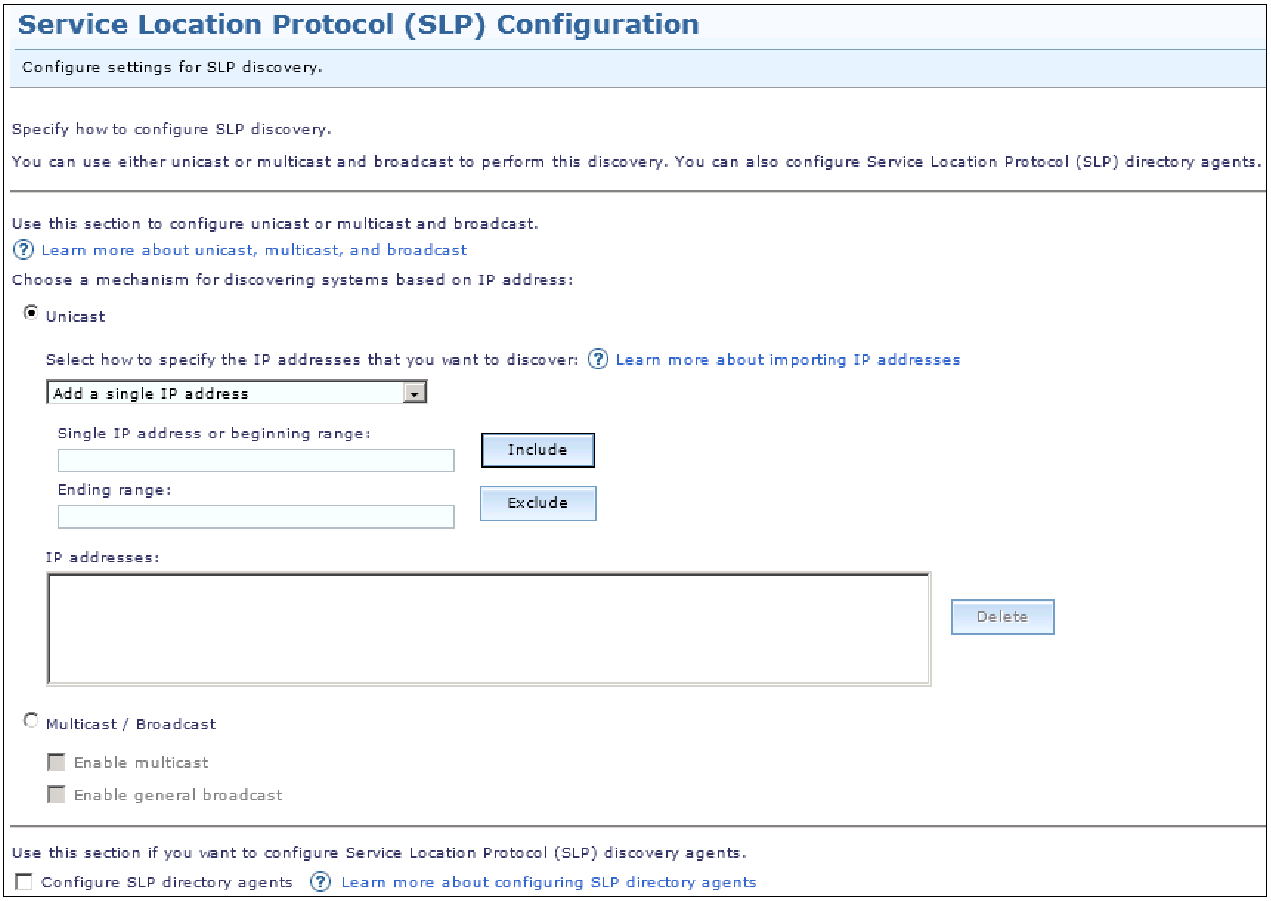

Service Location Protocol (SLP) Configuration

Configure settings for SLP discovery.

Specify how to configure SLP discovery.

You can use either unicast or multicast and broadcast to perform this discovery. You can also configure Service Location Protocol (SLP) directory agents.

Use this section to configure unicast or multicast and broadcast.

? Learn more about unicast, multicast, and broadcast

Choose a mechanism for discovering systems based on IP address:

● Unicast

Select how to specify the IP addresses that you want to discover: ? Learn more about importing IP addresses

Add a single IP address ▾

Single IP address or beginning range:

Ending range:

IP addresses:

○ Multicast / Broadcast

□ Enable multicast

□ Enable general broadcast

Use this section if you want to configure Service Location Protocol (SLP) discovery agents.

□ Configure SLP directory agents ? Learn more about configuring SLP directory agents

Figure 4-6 SLP Configuration page

In Figure 4-7, you can enter credentials that automatically request access to the endpoint after the discovery.

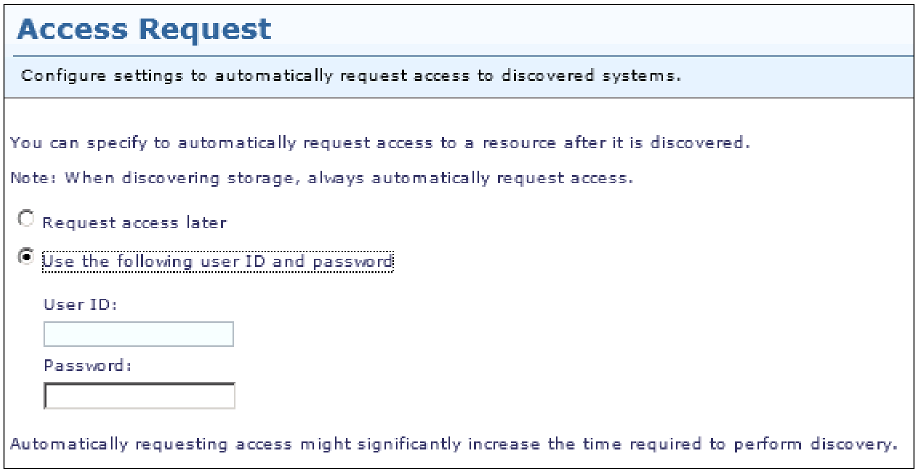

Access Request

Configure settings to automatically request access to discovered systems.

You can specify to automatically request access to a resource after it is discovered.

Note: When discovering storage, always automatically request access.

○ Request access later

● Use the following user ID and password

User ID:

Password:

Automatically requesting access might significantly increase the time required to perform discovery.

Figure 4-7 Access Request page

Figure 4-8 displays the new Discovery Profile wizard summary.

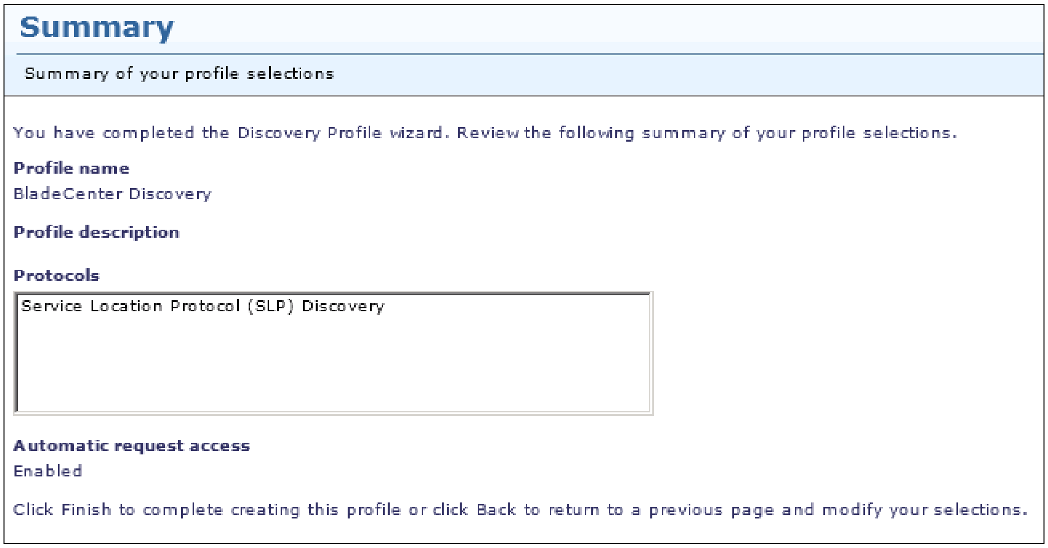

Summary

Summary of your profile selections

You have completed the Discovery Profile wizard. Review the following summary of your profile selections.

Profile name

BladeCenter Discovery

Profile description

Protocols

Service Location Protocol (SLP) Discovery

Automatic request access

Enabled

Click Finish to complete creating this profile or click Back to return to a previous page and modify your selections.

Figure 4-8 Summary

Figure 4-9 displays how to choose the profile that you created, which can be run immediately or scheduled.

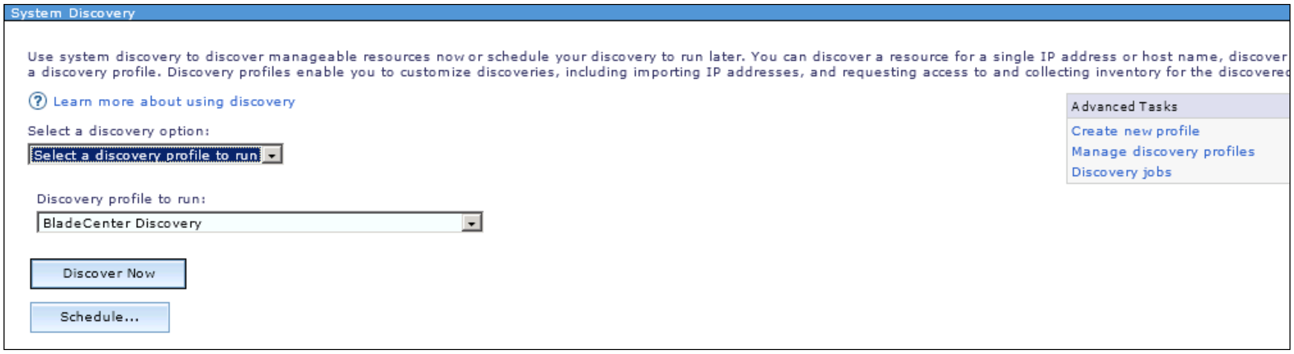

System Discovery

Use system discovery to discover manageable resources now or schedule your discovery to run later. You can discover a resource for a single IP address or host name, discover a discovery profile. Discovery profiles enable you to customize discoveries, including importing IP addresses, and requesting access to and collecting inventory for the discovered

Learn more about using discovery

Select a discovery option:

Select a discovery profile to run

Discovery profile to run:

BladeCenter Discovery

Figure 4-9 Selecting a discovery profile

4.1.2 BladeCenter discovery

To successfully discover a BladeCenter advanced management module (AMM), you must set several AMM prerequisites:

- *Increase the TCP command mode protocol to at least 10 connections.*

  Systems Director needs several concurrent connections to successfully communicate with a BladeCenter AMM.

Figure 4-10 shows the AMM page where the TCP command mode protocol is listed.

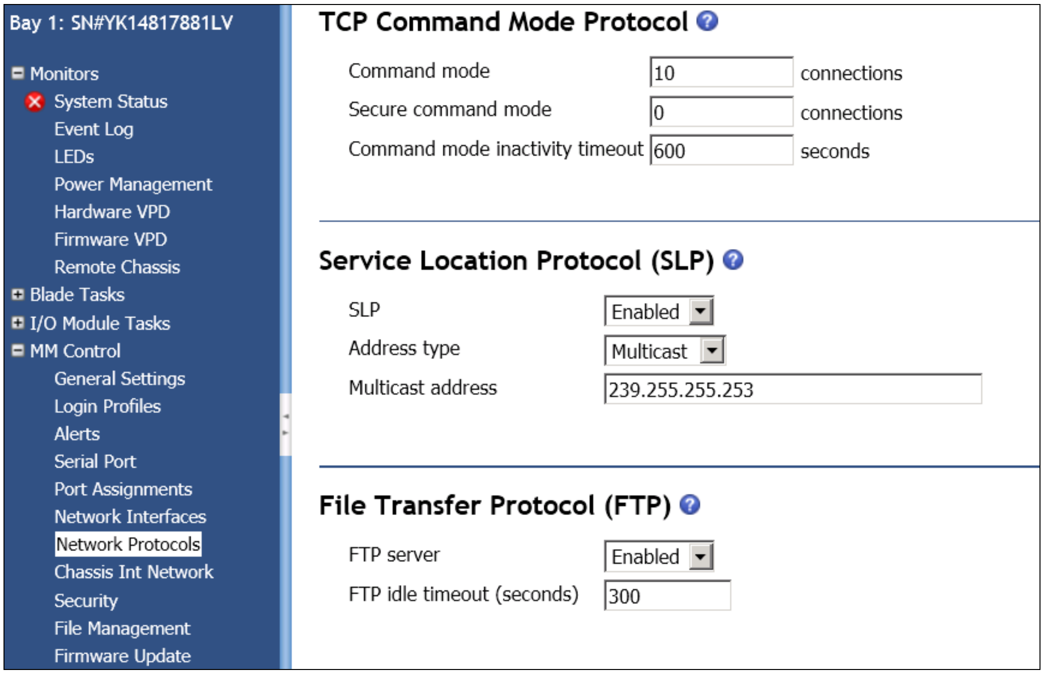

TCP Command Mode Protocol

Command mode 10 connections

Secure command mode 0 connections

Command mode inactivity timeout 600 seconds

Service Location Protocol (SLP)

SLP Enabled

Address type Multicast

Multicast address 239.255.255.253

File Transfer Protocol (FTP)

FTP server Enabled

FTP idle timeout (seconds) 300

Figure 4-10 TCP Command Mode Protocol page

- *Enable SNMPv1 and SNMPv3*.

  Set the trap destination to the IP address of the Systems Director server. Set the Access Type to **Get** or greater.

  By enabling SNMPv1 and SNMPv3, Systems Director uses the connection to collect a more comprehensive inventory. Figure 4-11 on page 84 shows the AMM SNMPv1 and SNMPv3 settings page. Both the AMM SNMPv1 and SNMPv3 agents need to be enabled.

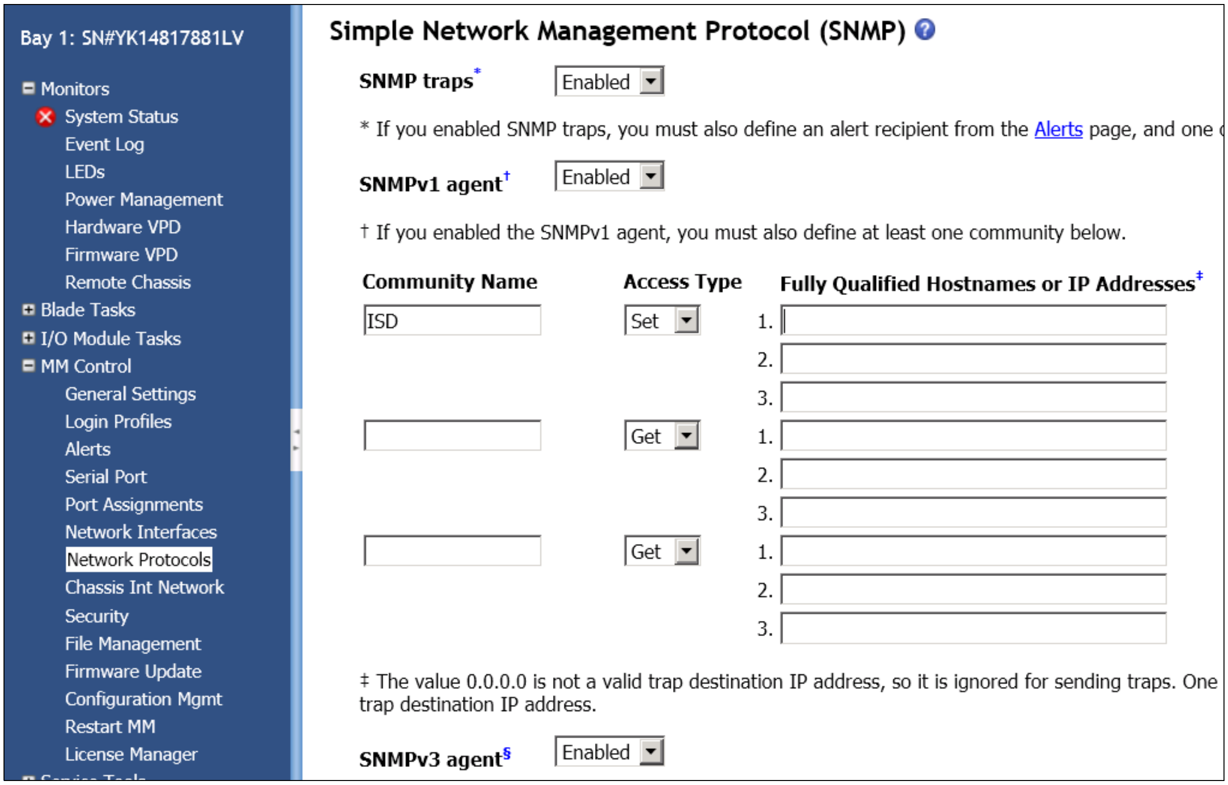

Simple Network Management Protocol (SNMP)

SNMP traps* Enabled

* If you enabled SNMP traps, you must also define an alert recipient from the Alerts page, and one

SNMPv1 agent† Enabled

† If you enabled the SNMPv1 agent, you must also define at least one community below.

Community Name	Access Type	Fully Qualified Hostnames or IP Addresses‡
ISD	Set	1.
		2.
		3.
	Get	1.
		2.
		3.
	Get	1.
		2.
		3.

‡ The value 0.0.0.0 is not a valid trap destination IP address, so it is ignored for sending traps. One trap destination IP address.

SNMPv3 agent§ Enabled

Figure 4-11 Simple Network Management Protocol (SNMP) configuration

- SNMPv3 needs a user profile that is associated to it. Set the Access type to **Set** (Figure 4-12).

  The user profile that is associated to SNMPv3 is used when you request access from the AMM.

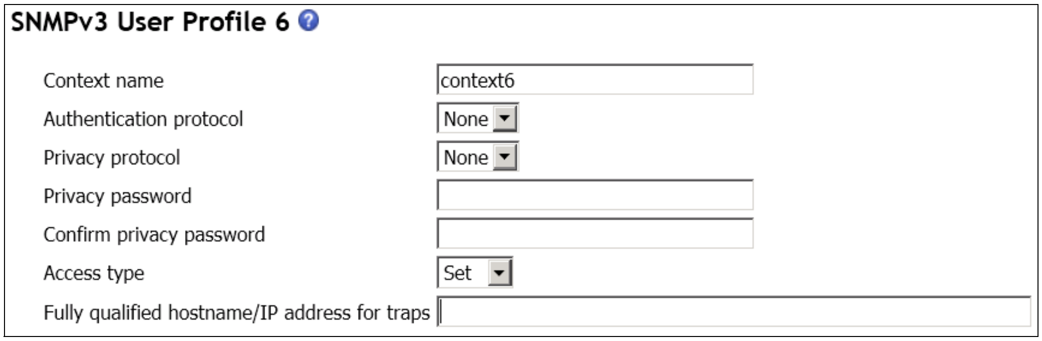

SNMPv3 User Profile 6

Context name	context6
Authentication protocol	None
Privacy protocol	None
Privacy password	
Confirm privacy password	
Access type	Set
Fully qualified hostname/IP address for traps	

Figure 4-12 SNMPv3 User Profile page

> **Tip:** Reboot AMM after these changes.

4.1.3 Best practices for discovery

Following are the best practices to use the discovery task efficiently:

- *Discover only the systems that you intend to manage.*

  Limiting discoveries to systems that you intend to manage speeds up the discovery process. You eliminate discovering other devices that might support the discovery protocols.

- *Keep the IP address ranges as small as possible.*

  Limiting the number of addresses in a single request can improve the reliability of the discovery.

- *Specify the types of resources to discover (avoid the use of All).*

  Systems Director can skip discovery protocols that are inappropriate for your resources. Skipping inappropriate discovery protocols results in shorter discovery time.

- *Schedule the discovery of large numbers of systems during off-hours.*

  Scheduling large discovery jobs off-hours improves the reliability of the discovery process and helps with the additional network traffic.

- *Where possible, use discovery profiles and specify individual IP addresses or use Service Location Protocol (SLP) Directory Agents.*

  Making the profile as specific as possible minimizes discovery time because Systems Director runs only the protocols that are configured.

  SLP Directory Agents reduce network traffic and increase discovery speed.

- *Prepare your hardware for discovery.*

4.2 Inventory

Inventory is one of the most important tasks. Inventory needs to be run on all systems that are managed by the Systems Director server. Inventory information provides the basis for much of the functionality in Systems Director.

The following elements are good examples of functionality that depends on inventory:

- Update Manager
- Compliance checks
- Dynamic groups

Inventory data for systems that are managed by Systems Director is stored in a database that is created and controlled by Systems Director. Since version 6.3, the default database format is IBM DB2.

Optionally, you can use external databases, such as IBM DB2, Oracle, and Microsoft SQL Server (the latter is for Windows platforms only). It is a preferred practice to use the built-in (local) IBM DB2 database, which is created and controlled by Systems Director at installation.

When a system is discovered by the Systems Director, a basic inventory scan runs for this system. This scan includes IP address, host name, OS, and if an agent is installed, the agent version. For additional information beyond these properties, the Systems Director needs full authorized access to the system.

When a system has access, run an inventory scan for this system to collect the complete inventory information. The complete inventory information includes hardware, software, and driver information from the system.

The following topics are described:

- 4.2.1, “Inventory data and collection profiles” on page 86
- 4.2.2, “Collecting inventory” on page 93
- 4.2.3, “Viewing inventory” on page 96
- 4.2.4, “Exporting inventory” on page 99

4.2.1 Inventory data and collection profiles

Systems Director uses inventory collection profiles to collect inventory data from discovered resources.

Systems Director uses profiles to manage the inventory collection tasks that you create and run. An *inventory collection profile* is a group of settings that are saved on the Systems Director server. The settings indicate the type of resources that are collected during the collection process.

By default, Systems Director includes the following inventory collection profiles:

- All Inventory

  This profile collects inventory from all resources and encompasses all the other inventory collection profiles.

  All inventory: The All Inventory profile is required if you intend to use Update Manager.

- All Hardware Inventory

  This profile collects inventory from physical and virtual devices.
- All Software Inventory

  This profile collects inventory from software resources.
- Basic System Information

  This profile collects inventory from system resources.

These predefined inventory profiles are read-only and cannot be deleted or edited. However, you can use these existing profiles to create your own profiles.

The use of inventory discovery profiles provides a predefined template to collect the inventory information that you need. This template is useful when you want to use the inventory information that Systems Director collects from the system for asset tools. Or, this template is useful if you need specific information from the system without going through all of the available inventory information.

To create your own inventory discovery profile, follow these steps:

1. Start on the View and Collect Inventory page. Select **Manage Profiles** next to the profile selection (Figure 4-13).

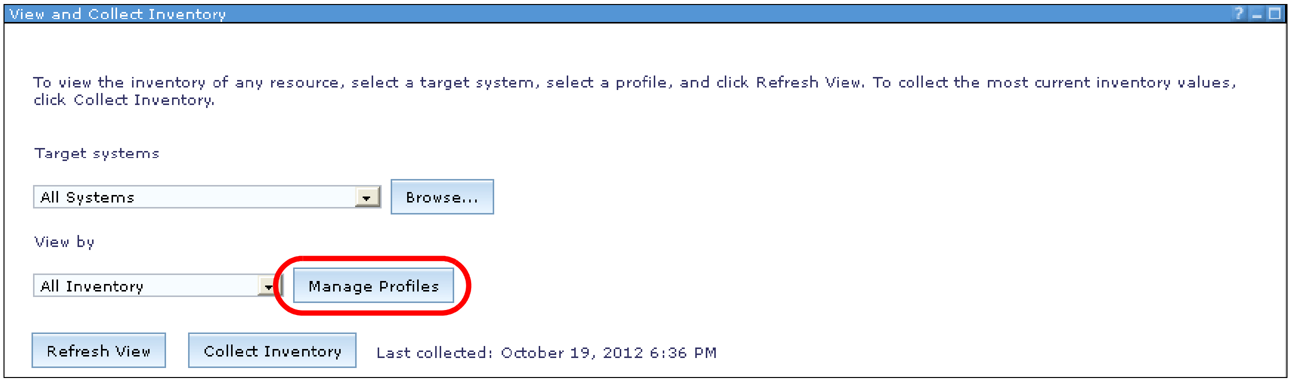

View and Collect Inventory

To view the inventory of any resource, select a target system, select a profile, and click Refresh View. To collect the most current inventory values, click Collect Inventory.

Target systems

All Systems

View by

All Inventory

Last collected: October 19, 2012 6:36 PM

Figure 4-13 Manage inventory discovery profiles

2. Select either to create a profile or copy an existing profile. When a function is selected, the Create Inventory Profile wizard opens. In our example, we create a profile. The first window is the Welcome page (Figure 4-14). Click **Next** at the bottom of the page (not shown) to continue.

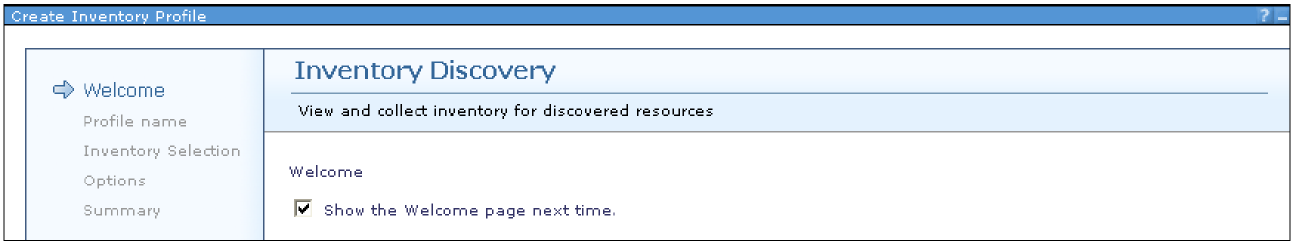

Create Inventory Profile

Inventory Discovery

View and collect inventory for discovered resources

Welcome

☑ Show the Welcome page next time.

Figure 4-14 Inventory Discovery Profile wizard: Welcome panel

3. In Figure 4-15, give your profile a name. If you chose to copy an existing profile, the default name is `copy_of_profilename`. In our example, we name the profile `book` as shown in Figure 4-15. A description is optional. Click **Next** to continue.

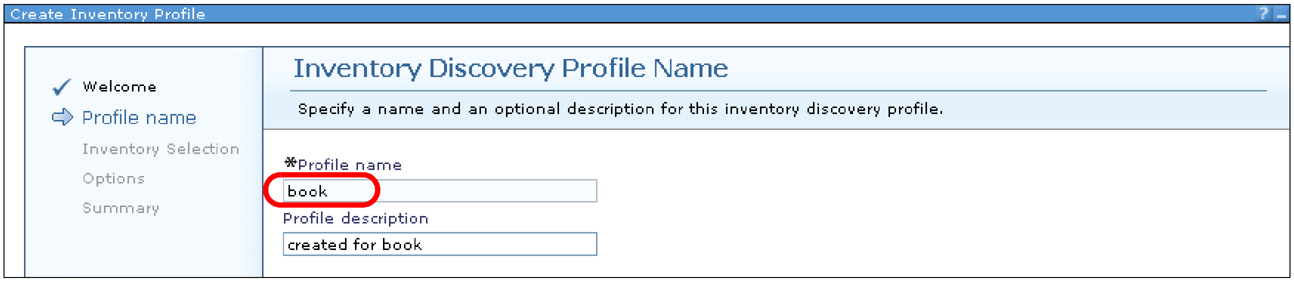

Create Inventory Profile

Inventory Discovery Profile Name

Specify a name and an optional description for this inventory discovery profile.

*Profile name: book

Profile description: created for book

Figure 4-15 Inventory Discovery Profile wizard profile name

4. In Figure 4-16, select which inventory resources to collect with your profile. To select an inventory resource, expand the resource groups on the left, make your selection, and click **Add** to copy your selection to the selected resources. You cannot copy complete resource groups. Instead, you must select each resource in a resource group to select the complete group. Click **Next** to continue.

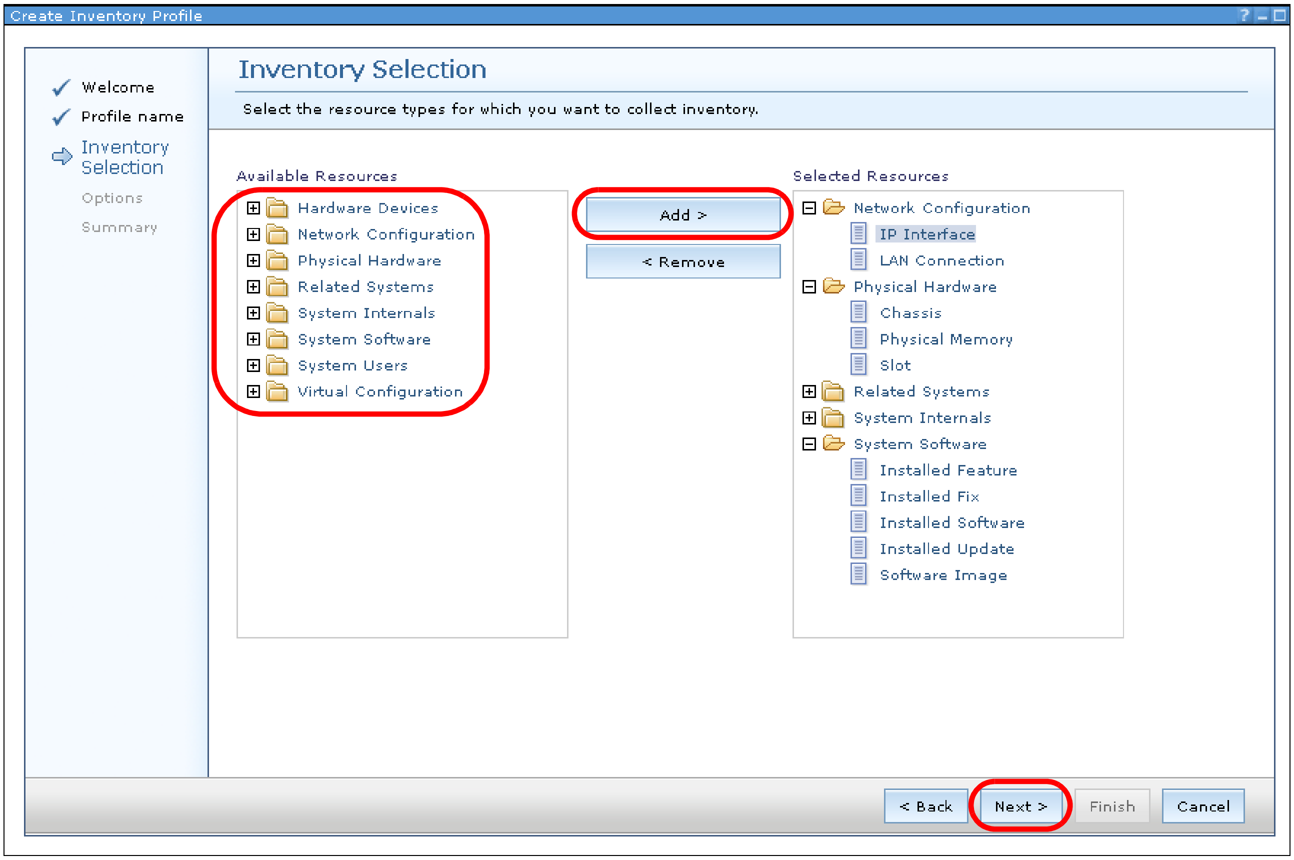

Create Inventory Profile

Inventory Selection

Select the resource types for which you want to collect inventory.

Available Resources

- Hardware Devices
- Network Configuration
- Physical Hardware
- Related Systems
- System Internals
- System Software
- System Users
- Virtual Configuration

Selected Resources

- Network Configuration
  - IP Interface
  - LAN Connection
- Physical Hardware
  - Chassis
  - Physical Memory
  - Slot
- Related Systems
- System Internals
- System Software
  - Installed Feature
  - Installed Fix
  - Installed Software
  - Installed Update
  - Software Image

Figure 4-16 Inventory Discovery Profile wizard: Inventory Selection page

5. In Figure 4-17, select the inventory service. You can either let the system select the inventory service or you can manually configure the discovery service.

   If you select **Let the system choose the discovery services**, click **Next** to see the option window as described in step 7 on page 91.

   If you select **Let me manually configure the discovery services** and click **Next**, you see what is shown in Figure 4-18. Go through the definition of the discovery services. In this window, you can select the available inventory profiles.

   In most cases, letting the system choose the discovery service is easier. With this method, no configuration mismatches occur and you include all necessary functions.

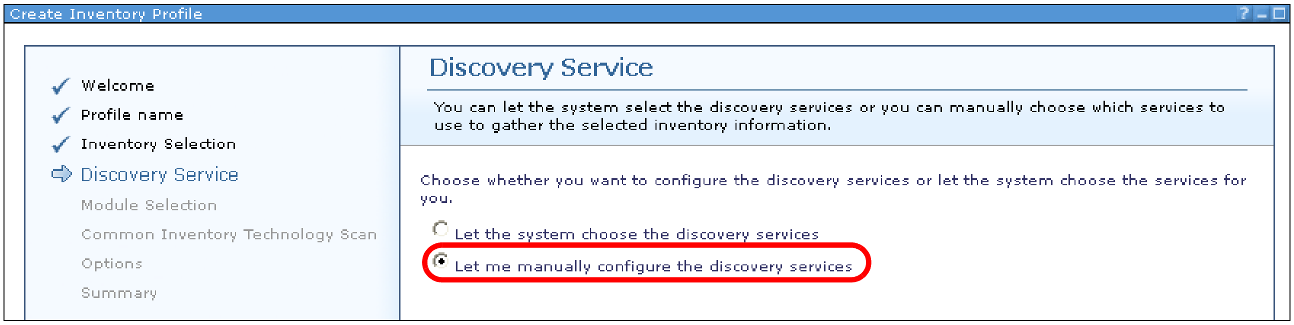

Discovery Service

You can let the system select the discovery services or you can manually choose which services to use to gather the selected inventory information.

Choose whether you want to configure the discovery services or let the system choose the services for you.

○ Let the system choose the discovery services

● Let me manually configure the discovery services

Figure 4-17 Select Discovery Service page

The example that is displayed in Figure 4-18 shows the available functions of this wizard. We choose the manual configuration. In the example, only one discovery service is available, the **CIT Software Discovery** module, which we select.

Select the modules (if more than one service is available) that you want to use and click **Next**.

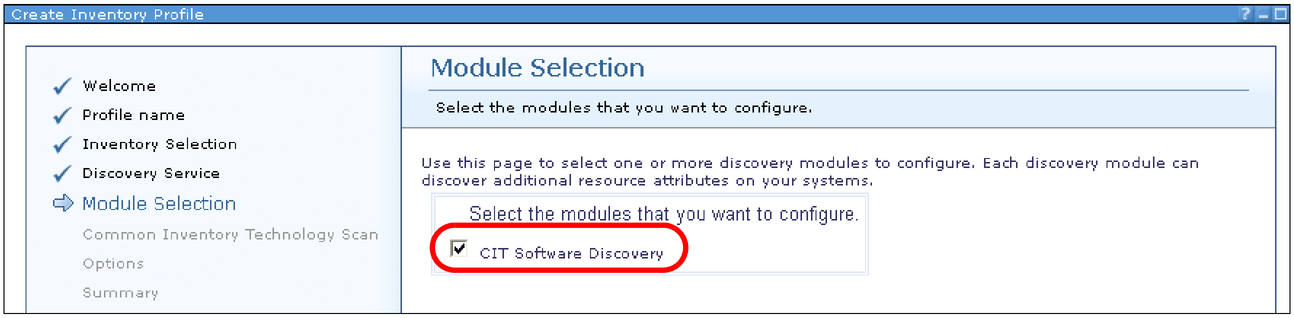

Module Selection

Select the modules that you want to configure.

Use this page to select one or more discovery modules to configure. Each discovery module can discover additional resource attributes on your systems.

Select the modules that you want to configure.

☑ CIT Software Discovery

Figure 4-18 Module Selection page

6. In our example, the option menu for the CIT Software Discovery module opens (Figure 4-20 on page 90). You can select whether you want to use the registry, the catalog, or both for the inventory collection. In our example, we select **Use both**.

   If you select the registry, the registry information from the system is used to collect software inventory information. If you select the catalog, the internal software catalog is used to collect software inventory information. When you choose both options for the software inventory collection, the CIT Software Discovery module checks the registry and the catalog to collect software information.

The default software signature file is the `softwaresignature.xml` file (Figure 4-19).

```
<!--
Licensed Materials - Property of IBM
(C) Copyright IBM Corp. 2010, 2011 All Rights Reserved
US Government Users Restricted Rights - Use, duplicate or disclosure
restricted by GSA ADP Schedule Contract with IBM Corp.
-->
<!-- IBM_COPYRIGHT_END -->
```

Figure 4-19   softwaresignature.xml file

If other modules are available, you see the option menus for these modules. Make your selection and click **Next**. Figure 4-20 shows the process of selecting the inventory collection method.

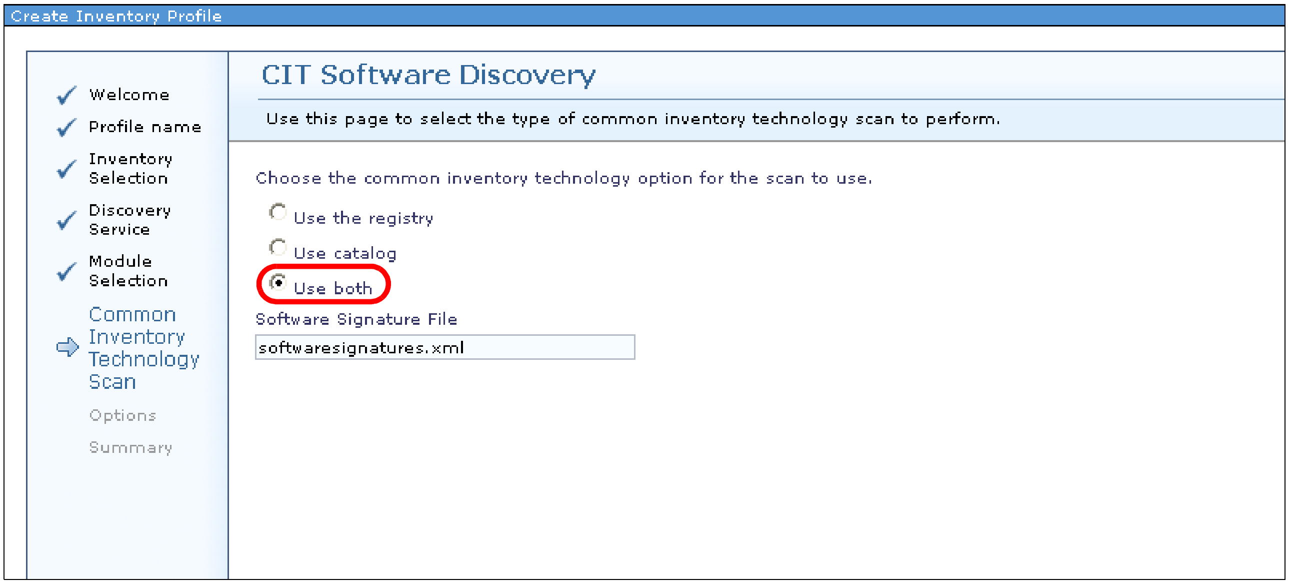

CIT Software Discovery

Use this page to select the type of common inventory technology scan to perform.

Choose the common inventory technology option for the scan to use.

○ Use the registry

○ Use catalog

● Use both

Software Signature File

softwaresignatures.xml

Figure 4-20   CIT Software Discovery options page

7. In the Options panel, Figure 4-21, you can define the timeout period and the number of simultaneous collections.

   Timeout period describes the length of time to wait for a response to inventory collection communications that are sent to systems. If the timeout value elapses before the response is received from the destination, no inventory data is collected from that target.

   Maximum simultaneous collections describe the maximum number of agents from which the Systems Director server can simultaneously collect inventory. To help reduce network traffic, specify the lowest possible number of agents.

   A check box asks whether you want to try failed agents again. If you select this function, Systems Director automatically tries again after failed collection attempts.

   Click **Next** to continue.

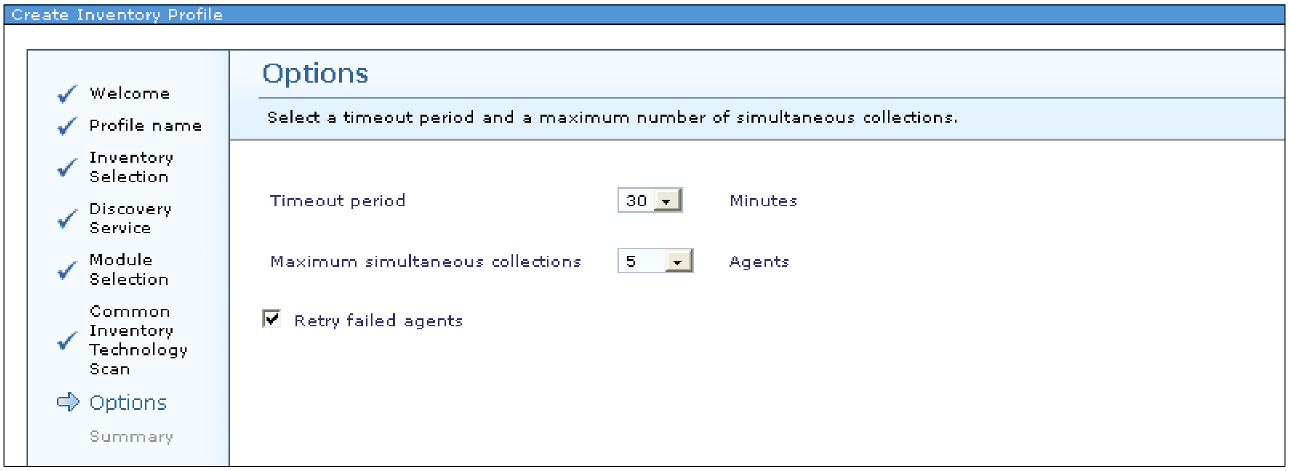
Create Inventory Profile

Options

Select a timeout period and a maximum number of simultaneous collections.

Timeout period 30 Minutes

Maximum simultaneous collections 5 Agents

☑ Retry failed agents

Figure 4-21 Options panel

8. The summary window for the wizard opens (Figure 4-22). This view is where you can verify the settings that you entered. Click **Finish** at the bottom of the window (not shown) to save the Inventory Discovery profile that you created.

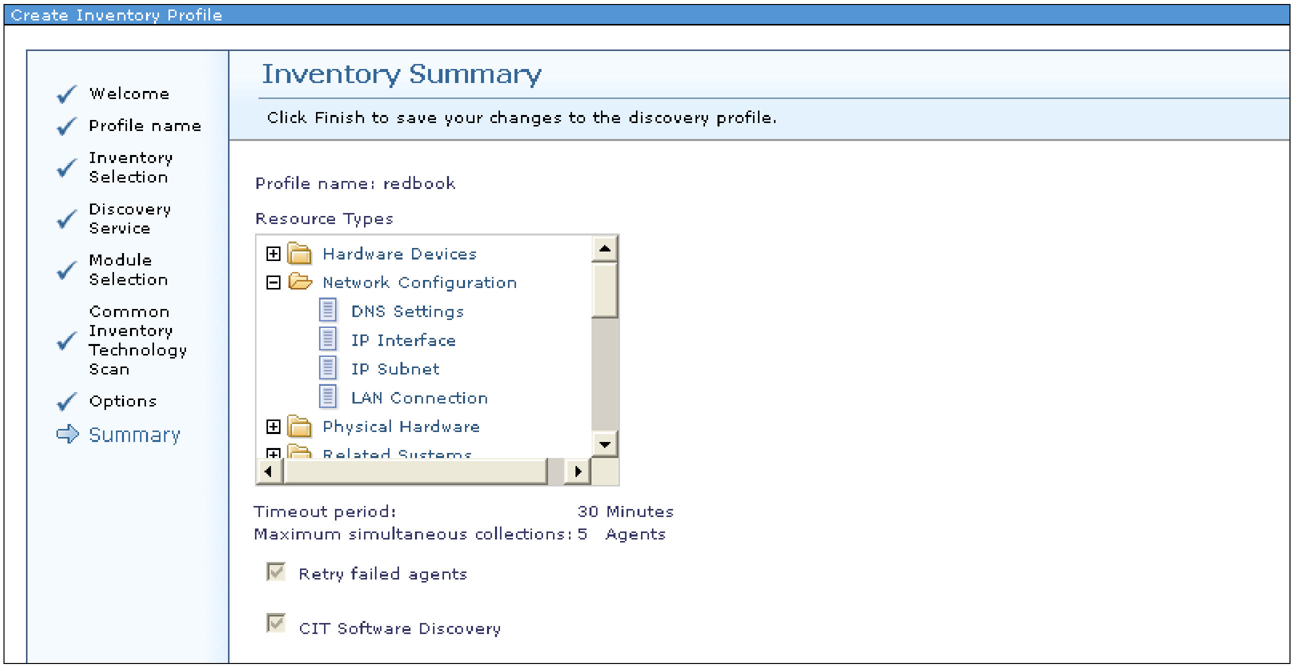

Inventory Summary

Click Finish to save your changes to the discovery profile.

Profile name: redbook

Resource Types

- ⊞ Hardware Devices
- ⊟ Network Configuration
  - DNS Settings
  - IP Interface
  - IP Subnet
  - LAN Connection
- ⊞ Physical Hardware
- ⊞ Related Systems

Timeout period: 30 Minutes
Maximum simultaneous collections: 5 Agents

☑ Retry failed agents

☑ CIT Software Discovery

Figure 4-22 Summary view

9. The list of available profiles for the inventory collection displays, including the profile that you created (Figure 4-23). In this window, you can edit or delete existing profiles.

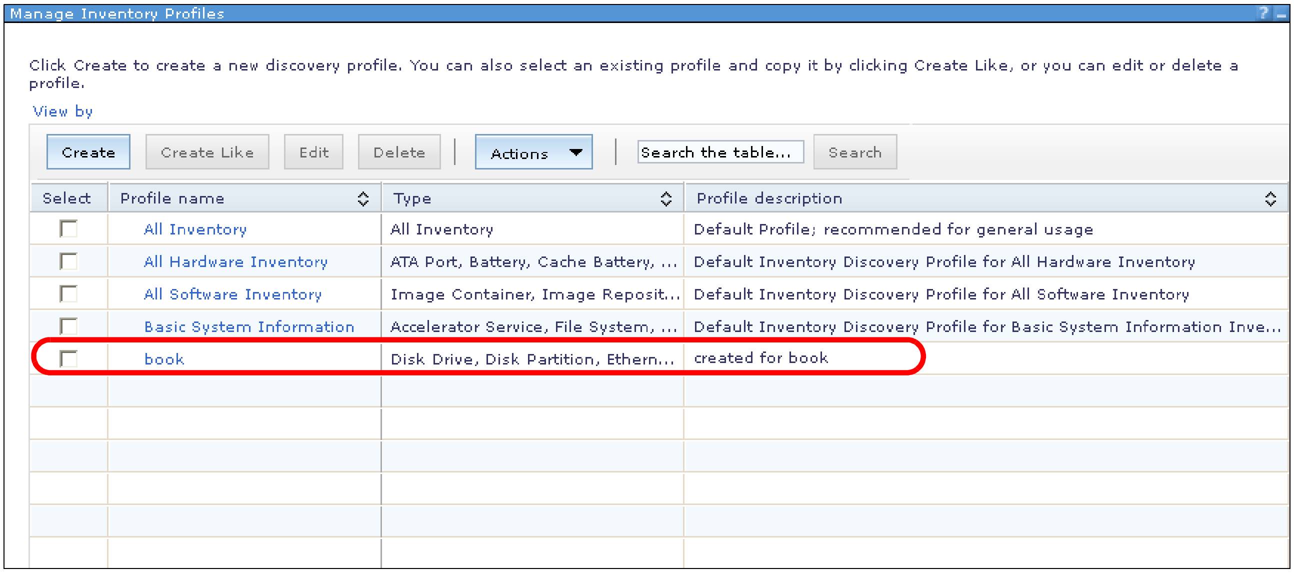

Click Create to create a new discovery profile. You can also select an existing profile and copy it by clicking Create Like, or you can edit or delete a profile.

Select	Profile name	Type	Profile description
□	All Inventory	All Inventory	Default Profile; recommended for general usage
□	All Hardware Inventory	ATA Port, Battery, Cache Battery, ...	Default Inventory Discovery Profile for All Hardware Inventory
□	All Software Inventory	Image Container, Image Reposit...	Default Inventory Discovery Profile for All Software Inventory
□	Basic System Information	Accelerator Service, File System, ...	Default Inventory Discovery Profile for Basic System Information Inve...
□	book	Disk Drive, Disk Partition, Ethern...	created for book

Figure 4-23 Inventory discovery profiles

4.2.2 Collecting inventory

Before you can view inventory for a resource, you must discover that resource by using discovery.

Inventory collection uses inventory collection profiles. You can use an existing profile to collect inventory for a system. If the inventory collection profile does not exist for the inventory data type that you want to collect, first create the inventory collection profile. Ensure that the inventory collection profile contains the appropriate settings.

Follow these steps to perform an inventory collection:

1. Launch *View and Collect Inventory*. Systems Director offers you various ways to initiate this task:

   – On the home page, on the Initial Setup tab, click **Collect Inventory**, as shown in Figure 4-24.

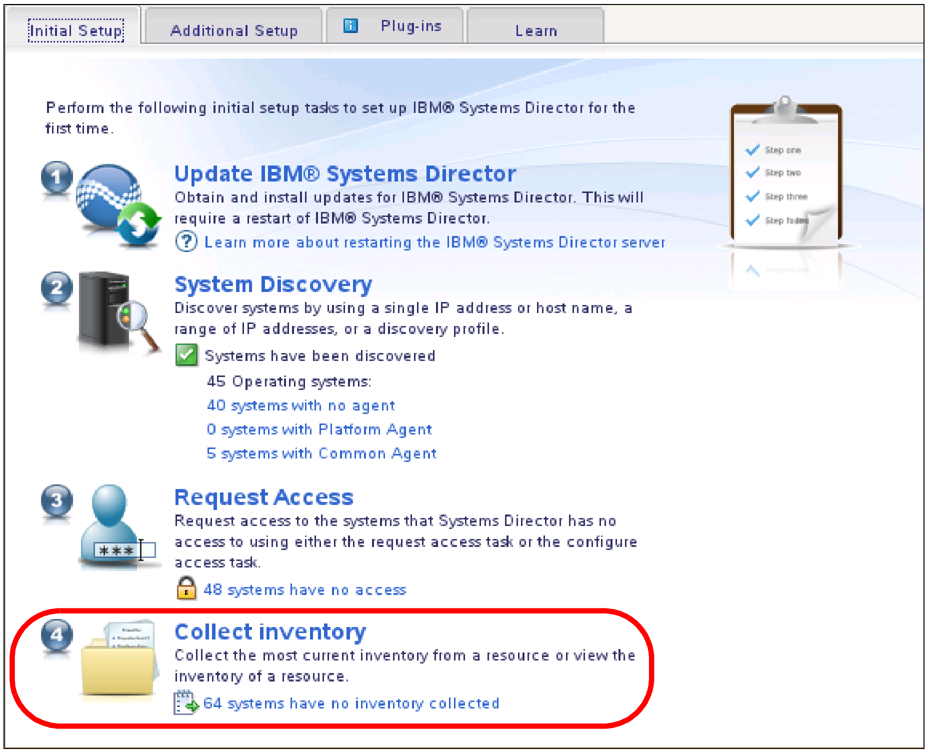

Initial Setup | Additional Setup | Plug-ins | Learn

Perform the following initial setup tasks to set up IBM® Systems Director for the first time.

1 **Update IBM® Systems Director**
Obtain and install updates for IBM® Systems Director. This will require a restart of IBM® Systems Director.
Learn more about restarting the IBM® Systems Director server

2 **System Discovery**
Discover systems by using a single IP address or host name, a range of IP addresses, or a discovery profile.
Systems have been discovered
45 Operating systems:
40 systems with no agent
0 systems with Platform Agent
5 systems with Common Agent

3 **Request Access**
Request access to the systems that Systems Director has no access to using either the request access task or the configure access task.
48 systems have no access

4 **Collect inventory**
Collect the most current inventory from a resource or view the inventory of a resource.
64 systems have no inventory collected

Figure 4-24 Inventory collection from the Initial Setup tab

– On the left-most tasks panel, click **Inventory** → **View and Collect Inventory**, as shown in Figure 4-25.

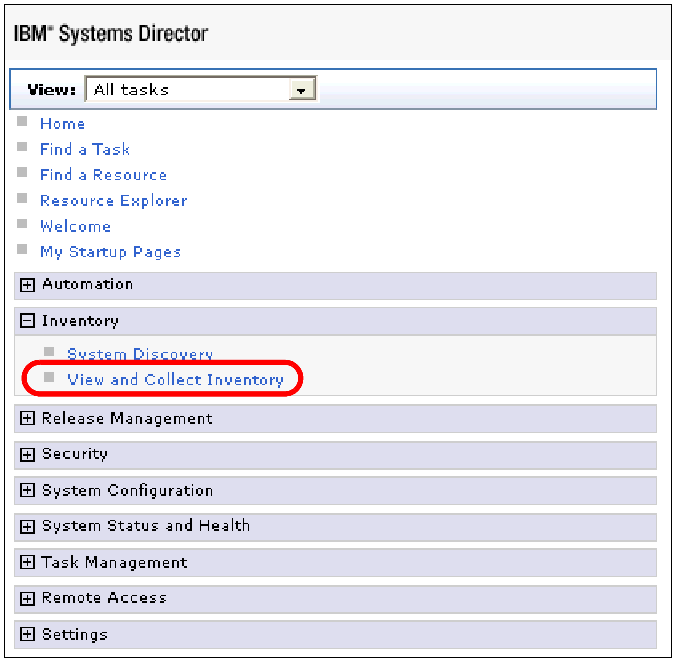

IBM® Systems Director

View: All tasks

- Home
- Find a Task
- Find a Resource
- Resource Explorer
- Welcome
- My Startup Pages

⊞ Automation

⊟ Inventory

- System Discovery
- View and Collect Inventory

⊞ Release Management

⊞ Security

⊞ System Configuration

⊞ System Status and Health

⊞ Task Management

⊞ Remote Access

⊞ Settings

Figure 4-25 Select the View and Collect Inventory option from the left pane

– From the Systems Director home page, click the **Plug-ins** tab and under Discovery Manager, click **View and Collect Inventory**, as shown in Figure 4-26.

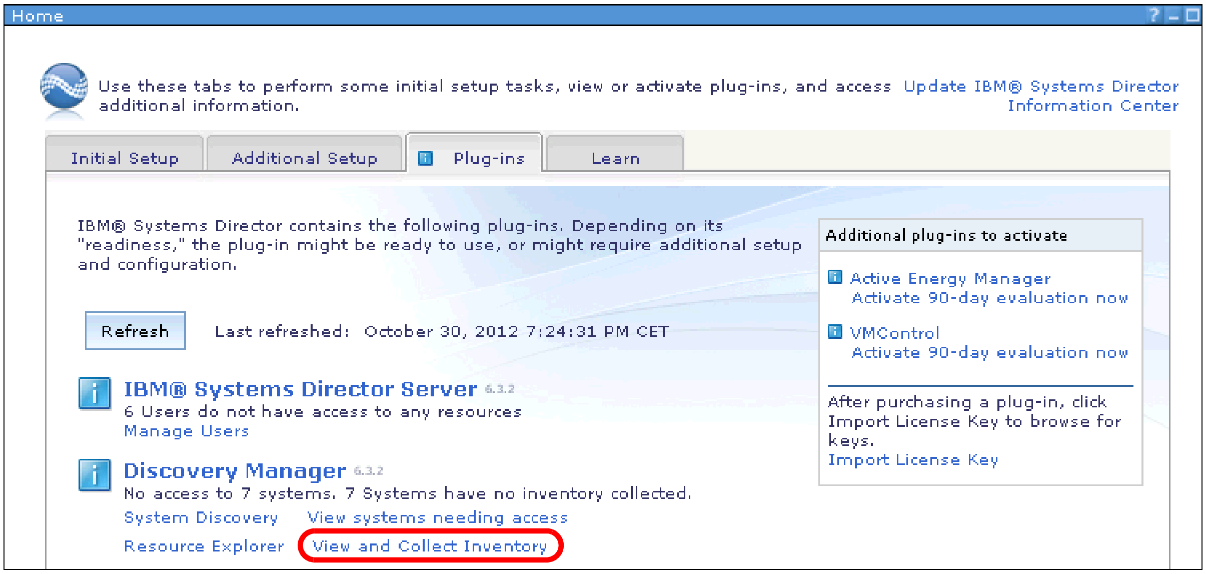

Home

Use these tabs to perform some initial setup tasks, view or activate plug-ins, and access additional information.

Update IBM® Systems Director Information Center

Initial Setup | Additional Setup | Plug-ins | Learn

IBM® Systems Director contains the following plug-ins. Depending on its "readiness," the plug-in might be ready to use, or might require additional setup and configuration.

Refresh   Last refreshed: October 30, 2012 7:24:31 PM CET

IBM® Systems Director Server 6.3.2
6 Users do not have access to any resources
Manage Users

Discovery Manager 6.3.2
No access to 7 systems. 7 Systems have no inventory collected.
System Discovery   View systems needing access
Resource Explorer   View and Collect Inventory

Additional plug-ins to activate

Active Energy Manager
Activate 90-day evaluation now

VMControl
Activate 90-day evaluation now

After purchasing a plug-in, click Import License Key to browse for keys.
Import License Key

Figure 4-26 View and Collect Inventory from home page

– In Resource Explorer, right-click a group or system. Then, from the menu, click **Inventory** → **View and Collect Inventory** (Figure 4-27).

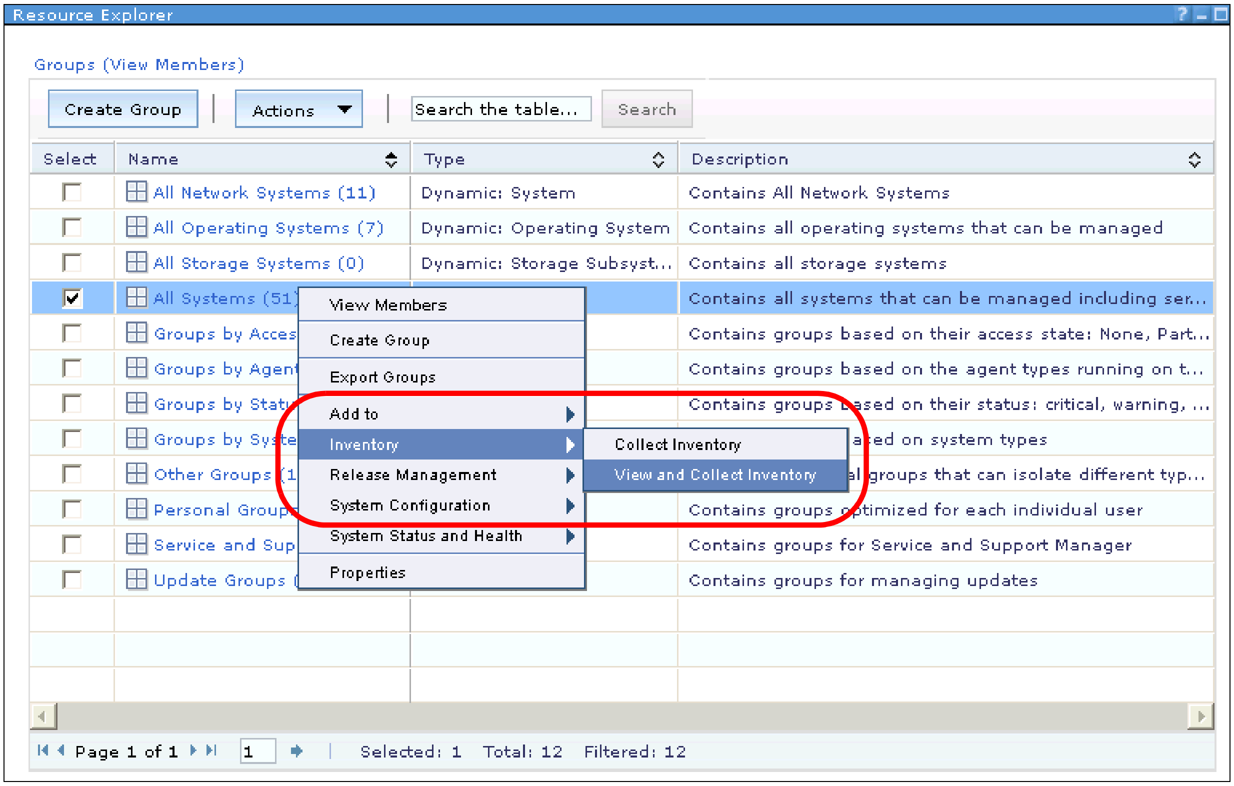

Select	Name	Type	Description
□	All Network Systems (11)	Dynamic: System	Contains All Network Systems
□	All Operating Systems (7)	Dynamic: Operating System	Contains all operating systems that can be managed
□	All Storage Systems (0)	Dynamic: Storage Subsyst...	Contains all storage systems
☑	All Systems (51		Contains all systems that can be managed including ser...
□	Groups by Acces		Contains groups based on their access state: None, Part...
□	Groups by Agent		Contains groups based on the agent types running on t...
□	Groups by Stat		Contains groups based on their status: critical, warning, ...
□	Groups by Syste		ased on system types
□	Other Groups (1		al groups that can isolate different typ...
□	Personal Group		Contains groups optimized for each individual user
□	Service and Sup		Contains groups for Service and Support Manager
□	Update Groups (		Contains groups for managing updates

Page 1 of 1 | 1 | Selected: 1 Total: 12 Filtered: 12

Figure 4-27 View and Collect Inventory from the Resource Explorer page

2. The inventory task then launches (Figure 4-28). If not preselected, select the system or group for which you want to run the inventory collection. Then, you can select an inventory discovery profile. We describe how to create a profile in 4.2.1, “Inventory data and collection profiles” on page 86. Click **Collect Inventory** to start the process.

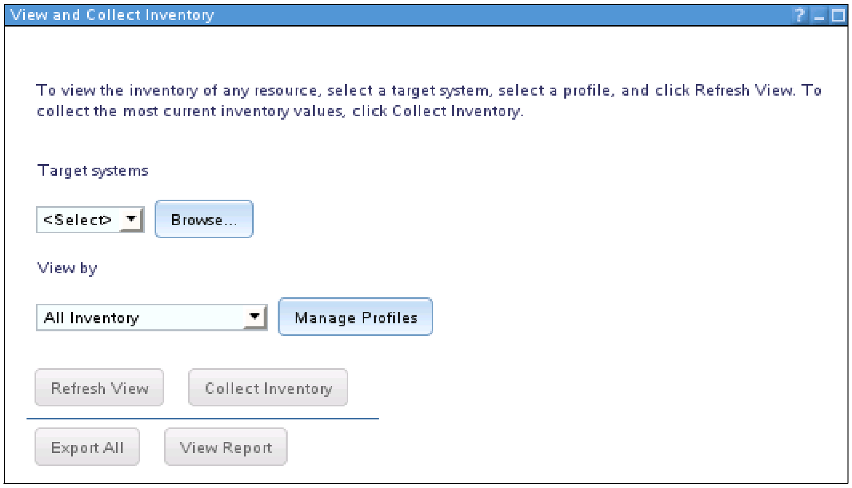

To view the inventory of any resource, select a target system, select a profile, and click Refresh View. To collect the most current inventory values, click Collect Inventory.

Target systems

<Select>

View by

All Inventory

Figure 4-28 Inventory collection

3. A scheduler window opens (Figure 4-29). Select **Run Now** or specify a time to run the inventory collection.

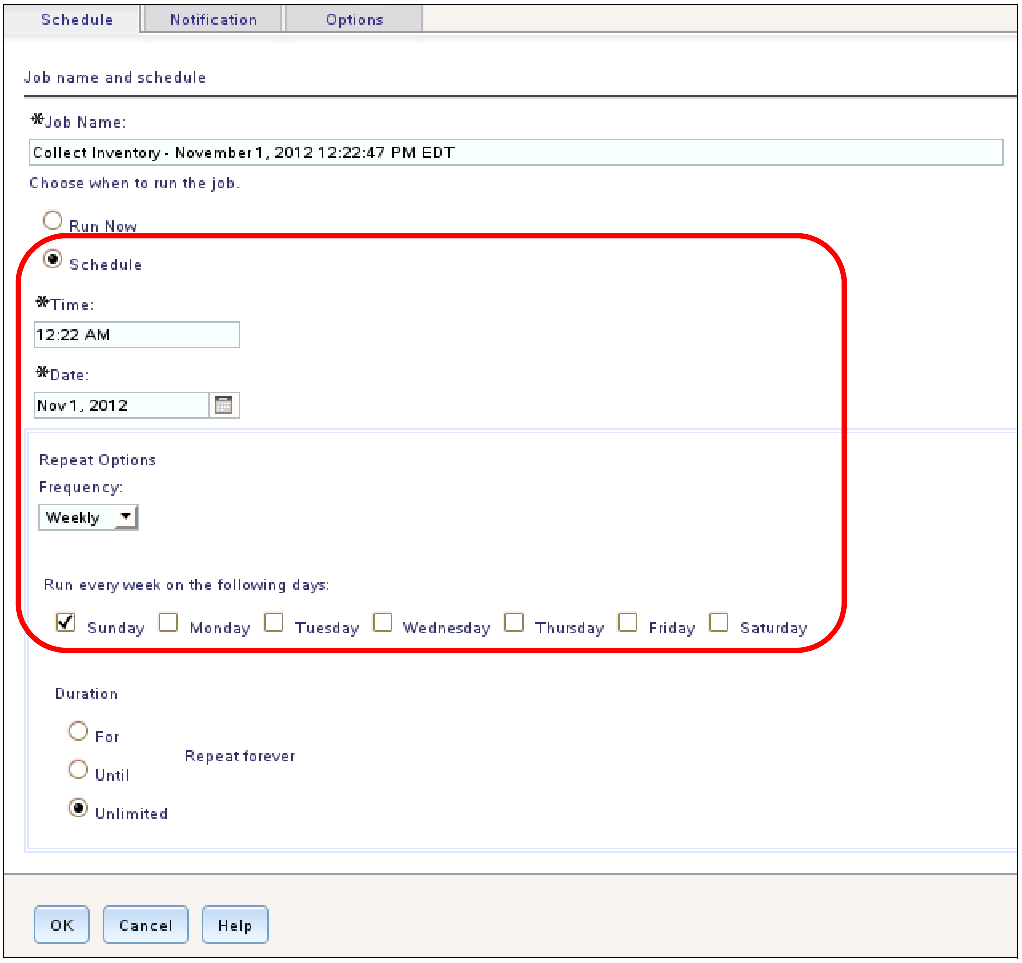

Job name and schedule

*Job Name:
Collect Inventory - November 1, 2012 12:22:47 PM EDT

Choose when to run the job.

○ Run Now

● Schedule

*Time:
12:22 AM

*Date:
Nov 1, 2012

Repeat Options

Frequency:
Weekly

Run every week on the following days:

☑ Sunday □ Monday □ Tuesday □ Wednesday □ Thursday □ Friday □ Saturday

Duration

○ For

○ Until

● Unlimited

Repeat forever

Figure 4-29 Scheduler example for weekly, Sunday 12:22 a.m. setting

Schedule an inventory collection once a week. Schedule this inventory collection in off-hours so that the inventory collection does not affect your daily business. Also, run an inventory scan when you plan to update systems or install agents on the system.

4. After the schedule is defined or you select **Run Now**, click **OK**.
5. You can see the status of the task in the left pane under **Task Management** → **Active and Scheduled Task**.

4.2.3 Viewing inventory

After the inventory task completes, you can view the results. The following examples show different profiles.

> **Tip:** When you select a single system, a summary shows at the top of the inventory information. If you select a group of systems, no summary is shown.

In our examples, we used the All Systems group to discover the inventory:

- Basic file system information (Figure 4-30)

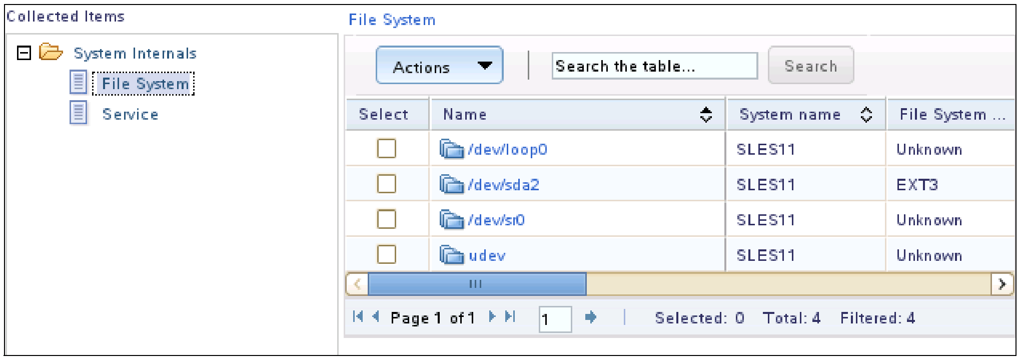

Collected Items

- System Internals
  - File System
  - Service

File System

Actions ▼ | Search the table... | Search

Select	Name	System name	File System ...
☐	/dev/loop0	SLES11	Unknown
☐	/dev/sda2	SLES11	EXT3
☐	/dev/sr0	SLES11	Unknown
☐	udev	SLES11	Unknown

Page 1 of 1 | 1 | Selected: 0 Total: 4 Filtered: 4

Figure 4-30 Basic file system information

- All software inventory (Figure 4-31)

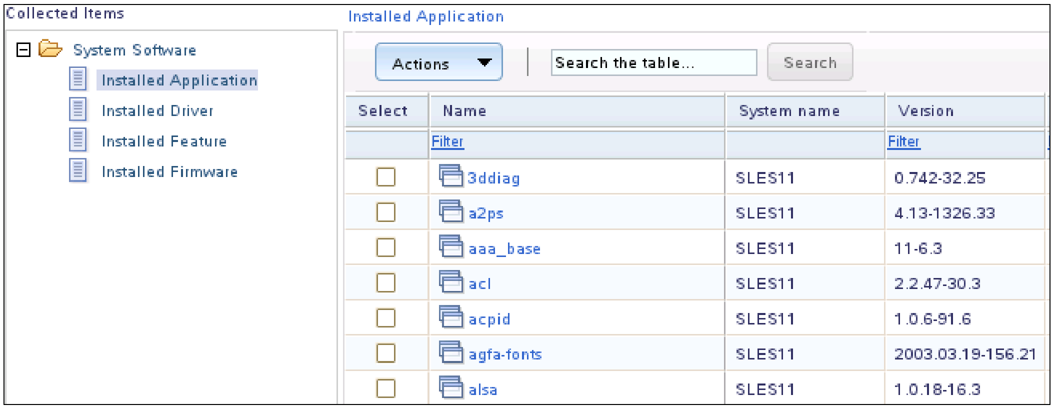

Collected Items

- System Software
  - Installed Application
  - Installed Driver
  - Installed Feature
  - Installed Firmware

Installed Application

Actions ▼ | Search the table... | Search

Select	Name	System name	Version
	Filter		Filter
☐	3ddiag	SLES11	0.742-32.25
☐	a2ps	SLES11	4.13-1326.33
☐	aaa_base	SLES11	11-6.3
☐	acl	SLES11	2.2.47-30.3
☐	acpid	SLES11	1.0.6-91.6
☐	agfa-fonts	SLES11	2003.03.19-156.21
☐	alsa	SLES11	1.0.18-16.3

Figure 4-31 All software inventory

- All hardware inventory (Figure 4-32)

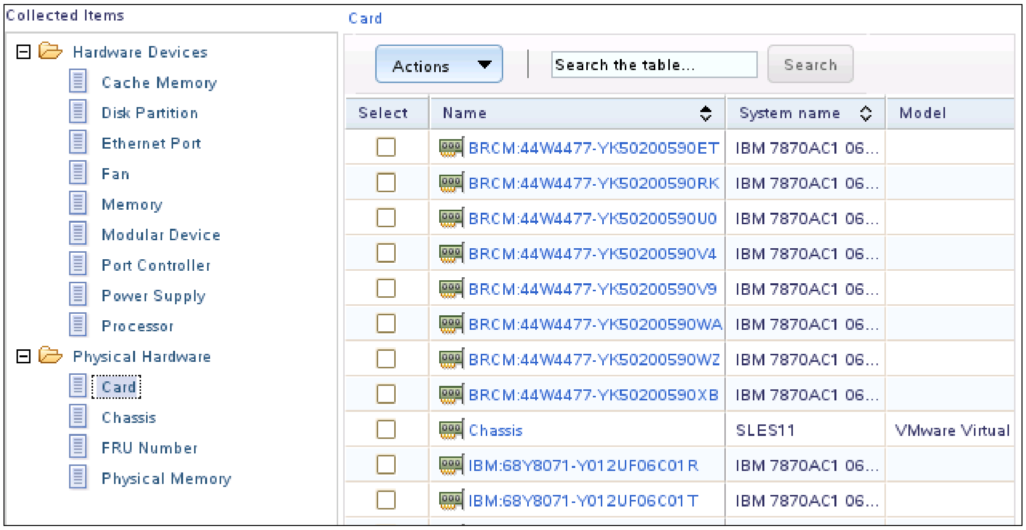

Collected Items

- Hardware Devices
  - Cache Memory
  - Disk Partition
  - Ethernet Port
  - Fan
  - Memory
  - Modular Device
  - Port Controller
  - Power Supply
  - Processor
- Physical Hardware
  - Card
  - Chassis
  - FRU Number
  - Physical Memory

Card

Actions ▼ | Search the table... | Search

Select	Name	System name	Model
☐	BRCM:44W4477-YK50200590ET	IBM 7870AC1 06...	
☐	BRCM:44W4477-YK50200590RK	IBM 7870AC1 06...	
☐	BRCM:44W4477-YK50200590U0	IBM 7870AC1 06...	
☐	BRCM:44W4477-YK50200590V4	IBM 7870AC1 06...	
☐	BRCM:44W4477-YK50200590V9	IBM 7870AC1 06...	
☐	BRCM:44W4477-YK50200590WA	IBM 7870AC1 06...	
☐	BRCM:44W4477-YK50200590WZ	IBM 7870AC1 06...	
☐	BRCM:44W4477-YK50200590XB	IBM 7870AC1 06...	
☐	Chassis	SLES11	VMware Virtual
☐	IBM:68Y8071-Y012UF06C01R	IBM 7870AC1 06...	
☐	IBM:68Y8071-Y012UF06C01T	IBM 7870AC1 06...	

Figure 4-32 All hardware inventory

- All inventory (Figure 4-33)

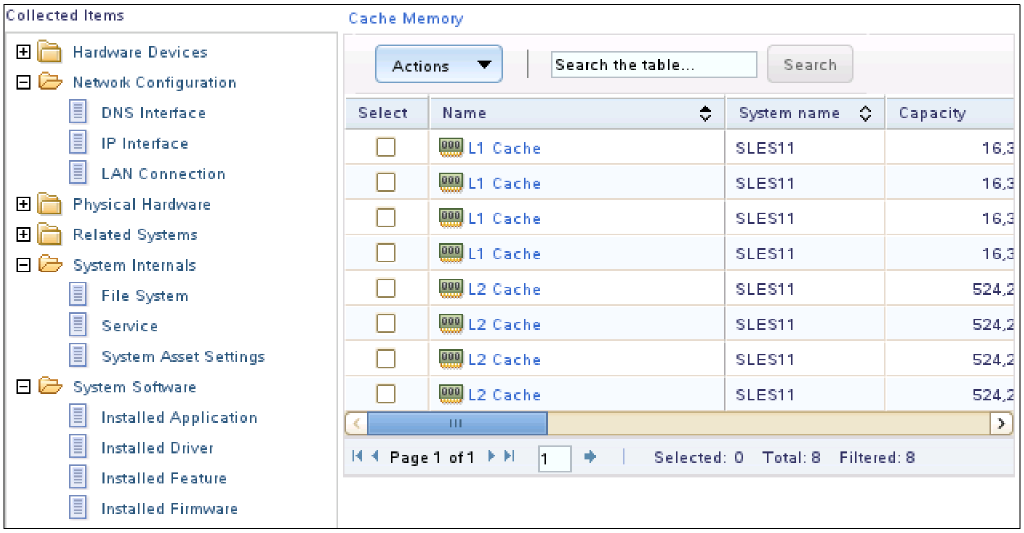

Collected Items

- ⊞ Hardware Devices
- ⊟ Network Configuration
  - DNS Interface
  - IP Interface
  - LAN Connection
- ⊞ Physical Hardware
- ⊞ Related Systems
- ⊟ System Internals
  - File System
  - Service
  - System Asset Settings
- ⊟ System Software
  - Installed Application
  - Installed Driver
  - Installed Feature
  - Installed Firmware

Cache Memory

Select	Name	System name	Capacity
□	L1 Cache	SLES11	16,3
□	L1 Cache	SLES11	16,3
□	L1 Cache	SLES11	16,3
□	L1 Cache	SLES11	16,3
□	L2 Cache	SLES11	524,2
□	L2 Cache	SLES11	524,2
□	L2 Cache	SLES11	524,2
□	L2 Cache	SLES11	524,2

Page 1 of 1 | 1 | Selected: 0 Total: 8 Filtered: 8

Figure 4-33 All inventory

If you select a single system, a system summary shows at the top of the inventory information. This summary provides an overview of the system information:

- Operating system summary
- Network configuration summary
- Systems Director Agent version that is installed on the system
- Access state
- Supported protocols
- Firmware information

In Figure 4-34, we show a system that runs SLES11. The Systems Director server 6.3.2 is installed. The system runs on a virtual machine that is hosted by VMware ESXi.

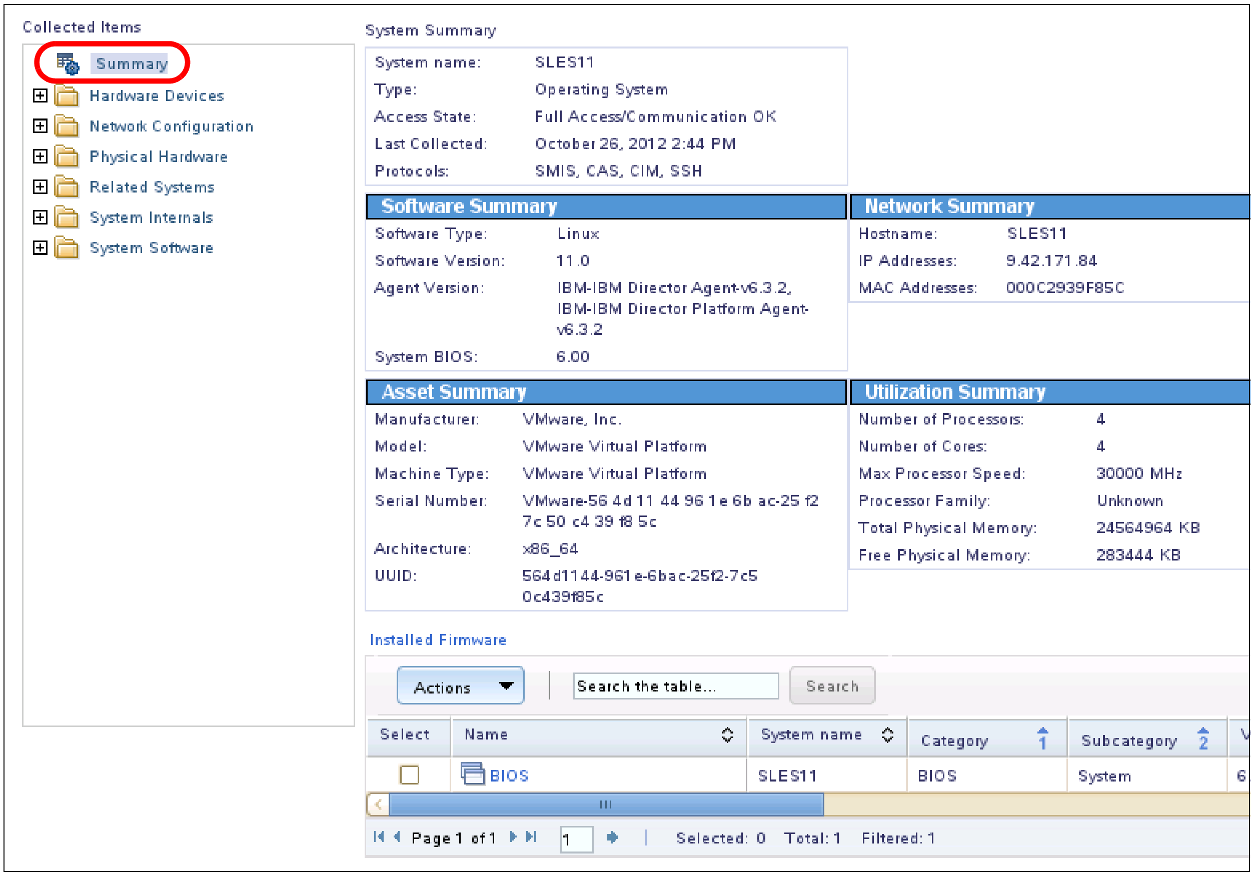

System Summary

System name: SLES11
Type: Operating System
Access State: Full Access/Communication OK
Last Collected: October 26, 2012 2:44 PM
Protocols: SMIS, CAS, CIM, SSH

Software Summary

Software Type: Linux
Software Version: 11.0
Agent Version: IBM-IBM Director Agent-v6.3.2, IBM-IBM Director Platform Agent-v6.3.2
System BIOS: 6.00

Network Summary

Hostname: SLES11
IP Addresses: 9.42.171.84
MAC Addresses: 000C2939F85C

Asset Summary

Manufacturer: VMware, Inc.
Model: VMware Virtual Platform
Machine Type: VMware Virtual Platform
Serial Number: VMware-56 4d 11 44 96 1e 6b ac-25 f2 7c 50 c4 39 f8 5c
Architecture: x86_64
UUID: 564d1144-961e-6bac-25f2-7c50c439f85c

Utilization Summary

Number of Processors: 4
Number of Cores: 4
Max Processor Speed: 30000 MHz
Processor Family: Unknown
Total Physical Memory: 24564964 KB
Free Physical Memory: 283444 KB

Installed Firmware

Select	Name	System name	Category	Subcategory	V
□	BIOS	SLES11	BIOS	System	6

Figure 4-34 Inventory summary view (only available for a single system selection)

4.2.4 Exporting inventory

To use the inventory information outside the Systems Director, you can export the inventory information for a system or group. This function might be useful to perform asset management tasks that are external to Systems Director. Or, this function can be useful if you want to print the inventory report for documentation.

Follow these steps to export the inventory data:

1. From the View and Collect Inventory page, click **Export All** (Figure 4-35).

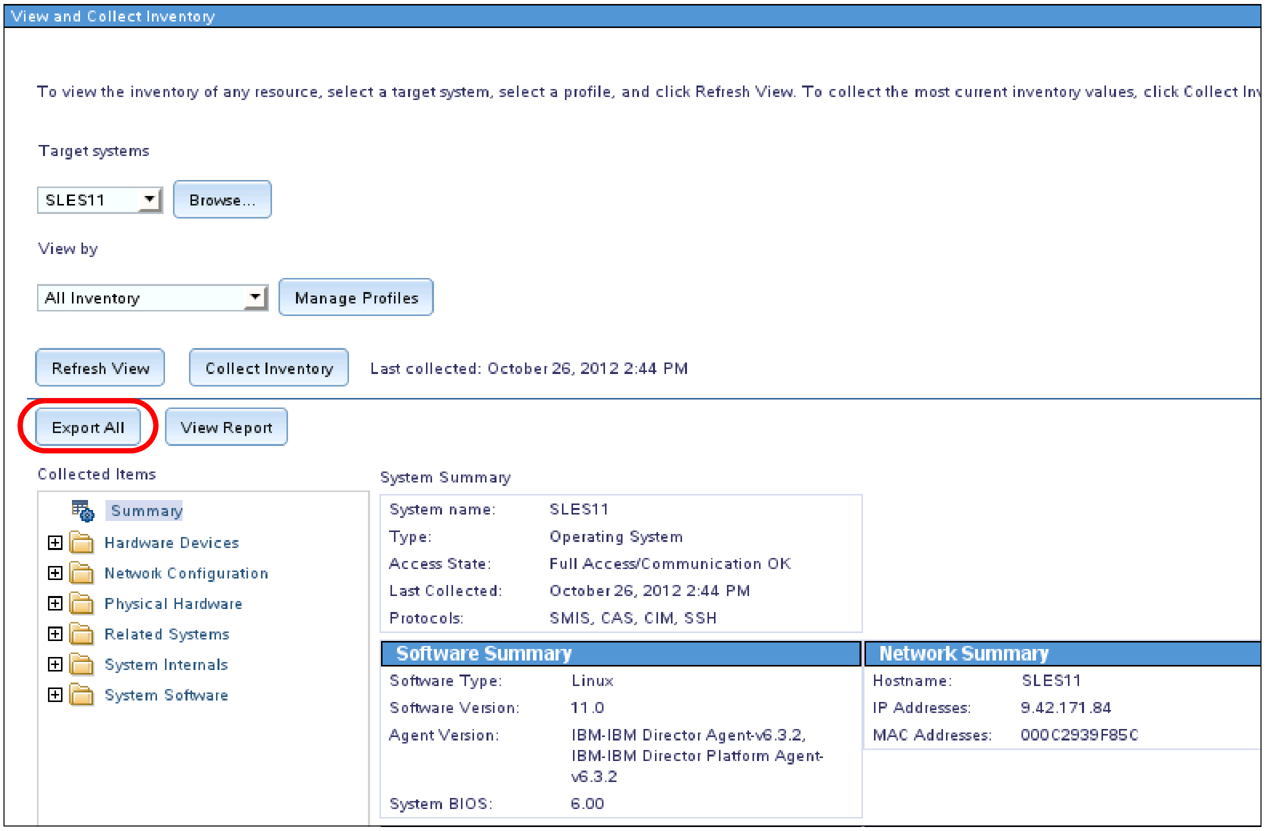

View and Collect Inventory

To view the inventory of any resource, select a target system, select a profile, and click Refresh View. To collect the most current inventory values, click Collect In

Target systems: SLES11

View by: All Inventory

Last collected: October 26, 2012 2:44 PM

System Summary

System name:	SLES11
Type:	Operating System
Access State:	Full Access/Communication OK
Last Collected:	October 26, 2012 2:44 PM
Protocols:	SMIS, CAS, CIM, SSH

Software Summary

Software Type:	Linux
Software Version:	11.0
Agent Version:	IBM-IBM Director Agent-v6.3.2, IBM-IBM Director Platform Agent-v6.3.2
System BIOS:	6.00

Network Summary

Hostname:	SLES11
IP Addresses:	9.42.171.84
MAC Addresses:	000C2939F85C

Figure 4-35 Select Export All to export inventory data

2. Choose the format to which to export the inventory data (Figure 4-36). Various formats are available to export your data:

   – Hypertext Markup Language (HTML)
   – Extensible Markup Language (XML)
   – Comma Separated Variable (CSV)

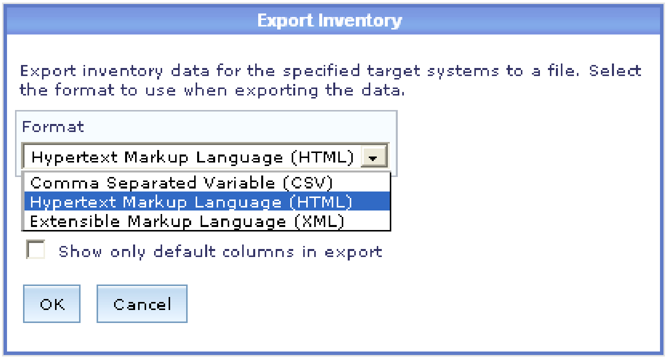

Export Inventory

Export inventory data for the specified target systems to a file. Select the format to use when exporting the data.

Format: Hypertext Markup Language (HTML)

Comma Separated Variable (CSV)
Hypertext Markup Language (HTML)
Extensible Markup Language (XML)
□ Show only default columns in export
OK
Cancel

Figure 4-36 Select the format for export

3. After you select the format in which you want to export the inventory data, click **OK**.

   If the HTML or XML format is selected, a web page that contains the data opens in your browser. You can save this data to a file or print the data, as needed. In this example, we use HTML as the file format (Figure 4-37).

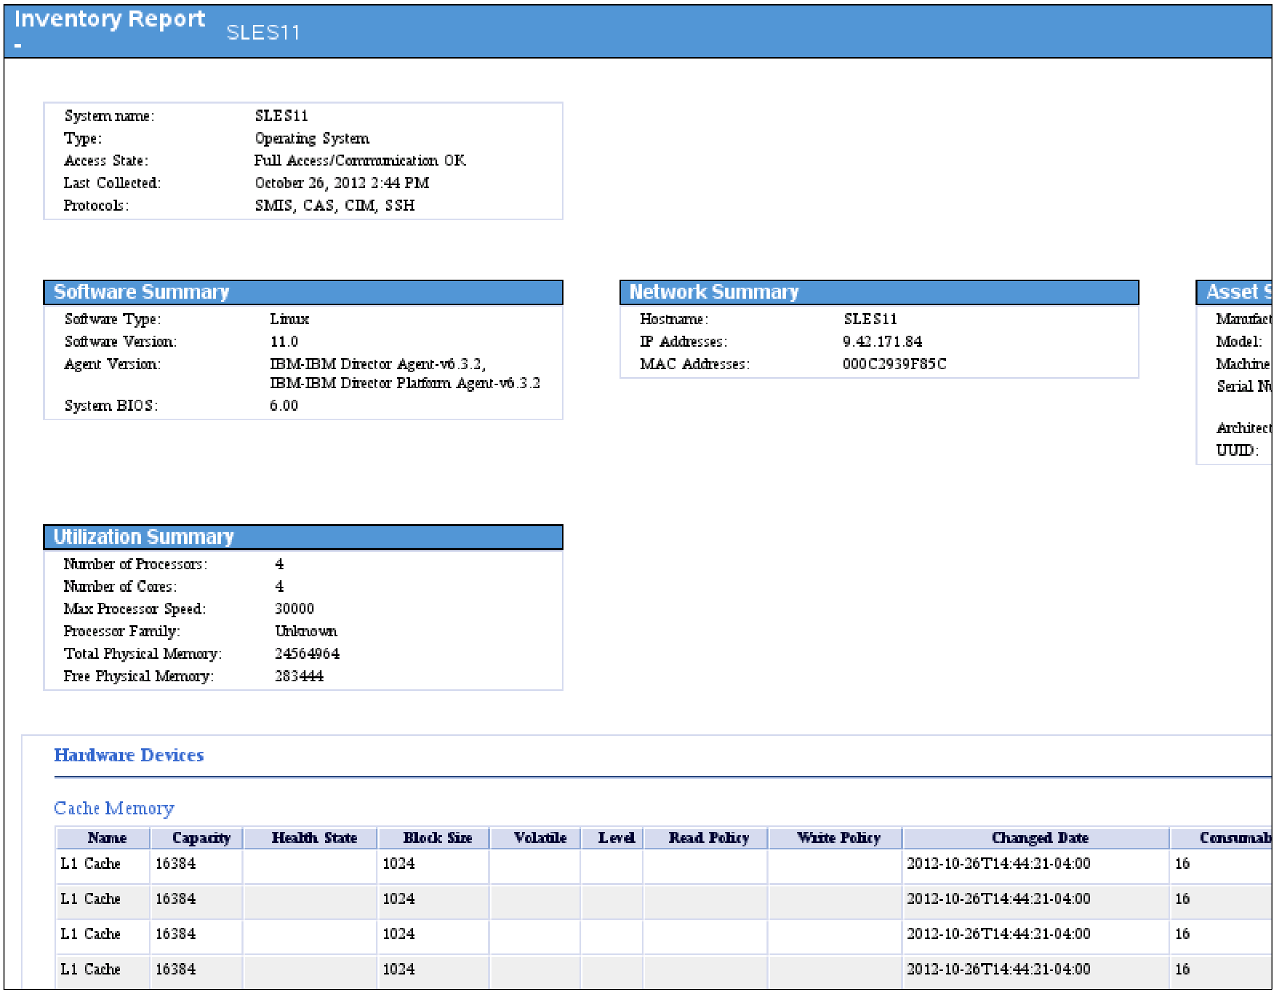

Inventory Report - SLES11

System name:	SLES11
Type:	Operating System
Access State:	Full Access/Communication OK
Last Collected:	October 26, 2012 2:44 PM
Protocols:	SMIS, CAS, CIM, SSH

Software Summary

Software Type:	Linux
Software Version:	11.0
Agent Version:	IBM-IBM Director Agent-v6.3.2, IBM-IBM Director Platform Agent-v6.3.2
System BIOS:	6.00

Network Summary

Hostname:	SLES11
IP Addresses:	9.42.171.84
MAC Addresses:	000C2939F85C

Asset S

Manufact
Model:
Machine
Serial N
Architect
UUID:

Utilization Summary

Number of Processors:	4
Number of Cores:	4
Max Processor Speed:	30000
Processor Family:	Unknown
Total Physical Memory:	24564964
Free Physical Memory:	283444

Hardware Devices

Cache Memory

Name	Capacity	Health State	Block Size	Volatile	Level	Read Policy	Write Policy	Changed Date	Consumab
L1 Cache	16384		1024					2012-10-26T14:44:21-04:00	16
L1 Cache	16384		1024					2012-10-26T14:44:21-04:00	16
L1 Cache	16384		1024					2012-10-26T14:44:21-04:00	16
L1 Cache	16384		1024					2012-10-26T14:44:21-04:00	16

Figure 4-37 HTML export

If the CSV format is selected, you can save or open the data with available applications as detected by your browser (Figure 4-38).

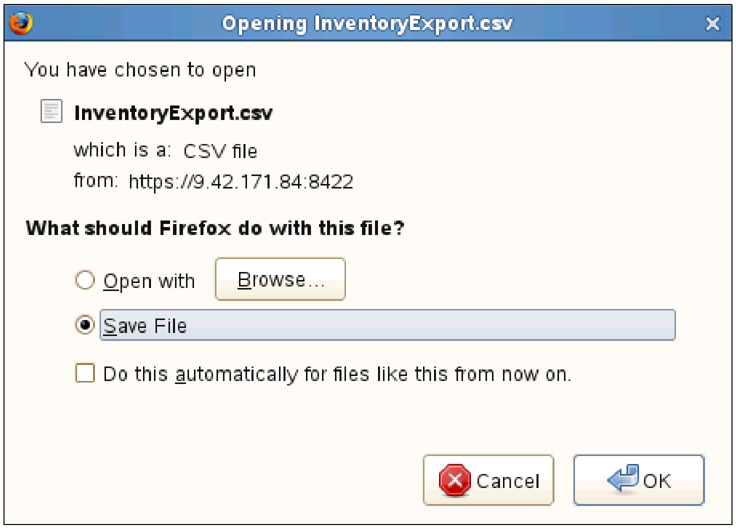

Opening InventoryExport.csv

You have chosen to open

InventoryExport.csv

which is a: CSV file

from: https://9.42.171.84:8422

What should Firefox do with this file?

○ Open with [Browse...]

● Save File

□ Do this automatically for files like this from now on.

Figure 4-38 Save the CSV file

The CSV file can be used, for example, to import this data into an Excel worksheet (Figure 4-39).

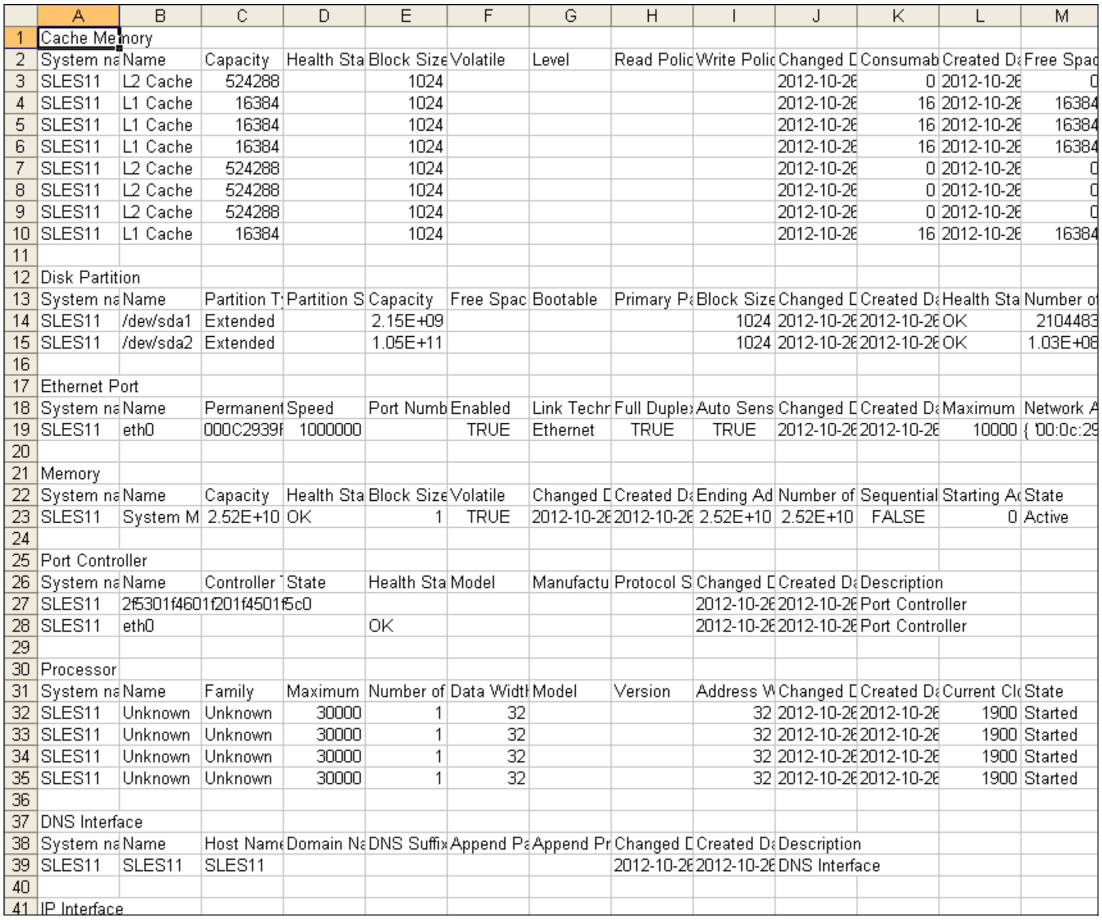

	A	B	C	D	E	F	G	H	I	J	K	L	M
1	Cache Memory												
2	System na	Name	Capacity	Health Sta	Block Size	Volatile	Level	Read Polic	Write Polic	Changed D	Consumab	Created Da	Free Spac
3	SLES11	L2 Cache	524288		1024					2012-10-2€	0	2012-10-2€	0
4	SLES11	L1 Cache	16384		1024					2012-10-2€	16	2012-10-2€	16384
5	SLES11	L1 Cache	16384		1024					2012-10-2€	16	2012-10-2€	16384
6	SLES11	L1 Cache	16384		1024					2012-10-2€	16	2012-10-2€	16384
7	SLES11	L2 Cache	524288		1024					2012-10-2€	0	2012-10-2€	0
8	SLES11	L2 Cache	524288		1024					2012-10-2€	0	2012-10-2€	0
9	SLES11	L2 Cache	524288		1024					2012-10-2€	0	2012-10-2€	0
10	SLES11	L1 Cache	16384		1024					2012-10-2€	16	2012-10-2€	16384
11													
12	Disk Partition												
13	System na	Name	Partition Ty	Partition S	Capacity	Free Spac	Bootable	Primary Pa	Block Size	Changed D	Created Da	Health Sta	Number of
14	SLES11	/dev/sda1	Extended		2.15E+09				1024	2012-10-2€	2012-10-2€	OK	2104483
15	SLES11	/dev/sda2	Extended		1.05E+11				1024	2012-10-2€	2012-10-2€	OK	1.03E+08
16													
17	Ethernet Port												
18	System na	Name	Permanent	Speed	Port Numb	Enabled	Link Techr	Full Duplex	Auto Sens	Changed D	Created Da	Maximum	Network A
19	SLES11	eth0	000C2939F	1000000		TRUE	Ethernet	TRUE	TRUE	2012-10-2€	2012-10-2€	10000	{ '00:0c:29
20													
21	Memory												
22	System na	Name	Capacity	Health Sta	Block Size	Volatile	Changed D	Created Da	Ending Ad	Number of	Sequential	Starting Ad	State
23	SLES11	System M	2.52E+10	OK	1	TRUE	2012-10-2€	2012-10-2€	2.52E+10	2.52E+10	FALSE	0	Active
24													
25	Port Controller												
26	System na	Name	Controller T	State	Health Sta	Model	Manufactu	Protocol S	Changed D	Created Da	Description		
27	SLES11	2f5301f4601f201f4501f5c0							2012-10-2€	2012-10-2€	Port Controller		
28	SLES11	eth0			OK				2012-10-2€	2012-10-2€	Port Controller		
29													
30	Processor												
31	System na	Name	Family	Maximum	Number of	Data Widtł	Model	Version	Address W	Changed D	Created Da	Current Clc	State
32	SLES11	Unknown	Unknown	30000	1	32			32	2012-10-2€	2012-10-2€	1900	Started
33	SLES11	Unknown	Unknown	30000	1	32			32	2012-10-2€	2012-10-2€	1900	Started
34	SLES11	Unknown	Unknown	30000	1	32			32	2012-10-2€	2012-10-2€	1900	Started
35	SLES11	Unknown	Unknown	30000	1	32			32	2012-10-2€	2012-10-2€	1900	Started
36													
37	DNS Interface												
38	System na	Name	Host Name	Domain Na	DNS Suffix	Append Pa	Append Pr	Changed D	Created Da	Description			
39	SLES11	SLES11	SLES11					2012-10-2€	2012-10-2€	DNS Interface			
40													
41	IP Interface												

Figure 4-39 Imported inventory information into an Excel worksheet by using the CSV file

4.2.5 Best practices for inventory

These are the best practices to use the inventory task efficiently:

- *Use the embedded DB2 database.*

  This database is managed by the IBM Systems Director. It is a full enterprise database.

- *Collect inventory as often as possible, and at a minimum once a week.*

  Inventory data is important for many functions in the IBM Systems Director. Only actual data keeps this function up to date and running.

- *If there are updates in your environment, run afterwards the inventory for these systems.*

  Inventory information about the installed driver and firmware are the basis for the compliance check. These checks can work successfully only if the inventory data is actual.

- *Run inventory collection scans in off-hours*.

  If you have many systems, schedule the inventory collection in off-hours to prevent problems in the network utilization in your environment.

- *Use the export function to save inventory data for reports.*

4.3 Updates

With Update Manager, a component of Systems Director, you keep the servers on your network at the software or firmware update levels that you want. Update Manager automatically checks for available updates and identifies which systems need attention. Update Manager also provides you with the ability to monitor your systems for needed updates. With Update Manager, you can schedule the updates at times that are convenient for you and your users.

Update Manager compares the update information that is loaded into it with the inventories of specified systems to determine whether updates are needed.

4.3.1 Prerequisites

Before you can start to use Update Manager to update your systems, ensure that an inventory of your system is performed. You can automate the collection of the inventory information as described in 4.2, “Inventory” on page 86.

To update your systems, the systems must be online and accessible. Therefore, you must have full access to the systems from the Systems Director server. The access state must be set to `OK`. Update Manager can be used to update agentless systems and systems with Platform Agent and Common Agent installed.

To update the BladeCenter AMM and server with Integrated Management Module I (IMMv1), you must configure a Trivial File Transfer Protocol (TFTP) server. Systems Director includes a TFTP server. See 4.3.3, “Settings for Update Manager” on page 106.

The best way to check whether your system is up-to-date or needs an update is to use the Compliance Check function, which is described in 4.3.6, “Compliance check” on page 114.

4.3.2 What can be updated

The following list shows the supported updates and the systems to which updates can be applied. Unless otherwise noted, the systems can be agentless-managed systems, Common Agent-managed systems, and Platform Agent-managed systems.

The following list shows the supported updates and systems:

- Systems Director:
  - 6.3.*x* (Common Agent, Platform Agent, and the Systems Director server)
  - 6.2.*x* and 6.1.*x* (Common Agent and Platform Agent)
  - IBM Director V5.20.*x* (IBM Director Agent version 5.20 and IBM Director Core Services version 5.20)
- Technology levels (TLs) and service packs (SPs):
  - AIX 5.3 TL6 SP5 and later (the Systems Director server or Common Agent only)
  - AIX 6.1 (the Systems Director server or Common Agent only)
- SUSE Linux
- Red Hat Enterprise Linux
- Cumulative program temporary fix (PTF) packages and PTF groups for IBM i (formerly IBM i5/OS™) 5.4 and later

- Hardware Management Console (HMC) systems at V7.3.3 SP2 or later
- Power Systems firmware for all systems that meet at least one of the following criteria:
  - Inband stand-alone (not managed by HMC or Integrated Virtualization Manager) Power Systems target systems that run AIX or Linux

    Required: These systems must have the Common Agent installed.

  - Out-of-band (managed by HMC) target systems

    No Common Agent: No Common Agent is required in this case because Secure Shell (SSH) performs the update.

  - Power Systems target systems that are managed by Integrated Virtualization Manager and that run Virtual I/O Server (VIOS) version 1.5.2.1 - Fix Pack (FP) 11.1 or later

    No Common Agent: No Common Agent is required in this case because SSH performs the update.

  - Migration, FPs, SPs, and interim fixes for VIOS version 1.5.2.1 - FP11.1 or later
- Device driver and firmware updates, or UpdateXpress System Pack updates, for System x servers that run Linux or Windows

  Support is provided for servers that run all available agent and agentless levels. No support is available for updating IMM V2 systems that run IBM Director Agent 5.*x*.

  Note: Using IBM Systems Director Update Manager on IBM System x or Flex System servers requires that you enable *LAN over USB* on the management processor. This is normally enabled by default and you can change the configuration by using the management processor web interface or by using the Advanced Settings Utility (ASU) for your operating system.

  Update Manager can update drivers only if devices have been properly discovered and installed in the operating system. For example, it will not install any adapter driver for adapters that are shown as unknown devices in the Windows device manager.

- IBM BladeCenter I/O module firmware

  Update to I/O modules must be installed from a TFTP or FTP server. The IBM Systems Director server can work as a TFTP server. (Need port 69 UDP inbound opened)

  Note: The SAS RAID Controller Module needs to be discovered by using the Storage Management Initiative Specification (SMI-S) provider to communicate with the module.

- IBM BladeCenter Management Modules, AMMs, and Pass-Thru Modules

  Updates to MMs, AMMs, and Pass-Thru Modules requires a TFTP or FTP server. The IBM Systems Director can work as a TFTP server (need port 69 UDP inbound opened)

  Note: The SNMP agent must be enabled in the BladeCenter Management Module (MM, AMM). The SNMP access type needs to set to **SET**.

Update Manager does not perform the following tasks:

- Installing new software products.
- Installing Systems Director agents on systems that currently do not have an agent.

  Instead, install Systems Director agents with the Agent Manager plug-in of Systems Director.
- Migrating to any version of Systems Director from any version of IBM Director.
- Performing actions on systems that are not accessible.

  You can perform update actions on those systems that are accessible only. To check whether the system is accessible, go to Resource Manager and check the access column. If there is a green circle icon, the access state is `OK`. If there is another icon (red, yellow, or gray), check the access state and return the system to the `OK` state.
- Uninstalling updates and rolling back updates are not supported.

Check several system-specific considerations before you use the Update Manager:

http://pic.dhe.ibm.com/infocenter/director/pubs/index.jsp?topic=%2Fcom.ibm.director.updates.helps.doc%2Ffqm0_c_um_platform_extensions.html

4.3.3 Settings for Update Manager

Before you start to use Update Manager, configure all of the necessary settings:

1. On the Update Manager page, click **Configure settings** as shown in Figure 4-40.

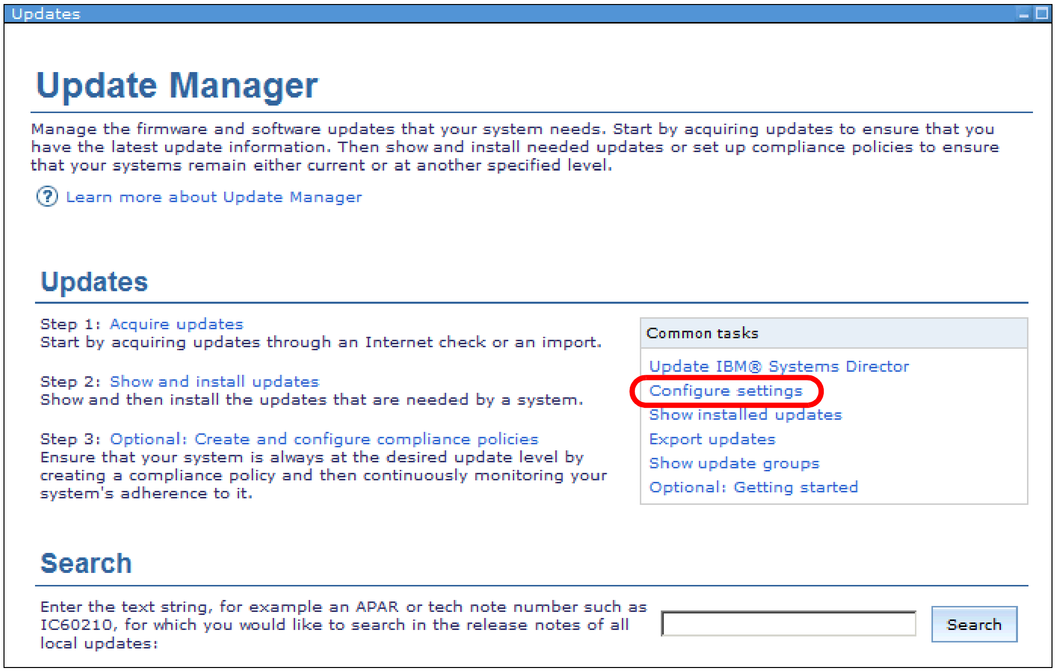

Updates

Update Manager

Manage the firmware and software updates that your system needs. Start by acquiring updates to ensure that you have the latest update information. Then show and install needed updates or set up compliance policies to ensure that your systems remain either current or at another specified level.

Learn more about Update Manager

Updates

Step 1: Acquire updates
Start by acquiring updates through an Internet check or an import.

Step 2: Show and install updates
Show and then install the updates that are needed by a system.

Step 3: Optional: Create and configure compliance policies
Ensure that your system is always at the desired update level by creating a compliance policy and then continuously monitoring your system's adherence to it.

Common tasks
- Update IBM® Systems Director
- Configure settings
- Show installed updates
- Export updates
- Show update groups
- Optional: Getting started

Search

Enter the text string, for example an APAR or tech note number such as IC60210, for which you would like to search in the release notes of all local updates:

Figure 4-40 Update Manager: Configure settings

A new window opens. Adjust these settings:

– Connection to the Internet (if the Systems Director accessed the Internet)
– Location for the local repository

  - Settings that are specific to System x and BladeCenter servers
  - Settings that are specific to AIX and VIOS systems

2. Select the **Connection** tab (Figure 4-41). You can configure a direct connection to the Internet or connect through an HTTP proxy server. After you make your selection, you can test the Internet connection by selecting **Test Internet Connection**.

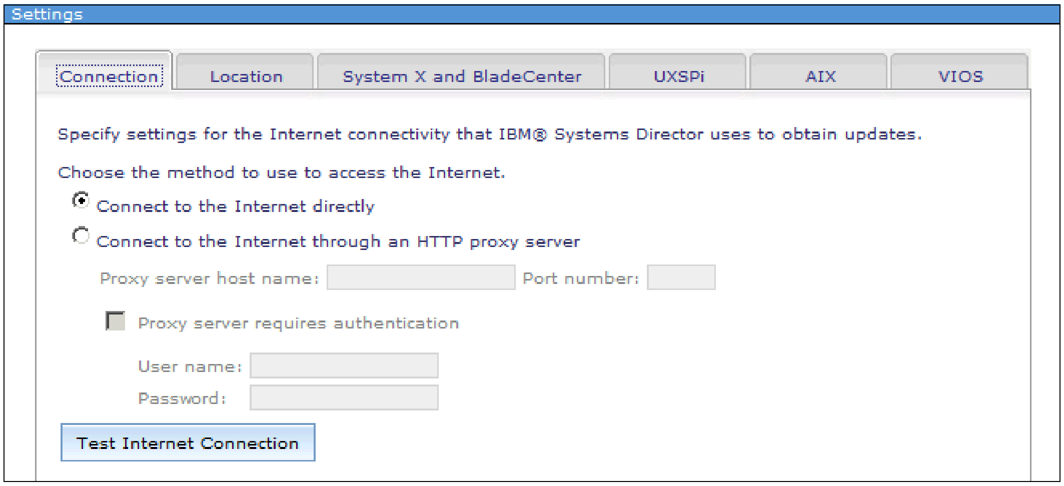

Specify settings for the Internet connectivity that IBM® Systems Director uses to obtain updates.

Choose the method to use to access the Internet.

● Connect to the Internet directly

○ Connect to the Internet through an HTTP proxy server

Proxy server host name: Port number:

□ Proxy server requires authentication

User name:

Password:

Figure 4-41 Connection tab

3. Select the **Location** tab (Figure 4-42). Define the size and location on disk of the local repository. The defaults are shown in Figure 4-42. The size might need to be increased. The size depends on the number of managed systems and the kinds of update packages that you want to deploy with the Systems Director server. The maximum size is 126 GB.

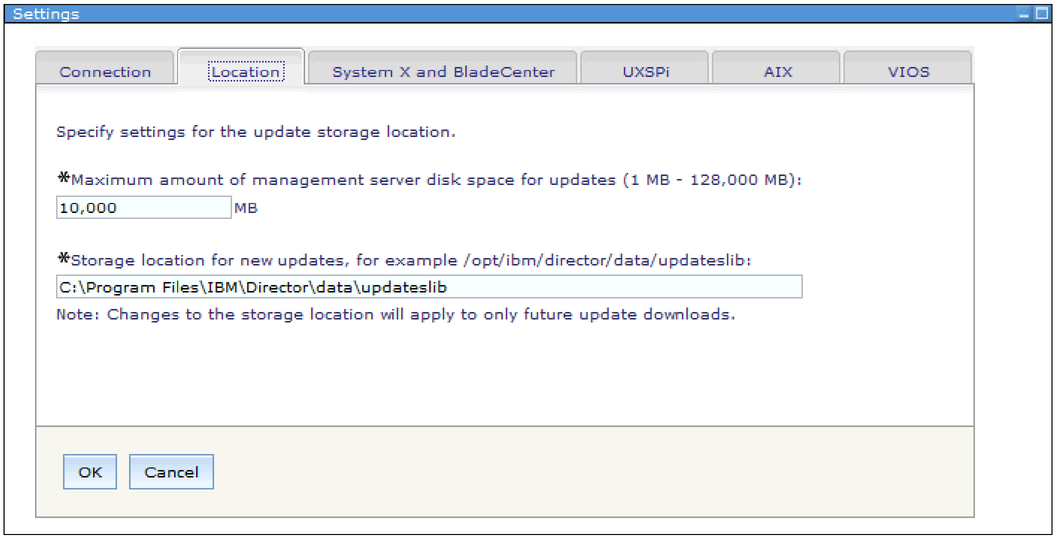

Specify settings for the update storage location.

*Maximum amount of management server disk space for updates (1 MB - 128,000 MB):
10,000 MB

*Storage location for new updates, for example /opt/ibm/director/data/updateslib:
C:\Program Files\IBM\Director\data\updateslib

Note: Changes to the storage location will apply to only future update downloads.

Figure 4-42 Location tab

4. In the System x and BladeCenter tab (Figure 4-43), you can define the use of a TFTP or FTP server for updates. This definition is necessary for updating the AMM and also for updating systems with IMMv1.

   For systems with IMMv2, this setting is not necessary. The service processor has enough internal memory to hold the update packages for updates to Unified Extensible Firmware Interface (UEFI), IMM, and preboot Dynamic System Analysis (pDSA).

   The Systems Director server can be used as a TFTP server, as indicated in Figure 4-43. Therefore, you do not need to install an external TFTP or FTP server.

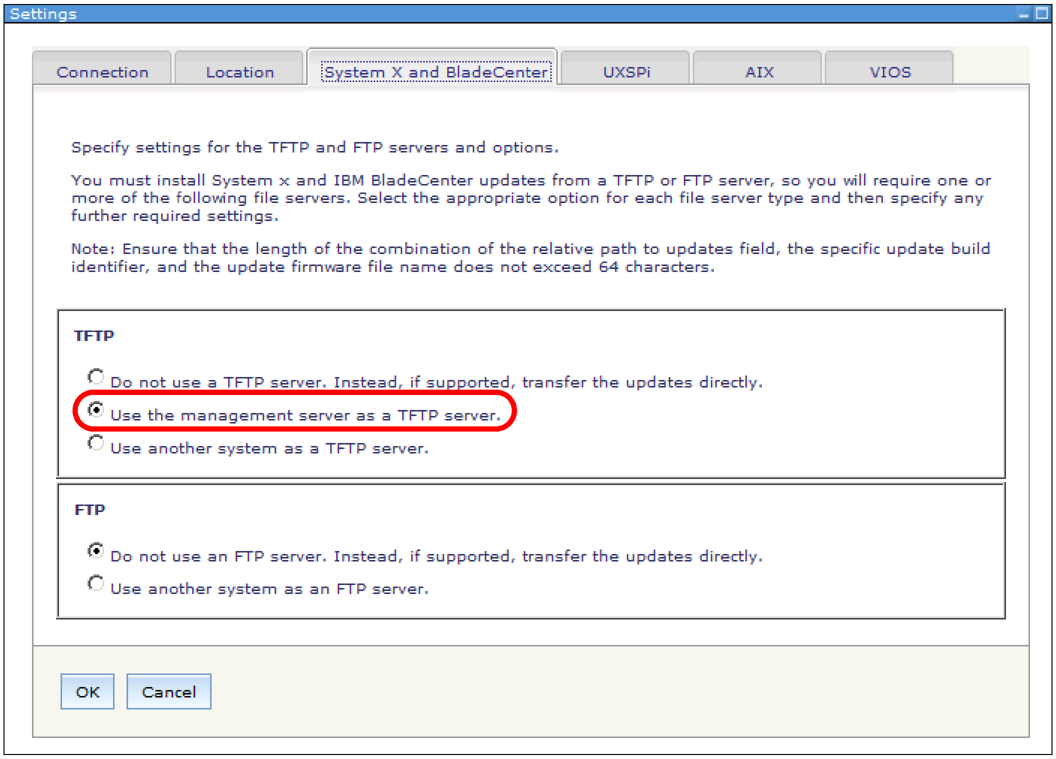

Specify settings for the TFTP and FTP servers and options.

You must install System x and IBM BladeCenter updates from a TFTP or FTP server, so you will require one or more of the following file servers. Select the appropriate option for each file server type and then specify any further required settings.

Note: Ensure that the length of the combination of the relative path to updates field, the specific update build identifier, and the update firmware file name does not exceed 64 characters.

TFTP

○ Do not use a TFTP server. Instead, if supported, transfer the updates directly.

● Use the management server as a TFTP server.

○ Use another system as a TFTP server.

FTP

● Do not use an FTP server. Instead, if supported, transfer the updates directly.

○ Use another system as an FTP server.

Figure 4-43 System x and BladeCenter: TFTP and FTP selection

If you wanted to use another TFTP or FTP server already in your infrastructure, you would have to discover, unlock, and inventory that operating system first before you can use it here.

5. The UXSPi tab, Figure 4-44, shows the installed UpdateXpress System Pack Installer (UXSPI) packages. Click **Import USXPi** to import these packages to Systems Director to deploy them.

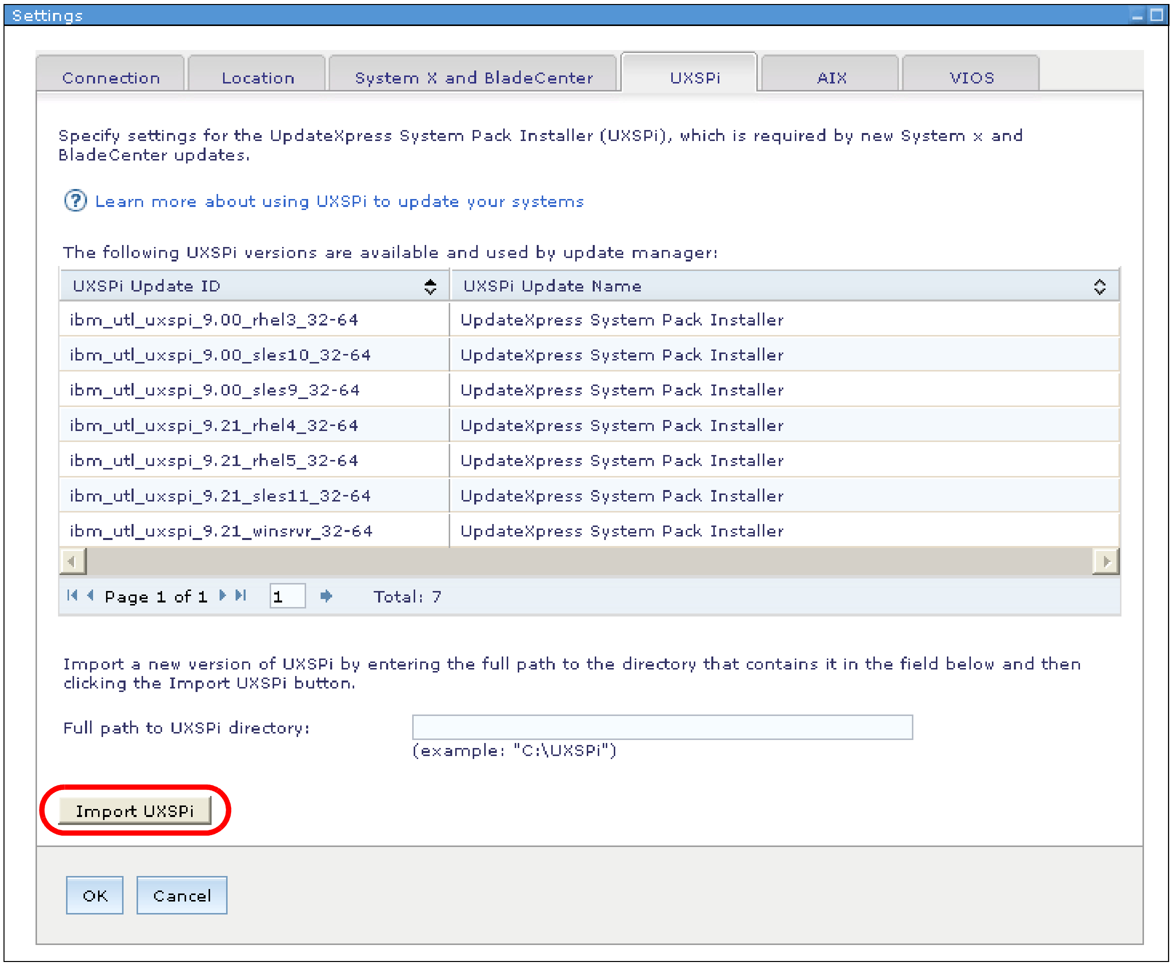

Specify settings for the UpdateXpress System Pack Installer (UXSPi), which is required by new System x and BladeCenter updates.

Learn more about using UXSPi to update your systems

The following UXSPi versions are available and used by update manager:

UXSPi Update ID	UXSPi Update Name
ibm_utl_uxspi_9.00_rhel3_32-64	UpdateXpress System Pack Installer
ibm_utl_uxspi_9.00_sles10_32-64	UpdateXpress System Pack Installer
ibm_utl_uxspi_9.00_sles9_32-64	UpdateXpress System Pack Installer
ibm_utl_uxspi_9.21_rhel4_32-64	UpdateXpress System Pack Installer
ibm_utl_uxspi_9.21_rhel5_32-64	UpdateXpress System Pack Installer
ibm_utl_uxspi_9.21_sles11_32-64	UpdateXpress System Pack Installer
ibm_utl_uxspi_9.21_winsrvr_32-64	UpdateXpress System Pack Installer

Page 1 of 1 1 Total: 7

Import a new version of UXSPi by entering the full path to the directory that contains it in the field below and then clicking the Import UXSPi button.

Full path to UXSPi directory:
(example: "C:\UXSPi")

Figure 4-44 UXSPi tab: Show version or import UXSPI

To import the UXSPI packages, you need the Subsystem Device Driver (SDD) file for each package. If the file is missing, you are prompted that the file is needed before you can continue.

6. The AIX tab (Figure 4-45) shows the selection for the AIX Network Installation Management (NIM) master. This NIM master is used for updates on AIX systems. Click **Browse** to select the AIX NIM master in your network. You must first discover, unlock, and inventory that system before you can use it here.

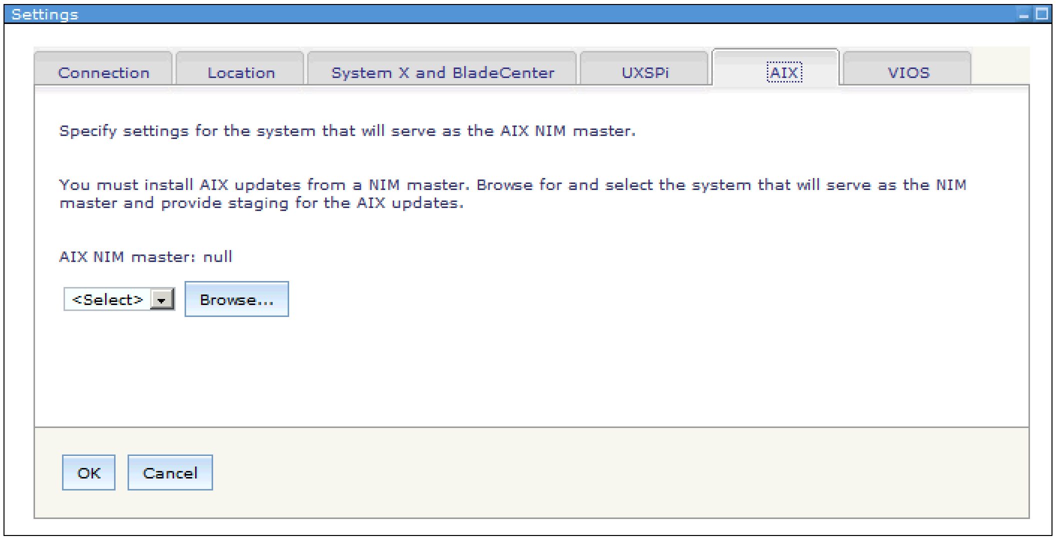

Specify settings for the system that will serve as the AIX NIM master.

You must install AIX updates from a NIM master. Browse for and select the system that will serve as the NIM master and provide staging for the AIX updates.

AIX NIM master: null

Figure 4-45 AIX: Define the AIX NIM master

7. The VIOS tab (Figure 4-46) shows the selection for the VIOS NIM master. This VIOS NIM master is used for a VIOS upgrade (migration). Before you can select a system here, you must have discovered, unlocked, and inventoried it.

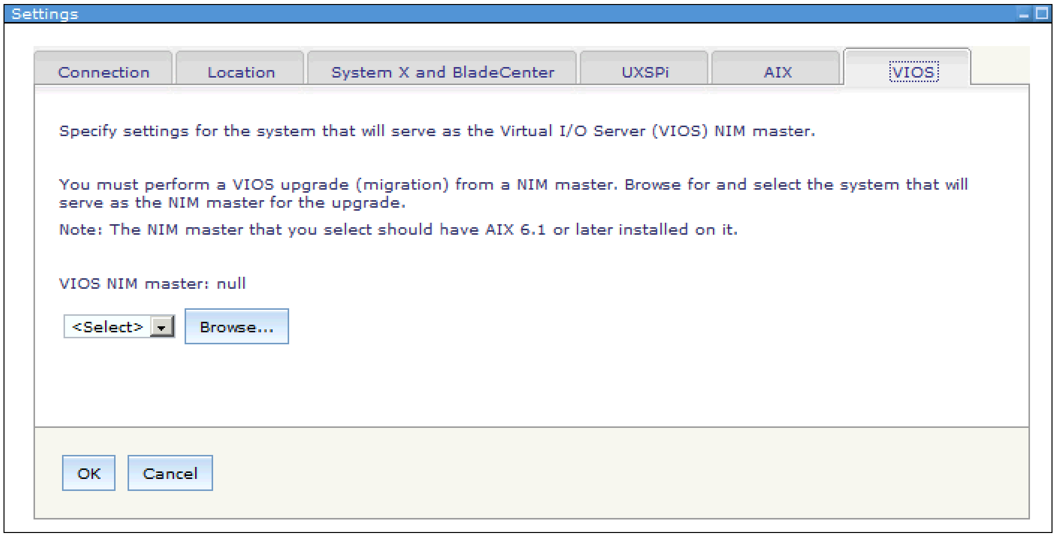

Specify settings for the system that will serve as the Virtual I/O Server (VIOS) NIM master.

You must perform a VIOS upgrade (migration) from a NIM master. Browse for and select the system that will serve as the NIM master for the upgrade.

Note: The NIM master that you select should have AIX 6.1 or later installed on it.

VIOS NIM master: null

Figure 4-46 VIOS: Define the VIOS NIM master

4.3.4 Update Manager with Internet connection

When Systems Director connects to the Internet, use Update Manager to check automatically and download the update information from a central IBM repository. Section 4.3.3, “Settings for Update Manager” on page 106 explains how to set up and verify the Internet connection.

Perform these steps to retrieve updates directly from the Internet:

1. On the Update Manager home page, click **Acquire updates**, as shown in Figure 4-47.

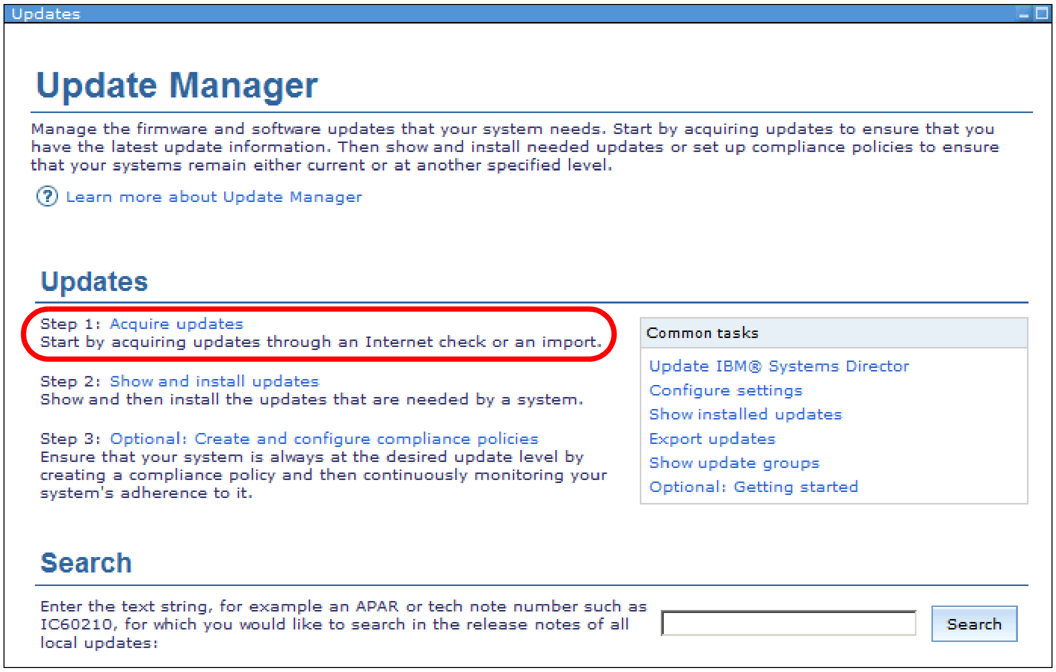

Updates

Update Manager

Manage the firmware and software updates that your system needs. Start by acquiring updates to ensure that you have the latest update information. Then show and install needed updates or set up compliance policies to ensure that your systems remain either current or at another specified level.

Learn more about Update Manager

Updates

Step 1: Acquire updates
Start by acquiring updates through an Internet check or an import.

Step 2: Show and install updates
Show and then install the updates that are needed by a system.

Step 3: Optional: Create and configure compliance policies
Ensure that your system is always at the desired update level by creating a compliance policy and then continuously monitoring your system's adherence to it.

Common tasks

- Update IBM® Systems Director
- Configure settings
- Show installed updates
- Export updates
- Show update groups
- Optional: Getting started

Search

Enter the text string, for example an APAR or tech note number such as IC60210, for which you would like to search in the release notes of all local updates:

Figure 4-47 Update Manager: Acquire updates

2. In the Acquire Updates window (Figure 4-48), click **Check for updates (Internet connection required)**.

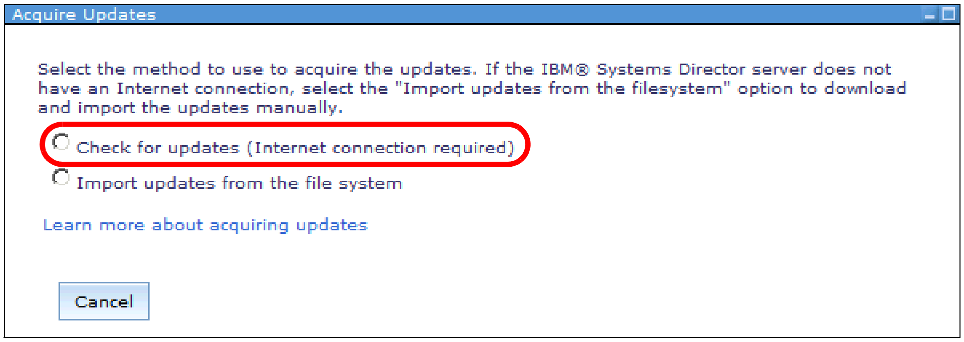

Acquire Updates

Select the method to use to acquire the updates. If the IBM® Systems Director server does not have an Internet connection, select the "Import updates from the filesystem" option to download and import the updates manually.

○ Check for updates (Internet connection required)

○ Import updates from the file system

Learn more about acquiring updates

Figure 4-48 Selection for Internet or import

3. The window then expands as shown in Figure 4-49. A selection window opens where you select the update that the Update Manager looks for in the IBM repository.
4. Expand **Available update types** on the left. Select the update types that you want and click **Add**.

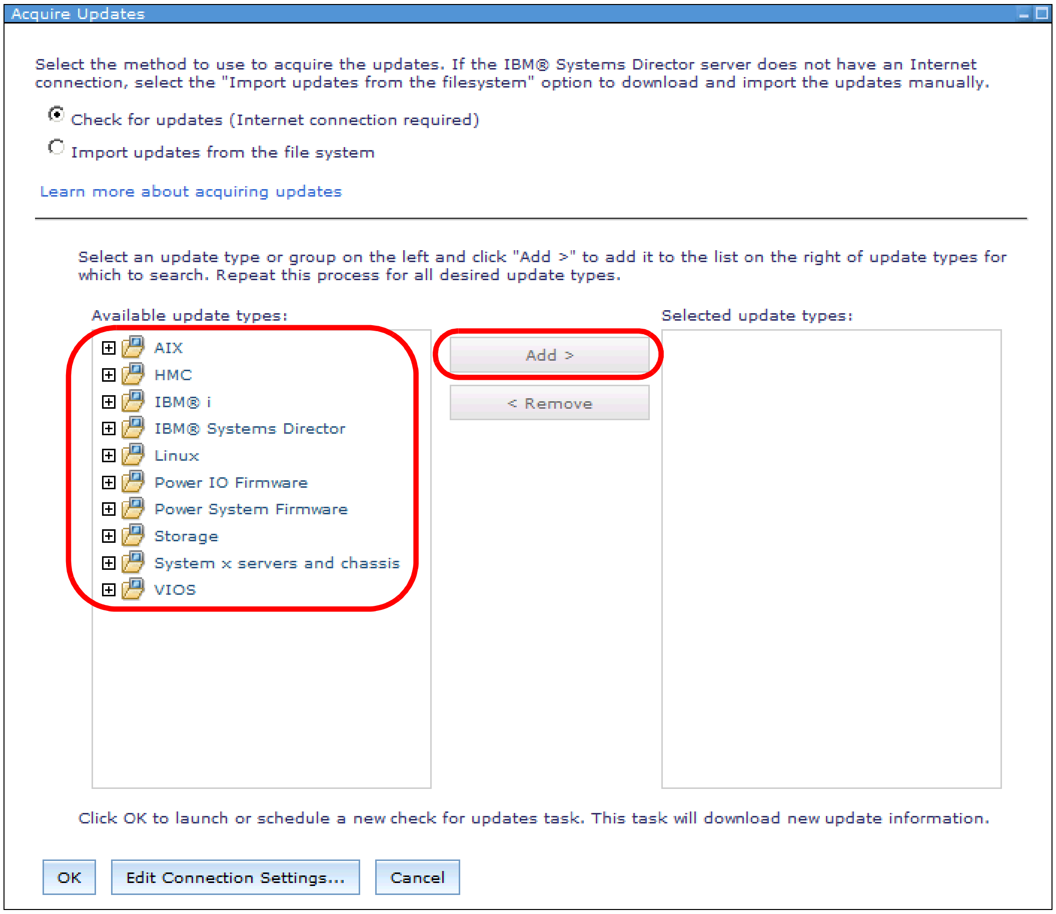

Acquire Updates

Select the method to use to acquire the updates. If the IBM® Systems Director server does not have an Internet connection, select the "Import updates from the filesystem" option to download and import the updates manually.

● Check for updates (Internet connection required)

○ Import updates from the file system

Learn more about acquiring updates

Select an update type or group on the left and click "Add >" to add it to the list on the right of update types for which to search. Repeat this process for all desired update types.

Click OK to launch or schedule a new check for updates task. This task will download new update information.

Figure 4-49 Select the updates for checking

5. After you select the updates that you want, click **OK**.
6. A scheduler window opens. Download the updates one time or define a recurring schedule. The best practice is that you perform the updates on a recurring schedule, such as once a week. The best time to update is off-hours (for example, weekend nights) so that this download traffic does not affect your daily business.

4.3.5 Update Manager with no Internet connection

If your Systems Director server does not connect to the Internet, you can use the Update Manager to import update packages. Download these update packages in advance from the IBM Fix Central website or another source for IBM updates.

IBM Fix Central is at the following link:

`http://ibm.com/support/fixcentral`

You can obtain single updates (latest updates) or you can also use the UXSPI packages for your system. UXSPI packages contain updates for your system that are tested and work together. The types of updates include updates for UEFI, IMM, drivers, or firmware.

UXSPI packages are easier to download. You do not need to locate a download for each update package for each component in your system separately.

Follow these steps to apply the updates that you previously downloaded to Update Manager:

1. On the Update Manager startup page, select **Acquire updates**, as shown in Figure 4-50.

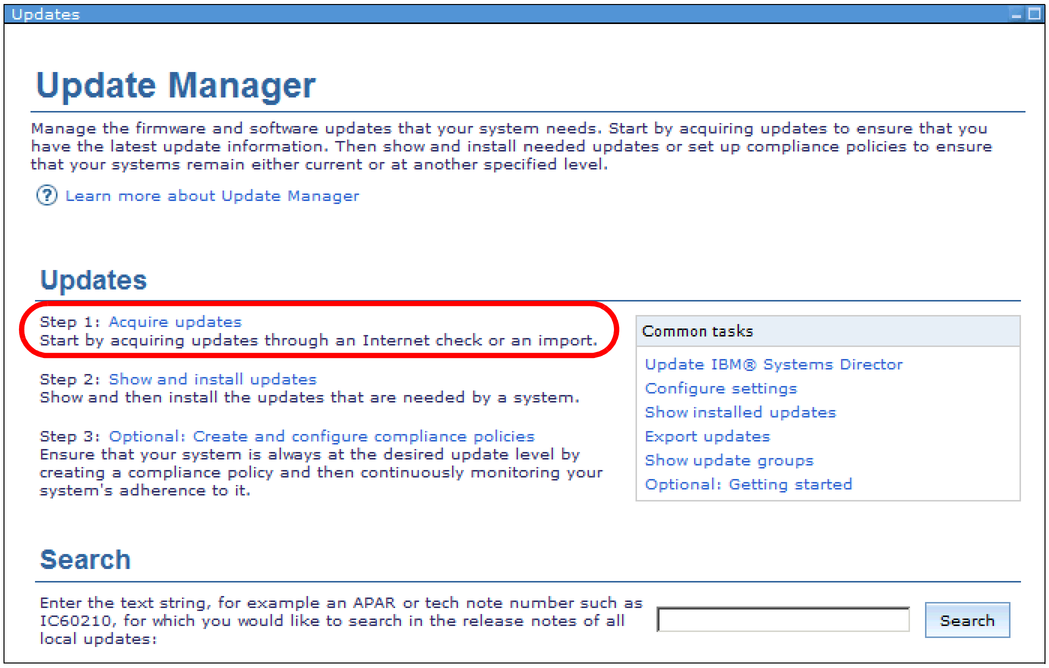

Updates

Update Manager

Manage the firmware and software updates that your system needs. Start by acquiring updates to ensure that you have the latest update information. Then show and install needed updates or set up compliance policies to ensure that your systems remain either current or at another specified level.

Learn more about Update Manager

Updates

Step 1: Acquire updates
Start by acquiring updates through an Internet check or an import.

Step 2: Show and install updates
Show and then install the updates that are needed by a system.

Step 3: Optional: Create and configure compliance policies
Ensure that your system is always at the desired update level by creating a compliance policy and then continuously monitoring your system's adherence to it.

Common tasks

- Update IBM® Systems Director
- Configure settings
- Show installed updates
- Export updates
- Show update groups
- Optional: Getting started

Search

Enter the text string, for example an APAR or tech note number such as IC60210, for which you would like to search in the release notes of all local updates:

Figure 4-50 Acquire updates

2. Click **Import updates from the file system**, as shown in Figure 4-51.

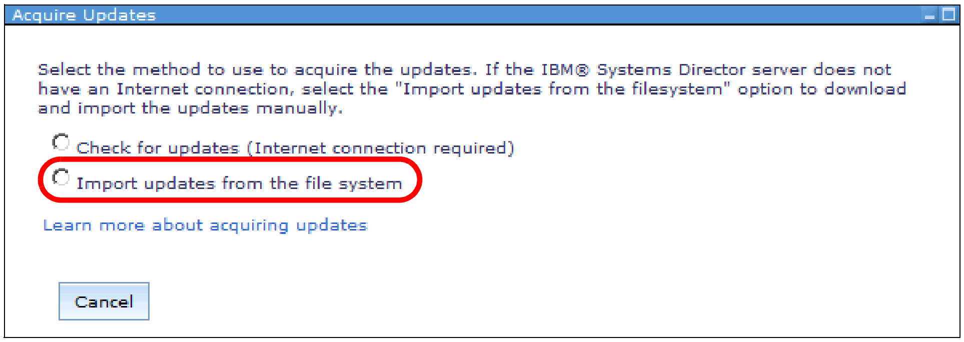

Acquire Updates

Select the method to use to acquire the updates. If the IBM® Systems Director server does not have an Internet connection, select the "Import updates from the filesystem" option to download and import the updates manually.

○ Check for updates (Internet connection required)

○ Import updates from the file system

Learn more about acquiring updates

Figure 4-51 Acquire Updates selection window

3. When you select to import the updates, the window expands (Figure 4-52). Select the directory where you downloaded the updates previously. This directory must be on or accessible from the management server.

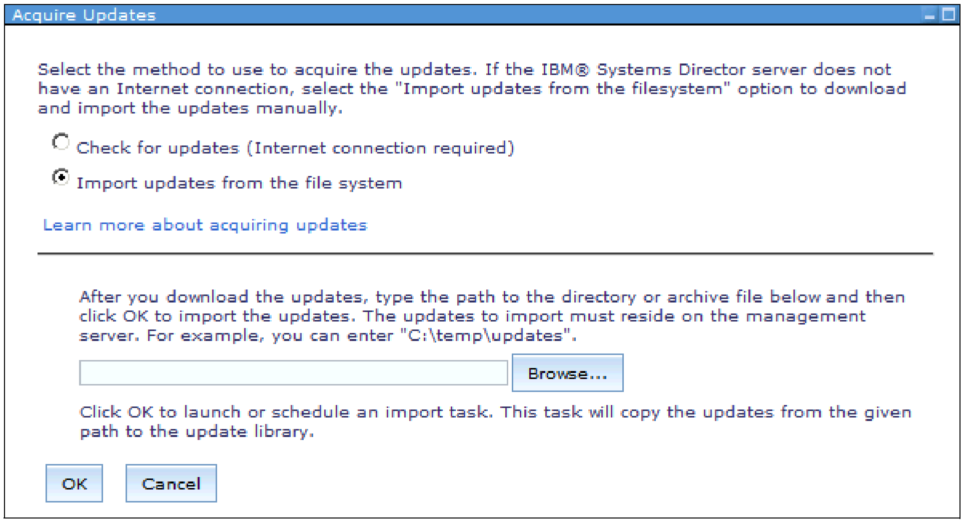

Select the method to use to acquire the updates. If the IBM® Systems Director server does not have an Internet connection, select the "Import updates from the filesystem" option to download and import the updates manually.

○ Check for updates (Internet connection required)

● Import updates from the file system

After you download the updates, type the path to the directory or archive file below and then click OK to import the updates. The updates to import must reside on the management server. For example, you can enter "C:\temp\updates".

Click OK to launch or schedule an import task. This task will copy the updates from the given path to the update library.

Figure 4-52 Import updates from the file system

4. After you select the directory, click **OK**. The updates are imported into the Systems Director server.

4.3.6 Compliance check

A *compliance check* compares information in the local repository of Systems Director with the inventory information that Systems Director collected from the managed systems. The compliance check process is a background process. After you define it, it runs automatically for each new update or for each new collection of system inventory information.

If the compliance check identifies new updates for a system, the system is marked as noncompliant by a status of *information*, *warning*, or *critical*. This status depends on the level of the system and the level of the available updates. This status can change if newer updates appear.

For the best results, schedule a regular download or import of the newest available updates and perform regular inventories for the available systems. You can automate this inventory collection for your systems as described in 4.2, “Inventory” on page 86.

You can set up a compliance check against a single system or a group of systems. If you use groups, define groups that contain systems of the same type or that share properties.

Follow these steps to set up the compliance check:

1. From the Update Manager main page, click **Optional: Create and configure compliance policies**, as shown in Figure 4-53.

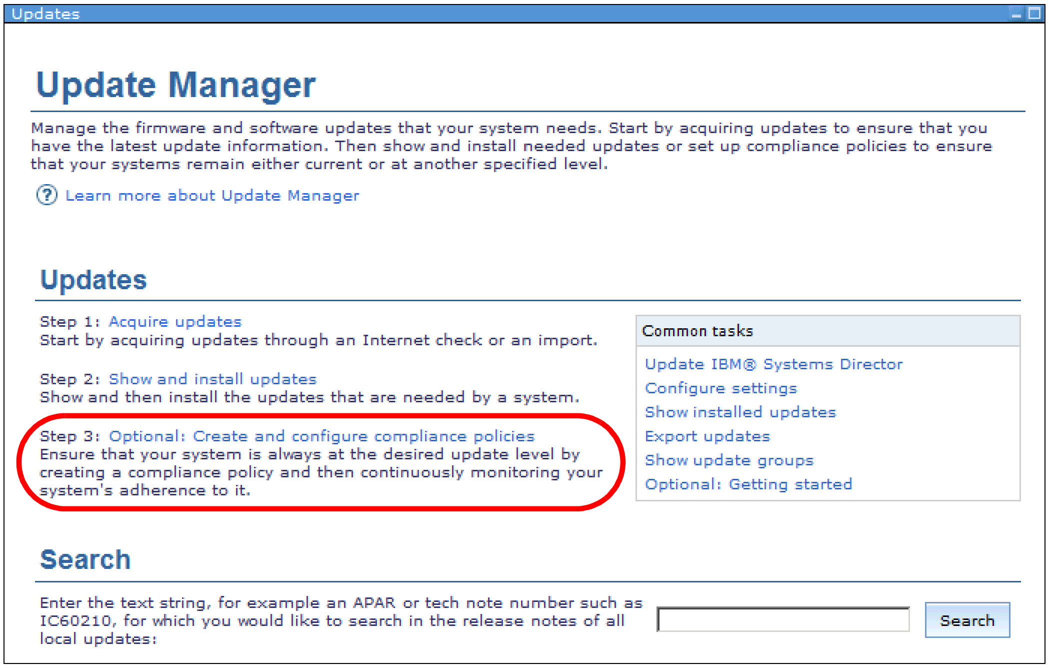

Updates

Update Manager

Manage the firmware and software updates that your system needs. Start by acquiring updates to ensure that you have the latest update information. Then show and install needed updates or set up compliance policies to ensure that your systems remain either current or at another specified level.

Learn more about Update Manager

Updates

Step 1: Acquire updates
Start by acquiring updates through an Internet check or an import.

Step 2: Show and install updates
Show and then install the updates that are needed by a system.

Step 3: Optional: Create and configure compliance policies
Ensure that your system is always at the desired update level by creating a compliance policy and then continuously monitoring your system's adherence to it.

Common tasks

- Update IBM® Systems Director
- Configure settings
- Show installed updates
- Export updates
- Show update groups
- Optional: Getting started

Search

Enter the text string, for example an APAR or tech note number such as IC60210, for which you would like to search in the release notes of all local updates:

Figure 4-53   Select, create, and configure compliance policies

2. In Figure 4-54, select systems or a group for which you want to create the compliance check. After you select these systems or a group, click **OK**.

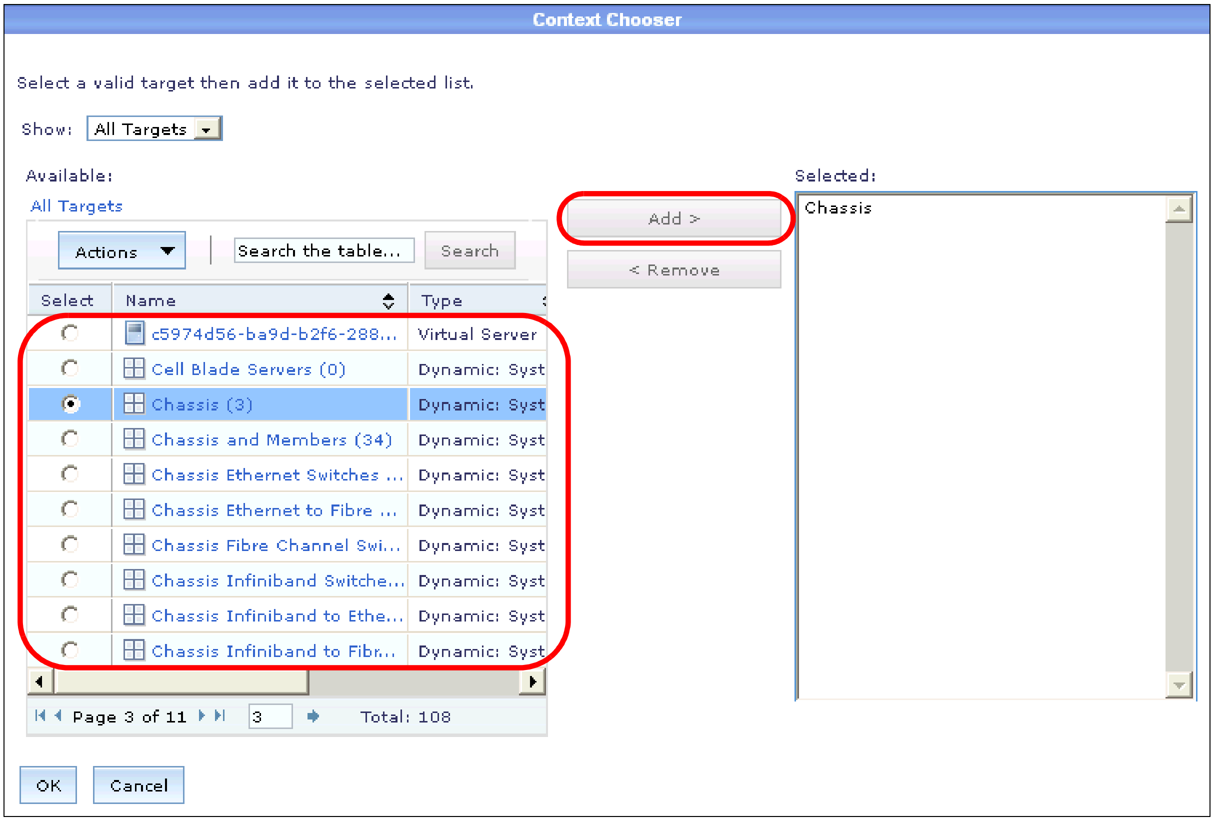

Select a valid target then add it to the selected list.

Show: All Targets

Available:

All Targets

Select	Name	Type
○	c5974d56-ba9d-b2f6-288...	Virtual Server
○	Cell Blade Servers (0)	Dynamic: Syst
●	Chassis (3)	Dynamic: Syst
○	Chassis and Members (34)	Dynamic: Syst
○	Chassis Ethernet Switches ...	Dynamic: Syst
○	Chassis Ethernet to Fibre ...	Dynamic: Syst
○	Chassis Fibre Channel Swi...	Dynamic: Syst
○	Chassis Infiniband Switche...	Dynamic: Syst
○	Chassis Infiniband to Ethe...	Dynamic: Syst
○	Chassis Infiniband to Fibr...	Dynamic: Syst

Page 3 of 11 3 Total: 108

Selected:

Chassis

Figure 4-54 Select systems or a group for the compliance check

If the group or systems that you select have no available inventory information, you see a message (Figure 4-55). The message states that no inventory is available and you need to perform an inventory collection by clicking **Collect Inventory**. After the inventory starts, click **Close Message**.

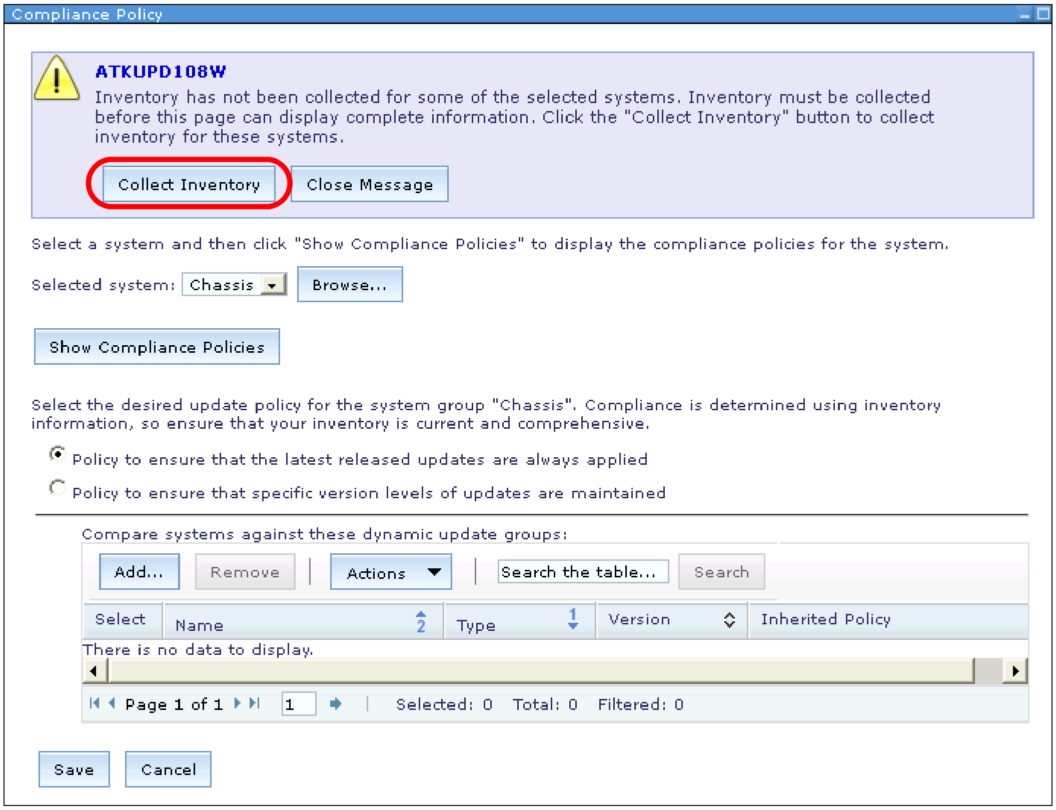

Compliance Policy

ATKUPD108W

Inventory has not been collected for some of the selected systems. Inventory must be collected before this page can display complete information. Click the "Collect Inventory" button to collect inventory for these systems.

Select a system and then click "Show Compliance Policies" to display the compliance policies for the system.

Selected system: Chassis

Select the desired update policy for the system group "Chassis". Compliance is determined using inventory information, so ensure that your inventory is current and comprehensive.

● Policy to ensure that the latest released updates are always applied

○ Policy to ensure that specific version levels of updates are maintained

Compare systems against these dynamic update groups:

Select	Name	Type	Version	Inherited Policy

There is no data to display.

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

Figure 4-55 Warning message that no inventory is available

3. After the inventory run completes, a message appears and you can add a compliance policy for the systems or group that you selected before. Click **Add** (Figure 4-56).

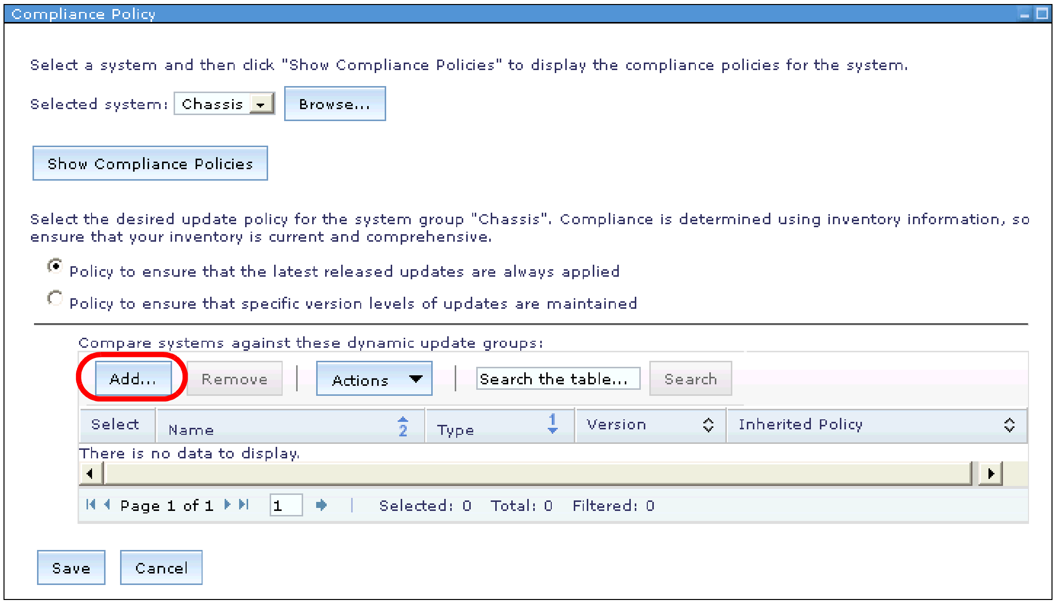

Compliance Policy

Select a system and then click "Show Compliance Policies" to display the compliance policies for the system.

Selected system: Chassis Browse...

Show Compliance Policies

Select the desired update policy for the system group "Chassis". Compliance is determined using inventory information, so ensure that your inventory is current and comprehensive.

● Policy to ensure that the latest released updates are always applied

○ Policy to ensure that specific version levels of updates are maintained

Compare systems against these dynamic update groups:

Select	Name	Type	Version	Inherited Policy

There is no data to display.

Selected: 0 Total: 0 Filtered: 0

Figure 4-56 After you collect the inventory, select Add

4. In Figure 4-57, select the update group and click **Add**. This group defines the type of updates for which your system runs the compliance check. In our example, we selected *All Critical IBM System x and BladeCenter Updates and All IBM Systems Director 6.3 Updates*. After you finish your selections, click **OK**.

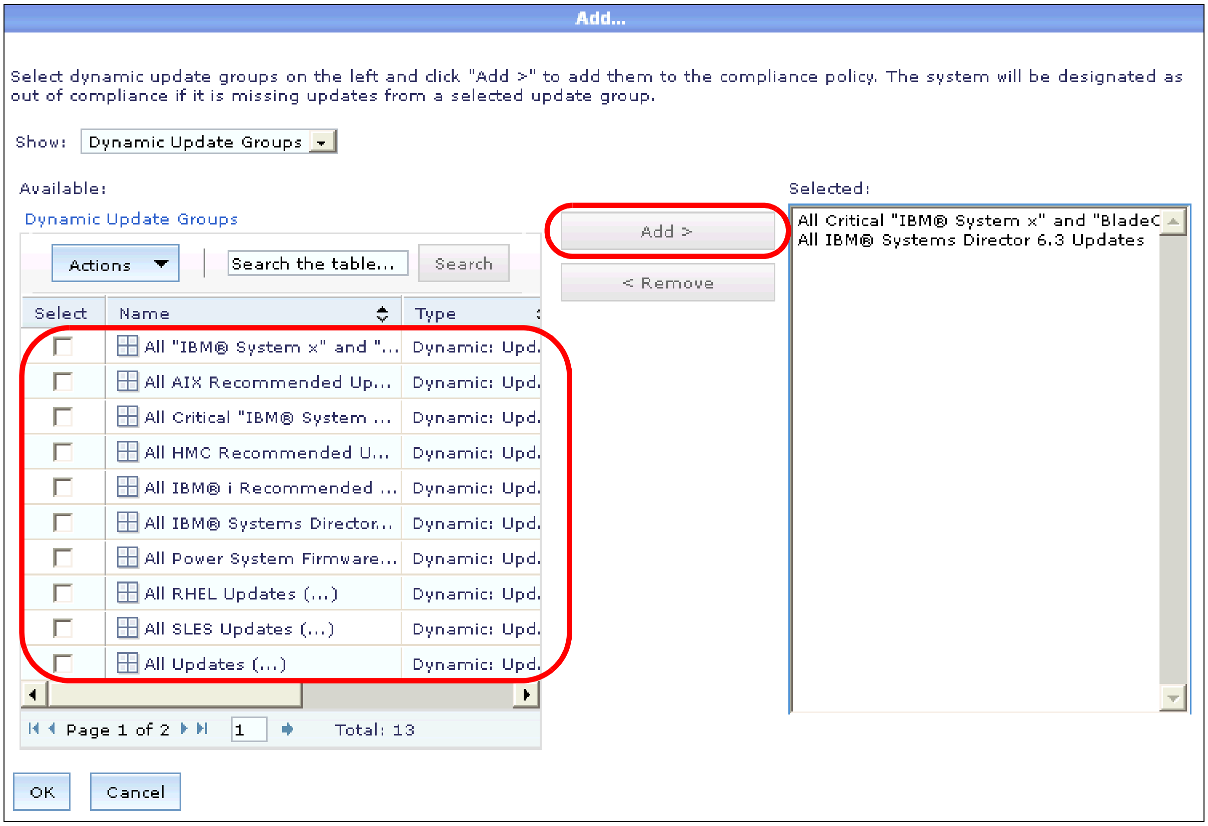

Add...

Select dynamic update groups on the left and click "Add >" to add them to the compliance policy. The system will be designated as out of compliance if it is missing updates from a selected update group.

Show: Dynamic Update Groups

Available:

Dynamic Update Groups

Select	Name	Type
☐	All "IBM® System x" and "...	Dynamic: Upd.
☐	All AIX Recommended Up...	Dynamic: Upd.
☐	All Critical "IBM® System ...	Dynamic: Upd.
☐	All HMC Recommended U...	Dynamic: Upd.
☐	All IBM® i Recommended ...	Dynamic: Upd.
☐	All IBM® Systems Director...	Dynamic: Upd.
☐	All Power System Firmware...	Dynamic: Upd.
☐	All RHEL Updates (...)	Dynamic: Upd.
☐	All SLES Updates (...)	Dynamic: Upd.
☐	All Updates (...)	Dynamic: Upd.

Page 1 of 2 1 Total: 13

Selected:

All Critical "IBM® System x" and "BladeC
All IBM® Systems Director 6.3 Updates

Figure 4-57 Select the update groups for the compliance check

5. You return to the previous window. The compliance policies that are defined for your systems are shown as seen in Figure 4-58. To finish the definition of the compliance policies, click **Save**.

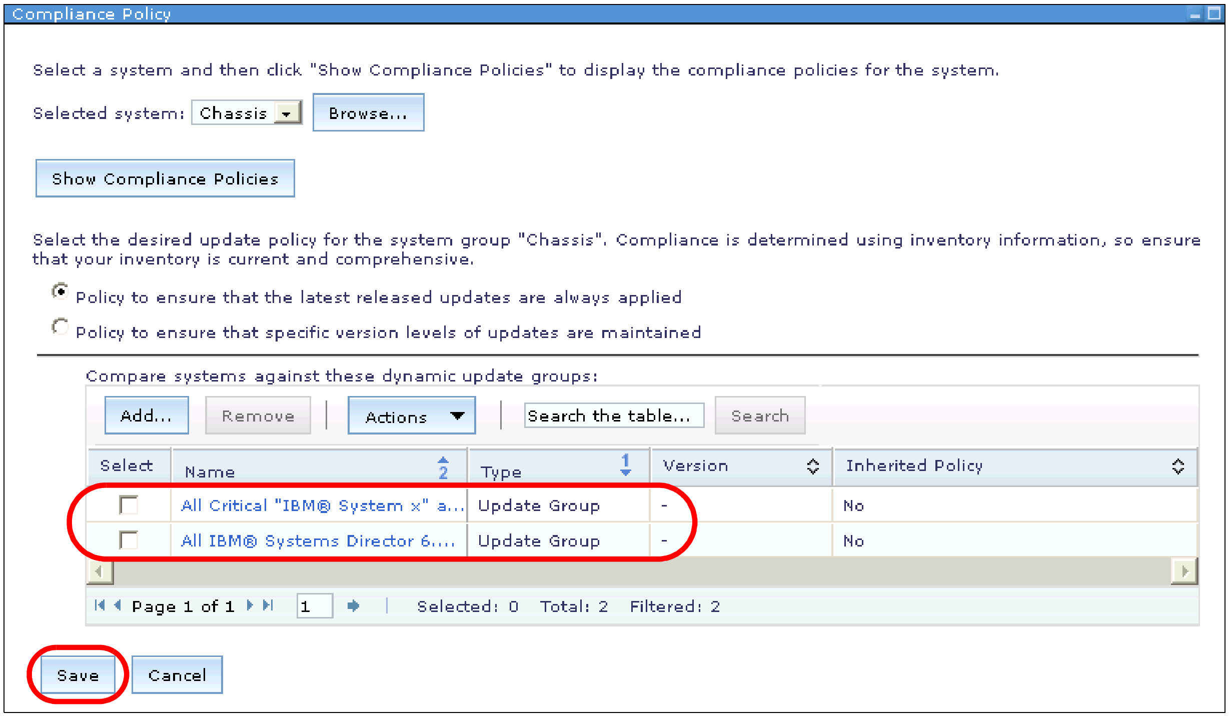

Compliance Policy

Select a system and then click "Show Compliance Policies" to display the compliance policies for the system.

Selected system: Chassis

Select the desired update policy for the system group "Chassis". Compliance is determined using inventory information, so ensure that your inventory is current and comprehensive.

● Policy to ensure that the latest released updates are always applied

○ Policy to ensure that specific version levels of updates are maintained

Compare systems against these dynamic update groups:

Select	Name	Type	Version	Inherited Policy
□	All Critical "IBM® System x" a...	Update Group	-	No
□	All IBM® Systems Director 6....	Update Group	-	No

Page 1 of 1   1   Selected: 0   Total: 2   Filtered: 2

Figure 4-58 Selected update groups for the systems

6. You return to the Update Manager home page where you see the Update Compliance section (Figure 4-59). You can see the compliance status of the systems that you specified.

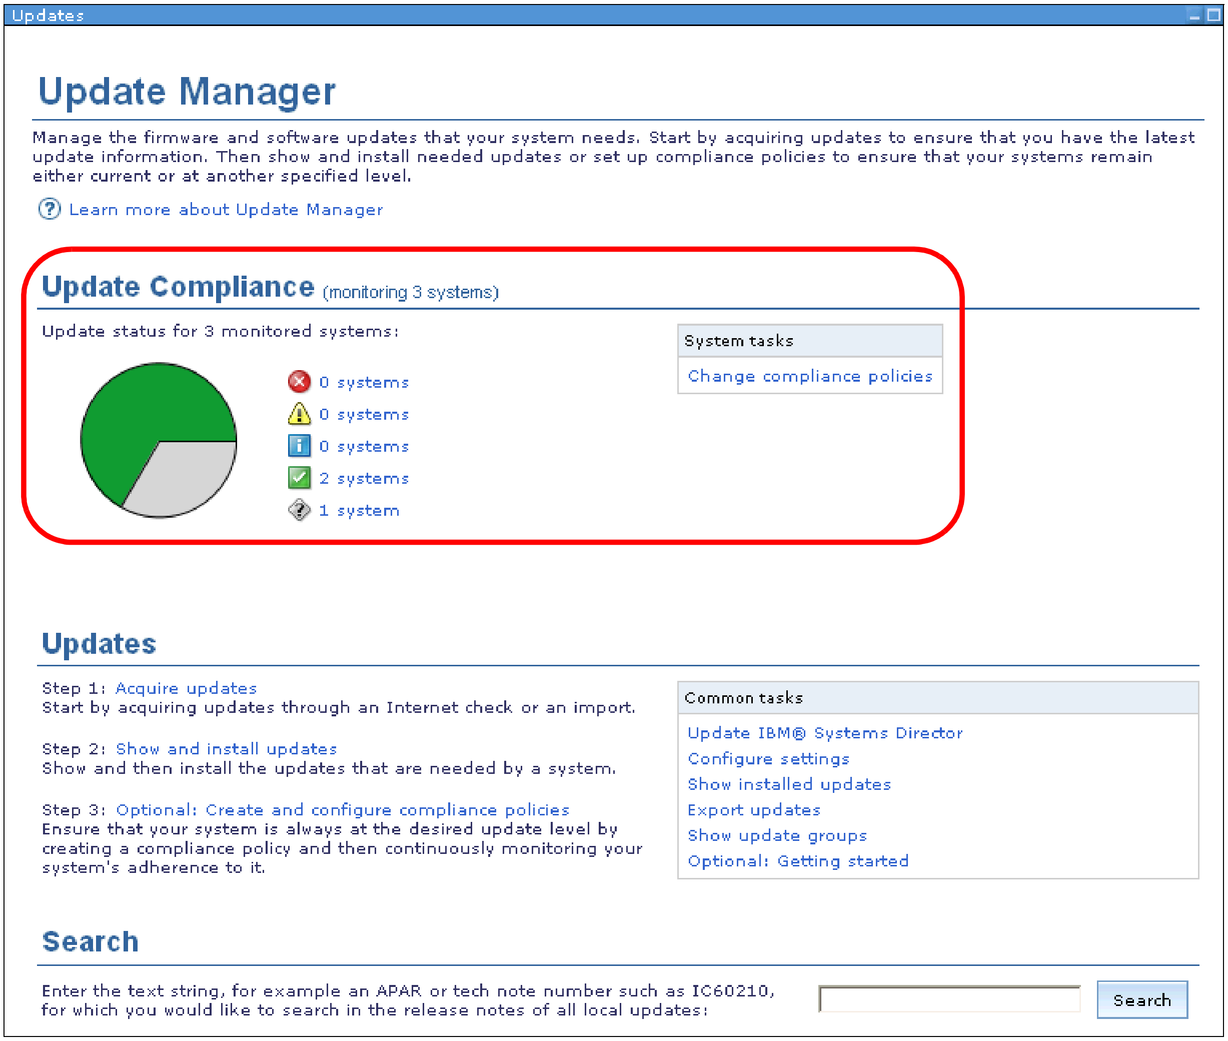

Updates

Update Manager

Manage the firmware and software updates that your system needs. Start by acquiring updates to ensure that you have the latest update information. Then show and install needed updates or set up compliance policies to ensure that your systems remain either current or at another specified level.

? Learn more about Update Manager

Update Compliance (monitoring 3 systems)

Update status for 3 monitored systems:

✖ 0 systems
⚠ 0 systems
ℹ 0 systems
✓ 2 systems
◆ 1 system

System tasks
Change compliance policies

Updates

Step 1: Acquire updates
Start by acquiring updates through an Internet check or an import.

Step 2: Show and install updates
Show and then install the updates that are needed by a system.

Step 3: Optional: Create and configure compliance policies
Ensure that your system is always at the desired update level by creating a compliance policy and then continuously monitoring your system's adherence to it.

Common tasks
Update IBM® Systems Director
Configure settings
Show installed updates
Export updates
Show update groups
Optional: Getting started

Search

Enter the text string, for example an APAR or tech note number such as IC60210, for which you would like to search in the release notes of all local updates:

Figure 4-59 Update Manager with defined compliance check

4.3.7 Update process

The Update Manager is configured and you have the required updates (either from the Internet or imported from a local directory to the Systems Director repository). Start the update process.

We describe two methods:

- “Using Update Manager: Show and install updates” on page 122
- “Using the compliance check to update” on page 129

Using Update Manager: Show and install updates

You can check whether updates are available to install for systems or a group of systems:

1. From Update Manager, click **Show and install updates**, as shown in Figure 4-60.

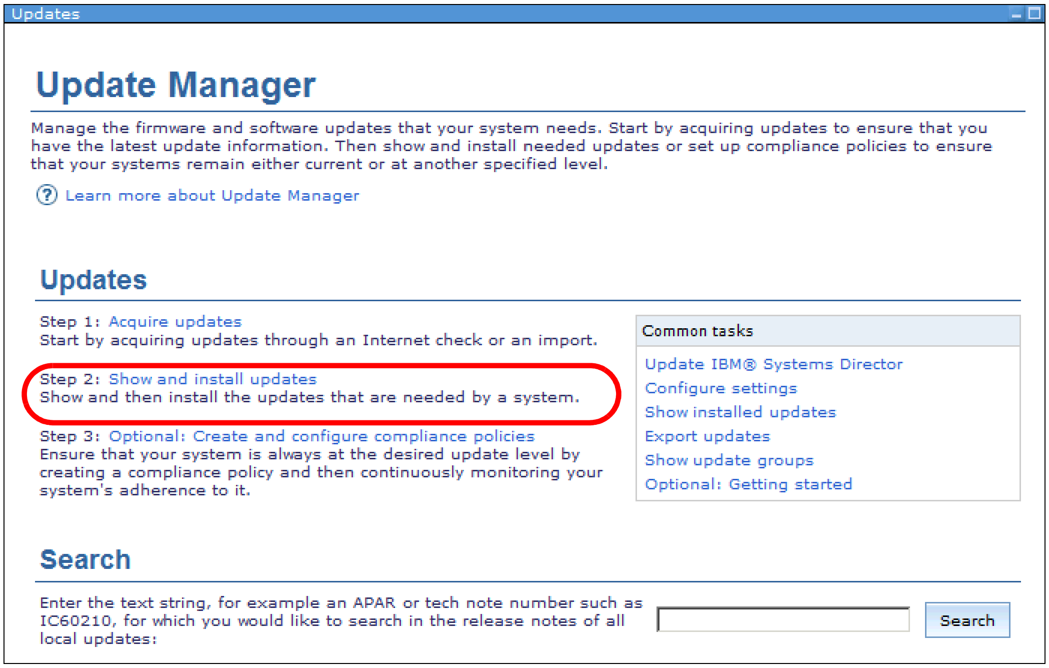

Update Manager

Manage the firmware and software updates that your system needs. Start by acquiring updates to ensure that you have the latest update information. Then show and install needed updates or set up compliance policies to ensure that your systems remain either current or at another specified level.

Learn more about Update Manager

Updates

Step 1: Acquire updates
Start by acquiring updates through an Internet check or an import.

Step 2: Show and install updates
Show and then install the updates that are needed by a system.

Step 3: Optional: Create and configure compliance policies
Ensure that your system is always at the desired update level by creating a compliance policy and then continuously monitoring your system's adherence to it.

Search

Enter the text string, for example an APAR or tech note number such as IC60210, for which you would like to search in the release notes of all local updates:

Figure 4-60 Update Manager: Show and install updates

2. Select the system or the group of systems for which you want to update. In our example, we select the *Chassis* group as shown in Figure 4-61. Click **Show and Install Updates**.

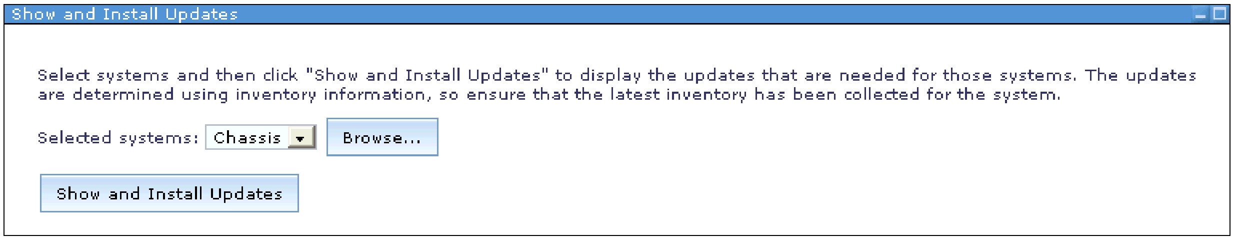

Select systems and then click "Show and Install Updates" to display the updates that are needed for those systems. The updates are determined using inventory information, so ensure that the latest inventory has been collected for the system.

Selected systems: Chassis

Figure 4-61 Show and Install Updates: System selection

3. The window expands to show the available updates for the systems that you selected. A message appears if no inventory is available for these systems. You can select to start the inventory collection.

   In our example, the inventory was run before and one update is available for the group, *Chassis*, which affects two systems (Figure 4-62).

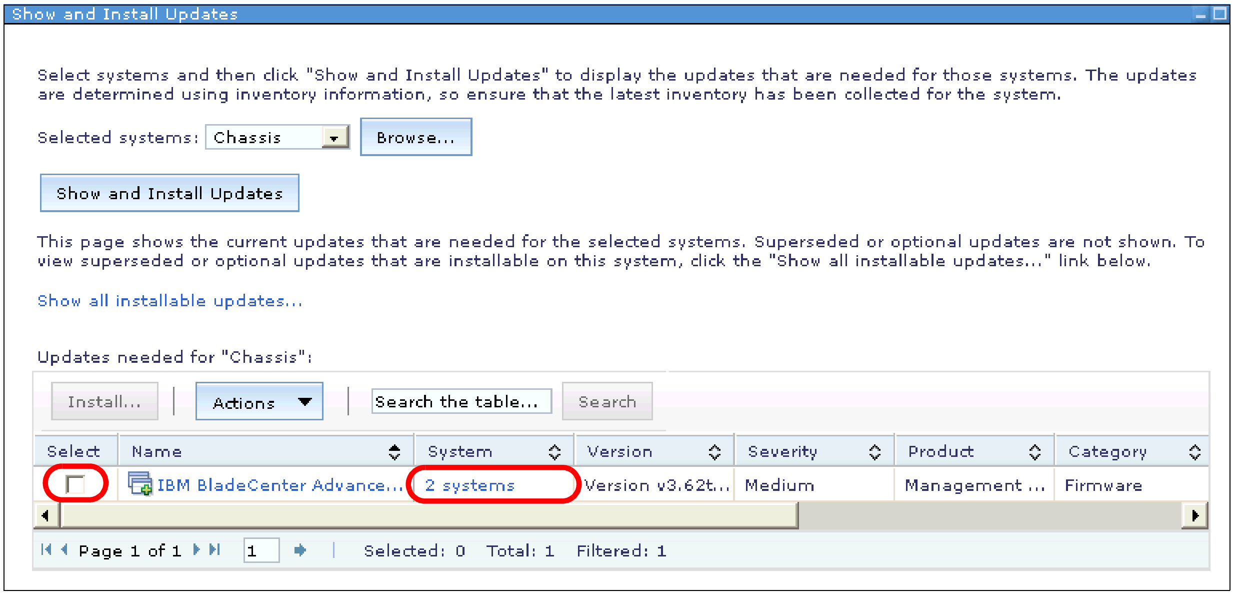

Select systems and then click "Show and Install Updates" to display the updates that are needed for those systems. The updates are determined using inventory information, so ensure that the latest inventory has been collected for the system.

Selected systems: Chassis

This page shows the current updates that are needed for the selected systems. Superseded or optional updates are not shown. To view superseded or optional updates that are installable on this system, click the "Show all installable updates..." link below.

Updates needed for "Chassis":

Select	Name	System	Version	Severity	Product	Category
□	IBM BladeCenter Advance...	2 systems	Version v3.62t...	Medium	Management ...	Firmware

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

Figure 4-62 Show and Install Updates window: Available update for selected systems

4. Multiple methods are available to install this update:

   – Click the check box next to the update package and click **Install**. This method is easiest and, in our example, updates two chassis.

   – But, because the update is available for two systems, you might want to know which system you are updating. Perhaps, you cannot update all systems at the same time for business reasons. You can click **2 systems** in the System column, as shown in Figure 4-62.

A new window opens where the two systems are listed, as shown in Figure 4-63.

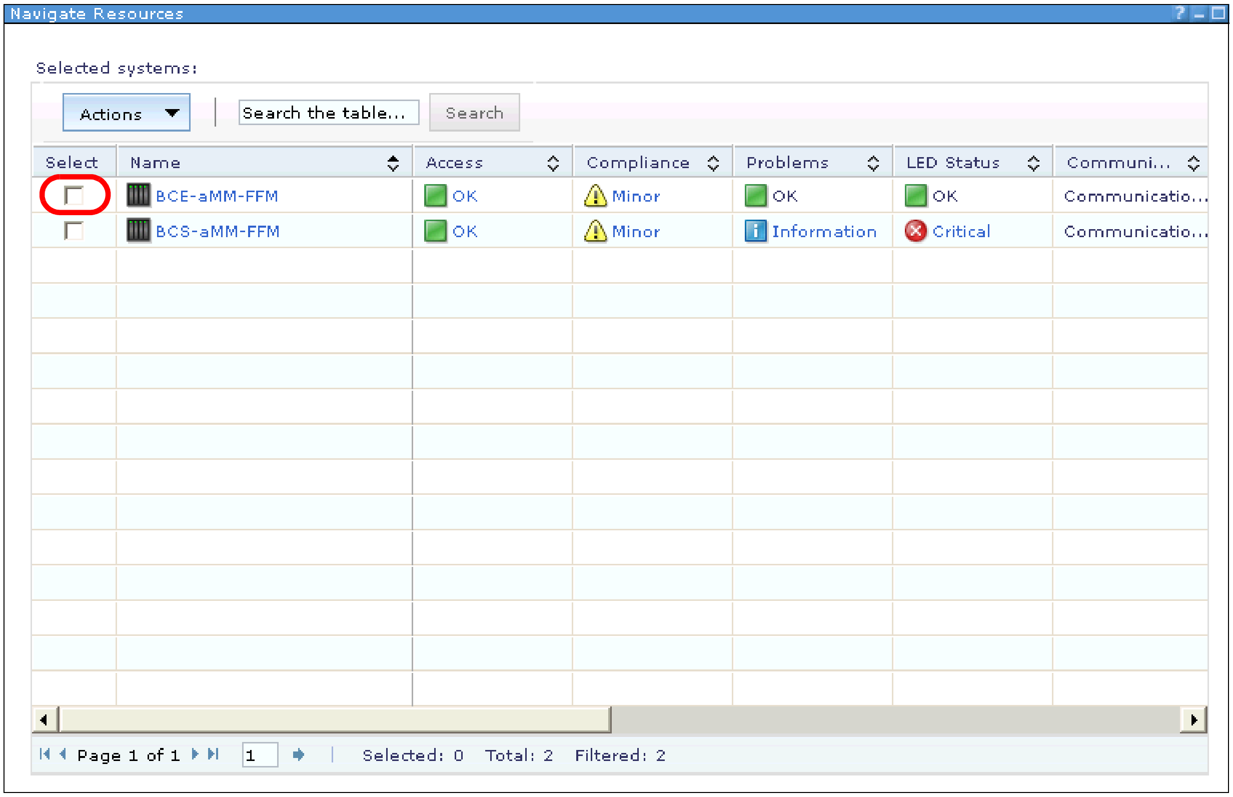

Navigate Resources

Selected systems:

Select	Name	Access	Compliance	Problems	LED Status	Communi...
□	BCE-aMM-FFM	🟩 OK	⚠ Minor	🟩 OK	🟩 OK	Communicatio...
□	BCS-aMM-FFM	🟩 OK	⚠ Minor	ℹ Information	✖ Critical	Communicatio...

Page 1 of 1 | 1 | Selected: 0 Total: 2 Filtered: 2

Figure 4-63 Updates for systems

Check the individual system or systems that you want to update and click **Actions** → **Release Management** → **Show and install updates**. This action returns you to the window that is shown in Figure 4-64. Click the check box next to the update and then click **Install** to start the upgrade installation process.

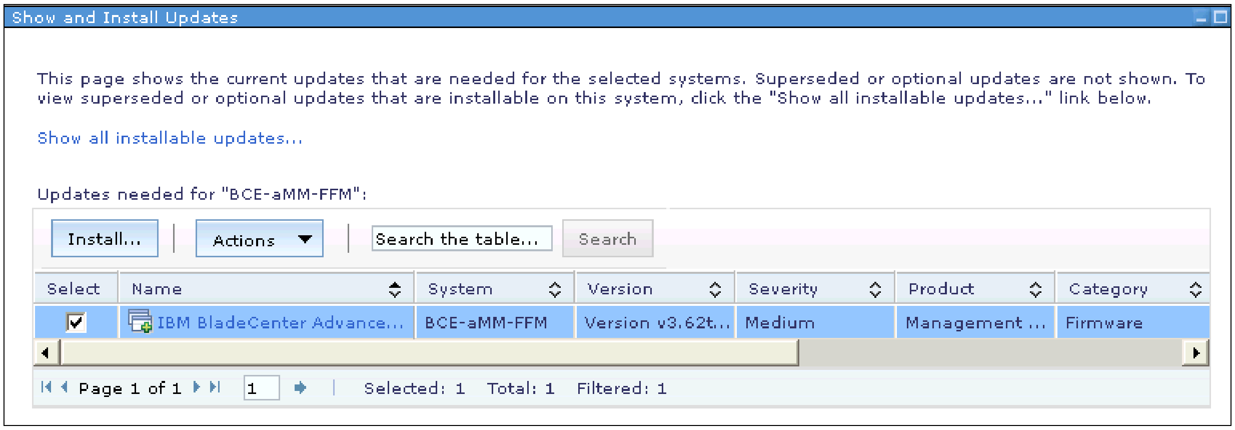

Show and Install Updates

This page shows the current updates that are needed for the selected systems. Superseded or optional updates are not shown. To view superseded or optional updates that are installable on this system, click the "Show all installable updates..." link below.

Show all installable updates...

Updates needed for "BCE-aMM-FFM":

Select	Name	System	Version	Severity	Product	Category
☑	IBM BladeCenter Advance...	BCE-aMM-FFM	Version v3.62t...	Medium	Management ...	Firmware

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

Figure 4-64 Select Install to start the installation process

5. For either method, the Install Wizard window opens (Figure 4-65). On the Welcome panel, click **Next** to proceed with the installation process.

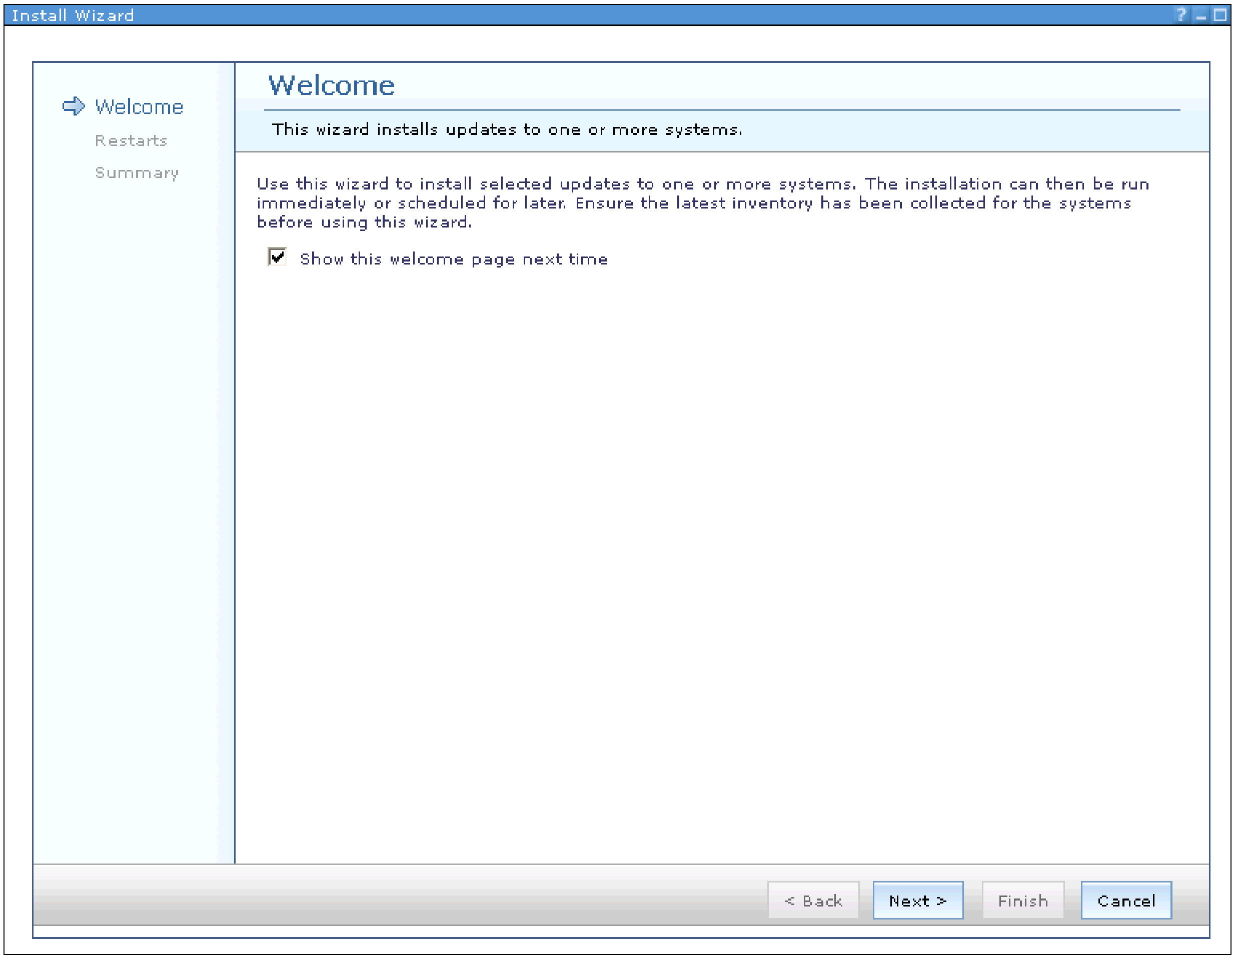

Welcome

This wizard installs updates to one or more systems.

Use this wizard to install selected updates to one or more systems. The installation can then be run immediately or scheduled for later. Ensure the latest inventory has been collected for the systems before using this wizard.

☑ Show this welcome page next time

Figure 4-65 Install wizard: Welcome panel

6. On Figure 4-66, you see the selected system. By default, the update process automatically restarts the systems, if needed. The specific systems to update are listed, and whether a restart is required is listed. If you clear the check box, the update is installed. However, you get an error message that the system is not restarted if a restart is required by the update process. You need to restart the system manually before the update takes effect. Click **Next** to continue.

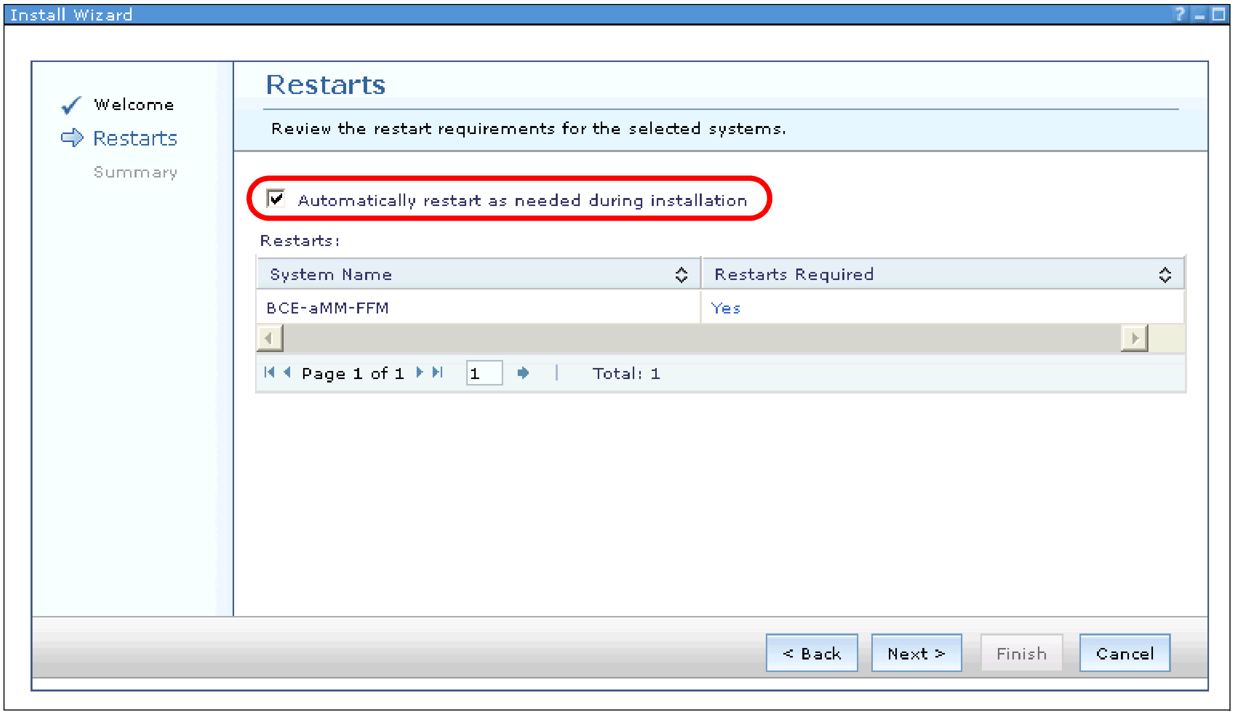

Restarts

Review the restart requirements for the selected systems.

☑ Automatically restart as needed during installation

Restarts:

System Name	Restarts Required
BCE-aMM-FFM	Yes

Page 1 of 1 | 1 | Total: 1

Figure 4-66 Install wizard: Restarts panel

7. A Summary window opens. Check the settings and the information of the update package. Click **Finish** to start the installation upgrade (Figure 4-67).

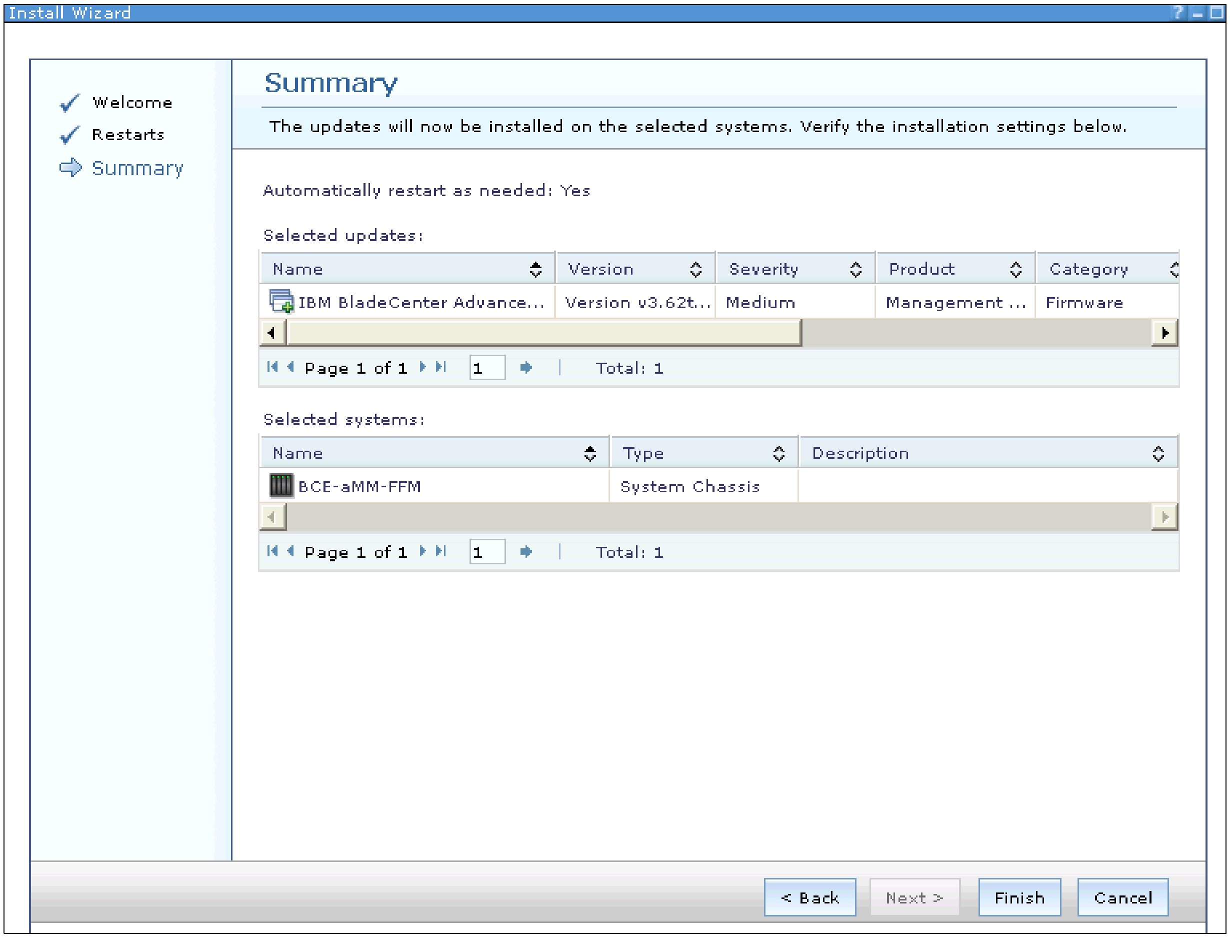

Summary

The updates will now be installed on the selected systems. Verify the installation settings below.

Automatically restart as needed: Yes

Selected updates:

Name	Version	Severity	Product	Category
IBM BladeCenter Advance...	Version v3.62t...	Medium	Management ...	Firmware

Page 1 of 1 | 1 | Total: 1

Selected systems:

Name	Type	Description
BCE-aMM-FFM	System Chassis	

Page 1 of 1 | 1 | Total: 1

Figure 4-67 Install wizard: Summary panel

8. The Schedule tab opens. Select whether to run the update now or at a defined time (Figure 4-68).

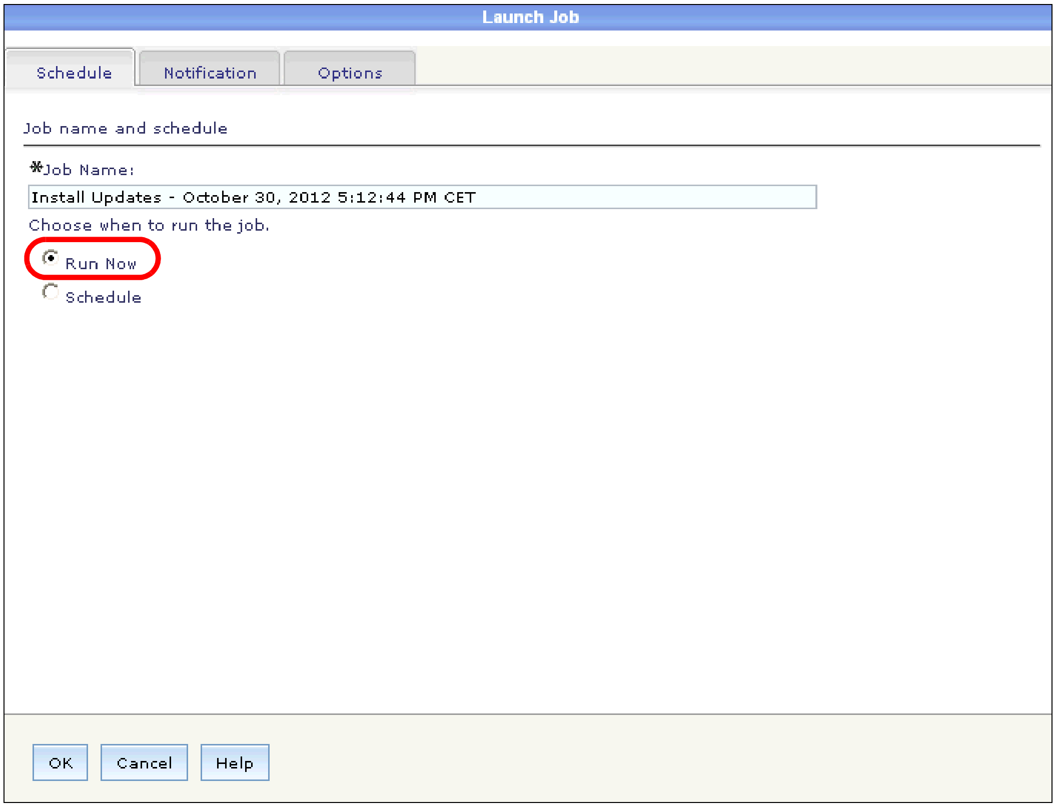

*Job Name:
Install Updates - October 30, 2012 5:12:44 PM CET

Choose when to run the job.

● Run Now
○ Schedule

Figure 4-68 Schedule tab

9. If you select Run Now, a window opens where you can see that the job is created and started, as shown in Figure 4-69.

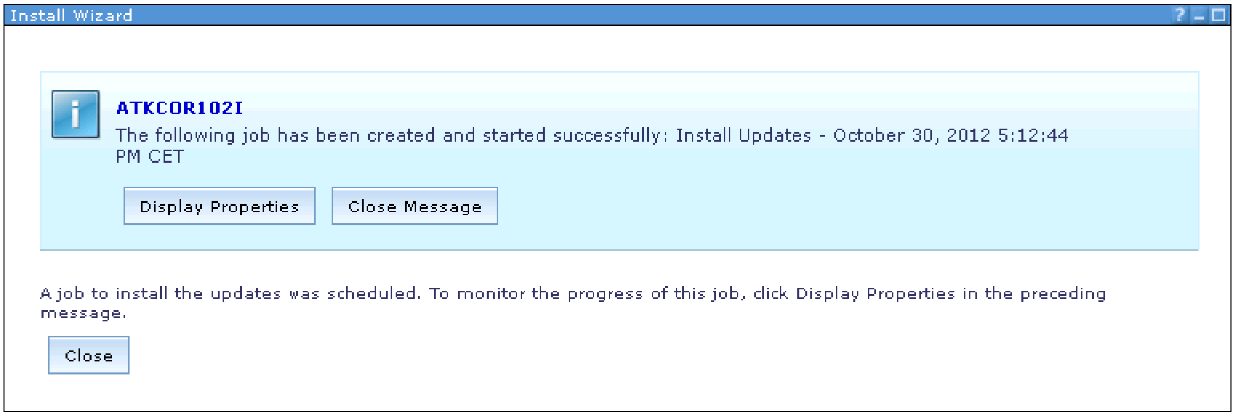

ATKCOR102I
The following job has been created and started successfully: Install Updates - October 30, 2012 5:12:44 PM CET

A job to install the updates was scheduled. To monitor the progress of this job, click Display Properties in the preceding message.

Figure 4-69 Job created and started

10. Click **Display Properties** to display the job properties where you can check the status and the log for the job (Figure 4-70). You can see that if the update was not downloaded before or imported, a download for the update package is started.

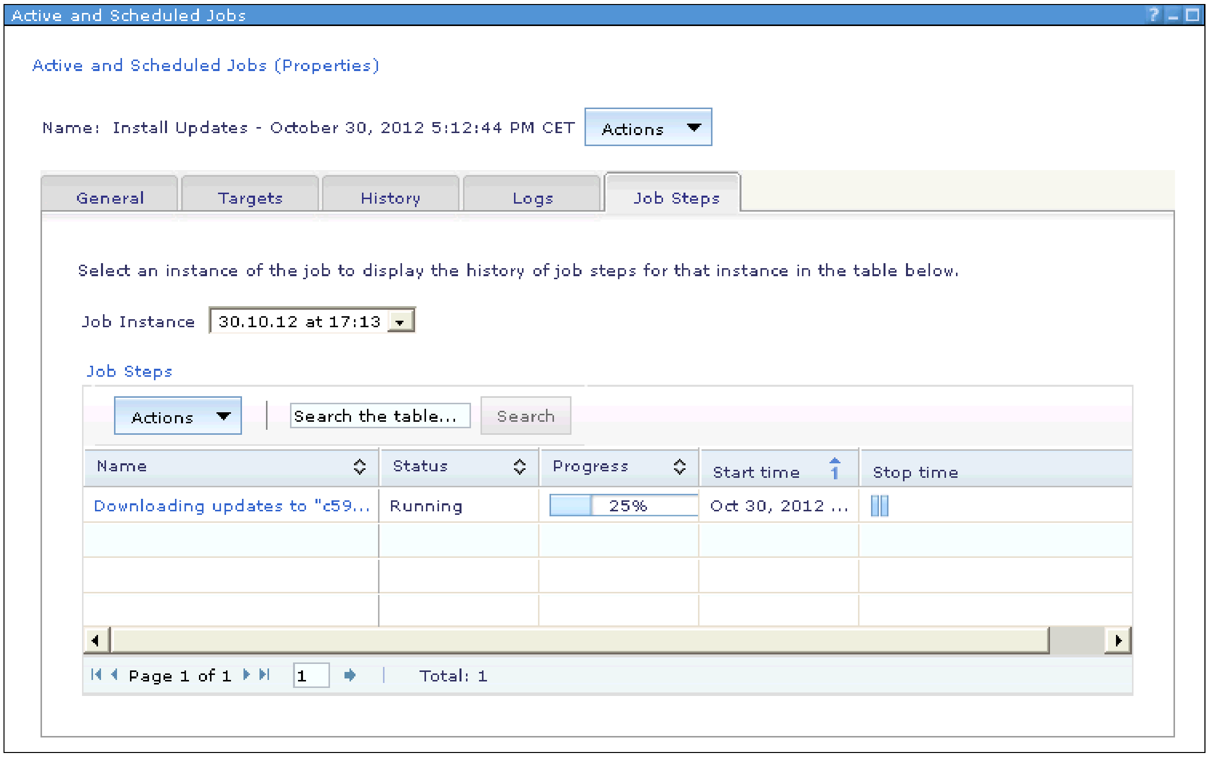

Active and Scheduled Jobs (Properties)

Name: Install Updates - October 30, 2012 5:12:44 PM CET

Select an instance of the job to display the history of job steps for that instance in the table below.

Job Instance 30.10.12 at 17:13

Job Steps

Search the table...

Name	Status	Progress	Start time	Stop time
Downloading updates to "c59...	Running	25%	Oct 30, 2012 ...	

Total: 1

Figure 4-70 Active and Scheduled Jobs (Properties) panel: Job Steps tab to download updates

Using the compliance check to update

Compliance checks are described in 4.3.6, “Compliance check” on page 114). If you set up a compliance check and a system is noncompliant, start here to update your system.

Follow these steps to update systems from the compliance check:

1. Click the link beside the red, yellow (in this example), or blue icon from the compliance check that is shown in Figure 4-71.

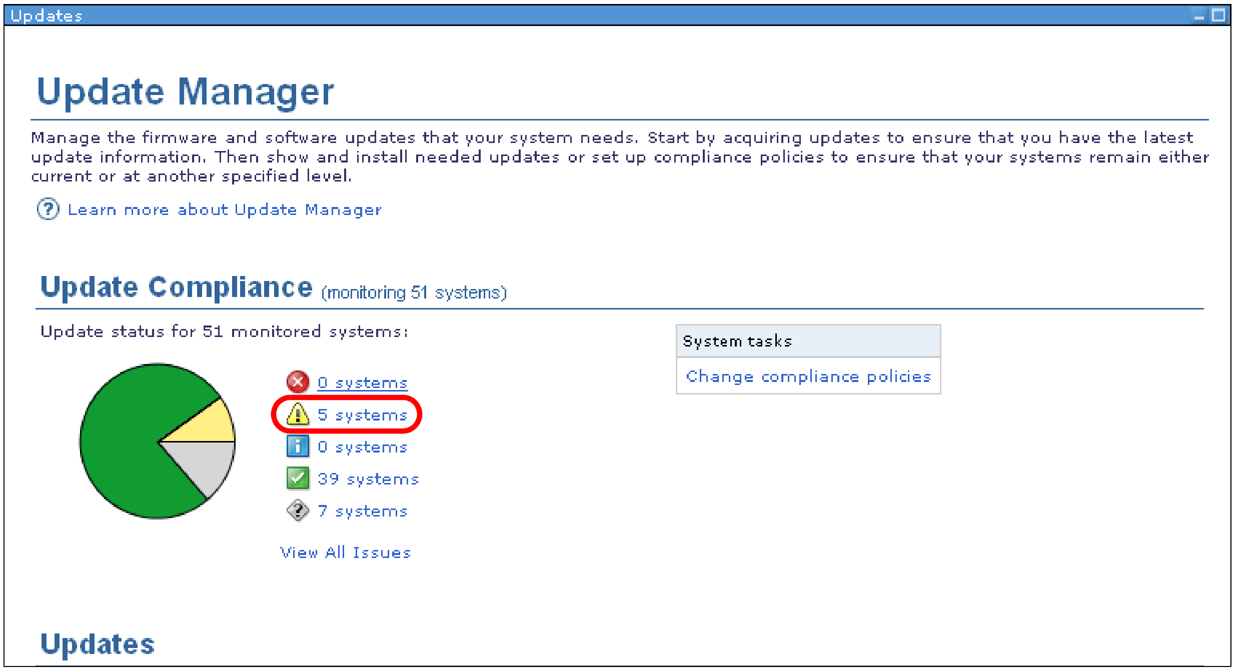

Updates

Update Manager

Manage the firmware and software updates that your system needs. Start by acquiring updates to ensure that you have the latest update information. Then show and install needed updates or set up compliance policies to ensure that your systems remain either current or at another specified level.

Learn more about Update Manager

Update Compliance (monitoring 51 systems)

Update status for 51 monitored systems:

- 0 systems
- 5 systems
- 0 systems
- 39 systems
- 7 systems

View All Issues

System tasks

Change compliance policies

Updates

Figure 4-71 Compliance check: Select a system with a problem

2. In the Navigate Resources window, you can see all the updates for the selected severity. Our example shows five systems with minor severity compliance issues (Figure 4-72). Click the check box next to the systems that you want to update. Then, click **Actions** → **Release Management** → **Show and Install updates**.

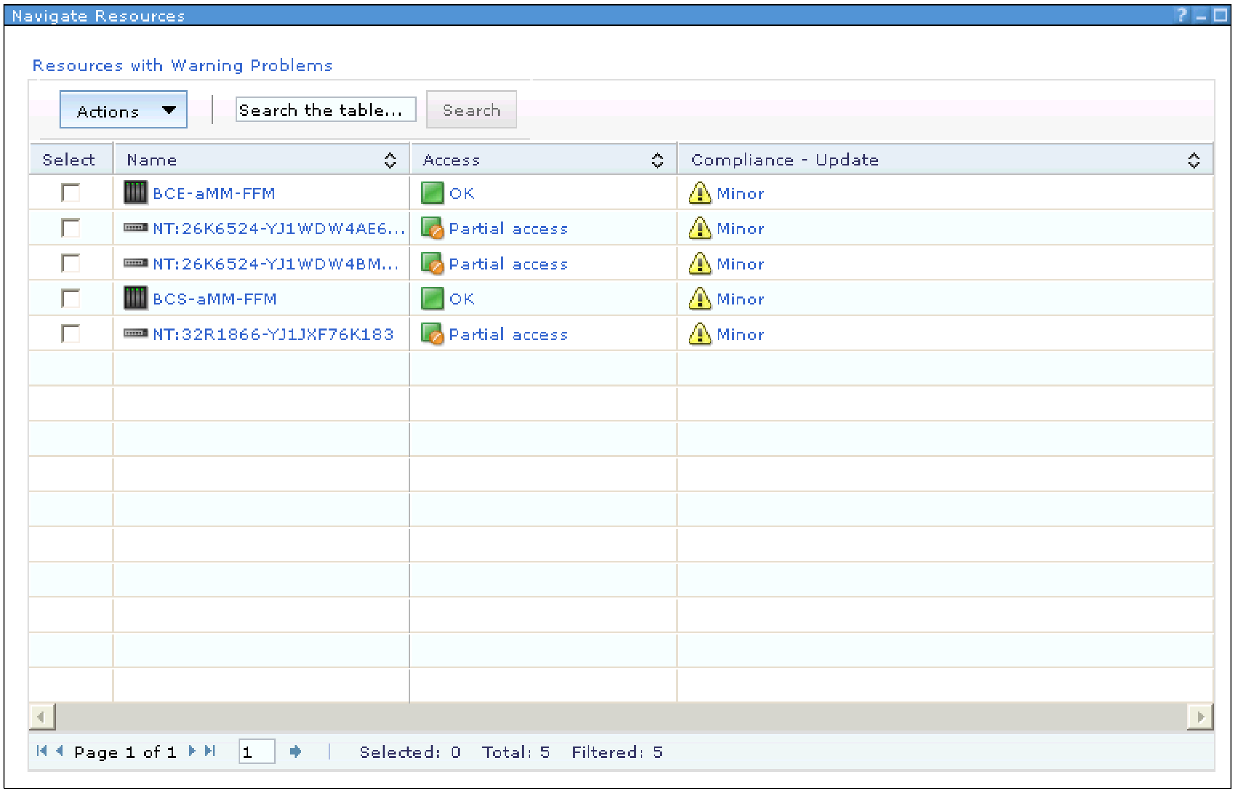

Select	Name	Access	Compliance - Update
□	BCE-aMM-FFM	OK	Minor
□	NT:26K6524-YJ1WDW4AE6...	Partial access	Minor
□	NT:26K6524-YJ1WDW4BM...	Partial access	Minor
□	BCS-aMM-FFM	OK	Minor
□	NT:32R1866-YJ1JXF76K183	Partial access	Minor

Page 1 of 1 | 1 | Selected: 0 Total: 5 Filtered: 5

Figure 4-72 Compliance check: Systems with minor severity updates

3. The Install wizard opens. The remaining steps are the same as described in “Using Update Manager: Show and install updates” on page 122, starting with step 5 on page 125.

Another way to update systems with compliance issues is to select the number beside the icon in the status bar on top of the Systems Director home page, as shown in Figure 4-73.

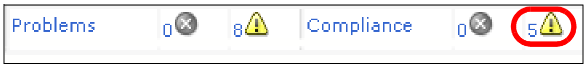

Figure 4-73 Status bar: Compliance status

The Active Status window opens (Figure 4-74) to show the systems with a compliance issue. This window differs from the window that you see when you select the compliance check from Update Manager (Figure 4-71 on page 130).

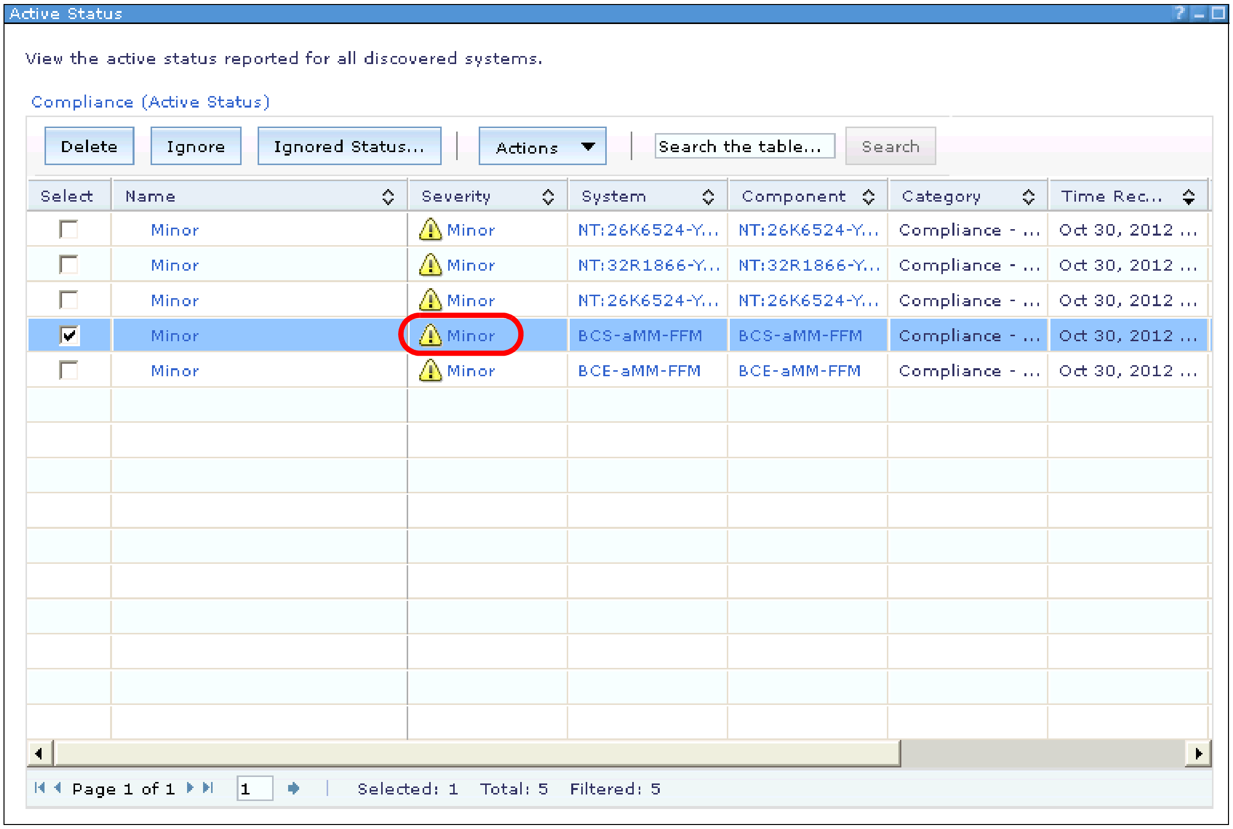

View the active status reported for all discovered systems.

Select	Name	Severity	System	Component	Category	Time Rec...
☐	Minor	⚠ Minor	NT:26K6524-Y...	NT:26K6524-Y...	Compliance - ...	Oct 30, 2012 ...
☐	Minor	⚠ Minor	NT:32R1866-Y...	NT:32R1866-Y...	Compliance - ...	Oct 30, 2012 ...
☐	Minor	⚠ Minor	NT:26K6524-Y...	NT:26K6524-Y...	Compliance - ...	Oct 30, 2012 ...
☑	Minor	⚠ Minor	BCS-aMM-FFM	BCS-aMM-FFM	Compliance - ...	Oct 30, 2012 ...
☐	Minor	⚠ Minor	BCE-aMM-FFM	BCE-aMM-FFM	Compliance - ...	Oct 30, 2012 ...

Page 1 of 1 | 1 | Selected: 1 Total: 5 Filtered: 5

Figure 4-74 System with minor severity compliance issues (example)

But, in our example, when you click the **Minor** link in the Severity column, you see the same window as shown before. The way to update is the same as described earlier.

4.3.8 Updating systems that run AIX and Linux

How to update AIX and Linux systems is described.

Updating Linux systems

Systems Director can update the Linux operating system with the latest patches from either Red Hat or SUSE. To use Update Manager in the Systems Director server for these updates, the system that you want to update must be registered with the provider.

To download the updates from the provider, follow steps 1 - 3 in 4.3.4, “Update Manager with Internet connection” on page 110.

In step 3, select the Red Hat and SUSE/Novell Linux updates in Update Manager that you need for your systems, as shown in Figure 4-75.

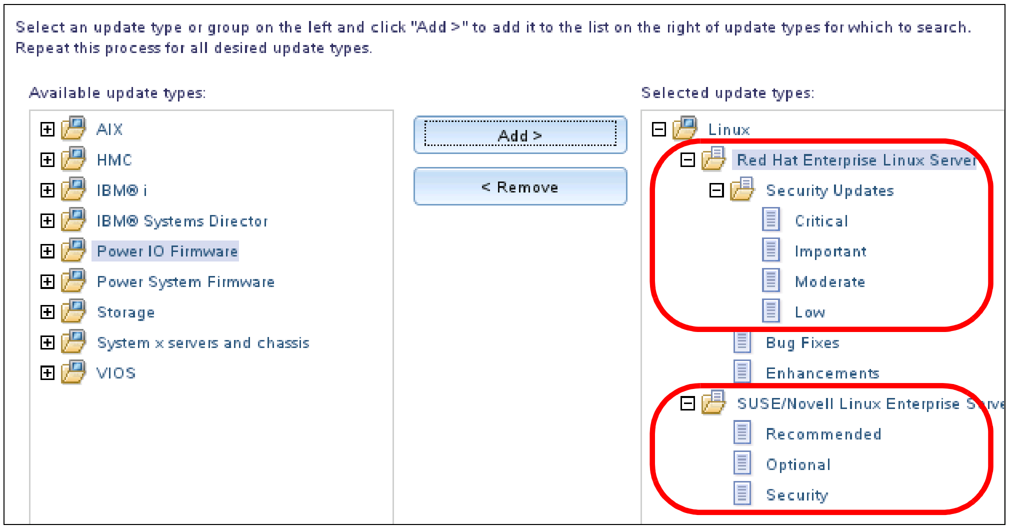

Select an update type or group on the left and click "Add >" to add it to the list on the right of update types for which to search. Repeat this process for all desired update types.

Figure 4-75 Select updates for Linux OS

You can also set a compliance check for your system for Linux updates. If the system discovers new Linux updates, use the same process for all other updates as described in 4.3.7, “Update process” on page 121.

Updating a system that runs AIX

The following page in the information center explains the requirements for updating AIX Systems:

http://pic.dhe.ibm.com/infocenter/director/pubs/index.jsp?topic=%2Fcom.ibm.director.updates.helps.doc%2Ffqm0_c_um_considerations_for_updating_aix_systems.html

AIX updates can be downloaded automatically from within Systems Director or manually from Fix Central:

http://www.ibm.com/support/fixcentral

When you download AIX fixes, ensure that you review your download options. This step is critical to ensure that you select **Include informational files, and files required by Systems Director for installation**. These files are needed to manually import AIX updates to Systems Director from the command line.

The syntax for importing AIX updates is shown in Figure 4-76.

```
#smcli importupd -vr /tmp/updates/aix/7100-01-06-1241
```

Figure 4-76 Importing AIX updates

For updates for AIX, Update Manager is the focal point for centralized management, and updates are performed by using NIM. Standard NIM troubleshooting procedures can be used, as needed.

Tip: If you import updates from a Network File System (NFS) mount, be careful with Secure Hash Algorithm (SHA). If the NFS mount is restricted, copy it locally to a temporary position on the Systems Director server and import to eliminate SHA warnings.

4.3.9 Updating the Systems Director server

If you want to update the Systems Director server, use the “Update IBM Systems Director” task. This task lets Update Manager use the defaults and run the update tasks for you automatically.

This task is accessible from the Systems Director home page (Figure 4-77) and also from the Update Manager page (Figure 4-78 on page 135).

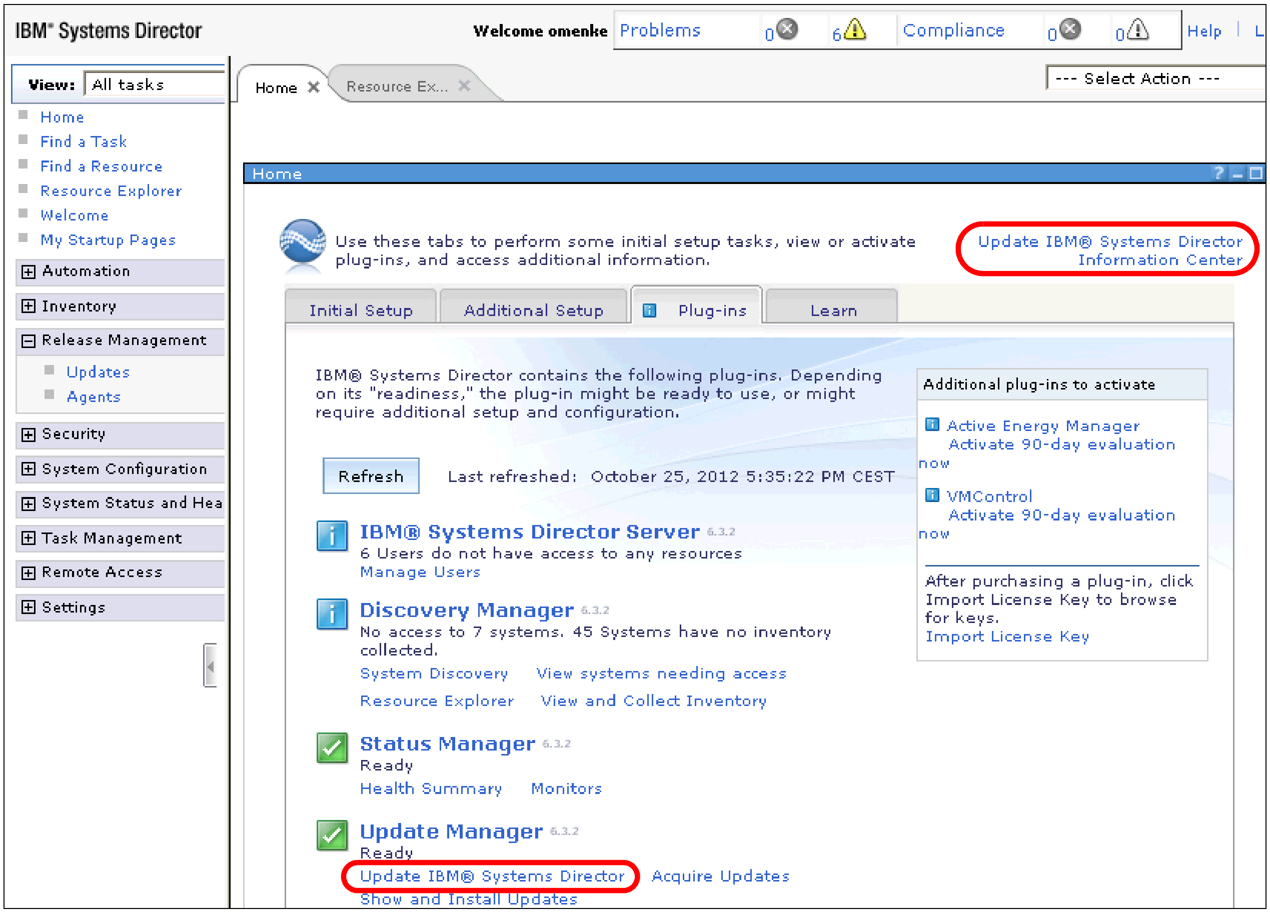

Use these tabs to perform some initial setup tasks, view or activate plug-ins, and access additional information.

Update IBM® Systems Director Information Center

IBM® Systems Director contains the following plug-ins. Depending on its "readiness," the plug-in might be ready to use, or might require additional setup and configuration.

Last refreshed: October 25, 2012 5:35:22 PM CEST

IBM® Systems Director Server 6.3.2
6 Users do not have access to any resources
Manage Users

Discovery Manager 6.3.2
No access to 7 systems. 45 Systems have no inventory collected.
System Discovery View systems needing access
Resource Explorer View and Collect Inventory

Status Manager 6.3.2
Ready
Health Summary Monitors

Update Manager 6.3.2
Ready
Update IBM® Systems Director Acquire Updates
Show and Install Updates

Additional plug-ins to activate

Active Energy Manager
Activate 90-day evaluation now

VMControl
Activate 90-day evaluation now

After purchasing a plug-in, click Import License Key to browse for keys.
Import License Key

Figure 4-77 Launching the Update IBM Systems Director task from the home page

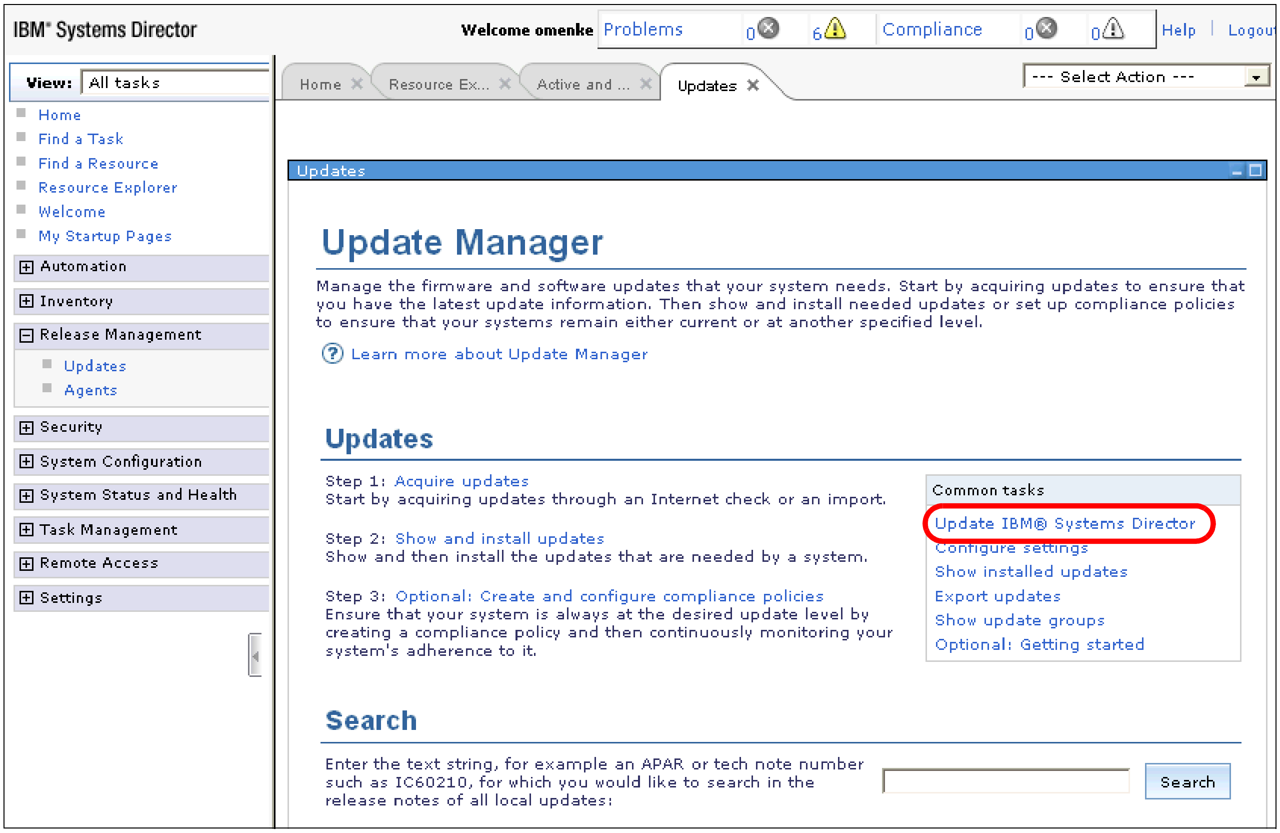

Update Manager

Manage the firmware and software updates that your system needs. Start by acquiring updates to ensure that you have the latest update information. Then show and install needed updates or set up compliance policies to ensure that your systems remain either current or at another specified level.

Learn more about Update Manager

Updates

Step 1: Acquire updates
Start by acquiring updates through an Internet check or an import.

Step 2: Show and install updates
Show and then install the updates that are needed by a system.

Step 3: Optional: Create and configure compliance policies
Ensure that your system is always at the desired update level by creating a compliance policy and then continuously monitoring your system's adherence to it.

Search

Enter the text string, for example an APAR or tech note number such as IC60210, for which you would like to search in the release notes of all local updates:

Figure 4-78 Launching the Update IBM Systems Director task from the Update Manager page

When you launch the update task, the system checks for available updates as shown in Figure 4-79.

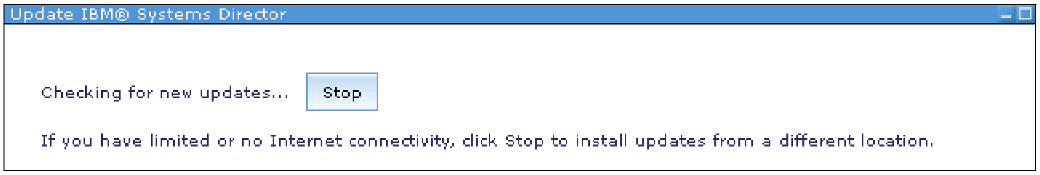

Checking for new updates...

If you have limited or no Internet connectivity, click Stop to install updates from a different location.

Figure 4-79 Update status for Systems Director update

If no new updates are available, you see a message that is similar to Figure 4-80.

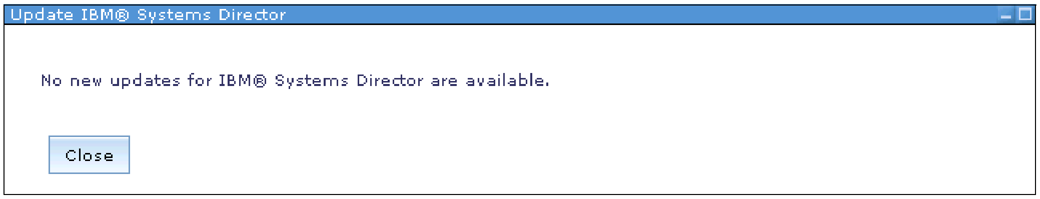

No new updates for IBM® Systems Director are available.

Figure 4-80 No updates are available

If new updates are available for your Systems Director server, you see a window similar to Figure 4-81.

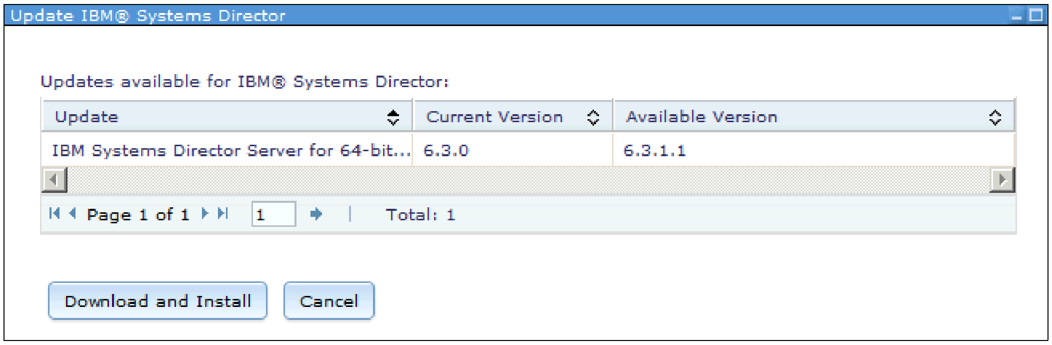

Update IBM® Systems Director

Updates available for IBM® Systems Director:

Update	Current Version	Available Version
IBM Systems Director Server for 64-bit...	6.3.0	6.3.1.1

Figure 4-81 Update Systems Director

Follow these steps to install the updates:

1. Click **Download and Install**, as shown in Figure 4-81. You are reminded to back up your server (Figure 4-82). We explain how to perform a backup in 4.7, “Backup” on page 233.

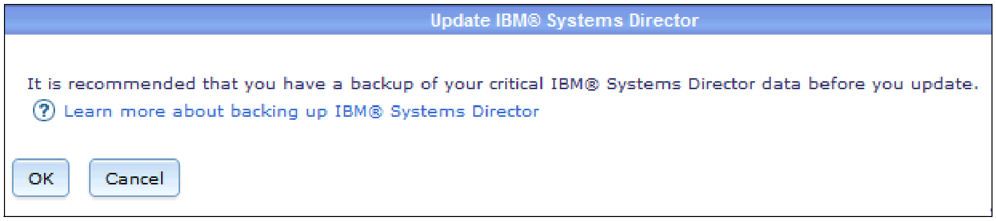

Update IBM® Systems Director

It is recommended that you have a backup of your critical IBM® Systems Director data before you update.

Figure 4-82 Information window about backup

2. Click **OK** to proceed to the Schedule window. Select **Run Now** and the updates are downloaded and installed. Or, you can schedule the update to be performed in off-hours. The status of the task is under **Task Management** → **Active and Scheduled Jobs**.
3. After the upgrade completes, the Systems Director server must be restarted. The best method is to use one of the following command-line commands:
   – Windows:

`net stop dirserver`	Stop the Systems Director server.
`net start dirserver`	Start the Systems Director server.

   – Linux and AIX:

`smstop`	Stop the Systems Director server.
`smstart`	Start the Systems Director server.

4. To check the status of the Systems Director, use the following command:
   – Windows:

     The status icon is on the Windows panel. Use the **`smstatus.bat (-r)`** command to see the status and the update of the status.

   – Linux and AIX:

     Use the **`smstatus (-r)`** command to see the status.

After the Systems Director server restarts, check whether the new version is installed and running. Check the version beside each manager on the home page. Or, check the `version.srv` file in the directory where the Systems Director server is installed.

4.3.10 Downloading and staging updates manually

When planning for maintenance on your servers, you usually want to keep the change window as small as possible so that your business does not have to suffer long service outages.

Some driver and firmware updates can be fairly large, which means that unless you have already installed them once, the night of the change, IBM Systems Director Update Manager will require additional time to fetch those update packages from IBM Fix Central.

One way to work around this issue is to manually download the updates ahead of time so that they are already available on your IBM Systems Director server when you need them:

1. First, figure out which updates you will require and, which are already downloaded. To do so, select the server that you want to update in the resource explorer, right-click it, and select **Related Resources** → **Updates**, as shown in Figure 4-83.

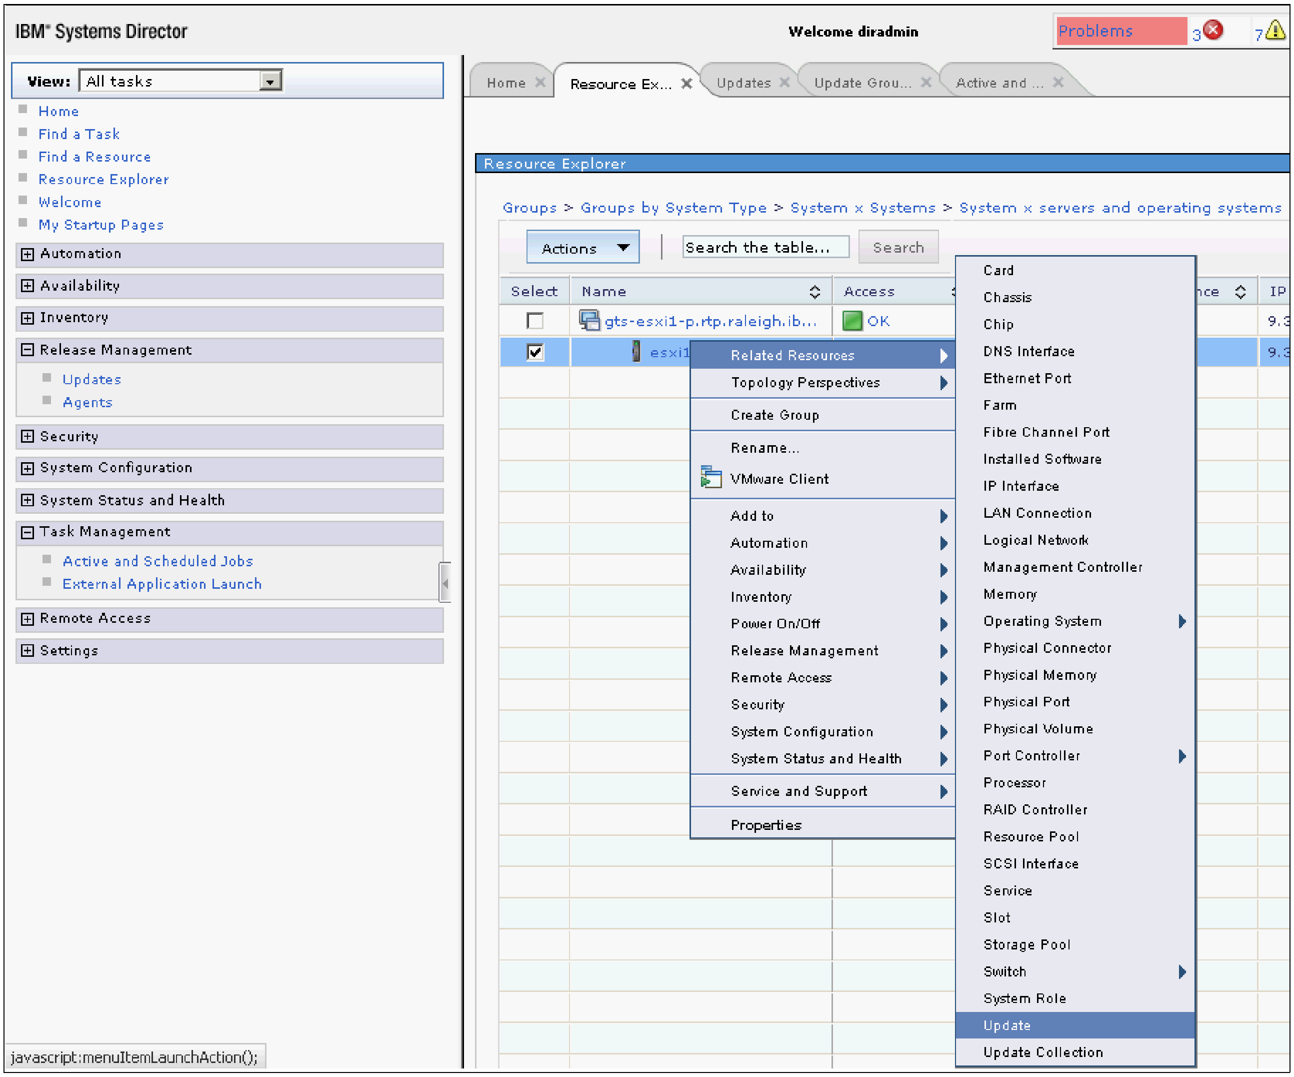

Figure 4-83 Finding out which updates might be required on a given server

2. Review the list of applicable updates and whether they are already downloaded or not, as shown in Figure 4-84.

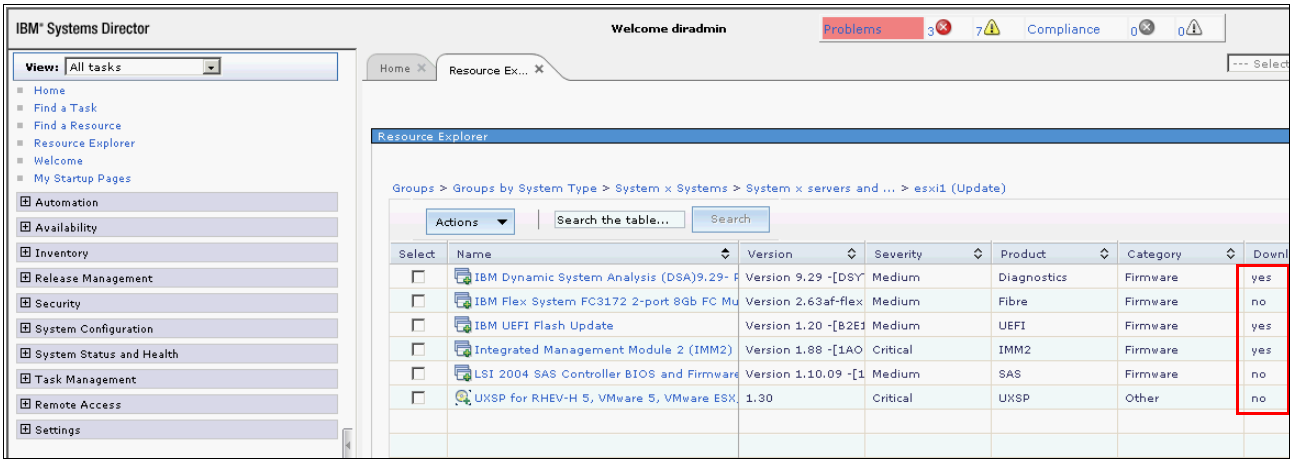

Select	Name	Version	Severity	Product	Category	Downl
□	IBM Dynamic System Analysis (DSA)9.29- F	Version 9.29 -[DSY	Medium	Diagnostics	Firmware	yes
□	IBM Flex System FC3172 2-port 8Gb FC Mu	Version 2.63af-flex	Medium	Fibre	Firmware	no
□	IBM UEFI Flash Update	Version 1.20 -[B2E1	Medium	UEFI	Firmware	yes
□	Integrated Management Module 2 (IMM2)	Version 1.88 -[1AO	Critical	IMM2	Firmware	yes
□	LSI 2004 SAS Controller BIOS and Firmware	Version 1.10.09 -[1	Medium	SAS	Firmware	no
□	UXSP for RHEV-H 5, VMware 5, VMware ESX	1.30	Critical	UXSP	Other	no

Figure 4-84 Reviewing applicable updates

3. You can then select updates that have not been downloaded yet, click the **Actions** menu, and select **Download**, as shown in Figure 4-85.

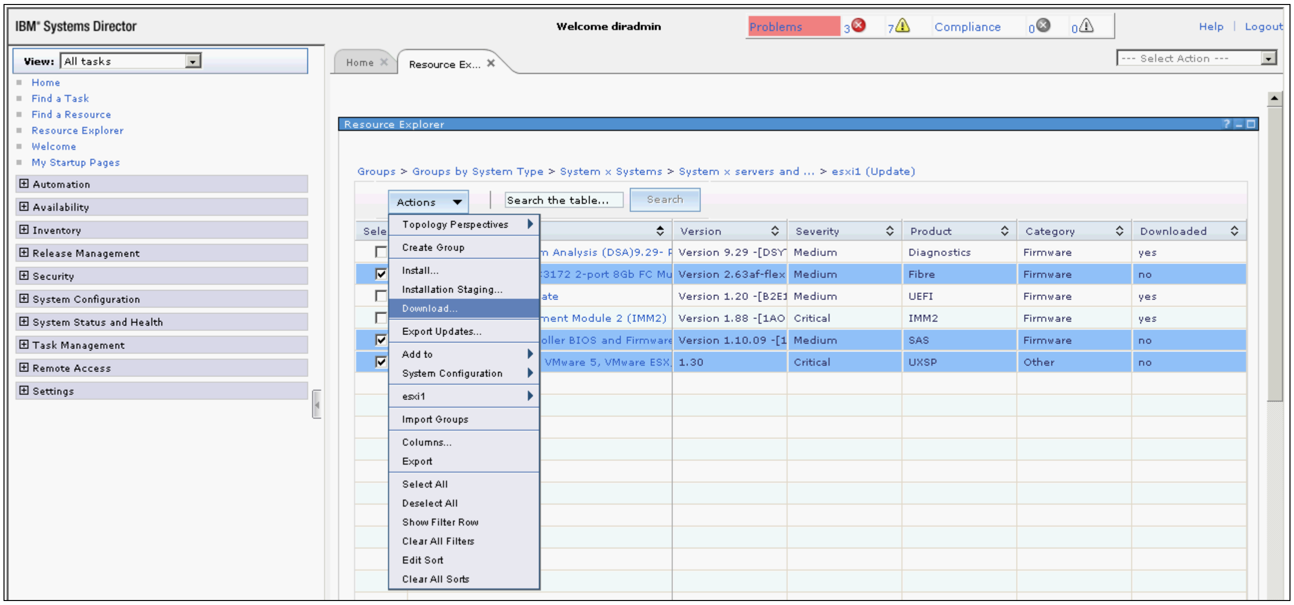

Sele		Version	Severity	Product	Category	Downloaded
□	n Analysis (DSA)9.29- F	Version 9.29 -[DSY	Medium	Diagnostics	Firmware	yes
☑	3172 2-port 8Gb FC Mu	Version 2.63af-flex	Medium	Fibre	Firmware	no
□	ate	Version 1.20 -[B2E1	Medium	UEFI	Firmware	yes
□	ment Module 2 (IMM2)	Version 1.88 -[1AO	Critical	IMM2	Firmware	yes
☑	oller BIOS and Firmware	Version 1.10.09 -[1	Medium	SAS	Firmware	no
☑	VMware 5, VMware ESX	1.30	Critical	UXSP	Other	no

Figure 4-85 Downloading updates

Alternatively, if you already know which updates you need and want to simply select them to download them, you can do so from the Update Groups view. This option is selectable from the Update Manager main window, as shown in Figure 4-86.

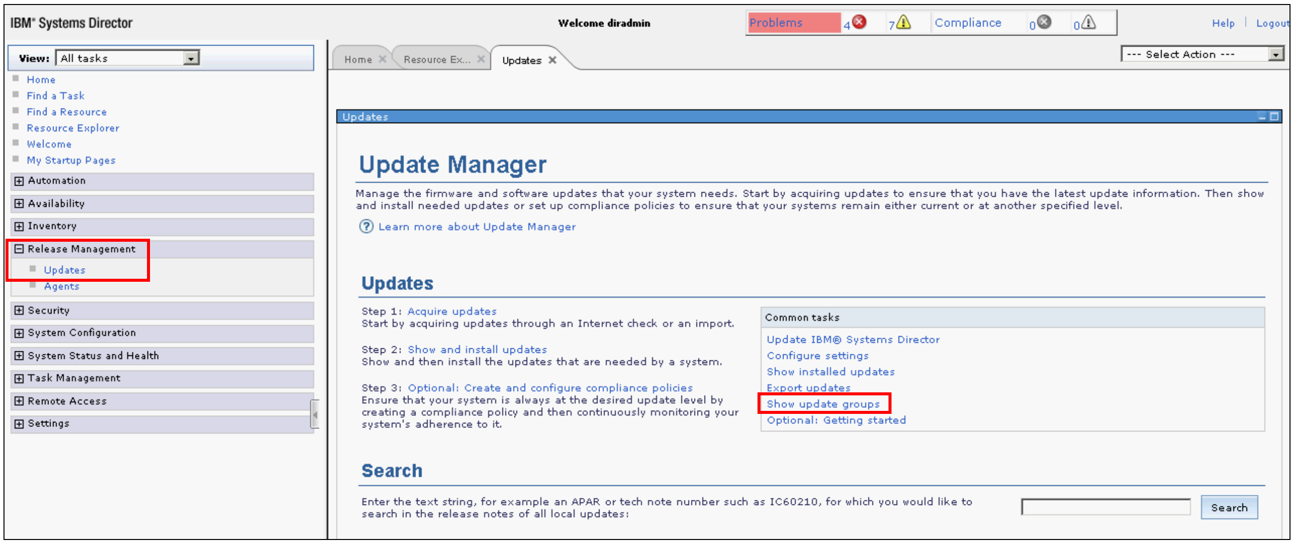

Update Manager

Manage the firmware and software updates that your system needs. Start by acquiring updates to ensure that you have the latest update information. Then show and install needed updates or set up compliance policies to ensure that your systems remain either current or at another specified level.

Learn more about Update Manager

Updates

Step 1: Acquire updates
Start by acquiring updates through an Internet check or an import.

Step 2: Show and install updates
Show and then install the updates that are needed by a system.

Step 3: Optional: Create and configure compliance policies
Ensure that your system is always at the desired update level by creating a compliance policy and then continuously monitoring your system's adherence to it.

Common tasks
Update IBM® Systems Director
Configure settings
Show installed updates
Export updates
Show update groups
Optional: Getting started

Search

Enter the text string, for example an APAR or tech note number such as IC60210, for which you would like to search in the release notes of all local updates:

Figure 4-86 The Update Manager main window

From this view, you are able to select and download all the updates that you want.

In addition to downloading updates, you can also pre-stage updates on a given server. This is useful if the server you want to update is in a remote location and you also want to cut the transfer time from the IBM Systems Director server to the managed endpoint.

To pre-stage updates, perform the following steps:

1. Discover applicable updates on your server by running an inventory and then navigating to the **Release Management** → **Show and Install Updates** view.
2. Select the updates that you want to stage and select **Installation Staging** from the **Actions** menu.
3. Follow the wizard to the end to start scheduling the staging job. The updates will be copied over to the managed endpoint and will be ready for installation.

Alternatively, you can select updates from the **Update Groups** view that is shown in Figure 4-86 and select **Installation Staging** from the **Actions** menu to stage them on multiple servers at once.

4.3.11 Command-line tools

Command-line tools are available for Update Manager as listed in Table 4-1.

Table 4-1 Command-line tools for Update Manager

Command	Description
`checkupd`	Check changed and superseding updates.
`cleanupd`	Clean (that is, delete) update files and information in the local update library.
`importupd`	Import updates into the update library on the management server. This command is used if no Internet access is available.

Command	Description
installneeded	Update the Systems Director server and agents or use this command to install other types of updates.
installupd	Install one or more updates to one or more systems.
lsupd	List the available updates and their attributes.
lsver	List the current version and, if you updated the product, the previous version of Systems Director that is installed on the system.
uninstallupd	Use to uninstall (roll back) an update on a specific system if the update package supports the rollback.

For detailed information about all the **smcli** command-line commands that are used for the Update Manager and their options, see the information center:

http://pic.dhe.ibm.com/infocenter/director/pubs/index.jsp?topic=%2Fcom.ibm.director.cli.helps.doc%2Ffqm0_r_cli_update_cmds.html

Certain functions are only available by using the command line. One example is cleaning or deleting the local update library. Use this function if multiple downloaded updates are not used or if you are running out of storage. After you clean the library, run a check for updates to fill the library with new update information. *The information from the library is used for the compliance check.*

In Example 4-1, we clean up the library and then start a check for new updates from the command line. You can also run the check for new updates in the browser interface as described in 4.3.4, “Update Manager with Internet connection” on page 110.

In Example 4-1, we first list the updates that are downloaded to the local update library with the **smcli lsupd** command. For brevity in this example, some output lines are not displayed (200 packages are available).

Example 4-1 smcli lsupd command

```
PS C:\Windows\system32> smcli lsupd
SysDir6_3_1_Platform_Agent_Windows
SysDir6_3_1_Platform_Agent_xLinux
SysDir6_3_Platform_Agent_AIX
SysDir6_3_Platform_Agent_Windows
SysDir6_3_Platform_Agent_pLinux
SysDir6_3_Platform_Agent_xLinux
SysDir6_3_Platform_Agent_zLinux
agentmanager.feature_6.3.1
bnt_fw_bcsw_110gup-6.3.1.1_anyos_noarch
bnt_fw_bcsw_110gup-7.2.2.0_anyos_noarch
bnt_fw_bcsw_24-10g-6.9.1.0_anyos_noarch
bnt_fw_bcsw_24-10g-7.2.2.0_anyos_noarch
bnt_fw_torsw_g8264-6.8.4.0_anyos_noarch
bnt_fw_torsw_g8316-6.8.4.0_anyos_noarch
brcd_fw_6.3.1-dcb2_anyos_noarch
brcd_fw_bcsw_sansm-505a_anyos_noarch
cigesm-i6q4l2-tar.121-22.ea13
com.ibm.aem.common_4.4.1
com.ibm.aem.console_4.4.1
.....
```

Then, we clean up the library by using the **smcli cleanupd -am** command, as listed in Example 4-2. This command needs time to delete all packages and remove the index file. After running the cleaning of the local library (**smcli cleanupd -am**), we check with the **smcli lsupd** command whether update packages are still available. You can see in the example that after cleaning, no installation package is available.

You can clean up only one or some of the installation packages that are in the local library. Use **-w %packagename%** instead of the **-am** option.

Example 4-2 smcli cleanupd -am command

```
PS C:\Windows\system32> smcli cleanupd -am
PS C:\Windows\system32> smcli lsupd

PS C:\Windows\system32>
```

We start to download new packages by using the **smcli checkupd -a** command, which checks the IBM repository for all updates. You can also use other options, such as the **-N groupname**, to check only for updates for a member of a defined group (Example 4-3).

Example 4-3 smcli checkupd -a command

```
PS C:\Windows\system32> smcli checkupd -a
PS C:\Windows\system32>
```

The **-a** option needs a long time to finish. No output is listed during the download process if it is run on the command line. You can see that the command is finished only when a new command prompt is visible.

When the command is finished, check which updates are downloaded by running the **smcli lsupd** command (Example 4-4). You can see examples from the listing to show the different types of updates that are downloaded. In this example, over 500 update packages are available (for Linux, AIX, firmware, driver, Director, and VIOS).

Example 4-4 smcli lsupd command with examples for different downloads

```
PS C:\Windows\system32> smcli lsupd
01AF743_100_100
01AF743_105_100
....
032512EE02F845008725779E00509382_AIX
032512EE02F845008725779E00509382_LNX
03F50ADC70A9CEA9872577B200727962_AIX
....
MH01084
MH01097
MH01101
MH01102
....
U823341
U824377
U824378
....
VIOS_2.2.1.3-FP25-SP01
VIOS_2.2.1.4-FP25-SP02
VIOS_2.2.2.1-FP26
....
agentmanager.feature_6.3.1
....
```

```
bnt_fw_bcsw_110gup-6.3.1.1_anyos_noarch
bnt_fw_bcsw_110gup-7.2.2.0_anyos_noarch
bnt_fw_bcsw_24-10g-6.9.1.0_anyos_noarch
.....
com.ibm.aem.common_4.4.1
com.ibm.aem.console_4.4.1
com.ibm.aem.discovery_4.4.1
....
com.ibm.director.storage.storagecontrol.member.AIX_4.2.2.build-00119
com.ibm.director.storage.storagecontrol.member.Linux_4.2.2.build-00095-20120516-iFix
....
csco_fw_bcio_12.2.50se1_anyos_noarch
ibm_fw_amm_bpet62t_anyos_noarch
ibm_fw_bcio_N4K_4.1.2.E1.1i_anyos_noarch
ibm_utl_uxspi_9.21_rhel5_32-64
ibm_utl_uxspi_9.21_sles11_32-64
ibm_utl_uxspi_9.21_winsrvr_32-64
```

4.3.12 Best practices for Update Manager

The following best practices enable you to use the Update Manager task efficiently:

- *Make all necessary configurations for the Update Manager before using it.*

  Set up the Internet connection, the settings for the local repository, and necessary settings for the BladeCenter, AIX, or VIOS to prevent problems during the installation process.

- *Test the Internet connection (when available).*

  The Internet connection is necessary for the download of the latest information about available updates.

- *Configure the Update Manager for scheduled checks* for new updates, at minimum once a week.

  This helps to stay up-to-date with the update information and is necessary for the compliance check.

- *Set the scheduled download in off-hours.*

  Prevent the environment from high network utilization and when the director download begins only update information.

- *Be aware that the inventory is up-to-date.*

  This is necessary for the compliance check as part of the Update Manager. The compliance check works as a background process and compares the information from the repository with the inventory information of the systems (see best practices for inventory).

- *Hold the IBM Systems Director up-to-date by using the IBM Systems Director update task.*

- For the IBM System x server and Flex System compute nodes, *ensure that you leave “LAN over USB” enabled* in the management processor configuration.

- *Download manually your updates ahead of time and consider staging* updates to keep your maintenance window as small as possible.

- *Configure the compliance check* that fits your requirements.

4.4 Event management

This section explains what type of events are collected by IBM Systems Director and covers best practices on how to integrate with your existing event management system.

To assist the system administrator, IBM Systems Director provides the Event Automation Plan wizard as a quick and easy way to create event automation plans that meet most system management needs.

By using the Event Automation Plan wizard, you can create plans that monitor for the most typical situations in systems management environments, including, but not limited to, the following examples:

- Critical hardware events from all systems in your environment.
- Processor (CPU) utilization in a specific group of systems, such as all servers running Linux.
- All Common Agent-managed systems to determine if Common Agent goes offline.
- The status of updates that are underway.
- Disk space usage in systems, such as those that store database data.

You can then configure actions to be performed based on the situation. For example, actions could send a page or email message, or start a program on a system.

4.4.1 Light path diagnostics

IBM x86 servers provide a diagnostic tool called *light path diagnostics* (LPDs) as an easy way to find hardware problems on the systems when they occur. LPDs consist of three components:

- A system warning LED on the front of the server.
- A panel of LEDs in a pop-out panel (or a panel inside or outside the server for some systems). This panel shows the status of major subsystems, for example, memory.
- Individual LEDs beside each component in the system, for example, each memory dual inline memory module (DIMM).

For information about the LPD LEDs that are available on your System x, BladeCenter, or Flex System hardware, see the server documentation.

IBM Systems Director can read this LPD information. This information is provided by the service processors that are integrated in the server:

- Integrated management module (IMM)
- IMMv2
- Advanced management module (AMM)
- Chassis management module (CMM)

You can view LPD status information from the Systems Director user interface (UI) or the command-line interface (CLI).

There are multiple ways to select the detailed LPD information:

- From the health summary, click **LED status** in the scoreboard.
- From the Resource Explorer window, in the LED Status column, click the red, yellow, or blue icon.

- From the right-click menu of a system, go to **System status and health** → **Lightpath**.
- From a command-line prompt, use the `smcli lsled` command.

Each method is described in the following sections.

LED status in the scoreboard

On the home page of Systems Director, the health summary scoreboard shows a summary of systems with LPD alerts under the name “LED Status”. Figure 4-87 shows an example.

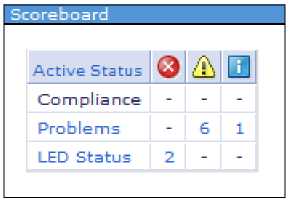

Scoreboard

Active Status	⊗	⚠	ℹ
Compliance	-	-	-
Problems	-	6	1
LED Status	2	-	-

Figure 4-87 Scoreboard that shows the LED status

When you click **LED Status**, a window opens that shows systems with alerts, as shown in Figure 4-88.

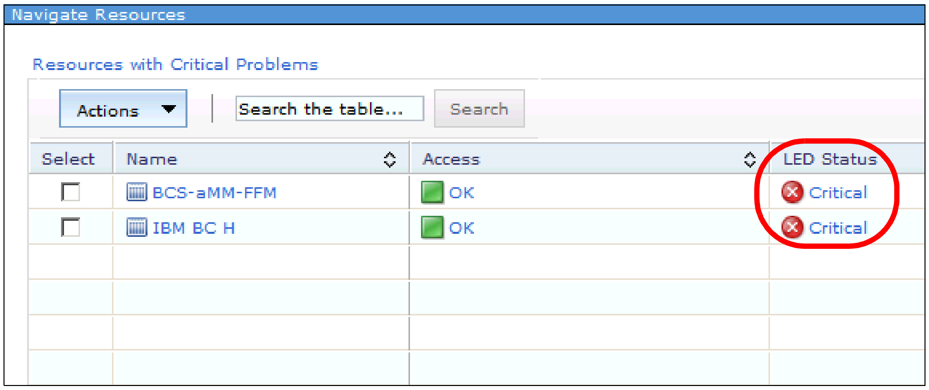

Navigate Resources

Resources with Critical Problems

Select	Name	Access	LED Status
□	BCS-aMM-FFM	🟩 OK	⊗ Critical
□	IBM BC H	🟩 OK	⊗ Critical

Figure 4-88 Systems with problems

If you click the LED status of the individual systems, you see the specific alert details as reported in the lightpath view, as shown in Figure 4-89.

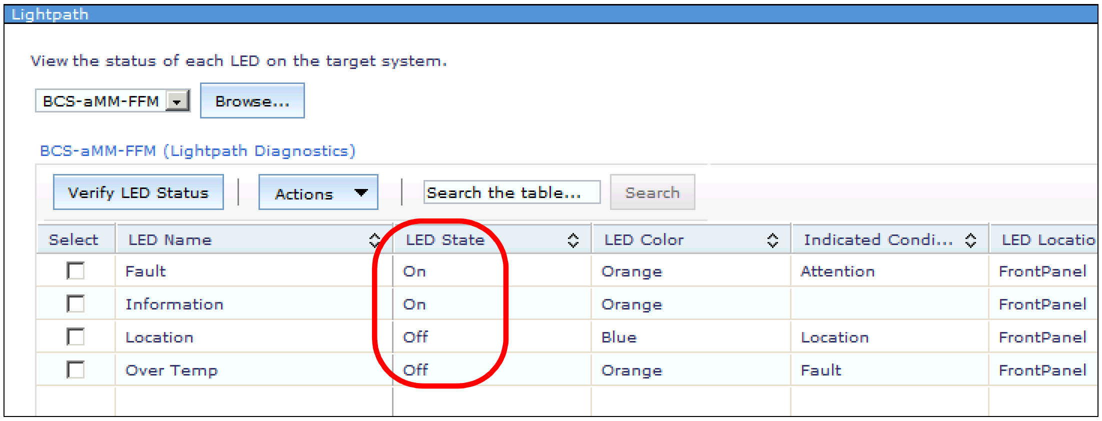

Lightpath

View the status of each LED on the target system.

BCS-aMM-FFM (Lightpath Diagnostics)

Select	LED Name	LED State	LED Color	Indicated Condi...	LED Locatio
□	Fault	On	Orange	Attention	FrontPanel
□	Information	On	Orange		FrontPanel
□	Location	Off	Blue	Location	FrontPanel
□	Over Temp	Off	Orange	Fault	FrontPanel

Figure 4-89 Lightpath detailed view

LED status in the Resource Explorer

To view the LPD status and information, first add a column for LED status to the Resource Explorer view for a group. When you are in a group in the Resource Explorer window, click **Actions** → **Columns**, as shown in Figure 4-90.

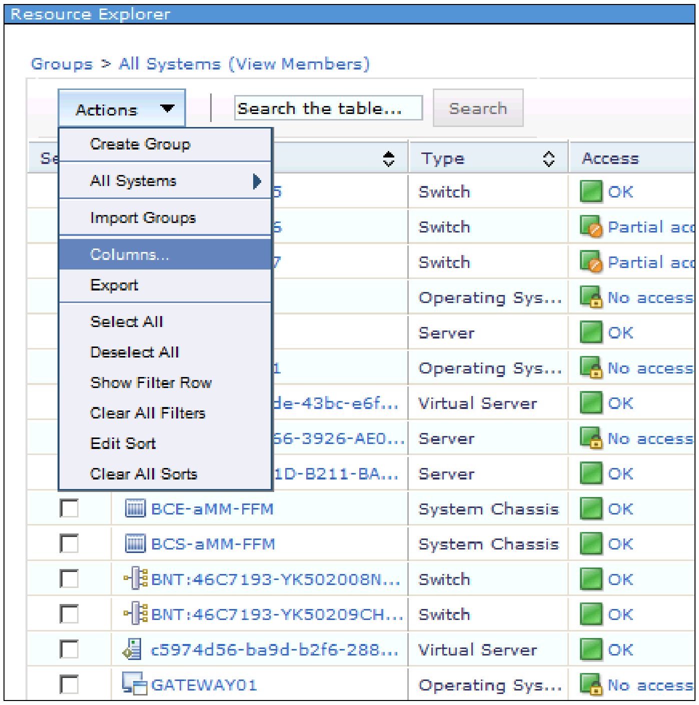

Se		Type	Access
	5	Switch	OK
	6	Switch	Partial acc
	7	Switch	Partial acc
		Operating Sys...	No access
		Server	OK
	1	Operating Sys...	No access
	de-43bc-e6f...	Virtual Server	OK
	56-3926-AE0...	Server	No access
	1D-B211-BA...	Server	OK
□	BCE-aMM-FFM	System Chassis	OK
□	BCS-aMM-FFM	System Chassis	OK
□	BNT:46C7193-YK502008N...	Switch	OK
□	BNT:46C7193-YK50209CH...	Switch	OK
□	c5974d56-ba9d-b2f6-288...	Virtual Server	OK
□	GATEWAY01	Operating Sys...	No access

Figure 4-90 Select to add columns to the group view

In Figure 4-91, select **LED Status** from the Available Columns list on the left. Click **Add** to add LED Status to the Selected Columns list. Use **Up** and **Down** to change the relative position of the column. Click **OK** to save the changes. See Figure 4-91.

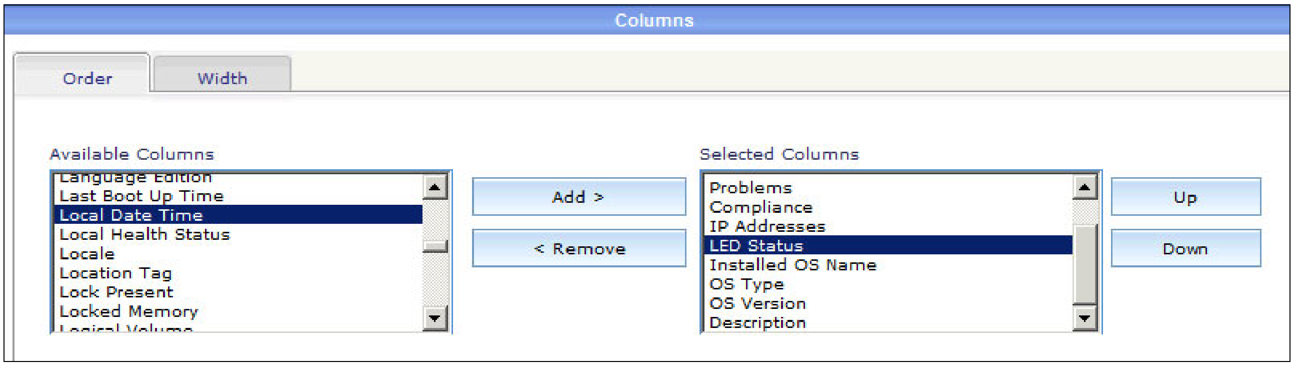

Available Columns

Language Edition
Last Boot Up Time
Local Date Time
Local Health Status
Locale
Location Tag
Lock Present
Locked Memory
Logical Volume

Selected Columns

Problems
Compliance
IP Addresses
LED Status
Installed OS Name
OS Type
OS Version
Description

Figure 4-91 Select LED Status to add a column to the Resource Explorer view

Now, the new LED Status column is displayed in the Resource Explorer window as shown in Figure 4-92.

If a critical, warning, or informational status message from a managed system exists on the LPD panel, you see the status. The status displays as a red, yellow, or blue icon in the LED Status column. See Figure 4-92.

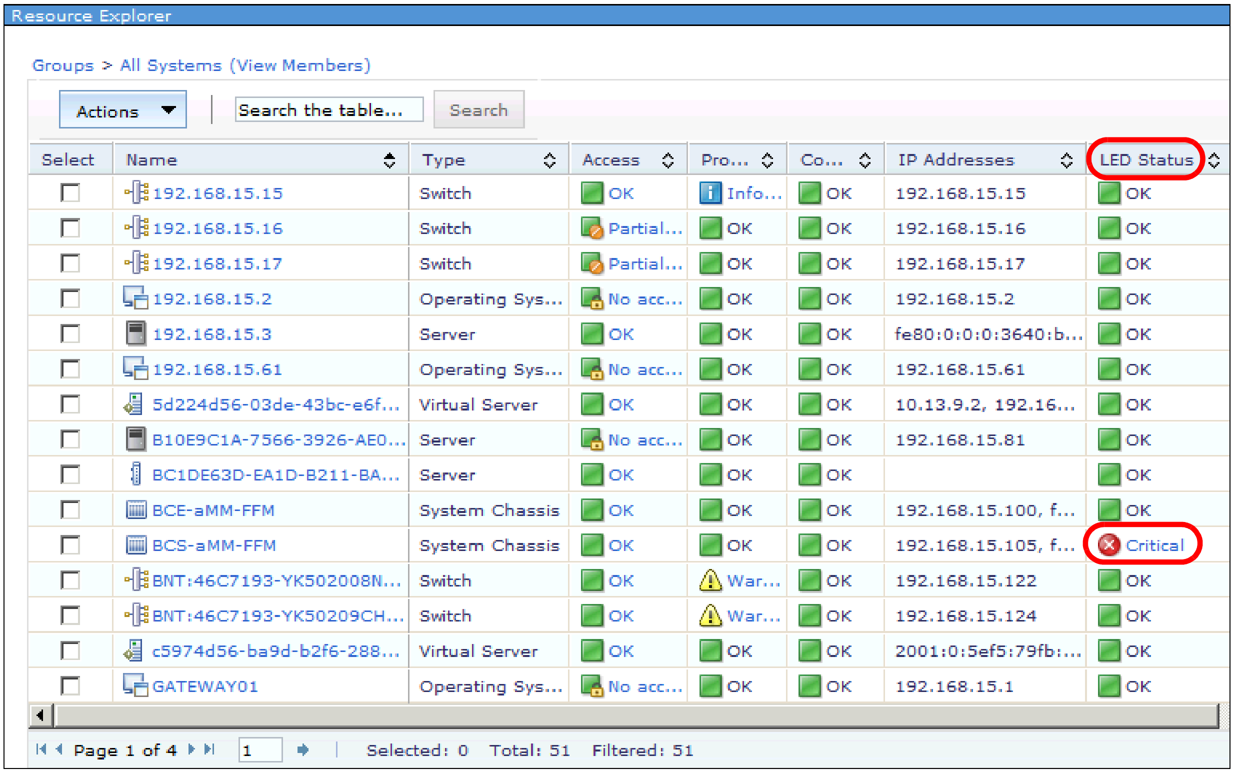
Resource Explorer
Groups > All Systems (View Members)
Actions ▼
Search the table...
Search
Page 1 of 4
1
Selected: 0 Total: 51 Filtered: 51

Select	Name	Type	Access	Pro...	Co...	IP Addresses	LED Status
□	192.168.15.15	Switch	OK	Info...	OK	192.168.15.15	OK
□	192.168.15.16	Switch	Partial...	OK	OK	192.168.15.16	OK
□	192.168.15.17	Switch	Partial...	OK	OK	192.168.15.17	OK
□	192.168.15.2	Operating Sys...	No acc...	OK	OK	192.168.15.2	OK
□	192.168.15.3	Server	OK	OK	OK	fe80:0:0:0:3640:b...	OK
□	192.168.15.61	Operating Sys...	No acc...	OK	OK	192.168.15.61	OK
□	5d224d56-03de-43bc-e6f...	Virtual Server	OK	OK	OK	10.13.9.2, 192.16...	OK
□	B10E9C1A-7566-3926-AE0...	Server	No acc...	OK	OK	192.168.15.81	OK
□	BC1DE63D-EA1D-B211-BA...	Server	OK	OK	OK		OK
□	BCE-aMM-FFM	System Chassis	OK	OK	OK	192.168.15.100, f...	OK
□	BCS-aMM-FFM	System Chassis	OK	OK	OK	192.168.15.105, f...	Critical
□	BNT:46C7193-YK502008N...	Switch	OK	War...	OK	192.168.15.122	OK
□	BNT:46C7193-YK50209CH...	Switch	OK	War...	OK	192.168.15.124	OK
□	c5974d56-ba9d-b2f6-288...	Virtual Server	OK	OK	OK	2001:0:5ef5:79fb:...	OK
□	GATEWAY01	Operating Sys...	No acc...	OK	OK	192.168.15.1	OK

Figure 4-92 LED Status column in Resource Explorer

After you add the column, you can click the status to see the detailed view (Figure 4-89 on page 144).

LED Status from the menu of a system

To see the LED status and detailed information for a single system, right-click the system and click **System Status and Health** → **Lightpath**. See Figure 4-93 on page 147. The window that opens is similar to the window from our example that is shown in Figure 4-89 on page 144.

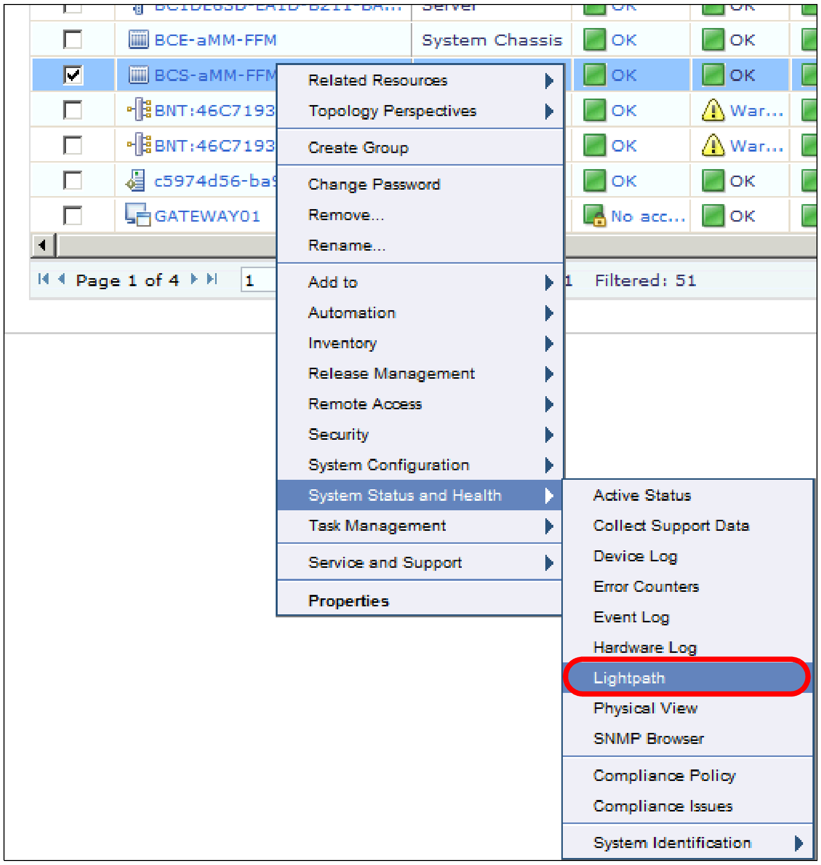

Figure 4-93 Lightpath menu for a single system

SMCLI command-line interface

You can also use the **smcli lsled** CLI command to see the LED status for a system. The results of our example are shown in Example 4-5.

Example 4-5 is the output from the command for a BladeCenter chassis that shows the status of all LEDs. This output shows that the Information and Fault LEDs for this system are on. The output also shows that the LEDs are on the front panel of the system.

Example 4-5 smcli lsled CLI command

```
PS C:\Windows\system32> smcli lsled -s all -i 9.42.171.73
-----------------------------------------------------------------
System Name: BC5AMM----------------------------------------------

Name                State               Color               Location

Over Temp           Off                 Orange              FrontPanel

Information         On                  Orange              FrontPanel

Location            Off                 Blue                FrontPanel

Fault               On                  Orange              FrontPanel
-----------------------------------------------------------------
```

If you want to see only the LEDs that are on or flashing, use the **`-s all`**, **`-s on`**, or **`-s flash`** option for this command. Or, use **`-s on, flash`** to see all LEDs that are on and blinking. In this example, we use the following command to make the LED on a remote server blink:

```
smcli runtask -i 9.42.171.173 "LED Flash"
```

Then, we run the **`smcli lsled`** command with the **`-s flash`** option again. Example 4-6 shows the result.

Example 4-6 smcli lsled -s flash option

```
PS C:\Windows\system32> smcli lsled -s flash -i 9.42.171.73
-----------------------------------------------------------------
System Name: BC5AMM----------------------------------------------


Name              State             Color             Location

Location          Blinking          Blue              FrontPanel

-----------------------------------------------------------------
```

For detailed information about the options for the **`smcli lsled`** command, see the information center:

http://pic.dhe.ibm.com/infocenter/director/pubs/index.jsp?topic=%2Fcom.ibm.director.cli.helps.doc%2Ffqm0_r_cli_lsled.html

4.4.2 Hardware logs

Hardware log information is provided by the service processors from the systems. The following service processors and management modules provide the hardware log information to the Systems Director:

- Baseboard management controller (BMC)
- Remote Supervisor Adapter II (RSA II)
- Remote Supervisor Adapter (RSA)
- Management module (MM)
- Integrated management module (IMMv1 and IMMv2)
- Advanced management module (AMM)
- Chassis management module (CMM)

Power Systems: LPD information for Power Systems is not accessible, except for IBM POWER® based servers in the BladeCenter and Flex System.

The information is provided by using an inband communication or out-of-band communication. The access path depends on the system hardware and configuration. See Table 4-2.

Table 4-2 Service processor hardware log access path

Systems	Service processor	Hardware log access path
BladeCenter	MM	Out-of-band communication
BladeCenter	AMM	Out-of-band communication
Flex System	CMM	Out-of-band communication, only if no Flex System Manager (FSM) is installed and used

Systems	Service processor	Hardware log access path
System x	RSA	Out-of-band communication
System x	RSA II	Out-of-band communication
System x BladeCenter	BMC	Out-of-band communication. Inband communication that uses Common Agent or Platform Agent
System x	IMMv1 or IMMv2 in rack or tower server	Out-of-band communication. Inband communication that uses Common Agent or Platform Agent
BladeCenter Flex System	Flex System and BladeCenter IMMv1 or IMMv2 in server	Out-of-band communication over AMM/CMM only. Inband communication that uses Common Agent or Platform Agent

Inband communication means that Systems Director accesses the agent on the system. This agent can read the hardware log information from the service processor (RSA, RSA II, IMM, IMMv2, or BMC) of the system. The agent uses a driver or other communication channels inside the system.

Out-of-band communication means that a direct connection exists from the Systems Director to the service processor over a TCP/IP communication. This communication is independent from the system state (power on/off). This communication is also independent from the operating system (running, starting, stopped). The minimum requirement is that the system has power and the Systems Director can access the service processor.

To access the hardware log information inband, you must have full access to the system and the system must be online. If you obtain the hardware log out-of-band, you can also access it from the System x server when this server is powered off. In Table 4-3, you can see which resource you must select to access the hardware log information.

Table 4-3 Selection of system resources for accessing the hardware log

Option	Description
Inband communication	Select the system that represents the Common Agent or Platform Agent.
Out-of-band communication with a system	Click the system, then select service processors or select the server.
Out-of-band communication with a BladeCenter chassis	Select the chassis.

Follow these steps to access the log information:

1. From the Resource Explorer window, right-click the system and click **Systems Status and Health** → **Hardware Log** (Figure 4-94).

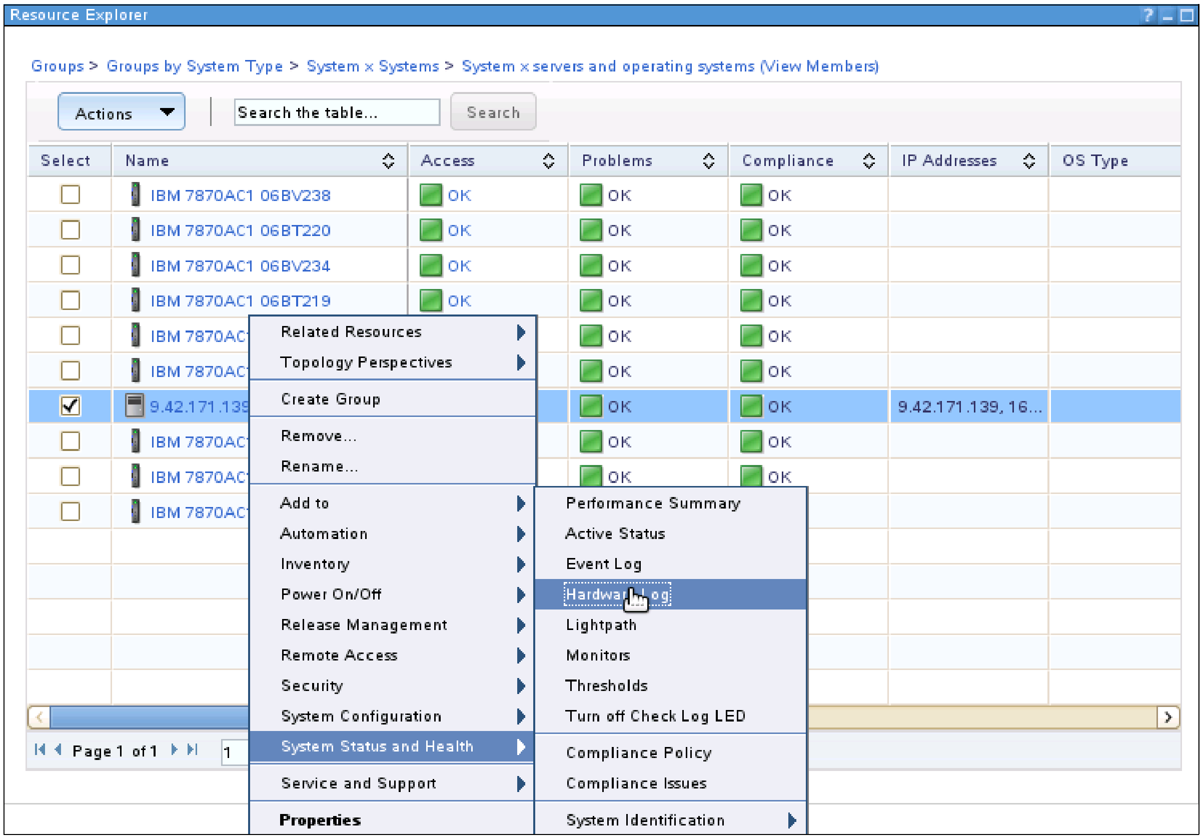

Groups > Groups by System Type > System x Systems > System x servers and operating systems (View Members)

Select	Name	Access	Problems	Compliance	IP Addresses	OS Type
□	IBM 7870AC1 06BV238	🟩 OK	🟩 OK	🟩 OK		
□	IBM 7870AC1 06BT220	🟩 OK	🟩 OK	🟩 OK		
□	IBM 7870AC1 06BV234	🟩 OK	🟩 OK	🟩 OK		
□	IBM 7870AC1 06BT219	🟩 OK	🟩 OK	🟩 OK		
□	IBM 7870AC		🟩 OK	🟩 OK		
□	IBM 7870AC		🟩 OK	🟩 OK		
☑	9.42.171.139		🟩 OK	🟩 OK	9.42.171.139, 16...	
□	IBM 7870AC		🟩 OK	🟩 OK		
□	IBM 7870AC		🟩 OK	🟩 OK		
□	IBM 7870AC					

Figure 4-94 Select hardware log

2. The window in Figure 4-95 opens to show the log entries that the Systems Director reads from the service processor.

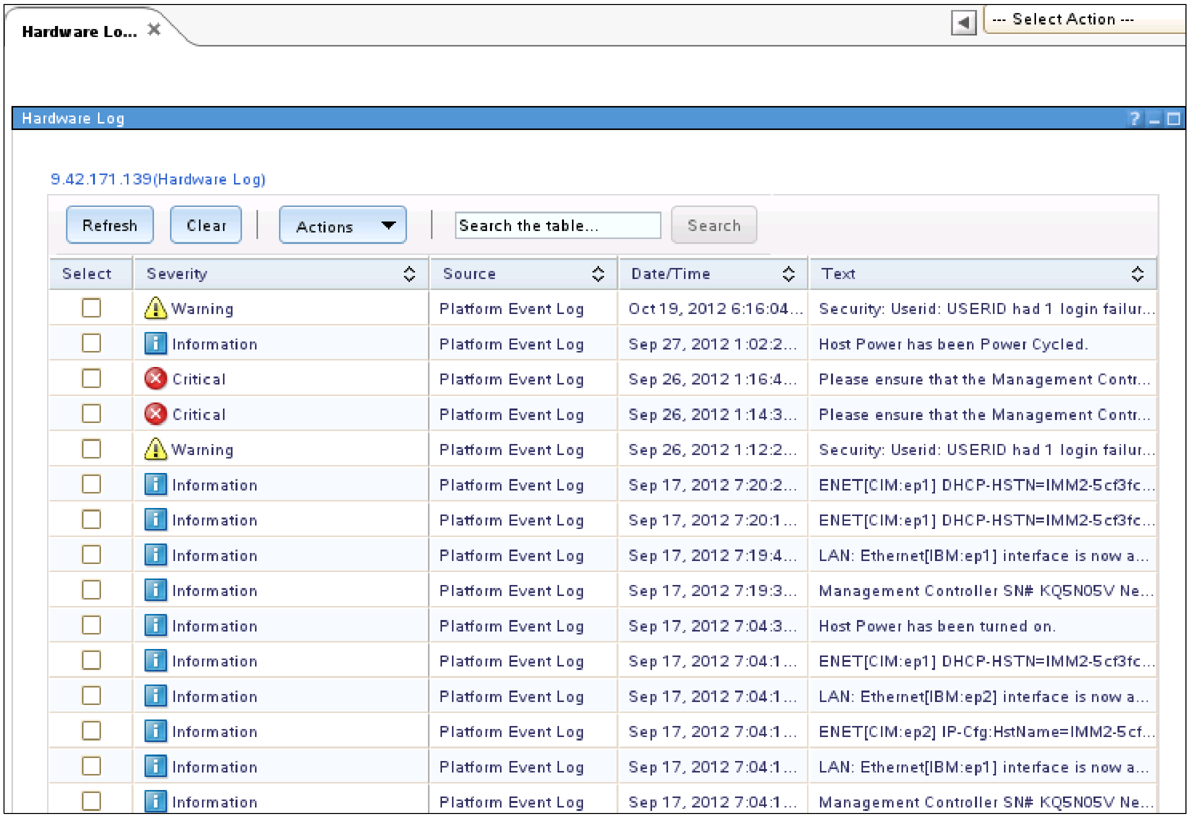

Select	Severity	Source	Date/Time	Text
□	⚠ Warning	Platform Event Log	Oct 19, 2012 6:16:04...	Security: Userid: USERID had 1 login failur...
□	Information	Platform Event Log	Sep 27, 2012 1:02:2...	Host Power has been Power Cycled.
□	Critical	Platform Event Log	Sep 26, 2012 1:16:4...	Please ensure that the Management Contr...
□	Critical	Platform Event Log	Sep 26, 2012 1:14:3...	Please ensure that the Management Contr...
□	⚠ Warning	Platform Event Log	Sep 26, 2012 1:12:2...	Security: Userid: USERID had 1 login failur...
□	Information	Platform Event Log	Sep 17, 2012 7:20:2...	ENET[CIM:ep1] DHCP-HSTN=IMM2-5cf3fc...
□	Information	Platform Event Log	Sep 17, 2012 7:20:1...	ENET[CIM:ep1] DHCP-HSTN=IMM2-5cf3fc...
□	Information	Platform Event Log	Sep 17, 2012 7:19:4...	LAN: Ethernet[IBM:ep1] interface is now a...
□	Information	Platform Event Log	Sep 17, 2012 7:19:3...	Management Controller SN# KQ5N05V Ne...
□	Information	Platform Event Log	Sep 17, 2012 7:04:3...	Host Power has been turned on.
□	Information	Platform Event Log	Sep 17, 2012 7:04:1...	ENET[CIM:ep1] DHCP-HSTN=IMM2-5cf3fc...
□	Information	Platform Event Log	Sep 17, 2012 7:04:1...	LAN: Ethernet[IBM:ep2] interface is now a...
□	Information	Platform Event Log	Sep 17, 2012 7:04:1...	ENET[CIM:ep2] IP-Cfg:HstName=IMM2-5cf...
□	Information	Platform Event Log	Sep 17, 2012 7:04:1...	LAN: Ethernet[IBM:ep1] interface is now a...
□	Information	Platform Event Log	Sep 17, 2012 7:04:1...	Management Controller SN# KQ5N05V Ne...

Figure 4-95 Hardware Log window

3. You can refresh the view, clear the entire log, or filter the view by using the Search function. You can also sort the view by clicking any of the column headings. From the Actions menu (Figure 4-96), you can export the log for problem determination. You can send the exported log information to IBM Support, if requested. The information is saved in CSV format.

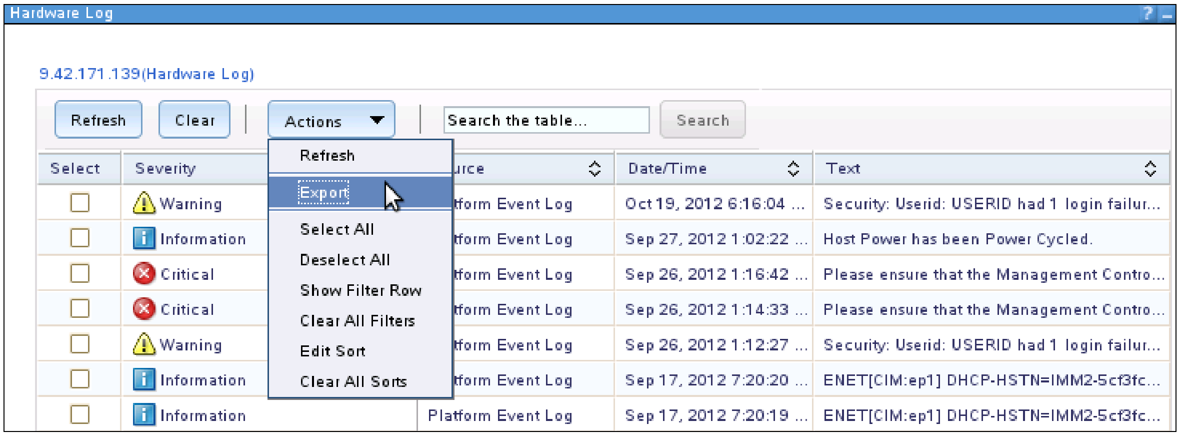

Select	Severity	urce	Date/Time	Text
□	⚠ Warning	tform Event Log	Oct 19, 2012 6:16:04 ...	Security: Userid: USERID had 1 login failur...
□	Information	tform Event Log	Sep 27, 2012 1:02:22 ...	Host Power has been Power Cycled.
□	Critical	tform Event Log	Sep 26, 2012 1:16:42 ...	Please ensure that the Management Contro...
□	Critical	tform Event Log	Sep 26, 2012 1:14:33 ...	Please ensure that the Management Contro...
□	⚠ Warning	tform Event Log	Sep 26, 2012 1:12:27 ...	Security: Userid: USERID had 1 login failur...
□	Information	tform Event Log	Sep 17, 2012 7:20:20 ...	ENET[CIM:ep1] DHCP-HSTN=IMM2-5cf3fc...
□	Information	Platform Event Log	Sep 17, 2012 7:20:19 ...	ENET[CIM:ep1] DHCP-HSTN=IMM2-5cf3fc...

Figure 4-96 Save hardware log

4.4.3 Event logs

You can work with the event logs that the Systems Director stores on the server. One predefined event action plan is available with the action “Log All Events”. This event action plan writes all events of the Systems Director server to a local event log.

Settings

Configure the settings for the event log in the left pane of the Systems Director web interface by clicking **Settings** → **Event Log Preferences** (Figure 4-97).

In the Settings window (Figure 4-97), you can select the time range that is reflected in the event log listings. Set a time range for hours, days, or weeks.

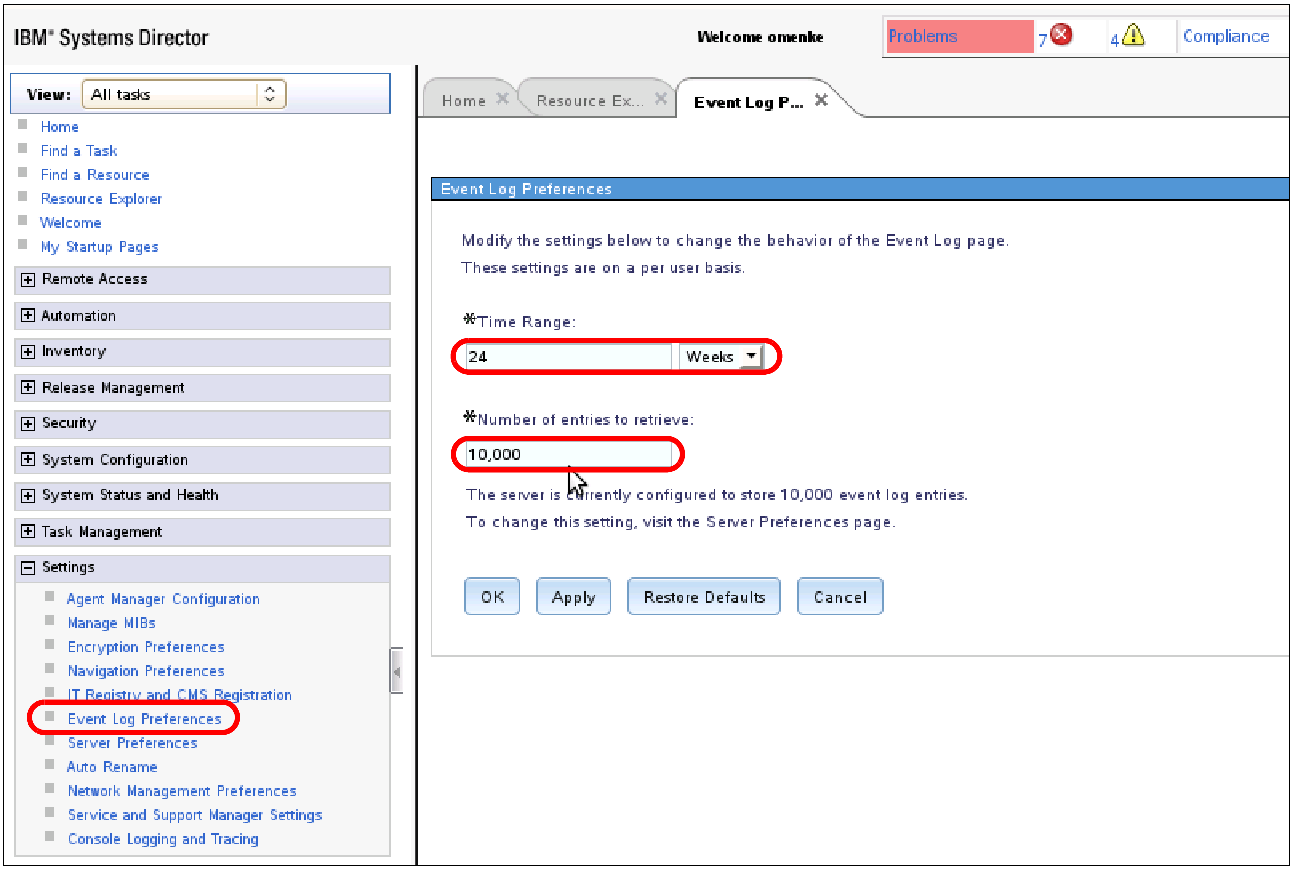

Event Log Preferences

Modify the settings below to change the behavior of the Event Log page.
These settings are on a per user basis.

*Time Range: 24 Weeks

*Number of entries to retrieve: 10,000

The server is currently configured to store 10,000 event log entries.
To change this setting, visit the Server Preferences page.

Figure 4-97 Settings for event log preferences

You can set the number of event log entries to retrieve. The maximum number for the server is 10,000 entries. If you set more than 10,000 entries, you see an error message (Figure 4-98).

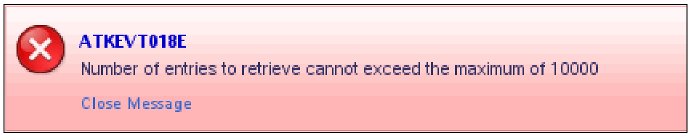

ATKEVT018E
Number of entries to retrieve cannot exceed the maximum of 10000

Figure 4-98 Error message when the number of entries exceeds the maximum number of 10,000

The default settings are 24 hours and 500 entries for the event log. You can go back to the default values by clicking **Restore Defaults**.

Launching the event log

You can access the Systems Director event log in a number of ways:

- From the left pane, click **System Status and Health** → **Event Log** (event log for **All Systems**).
- From the Resource Explorer menu, right-click a group (event log for the complete group) or a single system (event log for this system). Click **System Status and Health** → **Event Log** (Figure 4-99).
- Also, from the Resource Explorer, select the group or system and click **Actions** → **System Status and Health** → **Event Log** to access the event log.

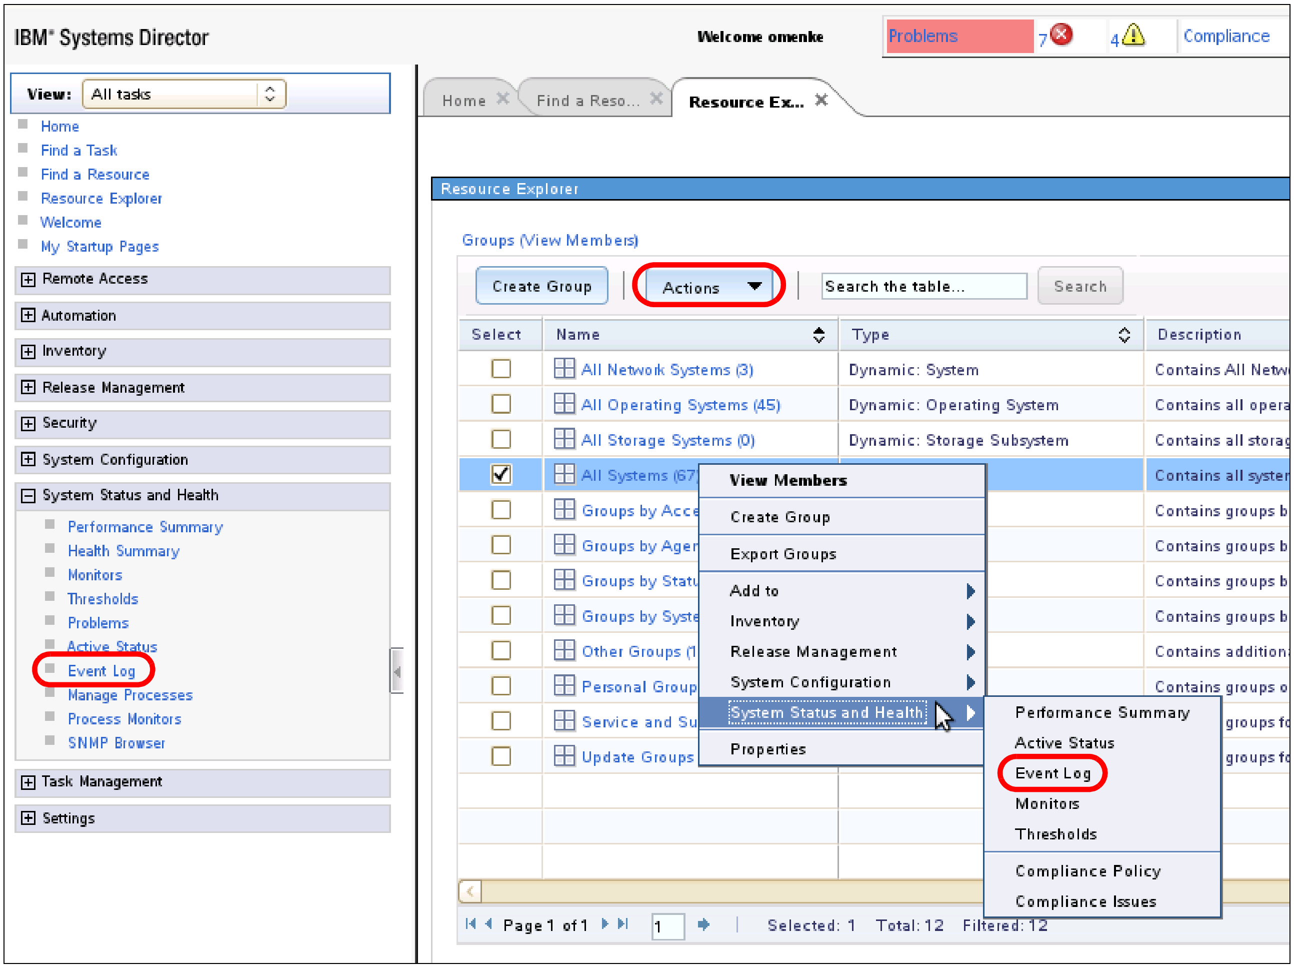

Select	Name	Type	Description
□	All Network Systems (3)	Dynamic: System	Contains All Netw
□	All Operating Systems (45)	Dynamic: Operating System	Contains all opera
□	All Storage Systems (0)	Dynamic: Storage Subsystem	Contains all storag
☑	All Systems (67		Contains all syster
□	Groups by Acce		Contains groups b
□	Groups by Ager		Contains groups b
□	Groups by Statu		Contains groups b
□	Groups by Syste		Contains groups b
□	Other Groups (1		Contains addition
□	Personal Group		Contains groups o
□	Service and Su		groups fo
□	Update Groups		groups fo

Page 1 of 1 | 1 | Selected: 1 Total: 12 Filtered: 12

Figure 4-99 Event log access

- Also, from the Resource Explorer, double-click a single system to see the properties of the system. Select the **Event Log** tab to access the event log for this system, as shown in Figure 4-100.

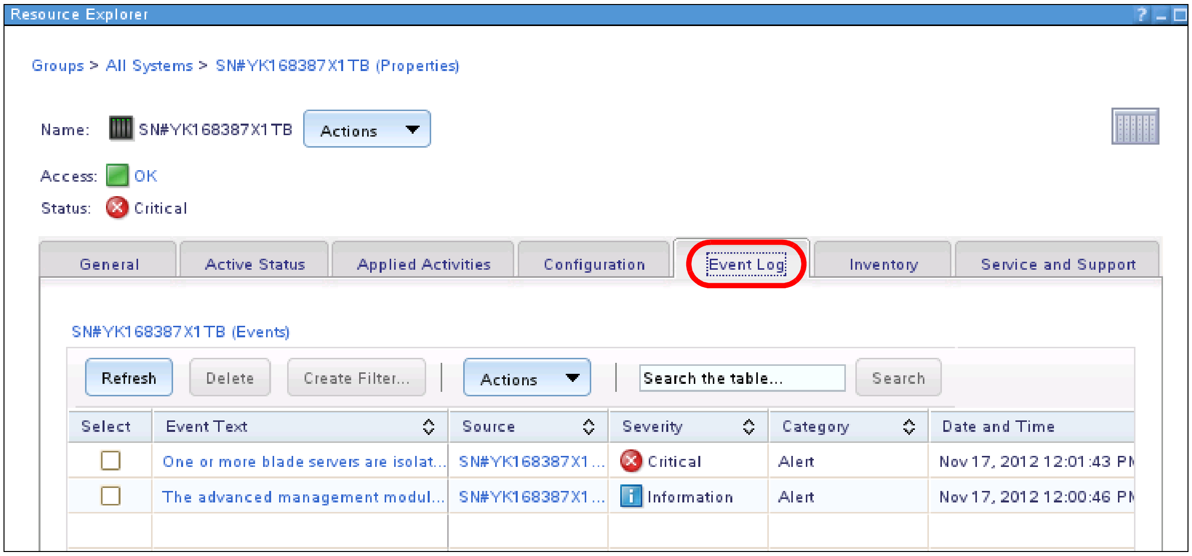

Name: SN#YK168387X1TB

Access: OK

Status: Critical

Select	Event Text	Source	Severity	Category	Date and Time
□	One or more blade servers are isolat...	SN#YK168387X1...	Critical	Alert	Nov 17, 2012 12:01:43 PM
□	The advanced management modul...	SN#YK168387X1...	Information	Alert	Nov 17, 2012 12:00:46 PM

Figure 4-100 Event Log access for a single system

Viewing the event log

When you elect to view the event logs of multiple systems, a new window opens that shows the event log (Figure 4-101 on page 155). All events are listed for the selected time range. Events are listed up to the maximum number that is set for the event log.

The event filter is at the top of the window (Figure 4-101). The use of the event filter is described in “Using event filters” on page 156. You can sort the events by date and time (default), by severity, or by source and category. Select the column and click the arrow in the top cell of the column. You can also use the search function to find specific events.

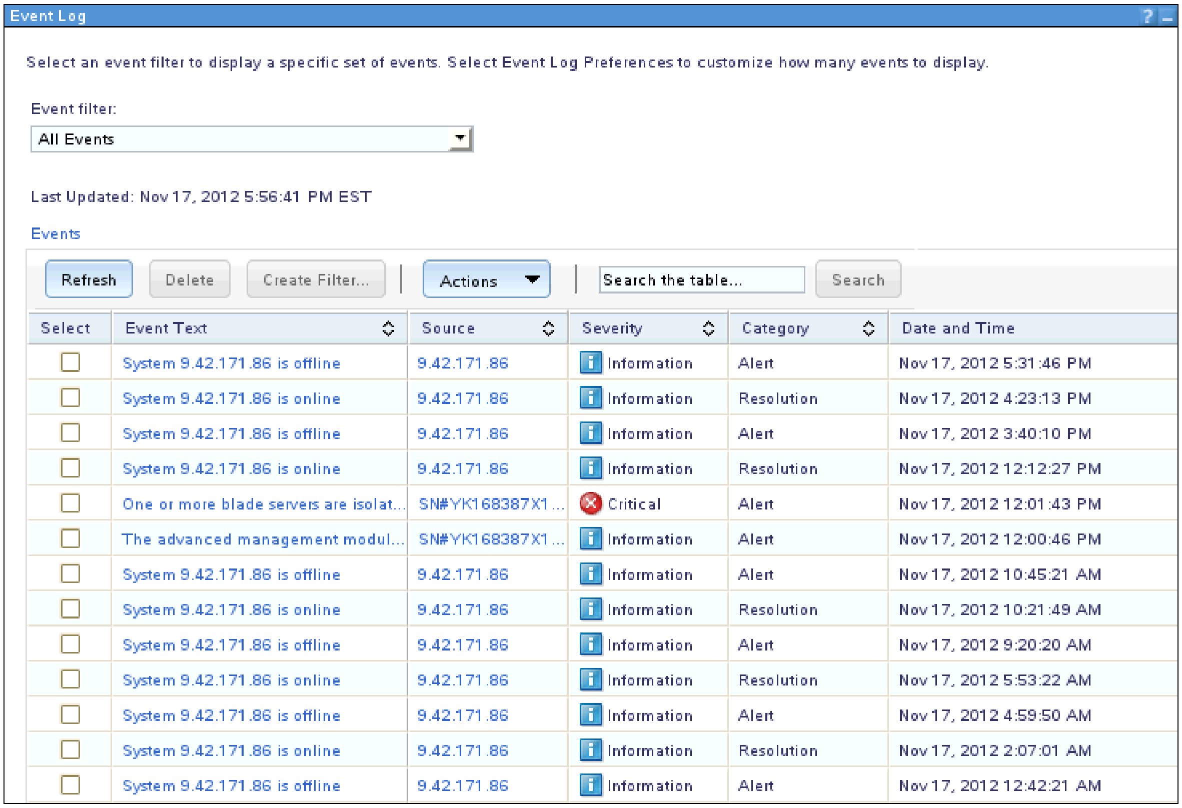

Select an event filter to display a specific set of events. Select Event Log Preferences to customize how many events to display.

Event filter: All Events

Last Updated: Nov 17, 2012 5:56:41 PM EST

Events

Select	Event Text	Source	Severity	Category	Date and Time
□	System 9.42.171.86 is offline	9.42.171.86	Information	Alert	Nov 17, 2012 5:31:46 PM
□	System 9.42.171.86 is online	9.42.171.86	Information	Resolution	Nov 17, 2012 4:23:13 PM
□	System 9.42.171.86 is offline	9.42.171.86	Information	Alert	Nov 17, 2012 3:40:10 PM
□	System 9.42.171.86 is online	9.42.171.86	Information	Resolution	Nov 17, 2012 12:12:27 PM
□	One or more blade servers are isolat...	SN#YK168387X1...	Critical	Alert	Nov 17, 2012 12:01:43 PM
□	The advanced management modul...	SN#YK168387X1...	Information	Alert	Nov 17, 2012 12:00:46 PM
□	System 9.42.171.86 is offline	9.42.171.86	Information	Alert	Nov 17, 2012 10:45:21 AM
□	System 9.42.171.86 is online	9.42.171.86	Information	Resolution	Nov 17, 2012 10:21:49 AM
□	System 9.42.171.86 is offline	9.42.171.86	Information	Alert	Nov 17, 2012 9:20:20 AM
□	System 9.42.171.86 is online	9.42.171.86	Information	Resolution	Nov 17, 2012 5:53:22 AM
□	System 9.42.171.86 is offline	9.42.171.86	Information	Alert	Nov 17, 2012 4:59:50 AM
□	System 9.42.171.86 is online	9.42.171.86	Information	Resolution	Nov 17, 2012 2:07:01 AM
□	System 9.42.171.86 is offline	9.42.171.86	Information	Alert	Nov 17, 2012 12:42:21 AM

Figure 4-101 Event Log window

Using event filters

You can use the event filter to select specific events. With filters, you can easily display only the event log entries that are important to you. With filters, you can easily export log entries for documentation.

The available filters are the same filters that are available in the event filter for the event automation plan. Any filters that you create for event automation plans are also visible and usable in this list. Figure 4-102 shows the filter list.

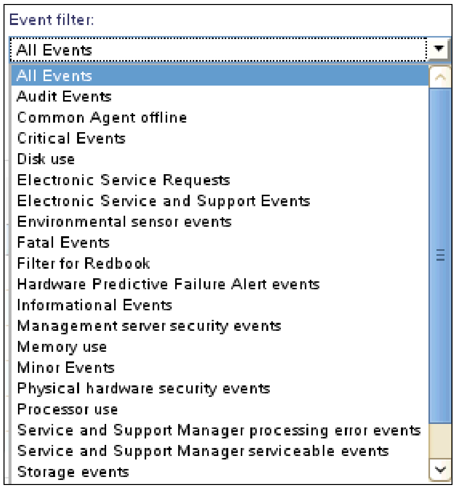

Event filter:

All Events

- All Events
- Audit Events
- Common Agent offline
- Critical Events
- Disk use
- Electronic Service Requests
- Electronic Service and Support Events
- Environmental sensor events
- Fatal Events
- Filter for Redbook
- Hardware Predictive Failure Alert events
- Informational Events
- Management server security events
- Memory use
- Minor Events
- Physical hardware security events
- Processor use
- Service and Support Manager processing error events
- Service and Support Manager serviceable events
- Storage events

Figure 4-102 Filter list

In our example, we select the critical events as a filter for the event log viewer. The result is shown in Figure 4-103 on page 157. You can see the critical events that are available in the event log of the Systems Director server.

You might see some events with `HIST:` in front, as indicated in Figure 4-103. These events are historical events. Historical events come from system logs from systems that are based on a time range before the actual Systems Director is active.

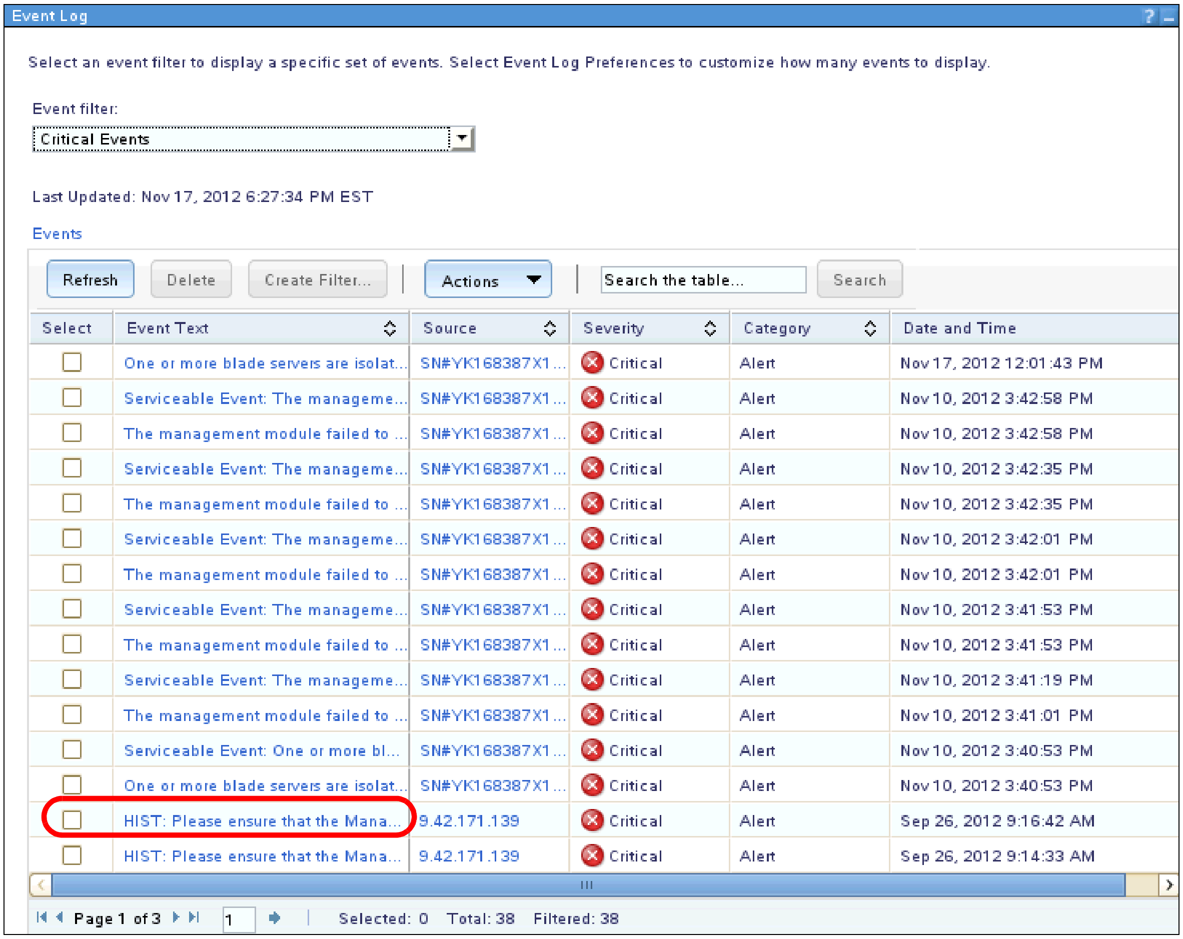

Event Log

Select an event filter to display a specific set of events. Select Event Log Preferences to customize how many events to display.

Event filter: Critical Events

Last Updated: Nov 17, 2012 6:27:34 PM EST

Events

Select	Event Text	Source	Severity	Category	Date and Time
□	One or more blade servers are isolat...	SN#YK168387X1...	Critical	Alert	Nov 17, 2012 12:01:43 PM
□	Serviceable Event: The manageme...	SN#YK168387X1...	Critical	Alert	Nov 10, 2012 3:42:58 PM
□	The management module failed to ...	SN#YK168387X1...	Critical	Alert	Nov 10, 2012 3:42:58 PM
□	Serviceable Event: The manageme...	SN#YK168387X1...	Critical	Alert	Nov 10, 2012 3:42:35 PM
□	The management module failed to ...	SN#YK168387X1...	Critical	Alert	Nov 10, 2012 3:42:35 PM
□	Serviceable Event: The manageme...	SN#YK168387X1...	Critical	Alert	Nov 10, 2012 3:42:01 PM
□	The management module failed to ...	SN#YK168387X1...	Critical	Alert	Nov 10, 2012 3:42:01 PM
□	Serviceable Event: The manageme...	SN#YK168387X1...	Critical	Alert	Nov 10, 2012 3:41:53 PM
□	The management module failed to ...	SN#YK168387X1...	Critical	Alert	Nov 10, 2012 3:41:53 PM
□	Serviceable Event: The manageme...	SN#YK168387X1...	Critical	Alert	Nov 10, 2012 3:41:19 PM
□	The management module failed to ...	SN#YK168387X1...	Critical	Alert	Nov 10, 2012 3:41:01 PM
□	Serviceable Event: One or more bl...	SN#YK168387X1...	Critical	Alert	Nov 10, 2012 3:40:53 PM
□	One or more blade servers are isolat...	SN#YK168387X1...	Critical	Alert	Nov 10, 2012 3:40:53 PM
□	HIST: Please ensure that the Mana...	9.42.171.139	Critical	Alert	Sep 26, 2012 9:16:42 AM
□	HIST: Please ensure that the Mana...	9.42.171.139	Critical	Alert	Sep 26, 2012 9:14:33 AM

Page 1 of 3 | 1 | Selected: 0 Total: 38 Filtered: 38

Figure 4-103 Critical event filter is used on the event log

Creating a filter by using an event from the event log

With Systems Director, you can create an event filter that is based on an existing event. This event filter can help you identify all of the events of the same type (for export, as an example).

To create a filter, go to an event and use the following steps:

1. From an existing event, right-click the event and in the pull-down menu, select **Create Filter**. Or, select the event by clicking the adjacent check box and click **Action** → **Create Filter**. See Figure 4-104.

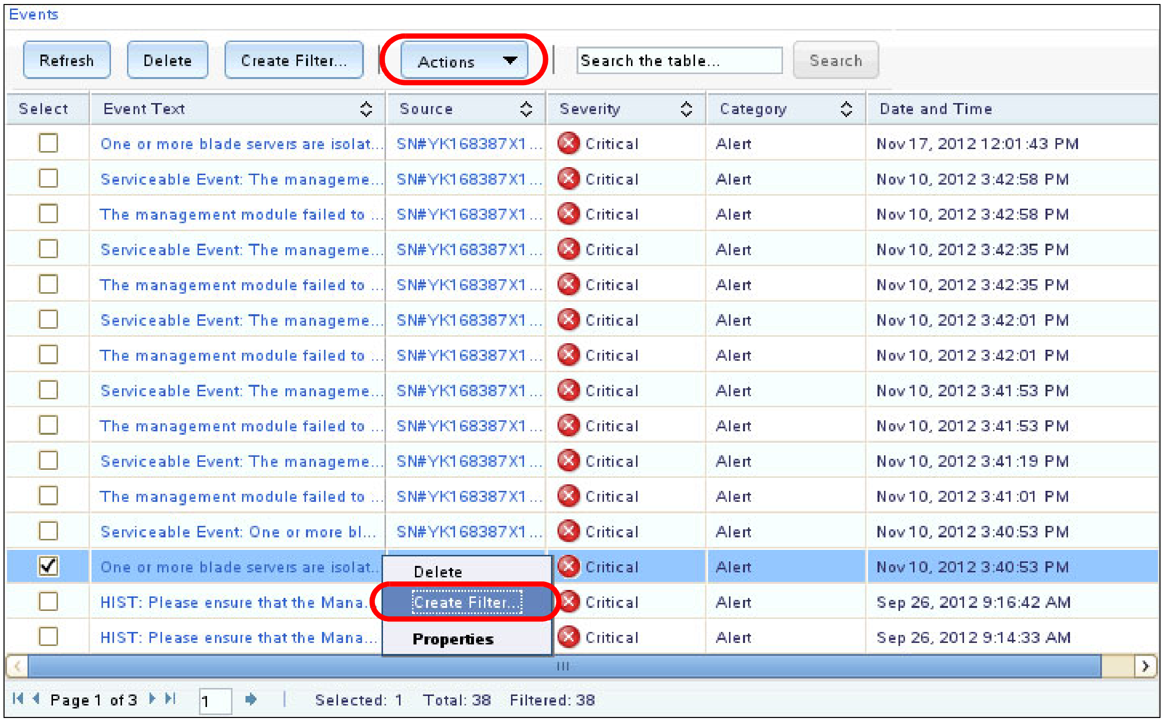

Select	Event Text	Source	Severity	Category	Date and Time
□	One or more blade servers are isolat...	SN#YK168387X1...	Critical	Alert	Nov 17, 2012 12:01:43 PM
□	Serviceable Event: The manageme...	SN#YK168387X1...	Critical	Alert	Nov 10, 2012 3:42:58 PM
□	The management module failed to ...	SN#YK168387X1...	Critical	Alert	Nov 10, 2012 3:42:58 PM
□	Serviceable Event: The manageme...	SN#YK168387X1...	Critical	Alert	Nov 10, 2012 3:42:35 PM
□	The management module failed to ...	SN#YK168387X1...	Critical	Alert	Nov 10, 2012 3:42:35 PM
□	Serviceable Event: The manageme...	SN#YK168387X1...	Critical	Alert	Nov 10, 2012 3:42:01 PM
□	The management module failed to ...	SN#YK168387X1...	Critical	Alert	Nov 10, 2012 3:42:01 PM
□	Serviceable Event: The manageme...	SN#YK168387X1...	Critical	Alert	Nov 10, 2012 3:41:53 PM
□	The management module failed to ...	SN#YK168387X1...	Critical	Alert	Nov 10, 2012 3:41:53 PM
□	Serviceable Event: The manageme...	SN#YK168387X1...	Critical	Alert	Nov 10, 2012 3:41:19 PM
□	The management module failed to ...	SN#YK168387X1...	Critical	Alert	Nov 10, 2012 3:41:01 PM
□	Serviceable Event: One or more bl...	SN#YK168387X1...	Critical	Alert	Nov 10, 2012 3:40:53 PM
☑	One or more blade servers are isolat...		Critical	Alert	Nov 10, 2012 3:40:53 PM
□	HIST: Please ensure that the Mana...		Critical	Alert	Sep 26, 2012 9:16:42 AM
□	HIST: Please ensure that the Mana...		Critical	Alert	Sep 26, 2012 9:14:33 AM

Page 1 of 3 1 Selected: 1 Total: 38 Filtered: 38

Figure 4-104 Select event to create a filter

2. Enter a name and a short description for the filter (Figure 4-105). Click **OK** to create the event filter that is based on the selected event. The definition for this event is at the bottom of the window.

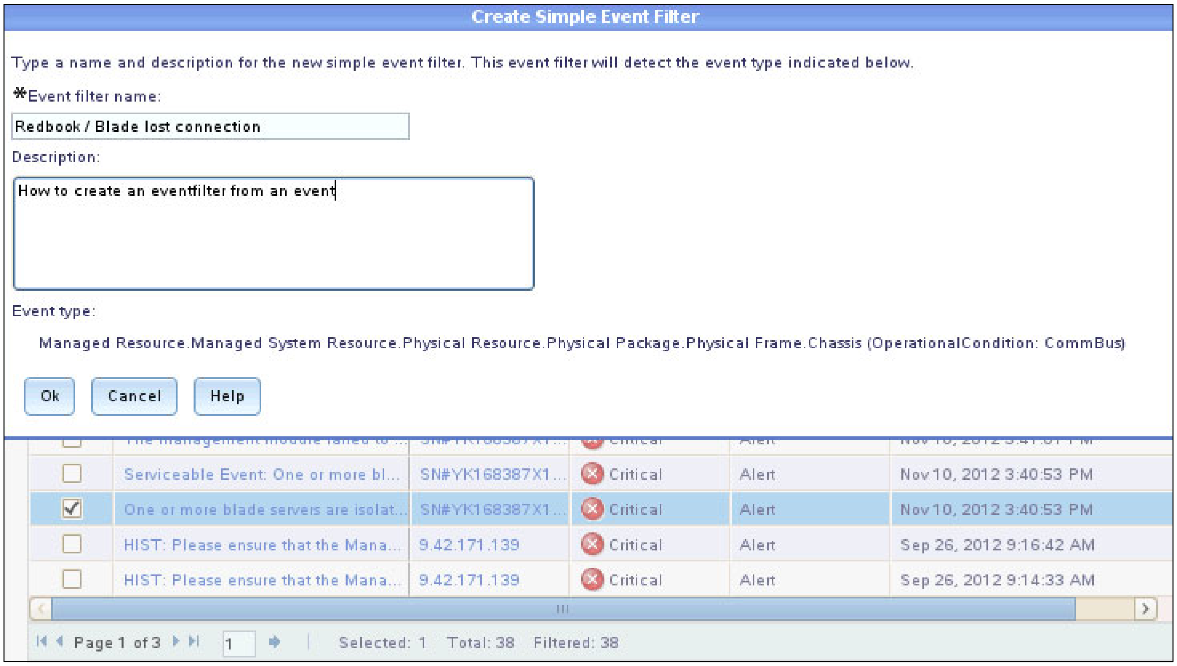

Create Simple Event Filter

Type a name and description for the new simple event filter. This event filter will detect the event type indicated below.

*Event filter name:

Redbook / Blade lost connection

Description:

How to create an eventfilter from an event

Event type:

Managed Resource.Managed System Resource.Physical Resource.Physical Package.Physical Frame.Chassis (OperationalCondition: CommBus)

□	Serviceable Event: One or more bl...	SN#YK168387X1...	Critical	Alert	Nov 10, 2012 3:40:53 PM
☑	One or more blade servers are isolat...	SN#YK168387X1...	Critical	Alert	Nov 10, 2012 3:40:53 PM
□	HIST: Please ensure that the Mana...	9.42.171.139	Critical	Alert	Sep 26, 2012 9:16:42 AM
□	HIST: Please ensure that the Mana...	9.42.171.139	Critical	Alert	Sep 26, 2012 9:14:33 AM

Page 1 of 3 | 1 | Selected: 1 Total: 38 Filtered: 38

Figure 4-105 Creating a filter from an event

3. A message appears that the event filter is created successfully (Figure 4-106).

   If there are problems, you see an error message instead. Fix the problem and create the event filter after you fix the problem.

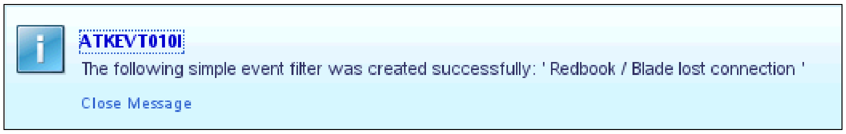

ATKEVT010I

The following simple event filter was created successfully: ' Redbook / Blade lost connection '

Figure 4-106 Creation of filter successful

Hardware event filters

If you are planning to use IBM Systems Director to perform hardware monitoring and you want to integrate it with an event management tool, there is a list below of recommended filters to configure your event filter from your automation plan. After it is configured, you can add or remove filters that are based on your needs:

- Managed Resource.Managed System Resource.Physical Resource
- Managed Resource.Managed System Resource.Logical Resource:
  - Service Access Point
  - System
  - Service
  - Logical Device:
    - Processor
    - Logical Module
    - Logical Port
    - Unknown
    - Sensor

- Storage Extent
- Battery
- Controller
- Media Access Device.OpticalDevice
- Media Access Device.Physical Volume
- Media Access Device.Disk Drive (RAID Subsystem Drive Bad Block: FullReassigned)
- Media Access Device.Disk Drive (RAID Subsystem Drive Synchronization: DetectedFailedCompletedStopped)
- Media Access Device.Disk Drive (Disk Drive Mounting Events: Mount as Global SpareUnmount Global SpareMount as Local SpareUnmount Local Spare)
- Media Access Device.Disk Drive (RAID Subsystem Drive Clear: DetectedFailedCompletedStoppedStartedProgressAborted)
- Media Access Device.Disk Drive (RAID Subsystem Drive Verify: DetectedFailedCompleted)
- Media Access Device.Disk Drive (Life Cycle: RemovedAdded)
- Media Access Device.Disk Drive (RAID Subsystem Drive Format: StartedCompleted)
- Media Access Device.Disk Drive (Operation: ActivatedDeactivated)
- Media Access Device.Disk Drive (RAID Subsystem Global Hot Spare: CreatedDeactivatedNot CoveringCommissioned)
- Media Access Device.Disk Drive (RAID Subsystem Hot-Spare Drive: AddedRemovedFailed)
- Media Access Device.Disk Drive (RAID Subsystem Drive Patrol Read: ProgressError)
- Media Access Device.Disk Drive (Operational Condition: CRC ErrorParity ErrorConfiguration ErrorFailedPredictive Failure Analysis (PFA))
- Media Access Device.Disk Drive (RAID Subsystem Life Cycle: RemovedAddedError)
- Media Access Device.Disk Drive (Power: OnOff)
- Media Access Device.Disk Drive (RAID Subsystem Configuration: Unsupported DriveDrive Too SmallDrives Missing)
- Media Access Device.Disk Drive (RAID Subsystem Microcode Update: DetectedFailedCompletedStopped)
- Media Access Device.Disk Drive (RAID Subsystem Drive Redundant Path: BrokenRestoredUnabled Accessed)
- Media Access Device.Disk Drive (RAID Subsystem Copyback: DetectedFailedCompletedStoppedProgress)
- Media Access Device.Disk Drive (RAID Subsystem Operational Condition: Predictive Failure Analysis (PFA)StatusInitialization CompletedInitialization StartedInitialization FailedFailed)
- Media Access Device.Disk Drive (RAID Subsystem Drive Security: ActivatedDeactivatedFailed)
- Media Access Device.Disk Drive (RAID Subsystem Dedicated Hot Spare: CreatedDeactivatedNot UsefulImported)
- Media Access Device.Disk Drive (RAID Subsystem Drive Security Key: CreatedBacked UpVerifiedChangedFailedInvalidDestroyed)
- Media Access Device.Disk Drive (RAID Subsystem SAS Port: Lost LinkRestored LinkError)
- Media Access Device.Disk Drive (Updates: InstallationTask Failed)
- Fan (Operational Condition: Power Predictive Error)
- Fan (Operational Condition: Power Unrecoverable Error)
- Fan (Operational Condition: Predictive Failure Analysis (PFA))
- Fan (Operational Condition: Configuration Error)
- Fan (Operational Condition: Returned To OK)

- Fan (Operational Condition: Degraded)
- Fan (Operational Condition: Speed)
- Fan (Operational Condition: Failed)
- Fan (Operational Condition: No Error or Informational)
- Fan (Operational Condition: Non-Serviceable)
- Fan (Operational Condition: Power Diagnostics Test)
- Fan (Operational Condition: Redundancy)
- Power Supply

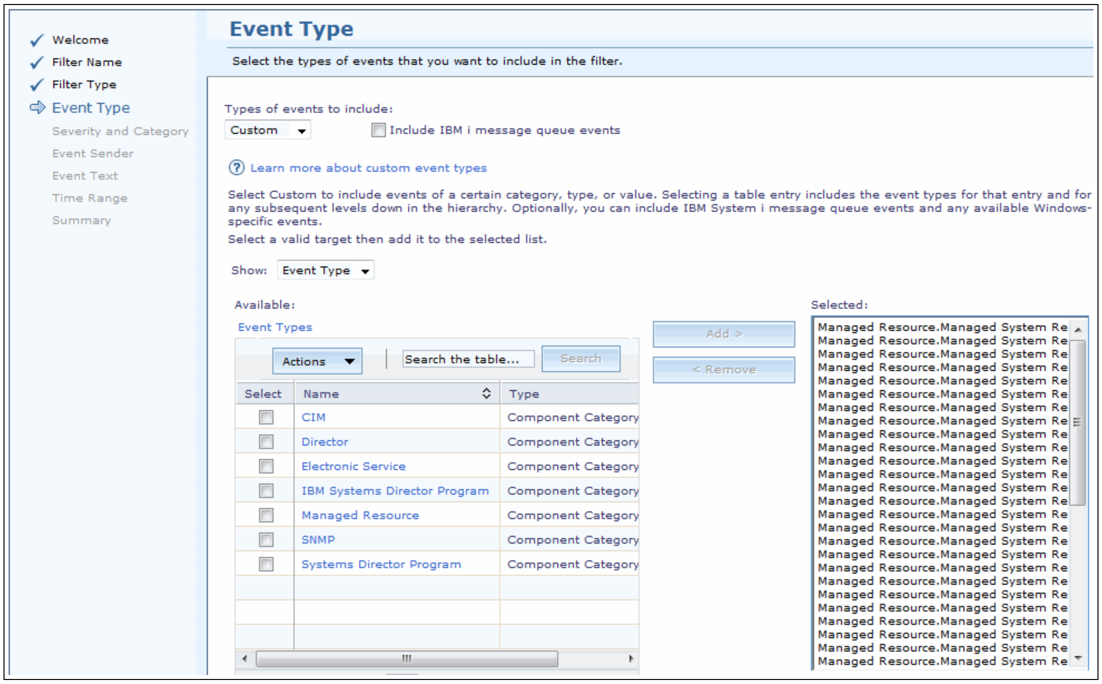

Event Type

Select the types of events that you want to include in the filter.

Types of events to include:

Custom ▾ □ Include IBM i message queue events

Learn more about custom event types

Select Custom to include events of a certain category, type, or value. Selecting a table entry includes the event types for that entry and for any subsequent levels down in the hierarchy. Optionally, you can include IBM System i message queue events and any available Windows-specific events.

Select a valid target then add it to the selected list.

Show: Event Type ▾

Available:

Event Types

Select	Name	Type
□	CIM	Component Category
□	Director	Component Category
□	Electronic Service	Component Category
□	IBM Systems Director Program	Component Category
□	Managed Resource	Component Category
□	SNMP	Component Category
□	Systems Director Program	Component Category

Selected:

Managed Resource.Managed System Re
Managed Resource.Managed System Re
Managed Resource.Managed System Re
Managed Resource.Managed System Re
Managed Resource.Managed System Re
Managed Resource.Managed System Re
Managed Resource.Managed System Re
Managed Resource.Managed System Re
Managed Resource.Managed System Re
Managed Resource.Managed System Re
Managed Resource.Managed System Re
Managed Resource.Managed System Re
Managed Resource.Managed System Re
Managed Resource.Managed System Re
Managed Resource.Managed System Re
Managed Resource.Managed System Re
Managed Resource.Managed System Re
Managed Resource.Managed System Re
Managed Resource.Managed System Re
Managed Resource.Managed System Re
Managed Resource.Managed System Re
Managed Resource.Managed System Re
Managed Resource.Managed System Re
Managed Resource.Managed System Re
Managed Resource.Managed System Re
Managed Resource.Managed System Re

Figure 4-107 Create a customized event filter

> **Note:** Multipath monitoring (MPIO) should be performed from the switch devices, not from IBM Systems Director.

Command-line tools

The following command-line commands are available to work with the Systems Director event log:

smcli evtlog	Set parameter for event log
smcli lsevtlog	List event log
smcli rmevtlog	Delete entries or complete event log

Example 4-7 shows the following information:

- The actual size with the **-s** option. Our example shows 791 entries in the event log.
- The setting for the maximum value with the **smcli evtlog -m** command.
- The command to set the maximum to a new value, in our example, 9500 with the **-M 9500** option.

Example 4-7 Example for the smcli evtlog

```
SLES11:/opt/ibm/director/bin #./smcli evtlog -s
791
SLES11:/opt/ibm/director/bin #./smcli evtlog -m
```

```
10000
SLES11:/opt/ibm/director/bin #./smcli evtlog -M 9500

SLES11:/opt/ibm/director/bin #./smcli evtlog -m
9500
SLES11:/opt/ibm/director/bin #
```

You can use the **smcli lsevtlog** command to read the event log. The following options are useful:

-e "*EventFilter_name*"	Show events only for this event filter
-s	Present a summary view of the event
-T	Filter for a time range (value in hours)
-o	Display the unique IDs that are associated with the event-log entries in addition to other information
-t	Filter on a system type, such as operating system

In Example 4-8, we use the **-e "CriticalEvents" -T 192 -o** parameters. These parameters list all critical events in the last 192 hours and the object identifiers (OIDs) for the events. Example 4-8 shows one critical event in the specified time range. The OID for this event (`0x2ab`) is listed.

Example 4-8 smcli lsevtlog command

```
SLES11:/opt/ibm/director/bin #./smcli lsevtlog -e "Critical Events" -T 192 -o
11/17/12 12:01 PM, One or more blade servers are isolated from the management
bus., SN#YK168387X1TB (0x175b), Managed Resource.Managed System Resource.Physical
Resource.Physical Package.Physical Frame.Chassis (OperationalCondition: CommBus),
Critical-Alert, 0x2ab
SLES11:/opt/ibm/director/bin #
```

Example 4-9 shows how to delete an event from the event log with the **smcli rmevtlog** command.

Use this command if many events of the same type are in the log and you want to clean up the log. Or, you generated test log entries that you want to delete from the log.

You can use the **-a** option to delete the complete log or use the **-e %event_oid%** option to delete specific events. In our example, we use the **-e** option to delete the event that is listed in Example 4-8. After we remove the event from the log with the `0x2ab` OID, we run the **smcli lsevtlog** command to list the event, as shown in Example 4-9. The example shows that this event does not exist anymore.

Example 4-9 Example for the smcli rmevtlog command

```
SLES11:/opt/ibm/director/bin #./smcli rmevtlog -e 0x2ab
SLES11:/opt/ibm/director/bin #
SLES11:/opt/ibm/director/bin #./smcli lsevtlog -e "Critical Events" -T 192 -o
SLES11:/opt/ibm/director/bin #
```

For detailed information about the commands in our example, see the Systems Director Information Center:

http://pic.dhe.ibm.com/infocenter/director/pubs/index.jsp?topic=%2Fcom.ibm.director.cli.helps.doc%2Ffqm0_r_cli_event_log_and_history_cmds.html

4.4.4 Automation Manager

Use Automation Manager to create and use event automation plans.

The event automation plan consists of two parts:

- Event filter
- Event action

You can create and assign the event automation plan, event filter, and event action through the graphical user interface (GUI) or through the command line.

The following topics are described:

- “Creating an event automation plan” on page 163
- “Creating an event filter” on page 174
- “Creating an event action” on page 184
- “Using the CLI for event automation plans” on page 190

Creating an event automation plan

To view and create the event automation plan, go to the Systems Director web interface. Select either **Automation** → **Event Automation Plans** on the left pane or the **Event Automation Plans** link under the Automation Manager, as shown in Figure 4-108.

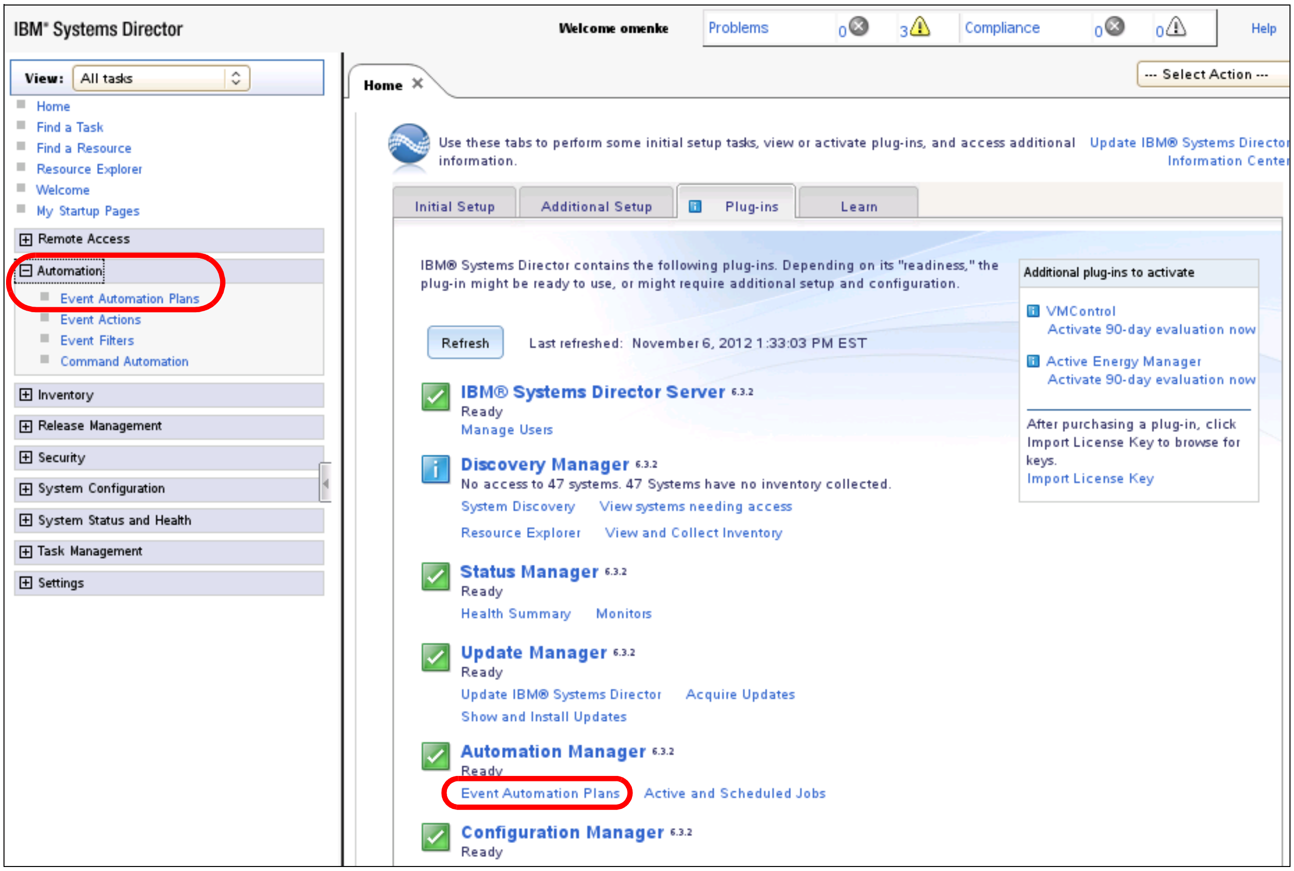

Use these tabs to perform some initial setup tasks, view or activate plug-ins, and access additional information.

IBM® Systems Director contains the following plug-ins. Depending on its "readiness," the plug-in might be ready to use, or might require additional setup and configuration.

Last refreshed: November 6, 2012 1:33:03 PM EST

After purchasing a plug-in, click Import License Key to browse for keys.

No access to 47 systems. 47 Systems have no inventory collected.

Figure 4-108 Launching Automation Manager

You can see the existing event automation plans, the targets to which they are assigned, the status of the plan, the time range that is defined for the plan, and a description, if available.

You can edit an existing event automation plan, create an event automation plan that is based on an existing event automation plan, or create an event automation plan.

In our example, we create an event automation plan that is named `book-EAP`:

1. In the Event Automation Plans window, click **Create** (Figure 4-109).

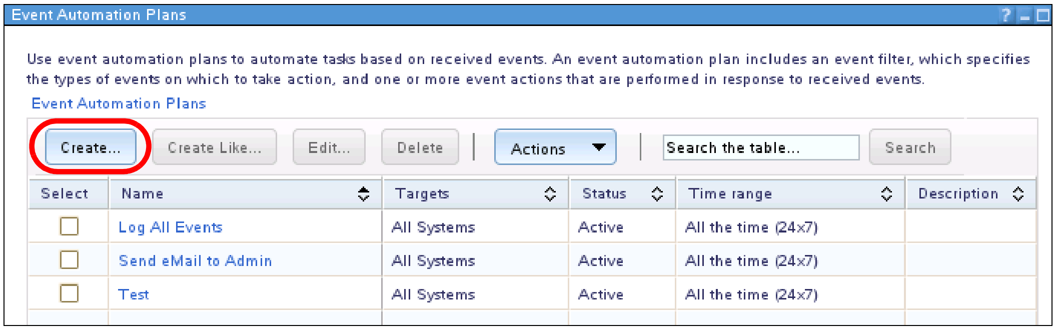

Use event automation plans to automate tasks based on received events. An event automation plan includes an event filter, which specifies the types of events on which to take action, and one or more event actions that are performed in response to received events.

Select	Name	Targets	Status	Time range	Description
□	Log All Events	All Systems	Active	All the time (24x7)	
□	Send eMail to Admin	All Systems	Active	All the time (24x7)	
□	Test	All Systems	Active	All the time (24x7)	

Figure 4-109 Event Automation Plans window

2. The Create Event Automation Plan wizard opens at the Welcome window (Figure 4-110). Click **Next** to continue.

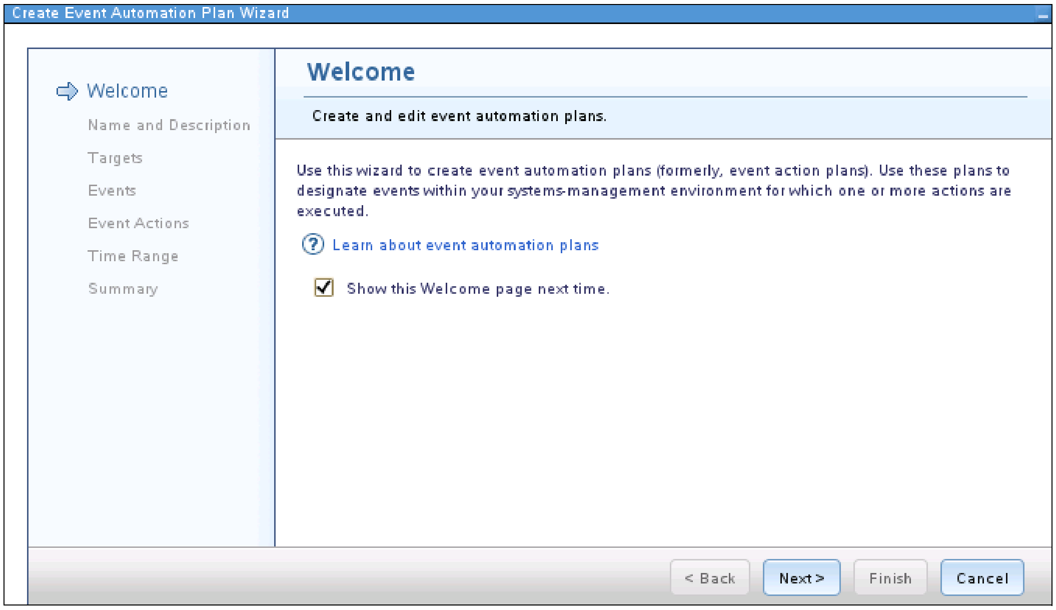

Create and edit event automation plans.

Use this wizard to create event automation plans (formerly, event action plans). Use these plans to designate events within your systems-management environment for which one or more actions are executed.

☑ Show this Welcome page next time.

Figure 4-110 Create Event Automation Plan wizard: Welcome panel

3. In the Name and Description step (Figure 4-111), give your event automation plan a name and suitable description. In our example, we name the event automation plan that we created, `book-EAP`. Click **Next** to continue.

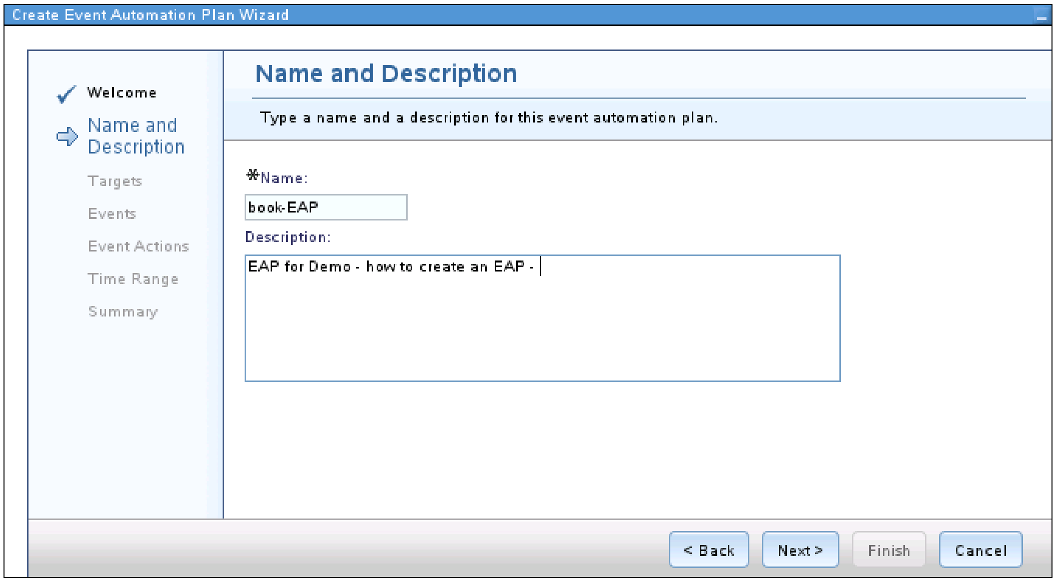

Create Event Automation Plan Wizard

Name and Description

Type a name and a description for this event automation plan.

*Name: book-EAP

Description: EAP for Demo - how to create an EAP -

Figure 4-111 Create Event Automation Plan wizard: Name and Description

4. In the next window (Figure 4-112), define the targets for which this event automation plan works. You can select groups or an individual system as the target for this event automation plan.

   Selecting groups can be helpful when you have different administrative or management groups. In our example, we select `All Systems`. To select systems or a group, select the system or group in the left column and then click **Add**. When finished, click **Next**.

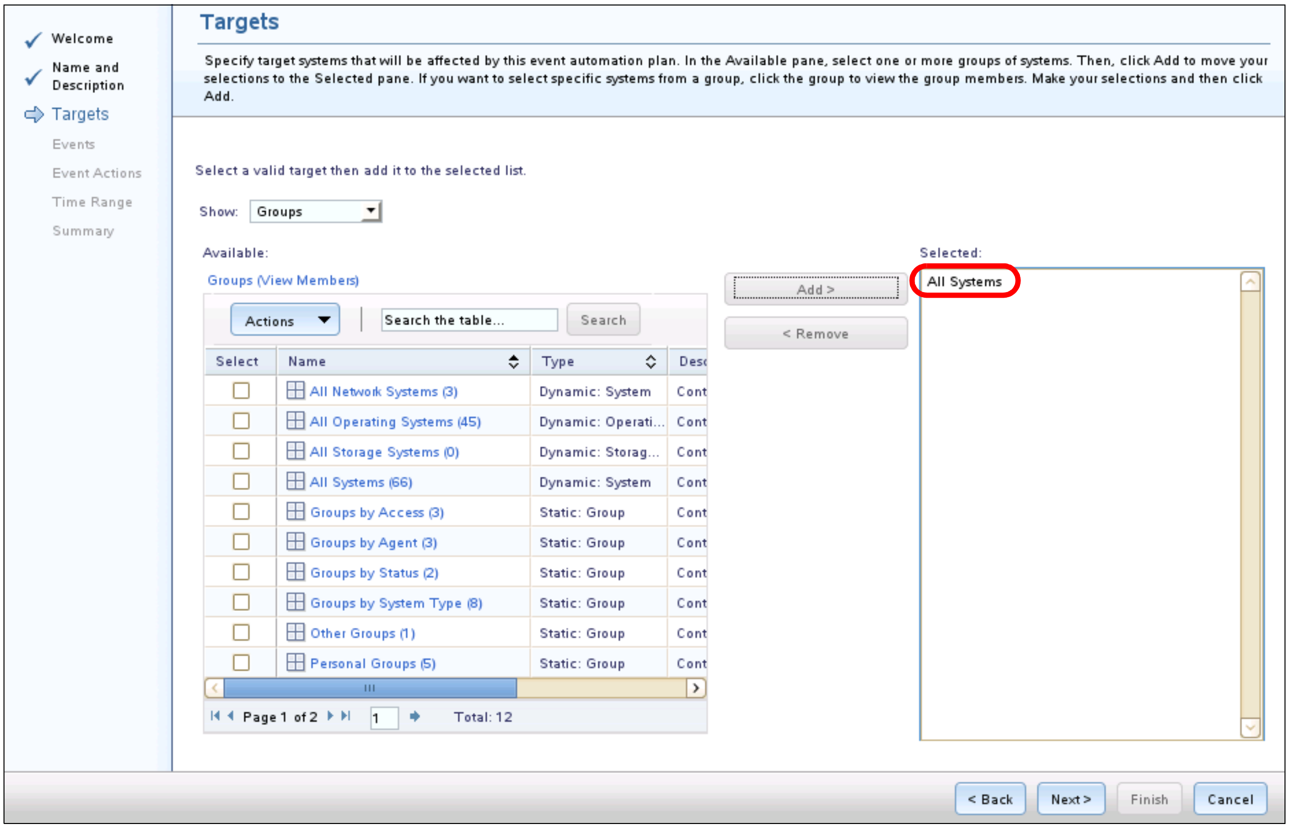

Targets

Specify target systems that will be affected by this event automation plan. In the Available pane, select one or more groups of systems. Then, click Add to move your selections to the Selected pane. If you want to select specific systems from a group, click the group to view the group members. Make your selections and then click Add.

Select a valid target then add it to the selected list.

Show: Groups

Available:

Groups (View Members)

Select	Name	Type	Desc
□	All Network Systems (3)	Dynamic: System	Cont
□	All Operating Systems (45)	Dynamic: Operati...	Cont
□	All Storage Systems (0)	Dynamic: Storag...	Cont
□	All Systems (66)	Dynamic: System	Cont
□	Groups by Access (3)	Static: Group	Cont
□	Groups by Agent (3)	Static: Group	Cont
□	Groups by Status (2)	Static: Group	Cont
□	Groups by System Type (8)	Static: Group	Cont
□	Other Groups (1)	Static: Group	Cont
□	Personal Groups (5)	Static: Group	Cont

Page 1 of 2 1 Total: 12

Selected:

All Systems

Figure 4-112 Create Event Automation Plan wizard: Target selection

5. In the Events window (Figure 4-113), select the filter for events.

   Various filter types are available:

   – *Common* event filters are predefined filters that monitor for common functions, such as hardware events. Examples are fan failures or processor usage (Figure 4-113). The common event filters are predefined and cannot be changed or enhanced. If you need more complex criteria, select *Advanced Event Filters*. You can select some of the common event filters to use in an event automation plan.

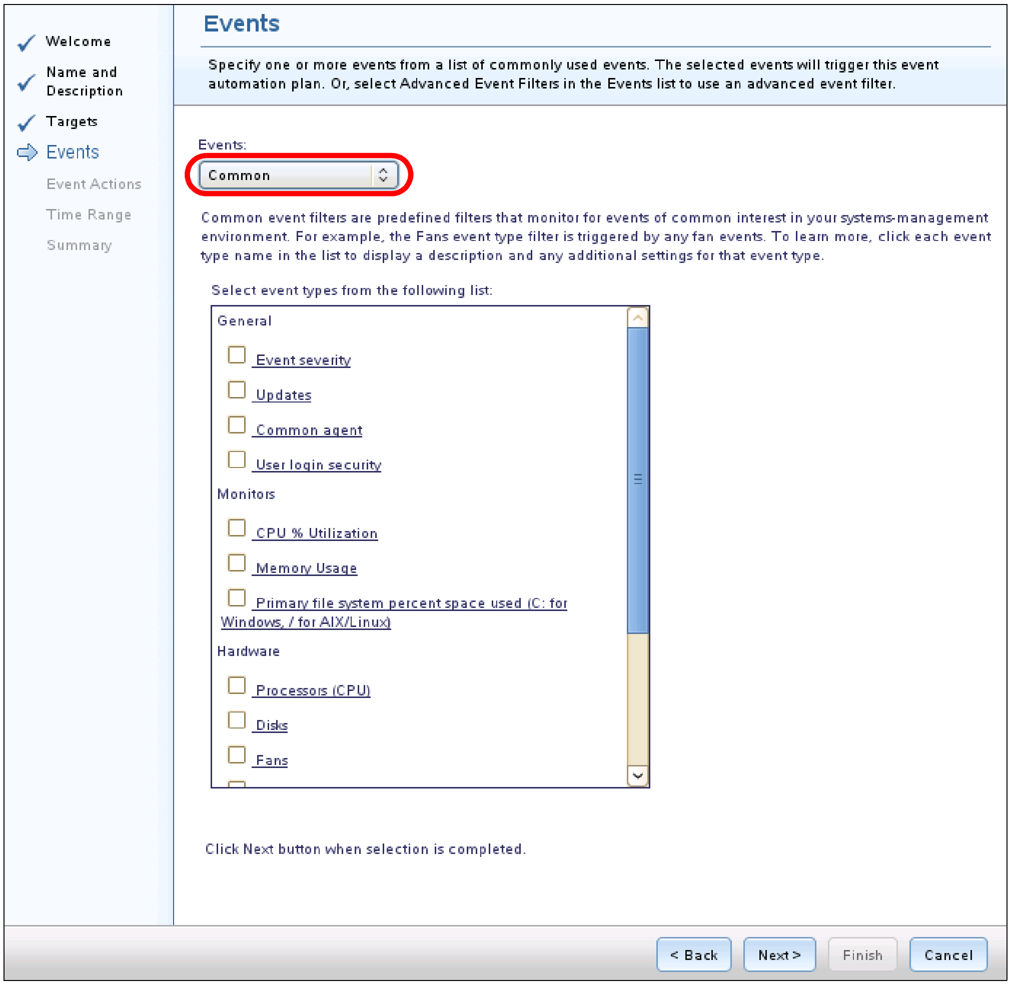

Events

Specify one or more events from a list of commonly used events. The selected events will trigger this event automation plan. Or, select Advanced Event Filters in the Events list to use an advanced event filter.

Events: Common

Common event filters are predefined filters that monitor for events of common interest in your systems-management environment. For example, the Fans event type filter is triggered by any fan events. To learn more, click each event type name in the list to display a description and any additional settings for that event type.

Select event types from the following list:

General

□ Event severity

□ Updates

□ Common agent

□ User login security

Monitors

□ CPU % Utilization

□ Memory Usage

□ Primary file system percent space used (C: for Windows, / for AIX/Linux)

Hardware

□ Processors (CPU)

□ Disks

□ Fans

Click Next button when selection is completed.

Figure 4-113 Create an event automation plan: Common filter

– Advanced event filters are used for monitoring specific events, single system events, or events that are based on severity (Figure 4-114). Predefined common event filters are available, but you can edit or enhance the advanced event filter. Only one advanced event filter can be selected for an event automation plan.

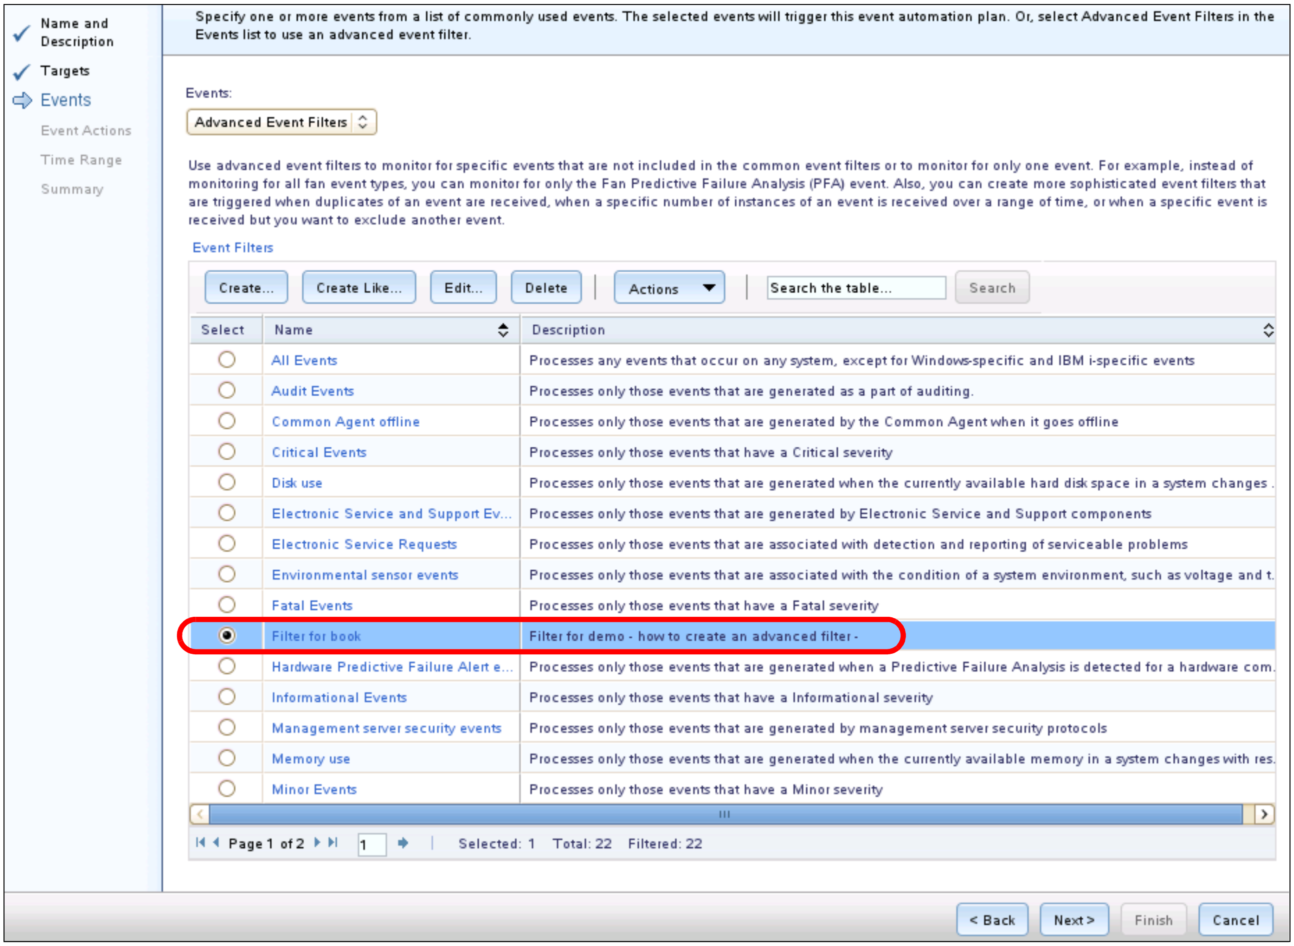

Specify one or more events from a list of commonly used events. The selected events will trigger this event automation plan. Or, select Advanced Event Filters in the Events list to use an advanced event filter.

Events: Advanced Event Filters

Use advanced event filters to monitor for specific events that are not included in the common event filters or to monitor for only one event. For example, instead of monitoring for all fan event types, you can monitor for only the Fan Predictive Failure Analysis (PFA) event. Also, you can create more sophisticated event filters that are triggered when duplicates of an event are received, when a specific number of instances of an event is received over a range of time, or when a specific event is received but you want to exclude another event.

Event Filters

Create... | Create Like... | Edit... | Delete | Actions | Search the table... | Search

Select	Name	Description
○	All Events	Processes any events that occur on any system, except for Windows-specific and IBM i-specific events
○	Audit Events	Processes only those events that are generated as a part of auditing.
○	Common Agent offline	Processes only those events that are generated by the Common Agent when it goes offline
○	Critical Events	Processes only those events that have a Critical severity
○	Disk use	Processes only those events that are generated when the currently available hard disk space in a system changes .
○	Electronic Service and Support Ev...	Processes only those events that are generated by Electronic Service and Support components
○	Electronic Service Requests	Processes only those events that are associated with detection and reporting of serviceable problems
○	Environmental sensor events	Processes only those events that are associated with the condition of a system environment, such as voltage and t.
○	Fatal Events	Processes only those events that have a Fatal severity
●	Filter for book	Filter for demo - how to create an advanced filter -
○	Hardware Predictive Failure Alert e...	Processes only those events that are generated when a Predictive Failure Analysis is detected for a hardware com.
○	Informational Events	Processes only those events that have a Informational severity
○	Management server security events	Processes only those events that are generated by management server security protocols
○	Memory use	Processes only those events that are generated when the currently available memory in a system changes with res.
○	Minor Events	Processes only those events that have a Minor severity

Page 1 of 2   1   Selected: 1   Total: 22   Filtered: 22

Figure 4-114 Create Event Automation Plan wizard: Advanced event filter

When you select the advanced event filter, you can use predefined filters or create your own. How to create your own filter is described in “Creating an event filter” on page 174.

In our example, we created the event filter named `Filter for book`. When you select the filter that you want to use, click **Next**.

6. In Figure 4-115, specify the event action for the event automation plan to perform on the target systems when the filter criteria is met. By default, Systems Director comes with only one predefined event action, which is named `Add to the event log`. Select an existing event action or create an event action. Creating an event action is described in “Creating an event action” on page 184.

   In our example, we use the created event action named `Mail to book` (Figure 4-115). You can select more than one action for an event automation plan. All of the selected actions run if the event that is monitored and filtered occurs.

   After you select the action that you want to use, click **Next**.

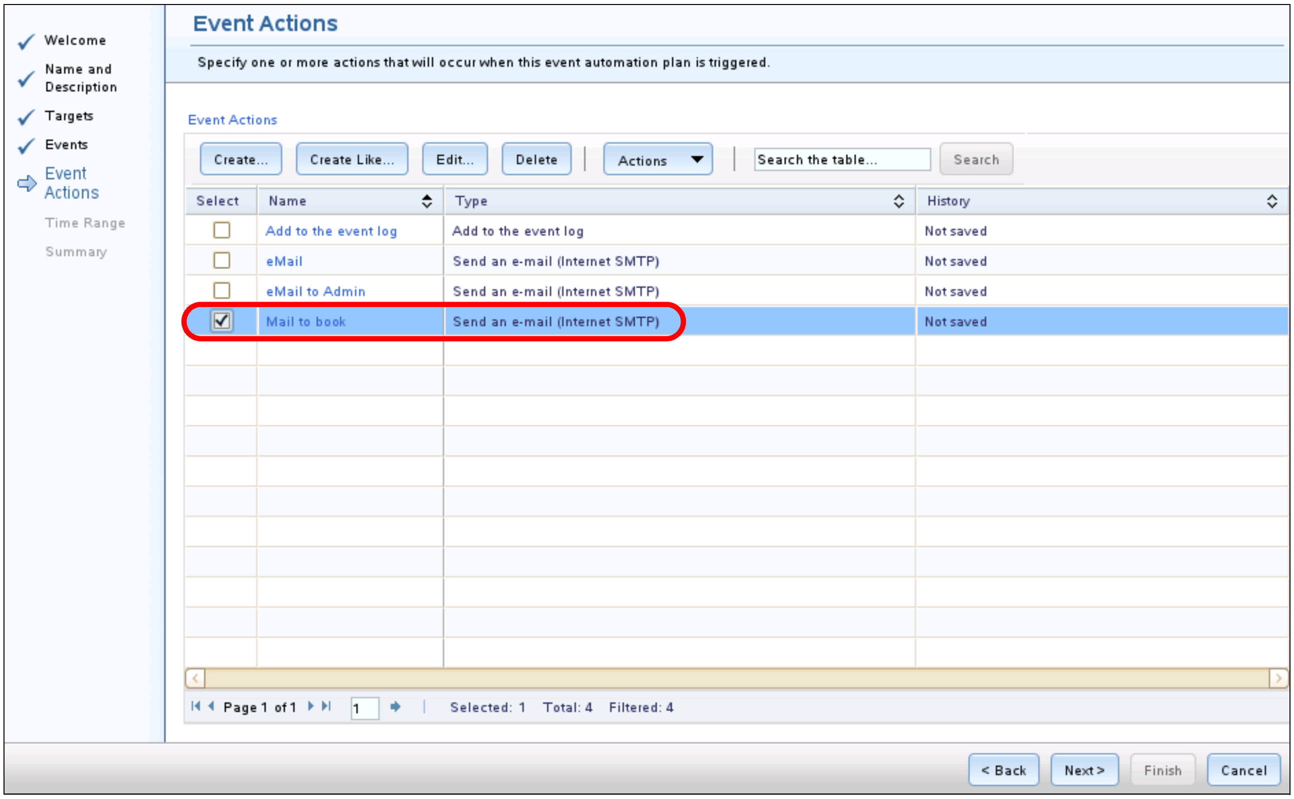

Event Actions

Specify one or more actions that will occur when this event automation plan is triggered.

Event Actions

Select	Name	Type	History
□	Add to the event log	Add to the event log	Not saved
□	eMail	Send an e-mail (Internet SMTP)	Not saved
□	eMail to Admin	Send an e-mail (Internet SMTP)	Not saved
☑	Mail to book	Send an e-mail (Internet SMTP)	Not saved

Selected: 1 Total: 4 Filtered: 4

Figure 4-115 Create Event Automation Plan wizard: Event Actions selection

7. On the next window, specify when the event automation plan can be activated. This time-range constraint can be helpful if you use one event automation plan for work days and another event automation plan for the weekend. You can create the event automation plan that works from Monday to Friday. And, you can create another event automation plan that works from Saturday to Sunday (Figure 4-116).

   In our example, we choose “All the time (24x7)”. But if you want, you can change it later by editing the created event automation plan. After you select the time range, click **Next**.

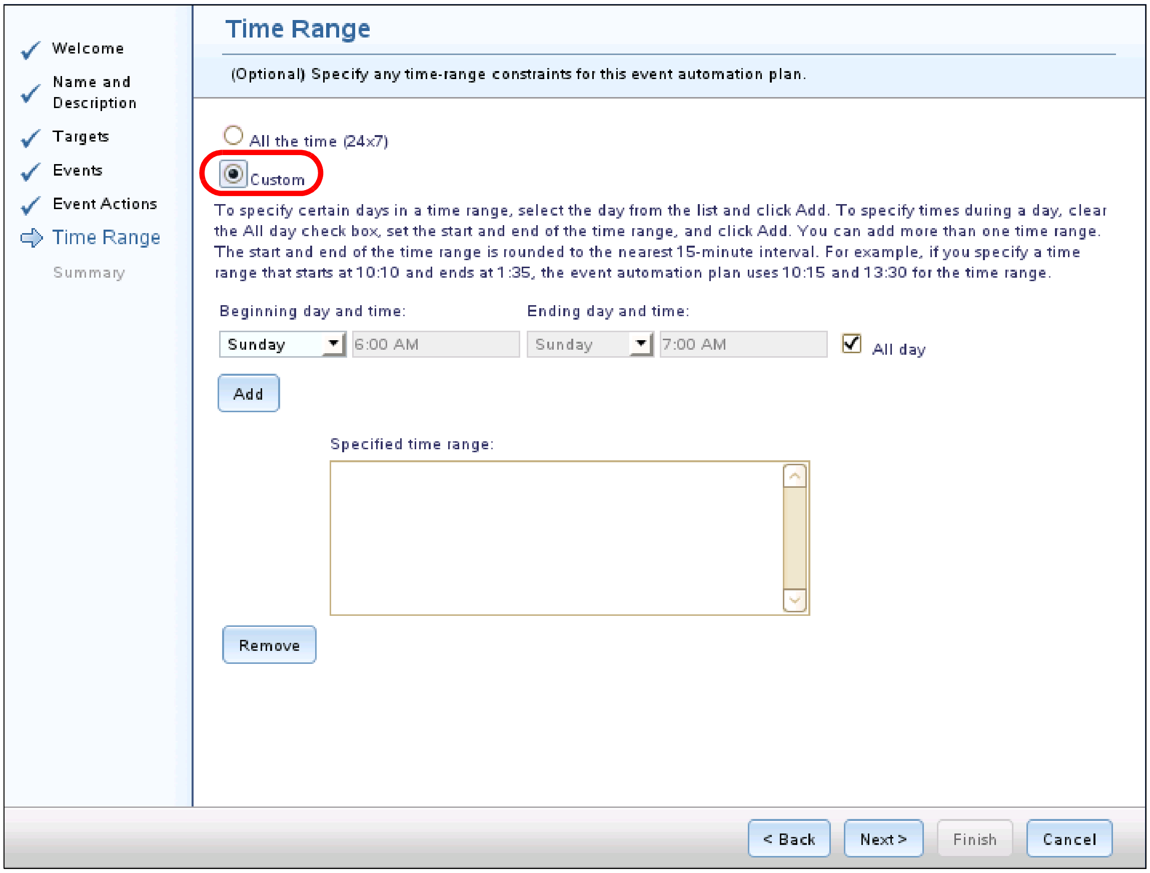

Time Range

(Optional) Specify any time-range constraints for this event automation plan.

○ All the time (24x7)

● Custom

To specify certain days in a time range, select the day from the list and click Add. To specify times during a day, clear the All day check box, set the start and end of the time range, and click Add. You can add more than one time range. The start and end of the time range is rounded to the nearest 15-minute interval. For example, if you specify a time range that starts at 10:10 and ends at 1:35, the event automation plan uses 10:15 and 13:30 for the time range.

Beginning day and time: Sunday 6:00 AM

Ending day and time: Sunday 7:00 AM

☑ All day

Specified time range:

Figure 4-116 Create Event Automation Plan wizard: Define a time range

8. In the next window (Figure 4-117), you see the summary for the event automation plan that you defined. Specify whether to activate the event automation plan by selecting the check box (default) after you click **Finish**.

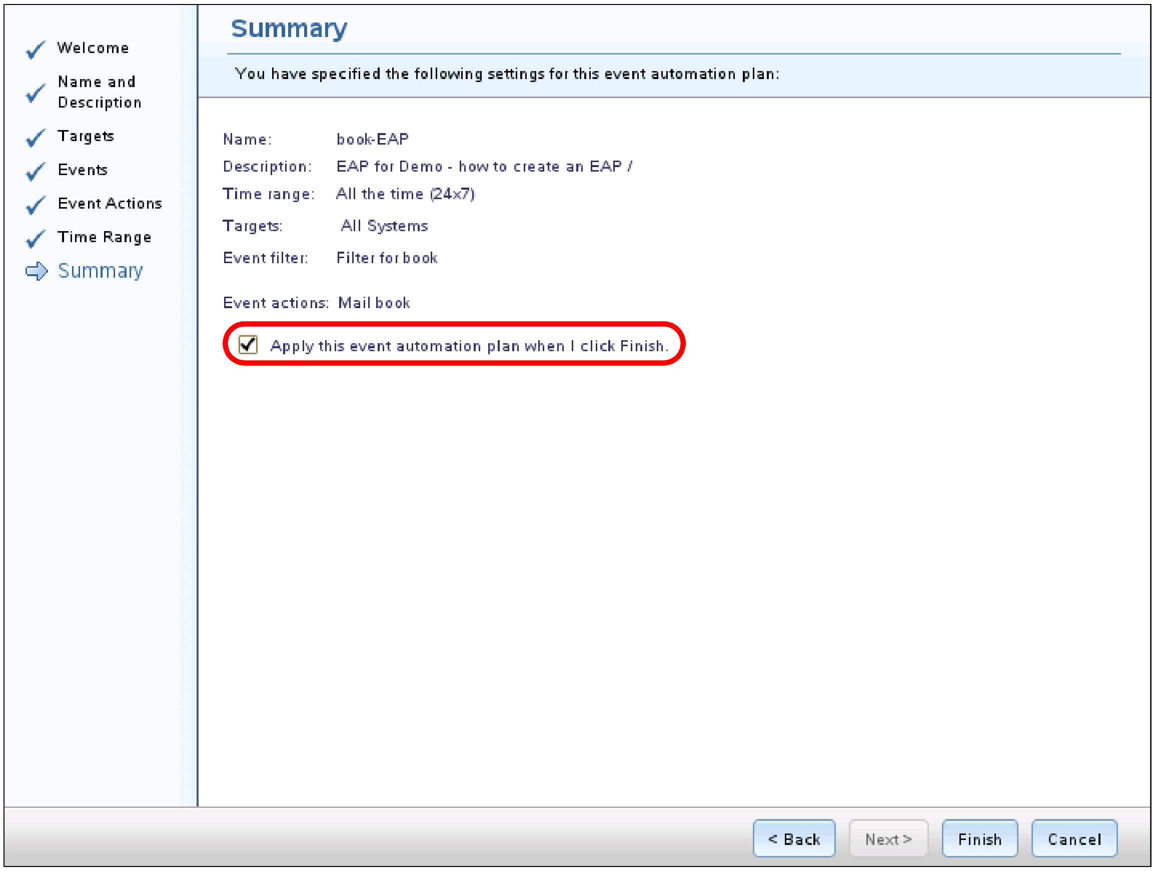

Summary

You have specified the following settings for this event automation plan:

Name: book-EAP

Description: EAP for Demo - how to create an EAP /

Time range: All the time (24x7)

Targets: All Systems

Event filter: Filter for book

Event actions: Mail book

☑ Apply this event automation plan when I click Finish.

Figure 4-117 Create Event Automation Plan wizard: Summary view

9. Now, you are back on the Event Automation Plans window. You can see the new event automation plan that you created (Figure 4-118).

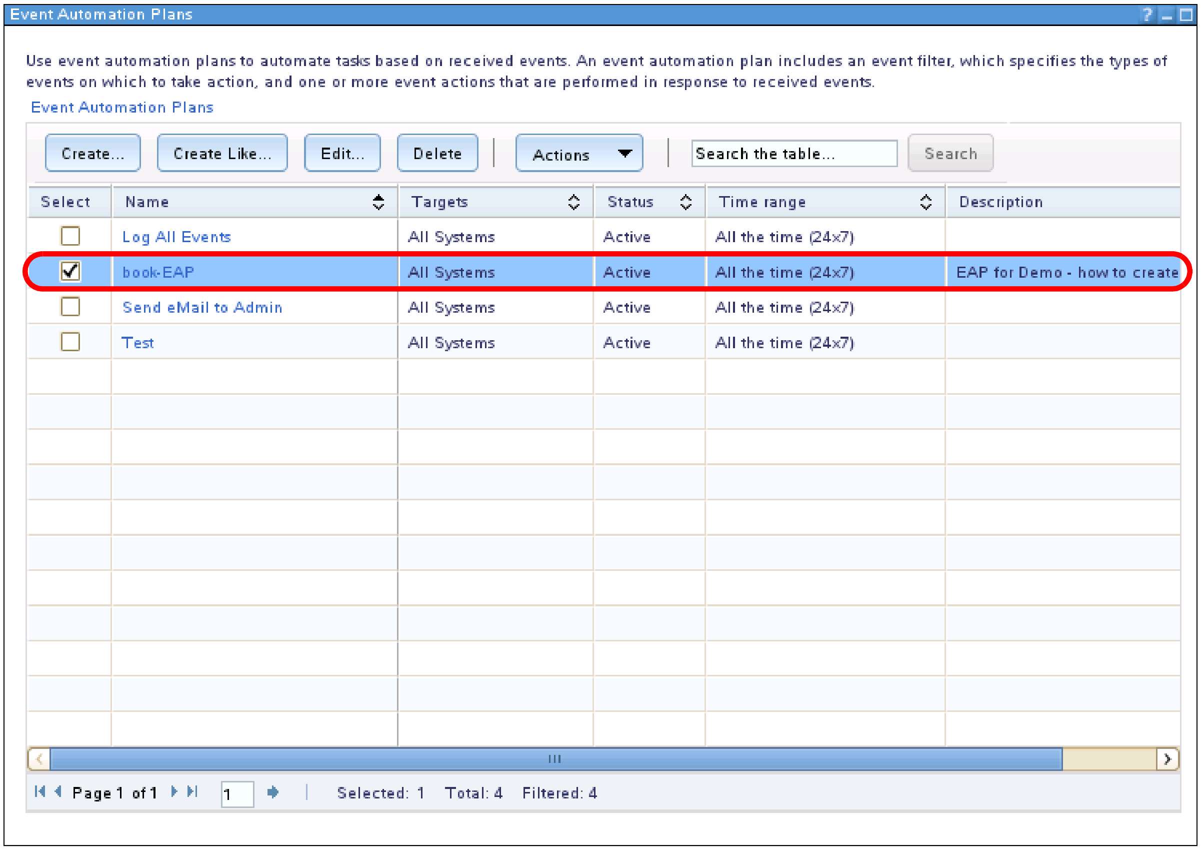

Event Automation Plans

Use event automation plans to automate tasks based on received events. An event automation plan includes an event filter, which specifies the types of events on which to take action, and one or more event actions that are performed in response to received events.

Event Automation Plans

Select	Name	Targets	Status	Time range	Description
□	Log All Events	All Systems	Active	All the time (24x7)	
☑	book-EAP	All Systems	Active	All the time (24x7)	EAP for Demo - how to create
□	Send eMail to Admin	All Systems	Active	All the time (24x7)	
□	Test	All Systems	Active	All the time (24x7)	

Page 1 of 1 | 1 | Selected: 1 Total: 4 Filtered: 4

Figure 4-118 Event Automation Plans window

The event automation plan is active and works for the defined systems and executes within its stated time period.

You can create as many event automation plans as you want, but keep the number of event automation plans to a minimum. If you have too many event automation plans, it can get confusing and you might get multiple alerts for each event.

You can use the GUI to export the event automation plan as a CSV file to use for documentation. On the Event Automation Plans window, select the plans that you want to export by clicking the associated check box. Then, select **Actions** → **Export**, as shown in Figure 4-119.

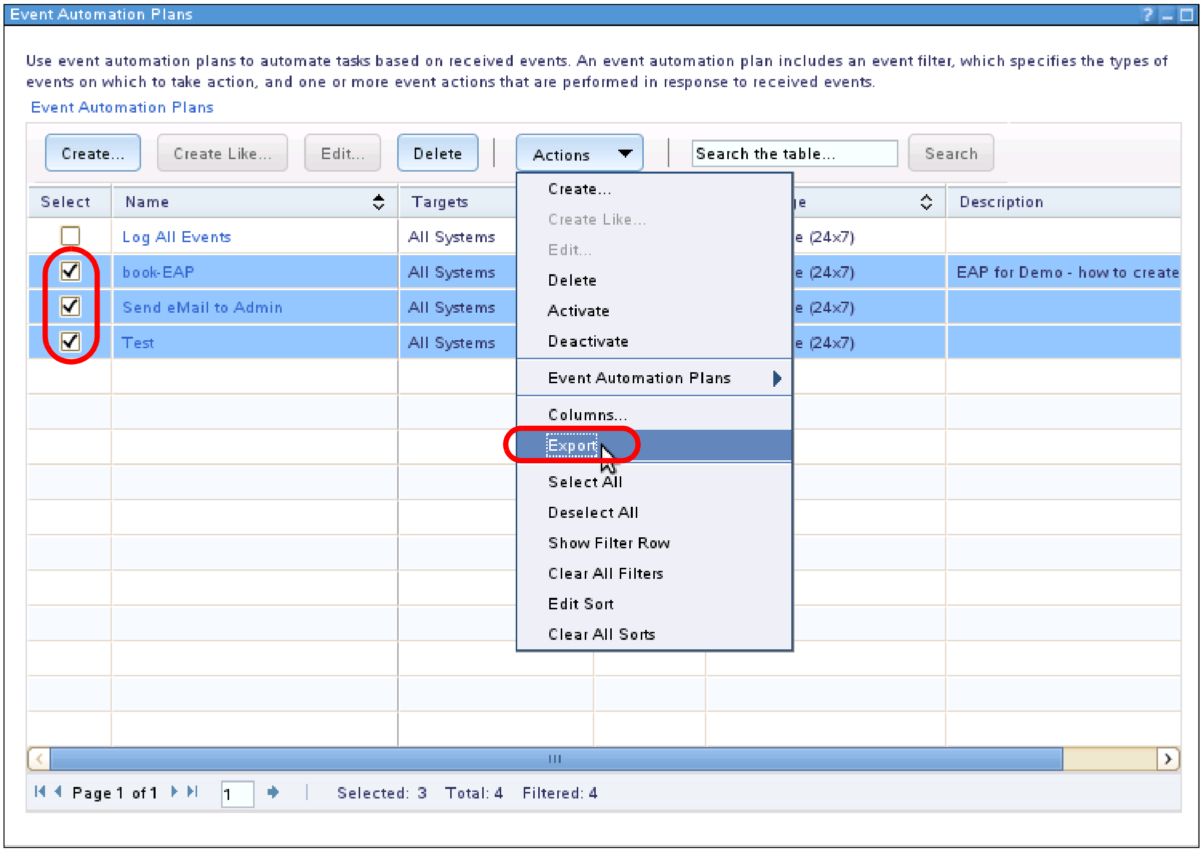
Event Automation Plans

Use event automation plans to automate tasks based on received events. An event automation plan includes an event filter, which specifies the types of events on which to take action, and one or more event actions that are performed in response to received events.

Event Automation Plans

Select	Name	Targets	…e	Description
□	Log All Events	All Systems	e (24x7)	
☑	book-EAP	All Systems	e (24x7)	EAP for Demo - how to create
☑	Send eMail to Admin	All Systems	e (24x7)	
☑	Test	All Systems	e (24x7)	

Page 1 of 1 | 1 | Selected: 3 Total: 4 Filtered: 4

Figure 4-119 Export the event automation plans

A window opens so that you can select the directory where you want to save the event automation plan (Figure 4-120). The name of the file is `Event_Automation_Plans.csv`.

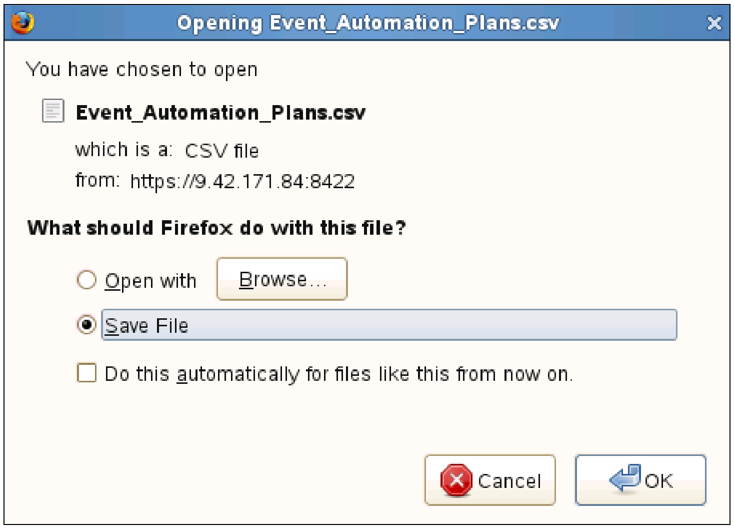
Opening Event_Automation_Plans.csv

You have chosen to open

Event_Automation_Plans.csv

which is a: CSV file

from: https://9.42.171.84:8422

What should Firefox do with this file?

○ Open with Browse...

● Save File

□ Do this automatically for files like this from now on.

Figure 4-120 Save the Event_Automation_Plans.csv file

If you want to export the event automation plan for use on another system or for backup and recovery, use the **smcli** command line. The use of the **smcli** command line is described in “Exporting systems and settings” on page 239, and in “Importing systems and settings” on page 242.

Creating an event filter

You can create an event filter from within the Create Event Automation Plan wizard. Or, select **Automation** → **Event Filters** on the left tab of the Systems Director home page, as shown in Figure 4-121.

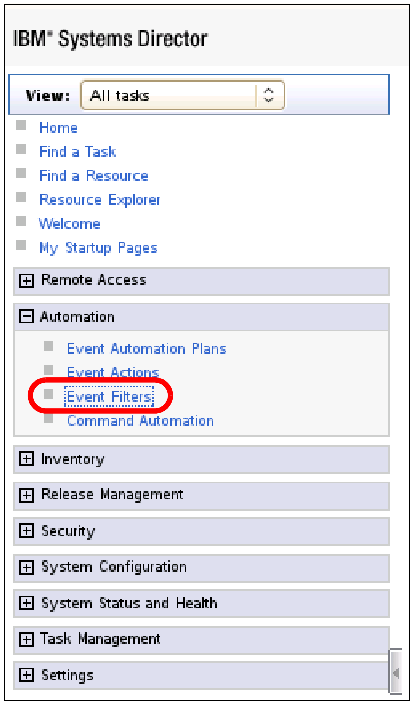

Figure 4-121 Selecting event filters

From this link, you are taken to the Event Filters page (Figure 4-122). Figure 4-122 shows the same list of filters in the Create Event Automation Plan wizard (Figure 4-114 on page 168).

Follow these steps to create an event filter:

1. As shown in Figure 4-122, click **Create** to create an event filter.

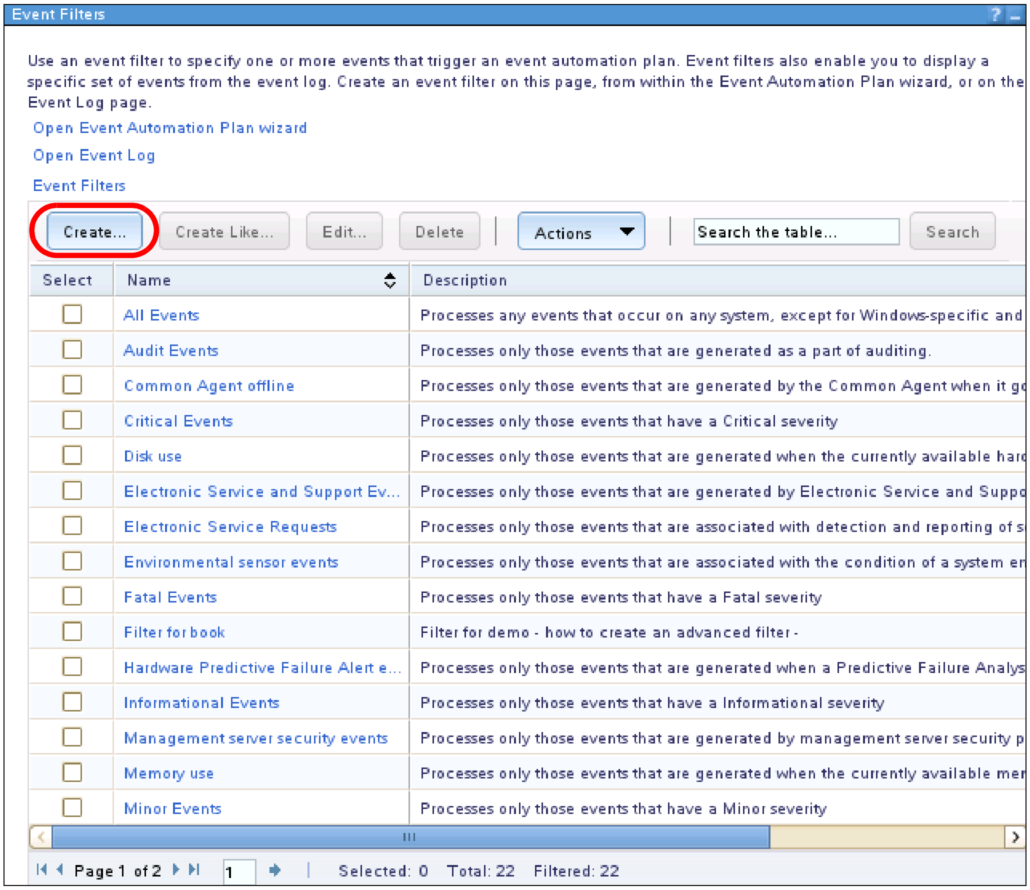

Use an event filter to specify one or more events that trigger an event automation plan. Event filters also enable you to display a specific set of events from the event log. Create an event filter on this page, from within the Event Automation Plan wizard, or on the Event Log page.

Open Event Automation Plan wizard

Open Event Log

Event Filters

Select	Name	Description
□	All Events	Processes any events that occur on any system, except for Windows-specific and
□	Audit Events	Processes only those events that are generated as a part of auditing.
□	Common Agent offline	Processes only those events that are generated by the Common Agent when it g
□	Critical Events	Processes only those events that have a Critical severity
□	Disk use	Processes only those events that are generated when the currently available har
□	Electronic Service and Support Ev...	Processes only those events that are generated by Electronic Service and Suppo
□	Electronic Service Requests	Processes only those events that are associated with detection and reporting of s
□	Environmental sensor events	Processes only those events that are associated with the condition of a system en
□	Fatal Events	Processes only those events that have a Fatal severity
□	Filter for book	Filter for demo - how to create an advanced filter -
□	Hardware Predictive Failure Alert e...	Processes only those events that are generated when a Predictive Failure Analys
□	Informational Events	Processes only those events that have a Informational severity
□	Management server security events	Processes only those events that are generated by management server security p
□	Memory use	Processes only those events that are generated when the currently available mer
□	Minor Events	Processes only those events that have a Minor severity

Page 1 of 2 1 Selected: 0 Total: 22 Filtered: 22

Figure 4-122 Listing of the event filters

2. The Create Event Filter wizard starts and displays the Welcome page (Figure 4-123). Click **Next** to continue.

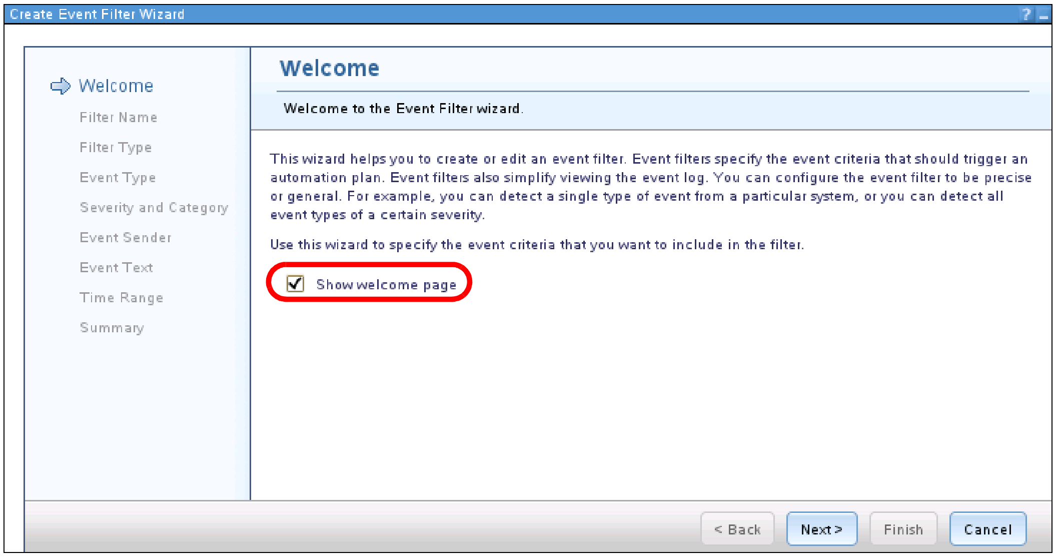

Welcome

Welcome to the Event Filter wizard.

This wizard helps you to create or edit an event filter. Event filters specify the event criteria that should trigger an automation plan. Event filters also simplify viewing the event log. You can configure the event filter to be precise or general. For example, you can detect a single type of event from a particular system, or you can detect all event types of a certain severity.

Use this wizard to specify the event criteria that you want to include in the filter.

☑ Show welcome page

Figure 4-123 Create Event Filter wizard: Welcome window

3. In the Filter Name window (Figure 4-124), enter the name and the description for the filter. In our example, we use the name `Filter for book`. You can also add a short description for the filter. Click **Next** to continue.

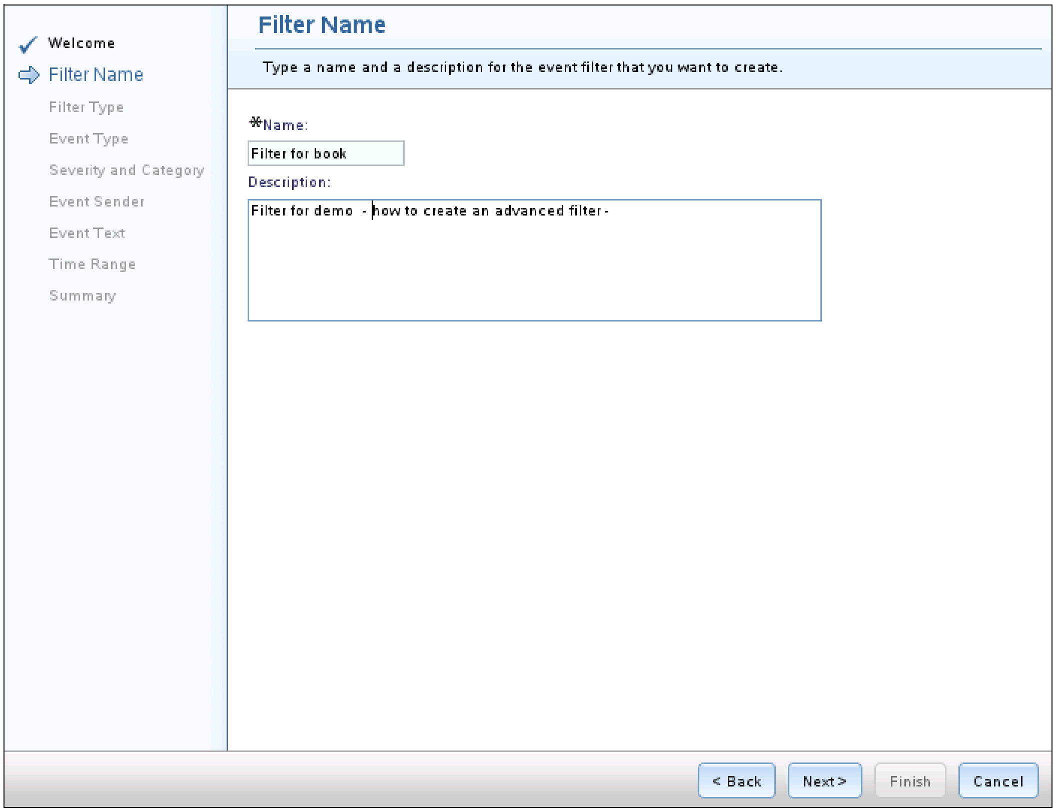

Filter Name

Type a name and a description for the event filter that you want to create.

*Name: Filter for book

Description: Filter for demo - how to create an advanced filter -

Figure 4-124 Create Event Filter wizard: Filter Name window

4. In the Filter Type window (Figure 4-125), select the type of filter that you want to create. We select “Simple event filter” for our example.

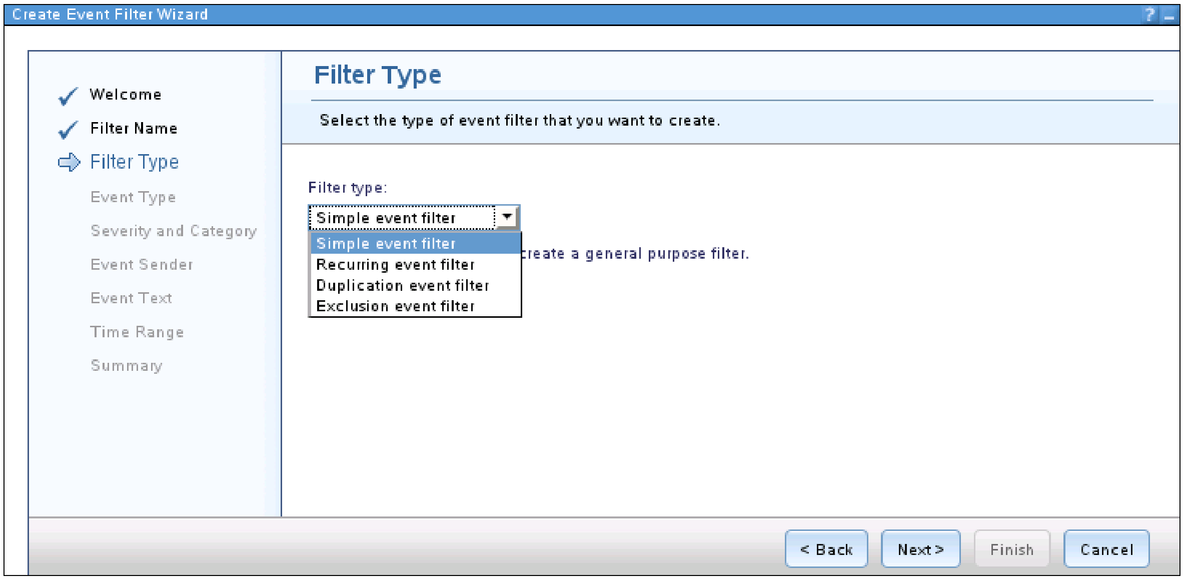

Filter Type

Select the type of event filter that you want to create.

Filter type: Simple event filter

reate a general purpose filter.

Figure 4-125 Create Event Filter wizard: Filter Type window

The following types of filters are available. Select the filter type and click **Next**.

- Simple event filter: Use the general-purpose filter to create your own filter.
- Recurring event filter: Use this filter to trigger only when the included event meets the filter criteria more than one time in the defined time range.
- Duplication event filter: Use this filter to ignore duplicate events.
- Exclusion event filter: Use this filter to exclude a specific event type from a larger list of event types that you included in the event.

5. In the Event Type window, Figure 4-126, select the filter type and define the filter type that you want use. The following event types are available:

   – Default: Include all events except IBM System i message queue events, which can be selected by clicking the check box, and Windows specific events. If you need to select the Windows specific events, use the Custom type.

   – Common: Include events that are often used in the custom environment. The custom environment events include general events, such as information about updates or user security events. General events also include hardware events, such as power, storage, fan, or processor events. You can add the system message events.

   – Custom: Include events of a certain category, type, or value. The available events depend on the system types, operating systems, or protocols that you use in your Systems Director environment.

   Our example that is shown in Figure 4-126 uses the default events. We select the “Default” filter type, which includes all events. Click **Next** to continue.

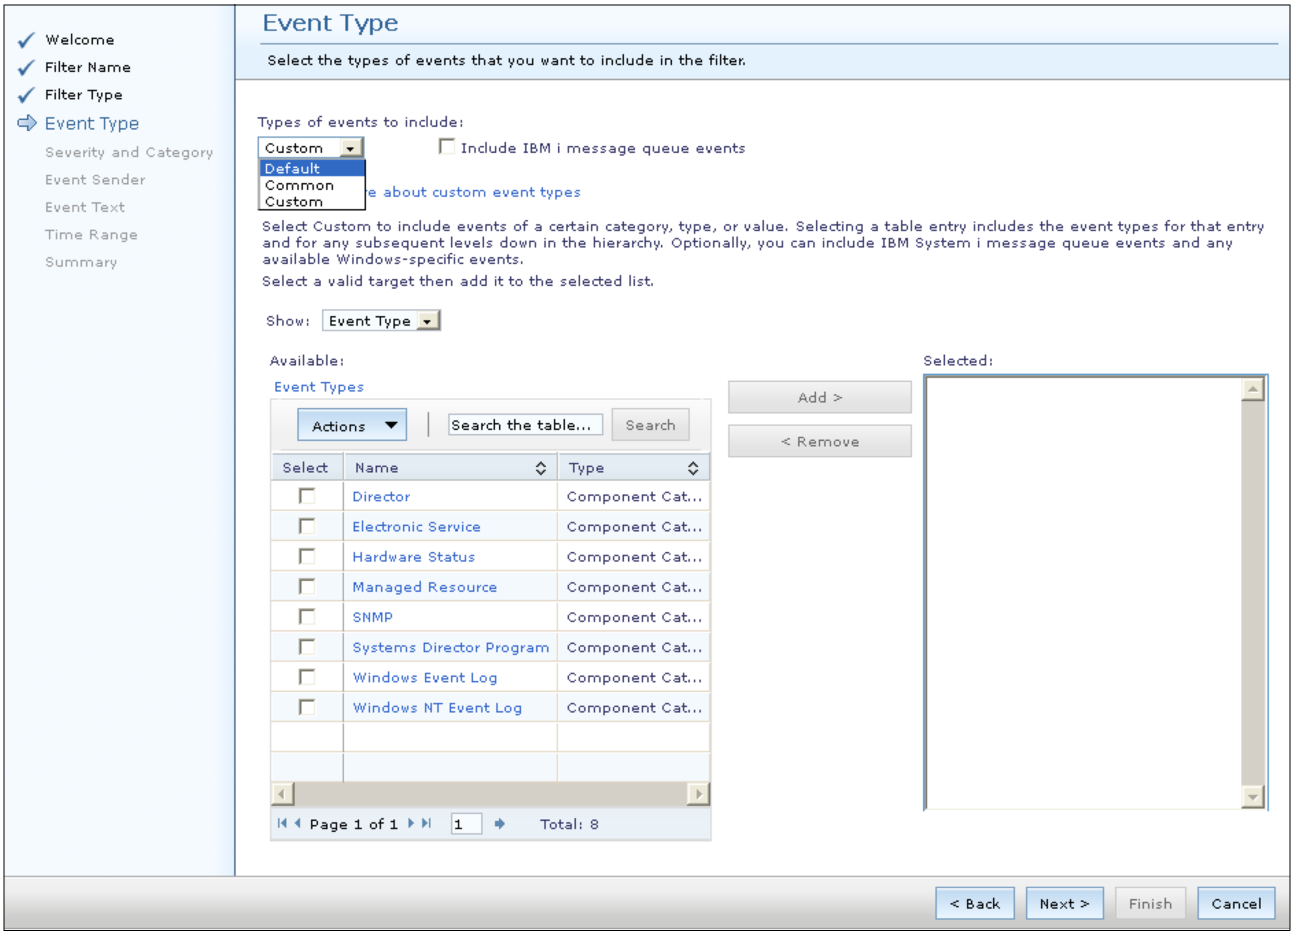

Event Type

Select the types of events that you want to include in the filter.

Types of events to include: Custom

☐ Include IBM i message queue events

e about custom event types

Select Custom to include events of a certain category, type, or value. Selecting a table entry includes the event types for that entry and for any subsequent levels down in the hierarchy. Optionally, you can include IBM System i message queue events and any available Windows-specific events.

Select a valid target then add it to the selected list.

Show: Event Type

Available:

Event Types

Select	Name	Type
☐	Director	Component Cat...
☐	Electronic Service	Component Cat...
☐	Hardware Status	Component Cat...
☐	Managed Resource	Component Cat...
☐	SNMP	Component Cat...
☐	Systems Director Program	Component Cat...
☐	Windows Event Log	Component Cat...
☐	Windows NT Event Log	Component Cat...

Page 1 of 1 | 1 | Total: 8

Selected:

Figure 4-126 Create Event Filter wizard: Event Type selection

6. In the Severity and Category window (Figure 4-127), select the severity and category for the filter. Various severities are available for events in Systems Director:

   – Fatal
   – Critical
   – Minor
   – Warning
   – Informational
   – Unknown

   Two event categories are available:

   – Alert
   – Resolution

   In our example, we use the Fatal, Critical, and Warning severities and the Alert category, as shown in Figure 4-127. Click **Next** to continue.

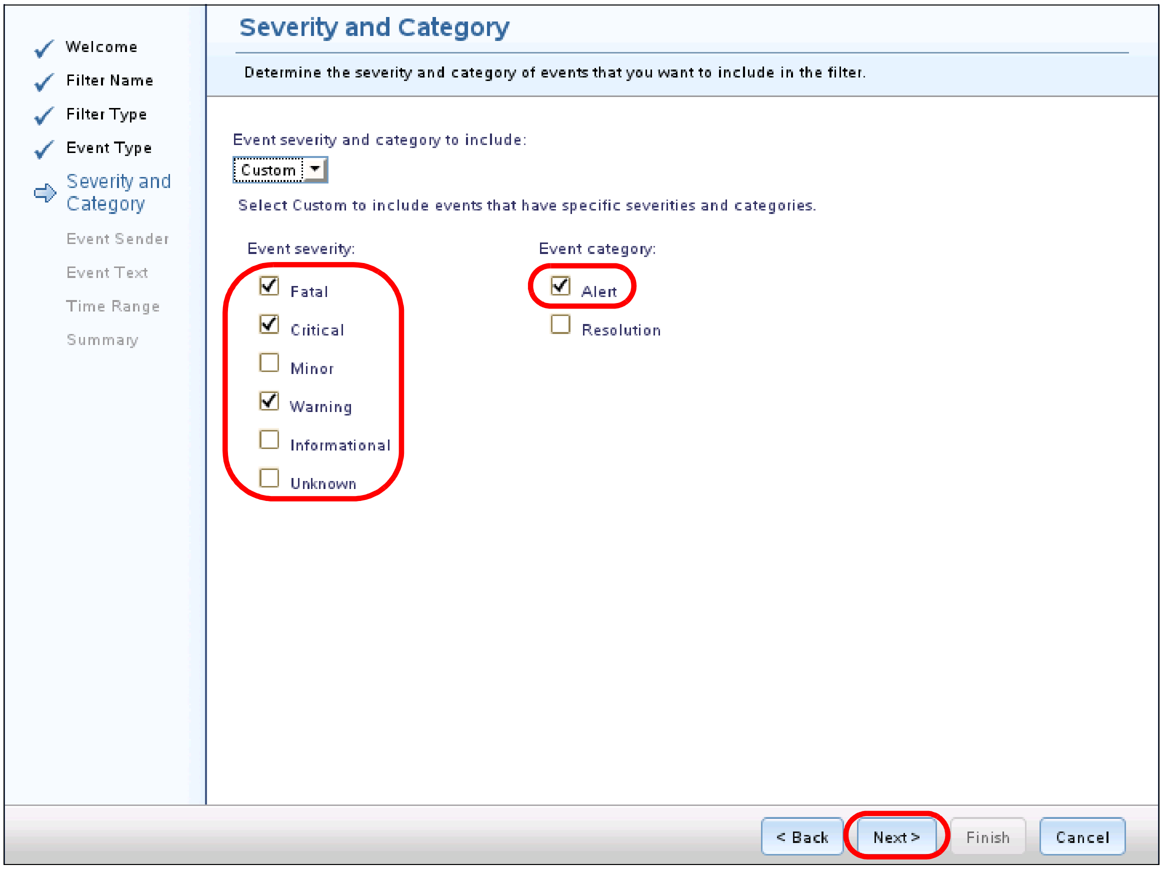

Severity and Category

Determine the severity and category of events that you want to include in the filter.

Event severity and category to include: Custom

Select Custom to include events that have specific severities and categories.

Event severity:

☑ Fatal
☑ Critical
□ Minor
☑ Warning
□ Informational
□ Unknown

Event category:

☑ Alert
□ Resolution

Figure 4-127 Create Event Filter wizard: Severity and Category selection

7. In the Event Sender window (Figure 4-128), select the system that you want to include in this filter:

   – Default: Includes all systems that Systems Director discovered or can access.
   – Custom: Select individual systems or groups to include in this filter.

   If you select “Custom”, the window expands (see Figure 4-128). On the left, you see a box to enter additional systems and a list of systems. Select the systems and click Add to add these systems to the “Selected senders” list. The filter works for only these systems.

   In our example, we use “Default” to select all systems because we use the filter in the event automation plan. In the event automation plan, you can also select the systems for which the event automation plan works. If you select specific systems on this window and different systems in the event automation plan, no events are handled through the event automation plan. Therefore, if you plan to use the filter in an event automation plan, leave the selection that is shown in Figure 4-128 as “Default”. Use Custom when you want to use the filter for event capture only. Do not use Custom in an event automation plan with event actions that use it.

   Click **Next** to continue.

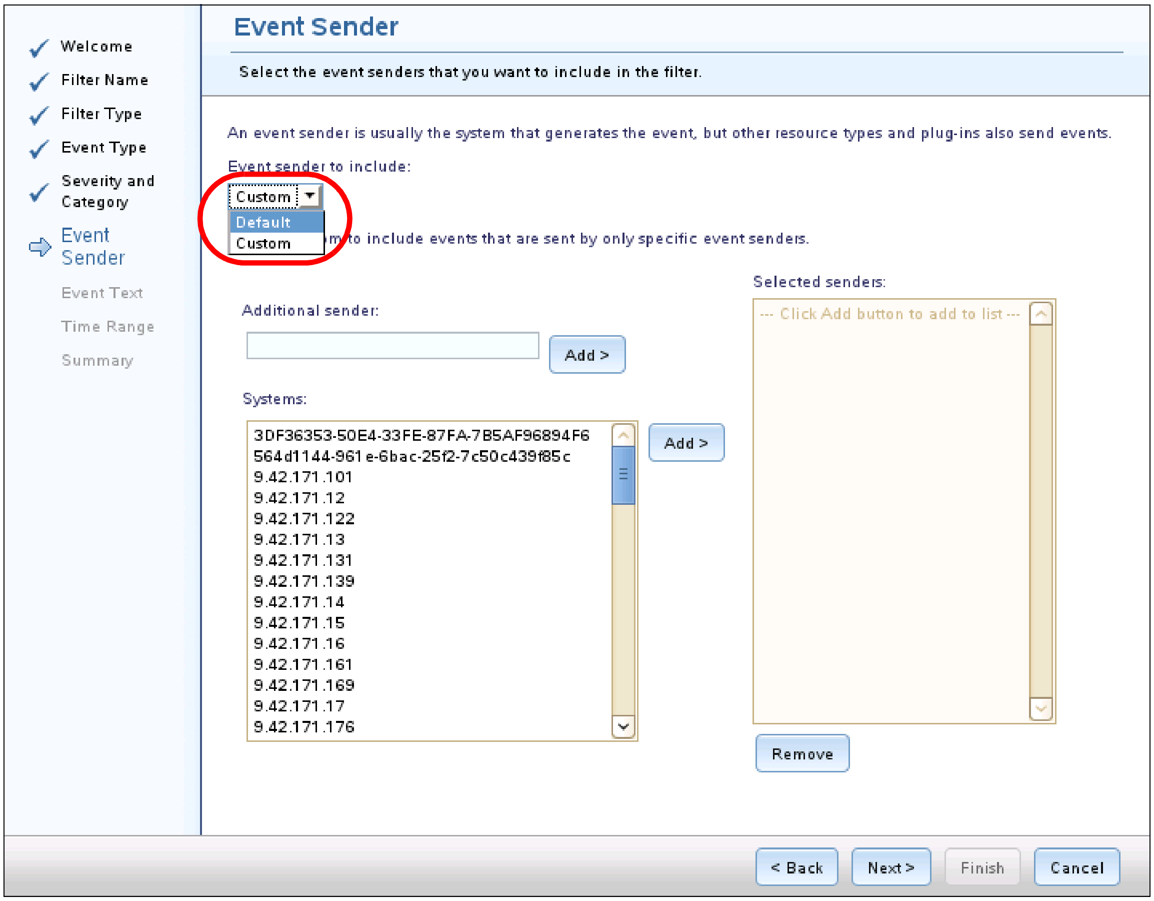

Event Sender

Select the event senders that you want to include in the filter.

An event sender is usually the system that generates the event, but other resource types and plug-ins also send events.

Event sender to include: Custom

m to include events that are sent by only specific event senders.

Additional sender:

Systems:

3DF36353-50E4-33FE-87FA-7B5AF96894F6
564d1144-961e-6bac-25f2-7c50c439f85c
9.42.171.101
9.42.171.12
9.42.171.122
9.42.171.13
9.42.171.131
9.42.171.139
9.42.171.14
9.42.171.15
9.42.171.16
9.42.171.161
9.42.171.169
9.42.171.17
9.42.171.176

Selected senders:

--- Click Add button to add to list ---

Figure 4-128 Create Event Filter wizard: Event Sender selection

8. In the Event Text window (Figure 4-129), select the event text. Two selections are possible:

   – Default: Include all event text.

   – Custom: Filter for specific event text. You might be interested in this option if you want only specific events from systems. Select Custom to specify a word, separate words, or a phrase that you want to include in the filter. The filter is triggered by only those events that you include in the filter that also contains the specified text.

   In our example, we leave the selection on Default. Therefore, we want to get all alerts from the alert type that we chose earlier. Click **Next** to continue.

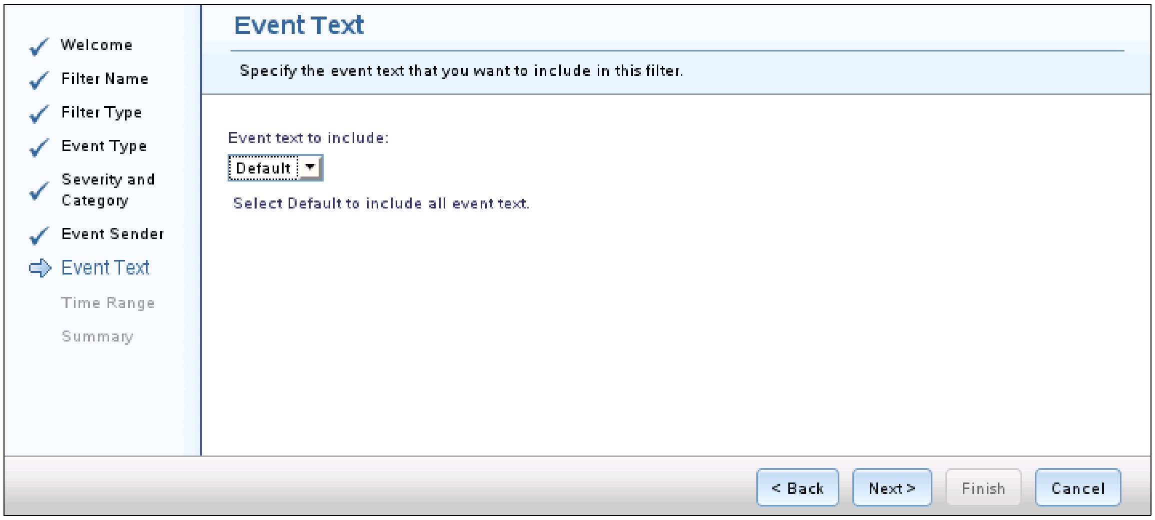

Event Text

Specify the event text that you want to include in this filter.

Event text to include: Default

Select Default to include all event text.

Figure 4-129 Create Event Filter wizard: Event Text selection

9. In the Time Range window (Figure 4-130), select a time range for the filter. You can either select “All” (the default), which is 24x7, or “Custom”. If you select Custom, you can define the days or hours that the filter works.

   In our example, we use the filter that is in the event automation plan; therefore, we keep the All setting on Figure 4-130. We also can set a time range in the Event Automation Plan wizard. We do not want a conflict between the settings in the filter and the settings in the event automation plan. The setting in Figure 4-130 is used if you use a filter only to capture events and not with an action.

   Click **Next** to continue.

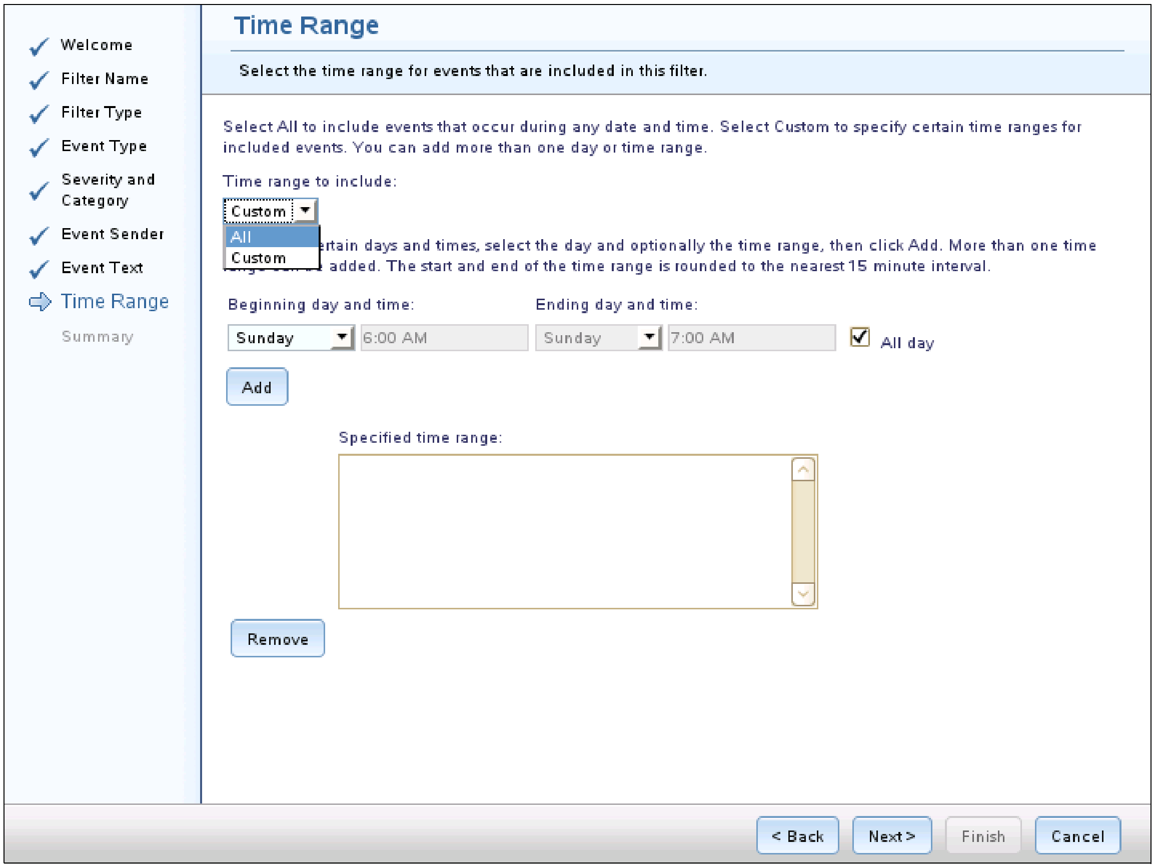

Time Range

Select the time range for events that are included in this filter.

Select All to include events that occur during any date and time. Select Custom to specify certain time ranges for included events. You can add more than one day or time range.

Time range to include: Custom

rtain days and times, select the day and optionally the time range, then click Add. More than one time added. The start and end of the time range is rounded to the nearest 15 minute interval.

Beginning day and time: Sunday 6:00 AM

Ending day and time: Sunday 7:00 AM

☑ All day

Specified time range:

Figure 4-130 Create Event Filter wizard: Time Range selection

10.The last window of the wizard shows the summary view for the filter that you defined (Figure 4-131). Check the settings and information and click **Finish** to create this filter.

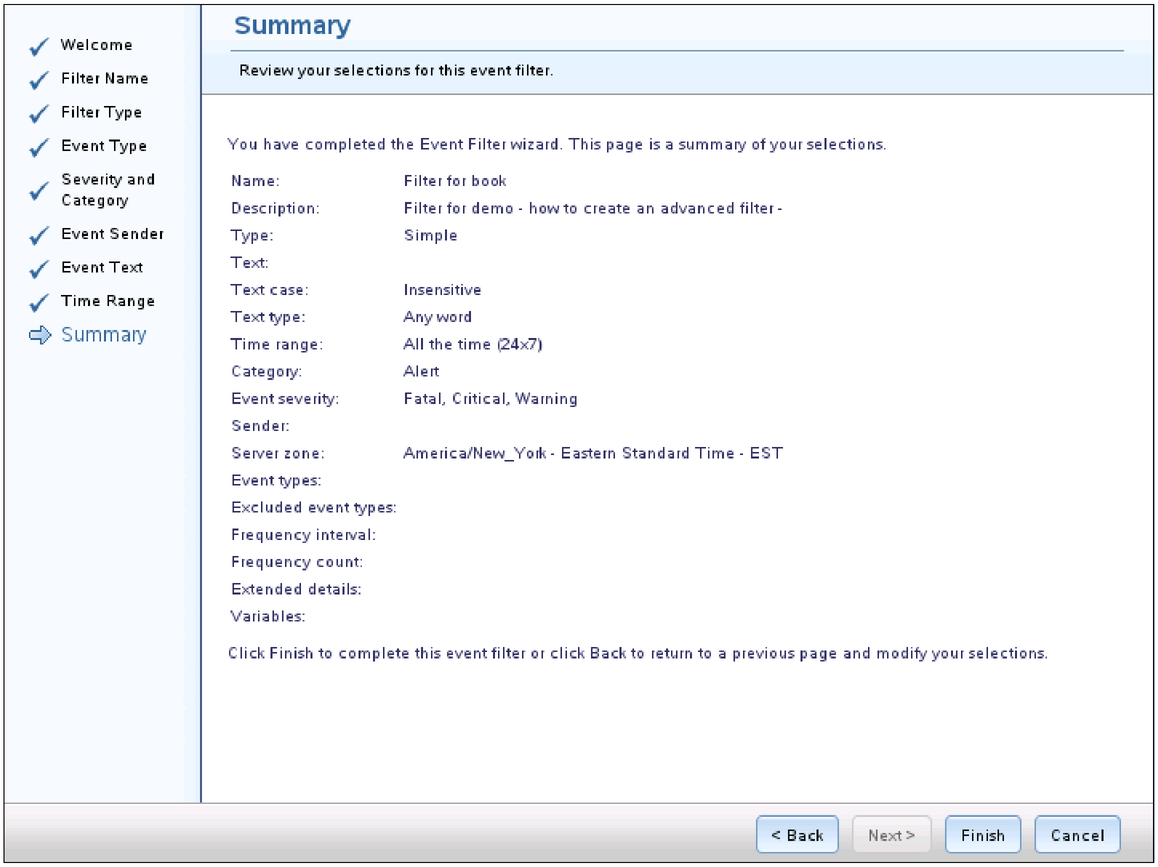

Summary

Review your selections for this event filter.

You have completed the Event Filter wizard. This page is a summary of your selections.

Name:	Filter for book
Description:	Filter for demo - how to create an advanced filter -
Type:	Simple
Text:	
Text case:	Insensitive
Text type:	Any word
Time range:	All the time (24x7)
Category:	Alert
Event severity:	Fatal, Critical, Warning
Sender:	
Server zone:	America/New_York - Eastern Standard Time - EST
Event types:	
Excluded event types:	
Frequency interval:	
Frequency count:	
Extended details:	
Variables:	

Click Finish to complete this event filter or click Back to return to a previous page and modify your selections.

Figure 4-131 Create Event Filter wizard: Summary view

The filter is created. The filter is available in the list of the filters and can be used in event automation plans.

Creating an event action

You can create event actions in the Create Event Automation Plans wizard or select the Event Actions link on the left panel of the Systems Director home page (Figure 4-132).

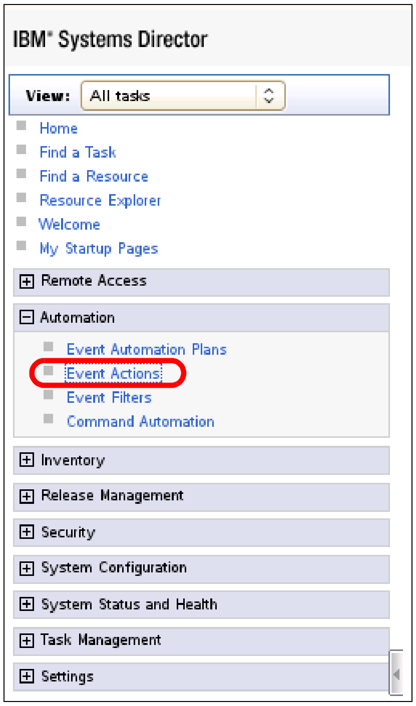

Figure 4-132 Selecting event actions

From this link, you are taken to the Event Actions page (Figure 4-133). This list shows the same actions in the Create Event Automation Plans wizard (Figure 4-115 on page 169).

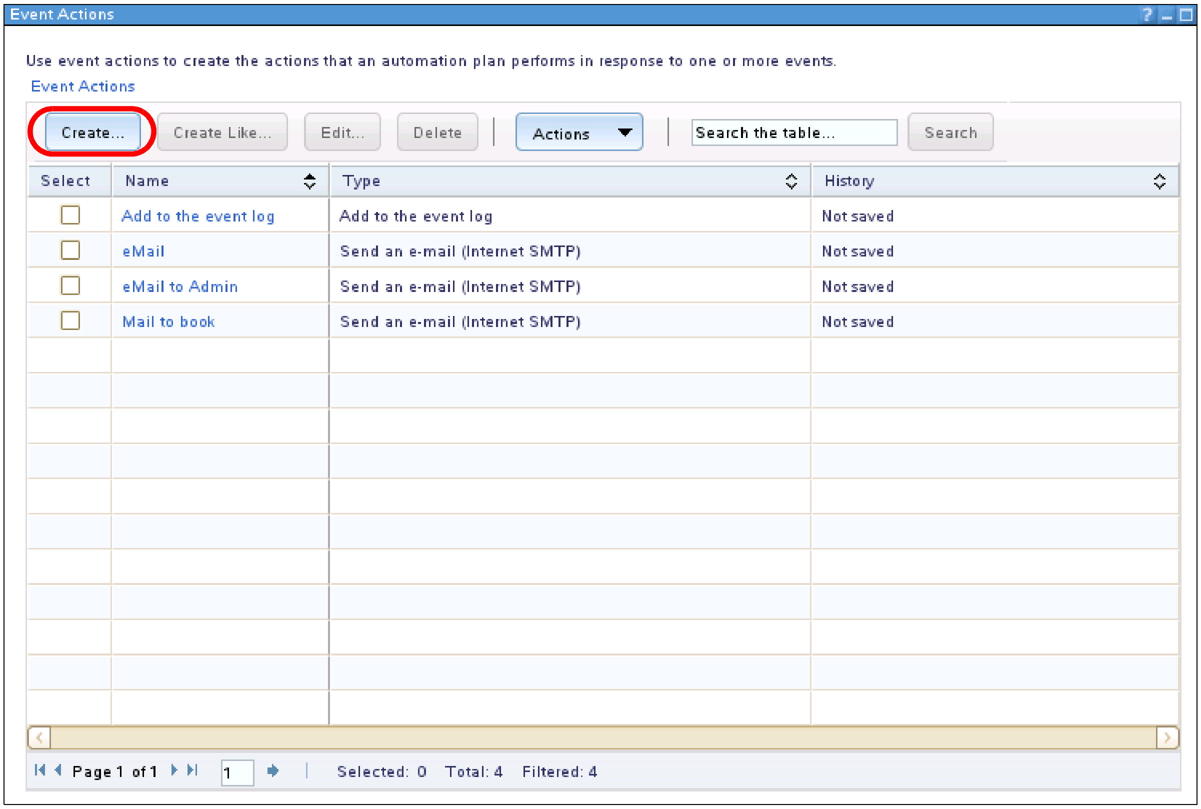

Use event actions to create the actions that an automation plan performs in response to one or more events.

Event Actions

Select	Name	Type	History
□	Add to the event log	Add to the event log	Not saved
□	eMail	Send an e-mail (Internet SMTP)	Not saved
□	eMail to Admin	Send an e-mail (Internet SMTP)	Not saved
□	Mail to book	Send an e-mail (Internet SMTP)	Not saved

Page 1 of 1 | 1 | Selected: 0 Total: 4 Filtered: 4

Figure 4-133 Event Actions

Follow these steps to create an event action:

1. Click **Create** (see Figure 4-133) to create an event action. Or, select an existing action and click **Edit** to change an existing action.

2. Figure 4-134 opens with the available event actions.

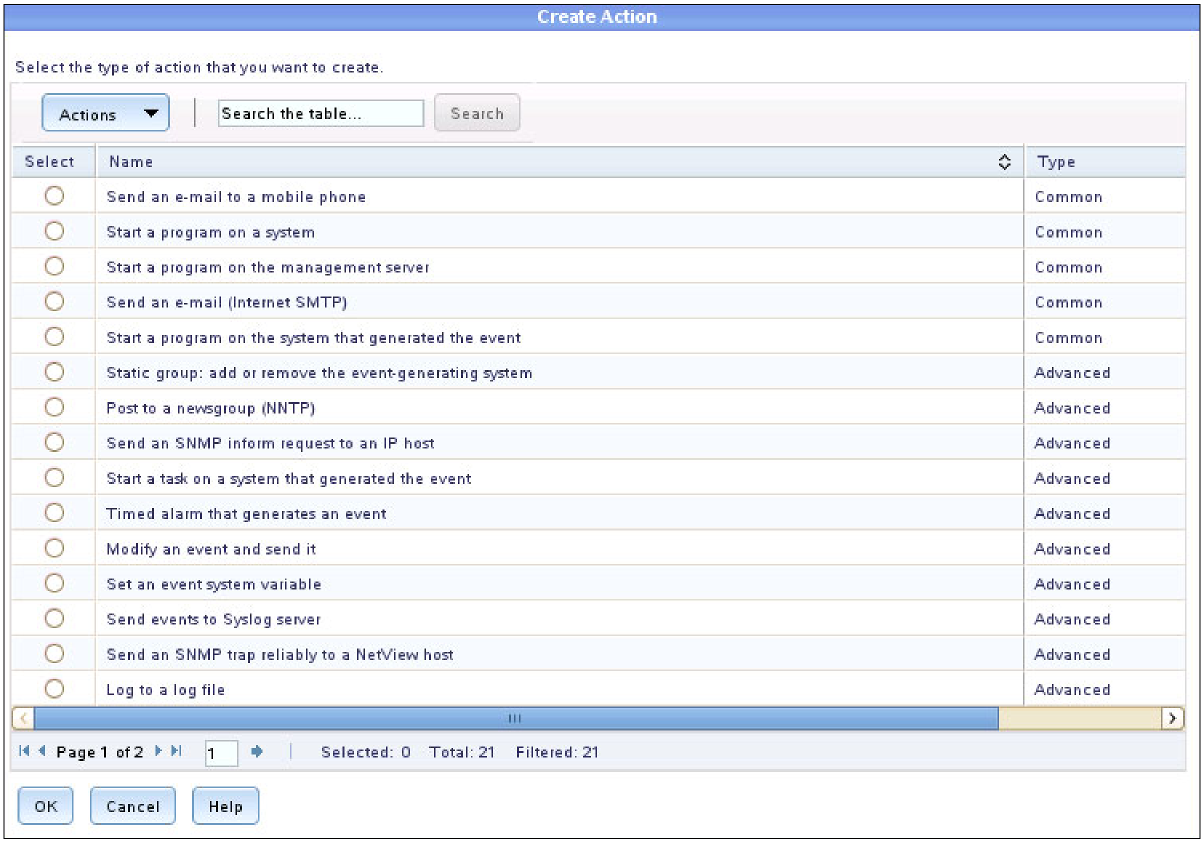

Create Action

Select the type of action that you want to create.

Select	Name	Type
○	Send an e-mail to a mobile phone	Common
○	Start a program on a system	Common
○	Start a program on the management server	Common
○	Send an e-mail (Internet SMTP)	Common
○	Start a program on the system that generated the event	Common
○	Static group: add or remove the event-generating system	Advanced
○	Post to a newsgroup (NNTP)	Advanced
○	Send an SNMP inform request to an IP host	Advanced
○	Start a task on a system that generated the event	Advanced
○	Timed alarm that generates an event	Advanced
○	Modify an event and send it	Advanced
○	Set an event system variable	Advanced
○	Send events to Syslog server	Advanced
○	Send an SNMP trap reliably to a NetView host	Advanced
○	Log to a log file	Advanced

Page 1 of 2 1 Selected: 0 Total: 21 Filtered: 21

Figure 4-134 Create Action window (Page 1 of 2)

Page 2 of the list of available actions is shown in Figure 4-135.

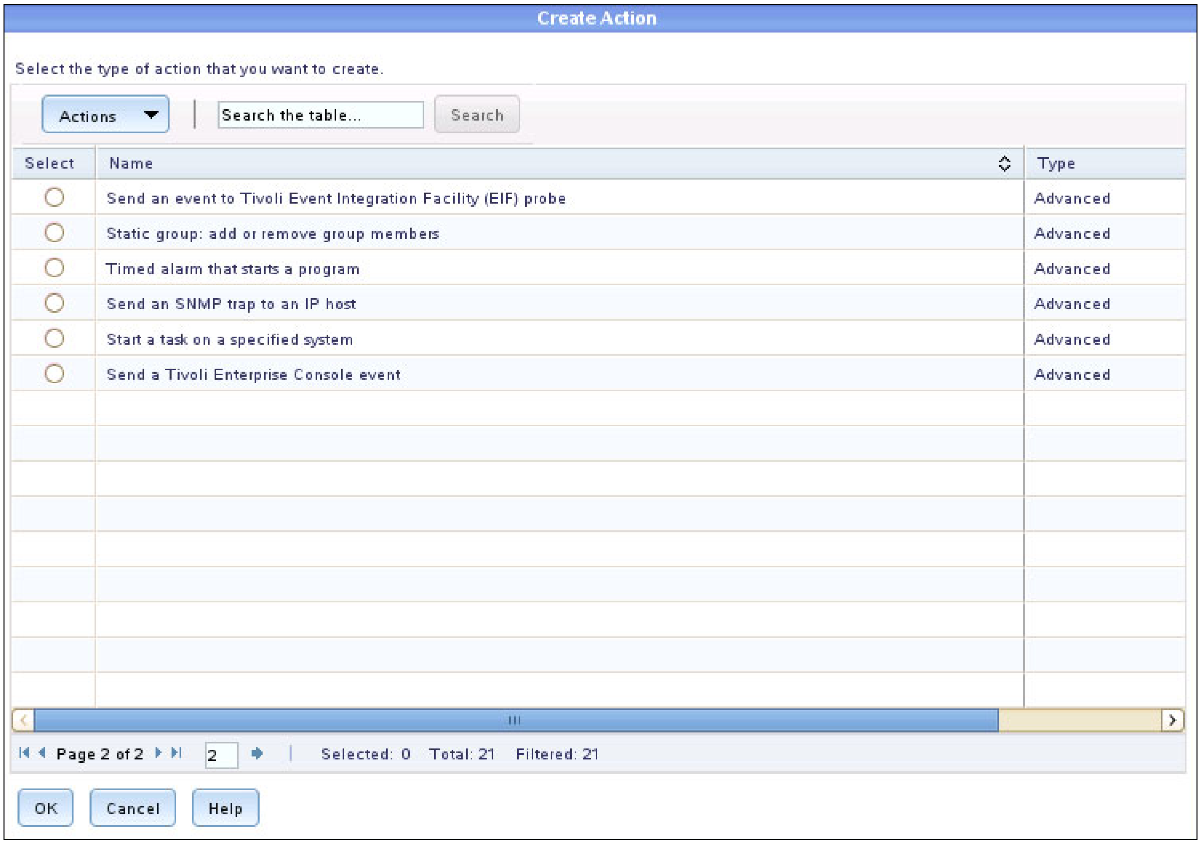
Create Action
Select the type of action that you want to create.
Actions
Search the table...
Search

Select	Name	Type
○	Send an event to Tivoli Event Integration Facility (EIF) probe	Advanced
○	Static group: add or remove group members	Advanced
○	Timed alarm that starts a program	Advanced
○	Send an SNMP trap to an IP host	Advanced
○	Start a task on a specified system	Advanced
○	Send a Tivoli Enterprise Console event	Advanced

Page 2 of 2
2
Selected: 0 Total: 21 Filtered: 21
OK
Cancel
Help

Figure 4-135 Create Action window (Page 2 of 2)

3. Select an action and click **OK**. In our example, we choose “Send an e-mail (Internet SMTP)”, as shown in Figure 4-136.

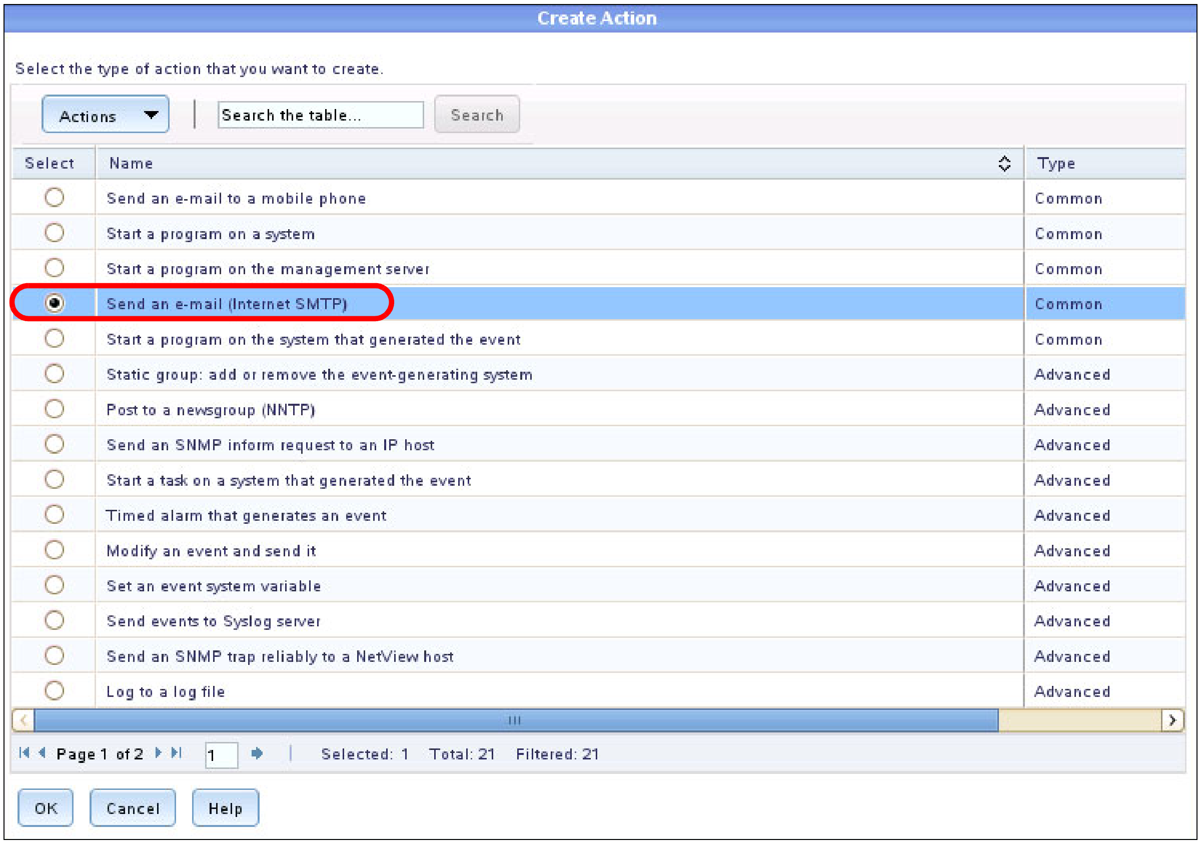

Select the type of action that you want to create.

Select	Name	Type
○	Send an e-mail to a mobile phone	Common
○	Start a program on a system	Common
○	Start a program on the management server	Common
●	Send an e-mail (Internet SMTP)	Common
○	Start a program on the system that generated the event	Common
○	Static group: add or remove the event-generating system	Advanced
○	Post to a newsgroup (NNTP)	Advanced
○	Send an SNMP inform request to an IP host	Advanced
○	Start a task on a system that generated the event	Advanced
○	Timed alarm that generates an event	Advanced
○	Modify an event and send it	Advanced
○	Set an event system variable	Advanced
○	Send events to Syslog server	Advanced
○	Send an SNMP trap reliably to a NetView host	Advanced
○	Log to a log file	Advanced

Page 1 of 2 1 Selected: 1 Total: 21 Filtered: 21

Figure 4-136 Select an action to configure

4. The configuration window opens for your selected action. The content of the window varies depending on the action that you select. Our example shows the settings for the “Send an e-mail (Internet SMTP)” action (Figure 4-137).

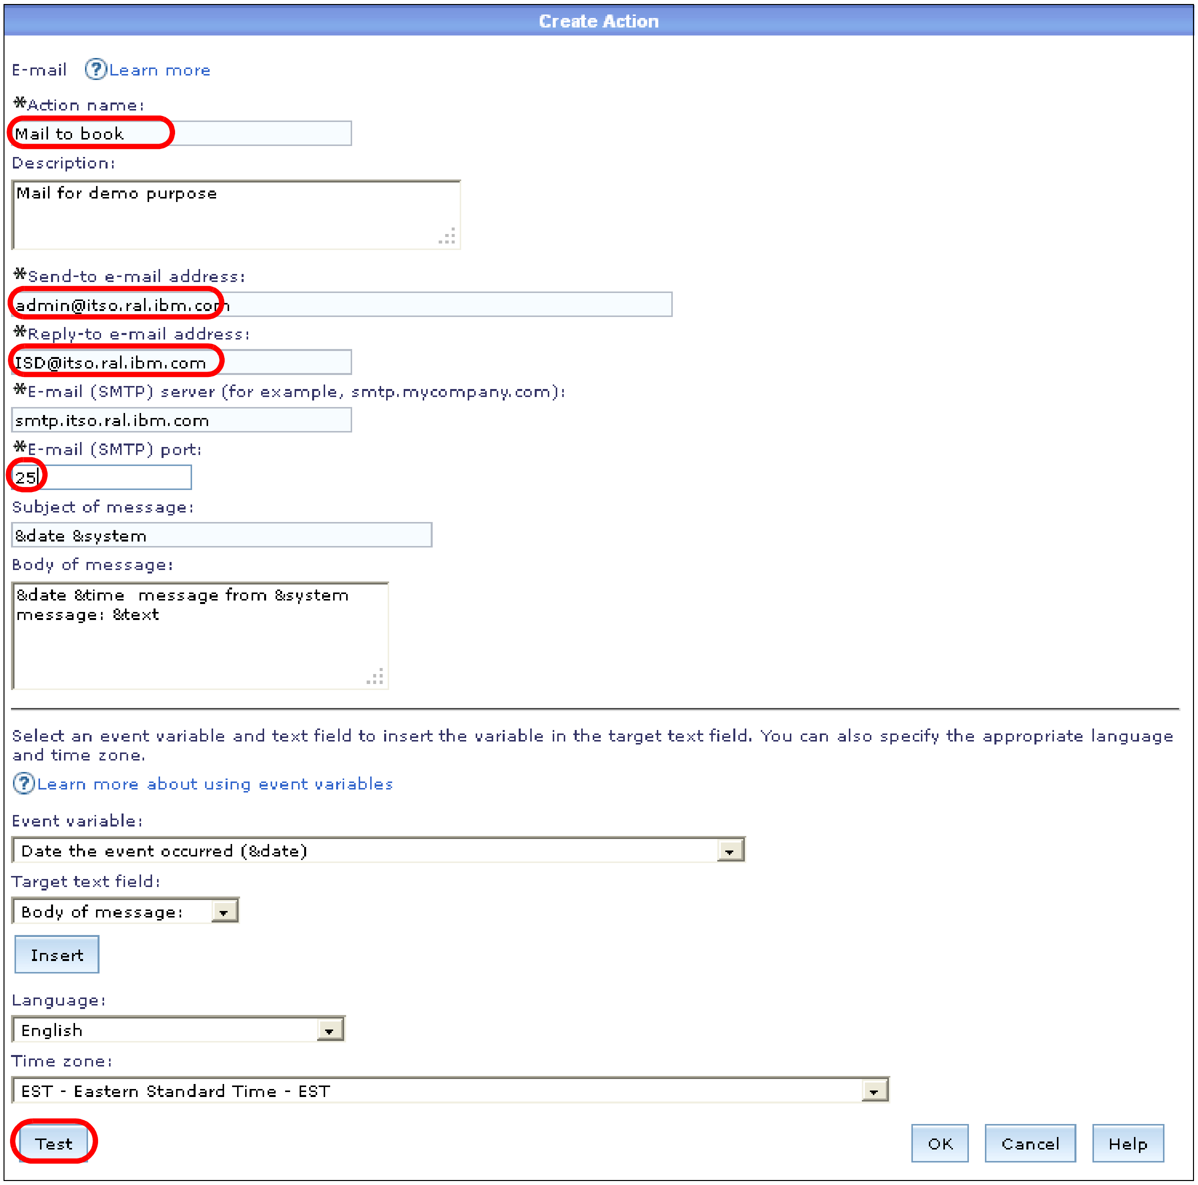

*Action name: Mail to book

Description: Mail for demo purpose

*Send-to e-mail address: admin@itso.ral.ibm.com

*Reply-to e-mail address: ISD@itso.ral.ibm.com

*E-mail (SMTP) server (for example, smtp.mycompany.com): smtp.itso.ral.ibm.com

*E-mail (SMTP) port: 25

Subject of message: &date &system

Body of message:

```
&date &time  message from &system
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: Body of message:

Language: English

Time zone: EST - Eastern Standard Time - EST

Figure 4-137 Configuration window for the “Send an e-mail (internet SMTP)” action

Set the following information for this example:

- Type a name for the action. In our example, we type `Mail to book`.
- Optional: Type a description for the action. We enter `Mail for demo purpose`.
- Enter the send-to email address. We enter `admin@itso.ral.ibm.com`.
- You must enter a Reply-to email address. This email address is listed as the sender of the email. This email address must be in the correct format but the email address does not need to exist. For example, you can use `noreply@example.com`. We use `ISD@itso.ral.ibm.com`.
- Enter the Simple Mail Transfer Protocol (SMTP) server that is used in your environment. We use `smtp.itso.ral.ibm.com`.

- Enter the port that is used by the SMTP server. The standard port for SMTP is port `25`.
- The next entry is for the subject of the message. The default subject line is `&date &system`, which prints the actual date and the system name that sent the event. You can add additional variables or write your own text. The complete list of variables is at this website:

  http://pic.dhe.ibm.com/infocenter/director/pubs/index.jsp?topic=%2Fcom.ibm.director.automation.helps.doc%2Ffqm0_c_ea_event_data_substitution_variables.html

  As an alternative to typing the variables, the window also includes two list boxes, “Event variable” and “Target text field”. Click **Insert** to insert the variable for you.
- Type the body of the message. The default is `&text` (the event text). In our example, we use a combination of text and variables:

  ```
  &date &time message from &system
  message: &text
  ```
- The last two fields specify the language and the time zone. Our example uses `English` and `EST- Eastern Standard Time- EST`.

Test the event action to confirm that the settings are valid. You can view the resulting email by clicking **Test**. The email that we received from our test is shown in Figure 4-138.

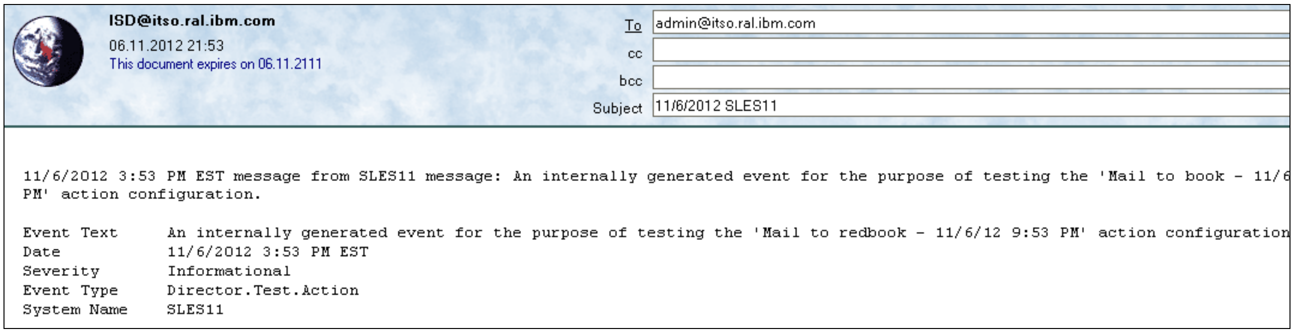

ISD@itso.ral.ibm.com
06.11.2012 21:53
This document expires on 06.11.2111

To admin@itso.ral.ibm.com
cc
bcc
Subject 11/6/2012 SLES11

```
11/6/2012 3:53 PM EST message from SLES11 message: An internally generated event for the purpose of testing the 'Mail to book - 11/6
PM' action configuration.

Event Text     An internally generated event for the purpose of testing the 'Mail to redbook - 11/6/12 9:53 PM' action configuration
Date           11/6/2012 3:53 PM EST
Severity       Informational
Event Type     Director.Test.Action
System Name    SLES11
```

Figure 4-138 Test email from an event action

After you confirm that the event action works correctly, click **OK** to save the changes. The action is then shown in the list of available actions.

Using the CLI for event automation plans

There are several available **`smcli`** commands that you can use to work with event automation plans (Table 4-4).

Table 4-4 Command-line tools for event automation

Command	Purpose
For event automation plans	
`lsevtautopln`	List information about an event automation plan.
`mkevtautopln`	Create an event automation plan.
`rmevtautopln`	Delete one or more event automation plans.
`evtautopln`	Apply one or more event automation plans to a system or a group. Use this command to remove systems or groups from an event automation plan or activate or deactivate an event automation plan.
`chevtautopln`	Change an existing event automation plan.

Command	Purpose
For event filters	
`lsevtfltr`	Display information about an event filter or list all available event filters.
`lsevttype`	List the event types.
`mkevtfltr`	Import event filters.
`rmevtfltr`	Remove event filter.
For event actions	
`lsevtact`	Display information about available event actions or export event actions to an XML file.
`mkevtact`	Import event actions.
`mkevtactemail`	Create a customized event action that sends email over an SMTP server.
`mkevtactstpgm`	Create a customized action that starts a program.
`mkevtacttask`	Create a customized action that starts a non-interactive task.
`rmevtact`	Remove a customized event action.
`testevtact`	Test a customized event action.

For detailed information about these commands and the options for these commands, see the information center:

http://pic.dhe.ibm.com/infocenter/director/pubs/index.jsp?topic=%2Fcom.ibm.director.cli.helps.doc%2Ffqm0_r_cli_automation_cmds.html

Example 4-10 creates the “`e-mail for test`” event action. Example 4-11 on page 192 creates the “`Email for Critical Events`” event automation plan that is used for critical events.

First, we list the available event action on our Systems Director server with the **`smcli lsevtact`** command. Then, we create an email event action with the **`smcli mkevtactemail`** command.

Example 4-10 Create action: Send an email

```
SLES11:/opt/ibm/director/bin#./smcli lsevtact
Add to the event log
Mail to book
eMail
eMail to admin
SLES11:/opt/ibm/director/bin#./smcli mkevtactemail -I -p 25 -s "&date &system" -m
"&date &time message form &system : &text "email for test" admin@itso.ral.ibm.com
ISD@ITSO.ral.ibm.com smtp.itso.ral.ibm.com
SLES11:/opt/ibm/director/bin#./smcli testevtact “email for test”
DNZEAP1073I: <informational> The test or the event action was successfully started
SLES11:/opt/ibm/director/bin#
SLES11:/opt/ibm/director/bin#./smcli lsevtact
Add to event log
Mail to book
eMail
eMail to admin
email for test
SLES11:/opt/ibm/director/bin#
```

We create our “`email for critical events`” event automation plan by using the **`smcli evtautopln`** command. First, we list the available event automation plans. Then, we create an event automation plan by using the “`Critical Events`” filter. Then, we create the “`email for test`” event action. We assign the new EAP to the “`All Systems`” group. After these steps, we list the event automation plans that are available now. You can see that the newly created event automation plan is in the list (Example 4-11).

Example 4-11 Create the event automation plan named “email for critical events”

```
SLES11:/opt/ibm/director/bin#
SLES11:/opt/ibm/director/bin#./smcli lsevtautopln
Log All Events
book-EAP
Send eMail to Admin
Test
SLES11:/opt/ibm/director/bin#
SLES11:/opt/ibm/director/bin# ./smcli mkevtautopln -D "Test" -e "Critical Events"
-x "email for test" -N "All Systems" “email for critical events”
SLES11:/opt/ibm/director/bin#
SLES11:/opt/ibm/director/bin#./smcli lsevtautopln
Log All Events
book-EAP
Send eMail to Admin
Test
email for critical events
SLES11:/opt/ibm/director/bin#
```

If you want to see the details or status for an event automation plan, use the **`smcli lsevtautopln -l “%EAP-Name%”`** command. Example 4-12 lists the detailed information for our newly created “`email for critical events`” event automation plan.

Example 4-12 List details for the “email for critical events” event automation plan

```
SLES11:/opt/ibm/director/bin# ./smcli lsevtautopln -l “email for critical events”
Name: email for critical events
Description: Test
Status: Active
Event Filter: Critical Events
Time Ranges:
    All the time (24x7)
Actions:
    email for test
Targets:
    Group Name: All Systems
SLES11:/opt/ibm/director/bin#
```

4.4.5 Best practice for event management

These are the best practices to use the event management efficiently:

- Use the health status to check your environment and to check the events.
- Create at minimum one event action plan, which contains an *event action* and an *event filter*.

- Use the appropriate action for your environment, which fits your requirements (like email notification). Define the settings for the event action carefully and test the action.
- Select the appropriate event filter or define your own filter that fits in your requirements.
- Assign the event automation plan to single systems or specific groups instead of all systems. This helps to get the right events for the right systems.

4.5 Hardware Management Console and AIX Launch-in-Context

Information about the discovery of a Hardware Management Console (HMC) and the HMC managed resources is described. The HMC Launch-in-Context (LiC) capability and extended tasks are shown.

The Systems Director server can be used for a wide range of tasks on systems that are under the control of a managed HMC:

- Creating virtual servers
- Editing virtual server resources
- Views
- Topology

Discovering an HMC by using the Systems Director server offers a *single-pane-of-glass* view to monitoring and supporting Power Systems hardware. Some dynamic logical partition (LPAR) (DLPAR) functions are embedded in the Systems Director UI, and other functions can be initiated by using LiC.

4.5.1 Before you begin

Before the discovery of the HMC, carefully determine what level of access is passed to the Systems Director server. This determination relates to the tasks that are required to manage the HMC.

Use the following links to the information center to set up:

- Setting up user access to the HMC:

  http://pic.dhe.ibm.com/infocenter/director/pubs/index.jsp?topic=%2Fcom.ibm.director.vim.helps.doc%2Feica7_t_setting_up_user_access_hmc.html

- Configuring the HMC:

  http://pic.dhe.ibm.com/infocenter/director/pubs/index.jsp?topic=%2Fcom.ibm.director.vim.helps.doc%2Feica7_t_configuring_hmc.html

- Managing systems that are controlled by HMC LiC:

  http://pic.dhe.ibm.com/infocenter/director/pubs/index.jsp?topic=%2Fcom.ibm.director.power.helps.doc%2Ffqm0_t_managing_hmc_ivm.html

- Preparing the HMC for discovery:

  http://pic.dhe.ibm.com/infocenter/director/pubs/index.jsp?topic=%2Fcom.ibm.director.install.helps.doc%2Ffqm0_t_preparing_hmc_for_discovery.html

4.5.2 Discovery

Discovery of the HMC is performed by way of the normal discovery process. Resources are discovered by either an IP address or host name.

The following steps show the discovery process:

1. From the UI on Systems Director, select **Inventory** → **System Discovery** (Figure 4-139).

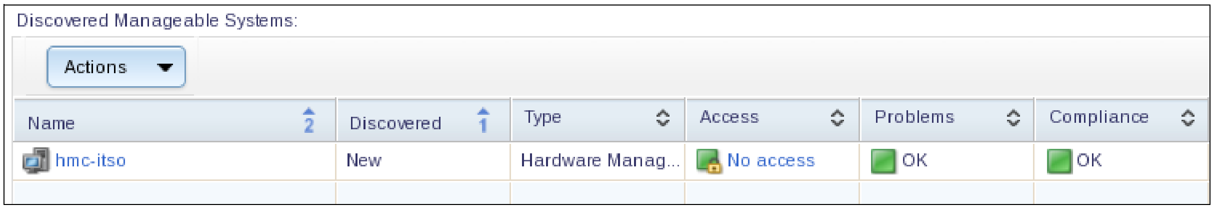

Discovered Manageable Systems:

Name	Discovered	Type	Access	Problems	Compliance
hmc-itso	New	Hardware Manag...	No access	🟩 OK	🟩 OK

Figure 4-139 Discovered appliances

2. Authenticate with the HMC by using a user ID, as shown in Figure 4-140.

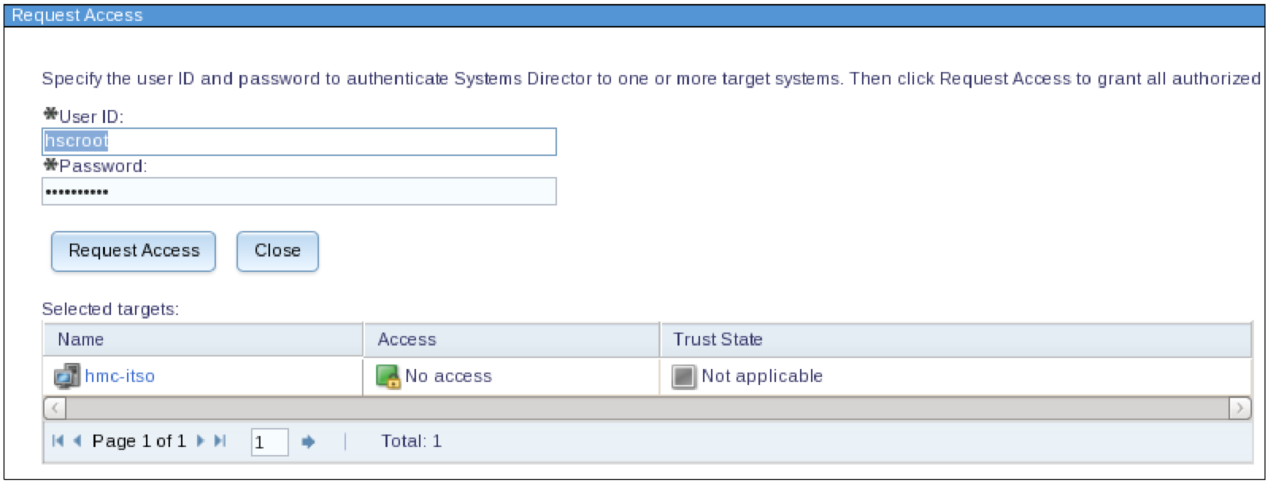

Request Access

Specify the user ID and password to authenticate Systems Director to one or more target systems. Then click Request Access to grant all authorized

*User ID: hscroot

*Password: ••••••••••

Selected targets:

Name	Access	Trust State
hmc-itso	No access	Not applicable

Total: 1

Figure 4-140 Request Access panel

3. If the HMC user access is set up correctly and the correct settings are enabled on the HMC network settings, the access column displays “OK”.

4. To view the recently discovered HMC and associated Power Systems, use the Systems Director UI and click **Inventory** → **Views** → **Platform Managers and Members**.

> **Terminology:** For a list of the terminology that is used in Systems Director for Power Systems users, see this website:
>
> `http://pic.dhe.ibm.com/infocenter/director/pubs/index.jsp?topic=%2Fcom.ibm.director.power.helps.doc%2Ffqm0_c_power_new_terms.html`

5. The view is automatically populated with the resources that are visible to the user ID on the HMC (Figure 4-141).

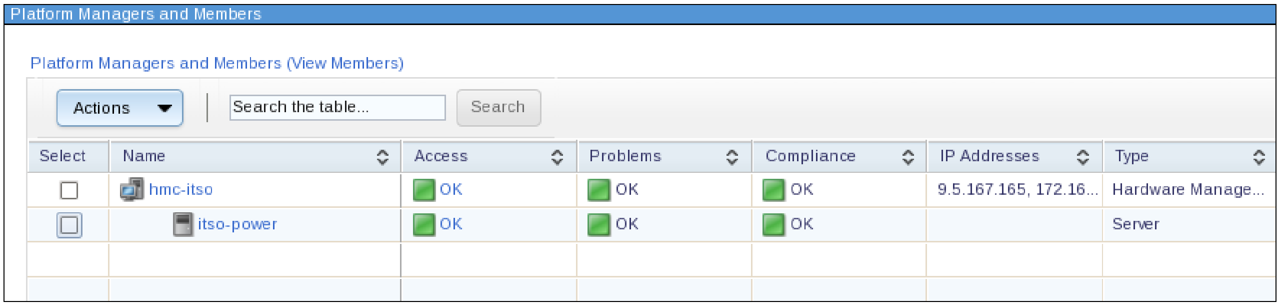
Platform Managers and Members
Platform Managers and Members (View Members)
Actions
Search the table...
Search

Select	Name	Access	Problems	Compliance	IP Addresses	Type
☐	hmc-itso	🟩 OK	🟩 OK	🟩 OK	9.5.167.165, 172.16...	Hardware Manage...
☐	itso-power	🟩 OK	🟩 OK	🟩 OK		Server

Figure 4-141 Platform Managers and Members panel

6. From the Systems Director UI, see an expanded view for Power Systems servers by clicking **Inventory** → **Views** → **Virtual Servers and Hosts**, as shown in Figure 4-142.

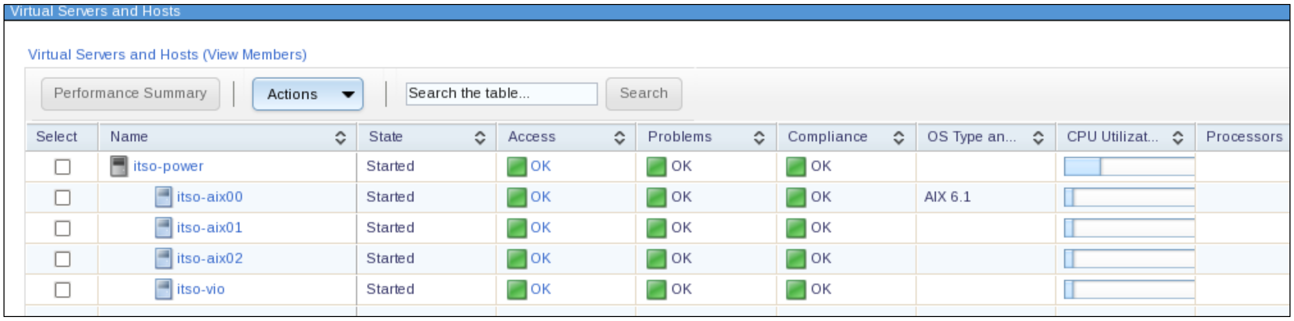
Virtual Servers and Hosts
Virtual Servers and Hosts (View Members)
Performance Summary
Actions
Search the table...
Search

Select	Name	State	Access	Problems	Compliance	OS Type an...	CPU Utilizat...	Processors
☐	itso-power	Started	🟩 OK	🟩 OK	🟩 OK			
☐	itso-aix00	Started	🟩 OK	🟩 OK	🟩 OK	AIX 6.1		
☐	itso-aix01	Started	🟩 OK	🟩 OK	🟩 OK			
☐	itso-aix02	Started	🟩 OK	🟩 OK	🟩 OK			
☐	itso-vio	Started	🟩 OK	🟩 OK	🟩 OK			

Figure 4-142 Virtual Servers and Hosts panel

7. Because the HMC is discovered, we get visibility to the HMC menu by Launch-in-Context, as shown in Figure 4-143.

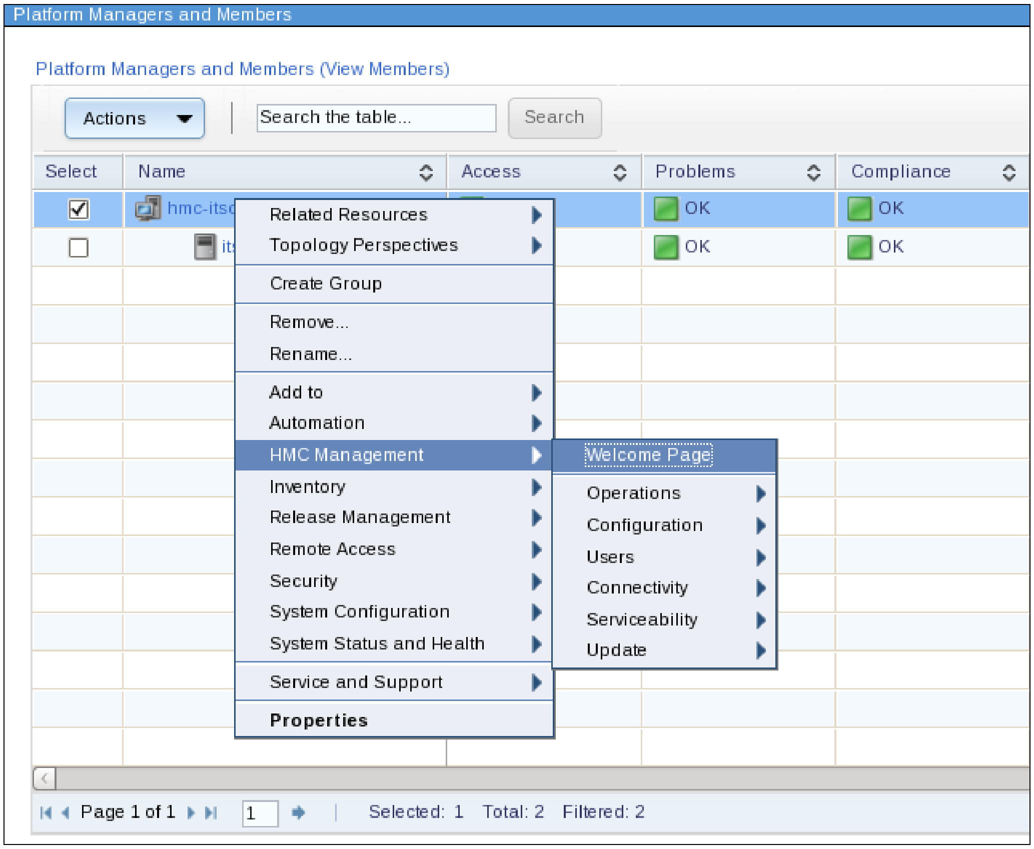

Select	Name	Access	Problems	Compliance
☑	hmc-itsc		🟩 OK	🟩 OK
☐	its		🟩 OK	🟩 OK

Selected: 1 Total: 2 Filtered: 2

Figure 4-143 HMC menu

8. If we select the Welcome Page from the HMC menu, it does not launch (Figure 4-144). We still do not have full authentication to use the HMC and need to configure single sign-on (SSO).

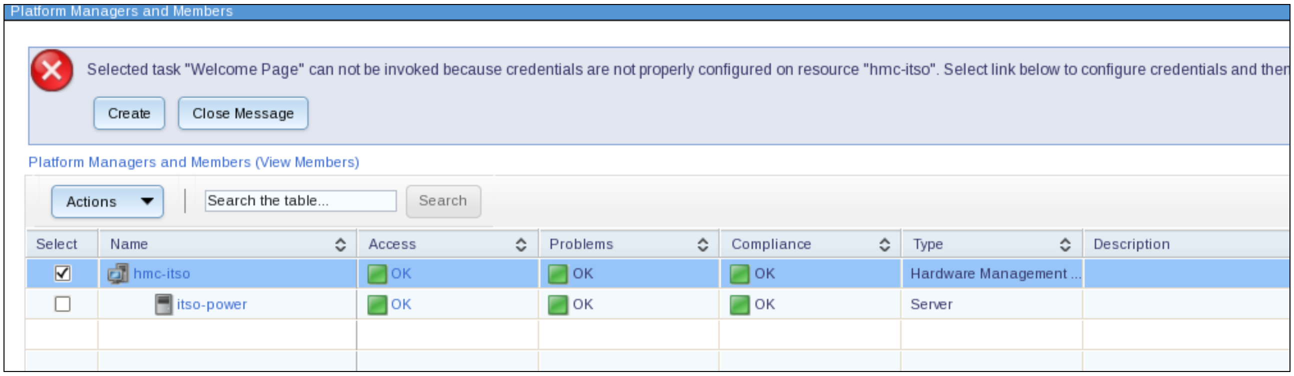
Platform Managers and Members
Create
Close Message
Platform Managers and Members (View Members)
Actions
Search the table...
Search

Selected task "Welcome Page" can not be invoked because credentials are not properly configured on resource "hmc-itso". Select link below to configure credentials and then

Select	Name	Access	Problems	Compliance	Type	Description
☑	hmc-itso	🟩 OK	🟩 OK	🟩 OK	Hardware Management ...	
☐	itso-power	🟩 OK	🟩 OK	🟩 OK	Server	

Figure 4-144 Configure SSO

9. Click **Create**. The Welcome wizard opens to enter valid SSO credentials (Figure 4-145).

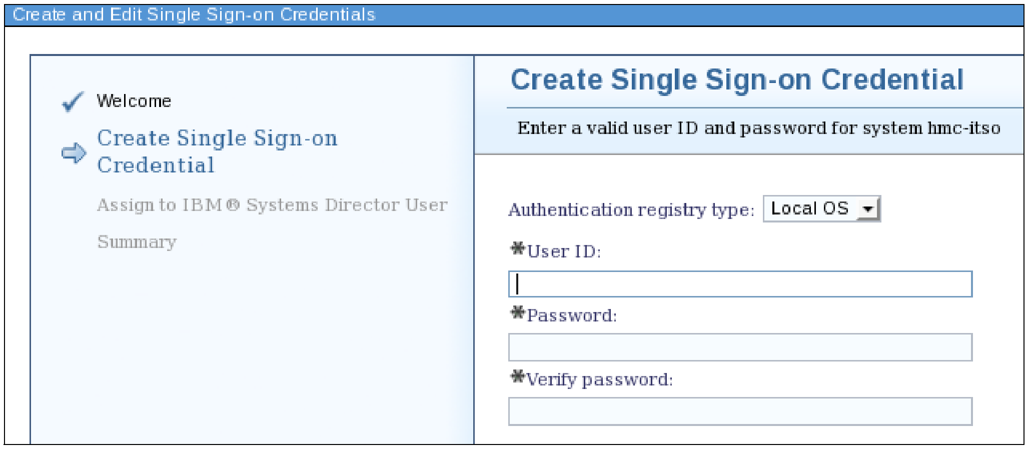

Create and Edit Single Sign-on Credentials

✓ Welcome

Create Single Sign-on Credential

Assign to IBM® Systems Director User

Summary

Create Single Sign-on Credential

Enter a valid user ID and password for system hmc-itso

Authentication registry type: Local OS

*User ID:

*Password:

*Verify password:

Figure 4-145 Create and edit SSO

10. Enter a user ID and password. Then, click **Next**.

11. "Assign to IBM Systems Director User", as shown in Figure 4-146.

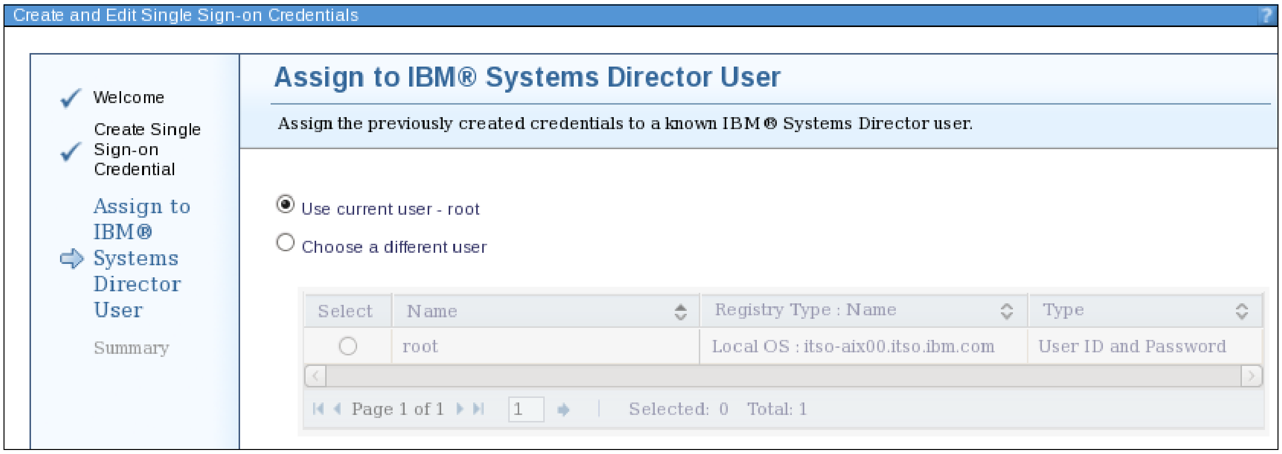

Create and Edit Single Sign-on Credentials

✓ Welcome

✓ Create Single Sign-on Credential

Assign to IBM® Systems Director User

Summary

Assign to IBM® Systems Director User

Assign the previously created credentials to a known IBM® Systems Director user.

● Use current user - root

○ Choose a different user

Select	Name	Registry Type : Name	Type
○	root	Local OS : itso-aix00.itso.ibm.com	User ID and Password

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

Figure 4-146 Assign to IBM Systems Director User

12. At the end of the wizard, click **Finish**.

13. Because the SSO for the HMC that you selected is configured, reattempt Launch-in-Context (Figure 4-143 on page 197). The HMC does not launch and you are brought to the requested page (Figure 4-147).

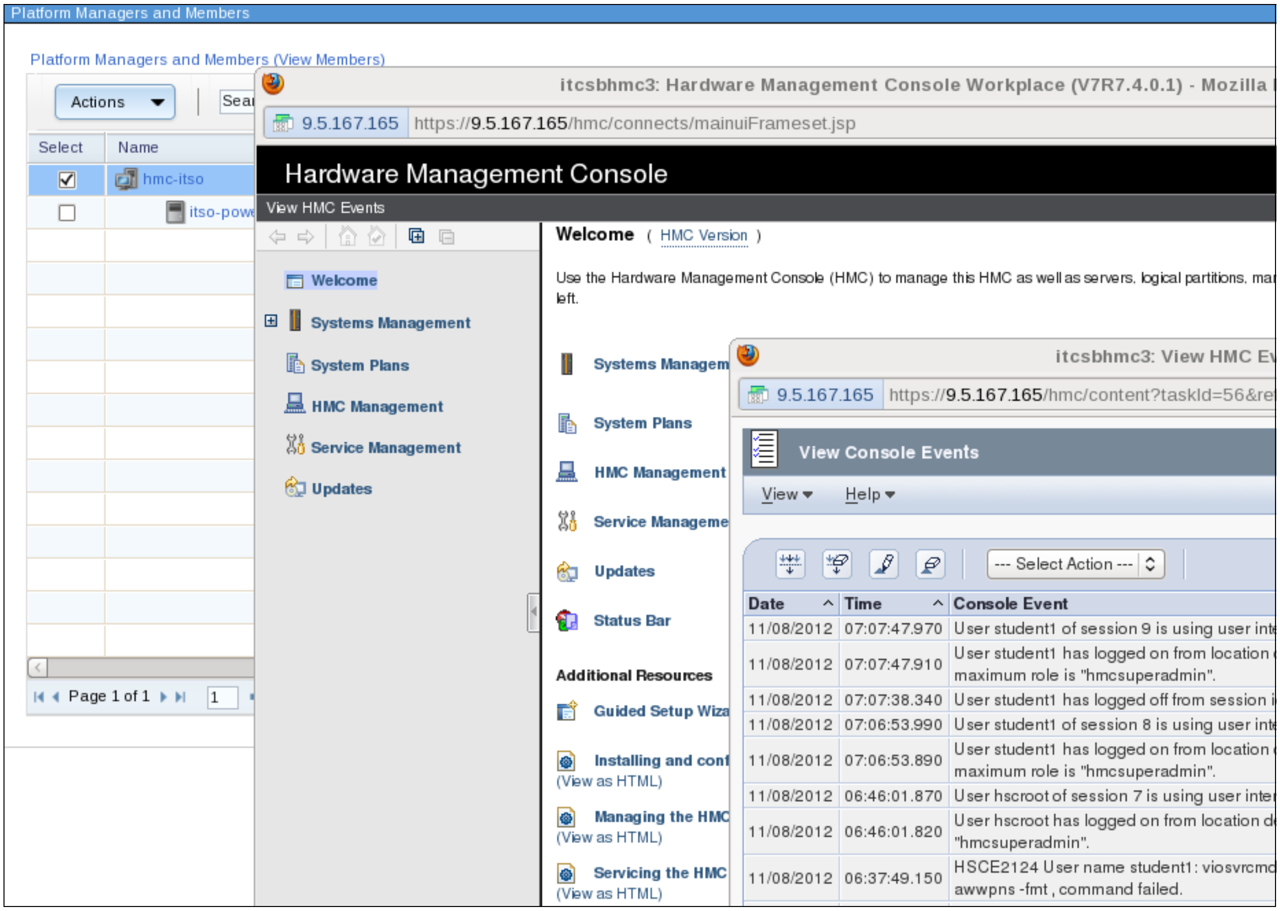

Use the Hardware Management Console (HMC) to manage this HMC as well as servers, logical partitions, mar left.

Date	Time	Console Event
11/08/2012	07:07:47.970	User student1 of session 9 is using user inte
11/08/2012	07:07:47.910	User student1 has logged on from location maximum role is "hmcsuperadmin".
11/08/2012	07:07:38.340	User student1 has logged off from session i
11/08/2012	07:06:53.990	User student1 of session 8 is using user inte
11/08/2012	07:06:53.890	User student1 has logged on from location maximum role is "hmcsuperadmin".
11/08/2012	06:46:01.870	User hscroot of session 7 is using user inter
11/08/2012	06:46:01.820	User hscroot has logged on from location d "hmcsuperadmin".
11/08/2012	06:37:49.150	HSCE2124 User name student1: viosvrcmd awwpns -fmt , command failed.

Figure 4-147 View Console Events panel

You can create a partition from the UI in two ways:

- Use LiC and launch the HMC UI (Figure 4-148 on page 200)
- Use the embedded view from the UI (Figure 4-150 on page 202)

By using LiC, you can also label the profile and assign these properties:

- Minimum processing units and memory
- Processing units and memory that you want
- Maximum processing units and memory
- Weighting for processing units

By using the standard LiC functionality, you can add virtual adapters while you build the virtual server. Depending on the complexity of the virtual server that you create, LiC gives you greater choice.

To create a virtual server by using LiC, follow these steps:

1. Right-click the server and click **Extended Management** → **Configuration** → **Create Logical Partition**. You can select to create an AIX or Linux partition, or an IBM i partition, as shown in Figure 4-148.

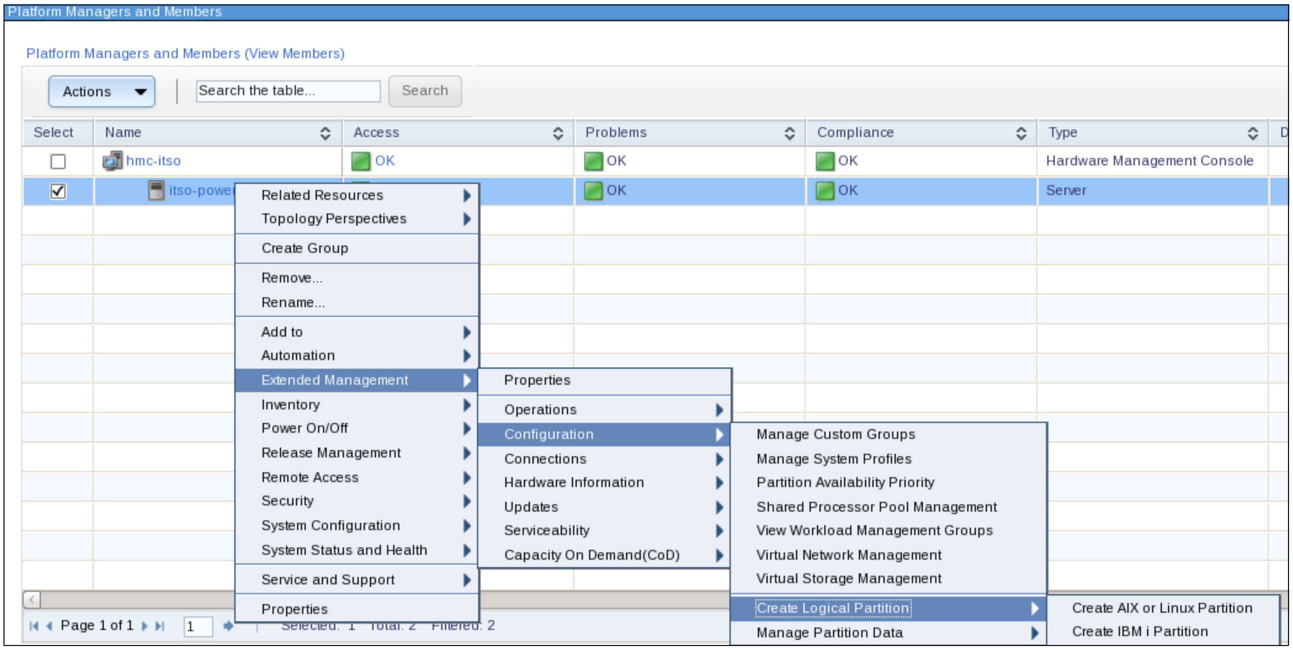

Platform Managers and Members

Platform Managers and Members (View Members)

Select	Name	Access	Problems	Compliance	Type	D
☐	hmc-itso	🟩 OK	🟩 OK	🟩 OK	Hardware Management Console	
☑	itso-powe		🟩 OK	🟩 OK	Server	

Figure 4-148 LiC virtual server creation

2. The HMC wizard opens in context (Figure 4-149).

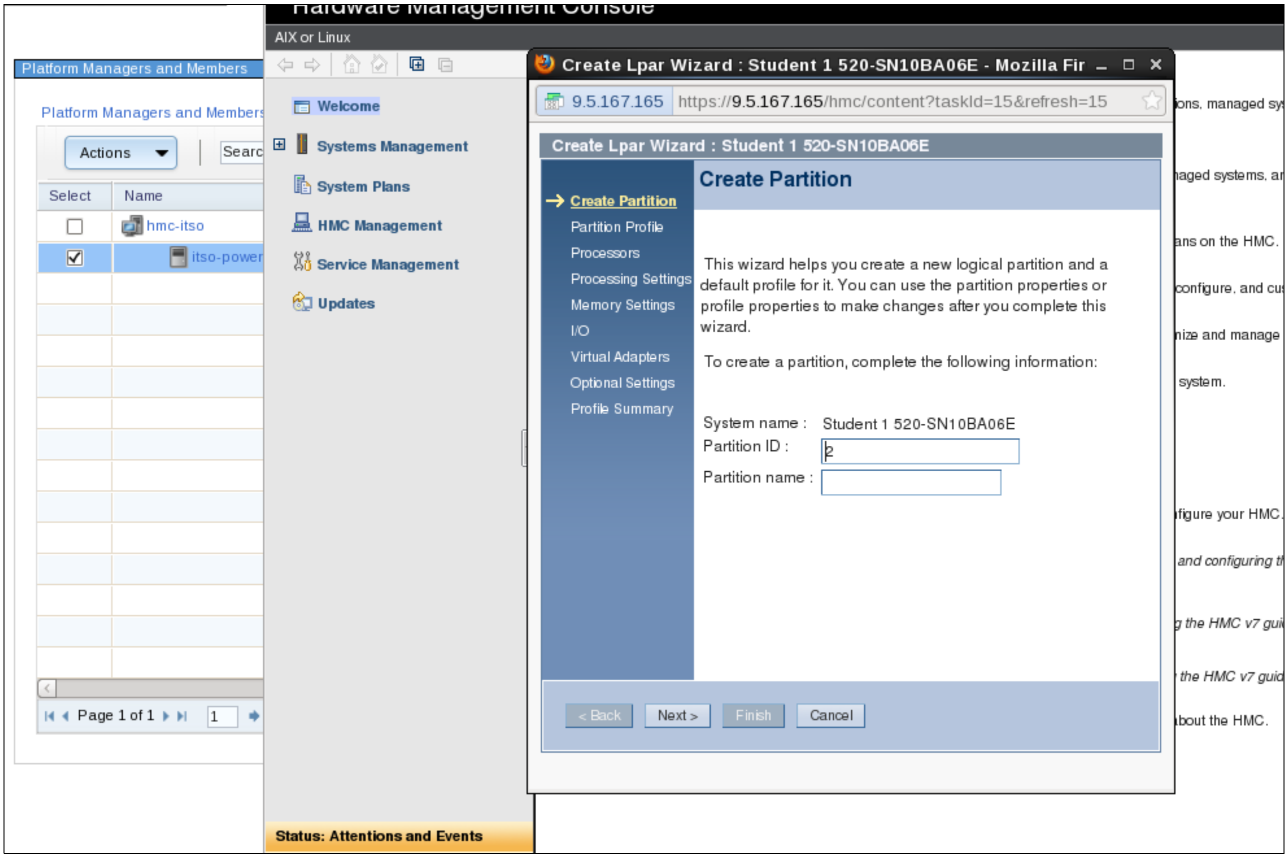

Create Lpar Wizard : Student 1 520-SN10BA06E - Mozilla Fir

https://9.5.167.165/hmc/content?taskId=15&refresh=15

Create Lpar Wizard : Student 1 520-SN10BA06E

Create Partition

This wizard helps you create a new logical partition and a default profile for it. You can use the partition properties or profile properties to make changes after you complete this wizard.

To create a partition, complete the following information:

System name : Student 1 520-SN10BA06E

Partition ID : 2

Partition name :

Figure 4-149 HMC LiC

Using the embedded menu simplifies the creation of the virtual server because you do not leave the UI, as shown in Figure 4-150. Follow these steps:

1. Right-click the server and click **System Configuration** → **Create Virtual Server**.

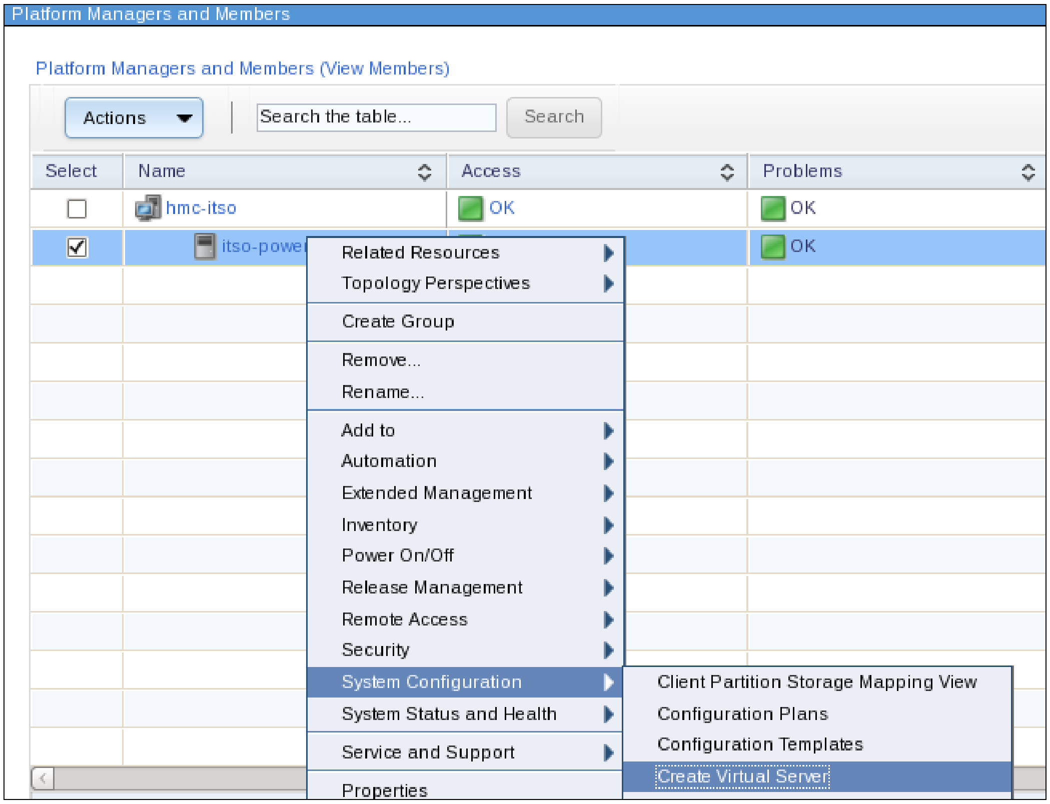

Platform Managers and Members

Platform Managers and Members (View Members)

Select	Name	Access	Problems
☐	hmc-itso	🟩 OK	🟩 OK
☑	itso-powe		🟩 OK

Figure 4-150 Virtual server creation embedded

2. The Create Virtual Server wizard starts. As shown in Figure 4-151, specify properties, such as name and source (AIX or Linux, and processors).

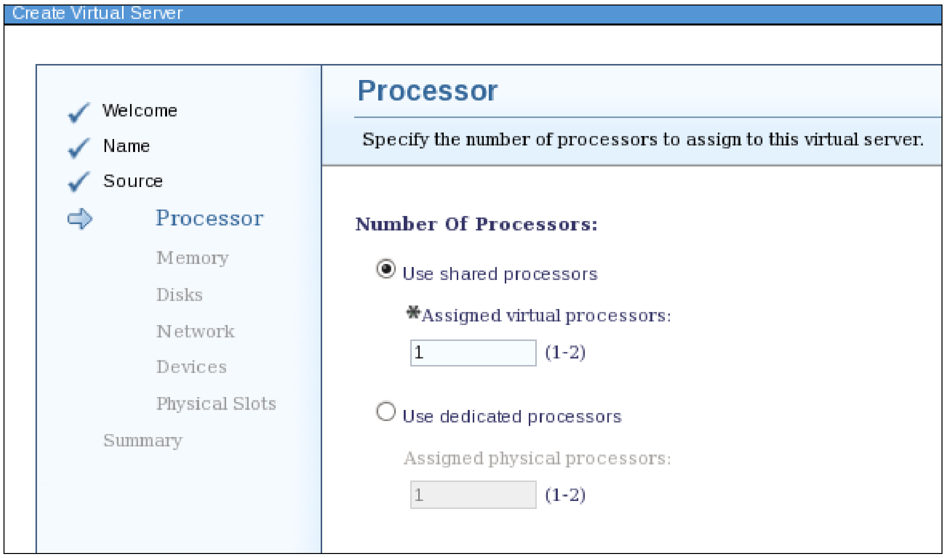

Create Virtual Server

Processor

Specify the number of processors to assign to this virtual server.

Number Of Processors:

● Use shared processors

*Assigned virtual processors: 1 (1-2)

○ Use dedicated processors

Assigned physical processors: 1 (1-2)

Figure 4-151 Create Virtual Server wizard

3. Specify the memory size as required, as shown in Figure 4-152.

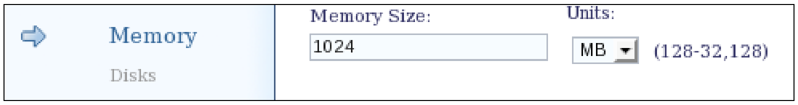

Memory Size: 1024

Units: MB (128-32,128)

Figure 4-152 Memory

4. Disks depend on storage visibility or whether disks are already mapped to the Virtual I/O Server (VIOS) (Figure 4-153). Disks can be virtual or physical system service processors (SSPs) or N-Port ID Virtualization (NPIV) volumes.

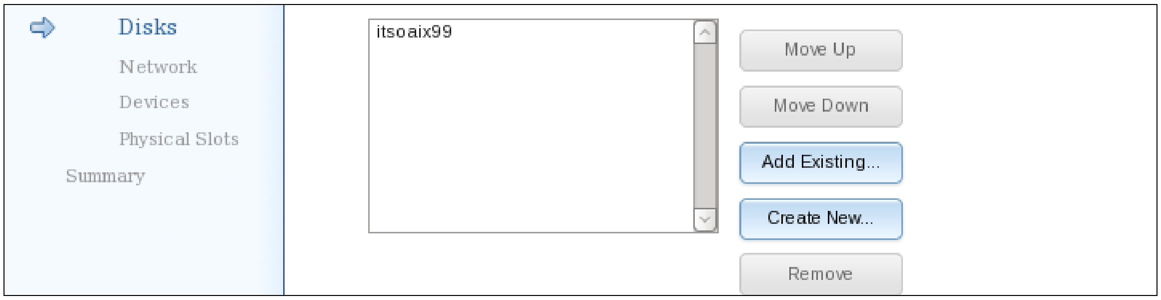

itsoaix99

Figure 4-153 Disks

5. Network definitions are shown in Figure 4-154.

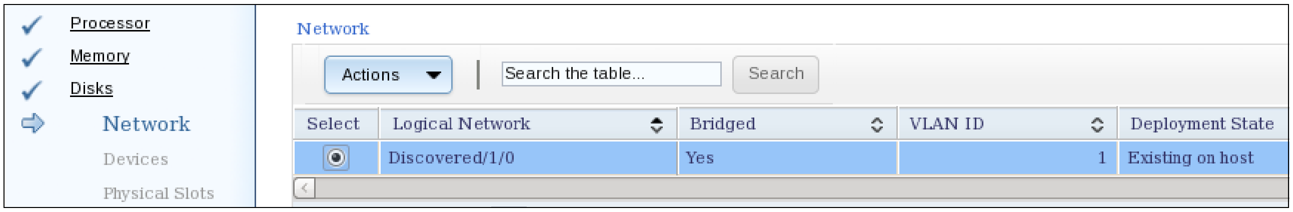

Select	Logical Network	Bridged	VLAN ID	Deployment State
●	Discovered/1/0	Yes	1	Existing on host

Figure 4-154 Networks

6. For devices, use the normal practices that you use for a regular HMC LPAR creation (Figure 4-155).

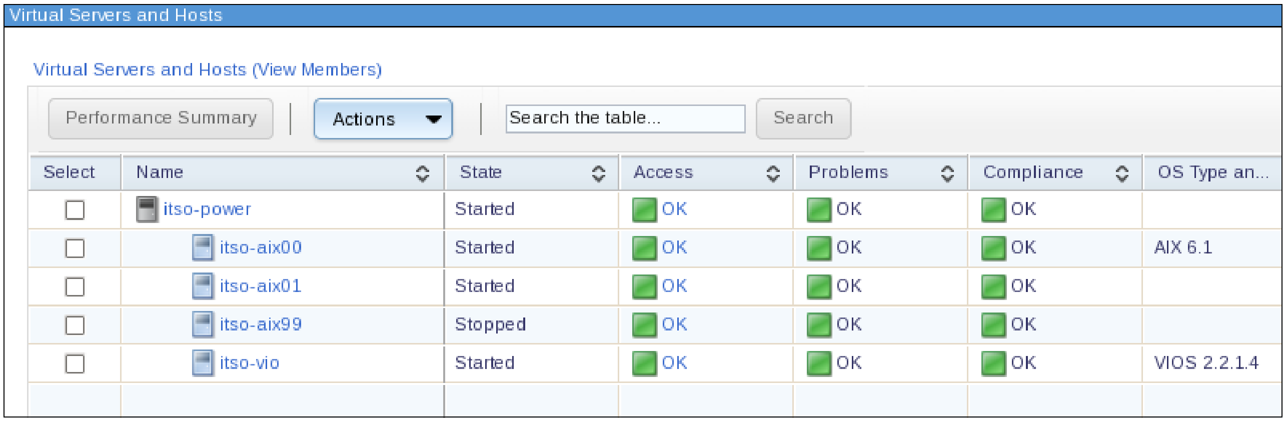
Virtual Servers and Hosts
Virtual Servers and Hosts (View Members)
Performance Summary
Actions
Search the table...
Search

Select	Name	State	Access	Problems	Compliance	OS Type an...
□	itso-power	Started	🟩 OK	🟩 OK	🟩 OK	
□	itso-aix00	Started	🟩 OK	🟩 OK	🟩 OK	AIX 6.1
□	itso-aix01	Started	🟩 OK	🟩 OK	🟩 OK	
□	itso-aix99	Stopped	🟩 OK	🟩 OK	🟩 OK	
□	itso-vio	Started	🟩 OK	🟩 OK	🟩 OK	VIOS 2.2.1.4

Figure 4-155 itso-aix99 completed

7. After the creation of a virtual server, additional networks can be added and CPU priority can be changed (Figure 4-156). Use this syntax:

   ```
   smcli chvs -A "networks=+Discovered-XX-0" -n itso-aix99
   ```

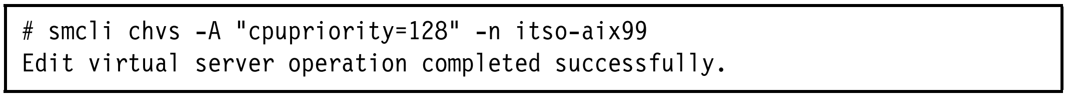

```
# smcli chvs -A "cpupriority=128" -n itso-aix99
Edit virtual server operation completed successfully.
```

Figure 4-156 chvs cpupriority

8. You can increase or decrease memory within the virtual server minimum and maximum range (Figure 4-157 and Figure 4-158).

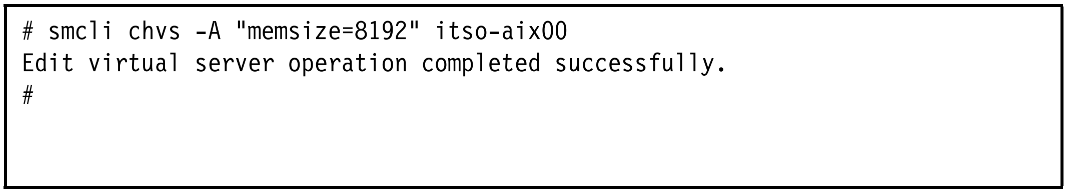

```
# smcli chvs -A "memsize=8192" itso-aix00
Edit virtual server operation completed successfully.
#
```

Figure 4-157 chvs memsize

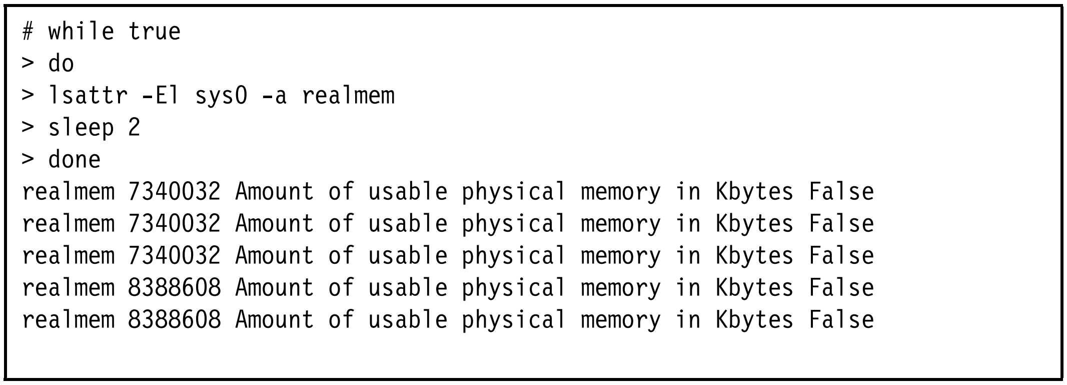

```
# while true
> do
> lsattr -El sys0 -a realmem
> sleep 2
> done
realmem 7340032 Amount of usable physical memory in Kbytes False
realmem 7340032 Amount of usable physical memory in Kbytes False
realmem 7340032 Amount of usable physical memory in Kbytes False
realmem 8388608 Amount of usable physical memory in Kbytes False
realmem 8388608 Amount of usable physical memory in Kbytes False
```

Figure 4-158 chvs memsize output

9. Memory, processor, and priority can be changed from the UI dynamically one time within the virtual server minimum and maximum range (Figure 4-159).

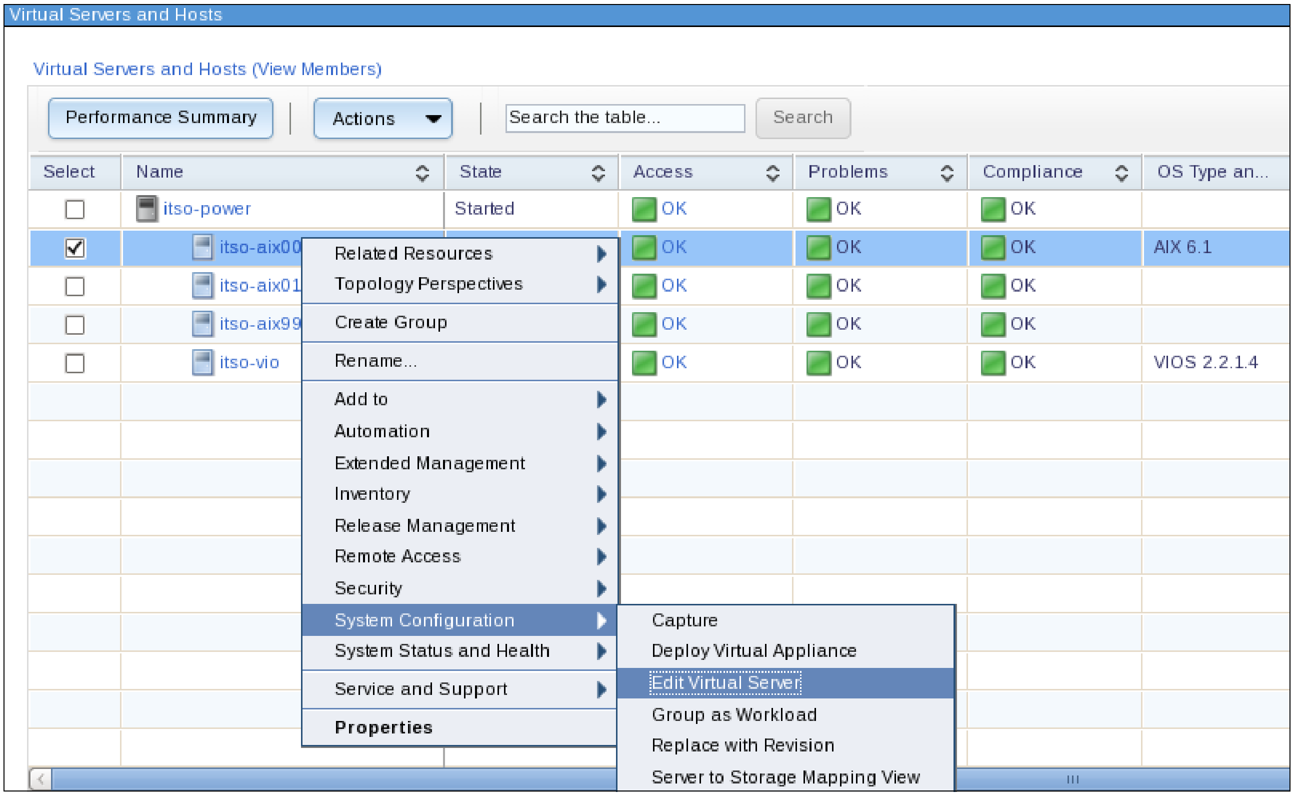

Select	Name	State	Access	Problems	Compliance	OS Type an...
☐	itso-power	Started	🟩 OK	🟩 OK	🟩 OK	
☑	itso-aix00		🟩 OK	🟩 OK	🟩 OK	AIX 6.1
☐	itso-aix01		🟩 OK	🟩 OK	🟩 OK	
☐	itso-aix99		🟩 OK	🟩 OK	🟩 OK	
☐	itso-vio		🟩 OK	🟩 OK	🟩 OK	VIOS 2.2.1.4

Figure 4-159 Edit Virtual Server option

10. Select the required tab and change the “Assigned” value (Figure 4-160). Click **OK**.

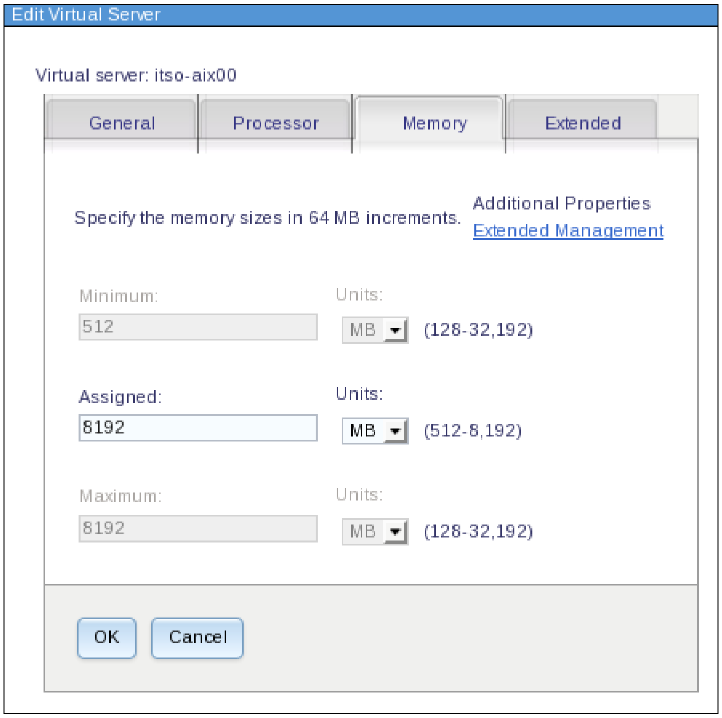

Virtual server: itso-aix00

Specify the memory sizes in 64 MB increments.

Minimum: 512 Units: MB (128-32,192)

Assigned: 8192 Units: MB (512-8,192)

Maximum: 8192 Units: MB (128-32,192)

Figure 4-160 Edit virtual server memory

11. On the submitted job, click **Display Properties**. Click **Complete (view log)** to view the output of the job. See Figure 4-161.

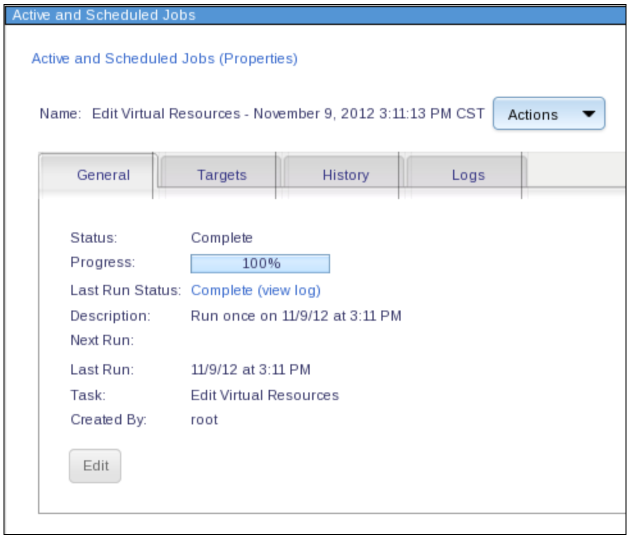

Active and Scheduled Jobs
Active and Scheduled Jobs (Properties)
Actions
General
Targets
History
Logs
Edit

Name: Edit Virtual Resources - November 9, 2012 3:11:13 PM CST

Status: Complete
Progress: 100%
Last Run Status: Complete (view log)
Description: Run once on 11/9/12 at 3:11 PM
Next Run:
Last Run: 11/9/12 at 3:11 PM
Task: Edit Virtual Resources
Created By: root

Figure 4-161 Job properties

12. Figure 4-162 shows the virtual server job output.

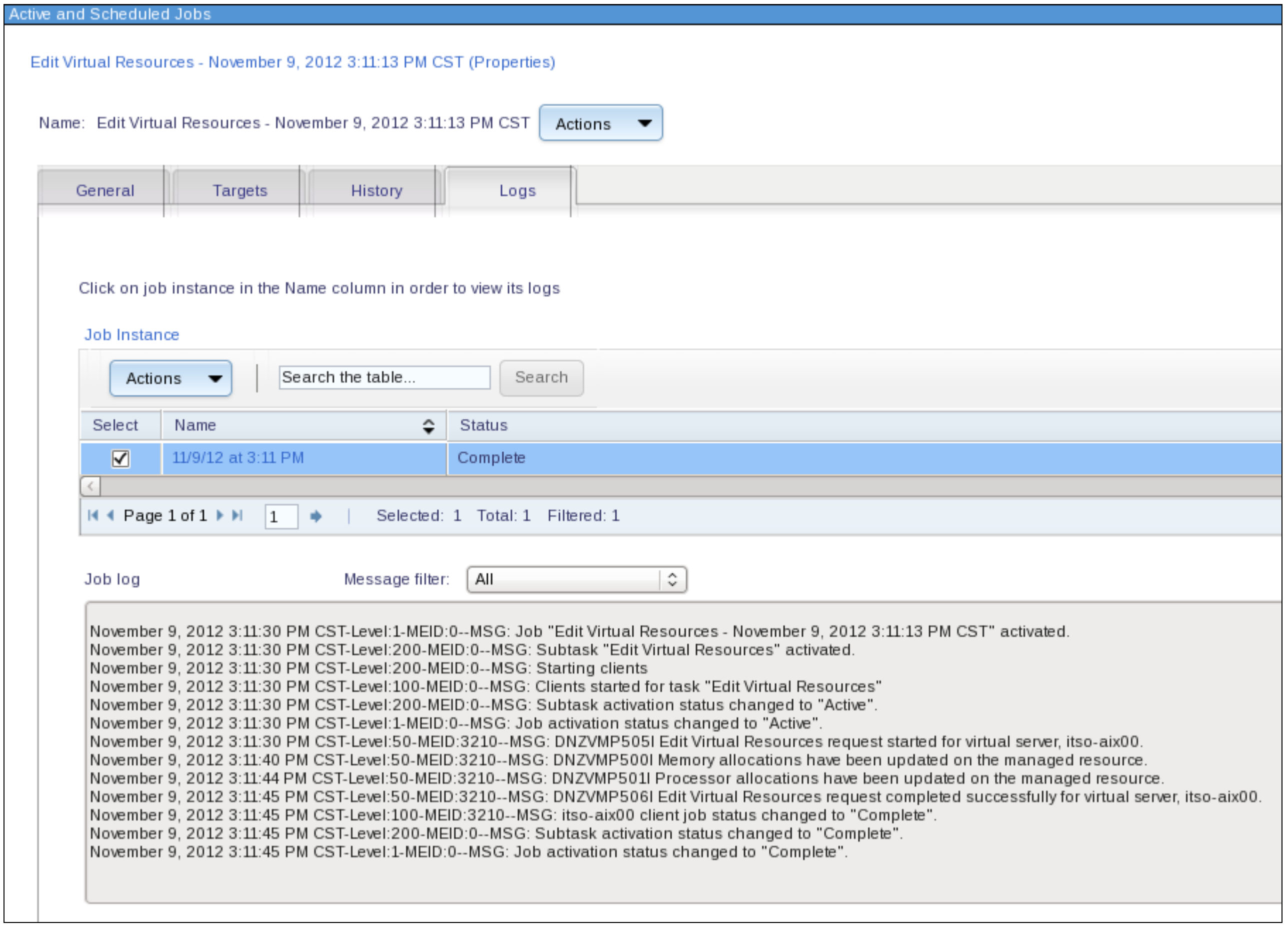

Edit Virtual Resources - November 9, 2012 3:11:13 PM CST (Properties)

Name: Edit Virtual Resources - November 9, 2012 3:11:13 PM CST

Click on job instance in the Name column in order to view its logs

Job Instance

Search the table...

Select	Name	Status
☑	11/9/12 at 3:11 PM	Complete

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

Job log   Message filter: All

```
November 9, 2012 3:11:30 PM CST-Level:1-MEID:0--MSG: Job "Edit Virtual Resources - November 9, 2012 3:11:13 PM CST" activated.
November 9, 2012 3:11:30 PM CST-Level:200-MEID:0--MSG: Subtask "Edit Virtual Resources" activated.
November 9, 2012 3:11:30 PM CST-Level:200-MEID:0--MSG: Starting clients
November 9, 2012 3:11:30 PM CST-Level:100-MEID:0--MSG: Clients started for task "Edit Virtual Resources"
November 9, 2012 3:11:30 PM CST-Level:200-MEID:0--MSG: Subtask activation status changed to "Active".
November 9, 2012 3:11:30 PM CST-Level:1-MEID:0--MSG: Job activation status changed to "Active".
November 9, 2012 3:11:30 PM CST-Level:50-MEID:3210--MSG: DNZVMP505I Edit Virtual Resources request started for virtual server, itso-aix00.
November 9, 2012 3:11:40 PM CST-Level:50-MEID:3210--MSG: DNZVMP500I Memory allocations have been updated on the managed resource.
November 9, 2012 3:11:44 PM CST-Level:50-MEID:3210--MSG: DNZVMP501I Processor allocations have been updated on the managed resource.
November 9, 2012 3:11:45 PM CST-Level:50-MEID:3210--MSG: DNZVMP506I Edit Virtual Resources request completed successfully for virtual server, itso-aix00.
November 9, 2012 3:11:45 PM CST-Level:100-MEID:3210--MSG: itso-aix00 client job status changed to "Complete".
November 9, 2012 3:11:45 PM CST-Level:200-MEID:0--MSG: Subtask activation status changed to "Complete".
November 9, 2012 3:11:45 PM CST-Level:1-MEID:0--MSG: Job activation status changed to "Complete".
```

Figure 4-162 Edit virtual server job output

13. Under the Systems Director UI, the **Inventory** → **Views** → **Virtual Servers and Hosts** view displays the changed memory value (Figure 4-163).

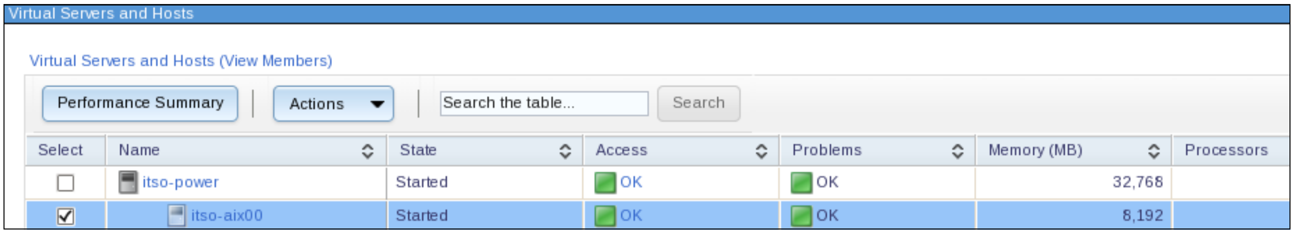

Virtual Servers and Hosts (View Members)

Search the table...

Select	Name	State	Access	Problems	Memory (MB)	Processors
□	itso-power	Started	🟩 OK	🟩 OK	32,768	
☑	itso-aix00	Started	🟩 OK	🟩 OK	8,192	

Figure 4-163 Updated server configuration

4.5.3 Best practices for HMC management

These are the best practices to use the HMC LiC efficiently:

- Properly configure user, access, and rights on the HMC for IBM Systems Director servers:
  - Read documentation
  - Discover the HMC
  - Create user and sufficient rights on the HMC
  - Configure access and single sign-on on IBM Systems Director servers to access the HMC
  - Run an inventory on the HMC

4.6 Security

Systems Director security is controlled by two interdependent processes: authentication and authorization.

Authentication is used to determine who can access the Systems Director server. *Authorization* determines the resources to which the user has access. Systems Director uses role-based access control (RBAC) where the administrator assigns roles and permissions to an authenticated user. On that basis, the user can work on resources that are based on the RBAC to which the user is assigned.

The security features of Systems Director enable an administrator to perform the following functions:

- Manage auditing
- View and manage authorized users and groups
- Assign roles and resources to users
- Manage user properties
- Create and modify roles
- Manage permissions that are grouped within a role
- Use roles to control access to a system
- Request access to a system
- Manage credentials and their associated mappings

The following flow allows a user to access or manage a system:

1. User must be authenticated.
2. User must be authorized to perform a task on the selected resource.

4.6.1 Users and groups for authentication

In Systems Director, users and user groups are based on users and groups that are defined in the configured registry. The registry is associated with either the operating system, directory services, such as Lightweight Directory Access Protocol (LDAP), or the domain controller. Systems Director uses the user and group information for authentication and authorization.

Access to particular resources or tasks is governed by restrictions. The restrictions are based on the user ID or user group membership and the roles that are defined for each user. For a user to access the Systems Director server, one of the following conditions must exist:

- The user is a member of a user group that is authorized for the Systems Director server.
- The user has administrator privileges on the Windows management server or Windows domain.
- The user is a root user on the AIX or Linux management server.

In a default Systems Director server installation scenario that uses the local operating system registry, four Systems Director user groups are automatically created. The user groups are created at the operating system level on the management server.

Table 4-5 lists the user groups, which are used for different access permissions to the Systems Director server.

Table 4-5 Default groups

Default groups	Role	Description
smadmin	SMAdministrator	Administrator group. Users in this group have administrative access to Systems Director and can perform all administrative tasks. These members can define the available privileges for the smmgr, smmon, smuser, and groupread groups. The privileges that are available to members of the smadmin group cannot be restricted.
smmgr	SMManager	Manager group. The supported operations are a subset of the SMAdministrator group. The members of this group have all rights except the rights to create or change user permissions and authorizations.
smmon	SMMonitor	Monitor group. This group supports some administrative functions, such as monitoring. The members of this group are restricted to read-only functionality.
smuser	SMUser	User group. Members of this group have, by default, no rights and no access to any system or functions.

Members of the root and Administrator group are authorized for all operations on all resources.

The only role that is automatically assigned is to the administrator user ID that installed Systems Director. So, initially, no other user is associated with a role.

If you want to use LDAP or another directory service that the user registry supports, you might need to manually create all the user groups and assign users to them.

The users for Systems Director must be added to one of the groups to get access to the Systems Director GUI.

Authenticating a local user

Systems Director can authenticate user login requests to the registry for the configured operating system. Systems Director uses the local operating system user registry by default.

Follow these steps to create a local operating system user account:

1. Create a user account in the user registry that is associated with the management server. The way that you create the user account depends on the operating system that you use.
2. Add the user as a member of one of the user groups that are defined for Systems Director at the user registry level. You can either use one of the predefined groups or create your own groups. If you create a custom group on the Systems Director server, you must authorize it. (Log in as a member of the smadmin group and then go to **Security** → **Users** → **Authorize Groups**. Or, use the `smcli authusergp` command).
3. Log in to the Systems Director web interface as a member of the smadmin group and go to **Security** → **Users**. The users that you configured in the previous steps are displayed in the list.

After users are authenticated to Systems Director, you can configure the authorizations for each user to Systems Director tasks and resources.

Authenticating a domain user

Systems Director can authenticate a user from an Active Directory domain to access the Systems Director GUI. Use the following steps to generate the access for the domain user:

1. Create a user account in the Active Directory user registry. For instructions about creating a user account in the domain server user registry, see the Active Directory documentation.
2. Add the Active Directory user to a defined Active Directory global security group. You must create your own Active Directory group if a suitable group does not exist.
3. Add the global group to an authorized local group of the Systems Director server, such as smadmin, smmgr, smmon, or smuser.

   Systems Director works best with Active Directory when its users are placed in global groups. Those global groups are then placed in the local groups of the Systems Director server.

   For preferred practices, do not add Active Directory users directly to the local groups of any Systems Director servers.
4. Log in to the Systems Director web interface as an administrator and go to **Security** → **Users**. Active Directory users that are managed as a group do not appear in the list. However, you see the group. Users that are local to the Systems Director server show on this list because they are managed as individuals.

You can now assign additional roles to users to access specific Systems Director tasks and resources.

Authenticating LDAP users

Systems Director can authenticate user login requests to an LDAP server. In section 4.6.5, “Lightweight Directory Access Protocol” on page 221, you can see an example of how to use LDAP for Systems Director.

4.6.2 Authorizing users

User authorization occurs when an authenticated user uses Systems Director to perform a task on a resource. The authorization mechanism compares the user account, or the group to which the user belongs, to the RBAC settings for that user or group. If a role exists that contains the necessary authorizations 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.

Roles

You can assign roles to Systems Director users to control their access to resources and limit the tasks that they can perform on those resources. The authorities that you configure for a role determine the level of access that is granted to each user who is assigned to that role. All users or groups of users that access Systems Director must have a user role assignment.

The Systems Director server uses an RBAC service with which an administrator can create custom sets of permissions. The administrator assigns these sets of permissions, which are known as *roles*, to individual users or groups. An *authorization role* is a set of tasks, CLI commands, and application permissions that is applied to one or more resources. Each role can be applied to many users, and each user can have many roles. Regulating user roles is an effective way to control security for your system. By regulating user roles, you can control access to every task and CLI command.

The following roles are available in the Systems Director server by default:

- SMAdministrator (Administrator role)

  The SMAdministrator role has full authority to perform all tasks and functions and full control over permissions. A user that is assigned to this role can perform all tasks (including security administration, product installation, and configuration) with any resource.

- SMManager (Manager role)

  The SMManager role can perform management operations, which are a subset of the functions that a member of the SMAdministrator role can perform. Typically, system administration, system health management, and system configuration tasks are available. This role cannot perform security administration or security configuration tasks. However, this role has full access to all the Systems Director functions that are included in a functional manager or feature.

- SMMonitor (Monitor role)

  The SMMonitor role can access the administrative functions that provide read-only access, such as monitoring, notification, and status. With this role, a user can complete tasks, such as monitoring a process, viewing and collecting inventory, and viewing hardware status.

- SMUser (User role)

  The SMUser role includes any authenticated user and includes the ability to perform only basic operations, such as viewing resources and properties.

- GroupRead (Group role)

  The GroupRead role has a single permission, which is known as group read, that defines the groups that are visible to each user. The administrator that assigns this role to a user can assign the groups that the user can view. The user then has access to see the groups but not necessarily the group contents.

These default user roles correspond directly with the groups that Systems Director installs at the operating system level. You cannot delete these roles and you cannot modify the permissions that are associated with them. However, you can add users and other groups to the system-defined roles as needed. You also can copy the system-defined roles or create new roles for your business needs.

Assigning a role to a user or user group

The roles that are assigned to a user or user group determine the tasks that the user has permission to access. From the Users page, you can assign one or more roles to a user or user group. When you assign a role, you also associate the specific resource groups to which that role applies to the selected user.

Before you can assign a role to a user, each user or group of users must have a valid user ID or group ID in the local operating system user registry on the management server. Also, ensure that the role that you want to assign to a user exists. If the role does not exist, you can create a role from the Roles page.

To assign a role to a user or group, complete the following steps:

1. In the Systems Director web interface navigation area, click **Security** → **Users and Groups**. Or, select **Manage Users** on the Home page Plug-ins tab. See Figure 4-164.

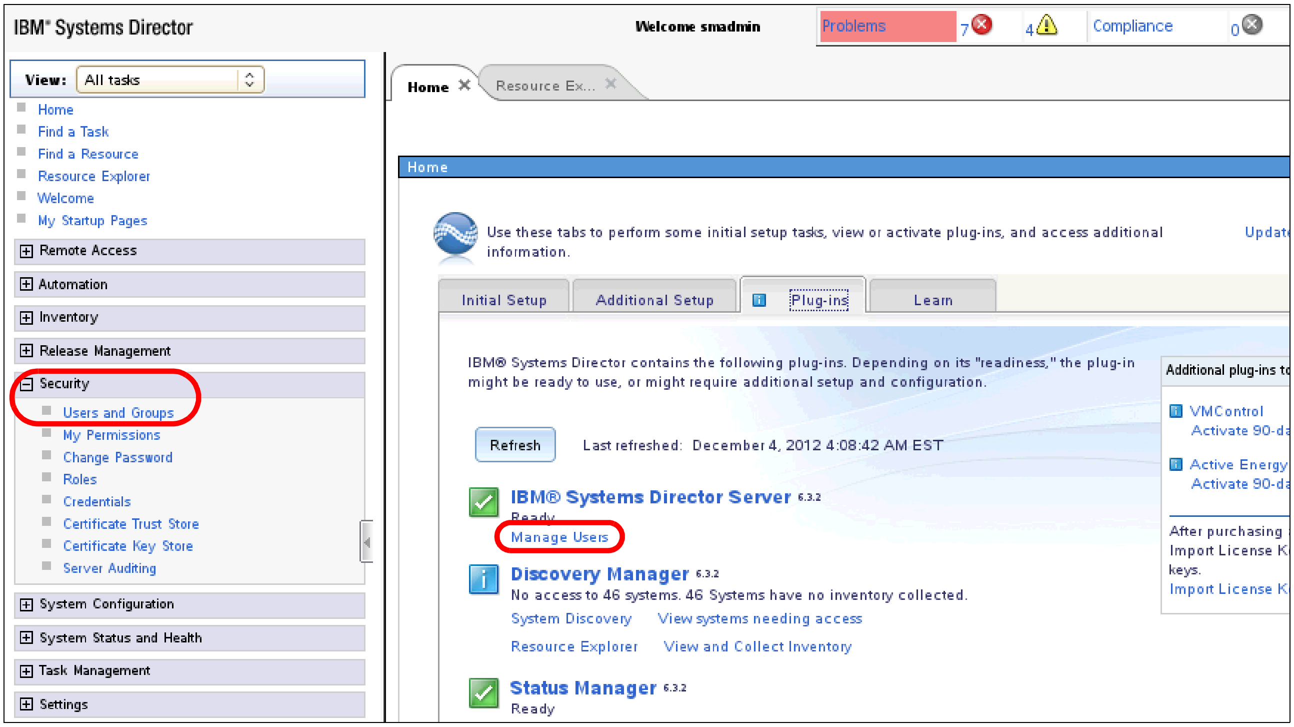

Use these tabs to perform some initial setup tasks, view or activate plug-ins, and access additional information.

IBM® Systems Director contains the following plug-ins. Depending on its "readiness," the plug-in might be ready to use, or might require additional setup and configuration.

Last refreshed: December 4, 2012 4:08:42 AM EST

IBM® Systems Director Server 6.3.2
Ready

Discovery Manager 6.3.2
No access to 46 systems. 46 Systems have no inventory collected.

Status Manager 6.3.2
Ready

After purchasing
Import License K
keys.

Figure 4-164 Security → Users and Groups or Manage Users

2. From the Users tab, select the user or group to which you want to assign a role. In our example, we select the user `smtest` (Figure 4-165).

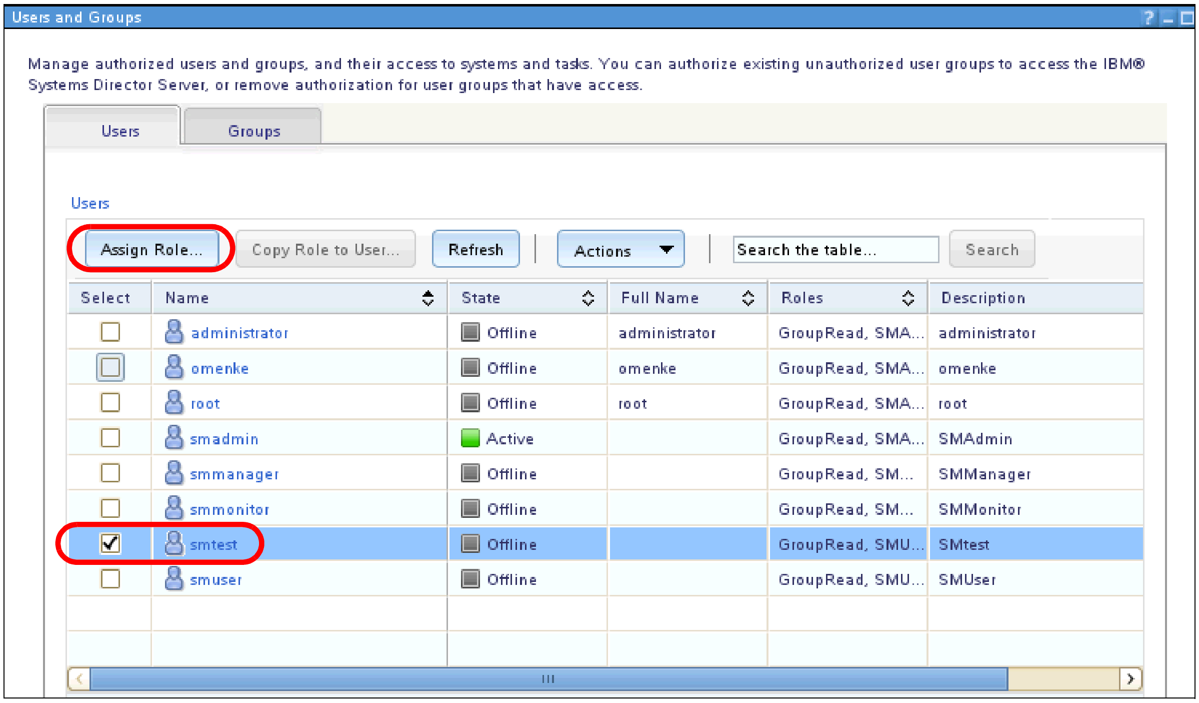

Manage authorized users and groups, and their access to systems and tasks. You can authorize existing unauthorized user groups to access the IBM® Systems Director Server, or remove authorization for user groups that have access.

Users

Select	Name	State	Full Name	Roles	Description
□	administrator	Offline	administrator	GroupRead, SMA...	administrator
□	omenke	Offline	omenke	GroupRead, SMA...	omenke
□	root	Offline	root	GroupRead, SMA...	root
□	smadmin	Active		GroupRead, SMA...	SMAdmin
□	smmanager	Offline		GroupRead, SM...	SMManager
□	smmonitor	Offline		GroupRead, SM...	SMMonitor
☑	smtest	Offline		GroupRead, SMU...	SMtest
□	smuser	Offline		GroupRead, SMU...	SMUser

Figure 4-165 Select user for assigning role

3. Click **Assign Role**. The Welcome page for the Assign Role wizard opens (Figure 4-166). Click **Next**.

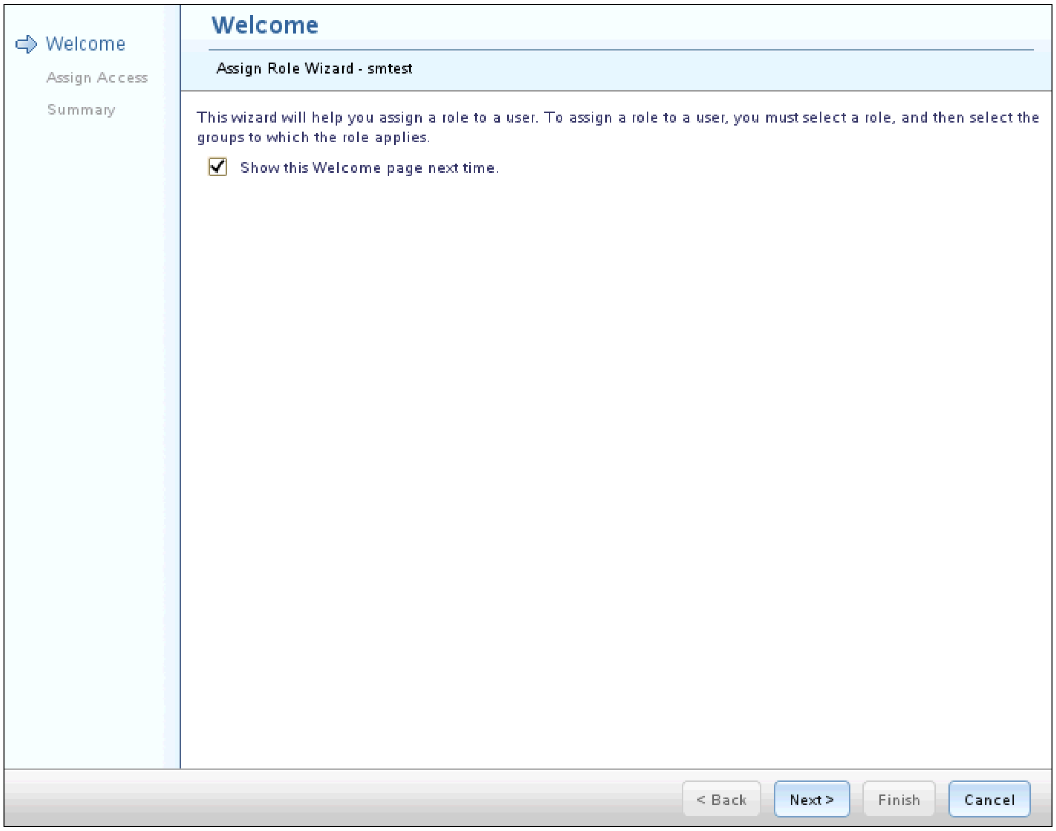

Welcome

Assign Role Wizard - smtest

This wizard will help you assign a role to a user. To assign a role to a user, you must select a role, and then select the groups to which the role applies.

☑ Show this Welcome page next time.

Figure 4-166 Welcome panel

4. The wizard lists the roles that are created. In our example, we select the `SMMonitor` role for the user `smtest` (Figure 4-167).

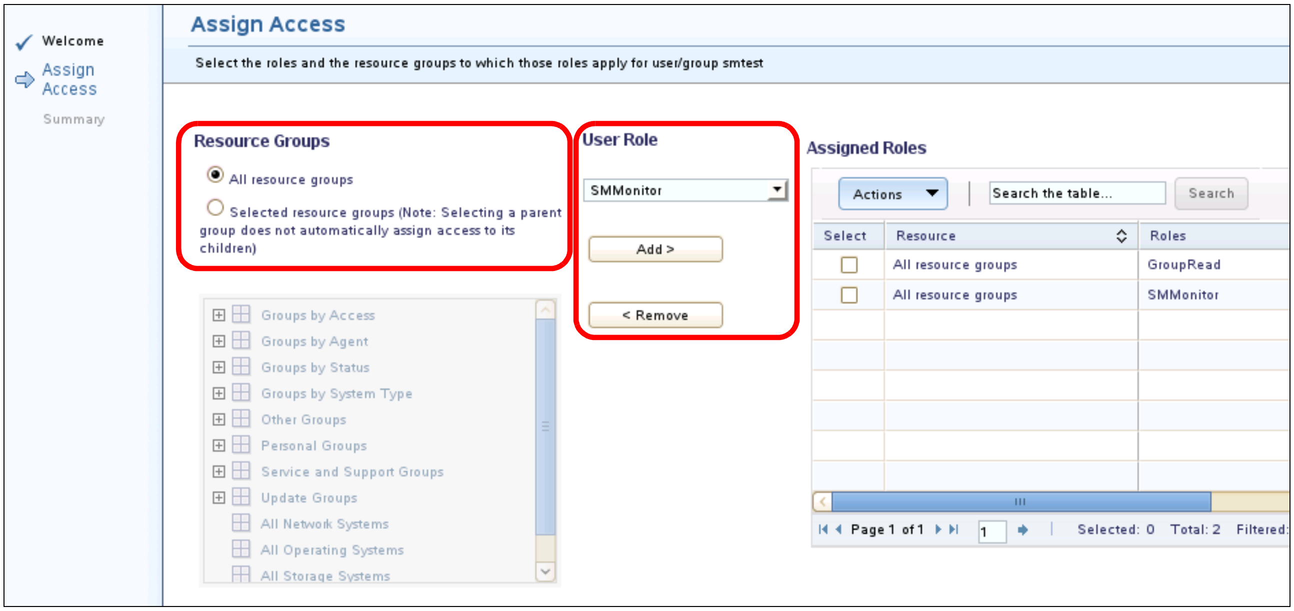

Assign Access

Select the roles and the resource groups to which those roles apply for user/group smtest

Resource Groups

● All resource groups

○ Selected resource groups (Note: Selecting a parent group does not automatically assign access to its children)

User Role

SMMonitor

Assigned Roles

Select	Resource	Roles
□	All resource groups	GroupRead
□	All resource groups	SMMonitor

Page 1 of 1 | 1 | Selected: 0 Total: 2 Filtered:

Figure 4-167 Assigning roles and resource groups

5. Select the role that you want to assign and click **Add**.
6. Select the resource groups that you want to associate with the role and the user. In our example, we select `All resource groups`.

> **Parent groups:** Selecting a parent group does not automatically assign access to its children.

7. Click **Next**. The Summary page opens (Figure 4-168). Click **Finish**.

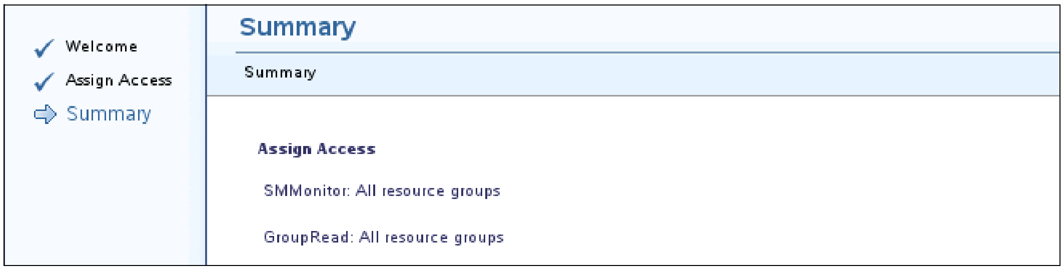

Summary

Summary

Assign Access

SMMonitor: All resource groups

GroupRead: All resource groups

Figure 4-168 Summary page

Working with roles

Use Systems Director to work with roles and assign individual users and user groups to those roles. From the Roles page (Figure 4-169 on page 215), you can view, copy, edit, or delete a role. To view, copy, edit, or delete a role, the role must exist. You can also use the Roles page to create a role that you can then manage.

Follow these steps to create a role:

1. In the navigation area, click **Security** → **Roles** (Figure 4-169).

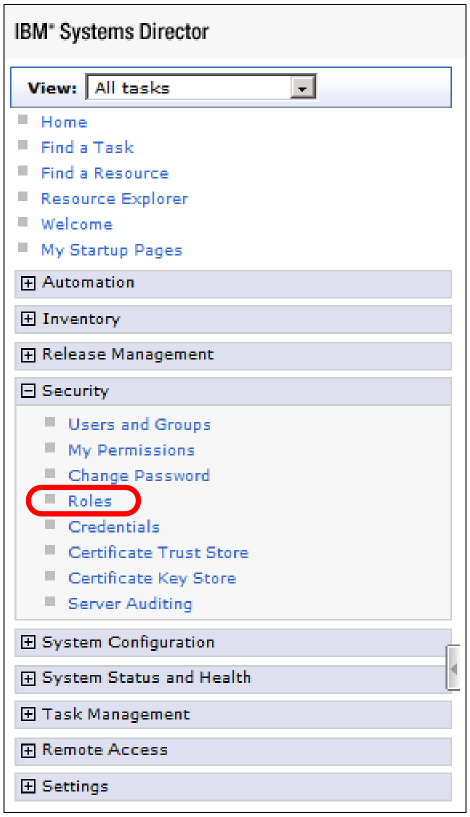

Figure 4-169 Select Roles from the Security section

2. On the Roles tab, click **Create** (Figure 4-170).

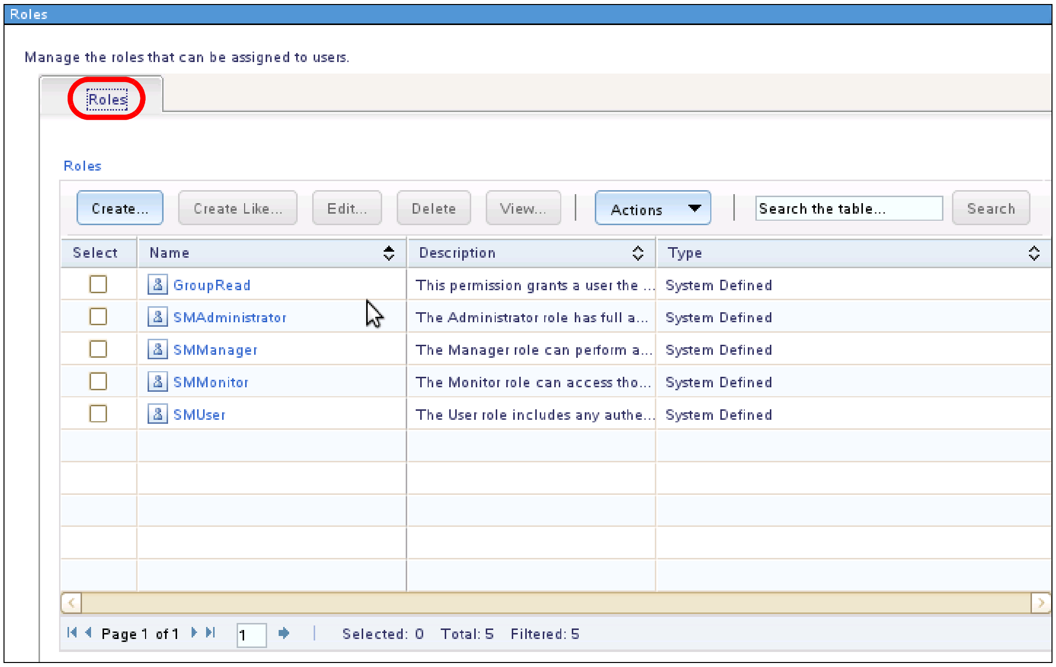

Select	Name	Description	Type
□	GroupRead	This permission grants a user the ...	System Defined
□	SMAdministrator	The Administrator role has full a...	System Defined
□	SMManager	The Manager role can perform a...	System Defined
□	SMMonitor	The Monitor role can access tho...	System Defined
□	SMUser	The User role includes any authe...	System Defined

Figure 4-170 Create a role

3. The Create Role wizard Welcome page opens. Click **Next**.
4. The Name page opens (Figure 4-171). In the Name field, type a name for the role that you want to create. In our example, we named the new role `book-Test`. In the Description field, type an optional brief description for the role. Click **Next**.

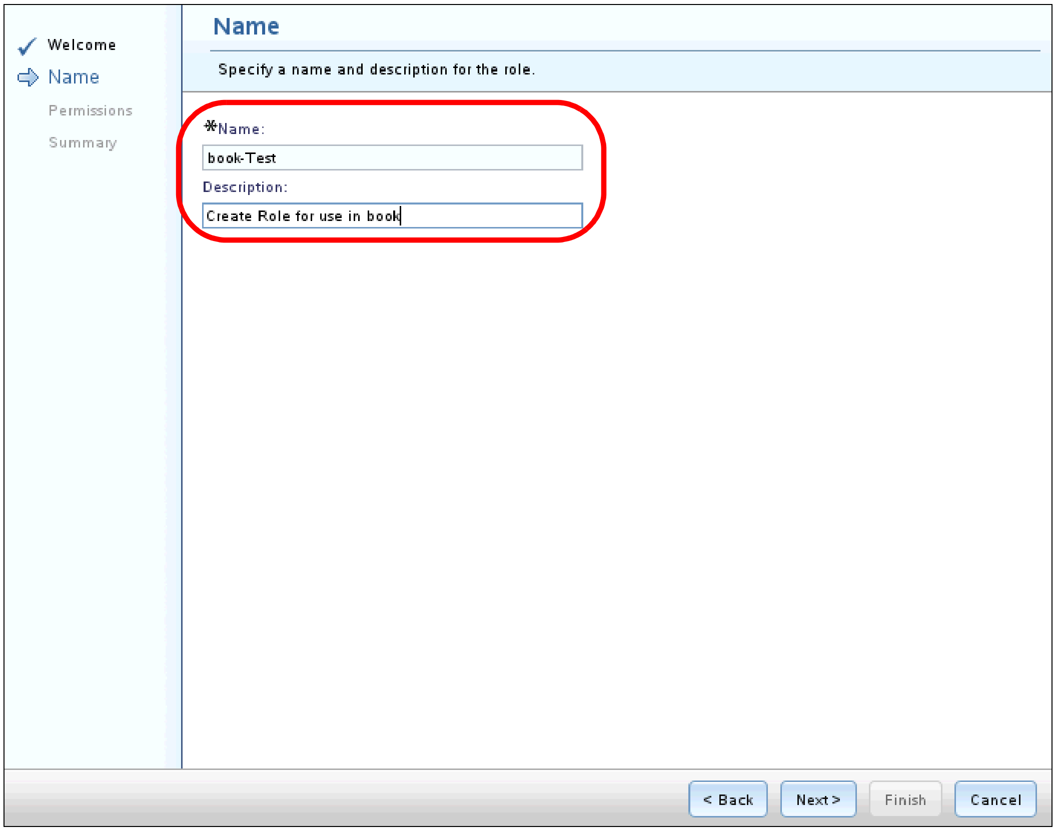

Specify a name and description for the role.

*Name:
book-Test

Description:
Create Role for use in book

Figure 4-171 Naming the new role

5. The Permissions page opens (Figure 4-172 on page 217). In the `Available permissions` list, select a permission that you want to add to the user role and then click **Add**. The selected permission is added to the `Selected permissions` list. Continue to add permissions until you add all permissions that are required for the role.

   In our example, we select Inventory and Task Management as permissions for the new role.

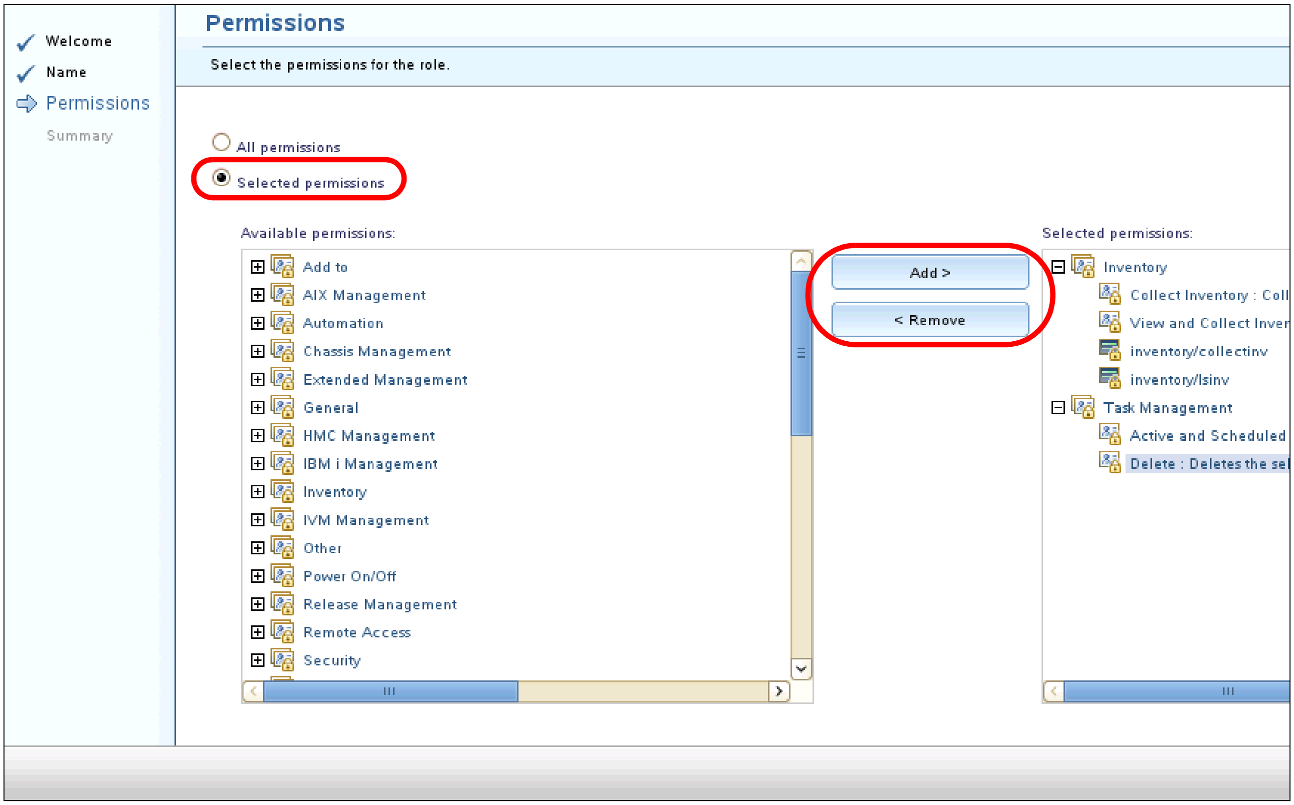

Permissions

Select the permissions for the role.

○ All permissions

● Selected permissions

Available permissions:

- Add to
- AIX Management
- Automation
- Chassis Management
- Extended Management
- General
- HMC Management
- IBM i Management
- Inventory
- IVM Management
- Other
- Power On/Off
- Release Management
- Remote Access
- Security

Selected permissions:

- Inventory
  - Collect Inventory : Coll
  - View and Collect Inver
  - inventory/collectinv
  - inventory/lsinv
- Task Management
  - Active and Scheduled
  - Delete : Deletes the sel

Figure 4-172 Select permissions

6. Click **Next**. The Summary page opens (Figure 4-173). Click **Finish**.

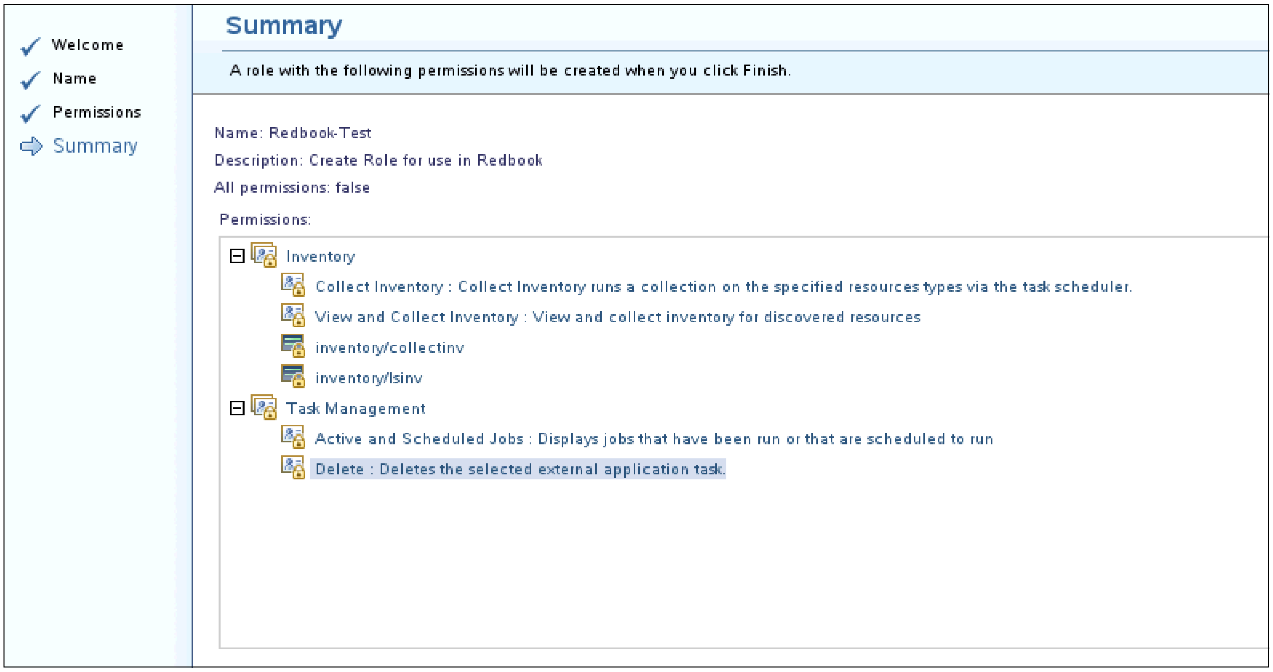

Summary

A role with the following permissions will be created when you click Finish.

Name: Redbook-Test

Description: Create Role for use in Redbook

All permissions: false

Permissions:

- Inventory
  - Collect Inventory : Collect Inventory runs a collection on the specified resources types via the task scheduler.
  - View and Collect Inventory : View and collect inventory for discovered resources
  - inventory/collectinv
  - inventory/lsinv
- Task Management
  - Active and Scheduled Jobs : Displays jobs that have been run or that are scheduled to run
  - Delete : Deletes the selected external application task.

Figure 4-173 Summary page

4.6.3 Access managed systems

Use Systems Director to configure credentials that are used to access managed systems. These credentials enable Systems Director to authenticate to and manage target systems by using the available protocols and access points on the managed system.

You can request access to and configure access options for systems in your environment by using these tasks:

- Request access task
- Configure access task
- Configure system credentials task

You can also revoke access to an accessed system.

Security protocols

Depending on the managed system, the following communication protocols are supported (Figure 4-6).

Encrypted protocols: Not all protocols are encrypted as indicated in the table.

Table 4-6 Supported communication protocols

Managed system type	Communication protocol	Encrypted	Encryption algorithm
Agentless-managed system	Distributed component object model (DCOM)	Yes	RC2
	SNMP v1 and v2	No	None
	Secure Shell (SSH)	Yes	Encrypted algorithm is negotiated
Platform Agent-managed system	Agentless	Yes	Supports the communication protocols and encryption algorithms that are listed for the agentless-managed system
	Common Information Model (CIM)	Yes	If configured, encryption is enabled by default by using Secure Sockets Layer (SSL)
Common Agent-managed system	IBM Director 5.*x* interprocess communication (IPC)	Yes	AES, DES, or 3DES
	Tivoli Common Agent Services 6.*x*	Yes	SSL
Other	Service Location Protocol (SLP)	No	None

Access secured systems

Use the Request Access page to request access to a secured system if the management server to which you connect is not yet authenticated to the system. You must be able to access the system before you can perform tasks or remotely access the system.

Ensure that you have the correct authorization to access the secured system.

Secured systems are displayed in the Systems Director web interface with a padlock icon in the Access field or column of the system details (Figure 4-174). After a system is accessed, the padlock disappears and additional tasks and status information are available.

The Access attribute for each resource shows the current access status. You cannot request access to the resources with the following types of access status:

- Offline: Use verify access instead.
- OK: No further action is required. You already have access to these resources.

To request access to secured managed systems, complete the following steps:

1. In the Systems Director web interface, click **Resource Explorer**.
2. Navigate to the system that you want to access.
3. Right-click the system for which you want to request access and click **Security** → **Request Access**.

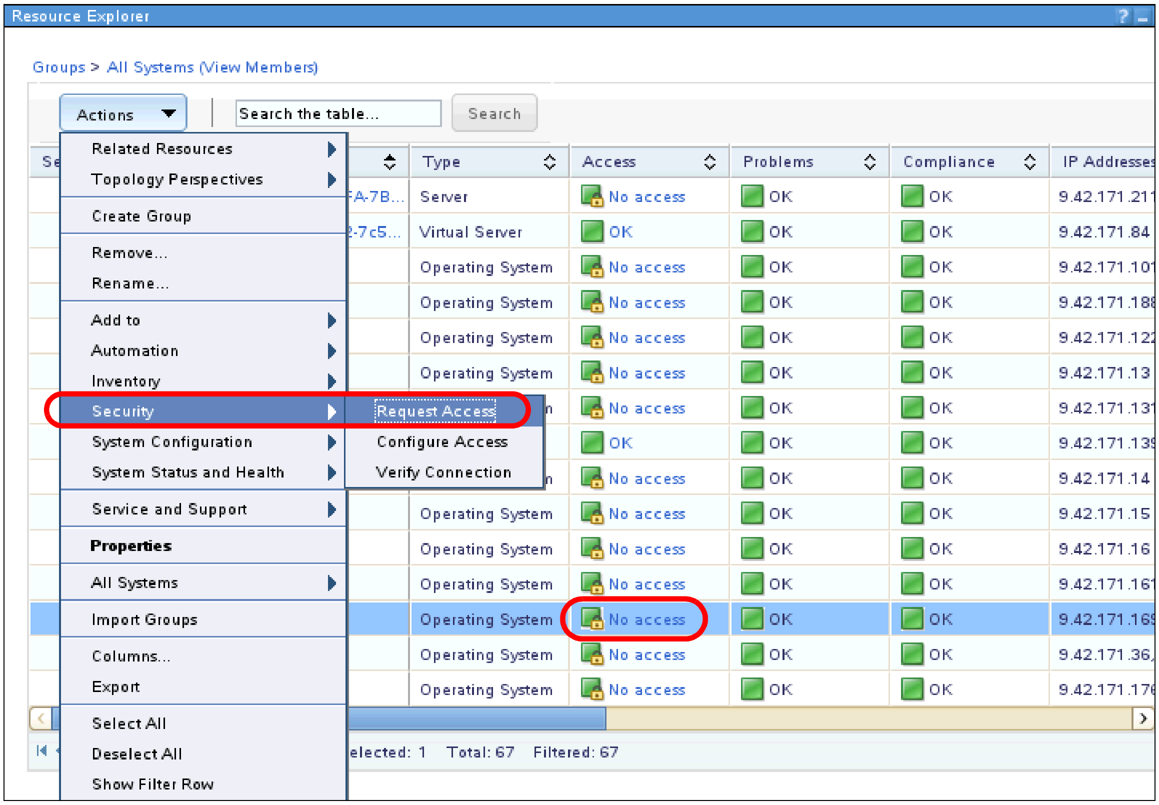

Se		Type	Access	Problems	Compliance	IP Addresses
	FA-7B...	Server	No access	OK	OK	9.42.171.211
	2-7c5...	Virtual Server	OK	OK	OK	9.42.171.84
		Operating System	No access	OK	OK	9.42.171.101
		Operating System	No access	OK	OK	9.42.171.188
		Operating System	No access	OK	OK	9.42.171.122
		Operating System	No access	OK	OK	9.42.171.13
		n	No access	OK	OK	9.42.171.131
			OK	OK	OK	9.42.171.139
		n	No access	OK	OK	9.42.171.14
		Operating System	No access	OK	OK	9.42.171.15
		Operating System	No access	OK	OK	9.42.171.16
		Operating System	No access	OK	OK	9.42.171.161
		Operating System	No access	OK	OK	9.42.171.169
		Operating System	No access	OK	OK	9.42.171.36,
		Operating System	No access	OK	OK	9.42.171.176

Figure 4-174 Request Access option

> **Tip:** Alternatively, you can click **Security** → **Configure Access** and then click **Request Access** on the Configure Access page.

4. On the Request Access page, type the user ID and password of a user that belongs to the system group (Figure 4-175). Only certain user accounts can be used to request access.

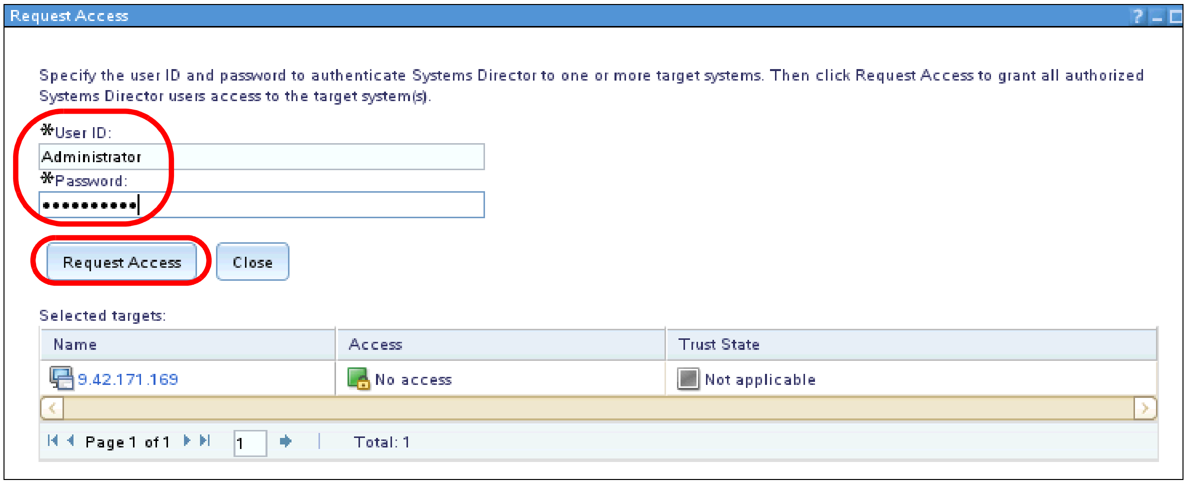

Request Access

Specify the user ID and password to authenticate Systems Director to one or more target systems. Then click Request Access to grant all authorized Systems Director users access to the target system(s).

*User ID: Administrator

*Password: ••••••••••

Selected targets:

Name	Access	Trust State
9.42.171.169	No access	Not applicable

Figure 4-175 Request Access panel

The following list shows the detailed requirements of the user accounts that can be used to request access for various types of agent systems:

– Common Agent:
  - Linux/AIX: Root or user in the system group
  - Windows: Administrator or user in the administrator group
– Platform Agent:
  - Linux/AIX: Root or user in the system group
  - Windows: Administrator or user in the administrator group
– Agentless systems:
  - Linux/AIX: Root or user in the system group. User that is configured with the **sudo** command.
  - Windows: Administrator or user in the administrator group

> **Tip:** To configure sudo and SSL keys for UNIX agentless systems, follow the procedure that is at the following link:
>
> `http://pic.dhe.ibm.com/infocenter/director/v6r2x/index.jsp?topic=/com.ibm.director.install.helps.doc/fqm0_c_sample_sudo_configuration_file.html`

Click **Request Access** (Figure 4-175). Credentials are created and authenticated to the managed system in an attempt to access it. If the access request is successful, the access status for the managed system changes to `OK`.

If the access status changes to `Partial Access`, the access request was unsuccessful for at least one protocol. Click **Configure Access** to see the list of available protocols for the system and their access states. If necessary, to create additional credentials, click an access point that does not have an access state of `OK` and repeat this procedure.

For information about accessing systems by using credentials, configuring access, or accessing CIM systems by using the x509 certificate, see the information center:

`http://pic.dhe.ibm.com/infocenter/director/pubs/index.jsp?topic=%2Fcom.ibm.director.security.helps.doc%2Ffqm0_t_managing_access.html`

4.6.4 Credentials

The Systems Director server uses credentials to implement SSO authentication. By using SSO with this authentication process, a user can access more than one system or application by entering a single user ID and password. The Systems Director server maps web interface user credentials to the necessary user credentials for authenticating to the target managed system. These credentials are saved in registries.

It is a preferred practice to use SSO because users are not required to type the user ID and password for the target system or resource each time that they or tasks access it. The Systems Director server automatically logs on as needed by retrieving the necessary credentials.

There are two types of credentials:

- Shared credentials

  Shared credentials are those credentials that exist in an authentication registry that is not specific to an access point.

- Targeted credentials

  Targeted credentials are each assigned to only one remote-service agent access point and are in an authentication registry that is specific to that access point.

For more information about credentials, see the Systems Director Information Center:

http://pic.dhe.ibm.com/infocenter/director/pubs/index.jsp?topic=%2Fcom.ibm.director.security.helps.doc%2Ffqm0_c_credentials.html

4.6.5 Lightweight Directory Access Protocol

The Systems Director server can authenticate users that are defined in the LDAP directory. Many benefits are possible if you use LDAP as the preferred authentication method:

- Ease of management
- Central administration
- Cross-platform synergies

The following LDAP servers are supported by Systems Director:

- Microsoft Active Directory
- IBM Lotus® Domino®
- IBM Tivoli Director Server
- Sun One
- OpenLDAP
- IBM Secure Way Server
- Novell eDirectory

Configuring OpenLDAP

We configure OpenLDAP on Red Hat Enterprise Linux Server 5.6. Install the Red Hat Package Manager (RPM) packages that are shown in Figure 4-176.

```
[root@xs-2120rhelppc ~]# rpm -qa | grep openld
compat-openldap-2.3.43_2.2.29-12.el5_5.3
openldap-devel-2.3.43-12.el5_5.3
openldap-2.3.43-12.el5_5.3
openldap-clients-2.3.43-12.el5_5.3
openldap-servers-2.3.43-12.el5_5.3
openldap-2.3.43-12.el5_5.3
openldap-devel-2.3.43-12.el5_5.3
[root@xs-2120rhelppc ~]
```

Figure 4-176 OpenLDAP RPM packages

Discuss the properties that need to be edited in the `securityLDAP.properties` file with your LDAP administrator. For more information about LDAP, see the information center at the following link:

http://pic.dhe.ibm.com/infocenter/director/pubs/index.jsp?topic=%2Fcom.ibm.director.security.helps.doc%2Ffqm0_t_ldap_authentication.html

> **File name change:** Rename the `security.ldap` file to the `securityLDAP.properties` file after you change the properties. If you use LDAP, Systems Director looks for the `securityLDAP.properties` file.

Table 4-7 shows the properties to be referenced or changed in the `securityLDAP.properties` file.

Table 4-7 OpenLDAP securityLDAP.properties file properties to change

Property	Value	Description
`com.ibm.lwi.LDAPHost`	IP or host name	Address of LDAP server
`com.ibm.lwi.LDAPAdminPassword`	Encrypted password	Read-only password for binding
`com.ibm.lwi.LDAPBase`	dc=itso,dc=ibm	Base distinguished name (DN) for LDAP server
`com.ibm.lwi.searchfilter`	(&(uid=%v)(objectclass=inetOrgPerson))	The user search filter for the LDAP server
`com.ibm.lwi.rolemanager.ldap.filters.usergroup`	(objectclass=posixGroup)	Authorized groups for the Systems Director server
`com.ibm.lwi.rolemanager.ldap.filters.users`	(|(objectClass=inetOrgPerson)(objectClass=posixAccount))	Group objects search
`com.ibm.lwi.rolemanager.ldap.names.memberAttribute`	uid	Member attribute role object
`com.ibm.lwi.rolemanager.ldap.names.loginName`	uid	Name of login attribute of user
`com.ibm.lwi.rolemanager.ldap.names.groupID`	gidNumber	Name of the group ID attribute of the group object
`com.ibm.lwi.rolemanager.ldap.names.userPrimaryGroupID`	gidNumber	Name of the group ID attribute of the group object

Property	Value	Description
`com.ibm.lwi.rolemanager.ldap.filters.usersByGroupId`	(&(gidNumber={0})(|(objectClass=inetOrgPerson)(objectClass=posixAccount)))	Users by gidNumber or member
`com.ibm.lwi.rolemanager.ldap.filters.groupsByMembers`	(&(|(gidNumber={0})(memberUid={1}))(objectclass=posixGroup))	Groups by gidNumber or posixGroup
`com.ibm.lwi.rolemanager.ldap.names.memberAttribute.isDN`	false	Specific to openLDAP

For information about the OpenLDAP slapd server configuration, see the following Red Hat web page:

https://access.redhat.com/knowledge/docs/en-US/Red_Hat_Enterprise_Linux/4/html/Reference_Guide/s1-ldap-quickstart.html#s2-ldap-files-slapd-conf

After you successfully install the required RPM packages, configure the `sldap.conf` file and encrypt the rootpw password by using the **`slappasswd`** command, as shown in Figure 4-177.

```
[root@xs-2120rhelppc openldap]# slappasswd
New password:
Re-enter new password:

{SSHA}4eb+Hf7KScesth8vftJ/Fdw8jKXV+mRL
```

Figure 4-177 slappasswd

The following configuration changes for the `slapd.conf` file are shown in Figure 4-178:

- suffix
- rootdn
- rootpw

```
database        bdb
suffix          "dc=itso,dc=ibm"
rootdn          "cn=root,dc=itso,dc=ibm"
# Cleartext passwords, especially for the rootdn, should
# be avoided.  See slappasswd(8) and slapd.conf(5) for details.
# Use of strong authentication encouraged.
# rootpw                secret
rootpw {SSHA}4eb+Hf7KScesth8vftJ/Fdw8jKXV+mRL
```

Figure 4-178 slapd.conf file configuration changes

From the command line, start the LDAP service and add an entry to start the service automatically on boot by using **`chkconfig`**, as shown in Figure 4-179.

```
[root@xs-2120rhelppc openldap]# service ldap start
Checking configuration files for slapd:  config file testing succeeded
                                                           [  OK  ]
Starting slapd:                                            [  OK  ]
[root@xs-2120rhelppc openldap]# chkconfig ldap on
[root@xs-2120rhelppc openldap]# chkconfig --list | grep ldap
ldap            0:off   1:off   2:on    3:on    4:on    5:on    6:off
[root@xs-2120rhelppc openldap]#
```

Figure 4-179 Service start: chkconfig check

Importing groups and users

From the LDAP server command line, import the ldif files for users and groups. To complete this task, we create a ldif file for importing. Create a `groups.ldif` file as shown in Figure 4-180.

```
dn: cn=smadmin,dc=itso,dc=ibm
cn: smadmin
objectClass: top
objectClass: posixGroup
gidNumber: 100
memberUid: root
memberUid: uid=root,cn=smadmin,dc=itso,dc=ibm

dn: uid=root,cn=smadmin,dc=itso,dc=ibm
cn: root
sn: root
uid: root
objectClass: top
objectClass: inetOrgPerson
objectClass: posixAccount
uidNumber: 100
gidNumber: 100
homeDirectory: /root
userPassword: operah09se

dn: cn=smmon,dc=itso,dc=ibm
cn: smmon
objectClass: top
objectClass: posixGroup
gidNumber: 101
memberUid: isduser
memberUid: uid=isduser,cn=smmon,dc=itso,dc=ibm

dn: cn=smmgr,dc=itso,dc=ibm
description: smmgr
cn: smmgr
objectClass: top
objectClass: posixGroup
gidNumber: 102
memberUid: uid=isdmgr,cn=smmgr,dc=itso,dc=ibm
memberUid: uid=isdmgr0,cn=smmgr,dc=itso,dc=ibm

dn: cn=smuser,dc=itso,dc=ibm
description: smuser
cn: smuser
objectClass: top
objectClass: posixGroup
gidNumber: 103
```

Figure 4-180 groups.ldif file

Create a `users.ldif` file as shown in Figure 4-181.

```
dn: uid=isduser,cn=smmon,dc=itso,dc=ibm
cn: isduser
sn: isduser
objectClass: top
objectClass: inetOrgPerson
objectClass: posixAccount
uid: isduser
uidNumber: 101
gidNumber: 101
homeDirectory: /home/isduser
userPassword: @Pa22w0rd

dn: uid=isdmgr0,cn=smmgr,dc=itso,dc=ibm
cn: isdmgr0
sn: isdmgr0
uid: isdmgr0
objectClass: top
objectClass: inetOrgPerson
objectClass: posixAccount
uidNumber: 102
gidNumber: 102
homeDirectory: /home/isdmgr0
userPassword: @Pa22w0rd

dn: uid=isdmgr1,cn=smmgr,dc=itso,dc=ibm
cn: isdmgr1
sn: isdmgr1
uid: isdmgr1
objectClass: top
objectClass: inetOrgPerson
objectClass: posixAccount
uidNumber: 103
gidNumber: 102
homeDirectory: /home/isdmgr1
userPassword: @Pa22w0rd
```

Figure 4-181 users.ldif file

The two ldif files are imported by using the `ldapadd` file as shown in Figure 4-182.

```
[root@xs-2120rhelppc ldapFiles]# ldapadd -H ldap://127.0.0.1 -x -D
"cn=root,dc=itso,dc=ibm" -f ~/ldapFiles/groups.ldif -w @Pa22w0rd
adding new entry "cn=smadmin,dc=itso,dc=ibm"

adding new entry "uid=root,cn=smadmin,dc=itso,dc=ibm"

adding new entry "cn=smmon,dc=itso,dc=ibm"

adding new entry "cn=smmgr,dc=itso,dc=ibm"

adding new entry "cn=smuser,dc=itso,dc=ibm"

[root@xs-2120rhelppc ldapFiles]# ldapadd -H ldap://127.0.0.1 -x -D
"cn=root,dc=itso,dc=ibm" -f ~/ldapFiles/users.ldif -w @Pa22w0rd
adding new entry "uid=isduser,cn=smmon,dc=itso,dc=ibm"

adding new entry "uid=isdmgr0,cn=smmgr,dc=itso,dc=ibm"

adding new entry "uid=isdmgr1,cn=smmgr,dc=itso,dc=ibm"
```

Figure 4-182 ldapadd file

To view the LDAP server, we use the LDAP command line (Figure 4-183).

```
[root@xs-2120rhelppc openldap]#ldapsearch -x -b 'dc=itso,dc=ibm'
# extended LDIF
#
# LDAPv3
# base <dc=itso,dc=ibm> with scope subtree
# filter: (objectclass=*)
# requesting: ALL
#
# itso.ibm
dn: dc=itso,dc=ibm
dc: itso
o: itso
objectClass: organization
objectClass: dcObject
# smadmin, itso.ibm
dn: cn=smadmin,dc=itso,dc=ibm
cn: smadmin
objectClass: top
objectClass: posixGroup
gidNumber: 100
memberUid: root
memberUid: uid=root,cn=smadmin,dc=itso,dc=ibm
# root, smadmin, itso.ibm
dn: uid=root,cn=smadmin,dc=itso,dc=ibm
cn: root
sn: root
uid: root
objectClass: top
objectClass: inetOrgPerson
objectClass: posixAccount
uidNumber: 100
gidNumber: 100
homeDirectory: /root
userPassword:: b3BlcmFoMDlzZQ==
# smmon, itso.ibm
dn: cn=smmon,dc=itso,dc=ibm
cn: smmon
objectClass: top
objectClass: posixGroup
gidNumber: 101
memberUid: isduser
# smmgr, itso.ibm
dn: cn=smmgr,dc=itso,dc=ibm
description: smmgr
cn: smmgr
objectClass: top
objectClass: posixGroup
gidNumber: 102
memberUid: uid=isdmgr,cn=smmgr,dc=itso,dc=ibm
memberUid: uid=isdmgr0,cn=smmgr,dc=itso,dc=ibm
memberUid: uid=isdmgr1,cn=smmgr,dc=itso,dc=ibm
```

Figure 4-183 ldapsearch

The LDAP server is now ready. The groups are defined. The users are defined. And, the `securityLDAP.properties` file is configured.

Therefore, we can make the last changes to the Systems Director server so that we can start to use LDAP as its authentication method (Figure 4-184). You need run the following commands on the IBM Systems Director server.

```
-bash-3.2# cd /opt/ibm/director/lwi/conf/overrides/
-bash-3.2# mv security.ldap securityLDAP.properties
-bash-3.2# mv security.properties security.properties.old
-bash-3.2# smstop;smstart;smstatus -r
```

Figure 4-184 File changes

From the Systems Director server home page, click **Plug-ins** and then **IBM Systems Director server**. A summary window opens and the authentication type is listed (Figure 4-185). The configuration is successful.

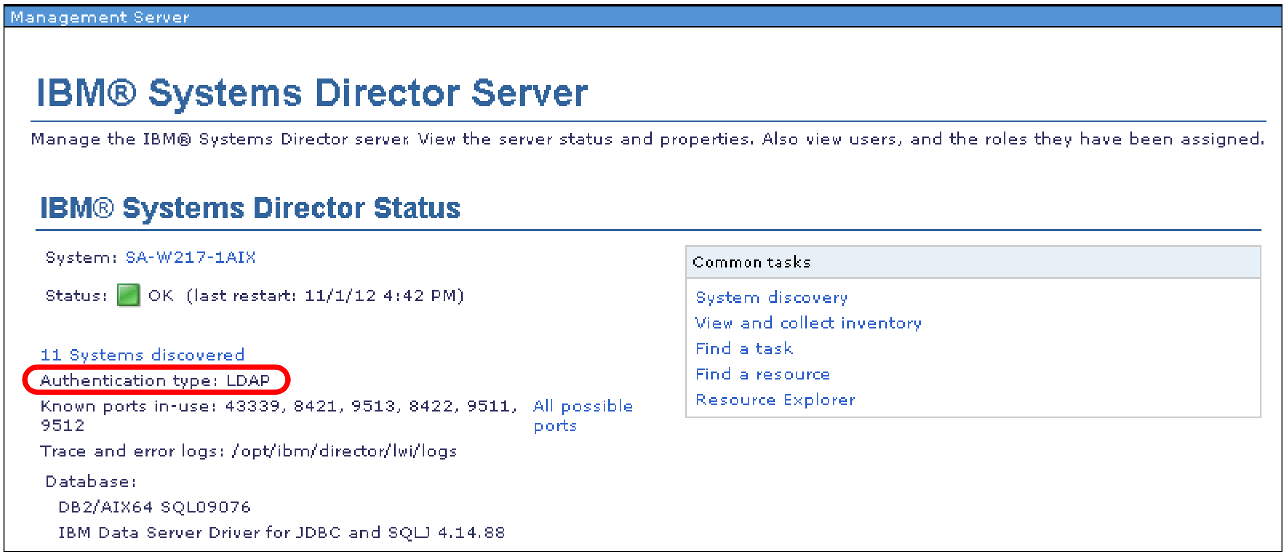

Management Server

IBM® Systems Director Server

Manage the IBM® Systems Director server. View the server status and properties. Also view users, and the roles they have been assigned.

IBM® Systems Director Status

System: SA-W217-1AIX

Status: 🟩 OK (last restart: 11/1/12 4:42 PM)

11 Systems discovered

Authentication type: LDAP

Known ports in-use: 43339, 8421, 9513, 8422, 9511, 9512 All possible ports

Trace and error logs: /opt/ibm/director/lwi/logs

Database:
DB2/AIX64 SQL09076
IBM Data Server Driver for JDBC and SQLJ 4.14.88

Figure 4-185 Confirming that LDAP is successfully configured

We now use OpenLDAP as the authentication type.

> **Important:** This example is a basic setup of openLDAP. Ensure that you take additional security measures with your configuration to further secure the openLDAP server and LDAP administration.

Authenticating users and groups

Additional groups and their associated users can be authorized to Systems Director. This authorization can be set up for new groups or groups that exist in the LDAP domain.

The following example imports a new Systems Director group with users to the OpenLDAP server by using the **ldapadd** command. This group is called the `isdgroup` (Figure 4-186).

```
[root@xs-2120rhelppc ldapFiles]# ldapadd -H ldap://127.0.0.1 -x -D
"cn=root,dc=itso,dc=ibm" -f ~/ldapFiles/addgroup.ldif  -w operah09se
adding new entry "cn=isdgroup,dc=itso,dc=ibm"

[root@xs-2120rhelppc ldapFiles]# ldapadd -H ldap://127.0.0.1 -x -D
"cn=root,dc=itso,dc=ibm" -f ~/ldapFiles/adduser.ldif  -w operah09se
adding new entry "uid=user0,cn=isdgroup,dc=itso,dc=ibm"

adding new entry "uid=user1,cn=isdgroup,dc=itso,dc=ibm"
```

Figure 4-186 Authorizing an additional group

Now that the group is added, we add the additional users because this task is an incremental installation. Then, we need to authorize the group to Systems Director. Follow these steps:

1. From the Systems Director home page, click **Security** → **Users and Groups** → **Groups**. Then, click **Authorize Groups**. The Authorize User Groups wizard starts.
2. On the Welcome page, click **Next**.
3. Figure 4-187 appears. The group that was imported to LDAP by using the **ldapadd** command is displayed. Place a check mark next to the group and click **Next**.

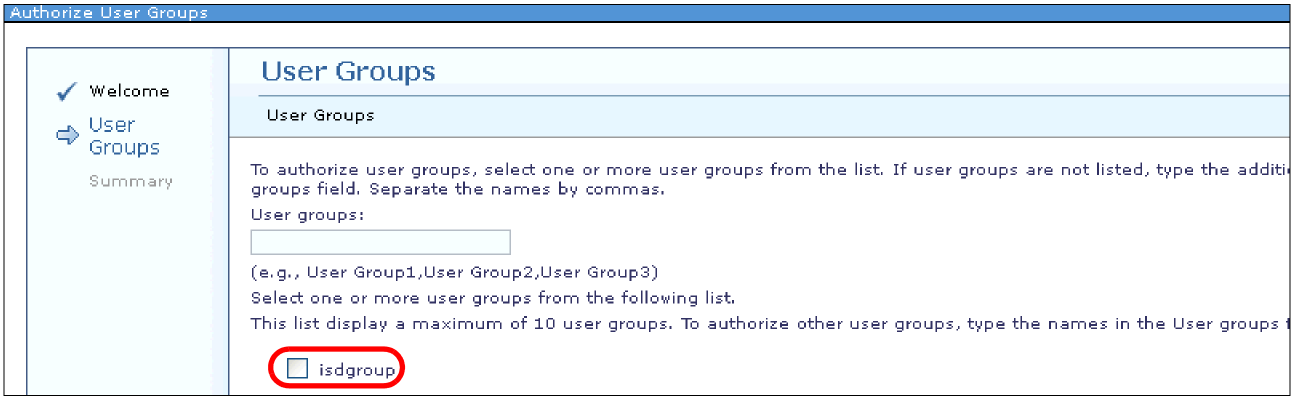

To authorize user groups, select one or more user groups from the list. If user groups are not listed, type the additi groups field. Separate the names by commas.

User groups:

(e.g., User Group1,User Group2,User Group3)

Select one or more user groups from the following list.

This list display a maximum of 10 user groups. To authorize other user groups, type the names in the User groups

□ isdgroup

Figure 4-187 Authorize User Groups window

4. Click **Finish** to authorize the `isdgroup` user group and complete the wizard.

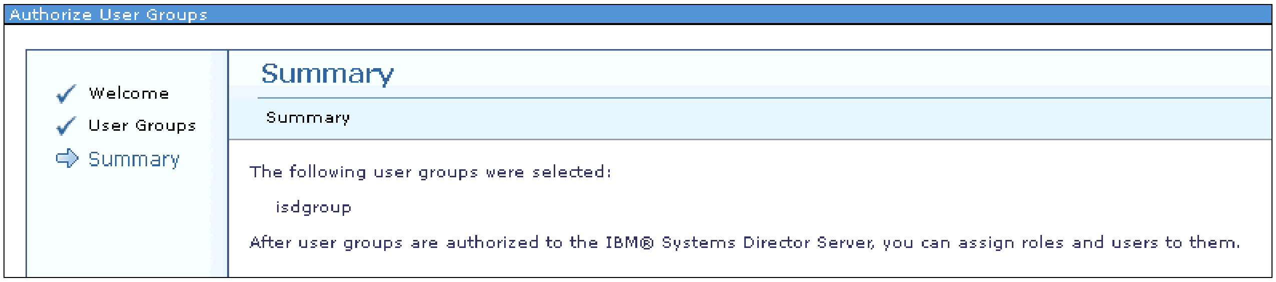

The following user groups were selected:

isdgroup

After user groups are authorized to the IBM® Systems Director Server, you can assign roles and users to them.

Figure 4-188 Authorize the isdgroup user group

5. We still need to assign a role to the group that is authorized. The group can be assigned to a custom role or one of the default roles:
   – SMAdministrator

  - SMManager
  - SMMonitor
  - SMUser

6. Select the group that we authorized (note that the `Roles` entry is empty for the `isdgroup` group in Figure 4-189) and click **Assign Role**.

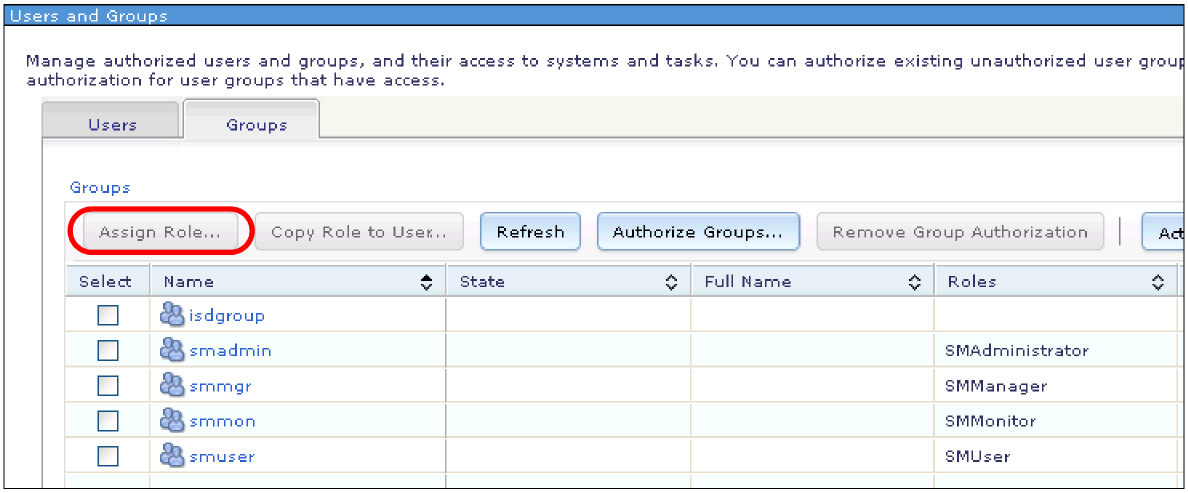

Users and Groups

Manage authorized users and groups, and their access to systems and tasks. You can authorize existing unauthorized user group authorization for user groups that have access.

Select	Name	State	Full Name	Roles
□	isdgroup			
□	smadmin			SMAdministrator
□	smmgr			SMManager
□	smmon			SMMonitor
□	smuser			SMUser

Figure 4-189 Assign Role task

7. For this group, we assign an `SMAdministrator` role to the group. The Assign Role wizard starts. Click **Next**.

8. Then, on the Assign Role window (Figure 4-190), choose **SMAdministrator** from the pull-down menu that is highlighted in Figure 4-190 and **Add**.

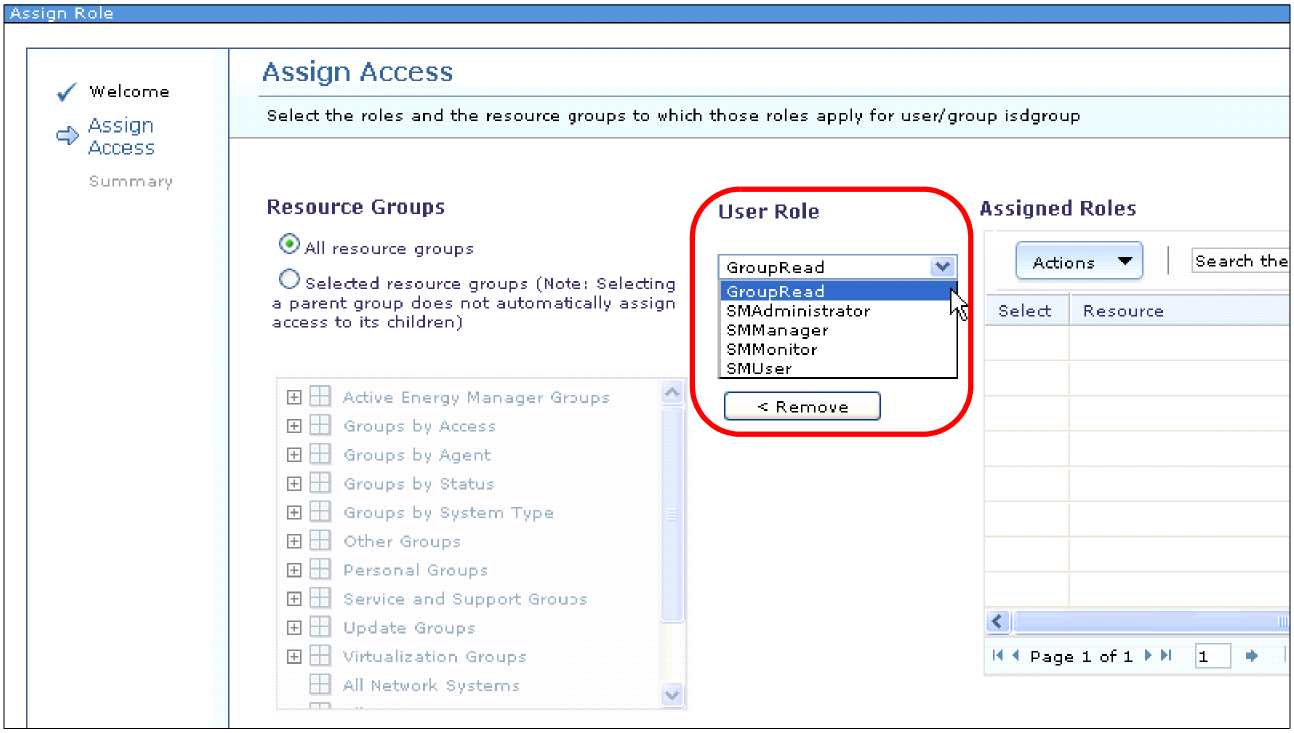

Assign Role

Assign Access

Select the roles and the resource groups to which those roles apply for user/group isdgroup

Resource Groups

● All resource groups

○ Selected resource groups (Note: Selecting a parent group does not automatically assign access to its children)

User Role: GroupRead

Figure 4-190 Assign Role window

9. On completion, click **Next** and **Finish**.

10. Now, the `isdgroup` group and its associated users have `SMAdministrator` access to all resources.

11. From the home page, click **Security** → **Users and Groups** → **Users** (Figure 4-191).

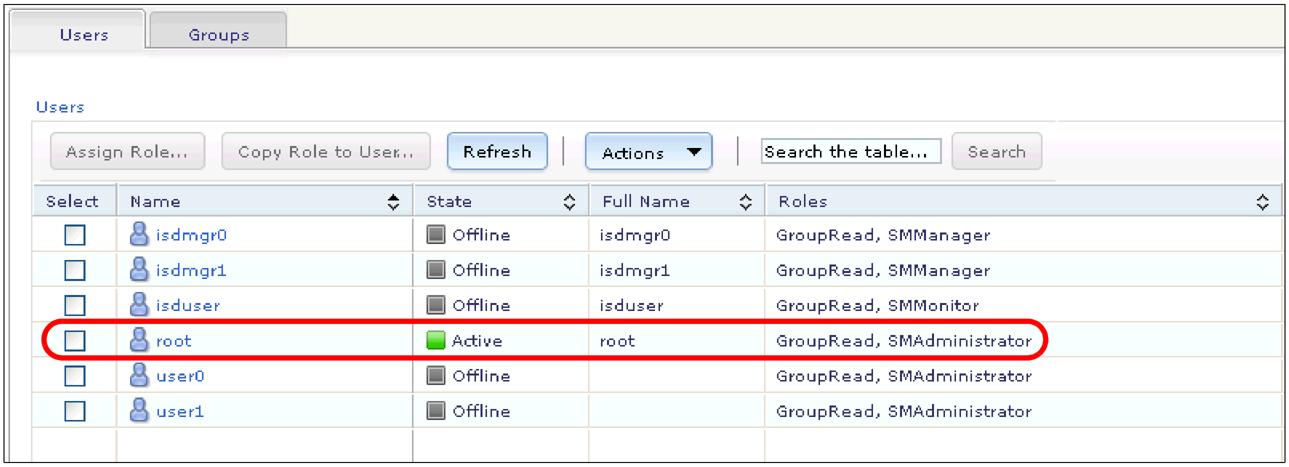

Select	Name	State	Full Name	Roles
□	isdmgr0	Offline	isdmgr0	GroupRead, SMManager
□	isdmgr1	Offline	isdmgr1	GroupRead, SMManager
□	isduser	Offline	isduser	GroupRead, SMMonitor
□	root	Active	root	GroupRead, SMAdministrator
□	user0	Offline		GroupRead, SMAdministrator
□	user1	Offline		GroupRead, SMAdministrator

Figure 4-191 User listing

You can now successfully log on as a user that is listed in the `isdgroup`.

> **Tip:** When you use the wizard to authorize groups, only the first 10 groups are returned in Figure 4-190 on page 230. If your group is not listed, type the name. If the typed name does not return the group, check your filters in the `securityLDAP.properties` file.

4.6.6 Using command-line tools for security

There are many available **`smcli`** command-line tools for security settings, as listed in Table 4-8.

Table 4-8 Command-line tools for security

Command	Description
`authusergp`	Authorize an existing user group to access the Systems Director server.
`cfgaccess`	Configure access for systems that are managed by Systems Director.
`cfgappcred`	Change the password that Systems Director uses to access particular associated applications.
`cfgcertpolicy`	View or configure the trust management certificate policy that IBM Systems Director uses.
`cfgcred`	Configure credentials for systems that are managed by Systems Director.
`cfgpwdpolicy`	Manage the password policies of users that Systems Director creates or manages.
`chaudit`	Modify audit settings.
`chcred`	Change credentials for systems that are managed by Systems Director.
`chrole`	Change the properties of a role.
`chuser`	Modify the properties of a user.
`chusergp`	Change attributes and access privileges for a user group.

Command	Description
exportcert	Export a certificate from a Systems Director keystore or truststore to a .pem file.
importcert	Import certificates into a Systems Director keystore or truststore.
lsaudit	List audit settings and categories.
lsauditlogs	List a specific number of audit log messages for one or more audit categories.
lscert	List the certificates in a Systems Director keystore or truststore.
scred	List credentials for systems that are managed by Systems Director.
lsperm	List the permissions.
lsrole	List the roles in Systems Director.
lsuser	List users.
lsusergp	List the Systems Director user groups.
mkrole	Create roles that contain a list of permissions for authorization to access Systems Director.
revokecert	Invalidate certificates in a Systems Director keystore or truststore.
rmauditlogs	Remove the audit log for one or more audit categories.
rmcert	Remove certificates from a Systems Director keystore or truststore.
rmcred	Remove credentials for systems that are managed by Systems Director.
rmrole	Delete roles.
rmusergp	Remove the access authorization for a user group or remove a user group.
unrevokecert	Revalidate revoked certificates in a Systems Director keystore or truststore.

For detailed information about the commands and all options for these commands, see the information center:

http://pic.dhe.ibm.com/infocenter/director/pubs/index.jsp?topic=%2Fcom.ibm.director.security.helps.doc%2Ffqm0_r_roles_required_to_run_commands.html

Accessing command-line tools by role/group

The command-line tools are restricted by the permissions. Table 4-9 lists which commands can be accessed by the roles.

Table 4-9 Command-line tools by role/group

Group	Restricted	Specifics
SMAdmin	Unrestricted	All commands
SMManager	Restricted	All commands except the security and system commands
SMMonitor	Restricted	All commands with list functions, such as **lscfgplan**, **lsinv**, **lsled**, **lsstatus**, **lsresmonlsps**, and **lsevtfltr** Also, commands such as **checkupd** and **lsupd**, or commands for SNMP, such as **get**, **walk**, and **getnext**
SMUser	Restricted	Only support for the following commands: **lssys**, **lsgp**, **lsjob**, **lsjobhistory**, **lstask**, **runjob**, and **runtask**

For a complete list of commands, see this website:

http://pic.dhe.ibm.com/infocenter/director/pubs/index.jsp?topic=%2Fcom.ibm.director.security.helps.doc%2Ffqm0_r_roles_required_to_run_commands.html

4.6.7 Best practices for security

These are the best practices to use the security setting in IBM Systems Director efficiently:

- Manage users that will connect to the IBM Systems Director console:
  - Define who can access the console.
  - Define authorization and roles for these users into IBM Systems Director.
  - Define method of authentication on IBM Systems Director server (LDAP, AD, or local).
  - For a UNIX agentless server, configure SSL keys and sudo configuration to avoid root utilization.
- When Using LDAP or Active Directory, configure the settings carefully.

4.7 Backup

Why you should back up IBM Systems Director server and how to recover from a Systems Director failure are described.

The following topics are included:

- 4.7.1, “Backup Q&A” on page 233
- 4.7.2, “Backup and recovery” on page 234
- 4.7.3, “Migration” on page 238

4.7.1 Backup Q&A

You might ask the following backup questions:

- Why back up?

  Backing up the Systems Director, including all data and settings, makes it easier to recover from a Systems Director server crash. Systems Director provides command-line tools to perform both the backup and the recovery. These tools are explained in section 4.7.2, “Backup and recovery” on page 234.

- When do I back up?

  The preferred practice is to back up your Systems Director server after installation and initial configuration. Also, back up after discovery and inventory but before you first install any updates. Then, you can quickly recover back to a fresh installation, if necessary. Always back up after you make any updates to the Systems Director server.

- How often do I back up?

  We suggest that you back up before you install plug-ins or advanced managers. We also recommend that you back up regularly, such as once a month.

- What information is backed up?
- The Systems Director backup routine creates a backup image of Systems Director persistent data. *Persistent data* includes file system data, which is also called *master data*, and database data:
  - Information about discovered systems and their access state
  - All event automation plans

  - All groups
  - All inventory data that is stored in the Systems Director database
  - Event logs

  The backup also includes data in the local repository, such as updates that you downloaded.

- What can the backup not be used for?

  Do not use the backup procedures to migrate to a new version or to switch over to another system with another version of Systems Director or another database. For a migration, use other tools that are described in section 4.7.3, “Migration” on page 238.

4.7.2 Backup and recovery

To protect your Systems Director 6.3.*x* data from a disaster, back up and restore your data. Use commands that are provided by Systems Director.

The following command-line tools are used for backup and recovery:

- **`smsave`** (backup)
- **`smrestore`** (restore)
- **`smreset`** (reset)

Systems Director backup

Use the **`smsave`** command to save a backup image of the Systems Director server. The command is in the `install_root\bin\` directory, where `install_root` is the root directory of your Systems Director installation.

Tip: The Systems Director server must be stopped before you run the **`smsave`** command.

A description of the options for the **`smsave`** command is at this website:

http://pic.dhe.ibm.com/infocenter/director/pubs/index.jsp?topic=%2Fcom.ibm.director.cli.helps.doc%2Ffqm0_r_cli_smsave.html

The backup image that is created by the **`smsave`** command is saved in the `install_root\backup\time_stamp` directory, unless you otherwise specified the output directory in the command.

The **`smsave`** command creates a backup image of Systems Director persistent data. Persistent data includes file system data (also called master data) and database data. The master data set contains information about the location of the database data set and uses that stored location when you run the restore operation. The database backup image is saved in the format that is specific to the database type. Backups cannot be moved from one database type and version to another database type or version.

When you run the command, the execution log is saved to `install_root\log\smave.log` and all status is updated in real time in that file. No output is posted to the command prompt.

Figure 4-192 on page 235 shows an example of using the **`smsave`** command on a Microsoft Windows 2008 R2 System.

First, you see that if the Systems Director server is not stopped before you run the **`smsave`** command, an information error is displayed. Stop the Systems Director server before you run the command by running the **`net stop dirserver`** command on Windows. For Linux, use the **`smstop`** command instead.

After you stop the Systems Director server, you can run the **smsave** command. Figure 4-192 shows the messages that you see during the procedure.

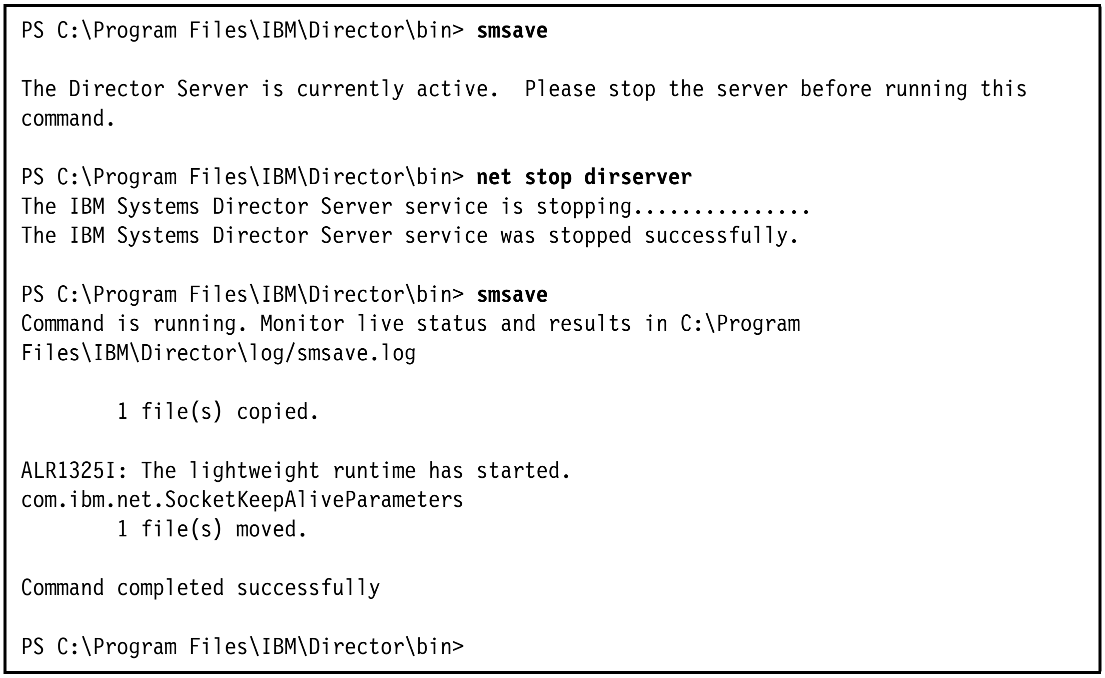

```
PS C:\Program Files\IBM\Director\bin> smsave

The Director Server is currently active.  Please stop the server before running this
command.

PS C:\Program Files\IBM\Director\bin> net stop dirserver
The IBM Systems Director Server service is stopping..............
The IBM Systems Director Server service was stopped successfully.

PS C:\Program Files\IBM\Director\bin> smsave
Command is running. Monitor live status and results in C:\Program
Files\IBM\Director\log/smsave.log

        1 file(s) copied.

ALR1325I: The lightweight runtime has started.
com.ibm.net.SocketKeepAliveParameters
        1 file(s) moved.

Command completed successfully

PS C:\Program Files\IBM\Director\bin>
```

Figure 4-192 smsave command

The log file, which is created during the backup process, is in the `install_root\log` directory. Figure 4-193 shows part of an example log file.

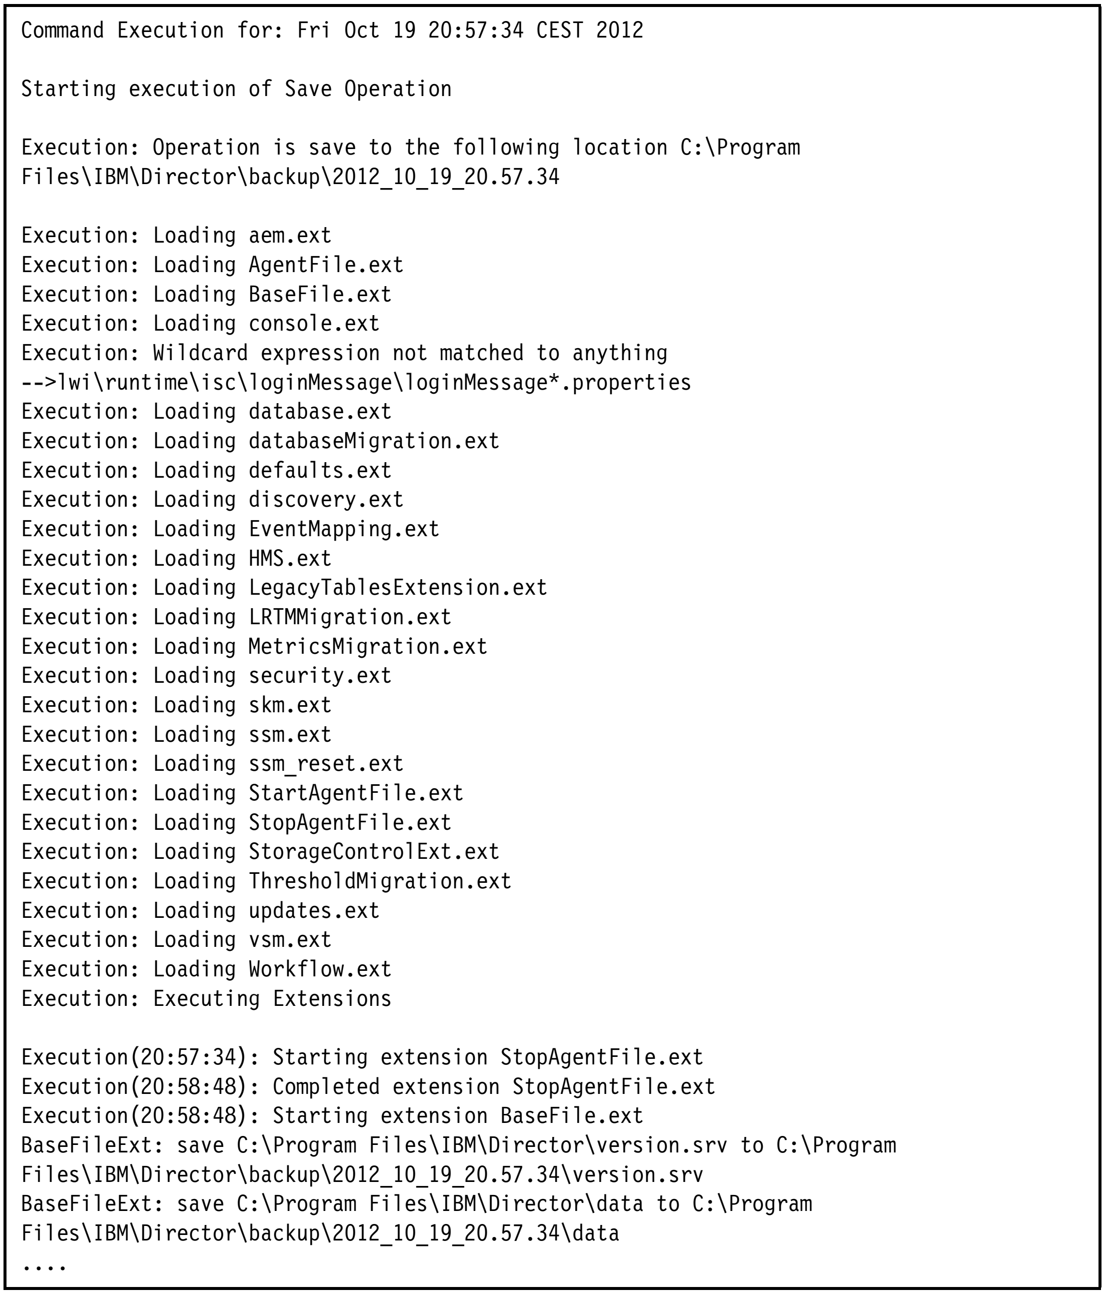

```
Command Execution for: Fri Oct 19 20:57:34 CEST 2012

Starting execution of Save Operation

Execution: Operation is save to the following location C:\Program
Files\IBM\Director\backup\2012_10_19_20.57.34

Execution: Loading aem.ext
Execution: Loading AgentFile.ext
Execution: Loading BaseFile.ext
Execution: Loading console.ext
Execution: Wildcard expression not matched to anything
-->lwi\runtime\isc\loginMessage\loginMessage*.properties
Execution: Loading database.ext
Execution: Loading databaseMigration.ext
Execution: Loading defaults.ext
Execution: Loading discovery.ext
Execution: Loading EventMapping.ext
Execution: Loading HMS.ext
Execution: Loading LegacyTablesExtension.ext
Execution: Loading LRTMMigration.ext
Execution: Loading MetricsMigration.ext
Execution: Loading security.ext
Execution: Loading skm.ext
Execution: Loading ssm.ext
Execution: Loading ssm_reset.ext
Execution: Loading StartAgentFile.ext
Execution: Loading StopAgentFile.ext
Execution: Loading StorageControlExt.ext
Execution: Loading ThresholdMigration.ext
Execution: Loading updates.ext
Execution: Loading vsm.ext
Execution: Loading Workflow.ext
Execution: Executing Extensions

Execution(20:57:34): Starting extension StopAgentFile.ext
Execution(20:58:48): Completed extension StopAgentFile.ext
Execution(20:58:48): Starting extension BaseFile.ext
BaseFileExt: save C:\Program Files\IBM\Director\version.srv to C:\Program
Files\IBM\Director\backup\2012_10_19_20.57.34\version.srv
BaseFileExt: save C:\Program Files\IBM\Director\data to C:\Program
Files\IBM\Director\backup\2012_10_19_20.57.34\data
....
```

Figure 4-193 smsave.log file

The data from the backup process is saved to the `install_root\backup` directory. The size depends on several components:

- Number of systems that are discovered
- Inventory that is collected
- Update packages that are downloaded
- All other settings

In our simple lab tests, the initial backup that we performed after installation and an inventory run is about 60 MB in size.

Systems Director restore

Use the **smrestore** command to restore the persistent data, including file system (master) data and databases, from a backup image.

You can run the **smrestore** command locally from the management server. Or, run the command remotely by accessing the management server by using a remote access utility, such as SSH or Telnet.

To run the **smrestore** command, go to the `install_root\bin` directory, where `install_root` is the root directory of your Systems Director installation. The Systems Director server must be stopped before you run the **smrestore** command.

A description of the options for the **smrestore** command is at this website:

http://pic.dhe.ibm.com/infocenter/director/pubs/index.jsp?topic=%2Fcom.ibm.director.cli.helps.doc%2Ffqm0_r_cli_smrestore.html

You can restore saved persistent data only on a management server with the same characteristics:

- Same operating system
- Same version of the Systems Director server from which the data was backed up
- Same database type and version

In addition, the Systems Director server and the database that you restore must be the same as the saved installation instances.

When you run the command, the execution log is recorded in the `install_root\log\smrestore.log` file, and all status is updated in real time in that file. Little information is posted to the command prompt.

Figure 4-194 shows the output from the command.

```
PS C:\Program Files\IBM\Director\bin> smrestore -sourceDir 'C:\Program
Files\IBM\Director\backup\201
2_10_19_20.57.34'
This operation will replace all current data with the specified backup set.
To continue, type "1" for yes or "0" for no.
1

Command is running. Monitor live status and results in C:\Program
Files\IBM\Director\log/smrestore.log

        1 file(s) copied.

ALR1325I: The lightweight runtime has started.
com.ibm.net.SocketKeepAliveParameters
        1 file(s) moved.

Command completed successfully

PS C:\Program Files\IBM\Director\bin>
```

Figure 4-194 smrestore command

Systems Director reset

The **smreset** command reinitializes the databases and clears all persistent data. The **smreset** command deletes local data on the file system where Systems Director is installed. The

smreset command deletes and rebuilds all database tables that are used by Systems Director.

A description of the options for the **smreset** command is at this website:

http://pic.dhe.ibm.com/infocenter/director/pubs/index.jsp?topic=%2Fcom.ibm.director.cli.helps.doc%2Ffqm0_r_cli_smreset.html

Use the **smreset** command to return the Systems Director server to its installation default values. This command must also be run immediately after you run the **cfgdbcmd** command to change to a new database. For example, run the **smreset** command when you upgrade from a managed IBM DB2 database to an enterprise database such as Oracle Database. Run the **smreset** command only when the Systems Director server is stopped.

The **smreset** command does not delete or reset Agent Manager information. The **smreset** command deletes the following data:

- Discovered resource data (except for 6.*x* Common Agents that were previously accessed)
- Inventory data
- Event data (event log, custom event filters, custom event actions, and custom event plans)
- Monitoring data
- Update data
- Status data
- Configuration templates
- Security configurations
- All other data that is associated with running and configuring Systems Director after the installation

The **smreset** command creates two log files, `smreset.log` and `reset.log`, which are in the `install_root\log` directory.

Figure 4-195 shows the output from the **smreset** command.

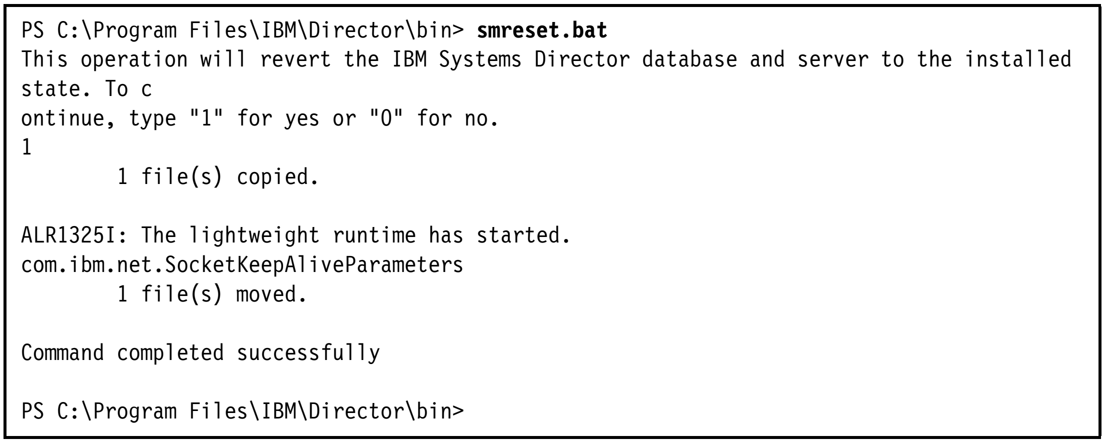

```
PS C:\Program Files\IBM\Director\bin> smreset.bat
This operation will revert the IBM Systems Director database and server to the installed
state. To c
ontinue, type "1" for yes or "0" for no.
1
        1 file(s) copied.

ALR1325I: The lightweight runtime has started.
com.ibm.net.SocketKeepAliveParameters
        1 file(s) moved.

Command completed successfully

PS C:\Program Files\IBM\Director\bin>
```

Figure 4-195 smreset command

4.7.3 Migration

The tools that can be used to migrate from one version of a Systems Director to a newer version or to another system are described.

A backup and restore process can be used only if the management server, where a backup is restored, runs the same operating system. And the version of the Systems Director server must be the same as the version from which the data was saved. The database type and version must be the same. In addition, the Systems Director server and the database that you restore must be the same as the saved installation instances.

You can switch to another operating system, another system, or another database and take the settings from the former Systems Director server with you. Use the commands in Table 4-10 to export and import settings and managed endpoints.

Table 4-10 Command-line tools for migration

Command	Description
`dircli lsmo`[a]	List managed objects - replaced by **`smcli lssys`**.
`dircli mkmo`[a]	Make managed objects.
`smcli lsevtautopln`	List information about event automation plans. You can also export one or more event automation plans to a file.
`smcli mkevtautopln`	Create an event automation plan or import one or more existing event automation plans.

[a] Before you use this command, issue the **`set CLILEGACY=1`** command.

For the complete list of the **`smcli`** commands and a description for each command, see the Systems Director Information Center:

http://pic.dhe.ibm.com/infocenter/director/pubs/index.jsp?topic=%2Fcom.ibm.director.cli.helps.doc%2Ffqm0_r_cli_smcli.html

Exporting systems and settings

How to export the systems and settings (for example, event automation plans) from an existing Systems Director server by using command-line tools is described.

Exporting systems

Use the following command for exporting the existing systems to a file. This information can be used later on a new Systems Director server to import the systems without discovery:

```
./smcli lssys -t "OperatingSystem" -d ";" -A
"Displayname,IPv4Address,Hostname,ManagementSoftware" >>OS.txt
```

In our lab example, the contents of the file that is created are shown in Figure 4-196.

```
9.42.171.194: 9.42.171.194;{ '9.42.171.194' };{ };{ '' }
9.42.171.195: 9.42.171.195;{ '9.42.171.195' };{ };{ '' }
9.42.171.196: 9.42.171.196;{ '9.42.171.196' };{ };{ '' }
9.42.171.197: 9.42.171.197;{ '9.42.171.197' };{ };{ '' }
9.42.171.198: 9.42.171.198;{ '9.42.171.198' };{ };{ '' }
9.42.171.199: 9.42.171.199;{ '9.42.171.199' };{ };{ '' }
9.42.171.203: 9.42.171.203;{ '9.42.171.203' };{ };{ '' }
9.42.171.22: 9.42.171.22;{ '9.42.171.22' };{ };{ '' }
9.42.171.23: 9.42.171.23;{ '9.42.171.23' };{ };{ 'IBM-IBM Director Agent-v6.3.2',
'IBM-IBM Director Platform Agent-v6.3.2' }
9.42.171.232: 9.42.171.232;{ '9.42.171.249', '9.42.171.232' };{ };{ ''}
9.42.171.244: 9.42.171.244;{ '9.42.171.244' };{ };{ '' }
9.42.171.25: 9.42.171.25;{ '9.42.171.25' };{ };{ '' }
9.42.171.254: 9.42.171.254;{ '9.42.171.254' };{ };{ '' }
9.42.171.26: 9.42.171.26;{ '9.42.171.32', '9.42.171.26', '9.42.171.29', '9.42.171.30',
'9.42.171.31', '9.42.171.33', '9.42.171.34' };{ };{ '' }
9.42.171.27: 9.42.171.27;{ '9.42.171.27' };{ };{ '' }
9.42.171.28: 9.42.171.28;{ '9.42.171.28' };{ };{ '' }
9.42.171.40: 9.42.171.40;{ '9.42.171.40' };{ };{ '' }
9.42.171.54: 9.42.171.54;{ '9.42.171.55', '9.42.171.54', '9.42.171.56' };{ };{ '' }
9.42.171.60: 9.42.171.60;{ '9.42.171.60' };{ };{ '' }
9.42.171.62: 9.42.171.62;{ '9.42.171.62' };{ };{ '' }
9.42.171.82: 9.42.171.82;{ '9.42.171.82' };{ };{ '' }
9.42.171.86: 9.42.171.86;{ '9.42.171.86' };{ };{ 'IBM-IBM Director Agent-v6.3.2',
'IBM-IBM Director Platform Agent-v6.3.2' }
9.42.171.97: 9.42.171.97;{ '9.42.171.97' };{ };{ 'IBM-IBM Director Core
Services-v6.2.1.2', 'IBM-IBM Director Agent-v6.2.1' }
9.42.171.99: 9.42.171.99;{ '9.42.171.99' };{ };{ 'IBM-IBM Director Agent-v6.3.2',
'IBM-IBM Director Core Services-v6.3.2' }
SLES11: SLES11;{ '9.42.171.84' };{ 'SLES11' };{ 'IBM-IBM Director Agent-v6.3.2',
'IBM-IBM Director Platform Agent-v6.3.2' }
```

Figure 4-196 Output of the smcli lssys command

Exporting settings

Use the Systems Director command-line tools to export settings, such as event automation plans and groups. You can export settings in two ways:

- Use the **smcli** command line, for example, to export event automation plans (Figure 4-197 on page 241).
- Use export functions in the Systems Director web interface (exporting groups as shown in Figure 4-198 on page 241).

To export the event automation plan, use the **smcli lsevtautopln** command. In our lab example, Figure 4-197, we exported the “Send eMail to Admin” event automation plan to the `EAPexport.xml` file. Use the **-o** attribute for an easier export because you can use the object identifier (OID) instead of the complete name.

```
SLES11:/opt/ibm/director/bin # ./smcli lsevtautopln
Log All Events
Send eMail to Admin
SLES11:/opt/ibm/director/bin # ./smcli lsevtautopln -o
Log All Events, 0x11
Send email to Admin, 0x11
SLES11:/opt/ibm/director/bin # ./smcli lsevtautopln -F xml 0x11 /tmp/EAPexport.xml
```

Figure 4-197 Export event automation plan

To export groups, you can use the Systems Director web interface. Start from the Resource Explorer Groups view and select the group that you want to export (Figure 4-198).

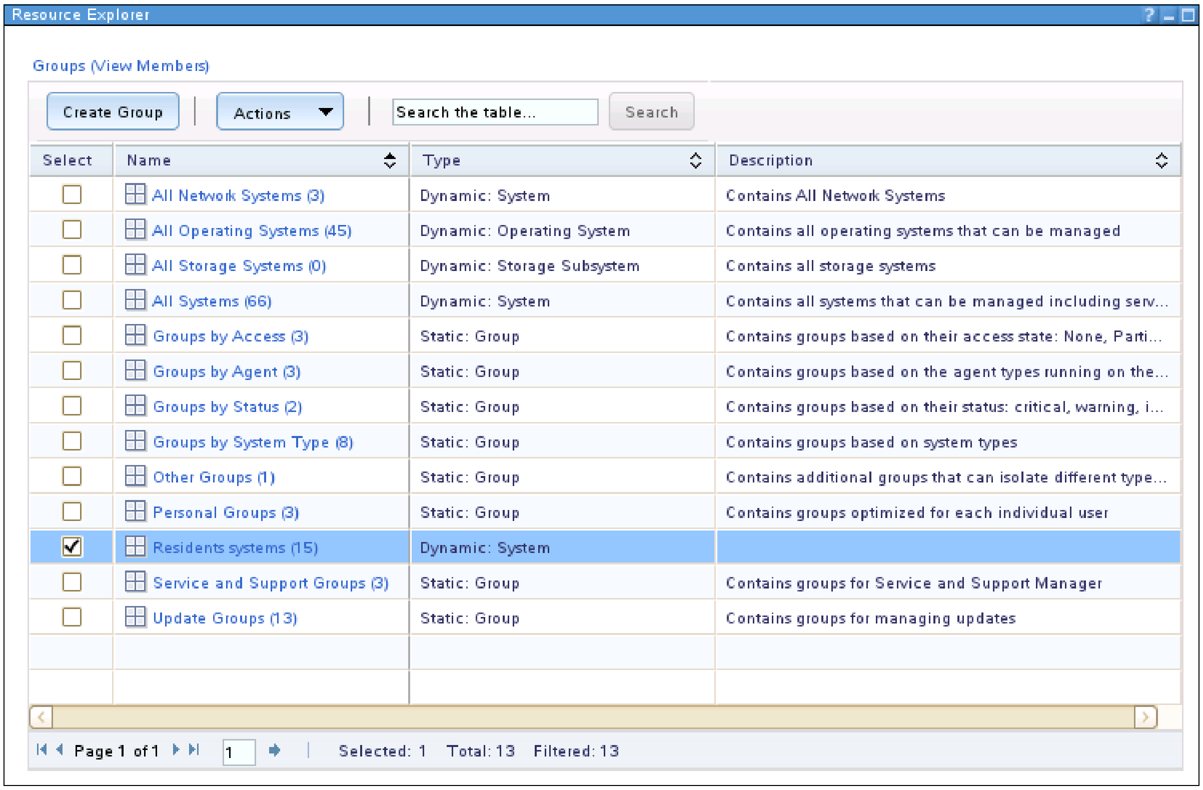

Select	Name	Type	Description
☐	All Network Systems (3)	Dynamic: System	Contains All Network Systems
☐	All Operating Systems (45)	Dynamic: Operating System	Contains all operating systems that can be managed
☐	All Storage Systems (0)	Dynamic: Storage Subsystem	Contains all storage systems
☐	All Systems (66)	Dynamic: System	Contains all systems that can be managed including serv...
☐	Groups by Access (3)	Static: Group	Contains groups based on their access state: None, Parti...
☐	Groups by Agent (3)	Static: Group	Contains groups based on the agent types running on the...
☐	Groups by Status (2)	Static: Group	Contains groups based on their status: critical, warning, i...
☐	Groups by System Type (8)	Static: Group	Contains groups based on system types
☐	Other Groups (1)	Static: Group	Contains additional groups that can isolate different type...
☐	Personal Groups (3)	Static: Group	Contains groups optimized for each individual user
☑	Residents systems (15)	Dynamic: System	
☐	Service and Support Groups (3)	Static: Group	Contains groups for Service and Support Manager
☐	Update Groups (13)	Static: Group	Contains groups for managing updates

Page 1 of 1 | 1 | Selected: 1 Total: 13 Filtered: 13

Figure 4-198 Groups in the Resource Explorer view

With the group selected, click **Actions** → **Export Group**.

The window that is shown in Figure 4-199 opens. In this view, you can select to which directory you want to save the data. The file is named `group _%username%.xml`. The *username* is the user that is logged on and creates the export file.

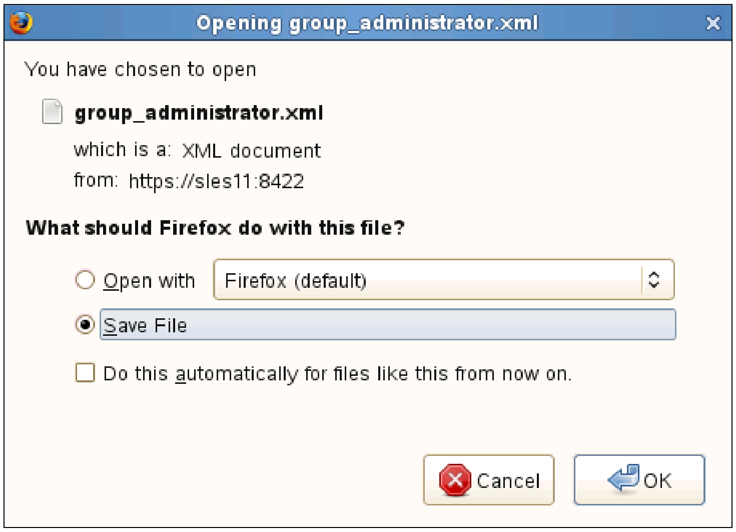

Opening group_administrator.xml

You have chosen to open

group_administrator.xml

which is a: XML document

from: https://sles11:8422

What should Firefox do with this file?

○ Open with Firefox (default)

● Save File

□ Do this automatically for files like this from now on.

Figure 4-199 Save group data

You can use the file to import the group and the members of the group to another Systems Director server or use the file for recovery options.

Importing systems and settings

With Systems Director, you can use command-line tools to import systems and to request access to those systems. The command to import systems is **`smcli mkmo`** (**`dircli mkmo`**).

> **Prerequisite**: Before you run the **`smcli mkmo`** (**`dircli mkmo`**) command, issue the following command:
>
> ```
> set CLILEGACY=1
> ```

Use the **`smcli mkmo`** command to create a managed object for the server and systems. The server represents the hardware service processor and the systems represent the operating system and agent.

In Figure 4-200, we first check whether the system is available. Then, we remove the system and check again whether the system exists. Then, we show the settings for **mkmo** and add the system to the Systems Director server again with the **smcli mkmo** command. We check again whether the system exists.

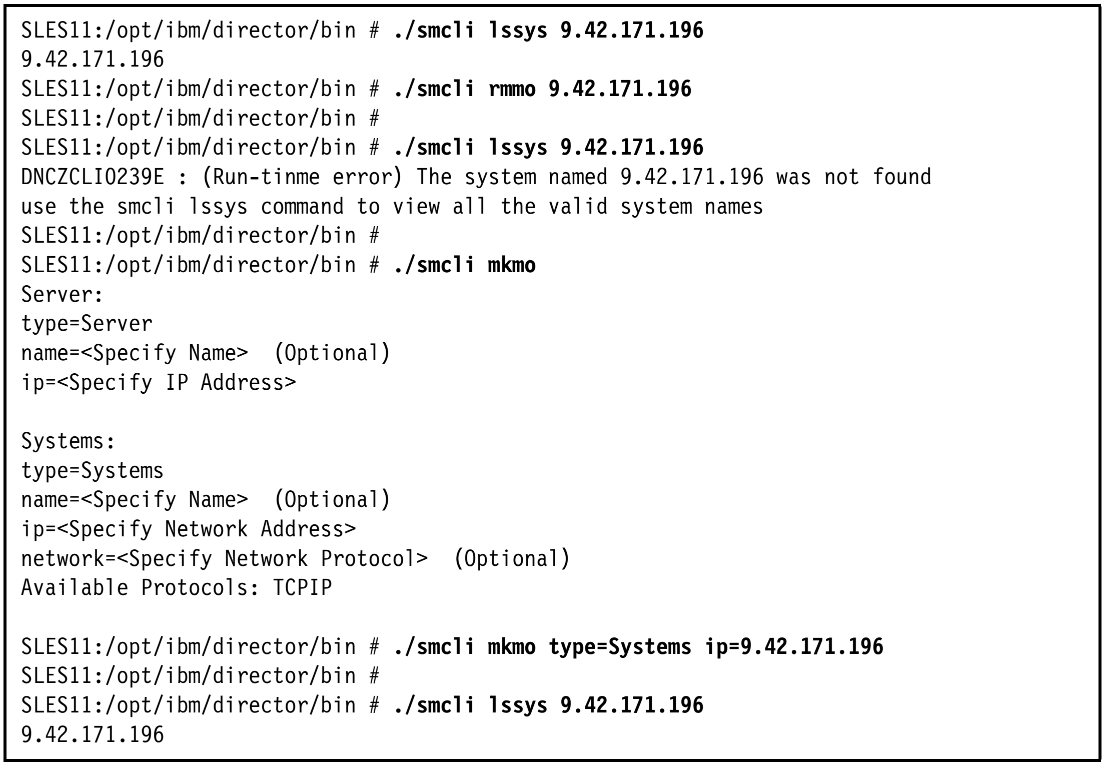

```
SLES11:/opt/ibm/director/bin # ./smcli lssys 9.42.171.196
9.42.171.196
SLES11:/opt/ibm/director/bin # ./smcli rmmo 9.42.171.196
SLES11:/opt/ibm/director/bin #
SLES11:/opt/ibm/director/bin # ./smcli lssys 9.42.171.196
DNCZCLI0239E : (Run-tinme error) The system named 9.42.171.196 was not found
use the smcli lssys command to view all the valid system names
SLES11:/opt/ibm/director/bin #
SLES11:/opt/ibm/director/bin # ./smcli mkmo
Server:
type=Server
name=<Specify Name>  (Optional)
ip=<Specify IP Address>

Systems:
type=Systems
name=<Specify Name>  (Optional)
ip=<Specify Network Address>
network=<Specify Network Protocol>  (Optional)
Available Protocols: TCPIP

SLES11:/opt/ibm/director/bin # ./smcli mkmo type=Systems ip=9.42.171.196
SLES11:/opt/ibm/director/bin #
SLES11:/opt/ibm/director/bin # ./smcli lssys 9.42.171.196
9.42.171.196
```

Figure 4-200 smcli mkmo command

You can also run the **smcli mkmo** command in a script. With a script, you can add many systems to the new Systems Director at the same time.

If you saved groups, you can import them to a new system. From the Resource Explorer view, follow these steps:

1. Click **Actions** → **Import Groups**, as shown in Figure 4-201.

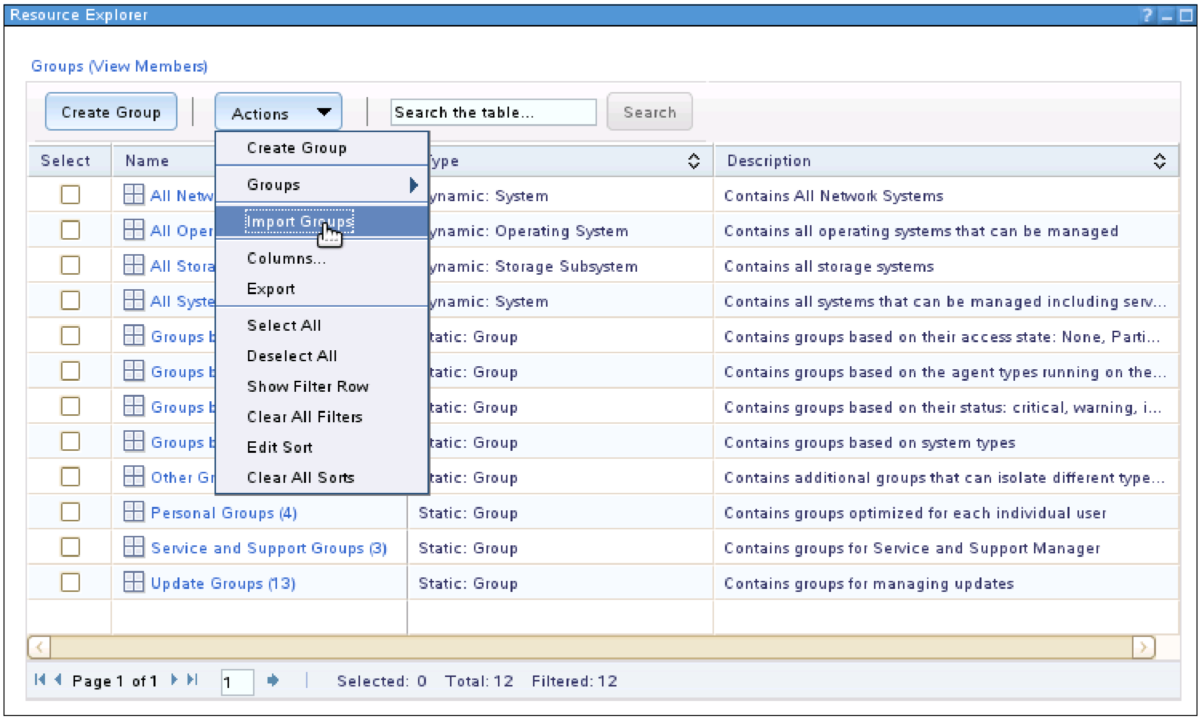

Select	Name	ype	Description
□	All Netw	ynamic: System	Contains All Network Systems
□	All Oper	ynamic: Operating System	Contains all operating systems that can be managed
□	All Stora	ynamic: Storage Subsystem	Contains all storage systems
□	All Syste	ynamic: System	Contains all systems that can be managed including serv...
□	Groups b	tatic: Group	Contains groups based on their access state: None, Parti...
□	Groups b	tatic: Group	Contains groups based on the agent types running on the...
□	Groups b	tatic: Group	Contains groups based on their status: critical, warning, i...
□	Groups b	tatic: Group	Contains groups based on system types
□	Other Gr	tatic: Group	Contains additional groups that can isolate different type...
□	Personal Groups (4)	Static: Group	Contains groups optimized for each individual user
□	Service and Support Groups (3)	Static: Group	Contains groups for Service and Support Manager
□	Update Groups (13)	Static: Group	Contains groups for managing updates

Page 1 of 1 | 1 | Selected: 0 Total: 12 Filtered: 12

Figure 4-201 Import Groups selection

2. Figure 4-202 opens where you can browse for the XML file that contains the group information.

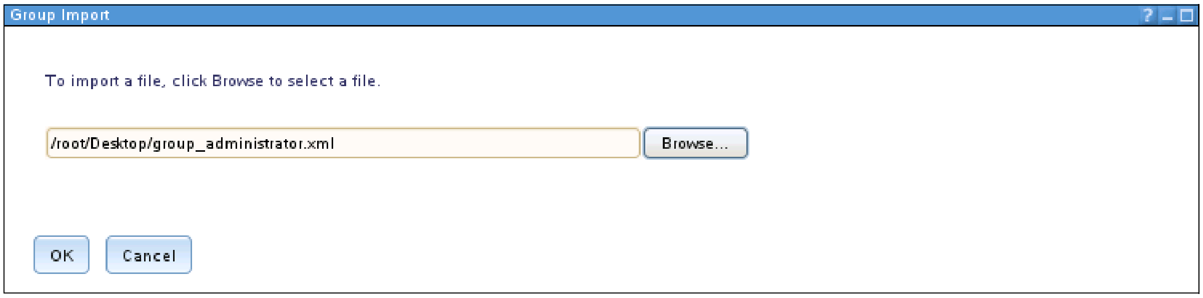

To import a file, click Browse to select a file.

/root/Desktop/group_administrator.xml

Figure 4-202 Select the file with the group information

3. The imported groups show under **Groups** → **Personal Groups**, as shown in Figure 4-203.

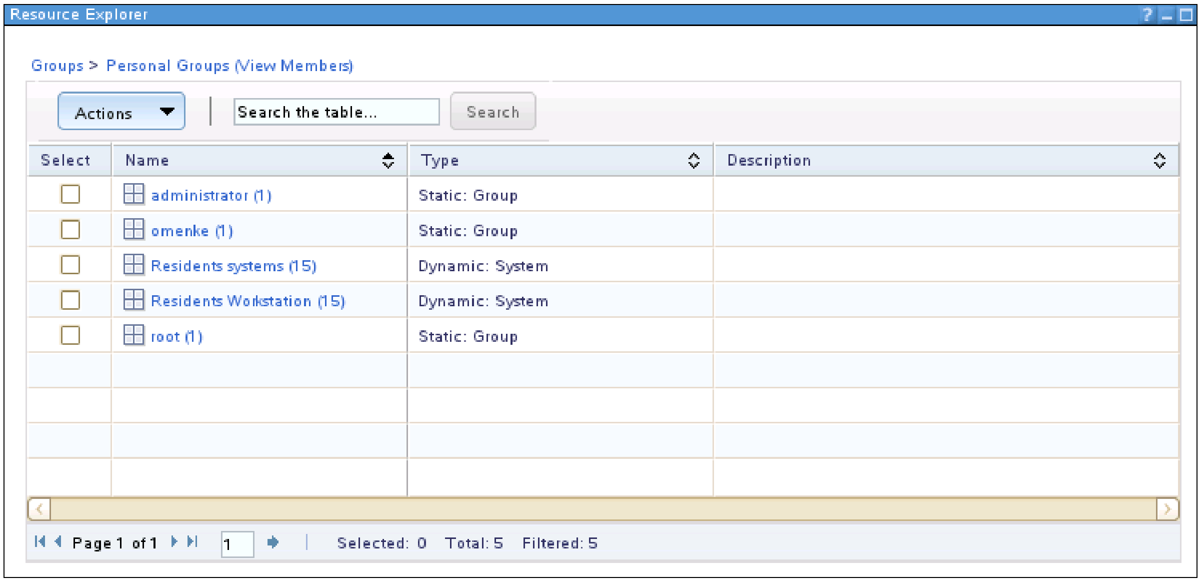
Resource Explorer
Groups > Personal Groups (View Members)
Actions
Search the table...
Search

Select	Name	Type	Description
□	administrator (1)	Static: Group	
□	omenke (1)	Static: Group	
□	Residents systems (15)	Dynamic: System	
□	Residents Workstation (15)	Dynamic: System	
□	root (1)	Static: Group	

Page 1 of 1
1
Selected: 0 Total: 5 Filtered: 5

Figure 4-203 Groups: Personal Groups view

To import a saved event automation plan, run the **smcli mkevtautopln /tmp/EAPexport.xml** command, as listed in Figure 4-204. In our example, we use the event automation plan that we exported before to the EAPexport.xml file.

```
SLES11:/opt/ibm/director/bin # ./smcli mkevtautopln /tmp/EAPexport.xml
Warning Number: 1
DNZEAP2068W: (Run-time warning)
The IP address or host name 'smtp.itso.ral.ibm.com' is not accessible.

Action name: eMail to Admin
Element name: EmailSmtpServer
Element value: smtp.itso.ral.ibm.com

Warning Number: 2
DNZEAP2059W: (Run-time warning)
The event filter named 'Critical Events' has the same name and definition as an existing
filter in the system.
The filter will be not be created again.

Filter name: Critical Events

Total number of warnings: 2

DNZEAP2064I: (Informational) Created event action 'eMail to Admin'.
DNZEAP2066I: (Informational) Created event automation plan 'Send eMail to Admin'.
DNZEAP2067I: (Informational) Targets 'All Systems', applied to event automation plan
'Send eMail to Admin'.
SLES11:/opt/ibm/director/bin #
```

Figure 4-204 Importing an event automation plan

During the import process, the event automation plan is checked. Warnings display if incorrect settings exist in the event automation plan or if event actions or filters exist on the system. The existing filter or event action is not created again. You can see that the event Action plan “Send eMail to Admin” is created. All systems, predefined in the event automation plan that we exported, are assigned to this event automation plan.

4.7.4 Best practices for backup

Keep in mind the following when managing IBM Systems Director backups:

- – Define a backup policy and configure it on your server. The **smsave** command must be used when the IBM Systems Director server is down (**smstop** command).
- – Before each backup, you should have a clean and recent inventory.
- – If you need to reset the IBM Systems Director server by using the **smreset** command, you should make a backup before doing the reset.
- – If you need to restore, use the **smrestore** command. The **smrestore** works only if the server has the same configuration, which has been taken by the **smsave** command (same operating system, same database, same version of IBM System Director).

5

VMControl

This chapter provides an overview of best practices when using IBM Systems Director VMControl. IBM Systems Director VMControl is a cross-platform product that assists you in rapidly deploying virtual appliances to create virtual servers that are configured with the operating system and software applications that you want. It also enables you to group resources into system pools, which enable you to centrally manage and control the different workloads in your environment.

This chapter contains the following topics:

5.1 About VMControl

To implement IBM Systems Director VMControl, you need to plan where you are going to use it and what purpose it should fill.

IBM Systems Director VMControl gives you the opportunity to control multi-hypervisor environments. However, some IBM components on upper layers of management stack give more advanced control features to hypervisors and especially deployment-related tasks. IBM Systems Director VMControl Enterprise version provides the most advanced features for Kernel-based Virtual Machine (KVM) on Red Hat Enterprise Linux (RHEL) and for IBM PowerVM hypervisor platforms.

Figure 5-1 shows a configuration with IBM Systems Director and IBM Systems Director VMControl where IBM Systems Director VMControl has control in all architectures and hypervisors.

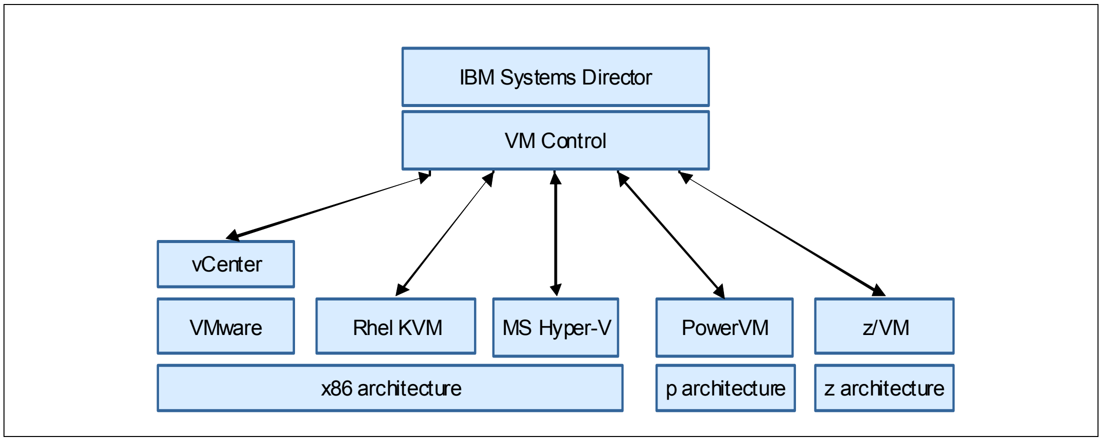

Figure 5-1 IBM Systems Director VMControl and hypervisors

Use this configuration when there is no need for self-service portal or advanced hypervisor-specific virtual machine deployment features.

Figure 5-2 on page 249 shows a configuration with IBM Systems Director, IBM Systems Director VMControl, and IBM SmartCloud® Entry (SCE). SCE communicates with IBM Systems Director VMControl and also manages VMware through IBM Systems Director VMControl (VMC) or vCenter application programming interface (API).

IBM SmartCloud Entry provides extra features like self service portal for end users, IP address pools, billing, and metering. When managing VMware in a cloud environment, it is a good practice to use SCE to communicate directly with vCenter API, as shown in Figure 5-2.

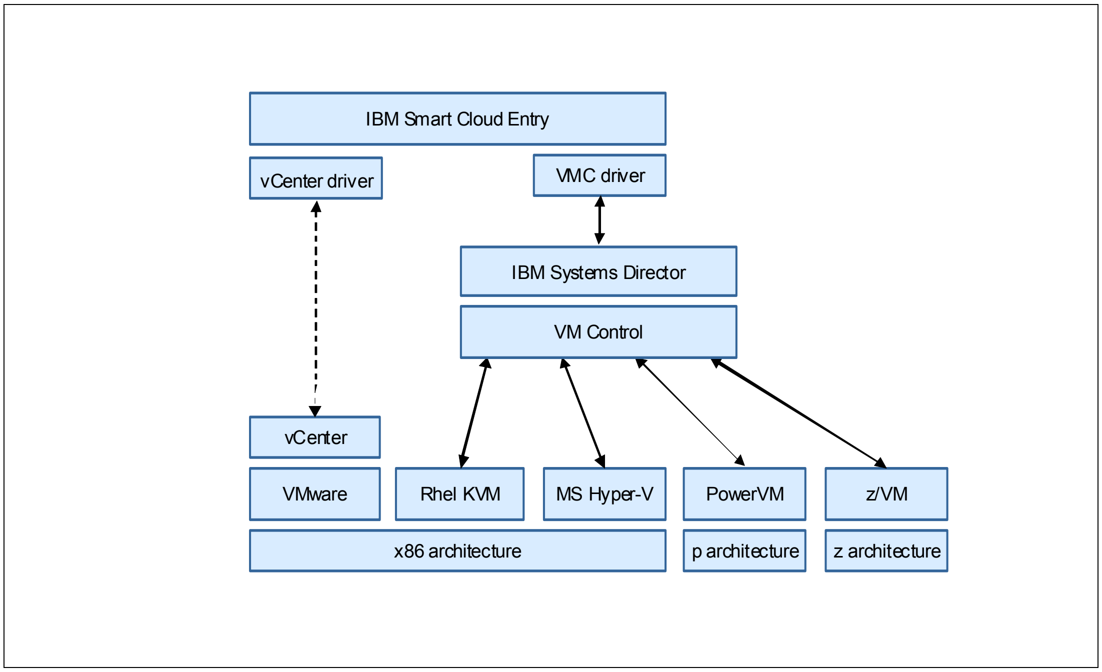

Figure 5-2 IBM Systems Director, VMControl, and IBM SmartCloud Entry

For more information about IBM SmartCloud Entry, see:

https://www.ibm.com/developerworks/community/wikis/home?lang=en#/wiki/W21ed5ba0f4a9_46f4_9626_24cbbb86fbb9/page/Documentation

5.2 Understanding the components of a VMControl environment

IBM Systems Director VMControl allows you to manage various components in your virtualization environment. VMControl has several components that need to be clarified before the KVM environment can be designed and configured. The following sections describe these components.

5.2.1 Platform managers

A platform manager manages one or more host systems and their associated virtual servers and operating systems. Examples of platform managers include:

- Hardware Management Console (HMC)
- IBM Flex System Manager™
- Integrated Virtualization Manager (IVM)
- VMware vCenter

IBM Systems Director does not recognize a managed system as a platform manager until the managed system has been unlocked. (The `padlock` icon in the Access column for a managed system indicates that it is secured.) To request access to the managed system, right-click the managed system and click **Request Access**. By providing a valid user name that has local administrative rights to that managed system and its password, you can unlock and access the system. In a KVM environment, IBM Systems Director VMControl behaves as a platform manager.

5.2.2 Hosts

In an IBM Systems Director environment, a host is a system that contains resources from which virtual servers are constructed. Hosts can be any of the following systems that are configured for the IBM Systems Director environment:

- A BladeCenter chassis
- IBM Flex System Enterprise Chassis
- An RHEL 6.0, 6.1, 6.2, or 6.3 server that has KVM virtualization support enabled
- IBM Power Systems that are under the control of an IBM HMC
- An IBM Power Systems server that is under the control of IBM IVM
- A system running VMware ESX Server or VMware ESXi that is under the control of VMware vCenter
- Windows Server 2008, Enterprise, Standard, and Datacenter x64 Editions with Hyper-V role enabled, Release 2
- A System z logical partition (LPAR) running z/VM hypervisor
- A Power Systems compute node

A host can manage multiple virtual servers and their guest operating systems.

5.2.3 Virtual servers or virtual machines

A virtual server is associated with a host system. The host must be part of a virtualization environment that is supported in IBM Systems Director. An operating system and other software can be installed on a virtual server. In a Power Systems environment, a virtual server is called a *logical partition* or *partition*. In z/VM and VMware environments, virtual servers are often called *virtual machines*.

A virtual server is the logical equivalent of a physical platform. After IBM Systems Director discovers a host, it continues the discovery process for all the virtual servers that are associated with the host. After virtual servers are discovered, they can be powered on and turned off through IBM Systems Director. In addition, you can edit resources that are assigned to virtual servers, and in some virtualization environments, you can relocate a virtual server from one host to another. You can also create additional virtual servers to meet your needs.

5.2.4 Guest-operating-systems

A guest-operating-system represents an operating system that is running on a virtual server on which Common Agent is installed or it has been discovered as agentless. A guest-operating-system is a particular type of managed system. The standard IBM Systems Director discovery process for managed systems can discover guest operating systems. However, if a guest operating system is not running Common Agent, it is not recognized as an agentless guest-operating-system object in IBM Systems Director.

5.2.5 Virtual farms

A virtual farm logically groups like-hosts and facilitates the relocation task — moving a virtual server from one host to another host within the virtual farm. A virtual farm can contain multiple hosts and their associated virtual servers. A virtual farm can contain only hosts of the same type. For example, a virtual farm that begins with a KVM host can contain only other KVM hosts. When a virtual farm is configured, you can relocate virtual servers between hosts in the farm.

You use the Create Virtual Farm wizard to group hosts together and enable specialized capabilities for the virtual servers running on the hosts. You can enable capabilities such as high availability, workload management, live relocation, and static relocation. Not all capabilities are supported on all platforms.

5.2.6 Virtual appliances

A virtual appliance is a ready-to-deploy operating system and software package that is stored by IBM Systems Director VMControl. A virtual appliance contains an image of a full operating system, and can contain software applications and middleware. A virtual appliance also contains metadata describing the virtual server that the image requires.

5.2.7 Workloads

A workload represents one or more virtual servers that can be monitored and managed as a single entity. For example, you can manage a workload that might contain both a web server and a database server. You can start and stop a workload, and thus the virtual servers it contains, as one entity. You can monitor the overall state and status of the workload by viewing the Workloads dashboard. A workload is automatically created when you deploy a virtual appliance. You can also create a workload by grouping one or more virtual servers that are not already part of an existing workload.

5.2.8 System pools

System pools group similar resources so that you can manage the resources within the system pool as a single unit. You can create storage system pools and server system pools. You can perform basic tasks such as creating and deleting system pools and adding and removing resources from the system pools. Additional system pool functions like Networking Pools are available, depending on the type of system pool that you created.

5.3 Managing KVM on Red Hat Enterprise Linux with VMControl

This section provides an overview of best practices when using IBM Systems Director VMControl with KVM on Red Hat Enterprise Linux.

There are two ways to use KVM on Red Hat Enterprise Linux with IBM Systems Director:

- Use a Network File System (NFS) storage-based solution.
- Use a storage area network (SAN) storage-based solution.

For information about installing an environment by using the NFS storage-based solution, see the following site:

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

The best practice is to use a SAN storage-based solution. Although this configuration looks more complex than the NFS solution, the block storage-based model offers better performance, as well as more functionality and flexibility.

Figure 5-3 shows the vital components of IBM Systems Director and KVM on a Red Hat Enterprise Linux platform.

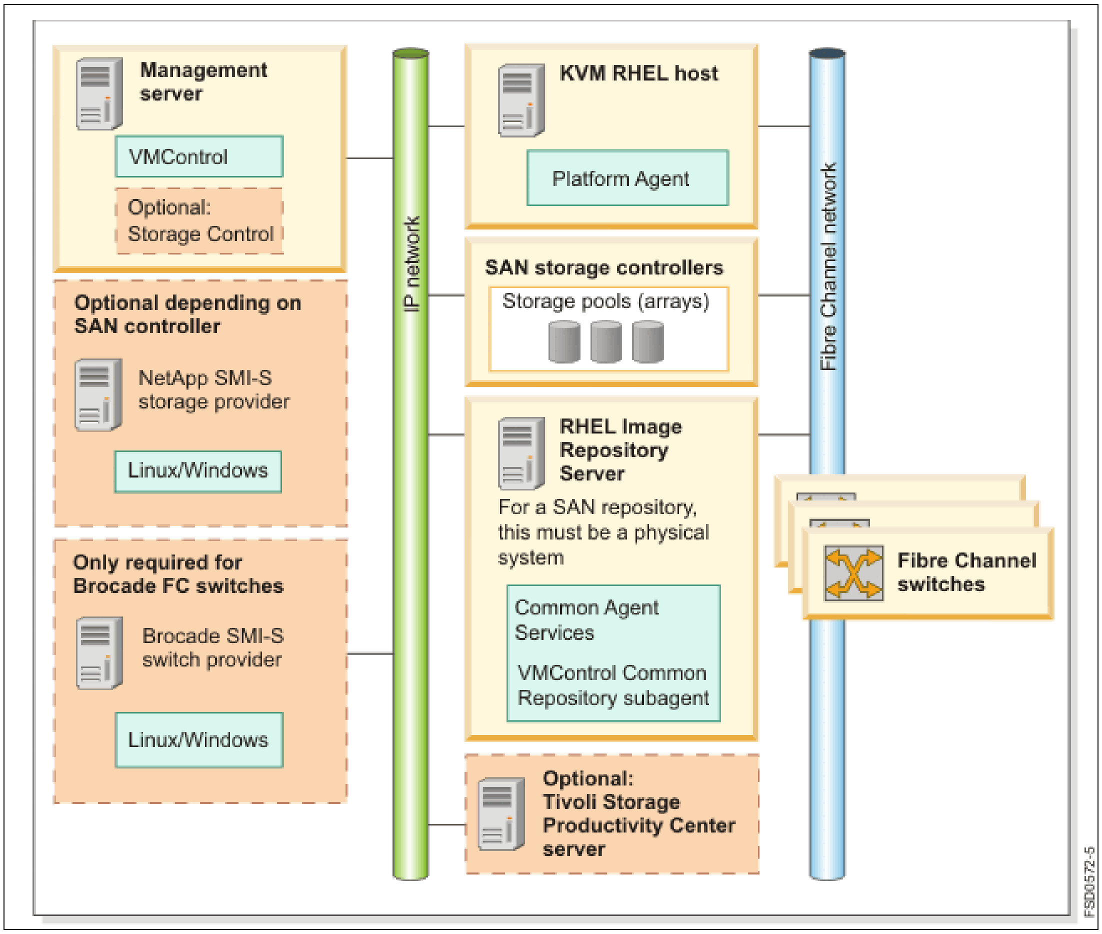
Management server
VMControl
Optional: Storage Control
Optional depending on SAN controller
NetApp SMI-S storage provider
Linux/Windows
Only required for Brocade FC switches
Brocade SMI-S switch provider
Linux/Windows
IP network
KVM RHEL host
Platform Agent
SAN storage controllers
Storage pools (arrays)
RHEL Image Repository Server
For a SAN repository, this must be a physical system
Common Agent Services
VMControl Common Repository subagent
Optional: Tivoli Storage Productivity Center server
Fibre Channel network
Fibre Channel switches
FSD0572-5

Figure 5-3 SAN storage-based KVM and VMControl environment

The details of KVM on Red Hat Enterprise Linux and IBM Systems Director Components are:

- IBM Systems Director server 6.3.2 or higher is installed on a supported server. This is the management server, which controls the KVM environment and behaves as the platform manager.
- If you are using Storage Control to manage the SAN storage, then a supported version of Storage Control 4.2.1 or higher must also be installed.
- IBM Systems Director VMControl 2.4 or greater is activated.
- Fibre Channel network for storage is in place. KVM virtualization with VMControl supports only SAN storage over Fibre Channel. It has to be correctly cabled and configured with the appropriate Fibre Channel switches. 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 Platform Agent for KVM is downloaded and installed.

- KVM hosts are discovered, accessed, and inventoried from your IBM Systems Director server.
- The SAN storage controllers (also called *storage subsystems*) are configured and storage pools are set up with the storage space and Redundant Array of Independent Disks (RAID) levels that you want for virtual disk images. VMControl and Storage Control does not provision these RAID storage pools for you. Best practice is to use IBM SAN Volume Controller (SVC) or IBM Storwize v7000 for Fibre Channel-based SAN storage subsystems.

  For information about the supported storage controllers with VMControl and KVM, see:

  http://pic.dhe.ibm.com/infocenter/director/pubs/index.jsp?topic=%2Fcom.ibm.director.vim.helps.doc%2Ffsd0_vim_c_learnmore_repositories_kvm.html
- A Fibre Channel switch provider is configured in the environment. 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 IBM Systems Director for shared access from endpoints in the KVM environment. These endpoints include KVM hosts and image repository servers. The image repository is used for storing and deploying virtual appliances.
- The physical server with Red Hat Enterprise Linux with Common Agent with a storage subagent is installed as an image repository server and is connected to the Fibre Channel network with a supported Fibre Channel host bus adapter (HBA).

  > **Note:** Skip this step if the IBM Systems Director server is running RHEL, is connected to the SAN Fibre Channel network, and will be used as the management server and for the repository server. However, the best practice is to use an external physical server when IBM Systems Director server can be virtualized itself.

- The image repository server is discovered and inventory is collected on it.
- The image repository is created from VMControl.
- All hosts have static IP-addresses and functional Domain Name System (DNS) services.
- KVM hosts have network interfaces for virtual servers and networks have been built inside to all hosts.

5.3.1 VMControl supported tasks and limitations

VMControl provides in KVM environments the most wide scale of tasks, which are available for hypervisor and virtual server management in IBM Systems Director VMControl.

VMControl provides the following supported tasks in a KVM environment:

- 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 Systems Director Network Control with VMControl).
- Adjust the virtualization monitor polling interval for KVM by using the `KvmPlatformPollingInterval` parameter.

VMControl supports the following network configurations in a KVM environment:

- Virtual Ethernet Bridging (VEB).
- Virtual Ethernet Port Aggregator (VEPA) network (Requires IBM Systems Director Network Control and that the host is in a network system pool).

The following limitations apply when using KVM on Red Hat Enterprise Linux.

The following limitations and restrictions apply to using the KVM virtualization environment:

- Restrictions when you manage IBM BladeCenter or System x systems.

  The following templates cannot be used to configure the operating system:

  - SNMP Agent Configuration template
  - Asset ID template
- A user account cannot be copied and used to create a user. Asset information cannot be configured for a managed system.
- Do not use the RHEL Virtual Machine Manager or other means to create virtual servers or to manage them directly on a KVM managed system. IBM Systems Director server does not receive events when operations are performed outside of VMControl in these cases. Results might vary if these external interfaces are used.
- Only NFS version 3 mounts are supported by VMControl. If both your NFS server and NFS client (hosting image repository) support version 4, you might need to either use the nfsvers=3 mount option to downgrade the mount or configure your NFS server to give only version 3 mounts.
- Red Hat Enterprise Linux Version 5.5 is not supported for the image repository server that hosts SAN repositories.
- Limited support for KVM hypervisor networks.

5.3.2 KVM virtual environment considerations

When creating an IBM Systems Director VMControl KVM virtual environment, the following factors need to be considered:

- Network design. Plan your network carefully because network design is one of the important things that needs to be set up properly:
  - Reserve enough physical network interfaces and size them correctly. Plan for the amount and speed of network interface needed by the network types and services.
  - If planning an environment that requires isolated networks, use the proper technology to provide not only hypervisor isolation, but also isolation that is done inside the switch.

  - Is there a need to implement VEB features?
  - Is there a need to implement VEPA features?
  - Use Virtio and e1000 model configurations for virtual network server adapters.
- KVM virtualized environments must run on System x servers, blades, or an IBM Flex System.
- You must use the following Linux versions for the KVM virtualization environment: Red Hat Enterprise Linux version 6.0, 6.1, 6.2, or 6.3 with KVM virtualization installed.
- Use para-virtualized (Virtio) drivers for enhanced performance.

5.3.3 KVM on Red Hat Enterprise Linux

This section provides information for installing KVM on Red Hat Enterprise Linux, including the basic steps that need to be performed to prepare virtual hosts and the image repository for IBM Systems Director VMControl.

To manage the KVM on a Red Hat Enterprise Linux host from IBM Systems Director VMControl, you must manually install the RHEL KVM Platform Agent. Always use the latest available packet. For successful installation, you need a minimum of three physical servers to create one KVM cluster and one image repository server for image management, in addition to IBM Systems Director with the VMControl Standard or Enterprise Edition activated.

Preparation

Install and configure RHEL 6.3 on the compute node using the Virtualization Host role. RHEL installation is not described in this book.

Notes:

- Red Hat installation steps and information can be found at the following site:

  https://access.redhat.com/site/documentation//en-US/Red_Hat_Enterprise_Linux/index.html
- When using SAN storage, remember to configure Red Hat multipathing, as discussed at the following site:

  https://access.redhat.com/site/documentation/en-US/Red_Hat_Enterprise_Linux/6/html/DM_Multipath/mpio_setup.html

After installing Red Hat Enterprise Linux, the host needs to be prepared for the RHEL KVM Platform Agent. Perform the following steps to allow installation and communication between IBM Systems Director and the KVM Platform Agent.

1. Change `SELINUX=enforcing` to `SELINUX=disabled` to disable SELinux, as shown in Example 5-1.

Example 5-1 Disable SELinux

```
[root@node11229 ~]# cat /etc/selinux/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 - SELinux is fully disabled.
SELINUX=disabled
# SELINUXTYPE= type of policy in use. Possible values are:
```

```
#       targeted - Only targeted network daemons are protected.
#       strict - Full SELinux protection.
SELINUXTYPE=targeted
```

Tip: You can also configure SELinux in *permissive* mode if required for security reasons.

2. Configure the iptables, as shown in Example 5-2.

Example 5-2 Configure iptables so that it saves the iptables settings

```
[root@TPMfImages ~]# cp /etc/sysconfig/iptables-config
/etc/sysconfig/iptables-config.old
[root@TPMfImages ~]# rm -f /etc/sysconfig/iptables-config
[root@TPMfImages ~]# echo 'IPTABLES_MODULES=""
> IPTABLES_MODULES_UNLOAD="yes"
> IPTABLES_SAVE_ON_STOP="yes"
> IPTABLES_SAVE_ON_RESTART="yes"
> IPTABLES_SAVE_COUNTER="no"
> IPTABLES_STATUS_NUMERIC="yes"
> IPTABLES_STATUS_VERBOSE="no"
> IPTABLES_STATUS_LINENUMBERS="yes"' > /etc/sysconfig/iptables-config
[root@TPMfImages ~]# cat /etc/sysconfig/iptables-config
IPTABLES_MODULES=""
IPTABLES_MODULES_UNLOAD="yes"
IPTABLES_SAVE_ON_STOP="yes"
IPTABLES_SAVE_ON_RESTART="yes"
IPTABLES_SAVE_COUNTER="no"
IPTABLES_STATUS_NUMERIC="yes"
IPTABLES_STATUS_VERBOSE="no"
IPTABLES_STATUS_LINENUMBERS="yes"
[root@TPMfImages ~]#
```

3. Open the required TCP/UDP ports on the hosts, as shown in Example 5-3.

Example 5-3 Open ports

```
[root@ISDNode1 ~]# iptables -A INPUT -p udp --dport 427 -j ACCEPT
iptables -A INPUT -p tcp --dport 15989 -j ACCEPT
[root@ISDNode1 ~]# iptables -A INPUT -p tcp --dport 22 -j ACCEPT
[root@ISDNode1 ~]# iptables -A INPUT -p tcp --dport 15988 -j ACCEPT
[root@ISDNode1 ~]# iptables -A INPUT -p tcp --dport 15989 -j ACCEPT
[root@ISDNode1 ~]# service iptables save
iptables: Saving firewall rules to /etc/sysconfig/iptables:[  OK  ]
```

Note: Configuring iptables is needed because you might face some issues during the inventory collection. You can also temporarily disable iptables for troubleshooting purposes.

4. Configure Yum on your Red Hat Enterprise Linux system and upgrade. See Example 5-4.

   For more information about how to configure Yum, see the following site:

   http://docs.redhat.com/docs/en-US/Red_Hat_Enterprise_Linux/6/html/Deployment_Guide/sec-Configuring_Yum_and_Yum_Repositories.html

   > **Tips:**
   >
   > - Configure Yum on your system because during the KVM Platform Agent installation, you might face RPM dependency requirements. You can save time if Yum is configured.
   > - When registering the RHEL subscription, ensure that only limited updates are available. Then, upgrade RHEL with Yum. Otherwise, Yum will upgrade RHEL to the newest version, which might not yet be supported by the IBM Systems Director KVM Platform Agent.

   Example 5-4 Upgrade environment with yum

   ```
   [root@ISDNode1 network-scripts]#sudo yum upgrade
   --setopt=protected_multilib=false --skip-broken
   ```

   This upgrade prevents the most common problems with dependencies.

5. Check that the date on your KVM node is the same as the other KVM hosts and the IBM Systems Director, as shown in Example 5-5.

   Example 5-5 Check date on the KVM host and IBM Systems Director host

   ```
   [root@node7391 network-scripts]# date
   Thu Apr 18 18:25:05 EEST 2013
   ```

   > **Tip:** Remember to check the time difference between the IBM Systems Director server and KVM hosts.
   >
   > Generally, the time difference between the agents and IBM Systems Director, when adjusted for the time zone, should not be more than one hour. Ensure that the system clocks remain synchronized on the systems that contain IBM Systems Director and agents.

6. Remove any unnecessary packages by using the command that is shown in Example 5-6. If packages exist, agent installation will fail.

   Example 5-6 Remove unnecessary packages

   ```
   [root@ISDNode1 ~]# yum -y erase tog-pegasus libcmpiutil libvirt-cim
   sblim-cmpi-nfsv3 sblim-cmpi-fsvol sblim-gather-provider sblim-gather
   sblim-cmpi-base openslp
   ```

7. Create the required data networks for virtual servers in all hosts that are part of the KVM server pool. See Example 5-7 on page 258.

   > **Note:** The server that acts as the image repository does not need bridging. You can skip this step when configuring the image repository host.

Example 5-7 shows an example of a bridge network file. You can create this file by copying the `ifcfg-eth0` file to create the `ifcfg-br0` file and after copying, edit it by using the vi editor. This example has been created for one of the Ethernet devices.

Example 5-7 How a bridge configuration file should look after modifying it

```
[root@ISDNode1 network-scripts]# cat ifcfg-br0
DEVICE=br0
BOOTPROTO=static
IPADDR=10.31.21.100
NETMASK=255.255.192.0
IPV6INIT=yes
MTU=1500
ONBOOT=yes
TYPE=Bridge
DELAY=0
UUID=c15c6764-f14c-426e-960e-3e7cfb1d65f3
```

Example 5-8 shows the file after it has been modified.

Example 5-8 Ethernet config file after modifying

```
[root@ISDNode1 network-scripts]# cat ifcfg-eth0
DEVICE="eth0"
HWADDR="34:40:B5:BE:79:60"
ONBOOT="yes"
BRIDGE="br0"
```

8. Restart the network-related services on the host. See Example 5-9.

Example 5-9 Restart network services

```
[root@ISDNode1 network-scripts]# service network restart
```

> **Note:** The Secure Shell (SSH) service must be configured and running on the KVM host. This configuration ensures that an SSH remote service access point for port 22 gets created for each host in addition to the Common Information Model (CIM) RSAP on ports 15988 and 15989.

KVM Platform Agent installation

Perform the following steps to install the KVM Platform Agent:

1. Download the latest KVM Platform Agent from the following IBM website:

   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

2. Install the KVM Platform Agent prerequisites. See Example 5-10.

Example 5-10 Install required packages

```
[root@TPMfImages ~]# yum install -y libconfig.x86_64 libsysfs.x86_64
libsysfs.i686 libicu.x86_64 lm_sensors-libs net-snmp.x86_64
net-snmp-libs.x86_64 redhat-lsb.x86_64
```

3. Place the downloaded agent in the `/tmp` folder of your KVM host by using Secure Copy Protocol (SCP) and an SCP tool. Uncompress the archive after transfer. Then, start the KVM Platform Agent installation.

Tip: You might receive error messages that are related to dependencies even though you have performed the upgrade. Use Yum to solve the dependencies issue and restart the KVM platform agent installation.

When your KVM Platform Agent installation is completed, repeat these steps in all hosts that are part of the KVM virtual host environment as hypervisors.

Note: A comprehensive step by step guide for installing the KVM Platform Agent can be found in the following IBM Redbooks publication:

http://www.redbooks.ibm.com/redbooks/pdfs/sg248060.pdf

This publication is written for the IBM Flex System Manager, but the tasks in section 8.2 are also suitable for KVM Platform Agent installation for IBM Systems Director environments.

Installing KVM image repository

One physical server is needed to act as an image repository. This server handles, for example, all ISO volumes, which enable virtual server installation from disk images. This server also tracks where captured appliances are. See Figure 5-4.

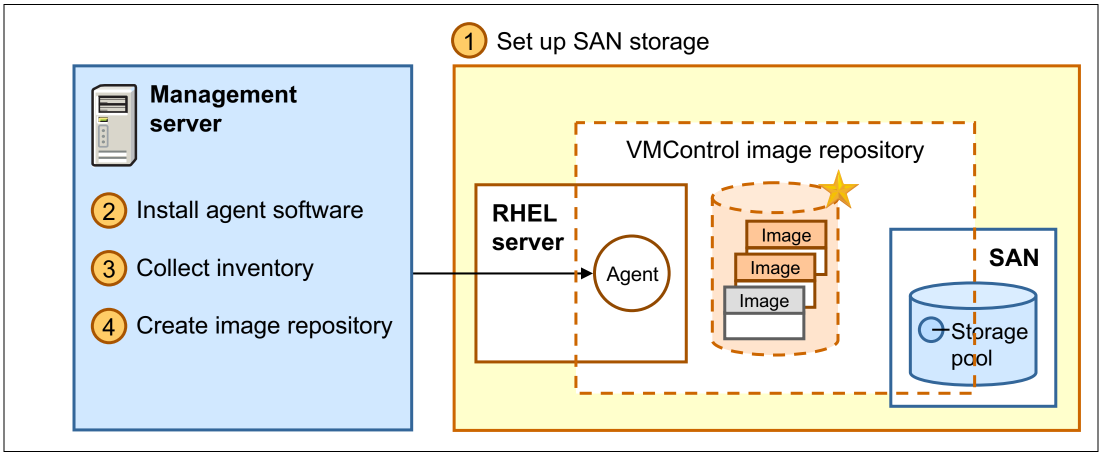

Figure 5-4 Procedure to create a SAN-based image repository

Installing the Common Agent

Perform the following steps to install the KVM Common Agent:

1. Download the latest Common Agent from the following IBM website:

   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&&dlmethod=dd

2. Install the Common Agent prerequisites. See Example 5-11.

   Example 5-11 Install required packages

```
[root@TPMfImages ~]# yum -y 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 libexpat.so.0
```

3. Turn off the iptables on the Common Agent host. See Example 5-12.

   Installing the Common Agent Services (CAS) agent correctly has been a problem in the past. More troubleshooting information for CAS agent problems is in section 12.1, “Troubleshooting the installation of IBM Systems Director components” on page 494. See Example 5-12.

Example 5-12 Turn iptables off before installing the CAS agent

```
[root@TPMfImages ~]# chkconfig iptables off
[root@TPMfImages ~]# service iptables stop
iptables: Saving firewall rules to /etc/sysconfig/iptables:[  OK  ]
iptables: Flushing firewall rules:                         [  OK  ]
iptables: Setting chains to policy ACCEPT: filter          [  OK  ]
iptables: Unloading modules:                               [  OK  ]
```

4. Place the downloaded agent in the `/tmp` folder of your KVM host by using the SCP protocol and an SCP tool. Uncompress the archive after transfer. Then, start the Common Agent installation.

> **Note:** Common Agent is installed by default in unmanaged mode. When IBM Systems Director server discovers an unmanaged agent, the agent becomes managed after requesting access.

If you want the agent to start out in managed mode, for example, because the IBM Systems Director server is configured with more than one agent manager and you want to choose the agent manager with which the agent associates, you can do so by using one of the following two methods:

- Install the agent in managed mode.

  To install the agent in managed mode, use the optional `diragent.rsp` file, as described in the following steps.

> **Tip:** More information about using the `diragent.rsp` file can be found at the following site:
>
> `http://pic.dhe.ibm.com/infocenter/director/pubs/index.jsp?topic=%2Fcom.ibm.director.install.helps.doc%2Ffqm0_t_installing_agents.html`

- Configure the agent for managed mode after installation.

  To configure the agent for managed mode after installation, run the following **`configure.sh`** command:

```
/opt/ibm/director/agent/runtime/agent/toolkit/bin/configure.sh -amhost
agentmanager_ip -passwd agentregistration_password -force
```

> **Tip:** You can use the **`configure.sh -unmanaged -force`** command to return the agent to unmanaged mode.

The agent will also change to managed mode automatically after you discover it, and request access to it through the IBM Systems Director server.

Note: A comprehensive step-by-step guide for installing Common Agent can be found at the following site:

http://www.redbooks.ibm.com/redbooks/pdfs/sg248060.pdf

This IBM Redbooks publication is for the Flex System Manager, but all tasks in section 8.3 are also suitable for the Common Agent installation for IBM Systems Director environments.

Install the Storage Subsystem Agent to the image repository server

Install the Storage Subsystem Agent from IBM Systems Director release management. Instructions can be found at the following information Center:

http://pic.dhe.ibm.com/infocenter/director/pubs/index.jsp?topic=%2Fcom.ibm.director.install.helps.doc%2Ffqm0_t_installing_agents.html

Notes:

- When installing, be sure to select the proper agent: CommonAgentSubagent_VMControl_commonrepository-2.4.1
- A comprehensive step-by-step guide for installing the Storage Subsystem Agent can be found at the following site:

  http://www.redbooks.ibm.com/redbooks/pdfs/sg248060.pdf

  This IBM Redbooks publication is written for Flex System Manager, but all tasks in section 8.3.3 are also suitable for Storage Subsystem Agent installation in IBM System Director environments.

Tip:

If you are having problems with requesting access to the Common Agent server, check troubleshooting steps in section 12.1, "Troubleshooting the installation of IBM Systems Director components" on page 494.

Create image repository in VMControl

To be able to create an image repository, IBM Systems Director VMControl needs connectivity to SAN storage and to Fibre Channel switches. A best practice is to use IBM SAN volume controller and or the v7000. Information on *how-to* create data sources for IBM Systems Director is found in Chapter 11, "Storage Management solutions and Storage Control" on page 451. Data sources for storage and Fibre Channel switches need to be in order before an image repository can be created.

Use the **smcli dumpstcfg** command to ensure that all is in place, as shown in Example 5-13.

Example 5-13 smcli dumpstcfg command

```
root@ISD632:~> smcli dumpstcfg
                                SAN Configuration
                        ------------------------------

Switches
--------
Name                           OID   Provider IP      Switch IP          WWPN

BRCD:32R1819-YK10UZ774015      18599 10.31.54.199     10.31.53.18        100000051E04FE7E
BRCD:32R1819-ZK11LV57B013      20681                  10.31.53.4
10.31.54.204                   7596                   10.31.54.204
10.31.54.203                   7585                   10.31.54.203
```

```
forum_even_B32                 18598 10.31.54.199    10.31.51.4         100000051E03845B
forum_odd_B32                  18600 10.31.54.199    10.31.51.3         100000051E03D8E5

Storage Subsystems
------------------
Name                                OID   Provider IP     Subsystem IP
Largest Slice(in GBs)   Fast Copy enabled?

SVC-2145-forumsvc-IBM               18339 -               { '10.31.51.9' }
247.25                  Yes(SSH)
Storwize V7000-2076-forumv7000-IBM  17798 -               { '10.31.51.20' }
1924.0                  Yes(SSH)
```

Example 5-13 on page 261 shows the connected Fibre Channel switches and storage subsystems. You can scroll down to see more results, such as the physical hosts that have connectivity to storage through Fibre Channel network zonings.

When zoning information and data sources are correct, you should see the storage pools that are available before creating the image repository, as shown in Example 5-14.

Example 5-14 Accessible storage subsystems and pools

```
Server Accessible Containers
----------------------------
NAME: STORAGE SUBSYSTEM/POOL

ISDKVM-node4:   SVC-2145-forumsvc-IBM/NOPEA_DS4700_10k_R5
  SVC-2145-forumsvc-IBM/FlashCopy target pool
  SVC-2145-forumsvc-IBM/NORMAALI_DS4700_10k_R6
  SVC-2145-forumsvc-IBM/NOPEIN_DS4700_15k_R5
  SVC-2145-forumsvc-IBM/HIDAS_DS4700_SATA_R5
  SVC-2145-forumsvc-IBM/MigrationPool_8192
  Storwize V7000-2076-forumv7000-IBM/Tier2
  Storwize V7000-2076-forumv7000-IBM/Hybrid_Pool
ISDKVM-Node3:   Storwize V7000-2076-forumv7000-IBM/Tier2
  Storwize V7000-2076-forumv7000-IBM/Hybrid_Pool
  SVC-2145-forumsvc-IBM/NOPEA_DS4700_10k_R5
  SVC-2145-forumsvc-IBM/FlashCopy target pool
  SVC-2145-forumsvc-IBM/NORMAALI_DS4700_10k_R6
  SVC-2145-forumsvc-IBM/NOPEIN_DS4700_15k_R5
  SVC-2145-forumsvc-IBM/HIDAS_DS4700_SATA_R5
  SVC-2145-forumsvc-IBM/MigrationPool_8192
ISDKVM-Node2:   Storwize V7000-2076-forumv7000-IBM/Tier2
  Storwize V7000-2076-forumv7000-IBM/Hybrid_Pool
  SVC-2145-forumsvc-IBM/NOPEA_DS4700_10k_R5
  SVC-2145-forumsvc-IBM/FlashCopy target pool
  SVC-2145-forumsvc-IBM/NORMAALI_DS4700_10k_R6
  SVC-2145-forumsvc-IBM/NOPEIN_DS4700_15k_R5
  SVC-2145-forumsvc-IBM/HIDAS_DS4700_SATA_R5
  SVC-2145-forumsvc-IBM/MigrationPool_8192
ImageRepository:   Storwize V7000-2076-forumv7000-IBM/Tier2
  Storwize V7000-2076-forumv7000-IBM/Hybrid_Pool
  SVC-2145-forumsvc-IBM/NOPEA_DS4700_10k_R5
  SVC-2145-forumsvc-IBM/FlashCopy target pool
  SVC-2145-forumsvc-IBM/NORMAALI_DS4700_10k_R6
  SVC-2145-forumsvc-IBM/NOPEIN_DS4700_15k_R5
  SVC-2145-forumsvc-IBM/HIDAS_DS4700_SATA_R5
```

```
SVC-2145-forumsvc-IBM/MigrationPool_8192
```

Example 5-14 on page 262 shows three KVM nodes and a host named as the image repository that has connectivity to two separate subsystems: IBM SAN Volume Controller and the IBM Storwize v7000.

> **Note:** You can add an image repository to VMControl using the IBM Systems Director server user interface. Instructions for doing this can be found at the following link:
>
> http://www.redbooks.ibm.com/redbooks/pdfs/sg248060.pdf
>
> This IBM Redbooks publication is written for the Flex System Manager, but all of the tasks in section 8.3.7 are suitable when working with the KVM Platform Agent installation tasks in IBM Systems Director environments.

Figure 5-5 shows the Image Repositories window in the VMControl user interface.

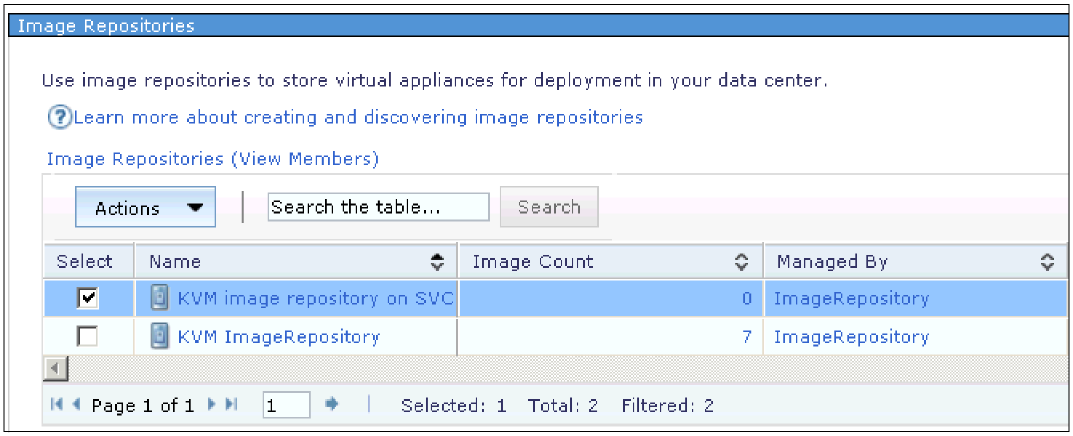

Image Repositories

Use image repositories to store virtual appliances for deployment in your data center.

Learn more about creating and discovering image repositories

Image Repositories (View Members)

Select	Name	Image Count	Managed By
☑	KVM image repository on SVC	0	ImageRepository
□	KVM ImageRepository	7	ImageRepository

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

Figure 5-5 Image repositories listed in IBM Systems Director VMControl

Create server system pools for KVM hosts

After zoning information and storage information have been checked and image repositories have been successfully created, you can create a server system pool.

Server system pools enable you to group similar hosts. With IBM Systems Director VMControl, you can create a system pool of selected hosts, add and remove hosts from the system pool, enter and exit maintenance mode for the hosts, and permanently delete a system pool.

If you have IBM Systems Director Network Control installed and licensed, you can also set up and manage network system pools (NSPs) to leverage automated network relocation and logical network provisioning.

Use the following steps to create a server system pool:

1. From the IBM Systems Director VMControl summary page, click the **System Pools** tab, ensure that the view is set to *Server system pools*, and click **Create**.
2. Follow the instructions in the Create Server System Pool wizard to create a server system pool that is composed of selected hosts and attached shared storage.

Note: A comprehensive step-by-step guide for creating server system pools in VMControl can be found at the following site:

http://www.redbooks.ibm.com/redbooks/pdfs/sg248060.pdf

This IBM Redbooks publication is written for the Flex System Manager, but all of the tasks in section 8.5 are suitable when working with server system pools in IBM Systems Director environments.

Server system pool optimization settings

When creating a server system pool, you can choose from two optimization choices:

- Manual optimization
- Automatic optimization

You can change the settings later or adjust the optimization interval. The difference between the two settings is that when automatic optimization is chosen, IBM Systems Director can relocate virtual servers automatically and choose the proper server automatically when new workloads are being deployed. See Figure 5-6.

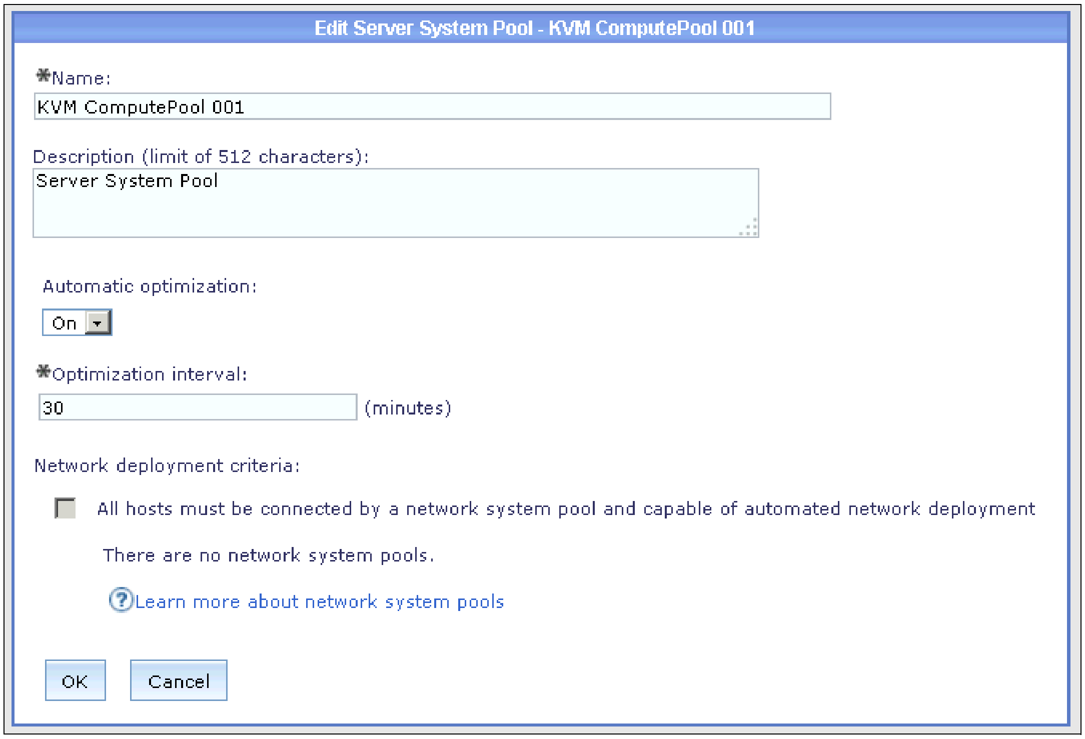

Edit Server System Pool - KVM ComputePool 001

*Name: KVM ComputePool 001

Description (limit of 512 characters): Server System Pool

Automatic optimization: On

*Optimization interval: 30 (minutes)

Network deployment criteria:

□ All hosts must be connected by a network system pool and capable of automated network deployment

There are no network system pools.

Figure 5-6 Optimization interval settings

Operating with VMControl in KVM in a Red Hat Enterprise Linux environment

After creating previous tasks, the KVM environment is ready for operations. The following tasks can be performed when you are operating in the KVM infrastructure:

- Create a new image in the image repository
- Create a virtual server
- Create an appliance from the virtual server

- Import an existing virtual appliance
- Deploy a virtual appliance to new workload
- Group virtual servers to run under one workload
- Relocate virtual servers and workloads

Create a new image in the image repository

Creating an image inside the image repository is a simple task. However, the principles for handling images are a little different from what users are use to doing with other virtualization environments.

VMControl needs an image repository host to manage images. Images are placed in individual SAN volumes, which are shown to a new virtual server as CD/DVD drives. So the difference between other hypervisor management tools is that there is no large image repository where to put all images. Every image is an individual disk, which can be mounted to a new virtual machine.

Follow these steps to create a new image in the image repository:

1. Transfer your image (for example, `rhel6.3.iso`) using SCP to the `/tmp` folder in the image repository host.
2. Create a new volume into the image repository for that ISO file.

   When creating virtual servers through the VMControl, the servers must have access to an ISO file. The best way to do this is to create a storage volume with an `_ISO` extension that the VMControl can mount and use to emulate as a CD/DVD device.
3. Create an ISO volume to the storage system to use this volume with the operating system (OS) installation on your first guest. In the Resource Explorer view, find and right-click the *ImageRepository server* (Figure 5-7 on page 266).

> **Note:** If you click the operating system, the Create Storage Volumes menu will not show. You need to click the physical server.

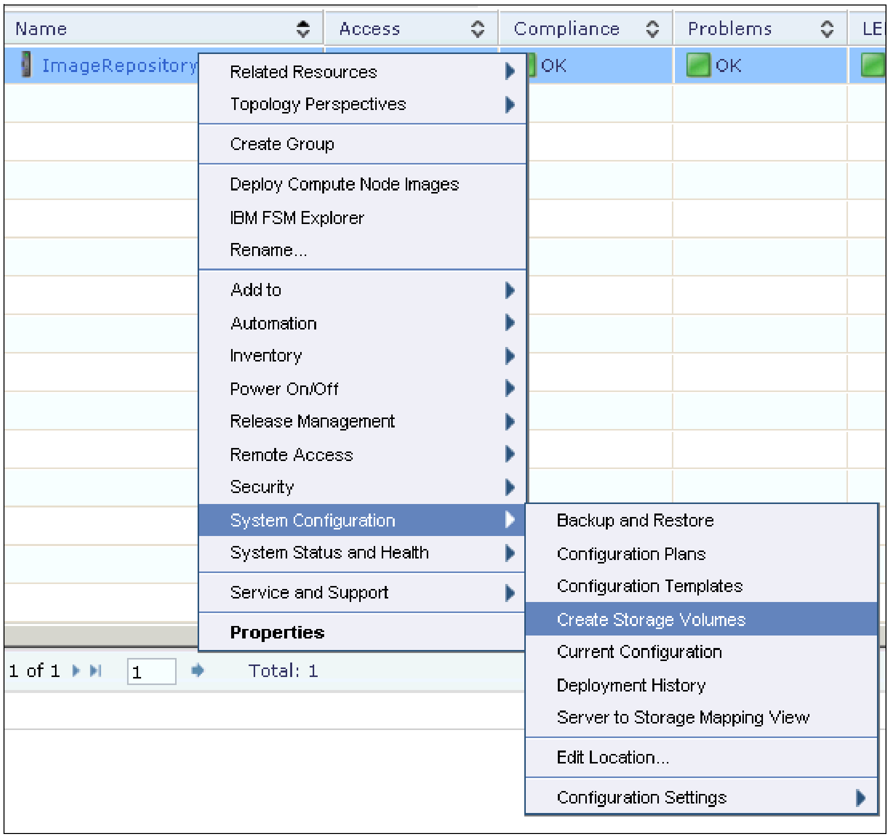

Name	Access	Compliance	Problems	LE
ImageRepository		OK	OK	

1 of 1 Total: 1

Figure 5-7 Choose image repository host and select Create Storage Volumes

When zoning information and data sources are properly installed, all available SAN storage systems are seen after clicking **System Configuration** and **Create Storage Volumes**. Then, select the storage pool where you want to place the new image (Figure 5-8).

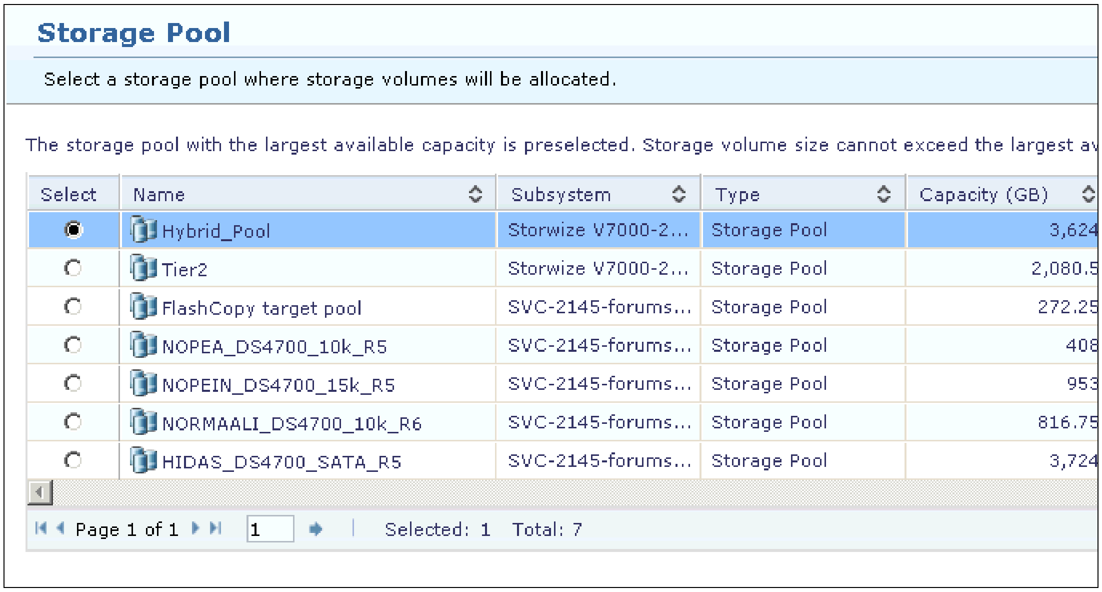

Storage Pool

Select a storage pool where storage volumes will be allocated.

The storage pool with the largest available capacity is preselected. Storage volume size cannot exceed the largest av

Select	Name	Subsystem	Type	Capacity (GB)
●	Hybrid_Pool	Storwize V7000-2...	Storage Pool	3,624
○	Tier2	Storwize V7000-2...	Storage Pool	2,080.5
○	FlashCopy target pool	SVC-2145-forums...	Storage Pool	272.25
○	NOPEA_DS4700_10k_R5	SVC-2145-forums...	Storage Pool	408
○	NOPEIN_DS4700_15k_R5	SVC-2145-forums...	Storage Pool	953
○	NORMAALI_DS4700_10k_R6	SVC-2145-forums...	Storage Pool	816.75
○	HIDAS_DS4700_SATA_R5	SVC-2145-forums...	Storage Pool	3,724

Page 1 of 1 Selected: 1 Total: 7

Figure 5-8 Select storage pool where the new volume for the ISO image will be located

After selecting **Storage Pool** and clicking **Next**, it is important to give access to the rest of the KVM hosts. You can see available hosts in the list. The host list becomes available from gathered zoning information.

Name the volume with the appropriate naming convention for the ISO image. For RHEL 6.3, for example, use `RHEL63_ISO`.

> **Note:** If the storage volume does not end with ISO, you are not able to use this as a CD/DVD emulated disk.

Size the disk correctly. An appropriate size for RHEL 6.3 x86 64 is 4 Gb. When volume is created, it is in a thick format. You cannot choose thin provisioned. If you want to use the storage system feature like compression, you have to create and name that volume in the storage system user interface and give access to all hosts through that user interface. When creating a disk for image, you use the *add existing disk* option.

After creating the new volume, go to the image repository server. You should see the new volume in the image repository disk list.

> **Note:** Do not mount that disk on to your image repository server and do not format it. Leave it as is.

4. Copy the `.ISO` file to a new ISO volume

   Go to the image repository server and list all disks. Select Create disk by listing disks and check which disk has been added as new into the inventory (Example 5-15).

Example 5-15 List disks to get device path

```
[root@ISDNode1 director]# fdisk -l |more

Disk /dev/sdc: 4294 MB, 4294967296 bytes
64 heads, 32 sectors/track, 4096 cylinders
Units = cylinders of 2048 * 512 = 1048576 bytes
Sector size (logical/physical): 512 bytes / 512 bytes
I/O size (minimum/optimal): 512 bytes / 512 bytes
Disk identifier: 0x023b007b

   Device Boot      Start         End      Blocks   Id  System
/dev/sdc1   *           1        3509     3593216   17  Hidden HPFS/NTFS
```

In this case, the new disk is mapped to `/dev/sdc1`. Next, copy the ISO image to that device. See Example 5-16.

Example 5-16 Copy image to device by using the dd command

```
[root@node11229 director]# dd if=/iso/RHEL6.3-Server-x86_64-DVD1.iso of=/dev/sdc
```

After the copying is done, the image is available for use as a CD/DVD device when creating the new virtual server.

Create a new virtual server

After creating CD/DVD devices, you can proceed to create a new virtual server. Creating a new virtual server is done from the VMControl user interface. Go to the Virtual Servers and Hosts tab. When all virtual servers and hosts are listed, select one of the KVM nodes.

Note: The new virtual server can be created only when the physical KVM host is selected, not the operating system.

1. Right-click and select **System configuration**. Then, select **Create Virtual Server** on the menu and type the name of the virtual server. See Figure 5-9, which involves naming the virtual server.

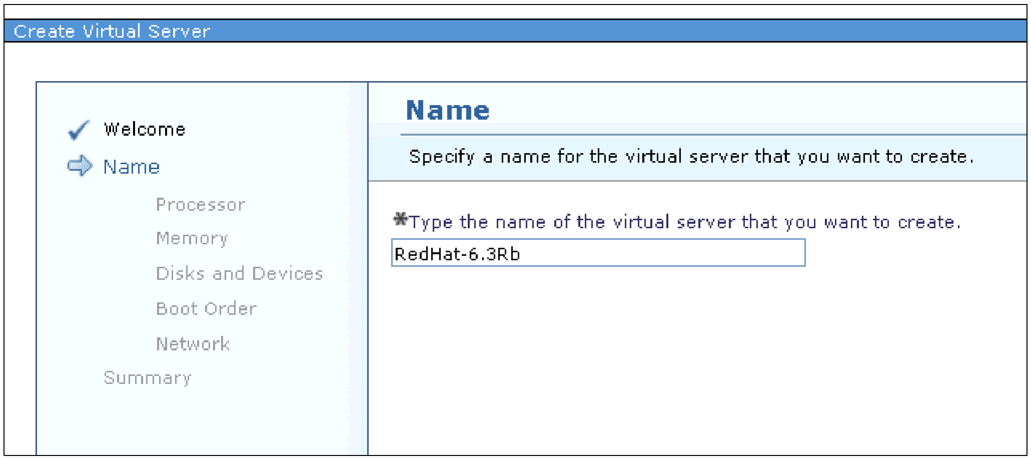

Create Virtual Server

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.

RedHat-6.3Rb

Figure 5-9 Name virtual server

2. Click **Next** and select the **Processor** amount. See Figure 5-10. Select the appropriate number of CPUs.

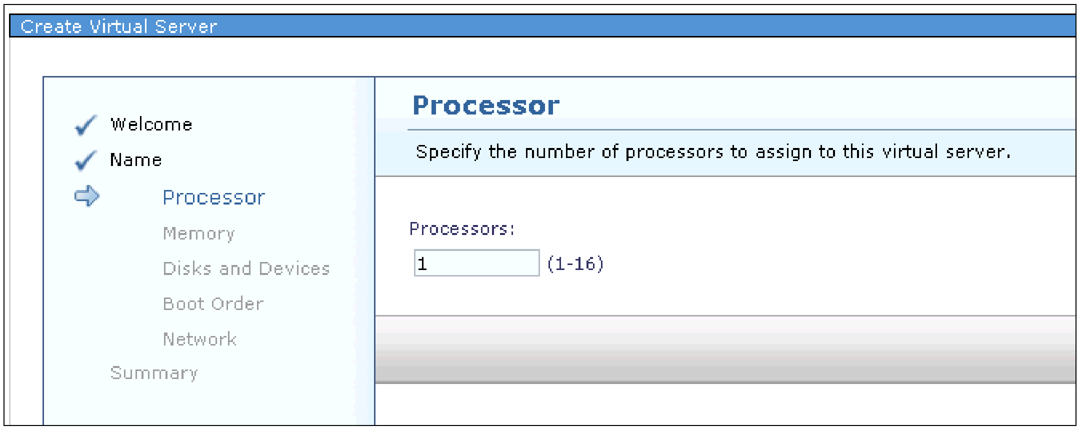

Create Virtual Server

Processor

Specify the number of processors to assign to this virtual server.

Processors: 1 (1-16)

Figure 5-10 Assign processor count

3. Click **Next** and assign the memory amount. See Figure 5-11. Specify memory amount.

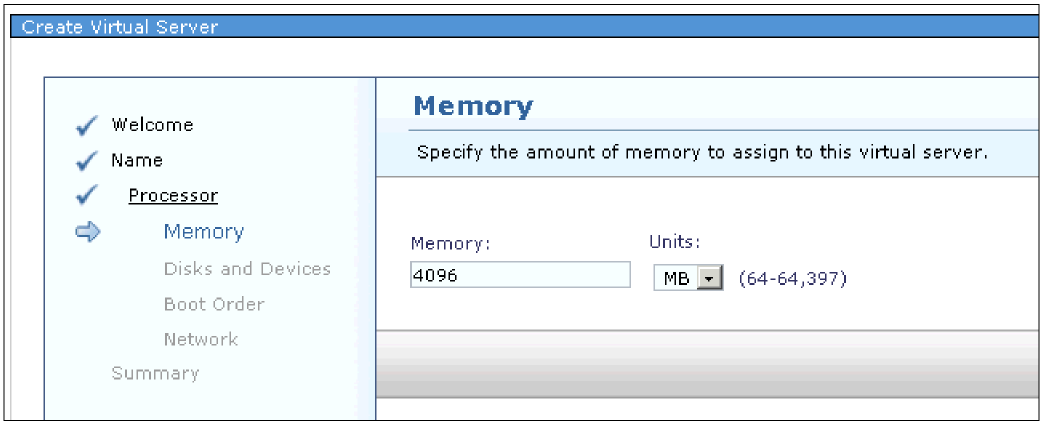

Memory

Specify the amount of memory to assign to this virtual server.

Memory: 4096

Units: MB (64-64,397)

Figure 5-11 Specify an appropriate amount of memory to the new virtual server

4. Click **Next** to see the Disks and Devices page. See Figure 5-12.

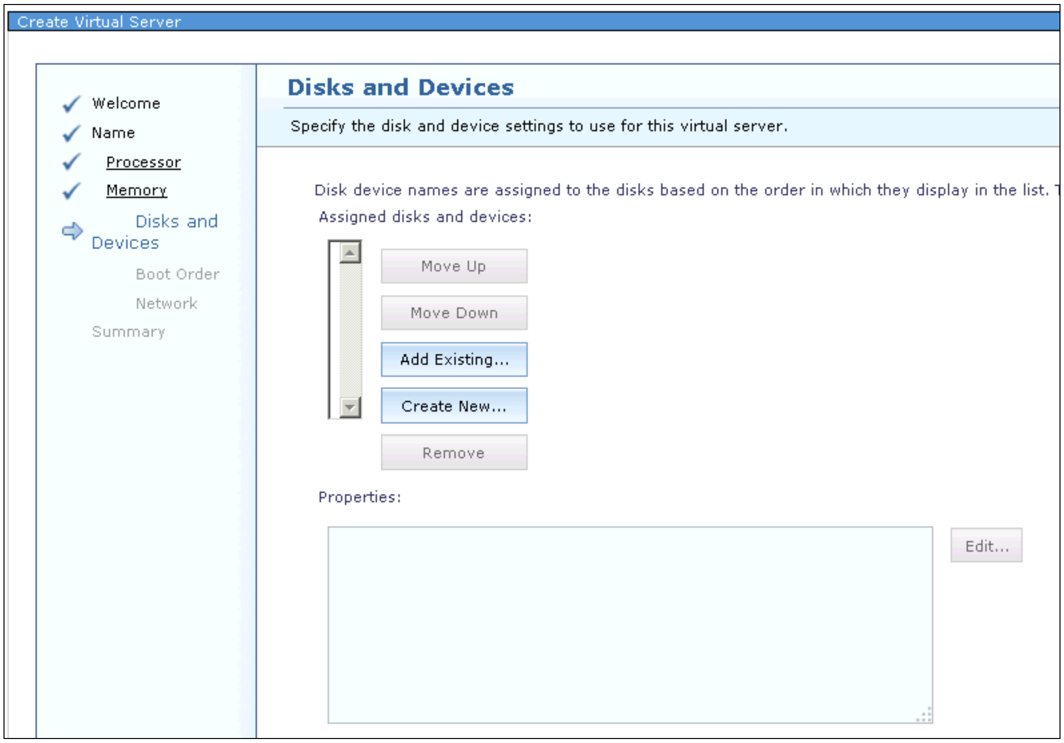

Disks and Devices

Specify the disk and device settings to use for this virtual server.

Disk device names are assigned to the disks based on the order in which they display in the list. T

Assigned disks and devices:

Properties:

Figure 5-12 Disks and Devices page

The virtual machine does not have any disks, so the list is empty. By clicking **Add Existing** disks, you can use existing disks that are available in the storage pools. By clicking **Create New**, you can create a new volume to the available storage pools.

When you are creating a new server, you must have the installation media or an existing PXE boot server. When using installation media, you need to add one of the existing ISO volumes. Otherwise, the server tries to locate the boot media from the PXE server. If the PXE server does not exist in your environment, the boot fails. See Figure 5-13.

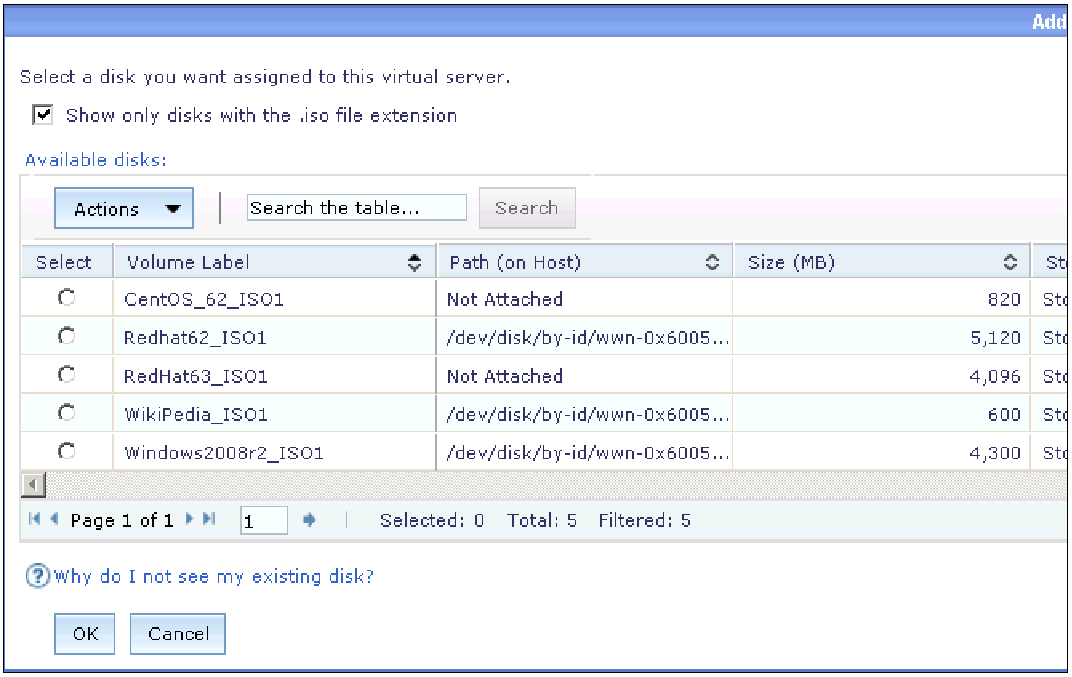

Select a disk you want assigned to this virtual server.

☑ Show only disks with the .iso file extension

Available disks:

Select	Volume Label	Path (on Host)	Size (MB)	St
○	CentOS_62_ISO1	Not Attached	820	Sto
○	Redhat62_ISO1	/dev/disk/by-id/wwn-0x6005...	5,120	Sto
○	RedHat63_ISO1	Not Attached	4,096	Sto
○	WikiPedia_ISO1	/dev/disk/by-id/wwn-0x6005...	600	Sto
○	Windows2008r2_ISO1	/dev/disk/by-id/wwn-0x6005...	4,300	Sto

Page 1 of 1 1 Selected: 0 Total: 5 Filtered: 5

Why do I not see my existing disk?

Figure 5-13 The page where you can select the right ISO volume to the virtual servers disk list

5. Select the appropriate ISO volume for the new operating system installation and press **OK**.

   For creating the new primary volume for the virtual server, you have two options:

   – Create a volume with VMControl
   – Create a volume in the storage system and use it when creating a virtual server

   When creating a volume with VMControl, it is automatically created to a storage system and mapped to selected hosts. This is the easiest way to assign new disks to the virtual server.

   When creating a disk in the VMControl user interface, the disk type will be *thick*, no matter what features your storage system might be providing.

   When creating a volume with the storage system, for example with the v7000 or SVC, you can use the storage system features like thin provisioning and compression. When the volume is created, it needs to be assigned to all hosts that are on the same server system pool.

   Creating disks for virtual servers in the storage system might be good when creating virtual servers that require a lot of disk space; or, when you want to use advanced features that are provided by that specific storage system.

> **Note:** When capturing the created virtual server, which uses a thin provisioned disk, it will be converted as a generic thick drive.

When creating a new virtual server, which will be used later as the virtual server appliance, the best practice is to use the VMControl user interface. See Figure 5-14.

All available storage pools are shown in the list in Figure 5-14 from all storage systems that are connected properly into the IBM Systems Director system.

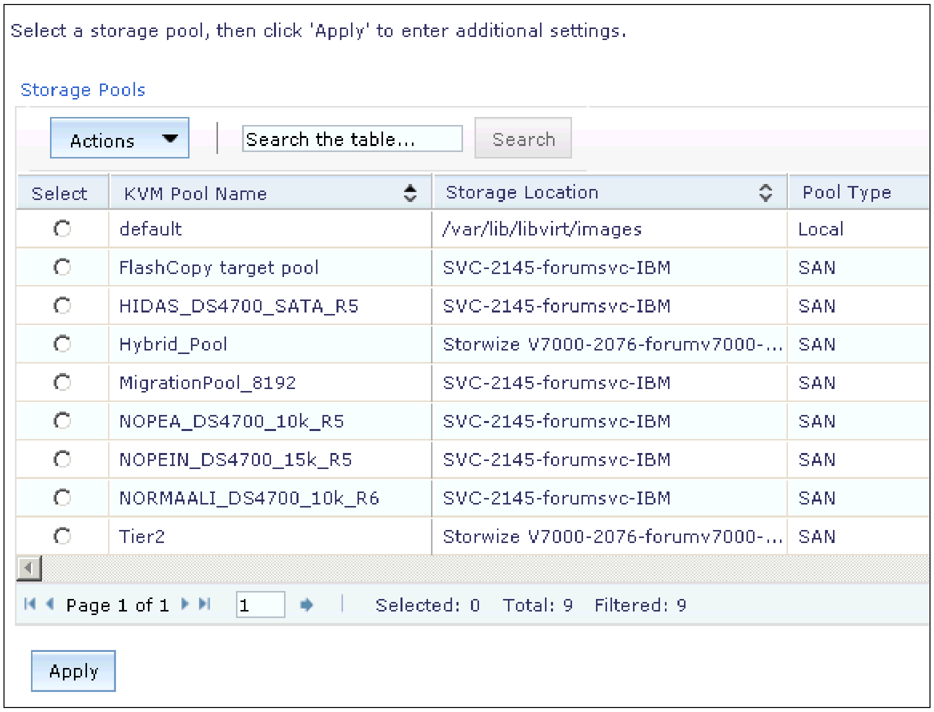

Select a storage pool, then click 'Apply' to enter additional settings.

Storage Pools

Select	KVM Pool Name	Storage Location	Pool Type
○	default	/var/lib/libvirt/images	Local
○	FlashCopy target pool	SVC-2145-forumsvc-IBM	SAN
○	HIDAS_DS4700_SATA_R5	SVC-2145-forumsvc-IBM	SAN
○	Hybrid_Pool	Storwize V7000-2076-forumv7000-...	SAN
○	MigrationPool_8192	SVC-2145-forumsvc-IBM	SAN
○	NOPEA_DS4700_10k_R5	SVC-2145-forumsvc-IBM	SAN
○	NOPEIN_DS4700_15k_R5	SVC-2145-forumsvc-IBM	SAN
○	NORMAALI_DS4700_10k_R6	SVC-2145-forumsvc-IBM	SAN
○	Tier2	Storwize V7000-2076-forumv7000-...	SAN

Figure 5-14 Select a storage pool for a virtual server new disk

6. Select a storage pool for the new virtual machine and click **Apply**. For our example, select **Hybrid_Pool**. See Figure 5-15 on page 272.

> **Note:** When creating the virtual appliance, remember that all new deployments of this appliance will be deployed to the same storage pool, which you select at this point.

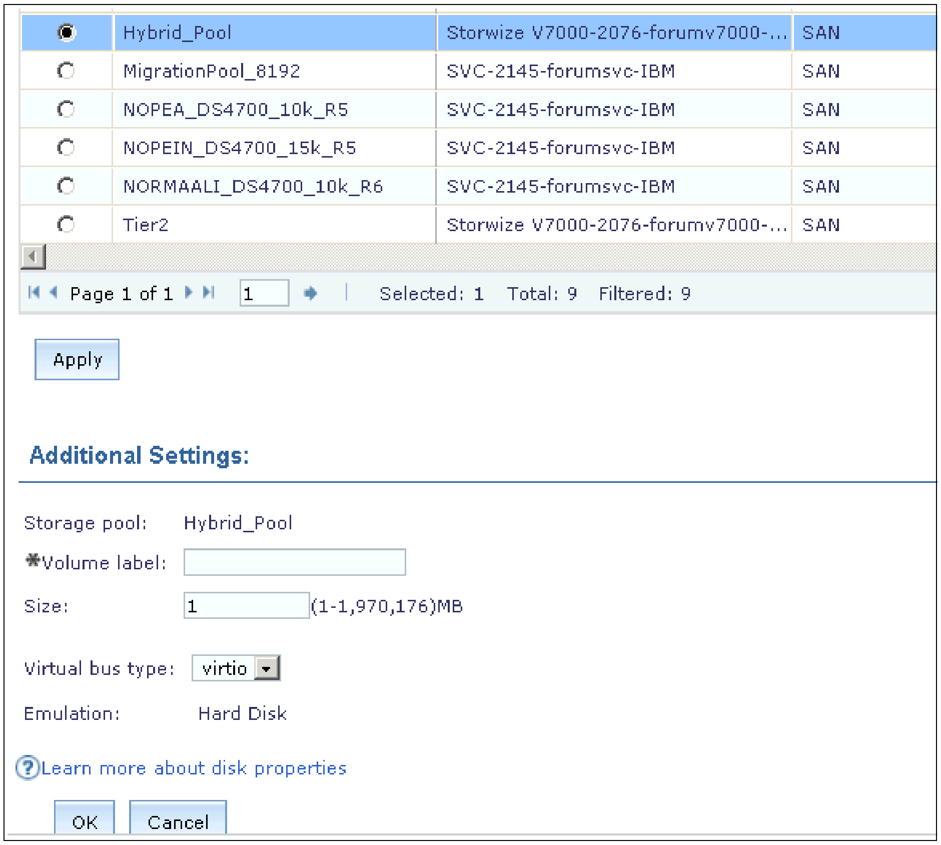

●	Hybrid_Pool	Storwize V7000-2076-forumv7000-...	SAN
○	MigrationPool_8192	SVC-2145-forumsvc-IBM	SAN
○	NOPEA_DS4700_10k_R5	SVC-2145-forumsvc-IBM	SAN
○	NOPEIN_DS4700_15k_R5	SVC-2145-forumsvc-IBM	SAN
○	NORMAALI_DS4700_10k_R6	SVC-2145-forumsvc-IBM	SAN
○	Tier2	Storwize V7000-2076-forumv7000-...	SAN

Page 1 of 1 | 1 | Selected: 1 Total: 9 Filtered: 9

Additional Settings:

Storage pool: Hybrid_Pool

*Volume label:

Size: 1 (1-1,970,176)MB

Virtual bus type: virtio

Emulation: Hard Disk

Figure 5-15 Additional settings for new volume

When the storage pool has been selected, name the volume. Proper naming convention is to not use spaces. Notice that the size is in MB, so if you want to have a disk size around 10 GB, put 10,000 in the Size text box.

After sizing the disk, select the virtual bus type. Select either VirtIO or integrated development environment (IDE) virtual bus types for the disk communication.

VMControl can use KVM on Red Hat Enterprise Linux para-virtualized VirtIO drivers. VirtIO is a common standard framework allowing guests to be more easily moved between hypervisor platforms, which are supporting VirtIO.

VirtIO allows for increased I/O performance for both network and block devices, when compared to regular emulated devices.

> **Tip:** If using the VirtIO bus with the new virtual server and you face problems, you can change it to use the IDE bus and problems should not occur.

7. After selecting parameters, click **OK** and you are returned to the disk page. Click **Next**. This is where you can define the boot order (Figure 5-16 on page 273).

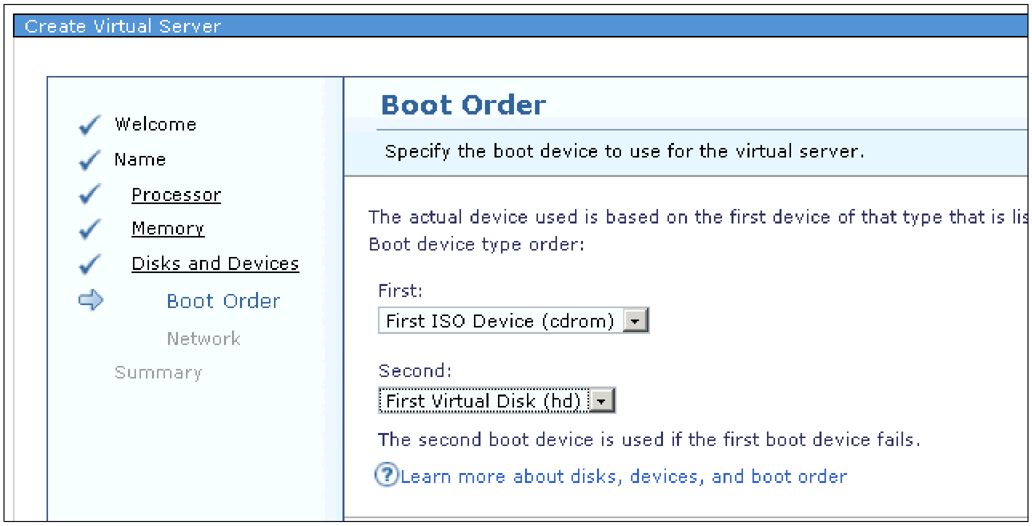

Boot Order

Specify the boot device to use for the virtual server.

The actual device used is based on the first device of that type that is lis

Boot device type order:

First:
First ISO Device (cdrom)

Second:
First Virtual Disk (hd)

The second boot device is used if the first boot device fails.

Learn more about disks, devices, and boot order

Figure 5-16 Specifying the boot order for a new virtual server

When creating a new virtual server without an operating system, the first boot device should be an ISO device that was created and selected in previous tasks. The second boot device is the primary disk for the virtual server that was previously selected and created.

8. Click **Next** and see the Network page (Figure 5-17).

   When creating or editing a KVM virtual server, specify the networks that you want to assign to the virtual server. However, only networks that are existing on that specific host are displayed in the list of networks. Be sure that the same networks in all hosts are available.

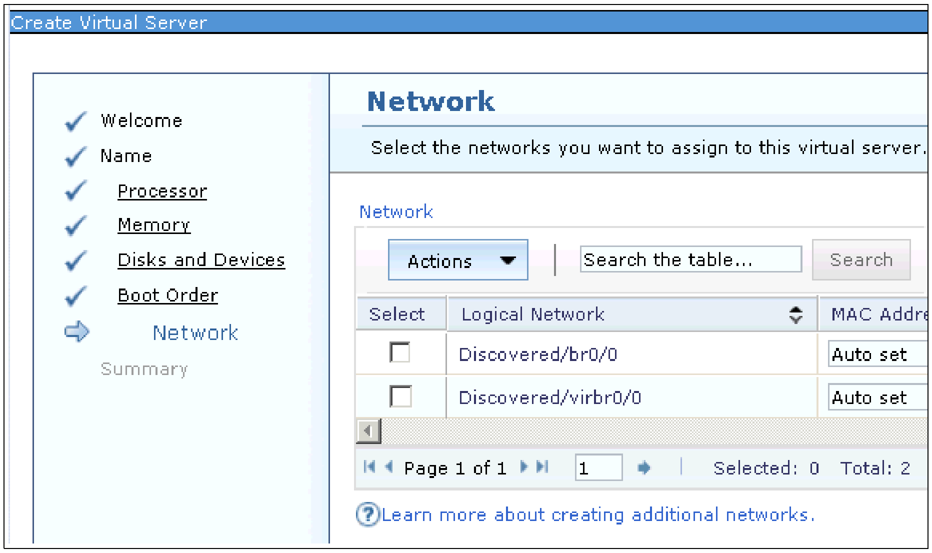

Network

Select the networks you want to assign to this virtual server.

Network

Select	Logical Network	MAC Addre
□	Discovered/br0/0	Auto set
□	Discovered/virbr0/0	Auto set

Page 1 of 1 | 1 | Selected: 0 Total: 2

Learn more about creating additional networks.

Figure 5-17 Select the proper network for your new virtual server

You can create additional networks for use with KVM by using IBM Systems Director Network Control or create them manually as shown in “Preparation” on page 255.

If you do not have Network Control installed, you can collect the inventory against the server and the server system pools to identify any predefined or new bridges that might exist.

As shown in the example in Figure 5-17 on page 273, these predefined bridges are listed because they were built in the preparation phase. If there is a need for more networks, they must be created either manually or through the IBM Systems Director.

After selecting the appropriate network connection for your new virtual server, click **Next**. The last page is a basic summary page. You can check to verify that all parameters are correct for the new virtual server.

9. After checking parameters, click **Finish**.

   You can run the job now or schedule it to run later. After the *Create Virtual Server* job has been executed, the virtual server is created and ready for the operating system installation.

   By checking the task log that is shown in Figure 5-18, you can see what steps are taken when creating the new virtual server and whether the task completed successfully or not. If there are errors during deployment, they will be shown in this log and are very helpful when solving potential problems.

```
April 19, 2013 6:48:52 PM EEST-Level:1-MEID:0--MSG: Job "Create Virtual Server - April 19, 2013 6:48:47 PM
EEST" activated.
April 19, 2013 6:48:53 PM EEST-Level:200-MEID:0--MSG: Subtask "Create Virtual Server" activated.
April 19, 2013 6:48:53 PM EEST-Level:200-MEID:0--MSG: Starting clients
April 19, 2013 6:48:53 PM EEST-Level:100-MEID:0--MSG: Clients started for task "Create Virtual Server"
April 19, 2013 6:48:53 PM EEST-Level:200-MEID:0--MSG: Subtask activation status changed to "Active".
April 19, 2013 6:48:53 PM EEST-Level:1-MEID:0--MSG: Job activation status changed to "Active".
April 19, 2013 6:48:53 PM EEST-Level:200-MEID:0--MSG: Subtask activation status changed to "Active".
April 19, 2013 6:48:54 PM EEST-Level:200-MEID:0--MSG: Create virtual server 'RedHat-6.3Rb' processing
started.
April 19, 2013 6:49:51 PM EEST-Level:200-MEID:0--MSG: Create virtual server process flow has been completed
successfully.
April 19, 2013 6:50:01 PM EEST-Level:200-MEID:0--MSG: A new virtual server has been created successfully:
RedHat-6.3Rb
April 19, 2013 6:50:01 PM EEST-Level:100-MEID:7344--MSG: ISDKVM-Node2 client job status changed to
"Complete".
April 19, 2013 6:50:01 PM EEST-Level:200-MEID:0--MSG: Subtask activation status changed to "Complete".
April 19, 2013 6:50:01 PM EEST-Level:1-MEID:0--MSG: Job activation status changed to "Complete".
```

Figure 5-18 Log for executing “Create Virtual Server” job

After executing the *Create Virtual Server* job, the new server can be seen in the list of virtual servers, similar to what is shown in Figure 5-19.

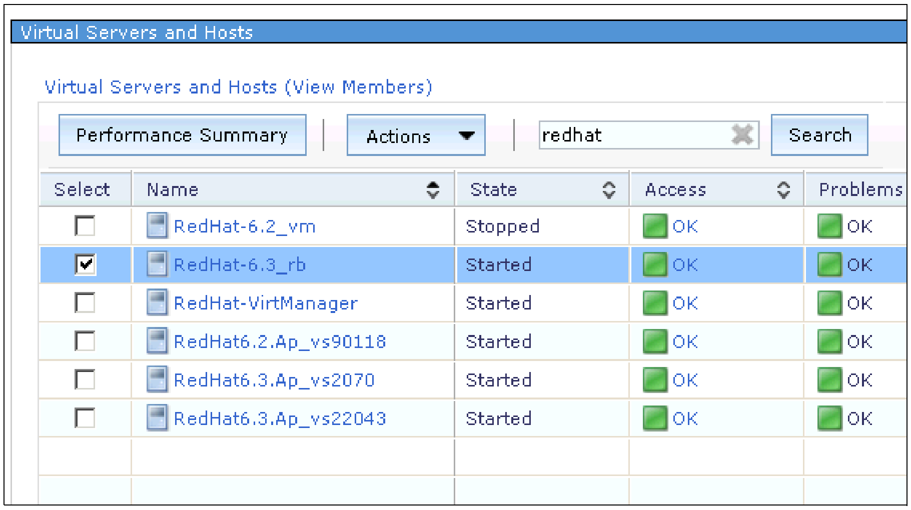
Virtual Servers and Hosts
Virtual Servers and Hosts (View Members)
Performance Summary
Actions
redhat
Search

Select	Name	State	Access	Problems
☐	RedHat-6.2_vm	Stopped	🟩 OK	🟩 OK
☑	RedHat-6.3_rb	Started	🟩 OK	🟩 OK
☐	RedHat-VirtManager	Started	🟩 OK	🟩 OK
☐	RedHat6.2.Ap_vs90118	Started	🟩 OK	🟩 OK
☐	RedHat6.3.Ap_vs2070	Started	🟩 OK	🟩 OK
☐	RedHat6.3.Ap_vs22043	Started	🟩 OK	🟩 OK

Figure 5-19 Virtual Servers and Hosts view: New virtual server can be found if search field is used

Installing the operating system

After creating the new server, start the server from the Systems Director user interface and install the operating system. You can use the image repository server virtual machine manager console for first time installations. In the image repository server, open connect virtual machine manager and connect that to KVM host. Then, open the host and see the virtual machines running on this host.

> **Note:** After installing the operating system, you can use the remote control virtual machines directly from IBM Systems Director with Virtual Network Computing (VNC). You must have the VNC viewer installed on your machine.

Creating an appliance from a new virtual server

Virtual appliance is the captured virtual server and it can be used as a template for new virtual servers. All virtual appliances that can be deployed are shown in the VMControl virtual appliances tab. To capture virtual servers, you need the VMControl Enterprise Edition.

To use the new virtual server as an appliance, it needs to be captured. To capture an existing virtual server for appliance use, it requires executing a specific activation engine tool, which is located in the IBM Systems Director server.

Activation engine files can be found from IBM Systems Director running on Linux or AIX from the following path:

```
/opt/ibm/director/proddata/activation-engine/
```

Activation engine files can be found from IBM Systems Director running on Windows from the following path:

```
<install root>\ibm\director\proddata\activation-engine\
```

Steps to run an activation engine in Linux virtual servers

Perform the following steps to run the activation engine in Linux virtual servers:

1. Copy `vmc.vsae.tar` from the IBM Systems Director server using SCP or WINSCP.
2. Copy `vmc.vsae.tar` and move it to a `/temp` directory on your new virtual server.
3. Take the SSH connection to your new virtual server and execute '`tar xvf /temp/vmc.vsae.tar`'
4. Execute '`/temp/linux-install.sh`'

When asked, if your new guest is running on a KVM host, answer `yes` and allow the installation to complete. Move back to the IBM Systems Director and now run an inventory of the new virtual server. See Figure 5-20 on page 276.

Once the inventory is completed, move back to the SSH session of the virtual server and execute 'AE.sh --reset'. Now the virtual server should do a graceful shutdown. See the SSH session in Figure 5-20.

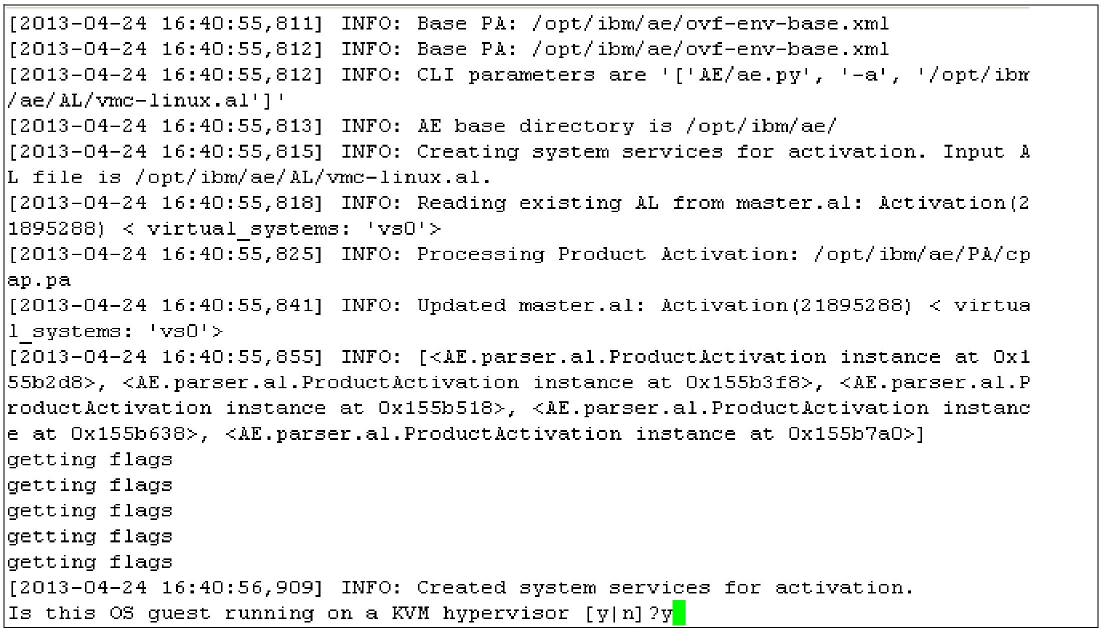

```
[2013-04-24 16:40:55,811] INFO: Base PA: /opt/ibm/ae/ovf-env-base.xml
[2013-04-24 16:40:55,812] INFO: Base PA: /opt/ibm/ae/ovf-env-base.xml
[2013-04-24 16:40:55,812] INFO: CLI parameters are '['AE/ae.py', '-a', '/opt/ibm
/ae/AL/vmc-linux.al']'
[2013-04-24 16:40:55,813] INFO: AE base directory is /opt/ibm/ae/
[2013-04-24 16:40:55,815] INFO: Creating system services for activation. Input A
L file is /opt/ibm/ae/AL/vmc-linux.al.
[2013-04-24 16:40:55,818] INFO: Reading existing AL from master.al: Activation(2
1895288) < virtual_systems: 'vs0'>
[2013-04-24 16:40:55,825] INFO: Processing Product Activation: /opt/ibm/ae/PA/cp
ap.pa
[2013-04-24 16:40:55,841] INFO: Updated master.al: Activation(21895288) < virtua
l_systems: 'vs0'>
[2013-04-24 16:40:55,855] INFO: [<AE.parser.al.ProductActivation instance at 0x1
55b2d8>, <AE.parser.al.ProductActivation instance at 0x155b3f8>, <AE.parser.al.P
roductActivation instance at 0x155b518>, <AE.parser.al.ProductActivation instanc
e at 0x155b638>, <AE.parser.al.ProductActivation instance at 0x155b7a0>]
getting flags
getting flags
getting flags
getting flags
getting flags
[2013-04-24 16:40:56,909] INFO: Created system services for activation.
Is this OS guest running on a KVM hypervisor [y|n]?y
```

Figure 5-20 Executing script for preparing image for capturing

Running the activation engine in Windows virtual servers

Perform the following steps to run the activation engine in Windows virtual servers:

1. Copy WinVSAE.zip from the IBM Systems Director server using the copy function.
2. Copy WinVSAE.zip and move it to a /temp directory on your new virtual server.
3. Take the remote desktop connection to your new virtual server and unpack WinVSAE.zip into the c:\temp directory.
4. Execute

   ```
   c:\temp\install-win-vsae.vbs command in command prompt window.
   ```

   This command installs required components for capturing. When that script has completed, run the **AE.bat --reset** command.

> **Note:** Comprehensive material for uninstalling and modifying parameters can be found at the following site:
>
> http://pic.dhe.ibm.com/infocenter/director/pubs/index.jsp?topic=%2Fcom.ibm.director.vim.helps.doc%2Ffsd0_vim_r_kvm_capture_reqs.html
>
> After running the activation engine, the virtual server must be powered off and can be captured from the IBM Systems Director VMControl user interface.

Capturing the virtual server

Select the virtual server. It should be already powered off after executing the reset script. Select *System Configuration* and *Capture*. See Figure 5-21.

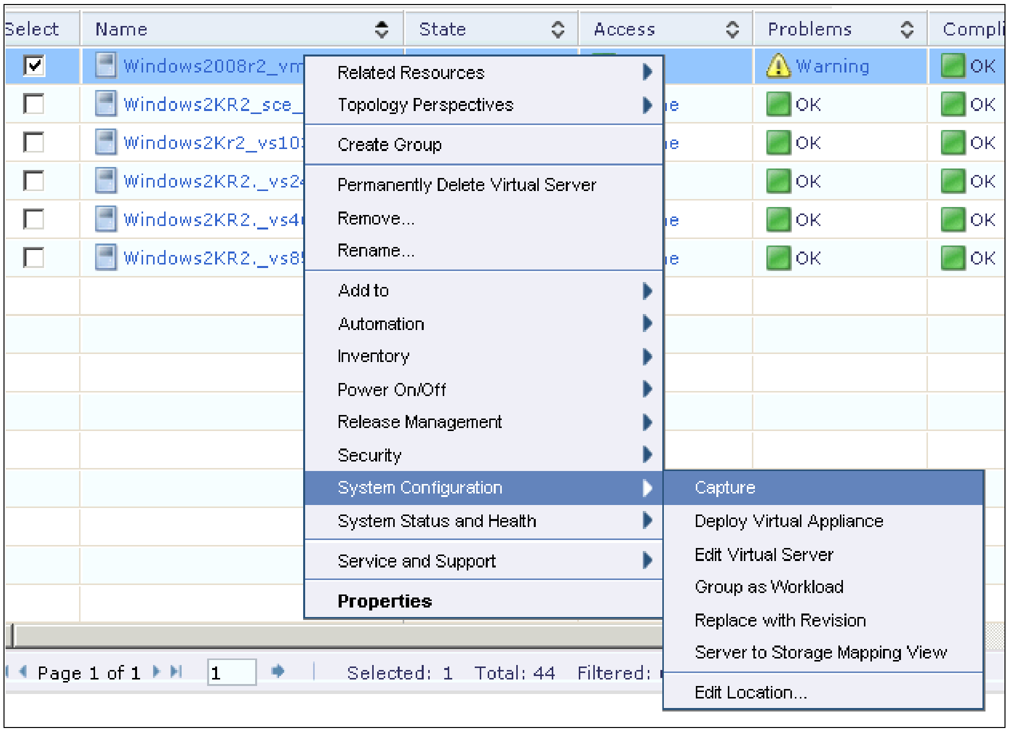

Select	Name	State	Access	Problems	Compli
☑	Windows2008r2_vm			⚠ Warning	🟩 OK
□	Windows2KR2_sce_		e	🟩 OK	🟩 OK
□	Windows2Kr2_vs10		e	🟩 OK	🟩 OK
□	Windows2KR2._vs2			🟩 OK	🟩 OK
□	Windows2KR2._vs4		e	🟩 OK	🟩 OK
□	Windows2KR2._vs8		e	🟩 OK	🟩 OK

Page 1 of 1 1 Selected: 1 Total: 44 Filtered:

Figure 5-21 Capturing a new appliance from an existing virtual server: select System Configuration and Capture

Note: A comprehensive step-by-step guide with screen captures that describes how to capture a virtual server and do operational tasks, such as relocating, is at the following site:

http://www.redbooks.ibm.com/redbooks/pdfs/sg248060.pdf

This IBM Redbooks publication is for the Flex System Manager, but all tasks are suitable when working with VMControl in IBM Systems Director environments.

Capture-related restrictions: The following restrictions apply to the capturing process:

- Only servers that are in the powered-off state can be captured.
- If you are using full virtualization, that is the IDE virtual bus type, select at most three disks to be captured. The fourth IDE disk slot is reserved for image activation with an ISO environment file during deployment.
- You cannot perform capture and relocation tasks concurrently for the same virtual server. Wait until one of these tasks is complete before beginning another.
- Virtual server information that is captured is limited to processors, memory, disks, network interfaces, and product activation properties.

  Other information such as consoles, input devices, and video devices are not captured to the appliance. Instead, default values for these items are used when the appliance is deployed and these defaults match how a virtual server is created through VMControl.

ISO disks are not captured from the virtual server.

Importing a virtual appliance

Virtual appliance packages are in the Open Virtualization Format (OVF), which is a platform independent and open packaging and distribution format for virtual servers. You can import the virtual appliance package from the Internet or from a system in your network. After importing the virtual appliance package, you can quickly deploy it into your environment.

Figure 5-22 on page 279 shows how a virtual appliance package can be imported to create a new virtual appliance:

- A user selects a virtual appliance package to import: *Virtual appliance package A*. The virtual appliance package can reside either on your IBM Systems Director or on a web server that is accessible to your IBM Systems Director.
- Importing *Virtual appliance package A* results in the virtual appliance, *Virtual appliance A*, which is stored by VMControl.
- *Virtual appliance A* contains the metadata that describes the virtual server, and a reference to *Image A*. Image A contains a fully configured and tested operating system and software applications for the virtual server. And, is stored in a VMControl image repository.
- After *Virtual appliance A* is imported, you can use VMControl to deploy the *Virtual appliance A* one or more times into your environment. See Figure 5-22 on page 279.

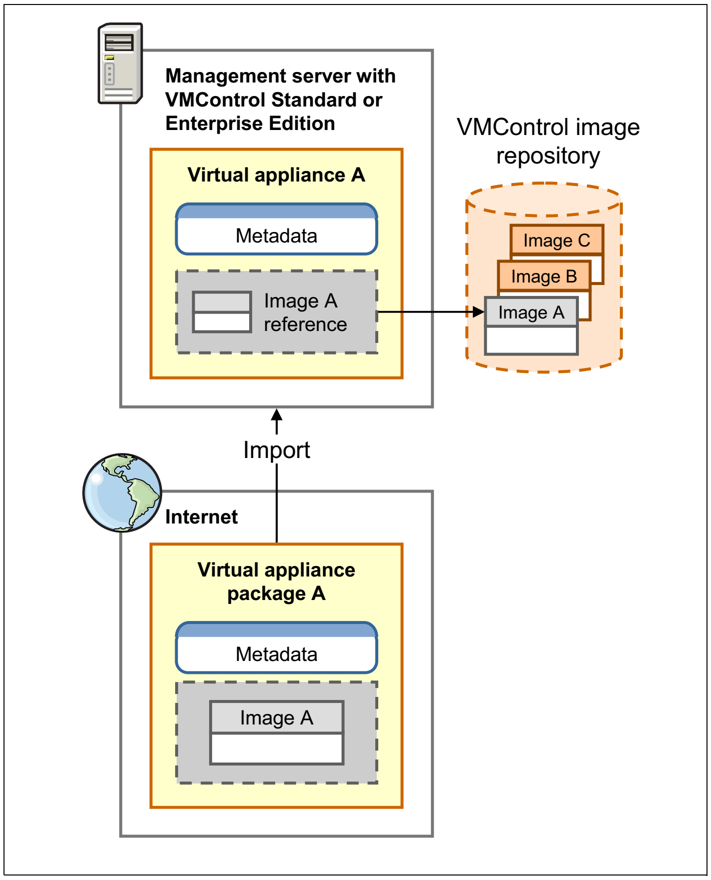
Management server with VMControl Standard or Enterprise Edition
Virtual appliance A
Metadata
Image A reference
VMControl image repository
Image C
Image B
Image A
Import
Internet
Virtual appliance package A
Metadata
Image A

Figure 5-22 Importing the appliance into IBM Systems Director

Note: A comprehensive step-by-step guide with screen captures, how to import a virtual appliance, and relocate a virtual server is found at the following site:

http://www.redbooks.ibm.com/redbooks/pdfs/sg248060.pdf

This IBM Redbooks publication is for the Flex System Manager, but all tasks are suitable when working with VMControl in IBM Systems Director environments.

Exporting a virtual appliance

This topic shares the best practice on how to export the created virtual appliance in the IBM Systems Director VMControl. The virtual appliance export runs on KVM and on Red Hat Enterprise Linux. Perform the following steps to export:

1. Map the virtual appliance volume from the storage GUI to the host, which will be used as the appliance export host. See Example 5-17.

Example 5-17 You can see appliance volume name while executing capture

```
April 26, 2013 6:22:38 PM EEST-Level:150-MEID:0--MSG: DNZCIR346I New disk group:
DG_04.26.2013-18:21:43:396
April 26, 2013 6:22:38 PM EEST-Level:150-MEID:0--MSG: DNZCIR353I The virtual appliance is
using disk group DG_04.26.2013-18:21:43:396 with the following SAN volumes:
[va_CentOS62-Export_1].
```

```
April 26, 2013 6:22:38 PM EEST-Level:150-MEID:0--MSG: DNZCIR354I The virtual server is
using disk group DG_04.21.2013-23:10:42:048 with the following SAN volumes:
[60050768028100B4CC00000000000214+9+00000200A0402D33+0].
```

2. Use for example, the Storwize v7000 user interface to map the appliance volume to a host.

   The selected host can be any host. But it is a good idea to select an existing image repository host with Red Hat Enterprise Linux, so the commands shown in the example will apply. See Figure 5-23.

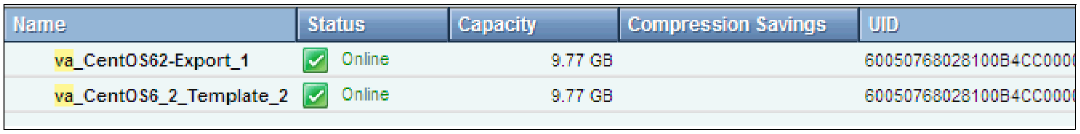

Name	Status	Capacity	Compression Savings	UID
va_CentOS62-Export_1	✓ Online	9.77 GB		60050768028100B4CC000
va_CentOS6_2_Template_2	✓ Online	9.77 GB		60050768028100B4CC000

Figure 5-23 Storwise v7000 user interface, where appliance va_CentOS62-Export1 volume is seen

   Ensure that you have enough disk space to capture the appliance.

3. Copy the raw drive to your image repository host directory.

   Use the **dd** command to capture the raw drive, for example, assuming that the virtual appliance drive becomes `/dev/sdk`. Use the command **dd if=/dev/sdk of=/root/export/<some_name>**.

   Example 5-18 shows an example that you might use for a disk, *centos*.

Example 5-18 Using the dd command to copy the disk to directory

```
[root@ImageRep ~]# dd if=/dev/sdk of=/root/export/centos
10485760+0 records in
10485760+0 records out
5368709120 bytes (5.4 GB) copied, 63.2503 s, 84.9 MB/s
[root@ImageRep ~]#
```

4. Use the **smcli lsva -o** command from the IBM Systems Directory. `Smcli` identifies the object identifier (OID) of the virtual appliance. See Example 5-19.

Example 5-19 smcli lsva -o command

```
USERID@PureFlex:~> smcli lsva -o
CentOS6.2_Template, 27599 (0x6bcf)
CentOS62-Export, 28032 (0x6d80)
USERID@PureFlex:~>
```

   In this case, the output from **smcli lsva -o**, is `CentOS6.2_Template, 27599 (0x6bcf) and CentOS62-Export, 28032 (0x6d80).`

5. After you know the ID for this appliance, use **smcli nimGetOVF <vaOID>** to get the OVF file.

   Execute **smcli nimGetOVF 28032 > CentOS62-Export.ovf**. You will find the `.ovf` file from that directory where you currently are.

6. Next, you need to edit the OVF file. See Example 5-20 on page 281.

   Find the line from the `.ovf` file:

```
<ovf:File
ovf:href="sanvolume://00000200A0402D33/@/60050768028100B4CC00000000000222"
ovf:id="file1" ovf:size="10485760000"/>
```

Example 5-20 Part of .ovf file: What you need to edit

```
http://schemas.dmtf.org/ovf/envelope/1 dsp8023_1.0.0.xsd
http://www.ibm.com/xmlns/ovf/extension/vim/2/rasd ibm-vim2-rasd_2.1.0.xsd"
xml:lang="en-US">
  <ovf:References>
    <ovf:File
ovf:href="sanvolume://00000200A0402D33/@/60050768028100B4CC00000000000222"
ovf:id="file1" ovf:size="10485760000"/>
  </ovf:References>
  <ovf:DiskSection>
    <ovf:Info>List of Virtual Disks used by this package</ovf:Info>
    <ovf:Disk ovf:capacity="10485760000" ovf:capacityAllocationUnits="byte"
ovf:diskId="disk1" ovf:fileRef="file1"
ovf:format="http://www.ibm.com/xmlns/ovf/diskformat/qemu.raw"
ovf:populatedSize="10485760000"/>
  </ovf:DiskSection>
```

Edit this `.ovf` file to reflect the name of the raw drive `<CentOS62>` made before to the image repository host:

```
<ovf:File ovf:href="Centos" ovf:id="file1" ovf:size="10485760000"/>
```

7. After editing the file, you have to unmount the appliance disk from the image repository if you intend to deploy more workloads from this exported appliance.
8. Copy the raw disk file and `.ovf` file to the host, which will be used as the repository for exported appliances. Keep the `.ovf` file and disk file together.

Now you should have the raw drive file and the `.ovf` file in edited format. You can import them using the IBM Systems Director UI. When importing, you have to know the full path and file name of the `.ovf` file.

Note: A comprehensive step-by-step guide with screen captures on how to import the virtual appliance and relocate the virtual server is found at the following website:

http://www.redbooks.ibm.com/redbooks/pdfs/sg248060.pdf

The IBM Redbooks publication is for the Flex System Manager, but all tasks are suitable when working with VMControl in IBM Systems Director environments.

Best Practices:

- VMControl provides a large selection of features that can be used with KVM hypervisor.
- Get to know restrictions and limitations, so you can design the purpose of that environment correctly.
- Use SAN (Fibre Channel storage, switches) to create your KVM environment.
- Plan your network carefully, draw a picture where you have all required network interfaces, VLANs, and networks.
- Build your network ready before deployment; reserve all required NICs to your servers.
- Plan your storage, calculate how much disk space you need, and what kind of storage you need.
- Create zones for servers and storage.
- Build data sources ready to IBM Systems Director for storage and Fibre Channel switches.
- Check that all is in order before installing Red Hat Enterprise Linux.
- Use Red Hat Enterprise Linux version 6.3.
- Follow instructions given in this Redbooks publication.
- Use IBM SmartCloud Entry to provide a self-service portal to your environment.

5.4 Managing VMware vSphere with VMControl

This topic describes the supported tasks with the VMware virtualization environment. It also describes basic requirements and support for the VMware virtualization environment with IBM Systems Director VMControl. To work with the VMware virtualization environment, you need IBM Systems Director VMControl Express version.

If there is a new or existing VMware and IBM Systems Director environment, before implementing VMControl to that environment, it is important to know the main goals for this integration, and what capabilities and restrictions it provides to your management environment.

VMware vCenter Server is the central management component for VMware ESX/ESXi hosts. vCenter is used in almost all VMware environments, and is required to use VMware cluster features. IBM Systems Director uses its VMControl plug-in to interact with vCenter. IBM Systems Director does not aim at replacing vCenter.

VMControl uses the robust and virtualization-specialized vCenter to run tasks that are targeted at the VMware vSphere infrastructure components.

IBM Systems Director provides an essential collection of most commonly used tasks by a privileged administrator. By using these tasks, an enterprise administrator with full privileges can manage all platforms in your chassis from the single IBM Systems Director interface. In addition, administrative and more advanced tasks can be performed also directly on vCenter.

Additionally, integrating IBM Systems Director 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 IBM Systems Director, you can operate your system from a single pane of glass from both a hardware and software perspective. See the infrastructure on Figure 5-24 on page 283.

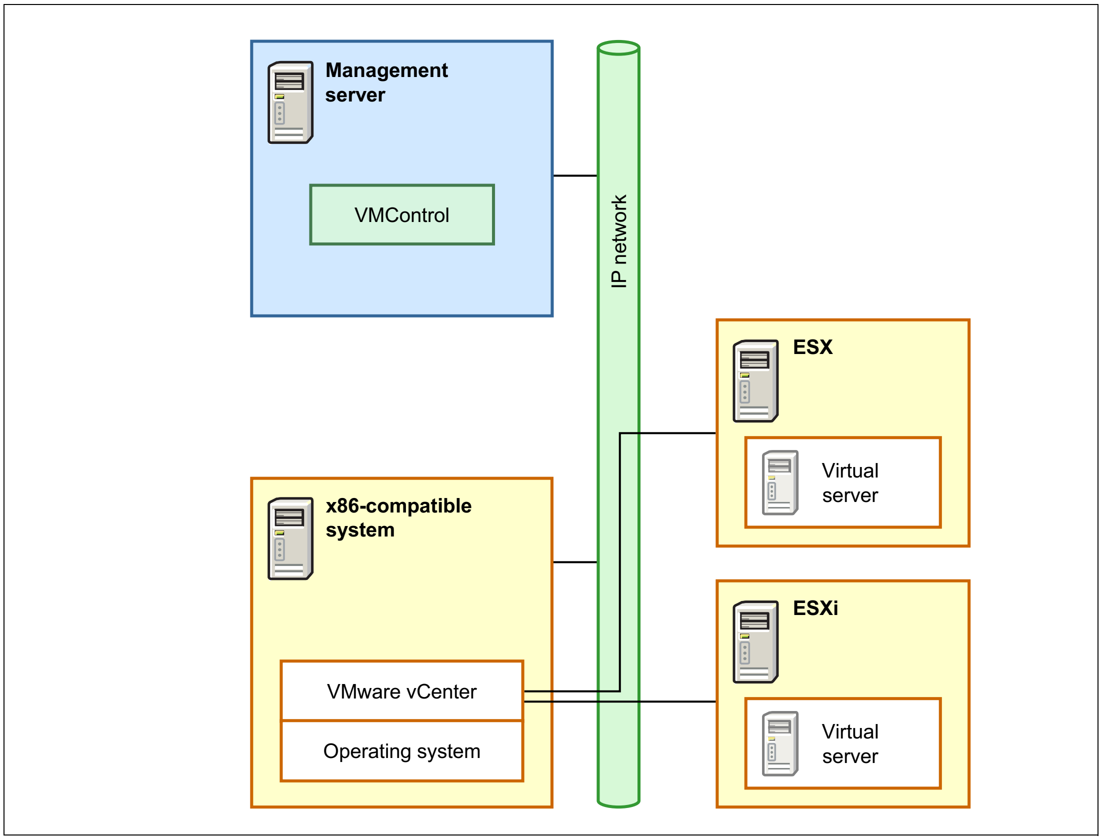

Figure 5-24 VMware virtualization environment with VMControl

The environment and requirements described at a high level:

- IBM Systems Director server is installed on a supported server.
- IBM Systems Director VMControl Express version is activated.

Note: To launch the VMware Infrastructure Client or the VMware vSphere Client from IBM Systems Director VMControl, the client must be installed on the IBM Systems Director server system and on any system that you use to log in to the IBM Systems Director web interface.

- VMware vCenter is installed on an x86-compatible system.

Note: IBM Systems Director and VMControl require that the OS that VMware vCenter is running on is an x86-compatible system with a Microsoft Windows-based OS.

- VMware ESXi exists to host virtual servers that you can manage using VMControl.

  Note: VMware ESXi is managed by VMware vCenter.
- The VMware vCenter system is discovered and the request access task has completed. After the request access task completes, the Configure Access task shows the vCenter protocol in an `OK` state.

Note: If you installed VMware vCenter with a non-default port number, you must create a VMware vCenter Server Discovery profile using the Discovery Profile wizard. Specify the unique port number in the profile that you create.

For more information about VMware requirements, see the following site:

http://pic.dhe.ibm.com/infocenter/director/pubs/index.jsp?topic=%2Fcom.ibm.director.vim.helps.doc%2Ffsd0_vim_r_vmware_vcenter.html

5.4.1 VMControl support for VMware vSphere

IBM Systems Director VMControl supports the following tasks with VMware vSphere:

- Create, edit, and delete virtual servers
- Create a data center or cluster using the Create Virtual Farm wizard
- Add a host to a data center or cluster using the “Add host to farm” function
- Remove a host from a data center or cluster using the “Remove host from farm” function
- Relocate virtual servers
- Put a host into maintenance mode
- Remove a host from maintenance mode

As listed, IBM Systems Director supports many general tasks with VMware. But if there is a need to use some advanced tasks that are not listed here, it is required to use vCenter for those actions.

Note: A source to work with IBM Systems Director and VMware virtualization with step-by-step screen captures can be found at the following site:

http://www.redbooks.ibm.com/redbooks/pdfs/sg248060.pdf

The IBM Redbooks publication is for the Flex System Manager, but all tasks related to VMware are suitable when working with IBM Systems Director.

The following tasks are supported by IBM Systems Director when working with VMware standard vSwitch and IBM Distributed Virtual Switch 5000V version 1.0.2:

- Discovery

Tip: You only need to discover the vCenter Server operating system. You do not need to discover the ESXi hosts directly because those will be added automatically during the vCenter Server managed endpoint (MEP) inventory. However, it is recommended to discover out-of-band management modules (such as IMM) manually for all your ESXi hosts in order to get accurate hardware status from your hosts and in order to be able to perform power operations out of band.

- Inventory
- Configuration management and automated logical network provisioning (ALNP)

Note: The management address for version 1.0.2 of the IBM System Networking Distributed Virtual Switch 5000V can be set up only in an IPv4 format.

Configuration management and automated logical network provisioning

You can work with the VLAN configuration and protocol configuration of some devices by using the IBM Systems Director Configuration Manager. Additionally, such devices can be members of network system pools, which will then support the automatic provisioning of logical network profiles, for VLAN settings, during virtual machine (VM) management.

For more information about automating logical network provisioning in a VMware environment, see the topic *Network Control*.

Note: The Disabling Protocol configuration on devices that are supported by configuration management will block communication between IBM Systems Director and the network device. In general, Protocol configuration should not be disabled.

You can use launch-in-context to access vendor management application for some devices directly from the IBM Systems Director interface. This function opens the vendor management application for the device you were viewing in IBM Systems Director. Tasks can then be completed from within the vendor management software.

Note: This IBM Systems Director Network Control task requires additional steps. See the topic “Set up IBM System Storage Data Center Fabric Manager (DCFM) to integrate with IBM Systems Director Network Control” for more information:

http://pic.dhe.ibm.com/infocenter/director/pubs/topic/com.ibm.sdnm.adv.helps.doc/fnc0_director_network_ctrl.pdf

Using IBM SmartCloud Entry with IBM Systems Director and VMware vSphere

If you want to publish your vCenter and VMware environment as a self-service portal, it is good to implement both IBM Systems Director and IBM SmartCloud Entry (SCE). In Figure 5-25 on page 286, IBM Systems Director, vCenter, and SCE components are shown, along with their discussion paths.

In Figure 5-25, IBM Systems Director monitors hardware, hypervisors, and vCenter. vCenter is connected to IBM SmartCloud Entry, which provides advanced deployment for virtual machines in VMware environments.

IBM Systems Director has a very important role in preventing hardware outages caused by service outages. IBM Systems Director commands vCenter to place the host in maintenance mode, which automates virtual server relocating before hardware problems occur.

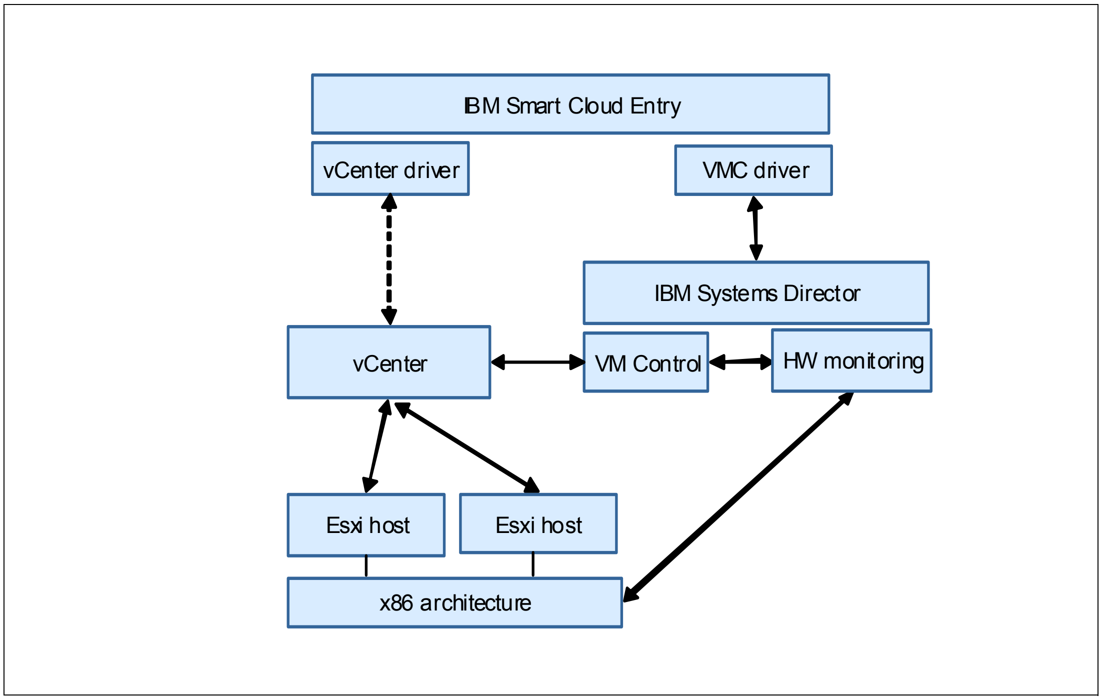

Figure 5-25 Component communication between Systems Director, vCenter, and SCE (if implemented)

Note: How to set up predictive failure alerts to do automated relocation in VMware is presented with step-by-step instructions in the following Redbooks publication:

`http://www.redbooks.ibm.com/redbooks/pdfs/sg248060.pdf`

Best Practice:

- Use IBM Systems Director for hardware management and monitoring.
- Use IBM Systems Director Predictive Failure Analysis (PFA) with VMWare to avoid hardware-caused outages.
- Use IBM SmartCloud Entry to provide a single-provisioning portal to create and manage virtual machines.

5.5 Managing Microsoft Hyper-V with VMControl

The management functionality provided by VMControl for Microsoft Hyper-V environments is limited to the most basic tasks, regardless of the VMControl edition you are using.

VMControl currently only supports Microsoft Windows Server 2008 R2 Hyper-V.

To understand the basic architecture of Hyper-V management with IBM Systems Director and VMControl, see the following information center site:

http://pic.dhe.ibm.com/infocenter/director/pubs/topic/com.ibm.director.vim.helps.doc/fsd0_vim_r_hyper_v.html

Hyper-V hosts that are managed by a Microsoft System Center Virtual Machine Manager (SCVMM) can also be managed at the same time by VMControl. Because VMControl only enables basic tasks, this might be appropriate for example when operators do not have privileged access to the SCVMM console.

In order to be able to manage Hyper-V from IBM Systems Director, perform the following steps:

1. Enable the VMControl plug-in in your IBM Systems Director server
2. Discover, unlock, and inventory the operating system of your Windows Server 2008 R2 server that has the Hyper-V role
3. Deploy the IBM Systems Director platform or common agent for Windows on your Hyper-V server, run an inventory, and unlock the physical server managed endpoint (MEP) that is discovered

> **Online Content:** For a video demonstration of the preceding tasks, see the following video:
>
> http://youtu.be/VUwp5JD1mws
>
> You can also scan the QR code that is displayed in the left margin to go directly to the video.

After those basic requirements have been met, you will be able to use the Virtual Servers and Hosts VMControl Inventory view and new basic actions will be added to the MEP actions menu.

In addition, if the Microsoft Windows Server 2008 R2 Hyper-v MEP is hosting virtual machines, those will be added to the IBM Systems Director resources when the host is inventoried. You will then be able to edit, stop, start, and remove those virtual servers on the host as well as view the relationships between hosts and virtual servers.

For more information about what tasks specifically are supported, refer to the following page in the IBM Systems Director information center:

http://pic.dhe.ibm.com/infocenter/director/pubs/index.jsp?topic=%2Fcom.ibm.director.vim.helps.doc%2Ffsd0_vim_r_virtualization_environments.html

In addition to the information that is available in the information center, here is a detailed list of supported operations on Hyper-v using VMControl:

1. Creating virtual servers: This action is performed on the Hyper-V physical server MEP or the Hyper-V OS MEP by right-clicking the MEP and selecting **System Configuration** → **Create Virtual Server**. When creating virtual servers with IBM Systems Director VMControl, you are limited to specifying the following information:
   a. The virtual server name
   b. The number of processors allocated to the virtual server
   c. The amount of memory allocated to the virtual server

d. The amount of virtual disk storage allocated to the virtual server. You can specify only one virtual disk. You cannot change the virtual disk location. It will use the Hyper-V server virtual storage default location.

e. To which network will the virtual server be attached? You can configure only one virtual network adapter. You will also not be able to specify a VLAN ID for that network connection.

Note: Creating a virtual server only creates the virtual machine and does not install the operating system in that virtual machine. Because you have no way of mounting ISO images to the newly created virtual server in VMControl, you will have to switch to the Hyper-V Manager management console in order to complete the virtual server installation.

2. Editing virtual servers: This action is performed on the virtual server MEP by right-clicking the MEP and selecting **System Configuration** → **Edit Virtual Server**. It is only available when the virtual server is powered off. You will be limited to changing the following components:

   a. The number of processors allocated to the virtual server

   b. The amount of memory allocated to the virtual server

3. Removing virtual servers: This action is performed on the virtual server MEP by right-clicking the MEP and selecting **Permanently Delete Virtual Server**. It is only available when the virtual server is powered off. This action will give you the option to delete the associated virtual disk storage, but it will not do so by default.

4. Viewing virtual servers and hosts relationship: You can view this relationship by selecting the **Virtual Servers and Hosts** inventory view, as shown in Figure 5-26 on page 289. The Hyper-V OS MEP where you installed the platform or common agent is only shown as an association with its physical server MEP in that view and not as an independent object.

5. Viewing virtual servers resource utilization: You can view CPU resource utilization for virtual servers from the **Virtual Servers and Hosts** inventory view, as shown in Figure 5-26 on page 289. You can also view the CPU and memory allocation from that window, but no other performance metric is supported.

6. Manage power operations on hosts and virtual servers: You can start, shut down, restart, or suspend virtual servers as well as start, shut down, and restart hosts. Power actions can be forced or performed cleanly, meaning that the operating system will gracefully shut down before the server or virtual server is powered off or restarted. From the host server MEP, you can force power off or start all virtual servers. If you need to gracefully shut down or restart multiple virtual servers, simply select each one in the **Virtual Servers and Hosts** inventory view (Figure 5-26 on page 289), then select **Actions** → **Power On/Off** → **Restart** or **Shutdown and power off**.

Warning: The **Power Off Now** and **Restart Now** actions will not perform a graceful shutdown and restart of the operating system (whereas the **Restart** and **Shutdown and power off** actions will). Only use those actions when the operating system is not installed or not responding.

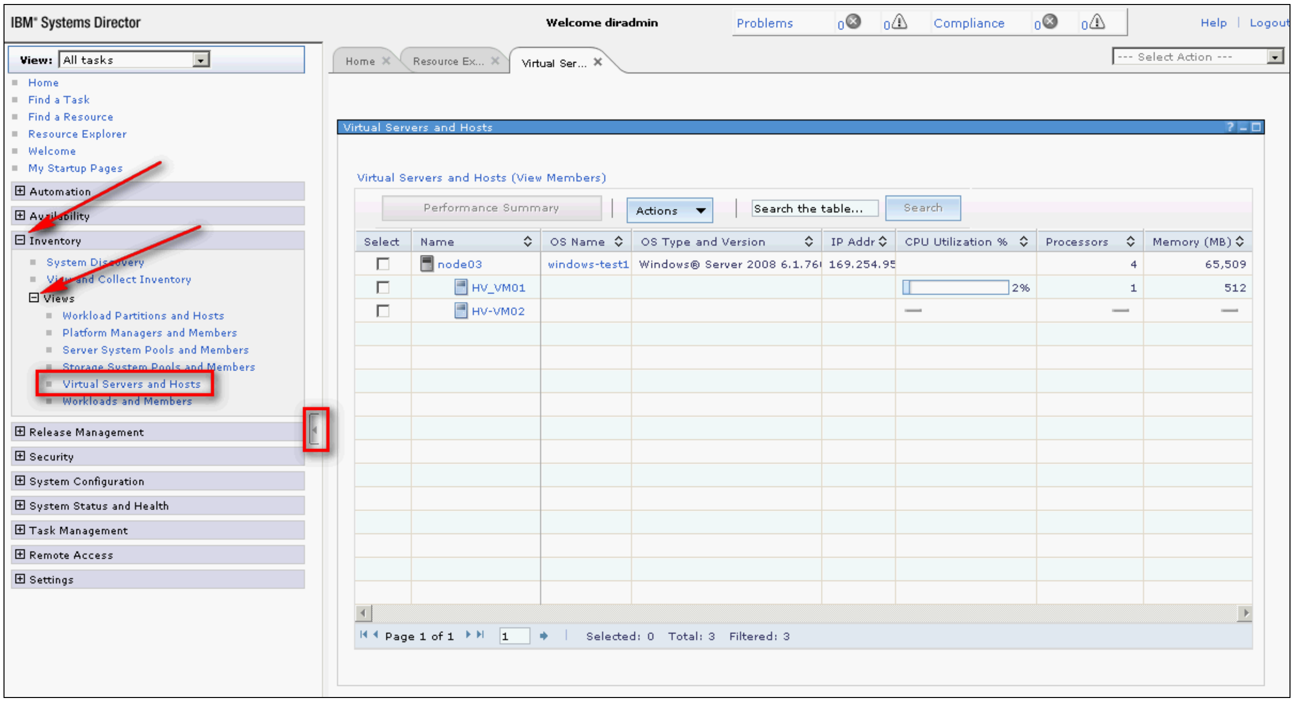

Select	Name	OS Name	OS Type and Version	IP Addr	CPU Utilization %	Processors	Memory (MB)
□	node03	windows-test1	Windows® Server 2008 6.1.76	169.254.95		4	65,509
□	HV_VM01				2%	1	512
□	HV-VM02				—	—	—

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

Figure 5-26 The Virtual Servers and Hosts inventory view

Online content: In addition, for a video demo of the preceding tasks, see the following link:

http://youtu.be/0KGu0nGw7nY

You can also scan the QR code that is displayed in the left margin to go directly to the video.

5.5.1 Best practices for managing Microsoft Hyper-V with VMControl

The following list contains recommendations of what you should and should not do while managing Microsoft Hyper-V with VMControl:

- *Always collect inventory on a Hyper-V server MEP before doing any action*: This helps ensure that you are not trying to take actions on objects that have been modified outside of VMControl and for which the information in the IBM Systems Director inventory is not up to date. Collecting up-to-date inventory on the virtual servers is not as critical because actions are taken at the hypervisor layer and not at the guest layer.
- *Always configure your Hyper-V host before managing it from VMContro*l: If you cannot create virtual machines on your Hyper-V host using the Microsoft Hyper-V Manager console, VMControl will also not be able to create virtual servers for you.
- Remember that *the Platform or Common Agent is only required on the host*, not in the guests.
- *There is no need to discover Hyper-V virtual machines directly*: They are discovered automatically when you inventory the host MEP (that is, the virtual server MEP that is being discovered, not the guest operating system MEP).
- *Always use the Virtual Servers and Hosts view* because this is the only view that updates dynamically. You can customize this view to add additional columns if it does not contain the information you require by default. The view refreshes every 30 seconds. The first time you access the view, it might take as long as 60 seconds to display current information.

- *Do not use "Remove" to delete virtual servers* because this will only remove the MEP from the IBM Systems Director inventory and leave the virtual machine on the Hyper-V host.

5.6 Using IBM SmartCloud Entry with VMControl

This topic describes IBM SmartCloud Entry main functionalities with IBM Systems Director VMControl. With IBM SmartCloud Entry 2.4, you can maintain control over the allocation of resources with a web-based application.

IBM SmartCloud Entry is implemented as a lightweight web-based application that runs as an Open Services Gateway initiative (OSGi) application.

You can perform with different hypervisors common public or private cloud operations, such as:

- Configuring multiple cloud support
- Provisioning and de-provisioning servers
- Drafting and cloning workloads
- Capturing workloads
- Starting and stopping servers as part of a workload
- Resizing existing servers
- Creating projects to give team-specific access to workloads
- Providing network configurations, which set unique network properties to different workloads
- Billing, accounting, and metering support
- Providing request and approval workflow support

5.6.1 IBM SmartCloud Entry hypervisor support

IBM SmartCloud Entry supports multiple hypervisors and multiple cloud connections. That means that you can use the most cost-efficient platform to produce your virtualization services.

Hardware, hypervisor, and hypervisor controller stack can then include multiple VMware environments and multiple VMControl environments. In the future, there will also be support for OpenStack driver for Microsoft Hyper-V. See Figure 5-27 on page 291.

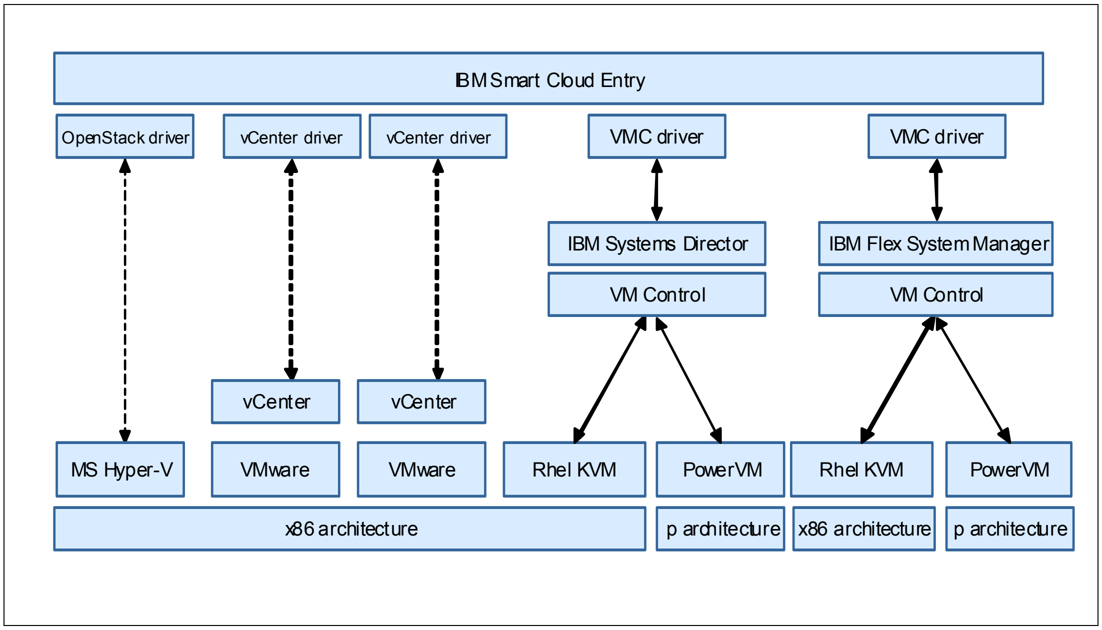

Figure 5-27 IBM SmartCloud Entry with multiple cloud sources and connections

IBM SmartCloud Entry can be a provisioning portal from small customers to large service providers that enables use of different cloud sources and a full view to the whole infrastructure from a single view.

Multi-architecture, multi-hypervisor, and multi-cloud connection capabilities are important features when calculating the total cost price for a virtual machine. Parameters should then include for example, architecture cost, hypervisor cost, guest operating system, and application cost.

These calculations are very useful for all and will guide future virtual server placements, where to deploy what virtual servers. The cost structure of hypervisors, guest OSs, and applications can be different when there is a shared infrastructure for multiple customers or it is dedicated. Remarkable cost savings can also be achieved when using, for example, different architecture with certain applications.

5.6.2 IBM SmartCloud Entry and VMControl

IBM SmartCloud Entry provides a web management portal for appliances that have been created with VMControl. To see all those workloads or virtual appliances, you need to first install IBM SmartCloud Entry and then configure a cloud connection to your VMControl application programming interface (API).

Cloud source configuration

To configure IBM Systems Director VMControl API against IBM SmartCloud Entry, you need to create a new cloud source from the administrator panel. Select a proper name for it and it is important to select the right connection type. When configuring VMControl, use the VMControl connection type.

The version is as important as the connection type. The version is dependent on which version your VMControl is. If you have the latest version of VMControl, use it. See Figure 5-28 on page 292.

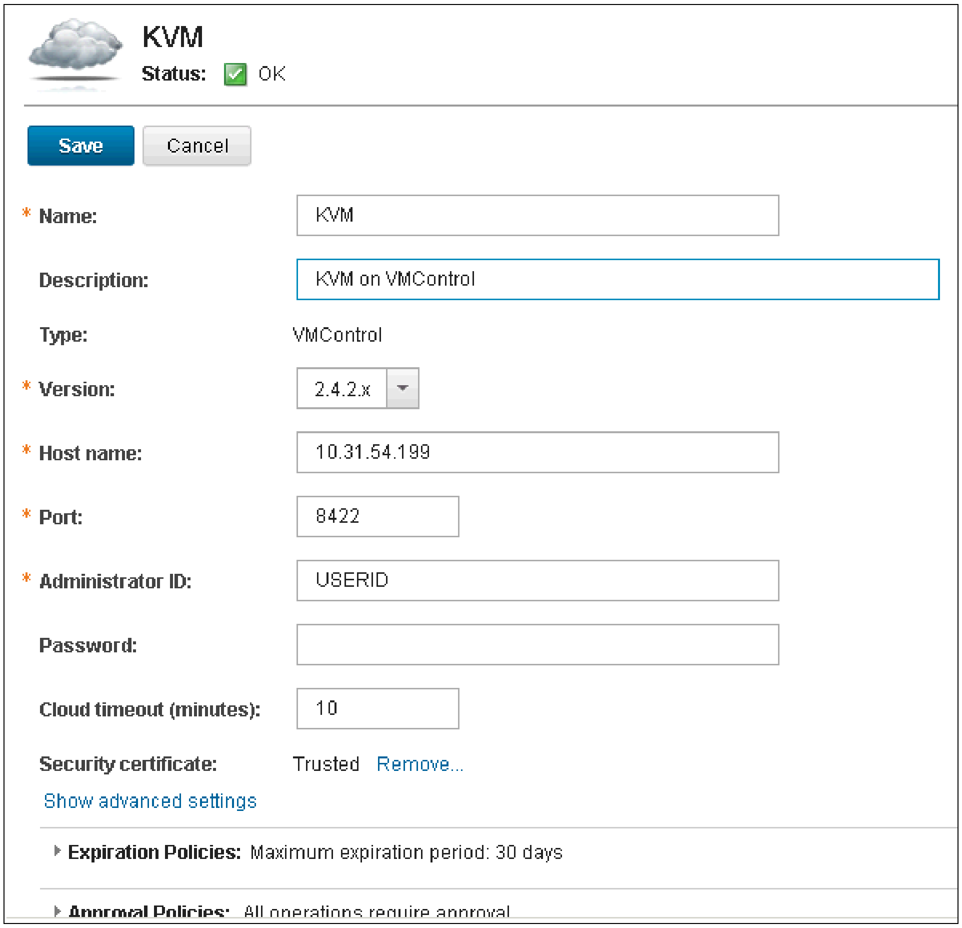

KVM
Status: ✓ OK

* Name:	KVM
Description:	KVM on VMControl
Type:	VMControl
* Version:	2.4.2.x
* Host name:	10.31.54.199
* Port:	8422
* Administrator ID:	USERID
Password:	
Cloud timeout (minutes):	10
Security certificate:	Trusted Remove...

▸ Expiration Policies: Maximum expiration period: 30 days

▸ Approval Policies: All operations require approval

Figure 5-28 IBM SmartCloud Entry cloud configuration for VMControl

After configuring this cloud source, IBM SmartCloud Entry shows running workloads and virtual appliances. Then, you can manage workloads and deploy new workloads through this web-based user interface. See Figure 5-29.

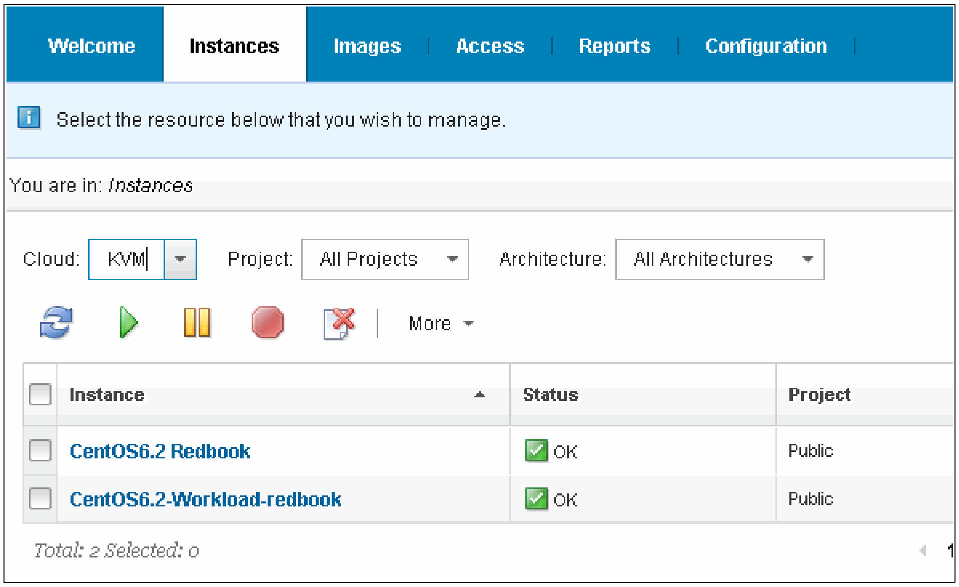

Select the resource below that you wish to manage.

You are in: *Instances*

Cloud: KVM Project: All Projects Architecture: All Architectures

□	Instance	Status	Project
□	CentOS6.2 Redbook	✓ OK	Public
□	CentOS6.2-Workload-redbook	✓ OK	Public

Total: 2 Selected: 0

Figure 5-29 IBM SmartCloud Entry and IBM Systems Director VMControl cloud source workloads

These two workloads are the same, which are running on the IBM Systems Director VMControl environment. See Figure 5-30.

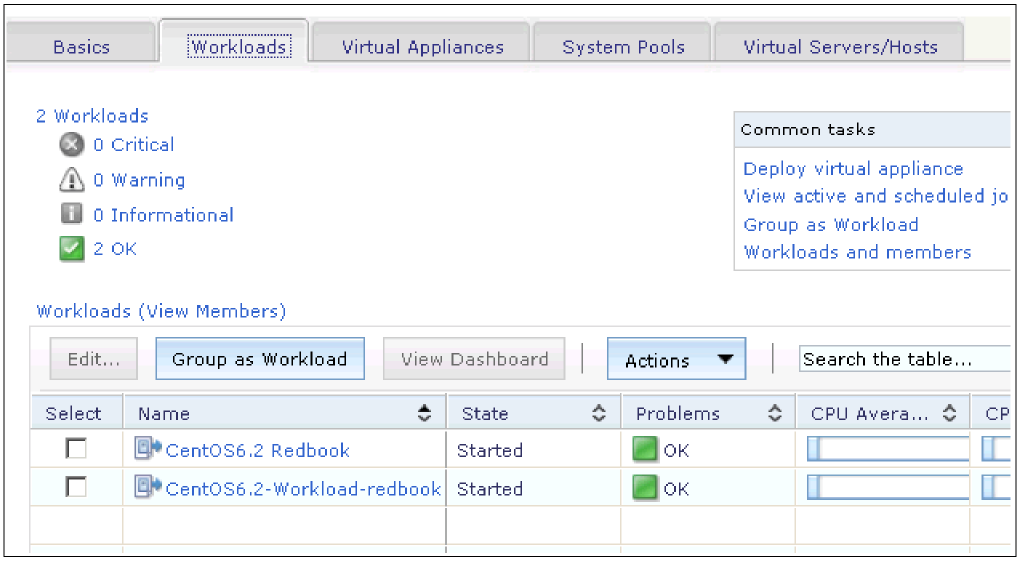

2 Workloads
- 0 Critical
- 0 Warning
- 0 Informational
- 2 OK

Workloads (View Members)

Select	Name	State	Problems	CPU Avera...	CP
□	CentOS6.2 Redbook	Started	OK		
□	CentOS6.2-Workload-redbook	Started	OK		

Figure 5-30 Workloads running on IBM Systems Director under VMControl

Deploying a new virtual server from an existing appliance

When you want to deploy a new workload from virtual appliances to deploy a new virtual server, you will see those also in IBM SmartCloud Entry. See Figure 5-31.

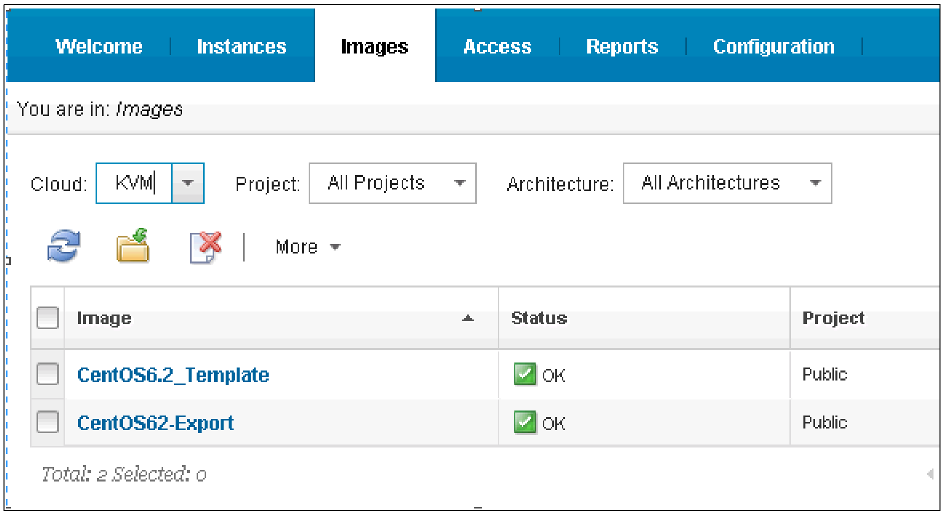

You are in: *Images*

Cloud: KVM Project: All Projects Architecture: All Architectures

□	Image	Status	Project
□	CentOS6.2_Template	OK	Public
□	CentOS62-Export	OK	Public

Total: 2 Selected: 0

Figure 5-31 Virtual appliances that are seen in IBM SmartCloud Entry

Virtual appliances come from the IBM Systems Director VMControl through an API. See Figure 5-32 on page 294.

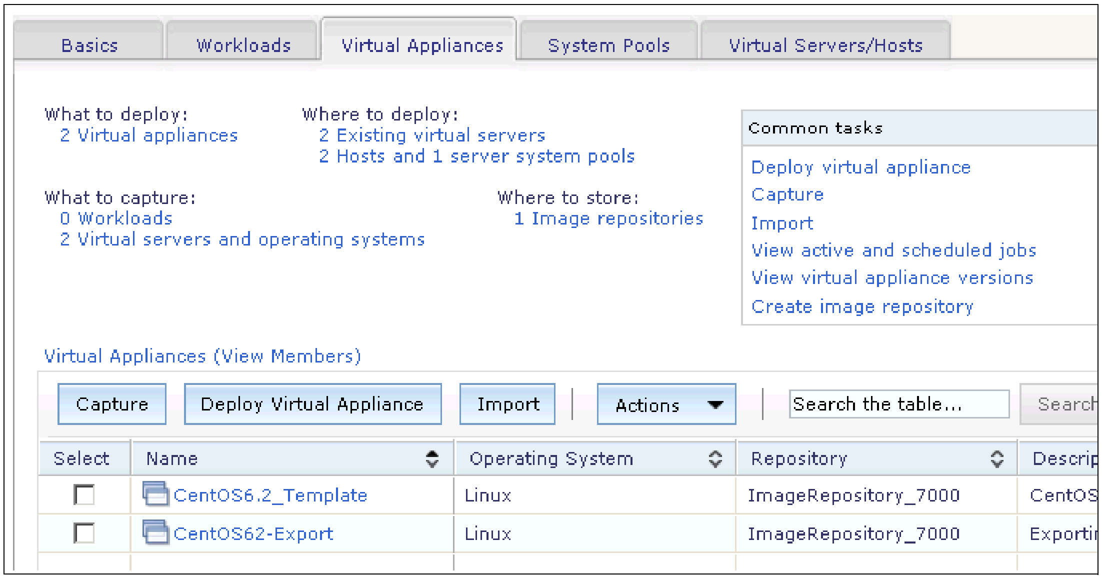

What to deploy:
2 Virtual appliances

Where to deploy:
2 Existing virtual servers
2 Hosts and 1 server system pools

What to capture:
0 Workloads
2 Virtual servers and operating systems

Where to store:
1 Image repositories

Virtual Appliances (View Members)

Select	Name	Operating System	Repository	Descrip
□	CentOS6.2_Template	Linux	ImageRepository_7000	CentOS
□	CentOS62-Export	Linux	ImageRepository_7000	Exporti

Figure 5-32 Appliances running under VMControl

Before deployment, you need to configure basic settings for the appliances. Configuration is done by using the IBM SmartCloud Entry user interface, and you will need to feed the basic parameters to each virtual appliance. Those parameters will guide the appliance, where, and how this appliance is deployed into the IBM Systems Director VMControl environment. See Figure 5-33 on page 295.

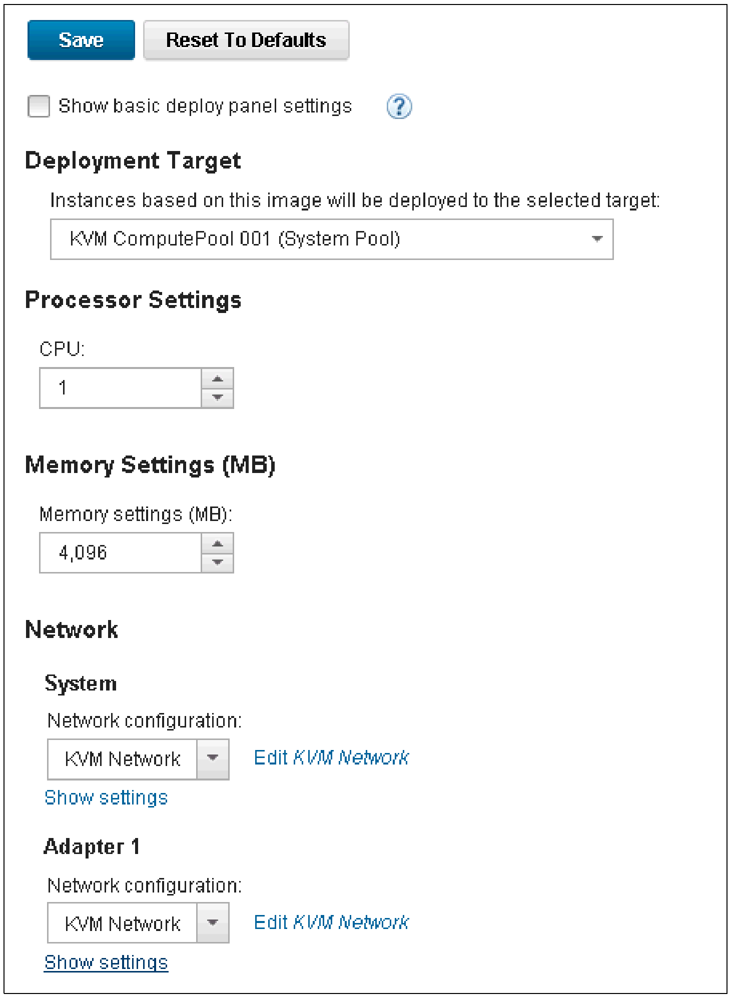

□ Show basic deploy panel settings

Deployment Target

Instances based on this image will be deployed to the selected target:
KVM ComputePool 001 (System Pool)

Processor Settings

CPU:
1

Memory Settings (MB)

Memory settings (MB):
4,096

Network

System

Network configuration:
KVM Network

Adapter 1

Network configuration:
KVM Network

Figure 5-33 Capture of some parameters related to the VMControl appliance

IBM Systems Director VMControl appliance deployment can be done from the IBM SmartCloud Entry user interface after configuring the appliance parameters.

You can give permissions from the IBM SmartCloud Entry configuration panels to gain access for end users, who can request new workloads and the administrator approves, rejects, or changes some deployment parameters. IBM SmartCloud Entry does have its own view for end users and administrators, and it supports approvals and different project views.

After placing parameters for a virtual appliance, to deploy a new server you need to choose the proper appliance and click **Deploy**. See Figure 5-34 on page 296.

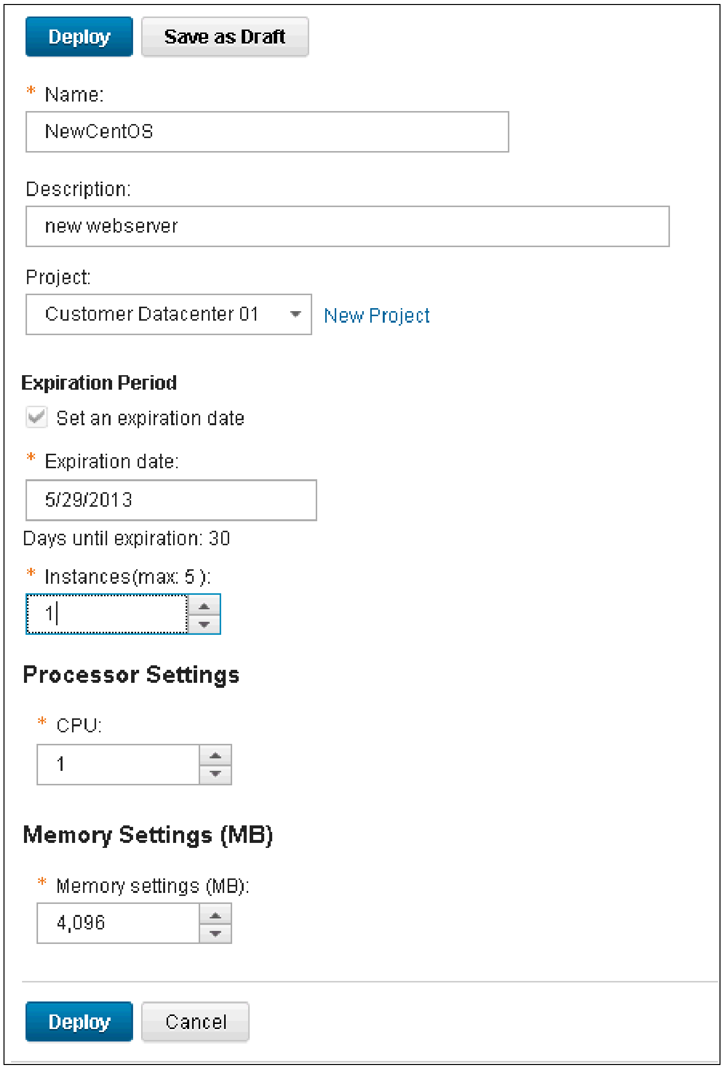

* Name: NewCentOS

Description: new webserver

Project: Customer Datacenter 01

Expiration Period

☑ Set an expiration date

* Expiration date: 5/29/2013

Days until expiration: 30

* Instances(max: 5): 1

Processor Settings

* CPU: 1

Memory Settings (MB)

* Memory settings (MB): 4,096

Figure 5-34 At the deployment phase, you can still adjust multiple parameters

When you have deployed a new workload from an existing appliance and the administrator has approved it, the new workload is displayed in VMControl. See Figure 5-35 on page 297.

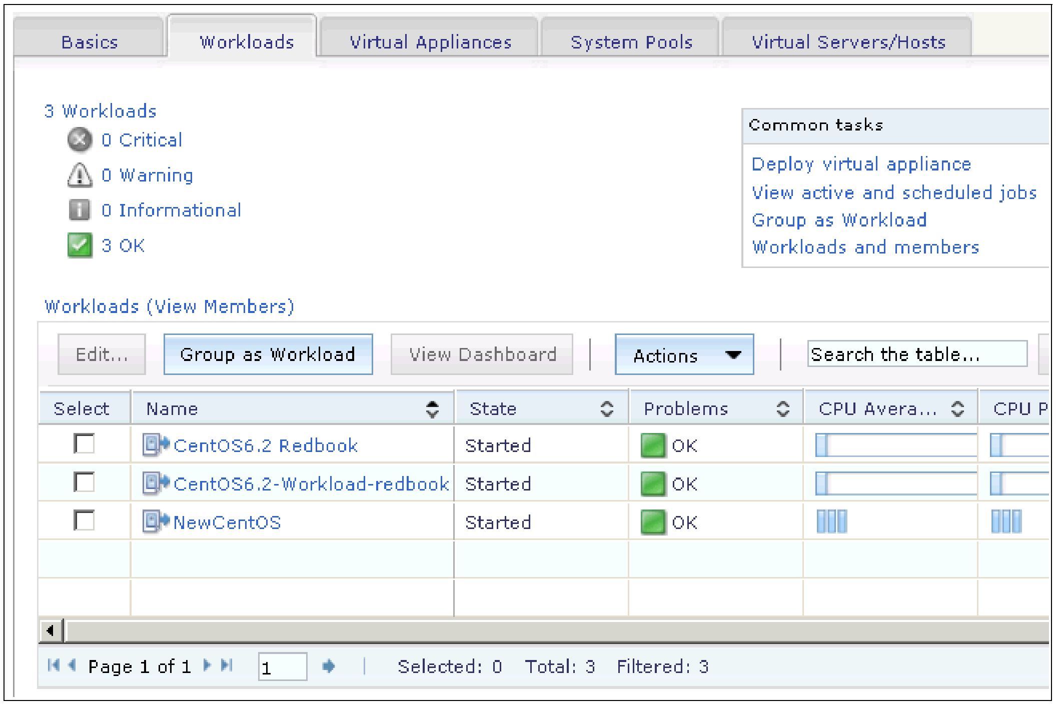

3 Workloads
- 0 Critical
- 0 Warning
- 0 Informational
- 3 OK

Workloads (View Members)

Select	Name	State	Problems	CPU Avera...	CPU P
□	CentOS6.2 Redbook	Started	🟩 OK		
□	CentOS6.2-Workload-redbook	Started	🟩 OK		
□	NewCentOS	Started	🟩 OK		

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

Figure 5-35 The new virtual server is displayed in VMControl after deployment

5.6.3 IBM SmartCloud Entry and VMware virtualization

This topic describes what extensions that the IBM SmartCloud Entry provides when working in VMware virtualization environments.

First, all basic features that are shown in the previous IBM SmartCloud Entry topic and IBM SmartCloud introduction chapter will apply also in the VMware environment. IBM SmartCloud Entry uses VMware vCenter API to communicate with the VMware environment. vCenter API brings more rich features to virtual server configurations such as resource pools, datastores, networks, or cluster selections.

Therefore, you can define very accurately where to deploy new workloads and what capabilities they have. See Figure 5-36 on page 298.

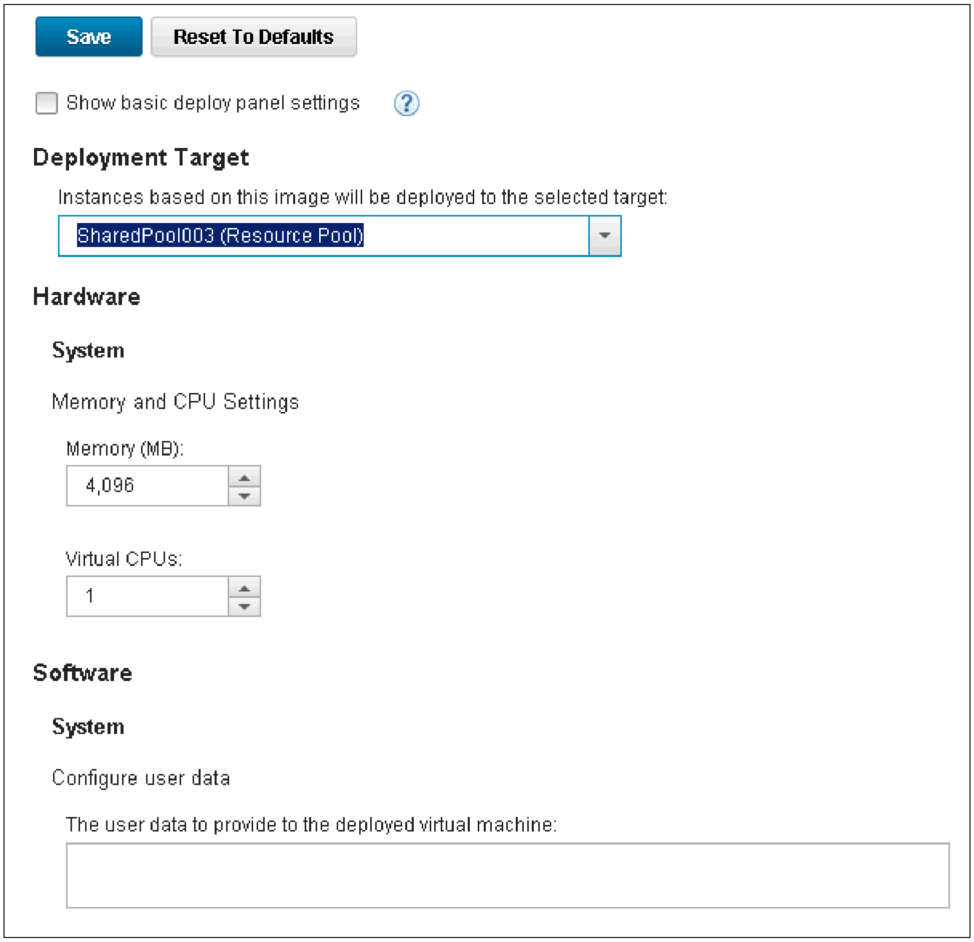

□ Show basic deploy panel settings

Deployment Target

Instances based on this image will be deployed to the selected target:

SharedPool003 (Resource Pool)

Hardware

System

Memory and CPU Settings

Memory (MB):

4,096

Virtual CPUs:

1

Software

System

Configure user data

The user data to provide to the deployed virtual machine:

Figure 5-36 IBM SmartCloud Entry VMware template parameters

After configuring the appliances, they are ready for deployment as explained in previous IBM SmartCloud Entry chapters.

> **Note:** More information about IBM SmartCloud Entry can be found at the following site:
>
> `https://www.ibm.com/developerworks/community/wikis/home?lang=en#!/wiki/W21ed5ba0f4a9_46f4_9626_24cbbb86fbb9`

5.7 Managing PowerVM with VMControl

PowerVM provides the industrial-strength virtualization solution for IBM Power Systems servers and blades. This solution provides proven workload consolidation that helps clients control costs and improves overall performance, availability, flexibility, and energy efficiency. PowerVM is a combination of hardware enablement and value-added software.

There are three versions of PowerVM, suited for various purposes:

- PowerVM Express Edition

  It provides more advanced virtualization features at a highly affordable price.

- PowerVM Standard Edition

It provides the most complete virtualization functionality for AIX, IBM i, and Linux OS.

- PowerVM Enterprise Edition

  It provides IBM Active Memory™ Sharing and Live Partitioning Mobility.

5.7.1 VMControl supported features on PowerVM

IBM Systems Director VMControl is a cross-platform product that assists you in rapidly deploying virtual appliances to create virtual servers that are configured with the operating system and software applications that you want. It also enables you to group resources into system pools, which enable you to centrally manage and control the different workloads in your environment.

Table 5-1 shows which tasks are supported in each virtualization environment on each Systems Director edition.

Table 5-1 VMControl supported features

IBM Systems Director edition	Supported Tasks	IBM Power Systems in a virtualized environment[1]	VMware vCenter and VMware ESX or ESXi, managed by VMware vCenter	Windows Server 2008, Ent., Std., and Datacenter x64 Editions with Hyper-V role enabled	Linux kernel-based virtual machine (KVM)
Express Editions	Create, edit, and delete virtual servers	Yes	Yes	Yes	Yes
	Create, edit, and delete virtual farms	Yes[2]	Yes	-	Yes
	Relocate virtual servers	Yes	Yes	-	Yes
	Put hosts into maintenance mode and remove hosts from maintenance mode	Yes	Yes	-	Yes
Standard Edition	Import virtual appliance packages	Yes	-	-	Yes
	Capture virtual servers	Yes	-	-	Yes
	Capture workloads	Yes	-	-	Yes
	Deploy virtual appliances	Yes	-	-	Yes
	Start, stop, and delete workloads	Yes	-	-	Yes
	Create, edit, and delete workloads	Yes	-	-	Yes[3]
Enterprise Edition	Create, edit, and delete server system pools	Yes[4]	-	-	Yes
	Create, edit, and delete storage system pools	Yes	-	-	-

1. For systems that are managed by Hardware Management Console, IBM Flex System Manager, or Integrated Virtualization Manager.
2. Not supported for IBM Flex System Manager.

3. This virtualization environment does not support encapsulating virtual servers into system pool management in groups as workload.
4. For systems that are capable of PowerVM live partition mobility, with the Active Mobility property set to `True`.

5.7.2 Supported Power virtualization environments

Described in this section are the virtualization environments that are supported by IBM Systems Director VMControl.

AIX using Network Installation Manager

Ensure that your environment satisfies the support and requirements for the components in a Network Installation Manager (NIM)-based Power Systems virtualization environment, including the management server, host, operating systems, and storage.

Figure 5-37 shows an example Power Systems virtualization environment for AIX virtual appliances, virtual servers, and workloads that rely on NIM. In this example, the Power Systems server is managed by the HMC.

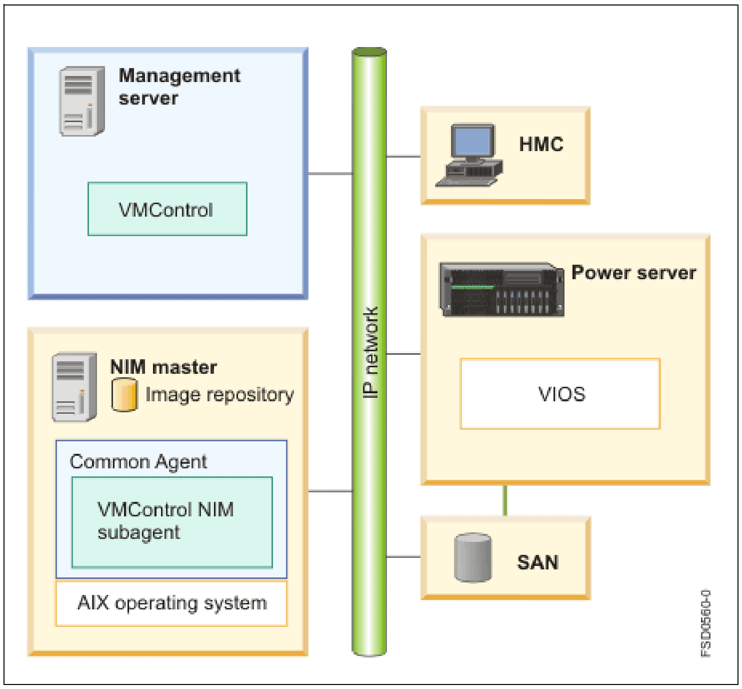

Figure 5-37 NIM-based Power Systems virtualization environment diagram

Following are the required components for a PowerVM environment using NIM:

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

> **Restriction:** You must install the `dsm.core` file set on the NIM master before the NIM subagent installs successfully.

IBM Systems Director server recognizes this NIM master as a VMControl image repository. The `/export/nim` file system 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 the export of a mounted file system.

> **Note:** 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 Power Systems 5, 6, or 7 server or blade exists to host virtual servers that you can capture from and deploy to using VMControl.

> **Notes:** If you plan to use N-Port ID Virtualization (NPIV) with Fibre Channel storage, the Power Systems server must be a POWER6 processor-based server, or higher.
>
> If manual or automated virtual server relocation capabilities are needed, multiple Power Systems 6 or 7 servers are required.

- The Power Systems server can be managed by the HMC, as shown in Figure 5-37 on page 300, or by IVM.

> **Note:** For blades and low-end Power Systems servers, you can use IVM on the Virtual I/O Server (VIOS) virtual server instead of the HMC.

- The Power Systems server is typically attached to a SAN as shown in the diagram. The SAN is used for the Fibre Channel storage or the virtual disks of the virtual servers that are hosted by the Power Systems server. A SAN is required to use the following capabilities:
  - VMControl Enterprise Edition server system pools
  - Relocation
  - NPIV

  If you do not plan to use these capabilities, a SAN is not required. Alternately, you can use disks that meet all of the following criteria:
  - Disks that are local to the Power Systems server
  - Disks that are virtualized by the VIOS
- Though not shown in the diagram, multiple VIOS virtual servers and multi-path I/O (MPIO) are supported.

For more information about supported AIX versions, firmware versions, and storage, see the following site:

http://pic.dhe.ibm.com/infocenter/director/pubs/index.jsp?topic=%2Fcom.ibm.director.vim.helps.doc%2Ffsd0_vim_r_power_component_reqs.html

Supported tasks

In this environment, you can perform the following tasks:

- Create, edit, and delete virtual servers
- Relocate virtual servers
- Import a virtual appliance package containing an AIX **mksysb** image

- Capture an AIX workload or virtual server, an AIX **mksysb** image file or NIM resource, or an AIX `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

AIX using Storage Copy Services

Ensure that your environment satisfies the support and requirements for the components in a Storage Copy Services (SCS)-based Power Systems virtualization environment, including the management server, host, operating systems, and storage.

Figure 5-38 shows an example Power Systems virtualization environment for AIX, IBM i, and Linux virtual appliances, virtual servers, and workloads that rely on SCS. In this example, the Power Systems server is managed by the HMC.

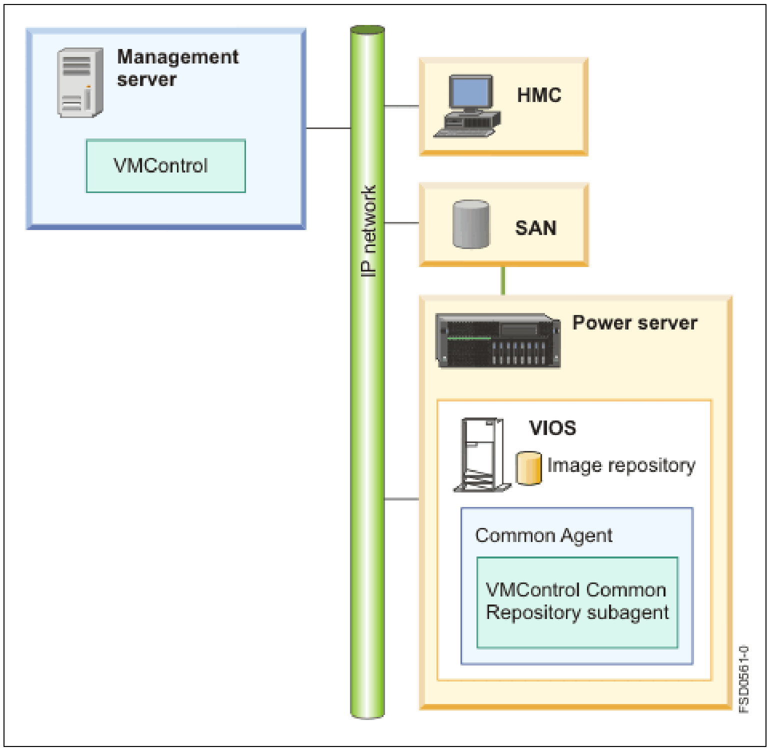

Figure 5-38 SCS-based Power Systems virtualization environment diagram

Following are the required components for a PowerVM environment using SCS:

- IBM Systems Director server is installed on a supported server.
- IBM Systems Director VMControl Standard Edition or IBM Systems Director VMControl Enterprise Edition is activated.
- A VIOS virtual server exists on a Power Systems server. The VIOS virtual server hosts the image repository used to store the raw disk images associated with your AIX, IBM i, and Linux virtual appliances.

Note: You can have multiple repositories. However, for repositories that are on separate Power Systems servers, 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 capture and deploy.

- The IBM Systems Director Common Agent and the VMControl Common Repository subagent are installed on the VIOS that you want to use as an image repository.
- At least one Power Systems 5, 6, or 7 server or blade exists to host virtual servers that you can capture from and deploy to using VMControl. If you plan to use NPIV with Fibre Channel storage, the Power Systems server must be a POWER6 processor-based server, or higher.
- The Power Systems server can be managed by the HMC as shown in Figure 5-38 on page 302, or by IVM.

Note: For blades and low-end Power Systems servers, you can use IVM on the VIOS virtual server instead of the HMC.

- All AIX, IBM i, and Linux virtual servers to be captured from or deployed to using VMControl:
  - Have their storage allocated from the SAN
  - Must use virtual Ethernet connections or Fibre Channel connections that are provided through one or more VIOS virtual servers. The AIX, IBM i, and Linux virtual servers must not have any physical devices allocated from the Power Systems server.
- For VIOS Version 2.2.2.0, any images for virtual servers that you capture, and that are using a shared storage pool, must be captured into a repository that uses the same shared storage pool. Any virtual servers that you deploy must use the same shared storage pool as the image repository in which you store the virtual appliance images.

For more information about supported operating systems and firmware versions, see the following site:

http://pic.dhe.ibm.com/infocenter/director/pubs/topic/com.ibm.director.vim.helps.doc/fsd0_vim_r_sb_aix_on_power_component_reqs.html

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 containing 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

Table 5-2 lists the virtualization tasks supported through HMC and IVM.

Table 5-2 Supported virtualization tasks through IVM and HMC

Virtualization tasks	IVM	HMC
Included in PowerVM	Yes	
Manage Power Blades	Yes	
Manage more than one server		Yes
Hardware monitoring	Yes	Yes
Service agent call home	Yes	Yes
Graphical interface	Yes	Yes
Requires a separate server to run on		Yes
Advances PowerVM features		Yes
High-end servers		Yes
Low-end and midrange servers	Yes	Yes
Redundant setup	Yes	Yes

5.7.3 How to activate IBM Systems Director VMControl

You must activate IBM Systems Director VMControl before it can be used. VMControl comes with an evaluation license, which enables use of the optional chargeable (fee-based) management functions.

There are two ways to activate VMControl:

Activating VMControl using the IBM Systems Director web interface

1. From the IBM Systems Director home page, click the Plug-ins tab.
2. In the “Additional plug-ins to activate” section (Figure 5-39), click **Activate now** under the IBM Systems Director VMControl heading.

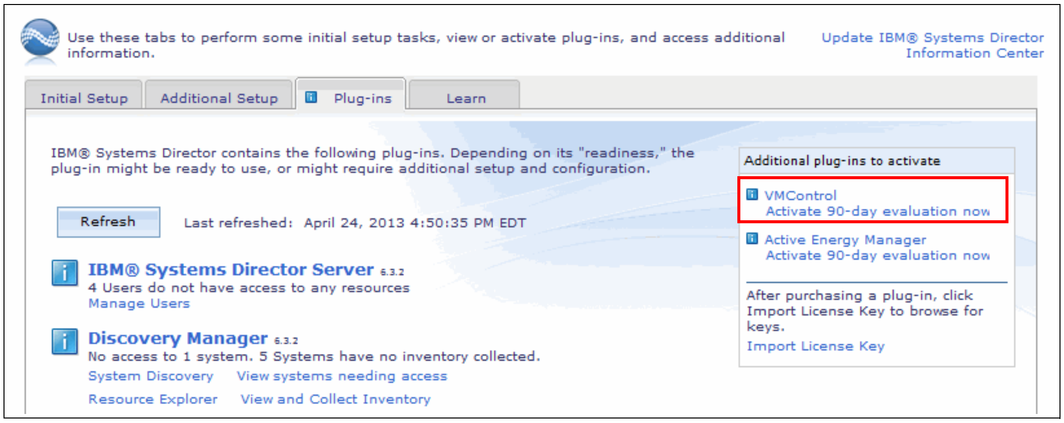

Use these tabs to perform some initial setup tasks, view or activate plug-ins, and access additional information.

Update IBM® Systems Director Information Center

IBM® Systems Director contains the following plug-ins. Depending on its "readiness," the plug-in might be ready to use, or might require additional setup and configuration.

Last refreshed: April 24, 2013 4:50:35 PM EDT

IBM® Systems Director Server 6.3.2
4 Users do not have access to any resources
Manage Users

Discovery Manager 6.3.2
No access to 1 system. 5 Systems have no inventory collected.
System Discovery View systems needing access
Resource Explorer View and Collect Inventory

Additional plug-ins to activate

VMControl
Activate 90-day evaluation now

Active Energy Manager
Activate 90-day evaluation now

After purchasing a plug-in, click Import License Key to browse for keys.
Import License Key

Figure 5-39 Activate VMControl

3. When the activation has completed, the following message is displayed (Figure 5-40). Then, restart IBM Systems Director server.

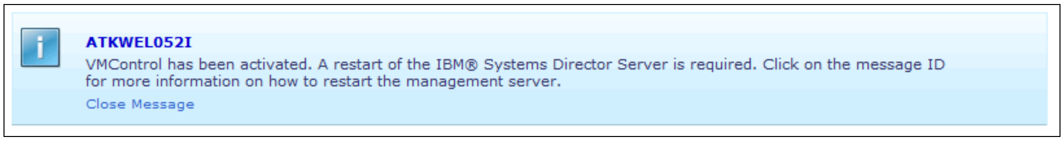

ATKWEL052I
VMControl has been activated. A restart of the IBM® Systems Director Server is required. Click on the message ID for more information on how to restart the management server.

Figure 5-40 VMControl activated

Activating VMControl using the command-line interface

1. From a command-line prompt, type the following string, and then press Enter:

```
smcli activatemgrs VMControl
```

2. When the activation has completed, restart IBM Systems Director server.

VMControl permanent license key

When you activate IBM Systems Director VMControl Express Edition, you are granted an evaluation period for IBM Systems Director VMControl Standard Edition and IBM Systems Director VMControl Enterprise Edition. When the evaluation period expires, you must purchase a license and install a license key to continue using these editions. You can install this license key by using an installation wizard. More information about license keys is shown at the following site:

http://pic.dhe.ibm.com/infocenter/director/pubs/index.jsp?topic=%2Fcom.ibm.director.vim.helps.doc%2Ffsd0_vim_t_installing.html

You can also install the license key silently by using a response file to specify information that the installation program requires. See the following site:

http://pic.dhe.ibm.com/infocenter/director/pubs/index.jsp?topic=%2Fcom.ibm.director.vim.helps.doc%2Ffsd0_vim_t_installing_silent_perm.html

5.7.4 Preparing the PowerVM environment for VMControl

Before you can use IBM Systems Director VMControl Standard Edition and IBM Systems Director VMControl Enterprise Edition, you must prepare your PowerVM environment, set up an image repository, and install the appropriate VMControl agent or subagent on the system that hosts the image repository.

Table 5-3 lists the AIX versions and firmware required for VMControl for each supported environment.

Table 5-3 Supported AIX and firmware versions for VMControl

	Network Installation Manager (NIM)	Storage Copy Services (SCS)
IBM Systems Director server	IBM Systems Director server V6.3.2 or later with IBM Systems Director VMControl V2.4 activated.	IBM Systems Director server V6.3.2 or later with VMControl Standard Edition or VMControl Enterprise Edition V2.4 activated.
NIM master	AIX 6.1 TL03 or newer.	-
HMC	HMC V7R3.5 and all available updates.	

<table>
<tr><th></th><th>Network Installation Manager (NIM)</th><th>Storage Copy Services (SCS)</th></tr>
<tr><td>VIOS</td><td colspan="2">► POWER5 and POWER6: Use a minimum of VIOS 2.1.2.0 and all available updates.
► POWER7: Use a minimum of VIOS 2.2.1.0 and all available updates.</td></tr>
<tr><td>Power Systems firmware</td><td>-</td><td>► POWER5 and POWER6 servers, use a minimum of FW3.5 and all available updates.
► POWER7 processor-based servers, use a minimum of FW7.2 and all available updates.</td></tr>
<tr><td>Operating systems</td><td>-</td><td>► AIX: AIX Version 5.3 TL9, or later; AIX Version 6.1 TL2, or later; or AIX Version 7.1.
► Linux on Power Systems: SUSE Linux Enterprise Server 10 SP3, or later; or SUSE Linux Enterprise Server 11.
► IBM i: IBM i 7.1 TR3, or later.</td></tr>
<tr><td>Virtual appliances</td><td>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, and you can import or deploy any AIX Version 5.3, AIX Version 6.1, or AIX Version 7.1 virtual appliance.</td><td>-</td></tr>
<tr><td>EMC requirements</td><td>VIOS 2.2.1.4 or later on your Power Systems hosts.</td><td>-</td></tr>
<tr><td>NPIV</td><td colspan="2">► NPIV is supported only if the following two conditions are met:
– The network switch must be an IBM Flex System FC5022 16 Gb SAN Scalable Switch.
– The SAN storage must be virtualized by IBM System Storage SAN Volume Controller (SVC) or the storage array is either an IBM Flex System V7000 Storage Node or an IBM Storwize V7000 system.
► You cannot manage the storage devices by using the SMI-S provider through IBM Systems Director.
► NPIV supports multi-disk virtual server disk attachment.
► Each disk that is allocated to the operating system on the virtual server must access its storage through a VIOS virtual Small Computer System Interface (SCSI) path or a VIOS NPIV path.</td></tr>
</table>

Discovering your storage and switch provider

See Chapter 11, “Storage Management solutions and Storage Control” on page 451 for detailed steps to run the discovery of the storage and the switch provider.

Discovering your Power Systems environment

Run the HMC discovery as a regular discovery operation. After the discovery, request access to the HMC and collect inventory. Figure 5-41 shows an example of an HMC discovered and inventoried.

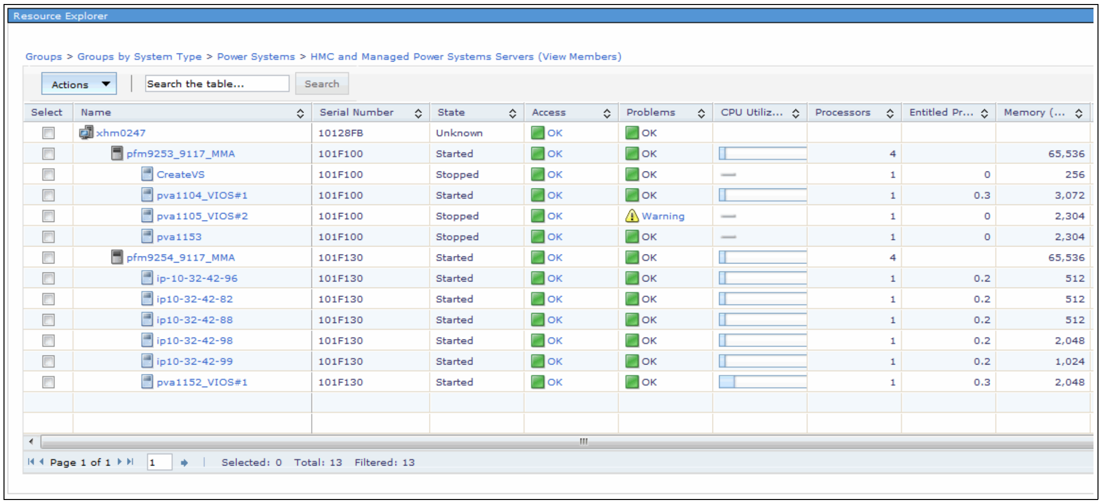
Resource Explorer
Groups > Groups by System Type > Power Systems > HMC and Managed Power Systems Servers (View Members)
Actions
Search the table...
Search

Select	Name	Serial Number	State	Access	Problems	CPU Utiliz...	Processors	Entitled Pr...	Memory (...
□	xhm0247	10128FB	Unknown	🟩 OK	🟩 OK				
□	pfm9253_9117_MMA	101F100	Started	🟩 OK	🟩 OK		4		65,536
□	CreateVS	101F100	Stopped	🟩 OK	🟩 OK	—	1	0	256
□	pva1104_VIOS#1	101F100	Started	🟩 OK	🟩 OK		1	0.3	3,072
□	pva1105_VIOS#2	101F100	Stopped	🟩 OK	⚠ Warning	—	1	0	2,304
□	pva1153	101F100	Stopped	🟩 OK	🟩 OK	—	1	0	2,304
□	pfm9254_9117_MMA	101F130	Started	🟩 OK	🟩 OK		4		65,536
□	ip-10-32-42-96	101F130	Started	🟩 OK	🟩 OK		1	0.2	512
□	ip10-32-42-82	101F130	Started	🟩 OK	🟩 OK		1	0.2	512
□	ip10-32-42-88	101F130	Started	🟩 OK	🟩 OK		1	0.2	512
□	ip10-32-42-98	101F130	Started	🟩 OK	🟩 OK		1	0.2	2,048
□	ip10-32-42-99	101F130	Started	🟩 OK	🟩 OK		1	0.2	1,024
□	pva1152_VIOS#1	101F130	Started	🟩 OK	🟩 OK		1	0.3	2,048

Page 1 of 1 | 1 | Selected: 0 Total: 13 Filtered: 13

Figure 5-41 Example of HMC discovered and inventoried

For more information about configuring HMC and IVM, see the following site:

http://pic.dhe.ibm.com/infocenter/director/pubs/index.jsp?topic=%2Fcom.ibm.director.vim.helps.doc%2Feica7_t_configuring_power_plat_managers.html

> **Note:** Ensure that the ID that will be used to request access to the HMC has `hmcsuperadmin` rights or `hmcoperator` rights.

Discover Virtual I/O Server

Before discovering your VIOS, check the following requisites:

- Ensure that your VIOS servers are not managed by any other IBM Systems Director instance. Check that there are no certificates in the following directory: `/opt/ibm/director/agent/runtime/agent/cert`
- Check if the CAS agent is running: `/opt/ibm/director/agent/runtime/agent/bin/endpoint.sh status`

After the VIOS is discovered, request access and collect inventory (Figure 5-42).

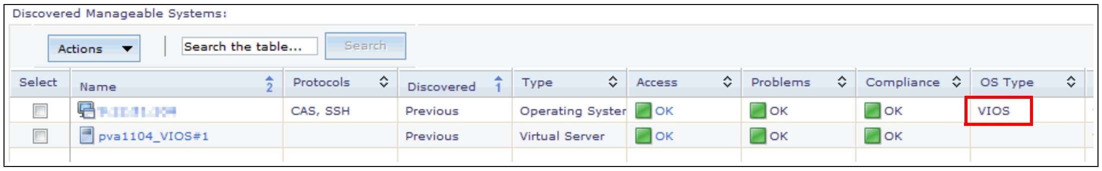

Select	Name	Protocols	Discovered	Type	Access	Problems	Compliance	OS Type
□	[redacted]	CAS, SSH	Previous	Operating Syster	🟩 OK	🟩 OK	🟩 OK	VIOS
□	pva1104_VIOS#1		Previous	Virtual Server	🟩 OK	🟩 OK	🟩 OK	

Figure 5-42 VIOS discovered

Install IBM Systems Director VMControl NIM subagent

You can install a VMControl subagent by using the installation wizard or you can manually install the subagent. Ensure that your system meets the following requirements:

- AIX 6.1.3 or later
- The following file set is installed: `dsm.core`

- Secure shell is installed and configured
- The systems must be configured as a NIM master system:

```
#lsnim -l master
Cstate = ready for a NIM operation
```

- Ensure that your NIM server has password-less access to the HMC/IVM using the SSH keys. For more information about how to configure your SSH keys, see the following site:

  http://www-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/TD101248

- Ensure that the system clocks on the IBM Systems Director server and the target system match as closely as possible. If the system clocks do not match closely enough, installation of the Common Repository subagent might fail.
- Check that **tftpd** and **bootps** are active on NIM:

```
#ps -ef | grep -i tftpd
#lssrc -t bootps
Service Command Description Status
bootps /usr/sbin/bootpd bootpd /etc/bootptab active
```

- The NIM repository must require additional configuration if the NIM repository (NIM master) connects to IBM Systems Director through one network adapter, and connects to the virtual server where you plan to deploy a virtual appliance through a different network adapter. For more details, see the following site:

  http://pic.dhe.ibm.com/infocenter/director/pubs/topic/com.ibm.director.vim.helps.doc/fsd0_vim_t_installing_agent_NIM.html

To install a VMControl subagent from the IBM Systems Director Release Management task, follow these steps:

1. In the IBM Systems Director navigation pane, expand **Release management**.
2. Click **Agents**.
3. On the Agents page, click **Common Agent Subagent Packages**.
4. From the Common Agent Subagent Packages view, select the subagent that you want to install.
5. Click **Actions** on the menu bar, and select **Release Management** → **Install Agent**.
6. Follow the instructions in the installation wizard to install the subagent for your virtualization environment.
7. When it is complete, the NIM server will be shown in IBM Systems Director (Figure 5-43).

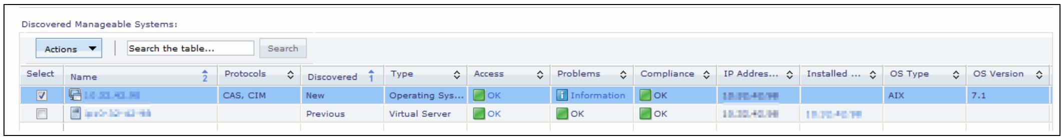

Discovered Manageable Systems:

Select	Name	Protocols	Discovered	Type	Access	Problems	Compliance	IP Addres...	Installed ...	OS Type	OS Version
☑	[redacted]	CAS, CIM	New	Operating Sys...	OK	Information	OK	[redacted]		AIX	7.1
□	[redacted]		Previous	Virtual Server	OK	OK	OK	[redacted]	[redacted]		

Figure 5-43 NIM server discovered

For more information about VMControl agents and subagents, see the following site:

http://pic.dhe.ibm.com/infocenter/director/pubs/index.jsp?topic=%2Fcom.ibm.director.vim.helps.doc%2Ffsd0_vim_t_installing_agent.html

8. Collect the inventory from your NIM server.
9. Finally, your VMControl is ready to be used. See Figure 5-44 on page 309.

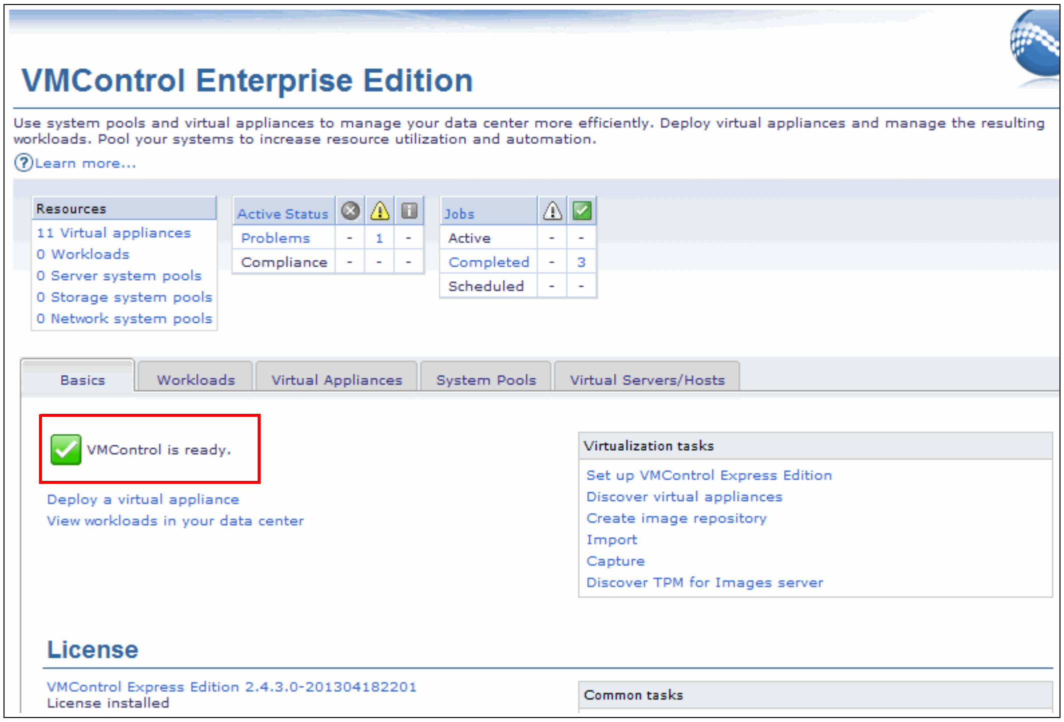

VMControl Enterprise Edition

Use system pools and virtual appliances to manage your data center more efficiently. Deploy virtual appliances and manage the resulting workloads. Pool your systems to increase resource utilization and automation.

Learn more...

Resources
11 Virtual appliances
0 Workloads
0 Server system pools
0 Storage system pools
0 Network system pools

Active Status	⊗	⚠	ℹ
Problems	-	1	-
Compliance	-	-	-

Jobs	⚠	✓
Active	-	-
Completed	-	3
Scheduled	-	-

✓ VMControl is ready.

Deploy a virtual appliance
View workloads in your data center

Virtualization tasks

Set up VMControl Express Edition
Discover virtual appliances
Create image repository
Import
Capture
Discover TPM for Images server

License

VMControl Express Edition 2.4.3.0-201304182201
License installed

Common tasks

Figure 5-44 VMControl ready

5.7.5 Storage allocation supported by IBM Systems Director VMControl

VMControl supports different types of storage allocation: Virtual SCSI, NPIV, and shared storage pool:

- **Virtual SCSI**

  Virtual Small Computer System Interface (SCSI) is used to refer to a virtualized implementation of the SCSI protocol. Virtual SCSI requires POWER5 or later hardware with the PowerVM feature activated. It provides virtual SCSI support for AIX, IBM i (requires POWER6 or later), and supported versions of Linux.

  Virtual SCSI is based on a client/server relationship. The Virtual I/O Server owns the physical resources and acts as a server or, in SCSI terms, *target device*. The client logical partitions access the virtual SCSI backing-storage devices provided by the Virtual I/O Server as clients.

  Three types of storage allocation that are supported by Systems Director can be grouped under virtual SCSI: virtual disk, physical volumes, and shared storage pool.

- **NPIV or virtual Fibre Channel**

  N-Port ID Virtualization (NPIV) is also known as *virtual Fibre Channel*. NPIV allows an NPIV-capable Fibre Channel adapter to be configured with multiple virtual world-wide port names (WWPNs). It is similar to the virtual SCSI functionality. The main difference is that the VIOS does not act as a SCSI emulator to its client partitions. It acts as a Fibre Channel pass-through for the Fibre Channel protocol I/O traffic through the Power Hypervisor.

- **Shared storage pool**

  A shared storage pool is a server-based storage virtualization that is clustered and is an extension of existing storage virtualization on the Virtual I/O Server.

  Shared storage pools simplify the aggregation and tasks of large numbers of disks across multiple VIOSs; it improves the utilization of the available storage.

After the physical volumes are allocated to a Virtual I/O Server in the shared storage pool environment, the physical volume management tasks, such as capacity management or an allocation of the volumes to a client partition, are performed by the Virtual I/O Server.

Figure 5-45 shows an abstract map of each storage implementation.

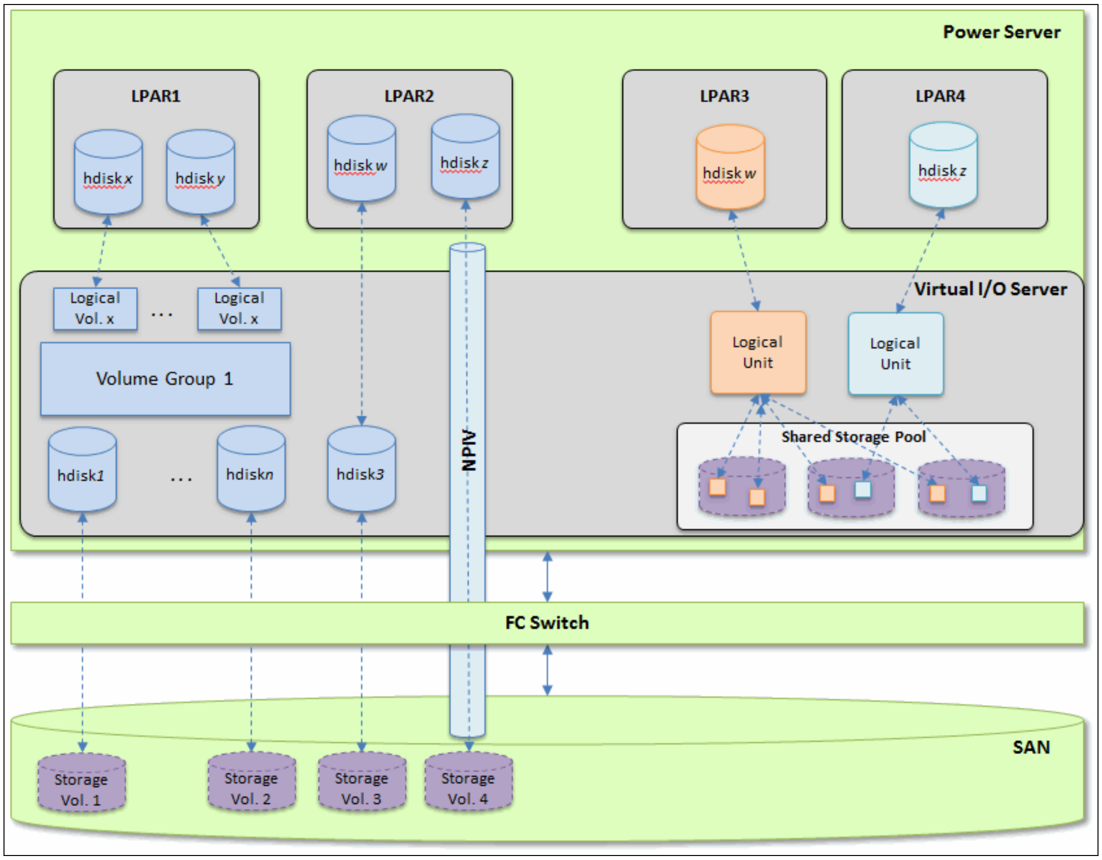

Figure 5-45 Virtual disks, NPIV, and shared storage pool supported by VMControl

For more information about storage virtualization on PowerVM, see the following Redbooks publication: *IBM PowerVM Virtualization Introduction and Configuration*, SG24-7940.

http://www.redbooks.ibm.com/abstracts/sg247940.html

5.7.6 Managing virtual servers with VMControl Express Edition

With IBM Systems Director VMControl Express Edition, you can edit, delete, and relocate virtual servers from one host to another. Also, it is possible to perform manual and automatic relocation.

Creating a virtual server

For IBM Systems Director VMControl in a Power Systems environment, a virtual server creation is tightly coupled with a virtual appliance creation or use. However, an empty virtual server can be created without a virtual appliance creation or input, and no workload is generated. The result is an empty virtual server into which the software image of the virtual appliance can be installed later.

The information requested is specific to the virtualization environment on which the virtual server is being created.

1. Go to VMControl and click the **Virtual Servers/Hosts** tab (Figure 5-46). You will see a list of your Power Systems servers and all the VMs that each server contains.

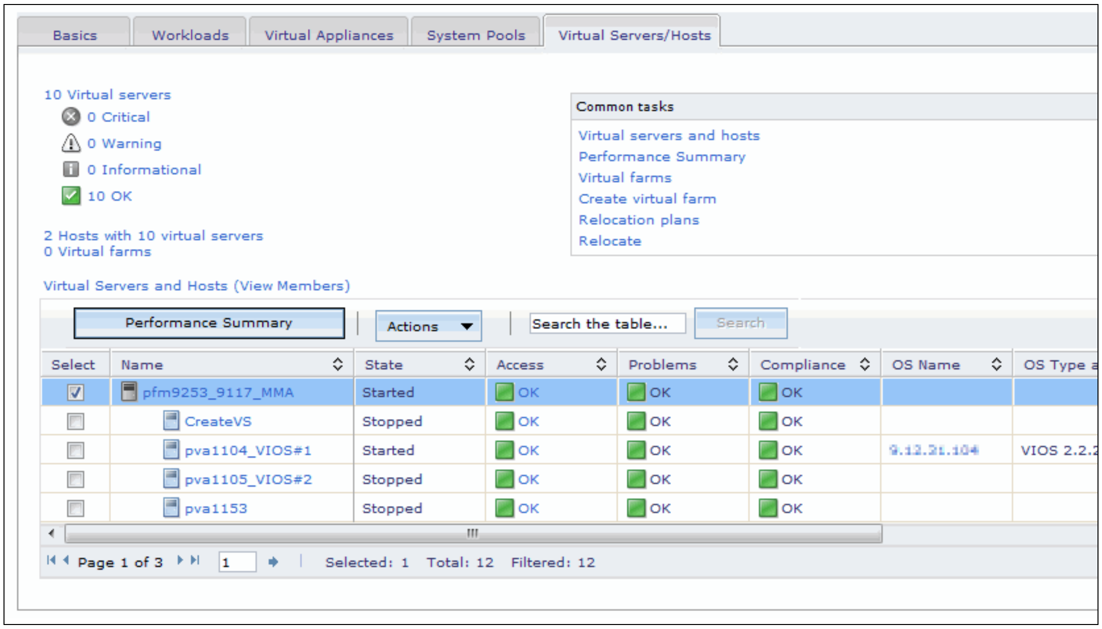

Select	Name	State	Access	Problems	Compliance	OS Name	OS Type a
☑	pfm9253_9117_MMA	Started	🟩 OK	🟩 OK	🟩 OK		
□	CreateVS	Stopped	🟩 OK	🟩 OK	🟩 OK		
□	pva1104_VIOS#1	Started	🟩 OK	🟩 OK	🟩 OK	9.12.21.104	VIOS 2.2.2
□	pva1105_VIOS#2	Stopped	🟩 OK	🟩 OK	🟩 OK		
□	pva1153	Stopped	🟩 OK	🟩 OK	🟩 OK		

Figure 5-46 VMControl: Virtual Servers/Hosts panel view

2. Right-click the Power Systems server where you want to create the VM, and select **System Configuration** → **Create Virtual Server** (Figure 5-47).

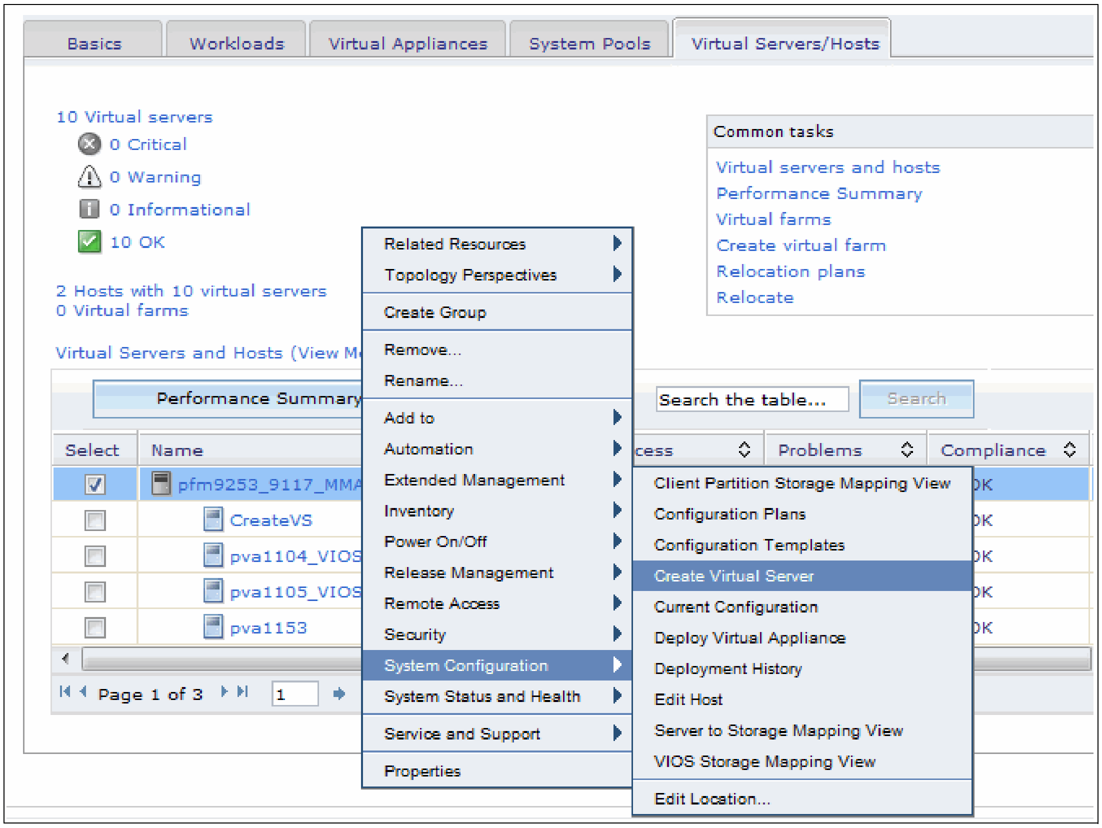

Figure 5-47 Creating a virtual server on your Power server

3. Type the name of your VM as shown in Figure 5-48.

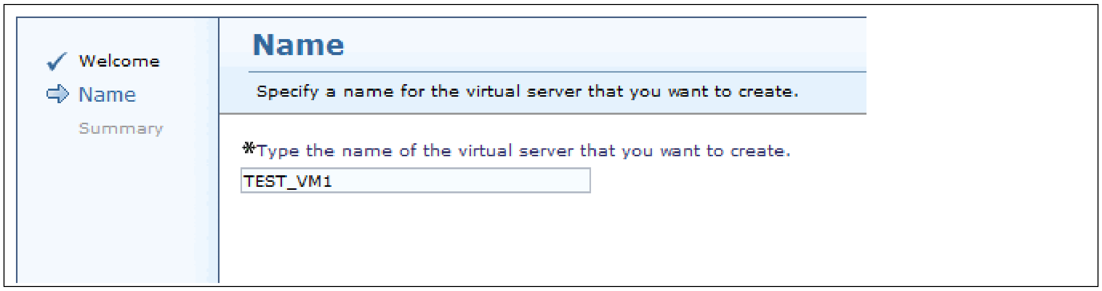

Name

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

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

TEST_VM1

Figure 5-48 Choose the name of the new VM

4. Select the operating system as shown in Figure 5-49.

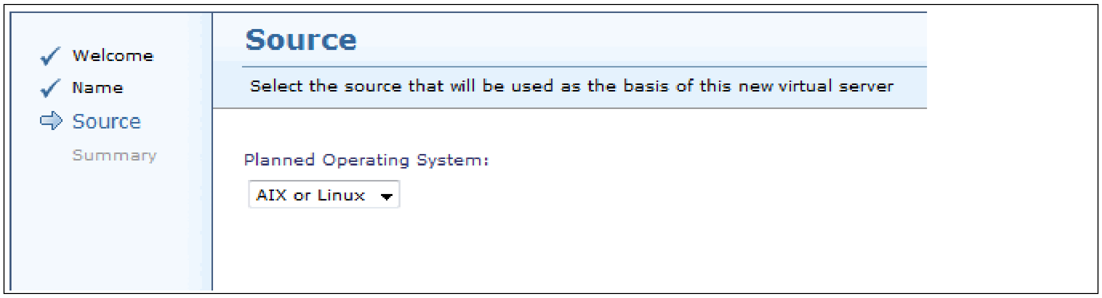

Source

Select the source that will be used as the basis of this new virtual server

Planned Operating System:

AIX or Linux ▾

Figure 5-49 Choose the OS that will run on the new VM

5. Configure the number of processors. This step allows you to assign your virtual processors or use dedicated processors for your virtual server (Figure 5-50).

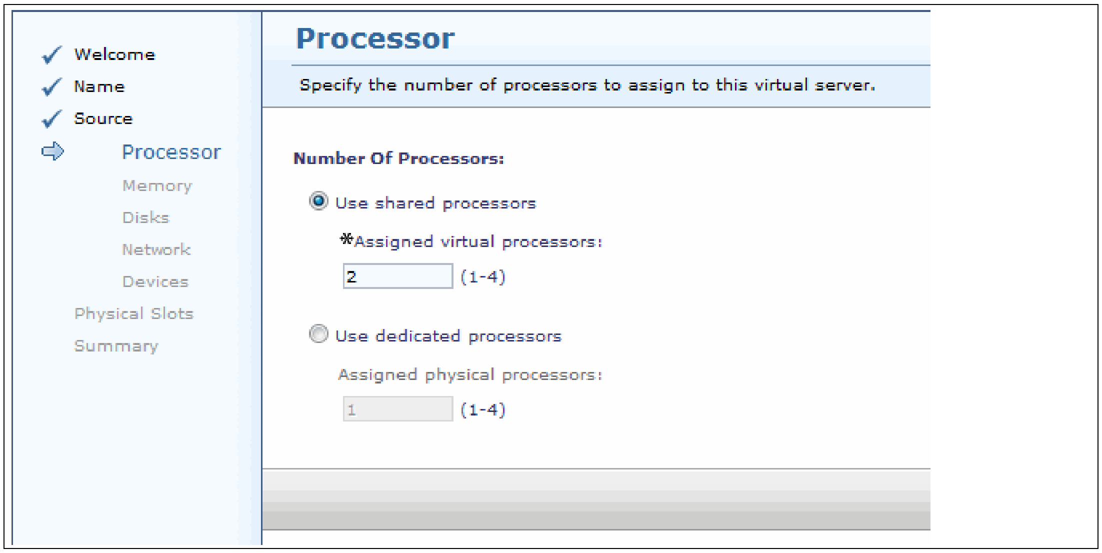

Processor

Specify the number of processors to assign to this virtual server.

Number Of Processors:

● Use shared processors

*Assigned virtual processors:

2 (1-4)

○ Use dedicated processors

Assigned physical processors:

1 (1-4)

Figure 5-50 Number of processor assignation

6. Configure the amount of memory. This step (Figure 5-51 on page 313) allows you to set up your memory size.

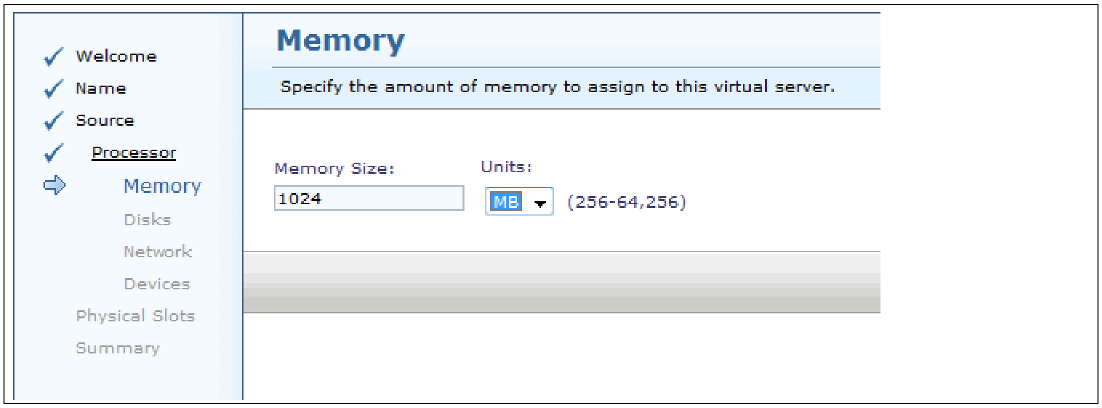

Memory

Specify the amount of memory to assign to this virtual server.

Memory Size: 1024

Units: MB (256-64,256)

Figure 5-51 Memory size configuration

7. Select your disk configuration (Figure 5-52). You can add an existing storage or create a new one. For this example, we will use an existing configuration.

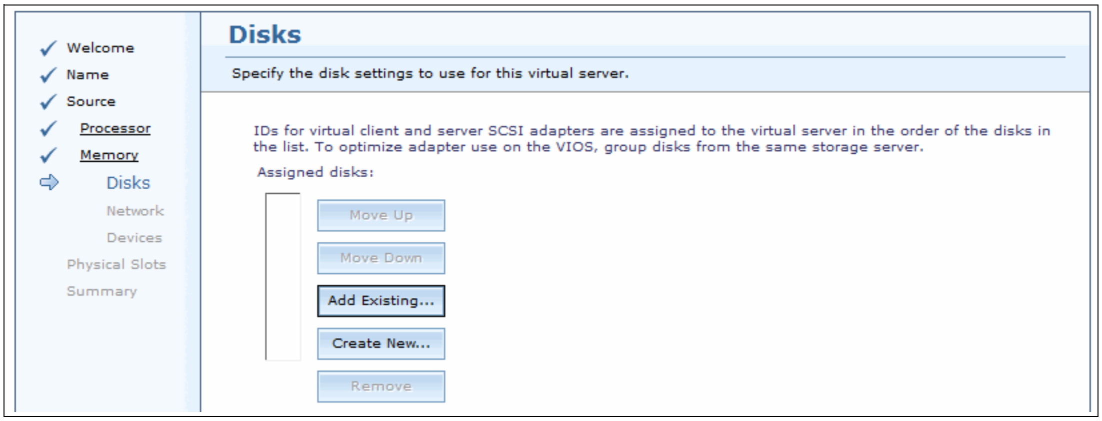

Disks

Specify the disk settings to use for this virtual server.

IDs for virtual client and server SCSI adapters are assigned to the virtual server in the order of the disks in the list. To optimize adapter use on the VIOS, group disks from the same storage server.

Assigned disks:

Figure 5-52 Disk configuration

8. In the Physical Volumes window (Figure 5-53 on page 314), select a hdisk*x* available to be used.

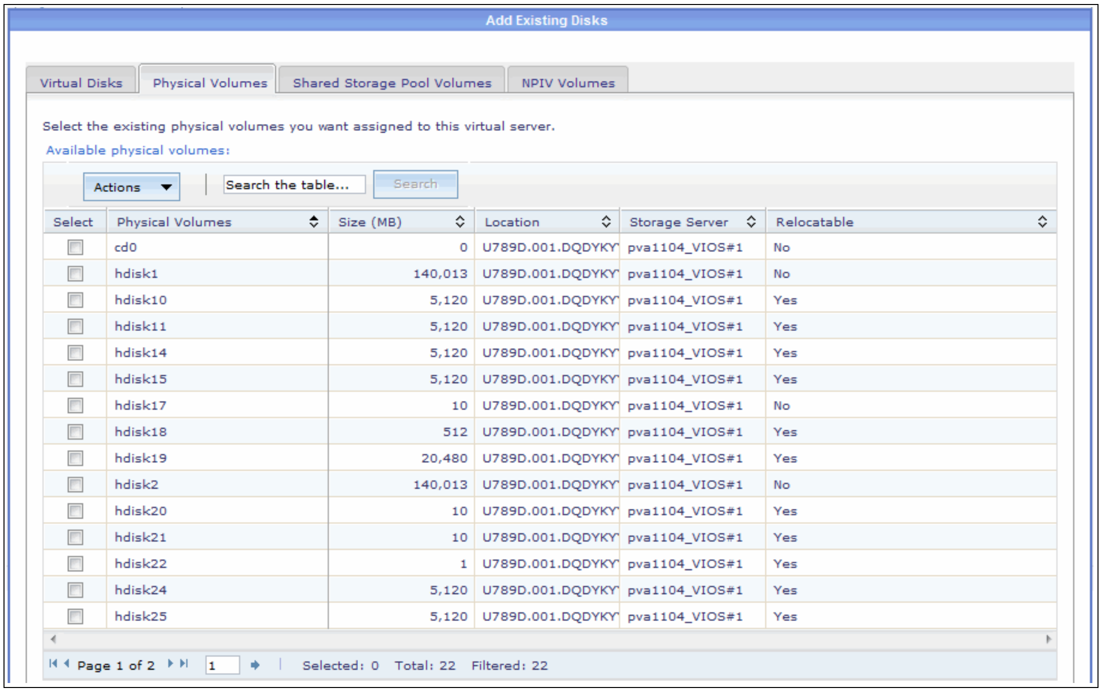

Select the existing physical volumes you want assigned to this virtual server.

Available physical volumes:

Select	Physical Volumes	Size (MB)	Location	Storage Server	Relocatable
□	cd0	0	U789D.001.DQDYKY	pva1104_VIOS#1	No
□	hdisk1	140,013	U789D.001.DQDYKY	pva1104_VIOS#1	No
□	hdisk10	5,120	U789D.001.DQDYKY	pva1104_VIOS#1	Yes
□	hdisk11	5,120	U789D.001.DQDYKY	pva1104_VIOS#1	Yes
□	hdisk14	5,120	U789D.001.DQDYKY	pva1104_VIOS#1	Yes
□	hdisk15	5,120	U789D.001.DQDYKY	pva1104_VIOS#1	Yes
□	hdisk17	10	U789D.001.DQDYKY	pva1104_VIOS#1	No
□	hdisk18	512	U789D.001.DQDYKY	pva1104_VIOS#1	Yes
□	hdisk19	20,480	U789D.001.DQDYKY	pva1104_VIOS#1	Yes
□	hdisk2	140,013	U789D.001.DQDYKY	pva1104_VIOS#1	No
□	hdisk20	10	U789D.001.DQDYKY	pva1104_VIOS#1	Yes
□	hdisk21	10	U789D.001.DQDYKY	pva1104_VIOS#1	Yes
□	hdisk22	1	U789D.001.DQDYKY	pva1104_VIOS#1	Yes
□	hdisk24	5,120	U789D.001.DQDYKY	pva1104_VIOS#1	Yes
□	hdisk25	5,120	U789D.001.DQDYKY	pva1104_VIOS#1	Yes

Page 1 of 2 | 1 | Selected: 0 Total: 22 Filtered: 22

Figure 5-53 Adding existing disks

9. After the disk configuration is completed, a summary is shown (Figure 5-54).

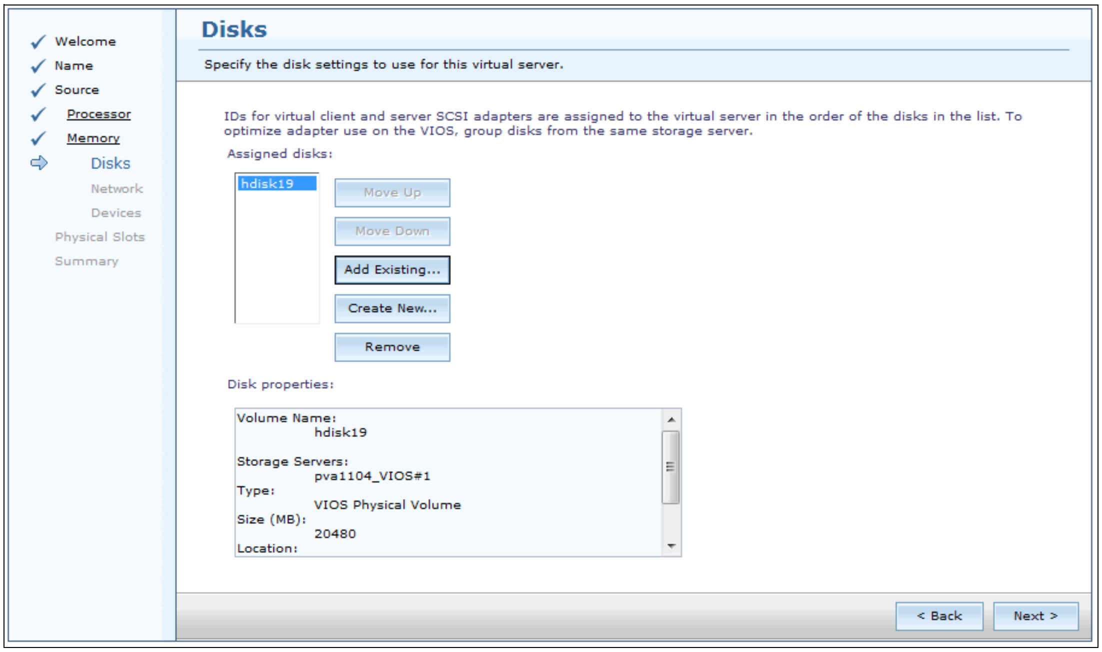

Disks

Specify the disk settings to use for this virtual server.

IDs for virtual client and server SCSI adapters are assigned to the virtual server in the order of the disks in the list. To optimize adapter use on the VIOS, group disks from the same storage server.

Assigned disks:

hdisk19

Disk properties:

```
Volume Name:
          hdisk19

Storage Servers:
          pva1104_VIOS#1
Type:
          VIOS Physical Volume
Size (MB):
          20480
Location:
```

Figure 5-54 Disk configuration summary

10. Select the Network interface that will be assigned to the virtual server (Figure 5-55 on page 315).

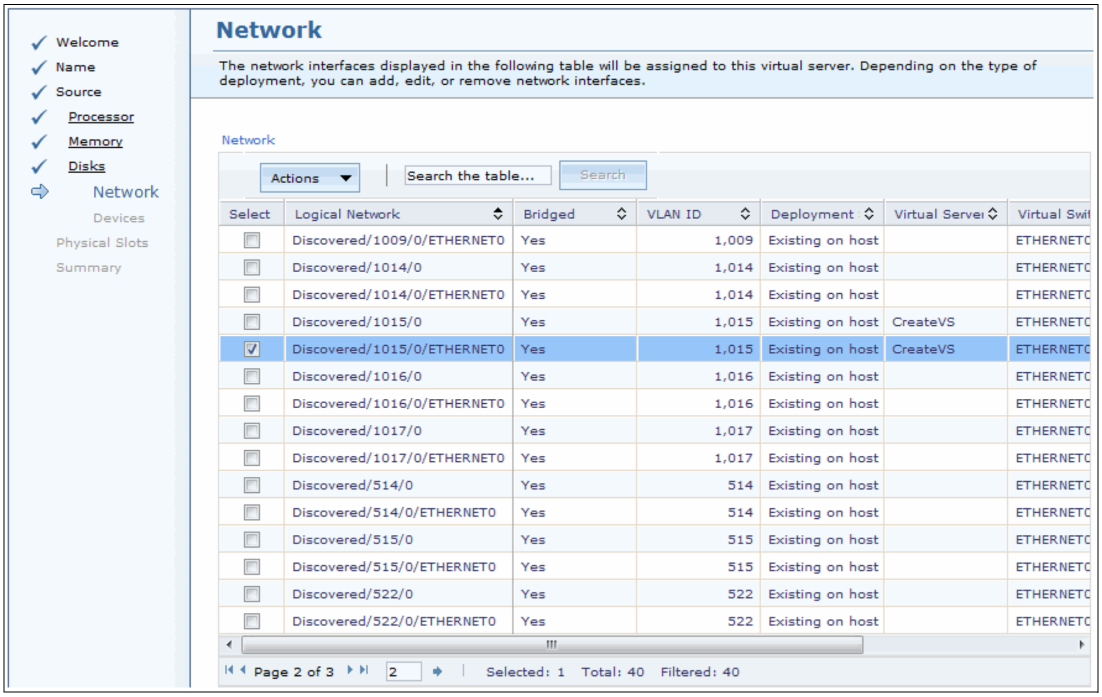

Network

The network interfaces displayed in the following table will be assigned to this virtual server. Depending on the type of deployment, you can add, edit, or remove network interfaces.

Network

Select	Logical Network	Bridged	VLAN ID	Deployment	Virtual Server	Virtual Swit
□	Discovered/1009/0/ETHERNET0	Yes	1,009	Existing on host		ETHERNETC
□	Discovered/1014/0	Yes	1,014	Existing on host		ETHERNETC
□	Discovered/1014/0/ETHERNET0	Yes	1,014	Existing on host		ETHERNETC
□	Discovered/1015/0	Yes	1,015	Existing on host	CreateVS	ETHERNETC
☑	Discovered/1015/0/ETHERNET0	Yes	1,015	Existing on host	CreateVS	ETHERNETC
□	Discovered/1016/0	Yes	1,016	Existing on host		ETHERNETC
□	Discovered/1016/0/ETHERNET0	Yes	1,016	Existing on host		ETHERNETC
□	Discovered/1017/0	Yes	1,017	Existing on host		ETHERNETC
□	Discovered/1017/0/ETHERNET0	Yes	1,017	Existing on host		ETHERNETC
□	Discovered/514/0	Yes	514	Existing on host		ETHERNETC
□	Discovered/514/0/ETHERNET0	Yes	514	Existing on host		ETHERNETC
□	Discovered/515/0	Yes	515	Existing on host		ETHERNETC
□	Discovered/515/0/ETHERNET0	Yes	515	Existing on host		ETHERNETC
□	Discovered/522/0	Yes	522	Existing on host		ETHERNETC
□	Discovered/522/0/ETHERNET0	Yes	522	Existing on host		ETHERNETC

Page 2 of 3 | 2 | Selected: 1 Total: 40 Filtered: 40

Figure 5-55 Network configuration

11.Select the additional devices to assign to the virtual server (Figure 5-56).

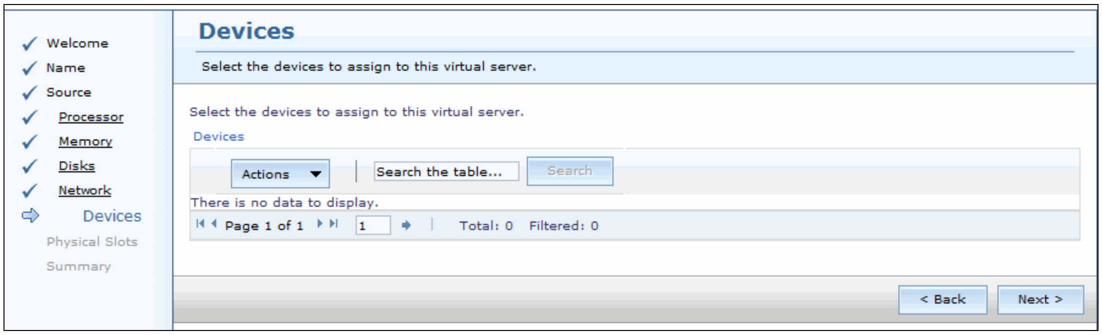

Devices

Select the devices to assign to this virtual server.

Select the devices to assign to this virtual server.

Devices

There is no data to display.

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

Figure 5-56 Additional devices configuration

12.Configure *Physical Slots* to assign to the virtual server (Figure 5-57 on page 316).

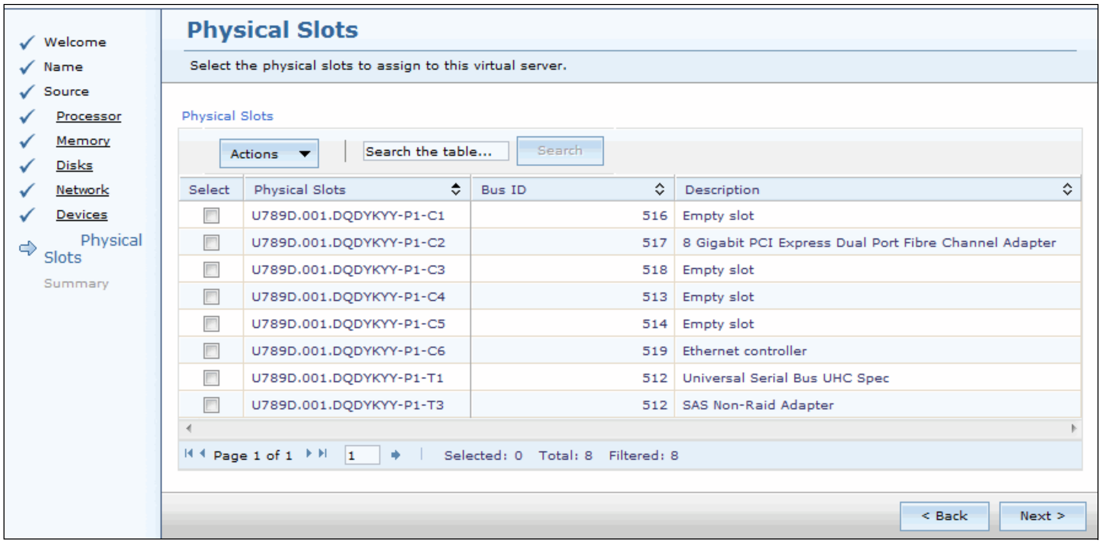

Physical Slots

Select the physical slots to assign to this virtual server.

Physical Slots

Select	Physical Slots	Bus ID	Description
□	U789D.001.DQDYKYY-P1-C1	516	Empty slot
□	U789D.001.DQDYKYY-P1-C2	517	8 Gigabit PCI Express Dual Port Fibre Channel Adapter
□	U789D.001.DQDYKYY-P1-C3	518	Empty slot
□	U789D.001.DQDYKYY-P1-C4	513	Empty slot
□	U789D.001.DQDYKYY-P1-C5	514	Empty slot
□	U789D.001.DQDYKYY-P1-C6	519	Ethernet controller
□	U789D.001.DQDYKYY-P1-T1	512	Universal Serial Bus UHC Spec
□	U789D.001.DQDYKYY-P1-T3	512	SAS Non-Raid Adapter

Page 1 of 1 | 1 | Selected: 0 Total: 8 Filtered: 8

Figure 5-57 Physical Slots configuration page

13. Once the VM creation wizard is complete, submit the job and check results. After it finishes, your new VM will be created under your physical server. See Figure 5-58.

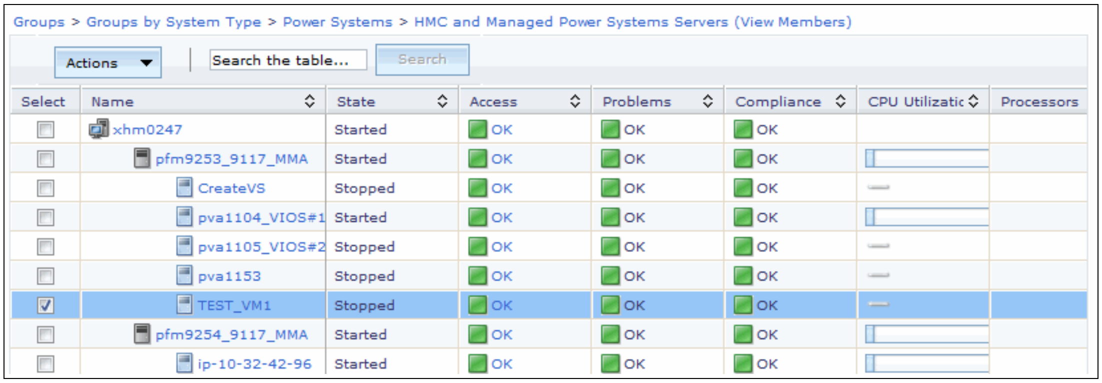

Groups > Groups by System Type > Power Systems > HMC and Managed Power Systems Servers (View Members)

Select	Name	State	Access	Problems	Compliance	CPU Utilizatic	Processors
□	xhm0247	Started	🟩 OK	🟩 OK	🟩 OK		
□	pfm9253_9117_MMA	Started	🟩 OK	🟩 OK	🟩 OK		
□	CreateVS	Stopped	🟩 OK	🟩 OK	🟩 OK		
□	pva1104_VIOS#1	Started	🟩 OK	🟩 OK	🟩 OK		
□	pva1105_VIOS#2	Stopped	🟩 OK	🟩 OK	🟩 OK		
□	pva1153	Stopped	🟩 OK	🟩 OK	🟩 OK		
☑	TEST_VM1	Stopped	🟩 OK	🟩 OK	🟩 OK		
□	pfm9254_9117_MMA	Started	🟩 OK	🟩 OK	🟩 OK		
□	ip-10-32-42-96	Started	🟩 OK	🟩 OK	🟩 OK		

Figure 5-58 New VM created

Editing a virtual server

After you use IBM Systems Director VMControl to create or deploy a virtual server, you can use IBM Systems Director to edit the virtual server. VMControl can be used to edit memory, processor, virtual or physical I/O slots, virtual or physical disks assigned to the virtual server, virtual Ethernet adapters, and optical devices. Most attributes can be changed only on a powered-off virtual server. In a running virtual server, only memory and processor allocation can be changed.

A running virtual server allows the changes to be either applied to the current running operating system and saved to the LPAR profile in the HMC. Or, it can be only saved to the LPAR profile in the HMC when it is activated after the next LPAR power cycle (shutdown/activate). Below are the steps to edit host resources:

1. In the IBM Systems Director navigation area, click **Resource Explorer** to locate the host that you want to edit.
2. On the Chassis Manager page, click **Resource Explorer** under General Actions.
3. Select the host, click **Actions** from the menu bar, and select **System Configuration** → **Edit Virtual Server**. The Edit Host Resources window opens.

4. After completing the changes you want to request, click **OK**.
5. In the scheduler window, click **OK** to run the task immediately. You also can schedule to run this task later.

Deleting a virtual server

Removing a virtual server causes the LPAR to be removed permanently. To permanently remove a virtual server, it needs to be powered off first and then removed.

Complete the following steps to delete a virtual server from its associated host:

1. In the IBM Systems Director navigation pane, click **Resource Explorer** to locate the virtual server.
2. Select the virtual server and click **Actions** → **Permanently Delete Virtual Server** from the menu bar, as shown in Figure 5-59.

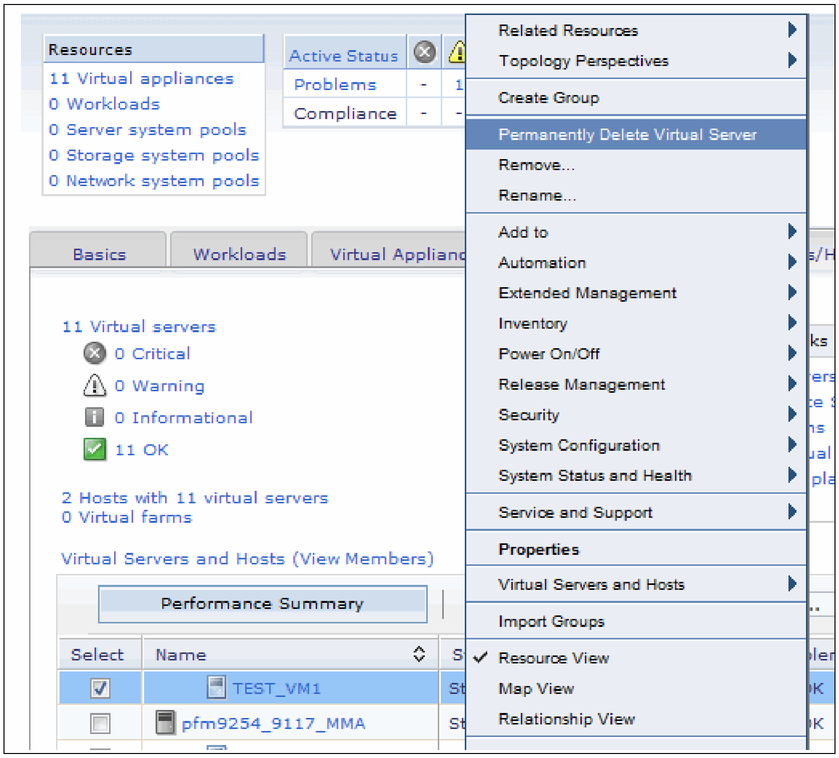

Figure 5-59 Permanently Delete Virtual Server task

3. In the scheduler window, click **OK** to run the task immediately. You also can schedule to run this task later.

Relocation

IBM Systems Director VMControl allows you, depending on the virtualization environment that you have and on the plug-ins you choose to install, to use virtual farms to enable either static relocation or live relocation. With *static relocation*, if the virtual server is powered on, the relocation operation powers off the virtual server at the beginning of the relocation process and powers the virtual server on when the relocation is complete. With *live relocation*, if the virtual server is powered on, the relocation occurs without powering the server off.

Server system pools provide live relocation. With server system pools, you can choose from three options for relocation:

- You can manually relocate virtual servers at any time.
- You can activate a resilience policy on your workload to relocate virtual servers automatically to prevent predicted hardware failures from affecting the availability of the workload.
- You can create an automation plan to relocate the virtual servers when certain events occur.

Best Practice: Use the HMC to validate if an LPAR can be relocated.

Relocating virtual servers manually

To relocate a server, a virtual farm must be created. Perform the following steps:

1. Go to VMControl **Virtual Servers/Hosts** view, and click **Create virtual farm**. See Figure 5-60.

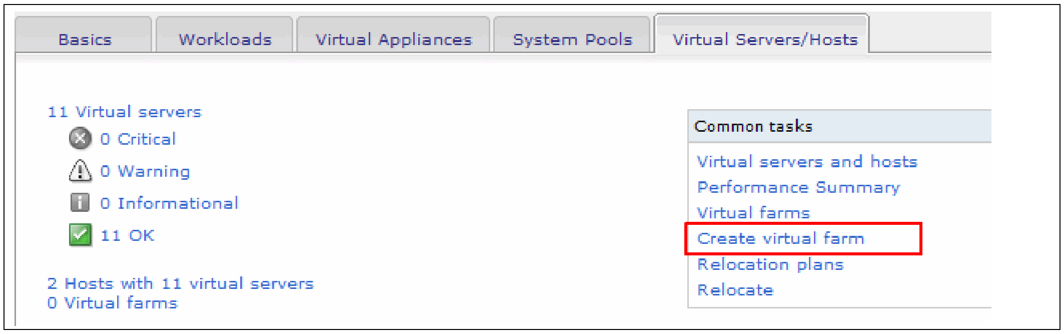

Figure 5-60 VMControl: Virtual Servers/Hosts tab

Note: To delete an existing virtual farm, first remove all the associated hosts.

2. Complete the Virtual Farm wizard. Ensure that you have enabled **Live relocation**, as shown in Figure 5-61.

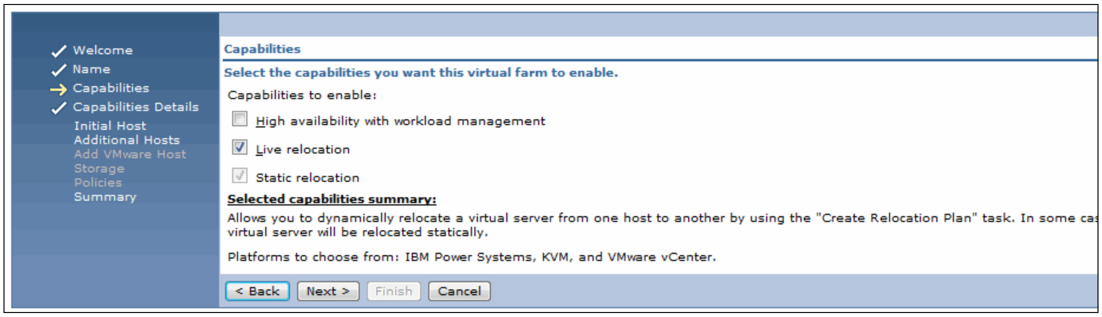

Capabilities

Select the capabilities you want this virtual farm to enable.

Capabilities to enable:

☐ High availability with workload management

☑ Live relocation

☑ Static relocation

Selected capabilities summary:

Allows you to dynamically relocate a virtual server from one host to another by using the "Create Relocation Plan" task. In some ca
virtual server will be relocated statically.

Platforms to choose from: IBM Power Systems, KVM, and VMware vCenter.

Figure 5-61 Virtual Farm wizard: Capabilities configuration

3. Select **Initial Host**, then **Additional Hosts** (Figure 5-62 on page 319).

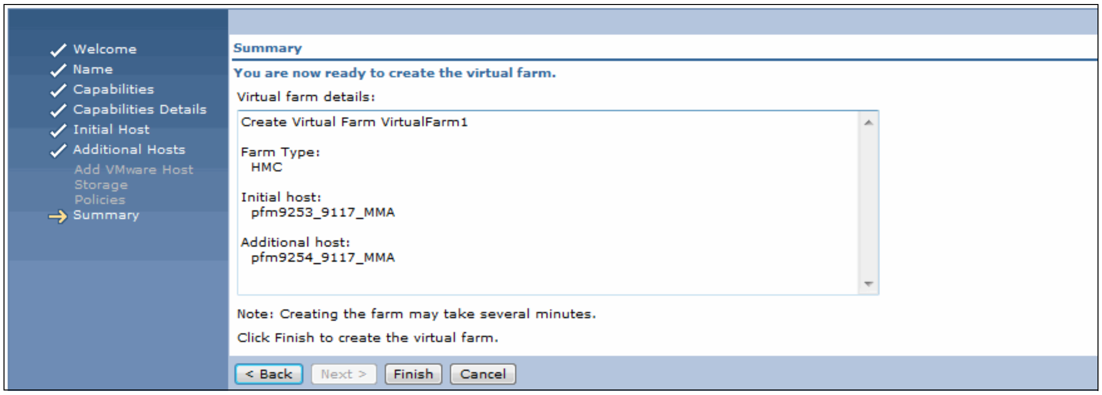

Summary

You are now ready to create the virtual farm.

Virtual farm details:

Create Virtual Farm VirtualFarm1

Farm Type:
HMC

Initial host:
pfm9253_9117_MMA

Additional host:
pfm9254_9117_MMA

Note: Creating the farm may take several minutes.

Click Finish to create the virtual farm.

Figure 5-62 Virtual Farm wizard summary

4. After the virtual farm is created, relocation is available. Go to the **Virtual Server and Host** view and right-click a virtual server. Click **Availability** → **Relocate** as shown in Figure 5-63.

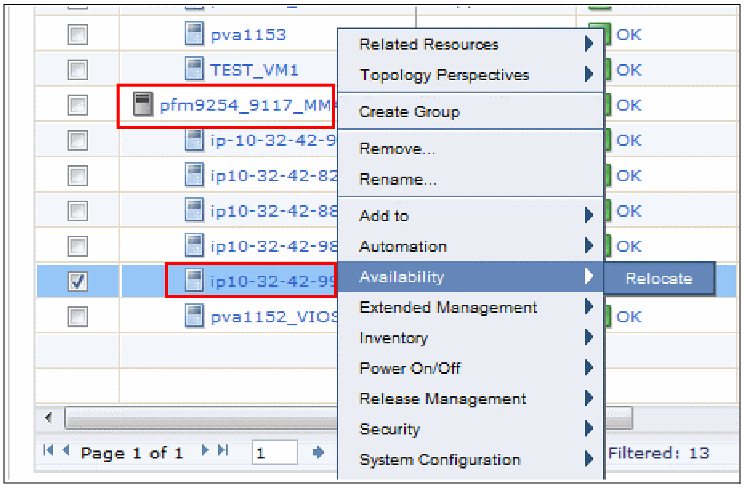

Figure 5-63 Relocate a virtual server

5. Complete the wizard steps. Review the summary (Figure 5-64 on page 320) and click **Finish**.

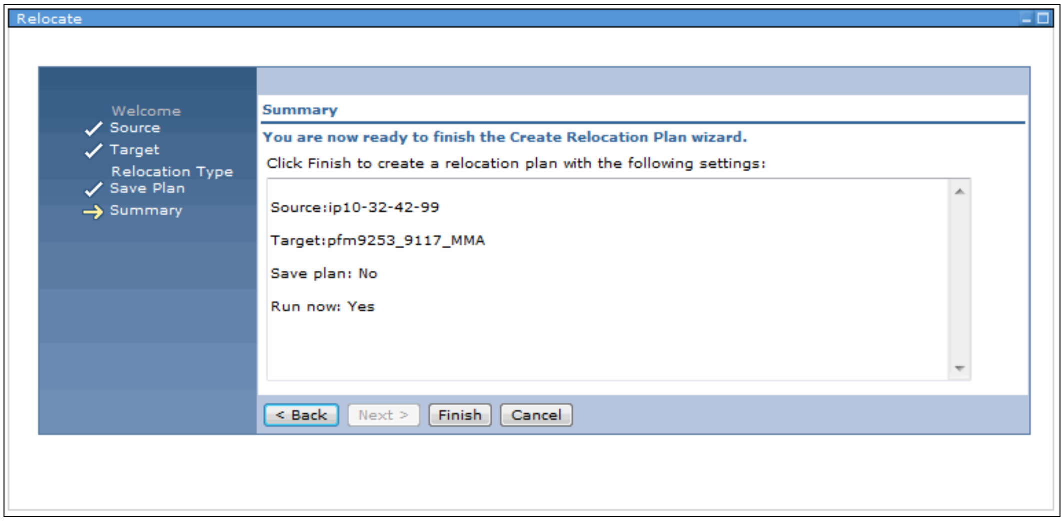

Summary

You are now ready to finish the Create Relocation Plan wizard.

Click Finish to create a relocation plan with the following settings:

Source:ip10-32-42-99

Target:pfm9253_9117_MMA

Save plan: No

Run now: Yes

Figure 5-64 Relocation wizard summary

6. The relocation process starts, as shown in Figure 5-65.

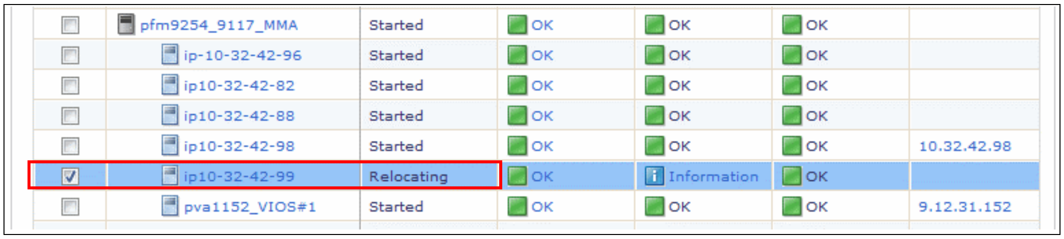

□	pfm9254_9117_MMA	Started	🟩 OK	🟩 OK	🟩 OK	
□	ip-10-32-42-96	Started	🟩 OK	🟩 OK	🟩 OK	
□	ip10-32-42-82	Started	🟩 OK	🟩 OK	🟩 OK	
□	ip10-32-42-88	Started	🟩 OK	🟩 OK	🟩 OK	
□	ip10-32-42-98	Started	🟩 OK	🟩 OK	🟩 OK	10.32.42.98
☑	ip10-32-42-99	Relocating	🟩 OK	ℹ Information	🟩 OK	
□	pva1152_VIOS#1	Started	🟩 OK	🟩 OK	🟩 OK	9.12.31.152

Figure 5-65 Relocating a virtual server

7. Finally, the selected virtual server was relocated to the target host. See Figure 5-66.

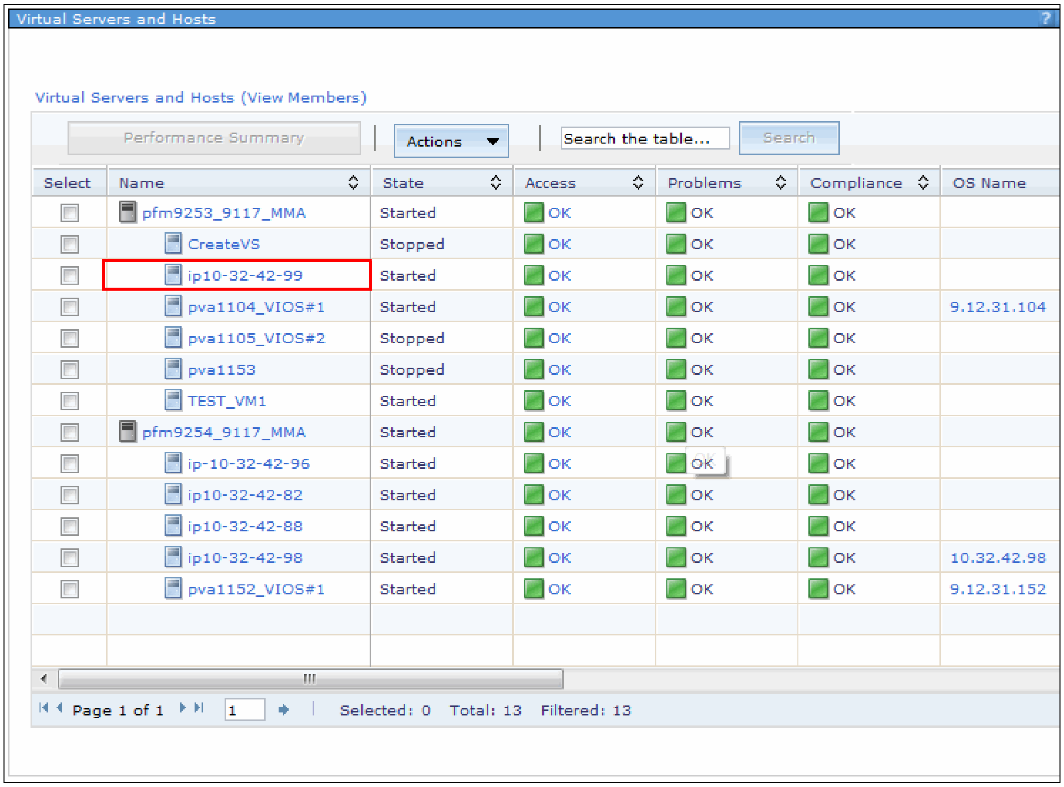

Virtual Servers and Hosts

Virtual Servers and Hosts (View Members)

Performance Summary | Actions ▼ | Search the table... | Search

Select	Name	State	Access	Problems	Compliance	OS Name
□	pfm9253_9117_MMA	Started	🟩 OK	🟩 OK	🟩 OK	
□	CreateVS	Stopped	🟩 OK	🟩 OK	🟩 OK	
□	ip10-32-42-99	Started	🟩 OK	🟩 OK	🟩 OK	
□	pva1104_VIOS#1	Started	🟩 OK	🟩 OK	🟩 OK	9.12.31.104
□	pva1105_VIOS#2	Stopped	🟩 OK	🟩 OK	🟩 OK	
□	pva1153	Stopped	🟩 OK	🟩 OK	🟩 OK	
□	TEST_VM1	Started	🟩 OK	🟩 OK	🟩 OK	
□	pfm9254_9117_MMA	Started	🟩 OK	🟩 OK	🟩 OK	
□	ip-10-32-42-96	Started	🟩 OK	🟩 OK	🟩 OK	
□	ip10-32-42-82	Started	🟩 OK	🟩 OK	🟩 OK	
□	ip10-32-42-88	Started	🟩 OK	🟩 OK	🟩 OK	
□	ip10-32-42-98	Started	🟩 OK	🟩 OK	🟩 OK	10.32.42.98
□	pva1152_VIOS#1	Started	🟩 OK	🟩 OK	🟩 OK	9.12.31.152

Page 1 of 1 | 1 | Selected: 0 Total: 13 Filtered: 13

Figure 5-66 Virtual Servers and Hosts view

Relocating automatically a virtual server that is based on events

IBM Systems Director VMControl Express Edition allows you to create an automated plan to relocate automatically a virtual server that is based on an event, for example when a Predictive Failure event is triggered. Follow the steps below to create an automation plan that can manage automatic relocation that is based on events.

1. Ensure that a virtual farm is created. This virtual farm must contain the source host and the destination host where the virtual server will be moved.

2. Create a custom task using the command-line interface. Use the **`mkrelocatetask`** command to create it on your IBM Systems Director server:

   ```
   smcli mkrelocatetask [-v] {-d destination} {-s source} [-A attribute_list]
   task_name
   ```

   Example: Move LPAR: ip10-32-42-86 to pfm9253_9117_MMA

   ```
   smcli mkrelocatetask -d pfm9253_9117_MMA -s ip10-32-42-86 -A Policy=Live -A
   RunMode=Save Relocate_VS
   ```

   For more information about this command, see the following site:

   `http://pic.dhe.ibm.com/infocenter/director/pubs/index.jsp?topic=%2Fcom.ibm.director.vim.helps.doc%2Ffqm0_r_cli_mkrelocatetask.html`

3. Create your event filter to customize what events will trigger the relocation task. Figure 5-67 shows an example.

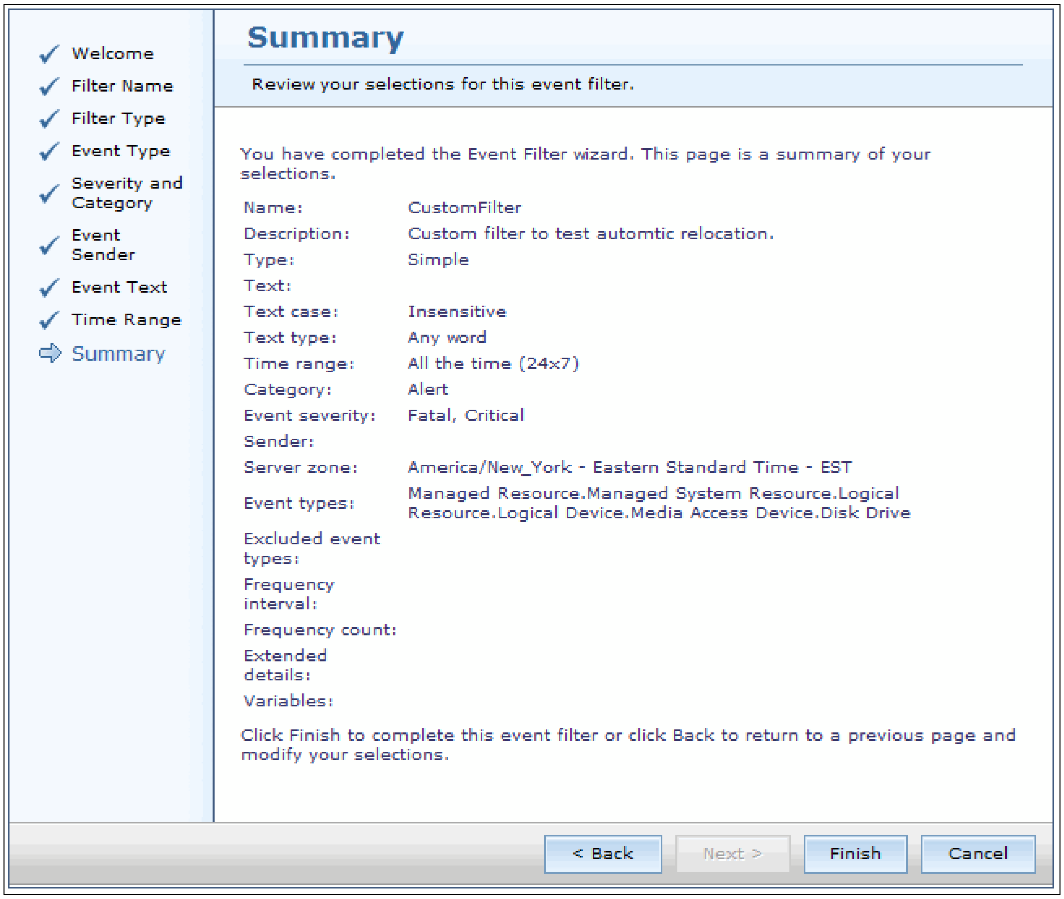

Summary

Review your selections for this event filter.

You have completed the Event Filter wizard. This page is a summary of your selections.

Name:	CustomFilter
Description:	Custom filter to test automtic relocation.
Type:	Simple
Text:	
Text case:	Insensitive
Text type:	Any word
Time range:	All the time (24x7)
Category:	Alert
Event severity:	Fatal, Critical
Sender:	
Server zone:	America/New_York - Eastern Standard Time - EST
Event types:	Managed Resource.Managed System Resource.Logical Resource.Logical Device.Media Access Device.Disk Drive
Excluded event types:	
Frequency interval:	
Frequency count:	
Extended details:	
Variables:	

Click Finish to complete this event filter or click Back to return to a previous page and modify your selections.

Figure 5-67 Customized event filter

Best Practice: After creating an Event action, configure it to save history.

4. Create an automation plan that uses the event filter and the relocation task that were created previously. See an example in Figure 5-68 on page 323.

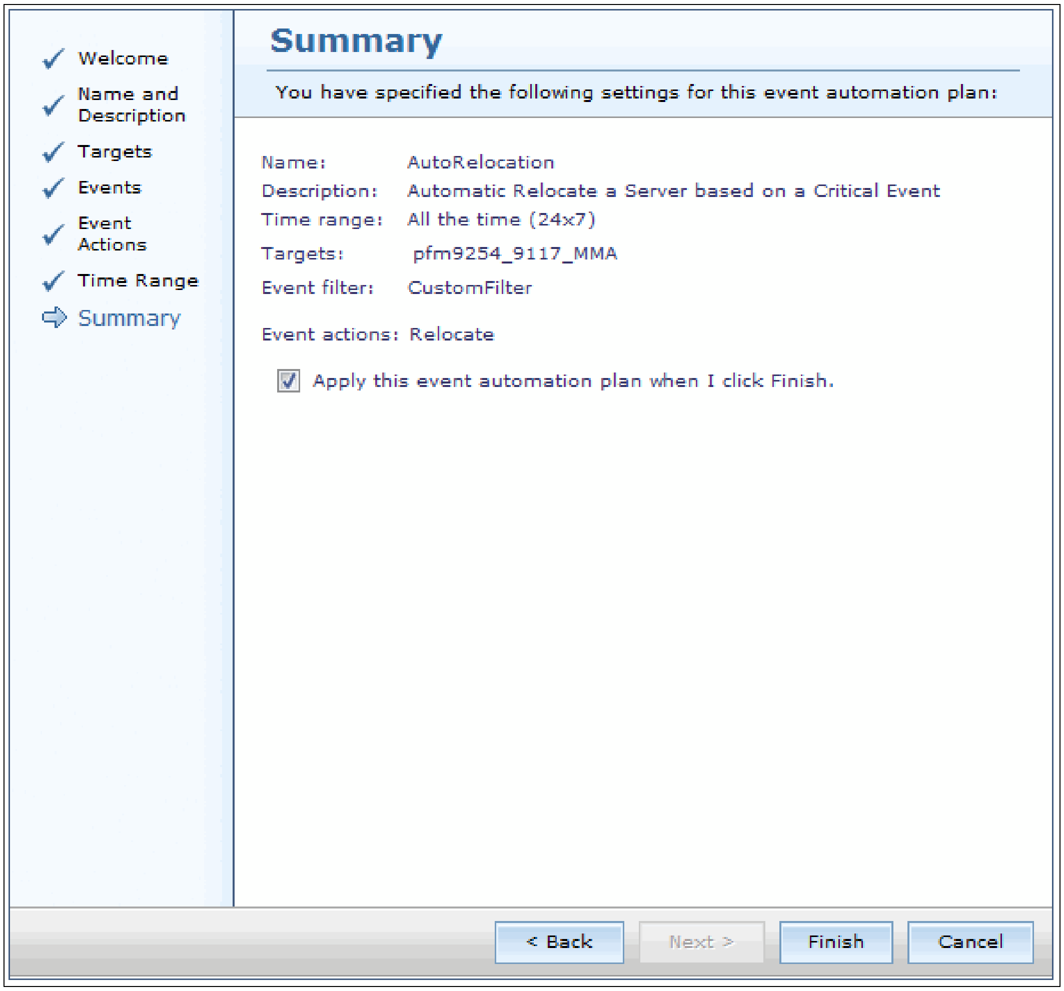

Summary

You have specified the following settings for this event automation plan:

Name: AutoRelocation
Description: Automatic Relocate a Server based on a Critical Event
Time range: All the time (24x7)
Targets: pfm9254_9117_MMA
Event filter: CustomFilter

Event actions: Relocate

☑ Apply this event automation plan when I click Finish.

Figure 5-68 New event automation plan

5. The automation plan is ready. For testing purposes, an event can be generated manually by running the **genevent** command from the command line:

```
smcli genevent /text:"Hard drive PFA event" /compcat:"Managed Resource.Managed
System Resource.Logical Resource.Logical Device.Media Access D
ice" /comptype:"Disk Drive" /mode:ALERT /sev:0 /MEID:37388
```

6. Check the event on the Event Log view. See Figure 5-69.

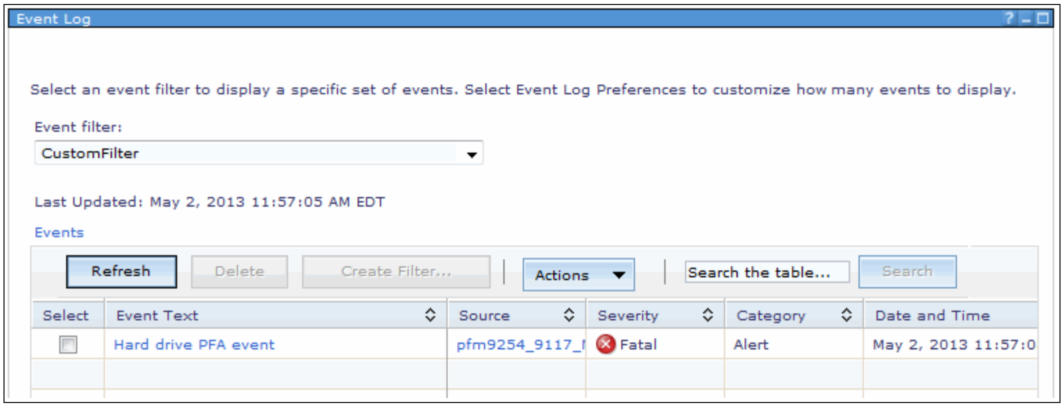

Select an event filter to display a specific set of events. Select Event Log Preferences to customize how many events to display.

Event filter:
CustomFilter

Last Updated: May 2, 2013 11:57:05 AM EDT

Events

Select	Event Text	Source	Severity	Category	Date and Time
□	Hard drive PFA event	pfm9254_9117_I	Fatal	Alert	May 2, 2013 11:57:0

Figure 5-69 Event Log filter

7. After the event is received, the relocating task is triggered automatically. From the Event Actions view, check if the relocation task was launched, as shown in Figure 5-70.

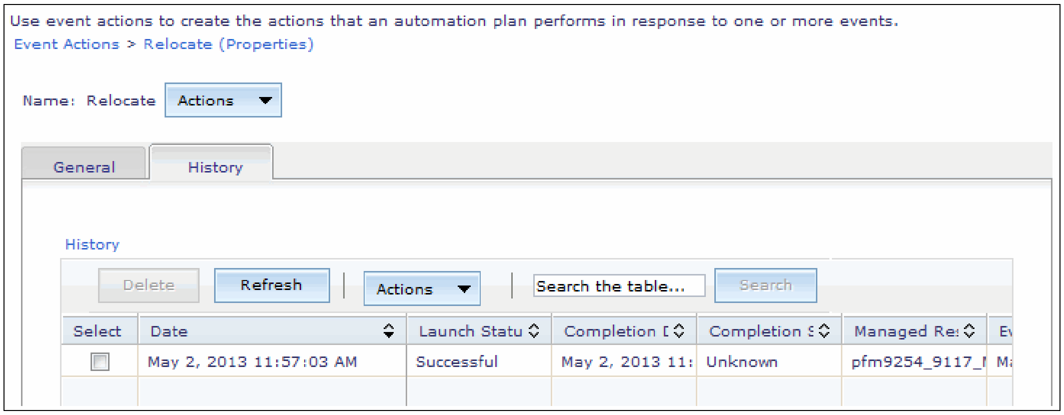

Use event actions to create the actions that an automation plan performs in response to one or more events.

Event Actions > Relocate (Properties)

Name: Relocate

History

Select	Date	Launch Statu	Completion D	Completion S	Managed Re:	Ev
□	May 2, 2013 11:57:03 AM	Successful	May 2, 2013 11:	Unknown	pfm9254_9117_I	Ma

Figure 5-70 Event Actions history

8. The status of the virtual server changes to *Relocating*. See Figure 5-71.

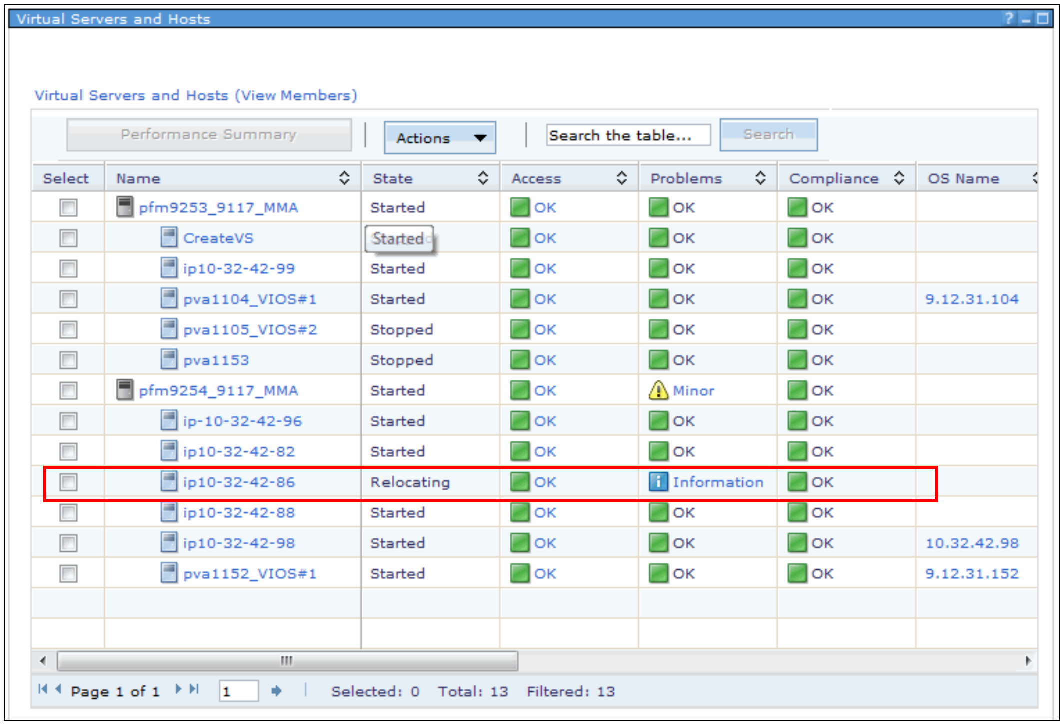

Virtual Servers and Hosts

Virtual Servers and Hosts (View Members)

Select	Name	State	Access	Problems	Compliance	OS Name
□	pfm9253_9117_MMA	Started	OK	OK	OK	
□	CreateVS	Started	OK	OK	OK	
□	ip10-32-42-99	Started	OK	OK	OK	
□	pva1104_VIOS#1	Started	OK	OK	OK	9.12.31.104
□	pva1105_VIOS#2	Stopped	OK	OK	OK	
□	pva1153	Stopped	OK	OK	OK	
□	pfm9254_9117_MMA	Started	OK	⚠ Minor	OK	
□	ip-10-32-42-96	Started	OK	OK	OK	
□	ip10-32-42-82	Started	OK	OK	OK	
□	ip10-32-42-86	Relocating	OK	ℹ Information	OK	
□	ip10-32-42-88	Started	OK	OK	OK	
□	ip10-32-42-98	Started	OK	OK	OK	10.32.42.98
□	pva1152_VIOS#1	Started	OK	OK	OK	9.12.31.152

Page 1 of 1 | 1 | Selected: 0 Total: 13 Filtered: 13

Figure 5-71 Relocation process started automatically

9. Finally, the virtual server has been relocated, as shown in Figure 5-72.

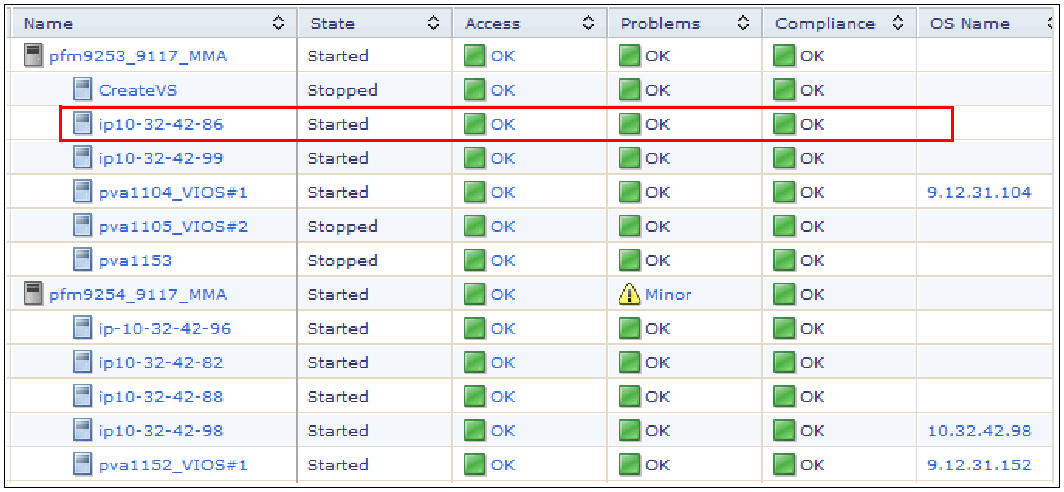

Name	State	Access	Problems	Compliance	OS Name
pfm9253_9117_MMA	Started	OK	OK	OK	
CreateVS	Stopped	OK	OK	OK	
ip10-32-42-86	Started	OK	OK	OK	
ip10-32-42-99	Started	OK	OK	OK	
pva1104_VIOS#1	Started	OK	OK	OK	9.12.31.104
pva1105_VIOS#2	Stopped	OK	OK	OK	
pva1153	Stopped	OK	OK	OK	
pfm9254_9117_MMA	Started	OK	⚠ Minor	OK	
ip-10-32-42-96	Started	OK	OK	OK	
ip10-32-42-82	Started	OK	OK	OK	
ip10-32-42-88	Started	OK	OK	OK	
ip10-32-42-98	Started	OK	OK	OK	10.32.42.98
pva1152_VIOS#1	Started	OK	OK	OK	9.12.31.152

Figure 5-72 The virtual server has been relocated to the target host

For more information about relocating virtual servers, see the following site:

http://pic.dhe.ibm.com/infocenter/director/pubs/index.jsp?topic=%2Fcom.ibm.director.vim.helps.doc%2Ffsd0_vim_t_relocating_vs.html

5.7.7 Managing a virtual appliance with VMControl Standard Edition

There are two ways to manage virtual appliances in AIX: NIM image repositories, and VIOS image repositories. You can set up your AIX environment with either type or with both types.

Creating and discovering NIM image repositories for AIX

Before working with virtual appliances that contain AIX mksysb and lpp_source images, you must ensure that the NIM repository has been properly configured and the NIM subagent installed.

Verify that the following prerequisites are met:

- Ensure that the NIM OS has been discovered, accessed, and inventoried by your IBM Systems Director server, and ensure that the IBM Systems Director 6.3 Common Agent has been installed.
- Ensure that the NIM server name is resolved.

Execute the following steps to discover your virtual appliances:

1. From the VMControl summary view, go to the **Basics** page and click **Discover virtual appliances** (Figure 5-73) to discover your repositories and virtual appliances.

   Virtual appliances already present in your repositories that have been imported or captured using VMControl are detected by VMControl. Additional virtual appliances can be added to your repositories by using the Capture and Import tasks in VMControl. See Figure 5-73.

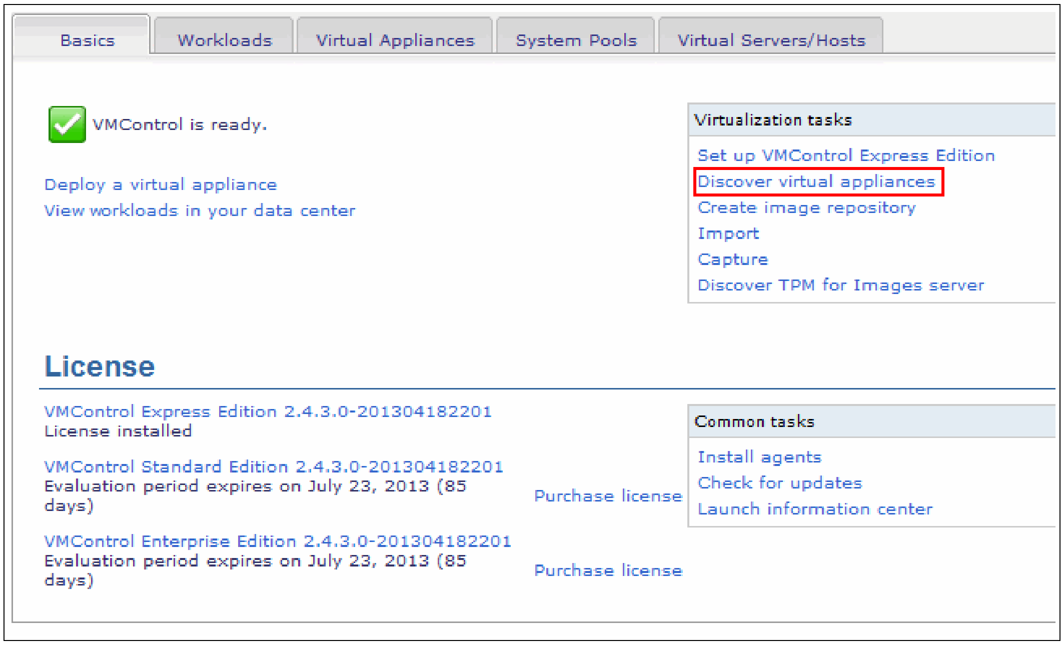

VMControl is ready.

Deploy a virtual appliance
View workloads in your data center

Virtualization tasks

Set up VMControl Express Edition
Discover virtual appliances
Create image repository
Import
Capture
Discover TPM for Images server

License

VMControl Express Edition 2.4.3.0-201304182201
License installed

VMControl Standard Edition 2.4.3.0-201304182201
Evaluation period expires on July 23, 2013 (85 days) Purchase license

VMControl Enterprise Edition 2.4.3.0-201304182201
Evaluation period expires on July 23, 2013 (85 days) Purchase license

Common tasks

Install agents
Check for updates
Launch information center

Figure 5-73 VMControl Basics page

2. Select the NIM server and run the job (Figure 5-74).

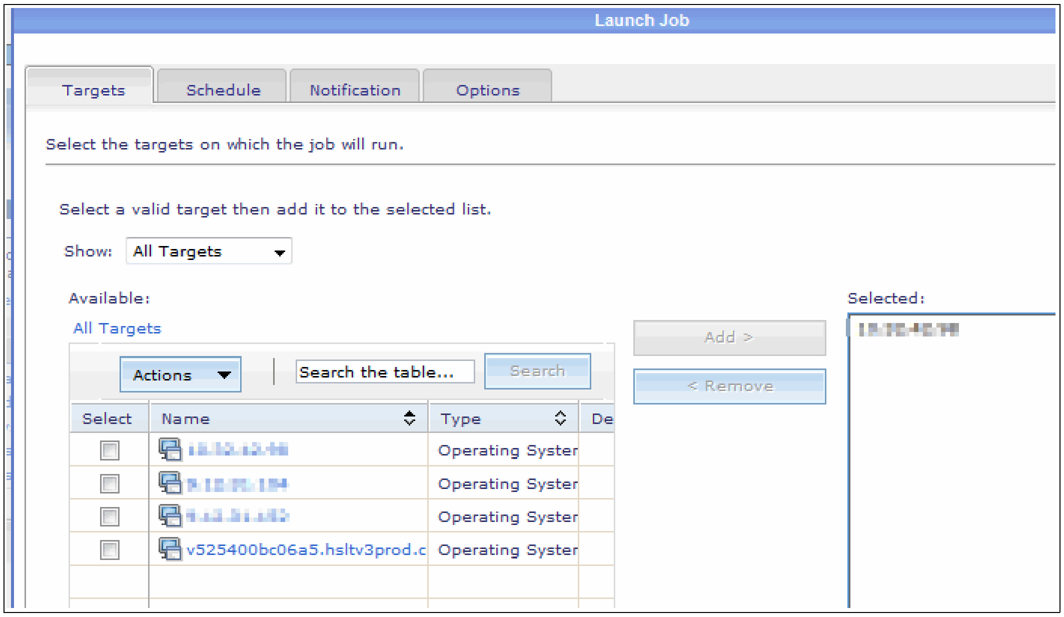

Select the targets on which the job will run.

Select a valid target then add it to the selected list.

Show: All Targets

Available:

All Targets

Search the table...

Selected:

[redacted]

Select	Name	Type	De
□	[redacted]	Operating Syster	
□	[redacted]	Operating Syster	
□	[redacted]	Operating Syster	
□	v525400bc06a5.hsltv3prod.c	Operating Syster	

Figure 5-74 Discover virtual appliances

3. After the job is completed, go back to the Basics page. The virtual appliance that was discovered is displayed under the Resources box, as shown in Figure 5-75.

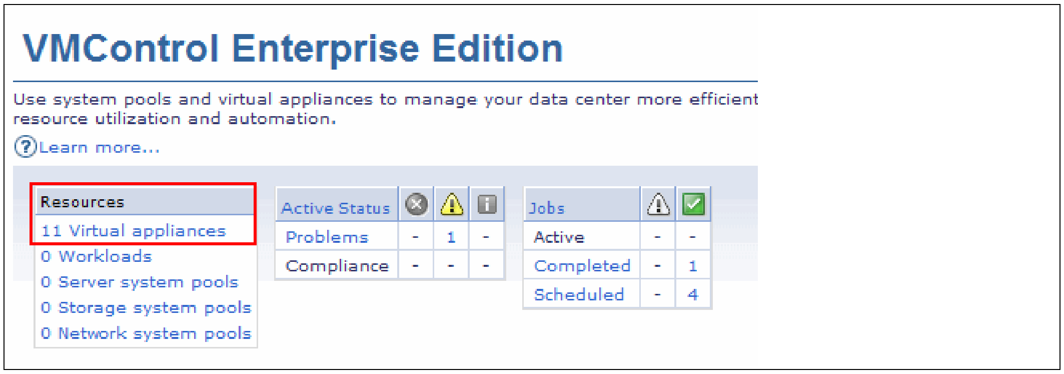

VMControl Enterprise Edition

Use system pools and virtual appliances to manage your data center more efficient resource utilization and automation.

?Learn more...

Resources
11 Virtual appliances
0 Workloads
0 Server system pools
0 Storage system pools
0 Network system pools

Active Status	✗	⚠	i
Problems	-	1	-
Compliance	-	-	-

Jobs	⚠	✓
Active	-	-
Completed	-	1
Scheduled	-	4

Figure 5-75 VMControl summary page

Creating and discovering VIOS image repositories on Power Systems

Before setting up a VIOS image repository for virtual appliances containing AIX, IBM i, or Linux on Power Systems raw images, ensure that all prerequisites are met:

- A SAN storage pool has been configured and the VIOS has access to it.
- Ensure that IBM Systems Director Common Agent is started on the VIOS.
- Discover and request access to the storage and the operating system of the VIOS mentioned.

To create an image repository, follow the steps below:

1. From the VMControl summary page, go to the **Virtual Appliances** tab and click **Create image repository** (Figure 5-76).

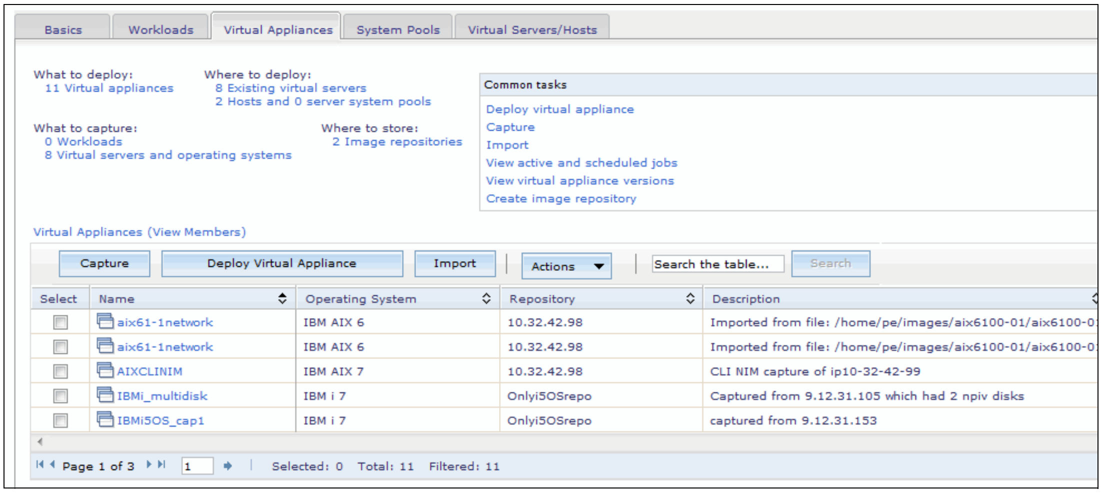

What to deploy:
11 Virtual appliances

Where to deploy:
8 Existing virtual servers
2 Hosts and 0 server system pools

What to capture:
0 Workloads
8 Virtual servers and operating systems

Where to store:
2 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)

Select	Name	Operating System	Repository	Description
□	aix61-1network	IBM AIX 6	10.32.42.98	Imported from file: /home/pe/images/aix6100-01/aix6100-0
□	aix61-1network	IBM AIX 6	10.32.42.98	Imported from file: /home/pe/images/aix6100-01/aix6100-0
□	AIXCLINIM	IBM AIX 7	10.32.42.98	CLI NIM capture of ip10-32-42-99
□	IBMi_multidisk	IBM i 7	Onlyi5OSrepo	Captured from 9.12.31.105 which had 2 npiv disks
□	IBMi5OS_cap1	IBM i 7	Onlyi5OSrepo	captured from 9.12.31.153

Page 1 of 3 | 1 | Selected: 0 Total: 11 Filtered: 11

Figure 5-76 IBM Systems Director VMControl: Virtual Appliances tab

2. Follow the instructions in the Create Image Repository wizard (Figure 5-77). Only systems that satisfy the requirement for hosting an image repository will be available to select.

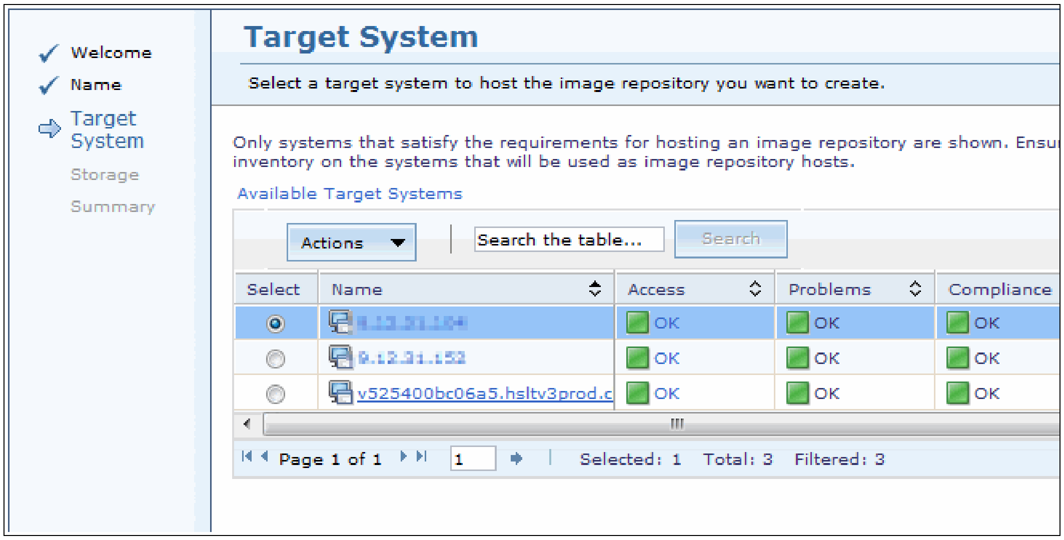

✓ Welcome
✓ Name
⇨ Target System
Storage
Summary

Target System

Select a target system to host the image repository you want to create.

Only systems that satisfy the requirements for hosting an image repository are shown. Ensu
inventory on the systems that will be used as image repository hosts.

Available Target Systems

Actions ▾ | Search the table... | Search

Select	Name	Access	Problems	Compliance
●	[redacted]	🟩 OK	🟩 OK	🟩 OK
○	[redacted]	🟩 OK	🟩 OK	🟩 OK
○	v525400bc06a5.hsltv3prod.c	🟩 OK	🟩 OK	🟩 OK

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

Figure 5-77 Select your VIOS

3. Select the storage to use for the image repository (Figure 5-78). Take note that this storage was previously discovered and the inventory was collected.

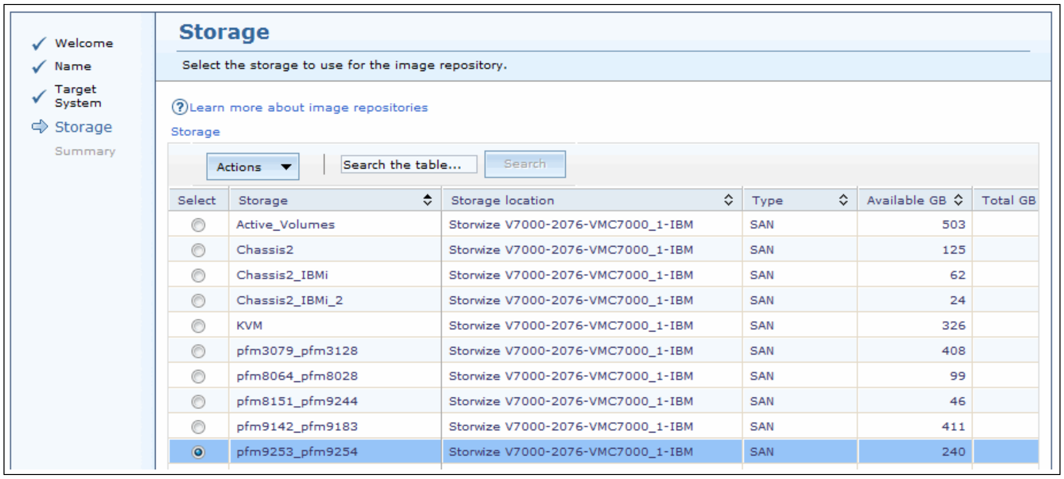

✓ Welcome
✓ Name
✓ Target System
⇨ Storage
Summary

Storage

Select the storage to use for the image repository.

?Learn more about image repositories

Storage

Actions ▾ | Search the table... | Search

Select	Storage	Storage location	Type	Available GB	Total GB
○	Active_Volumes	Storwize V7000-2076-VMC7000_1-IBM	SAN	503	
○	Chassis2	Storwize V7000-2076-VMC7000_1-IBM	SAN	125	
○	Chassis2_IBMi	Storwize V7000-2076-VMC7000_1-IBM	SAN	62	
○	Chassis2_IBMi_2	Storwize V7000-2076-VMC7000_1-IBM	SAN	24	
○	KVM	Storwize V7000-2076-VMC7000_1-IBM	SAN	326	
○	pfm3079_pfm3128	Storwize V7000-2076-VMC7000_1-IBM	SAN	408	
○	pfm8064_pfm8028	Storwize V7000-2076-VMC7000_1-IBM	SAN	99	
○	pfm8151_pfm9244	Storwize V7000-2076-VMC7000_1-IBM	SAN	46	
○	pfm9142_pfm9183	Storwize V7000-2076-VMC7000_1-IBM	SAN	411	
●	pfm9253_pfm9254	Storwize V7000-2076-VMC7000_1-IBM	SAN	240	

Figure 5-78 Storage selection for VIOS image repository

4. Review the summary as shown in Figure 5-79.

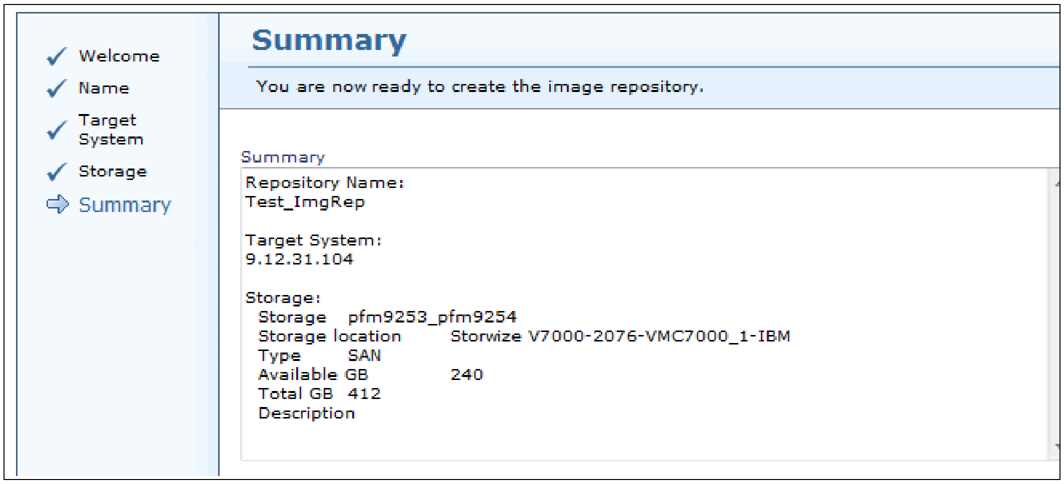

Summary

You are now ready to create the image repository.

Summary

```
Repository Name:
Test_ImgRep

Target System:
9.12.31.104

Storage:
  Storage   pfm9253_pfm9254
  Storage location      Storwize V7000-2076-VMC7000_1-IBM
  Type      SAN
  Available GB          240
  Total GB  412
  Description
```

Figure 5-79 New VIOS image repository summary

5. IBM Systems Director automatically discovers the image repository after it is created. See Figure 5-80.

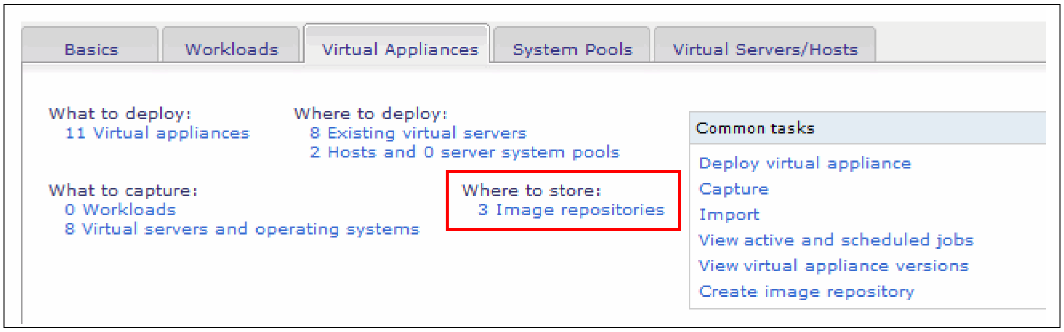

What to deploy:
11 Virtual appliances

Where to deploy:
8 Existing virtual servers
2 Hosts and 0 server system pools

What to capture:
0 Workloads
8 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

Figure 5-80 New VIOS image repository created

6. On the Virtual Appliances page, click the **Image Repositories** view (Figure 5-81).

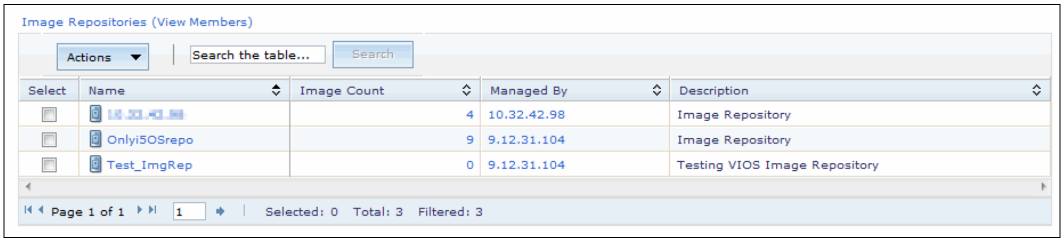

Select	Name	Image Count	Managed By	Description
□	[redacted]	4	10.32.42.98	Image Repository
□	Onlyi5OSrepo	9	9.12.31.104	Image Repository
□	Test_ImgRep	0	9.12.31.104	Testing VIOS Image Repository

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

Figure 5-81 Image Repositories view

You can add virtual appliances to your repositories by using the Capture and Import tasks in VMControl. The metadata that is associated with the virtual appliance is stored in the image repository, and the image is stored in the storage pool.

5.7.8 Managing server system pools with VMControl Enterprise Edition

As previously explained, IBM Systems Director VMControl Enterprise Edition can work with server systems pools. A server system pool allows you to group similar hosts and define specialized capabilities for the virtual servers running on the hosts. A server system pool has potential for capabilities such as live relocation, static relocation, and automated network relocation with network system pools (NSPs), if the hosts are enabled for these features.

Creation of storage system pools

A server system pool cannot be created without specifying the storage. In addition, hosts that are not connected to the same shared storage as the server system pool cannot be added to the pool. You can configure your SAN Storage Provider to be used with IBM Systems Director VMControl Power Systems server system pools.

To create a storage system pool, perform the following steps:

1. Go to IBM Systems Director VMControl and click the **System Pools** page. Then, select the **Storage system pools** view, and click **Create**. See Figure 5-82.

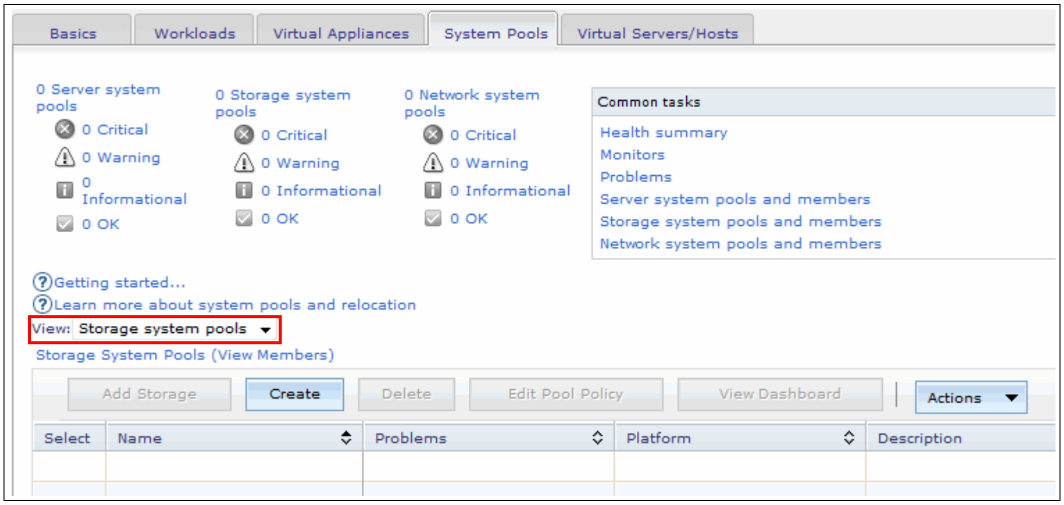

View: Storage system pools

Figure 5-82 VMControl: System Pools view

2. Complete the wizard steps. Enter the storage system pool name. Then, select the storage subsystem that will be used (Figure 5-83 on page 331).

> **Note:** When selecting multiple storages, they must be in the same zone.

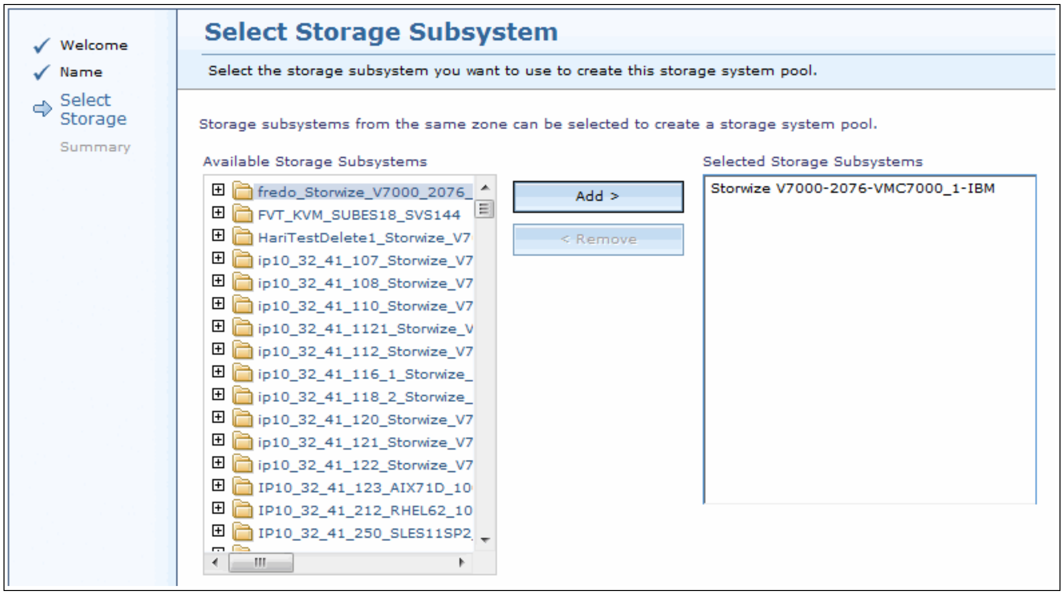

✓ Welcome
✓ Name
Select Storage
Summary

Select Storage Subsystem

Select the storage subsystem you want to use to create this storage system pool.

Storage subsystems from the same zone can be selected to create a storage system pool.

Available Storage Subsystems

- fredo_Storwize_V7000_2076_
- FVT_KVM_SUBES18_SVS144
- HariTestDelete1_Storwize_V7
- ip10_32_41_107_Storwize_V7
- ip10_32_41_108_Storwize_V7
- ip10_32_41_110_Storwize_V7
- ip10_32_41_1121_Storwize_V
- ip10_32_41_112_Storwize_V7
- ip10_32_41_116_1_Storwize_
- ip10_32_41_118_2_Storwize_
- ip10_32_41_120_Storwize_V7
- ip10_32_41_121_Storwize_V7
- ip10_32_41_122_Storwize_V7
- IP10_32_41_123_AIX71D_10
- IP10_32_41_212_RHEL62_10
- IP10_32_41_250_SLES11SP2_

Add >
< Remove

Selected Storage Subsystems

Storwize V7000-2076-VMC7000_1-IBM

Figure 5-83 Create storage system pools

3. After the wizard is complete, the new storage system pool is listed, as shown in Figure 5-84.

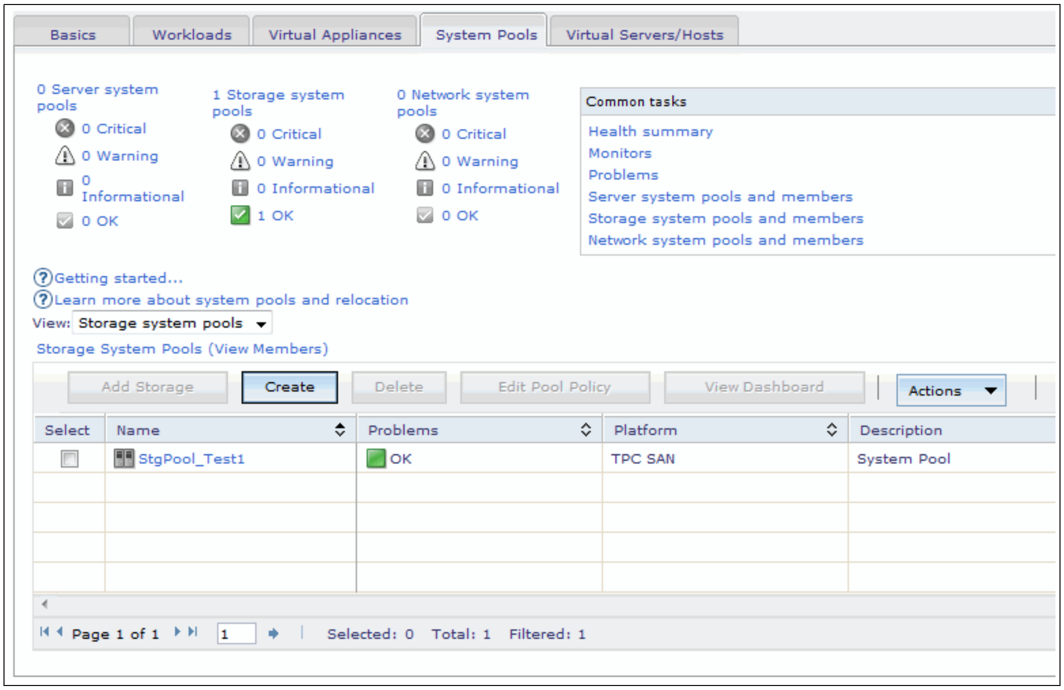

Basics | Workloads | Virtual Appliances | System Pools | Virtual Servers/Hosts

0 Server system pools
0 Critical
0 Warning
0 Informational
0 OK

1 Storage system pools
0 Critical
0 Warning
0 Informational
1 OK

0 Network system pools
0 Critical
0 Warning
0 Informational
0 OK

Common tasks
Health summary
Monitors
Problems
Server system pools and members
Storage system pools and members
Network system pools and members

Getting started...
Learn more about system pools and relocation
View: Storage system pools

Storage System Pools (View Members)

Add Storage | Create | Delete | Edit Pool Policy | View Dashboard | Actions

Select	Name	Problems	Platform	Description
□	StgPool_Test1	🟩 OK	TPC SAN	System Pool

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

Figure 5-84 VMControl: Storage system pools view

Creation of server system pools

After creating the storage system pool, a server system pool can be created. To create a server system pool, do the following steps:

1. On the VMControl System Pools page, select the **Server system pools** view and click **Create** (Figure 5-85).

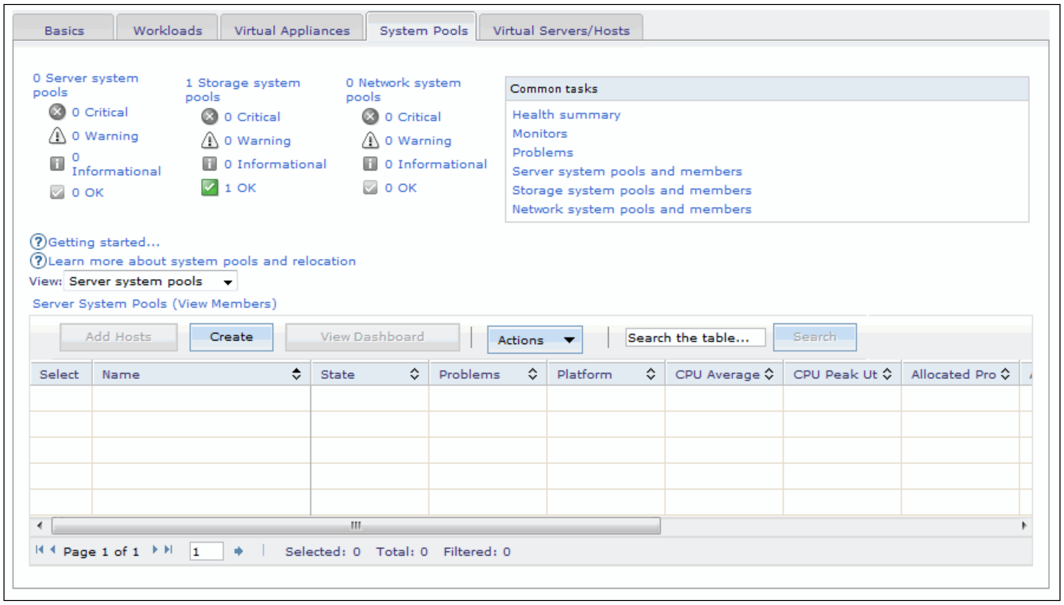

Basics | Workloads | Virtual Appliances | System Pools | Virtual Servers/Hosts

0 Server system pools
- 0 Critical
- 0 Warning
- 0 Informational
- 0 OK

1 Storage system pools
- 0 Critical
- 0 Warning
- 0 Informational
- 1 OK

0 Network system pools
- 0 Critical
- 0 Warning
- 0 Informational
- 0 OK

Common tasks
- Health summary
- Monitors
- Problems
- Server system pools and members
- Storage system pools and members
- Network system pools and members

Getting started...

Learn more about system pools and relocation

View: Server system pools

Server System Pools (View Members)

Add Hosts | Create | View Dashboard | Actions | Search the table... | Search

Select	Name	State	Problems	Platform	CPU Average	CPU Peak Ut	Allocated Pro

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

Figure 5-85 VMControl: Server system pools view

2. Complete the name and the description. In the Pooling Criteria window (Figure 5-86), ensure that the check box “Only add hosts capable of live virtual server relocation” is checked to be sure that you will be grouping servers with relocation capabilities. If a network system pool was previously configured, the “Network deployment criteria” check box will be enabled to configure.

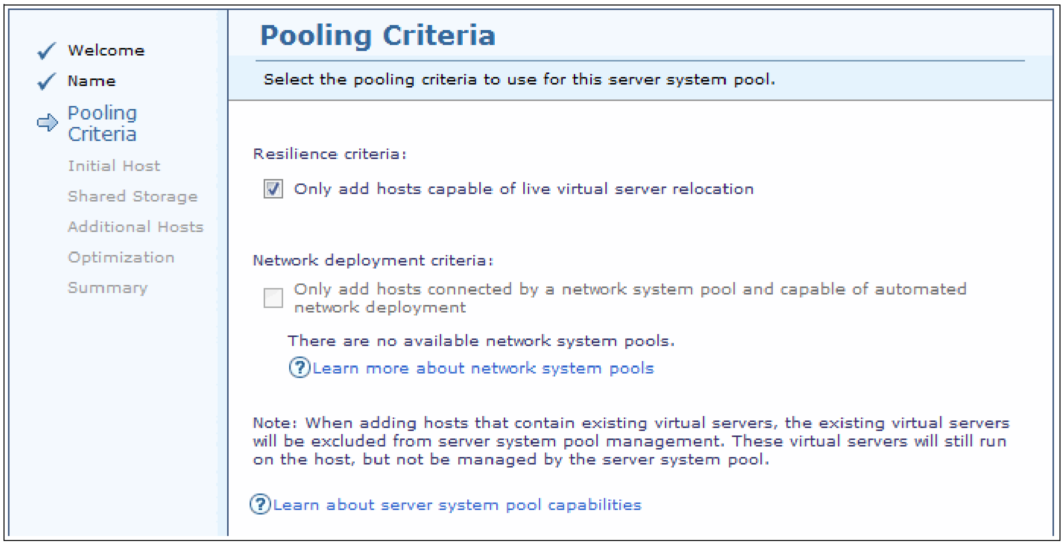

✓ Welcome
✓ Name
Pooling Criteria
Initial Host
Shared Storage
Additional Hosts
Optimization
Summary

Pooling Criteria

Select the pooling criteria to use for this server system pool.

Resilience criteria:

☑ Only add hosts capable of live virtual server relocation

Network deployment criteria:

☐ Only add hosts connected by a network system pool and capable of automated network deployment

There are no available network system pools.

Learn more about network system pools

Note: When adding hosts that contain existing virtual servers, the existing virtual servers will be excluded from server system pool management. These virtual servers will still run on the host, but not be managed by the server system pool.

Learn about server system pool capabilities

Figure 5-86 VMControl: Server system pools pooling criteria

> **Tip:** The existing virtual servers on the host that are resilient capable can be grouped as a workload to bring them under server system pool management.

3. Select your **Initial Host** (Figure 5-87). This host is used to find similar hosts that support the required capabilities for this server system pool.

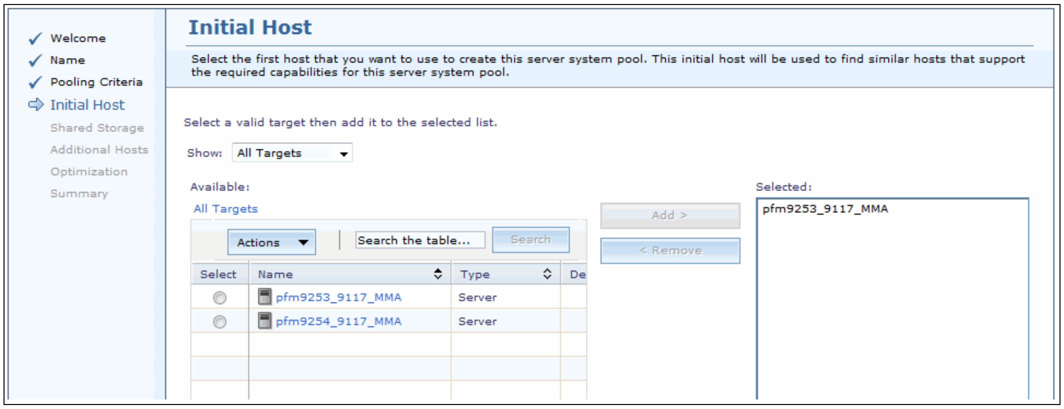

Initial Host

Select the first host that you want to use to create this server system pool. This initial host will be used to find similar hosts that support the required capabilities for this server system pool.

Select a valid target then add it to the selected list.

Show: All Targets

Available:

All Targets

Search the table...

Select	Name	Type	De
○	pfm9253_9117_MMA	Server	
○	pfm9254_9117_MMA	Server	

Selected:

pfm9253_9117_MMA

Figure 5-87 VMControl: Server system pools Initial Host configuration

4. Configure the **Shared Storage** (Figure 5-88). In the available shared storage, the storage system pool previously created will be available to be selected.

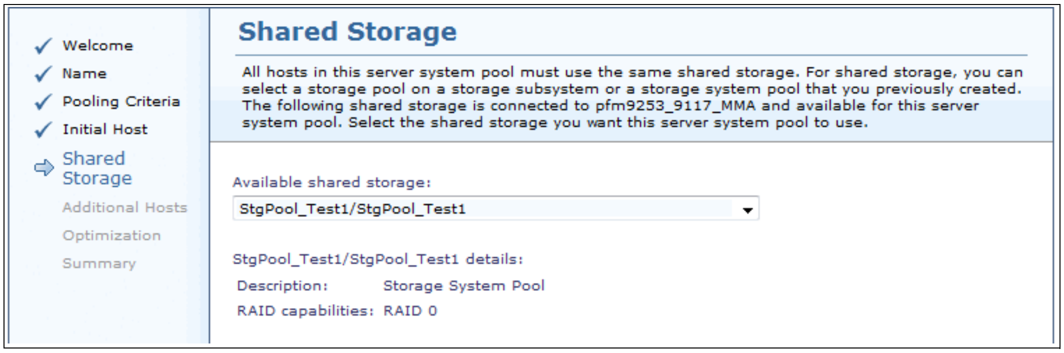

Shared Storage

All hosts in this server system pool must use the same shared storage. For shared storage, you can select a storage pool on a storage subsystem or a storage system pool that you previously created. The following shared storage is connected to pfm9253_9117_MMA and available for this server system pool. Select the shared storage you want this server system pool to use.

Available shared storage:

StgPool_Test1/StgPool_Test1

StgPool_Test1/StgPool_Test1 details:

Description: Storage System Pool

RAID capabilities: RAID 0

Figure 5-88 VMControl: Server system pools Shared Storage configuration

5. Add additional hosts to your server system pool. Remember that only the ones that are compatible with your initial selected host will be displayed (Figure 5-89).

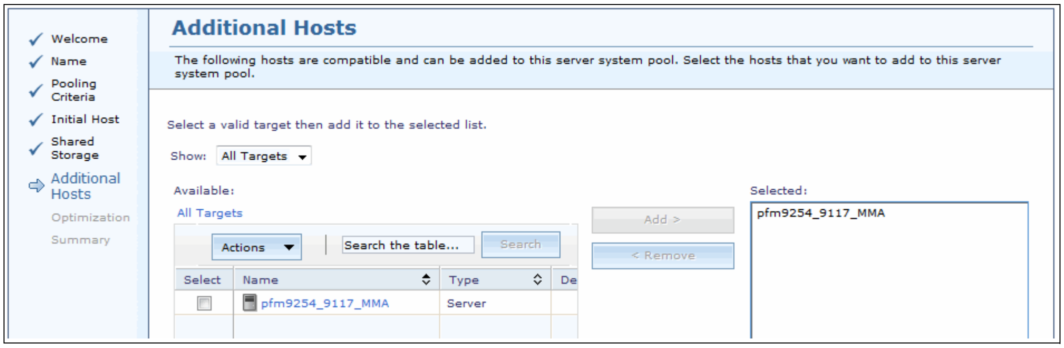

Additional Hosts

The following hosts are compatible and can be added to this server system pool. Select the hosts that you want to add to this server system pool.

Select a valid target then add it to the selected list.

Show: All Targets

Available:

All Targets

Search the table...

Select	Name	Type	De
□	pfm9254_9117_MMA	Server	

Selected:

pfm9254_9117_MMA

Figure 5-89 VMControl: Add hosts to the server system pool

6. Configure the optimization settings (Figure 5-90); manual or automatic optimization can be configured. Optimization enables the analysis and periodic performance improvement of all virtual servers within your server system pool.

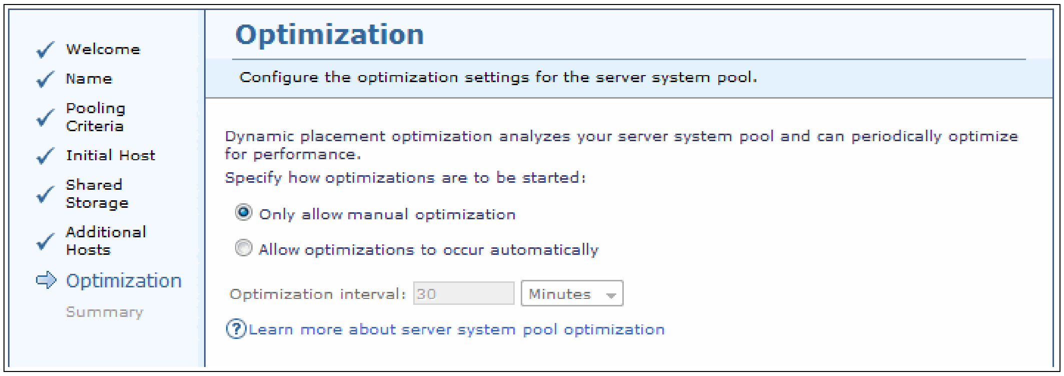

Optimization

Configure the optimization settings for the server system pool.

Dynamic placement optimization analyzes your server system pool and can periodically optimize for performance.

Specify how optimizations are to be started:

● Only allow manual optimization

○ Allow optimizations to occur automatically

Optimization interval: 30 Minutes

?Learn more about server system pool optimization

Figure 5-90 VMControl: Server system pool optimization settings

For more information about server system pool optimization, see the following site:

http://pic.dhe.ibm.com/infocenter/director/pubs/topic/com.ibm.director.vim.helps.doc/fsd0_vim_c_learnmore_system_pool_optimization_settings.html

7. When the wizard is completed, the new server system pool is listed in the “Server system pools” view, as shown in Figure 5-91.

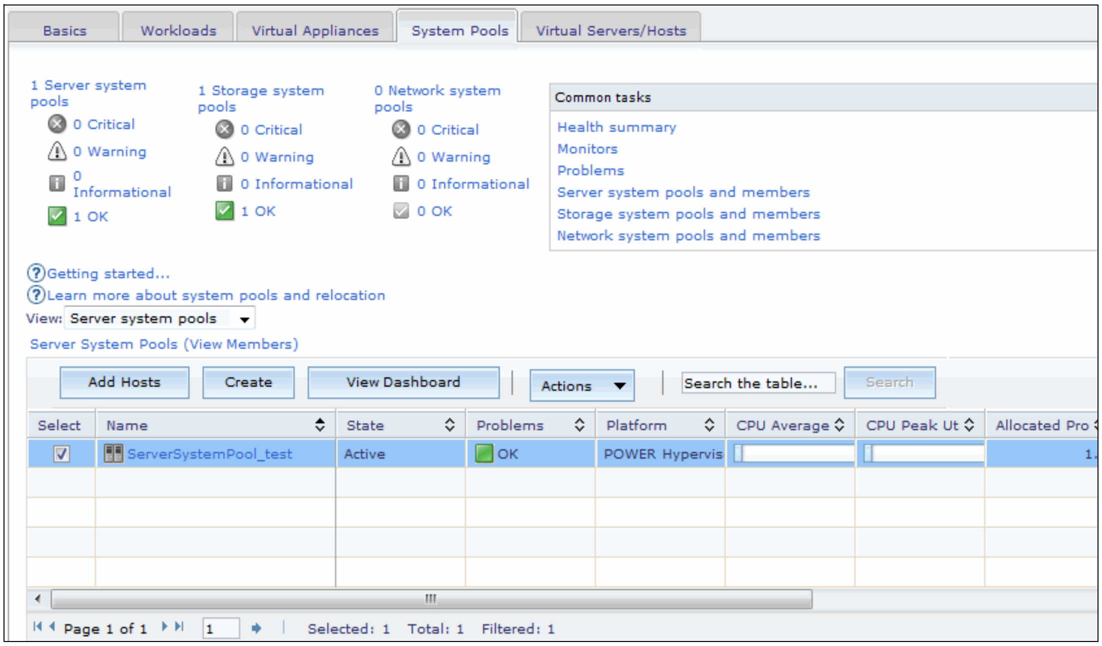

1 Server system pools
0 Critical
0 Warning
0 Informational
1 OK

1 Storage system pools
0 Critical
0 Warning
0 Informational
1 OK

0 Network system pools
0 Critical
0 Warning
0 Informational
0 OK

Common tasks
Health summary
Monitors
Problems
Server system pools and members
Storage system pools and members
Network system pools and members

?Getting started...
?Learn more about system pools and relocation
View: Server system pools
Server System Pools (View Members)

Select	Name	State	Problems	Platform	CPU Average	CPU Peak Ut	Allocated Pro
☑	ServerSystemPool_test	Active	OK	POWER Hypervis			1

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

Figure 5-91 VMControl: Server system pools view

Managing workloads in a server system pool environment

A workload is a deployed virtual appliance that allows you to monitor and manage one or more virtual servers as a single entity. With IBM Systems Director VMControl Standard Edition, you can capture, import, and deploy virtual appliances. Workloads are created by deploying virtual appliances or by grouping existing virtual servers as a workload. The first step in managing virtual appliances and workloads is creating an image repository.

Capture a virtual appliance using NIM

By capturing a virtual appliance, the operations result in a virtual appliance that you can deploy to create a new virtual server. There are three types of sources that can be captured:

- A virtual server (LPAR) or a workload
- An existing mksysb image file
- An existing NIM **`mksysb`** resource, **`lpp_source`** directory, or **`lpp_source`** resource on your NIM master

Before capturing virtual appliances, check the requirements listed in the information center:

`http://pic.dhe.ibm.com/infocenter/director/pubs/index.jsp?topic=%2Fcom.ibm.director.vim.helps.doc%2Ffsd0_vim_r_power_capture_reqs.html`

To capture a running virtual server using NIM, execute the following steps:

1. Go to the VMControl Virtual Appliances page and click **Capture** (Figure 5-92).

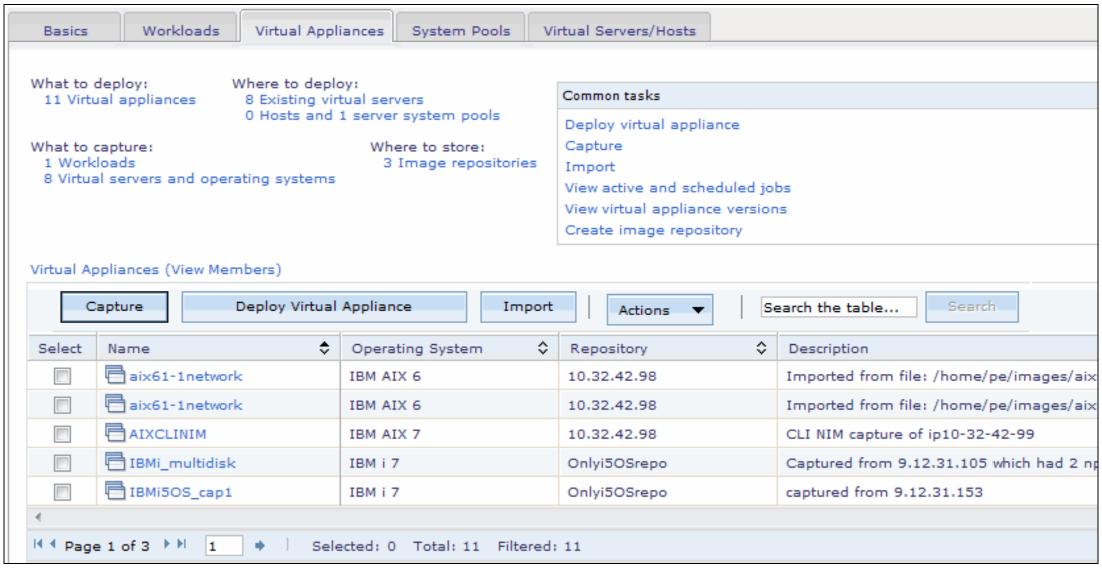

What to deploy:
11 Virtual appliances

Where to deploy:
8 Existing virtual servers
0 Hosts and 1 server system pools

What to capture:
1 Workloads
8 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)

Select	Name	Operating System	Repository	Description
□	aix61-1network	IBM AIX 6	10.32.42.98	Imported from file: /home/pe/images/aix
□	aix61-1network	IBM AIX 6	10.32.42.98	Imported from file: /home/pe/images/aix
□	AIXCLINIM	IBM AIX 7	10.32.42.98	CLI NIM capture of ip10-32-42-99
□	IBMi_multidisk	IBM i 7	Onlyi5OSrepo	Captured from 9.12.31.105 which had 2 np
□	IBMi5OS_cap1	IBM i 7	Onlyi5OSrepo	captured from 9.12.31.153

Page 1 of 3 | 1 | Selected: 0 Total: 11 Filtered: 11

Figure 5-92 VMControl: Virtual Appliances page

2. Enter the name, description, and search tags for the new virtual appliance, as shown in Figure 5-93.

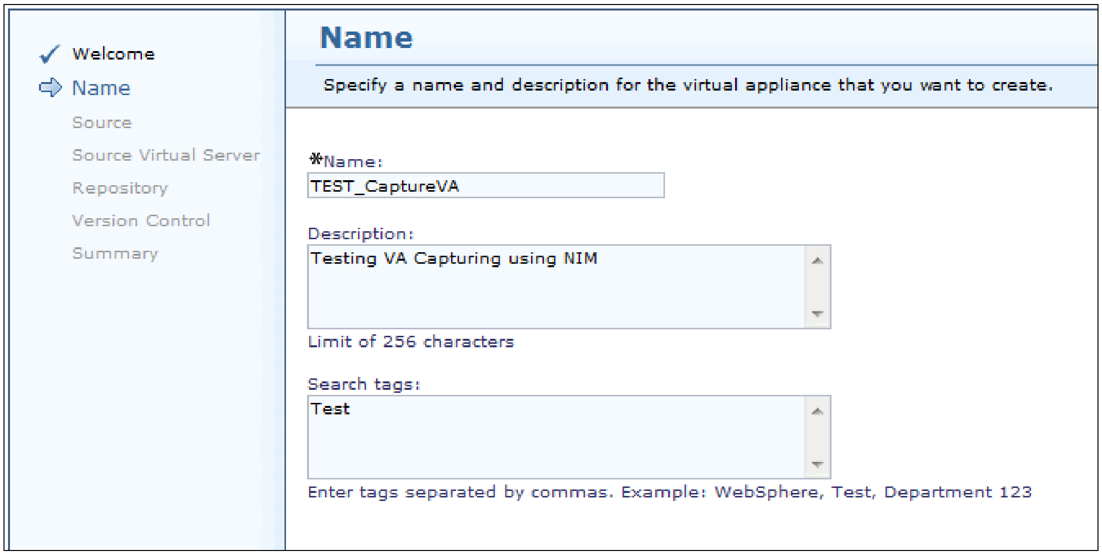

Name

Specify a name and description for the virtual appliance that you want to create.

*Name: TEST_CaptureVA

Description: Testing VA Capturing using NIM

Limit of 256 characters

Search tags: Test

Enter tags separated by commas. Example: WebSphere, Test, Department 123

Figure 5-93 VMControl: Capture a virtual appliance

3. Select the source. In this case, select **Virtual Server** (Figure 5-94).

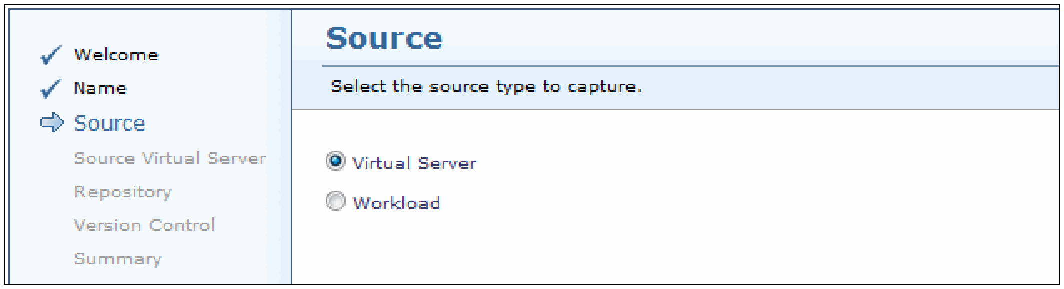

Source

Select the source type to capture.

● Virtual Server

○ Workload

Figure 5-94 VMControl: Virtual appliance source configuration

4. Select the source **Virtual Server** that will be captured (Figure 5-95 on page 337).

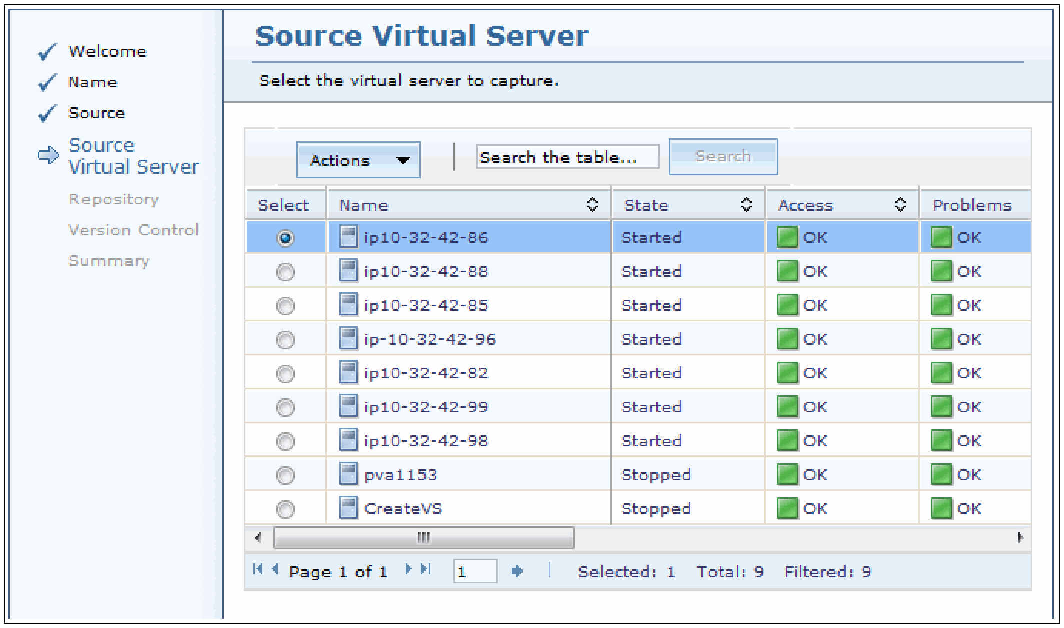

Source Virtual Server

Select the virtual server to capture.

Select	Name	State	Access	Problems
●	ip10-32-42-86	Started	🟩 OK	🟩 OK
○	ip10-32-42-88	Started	🟩 OK	🟩 OK
○	ip10-32-42-85	Started	🟩 OK	🟩 OK
○	ip-10-32-42-96	Started	🟩 OK	🟩 OK
○	ip10-32-42-82	Started	🟩 OK	🟩 OK
○	ip10-32-42-99	Started	🟩 OK	🟩 OK
○	ip10-32-42-98	Started	🟩 OK	🟩 OK
○	pva1153	Stopped	🟩 OK	🟩 OK
○	CreateVS	Stopped	🟩 OK	🟩 OK

Page 1 of 1   1   Selected: 1 Total: 9 Filtered: 9

Figure 5-95 VMControl: Virtual appliance Source Virtual Server selection

> **Note:** The operating system on the virtual server that you want to capture must be discovered, accessed, and inventoried with IBM Systems Director.

5. Select the repository to be used. In this step, decide whether to use NIM or SCS. In this case, select your NIM server (Figure 5-96).

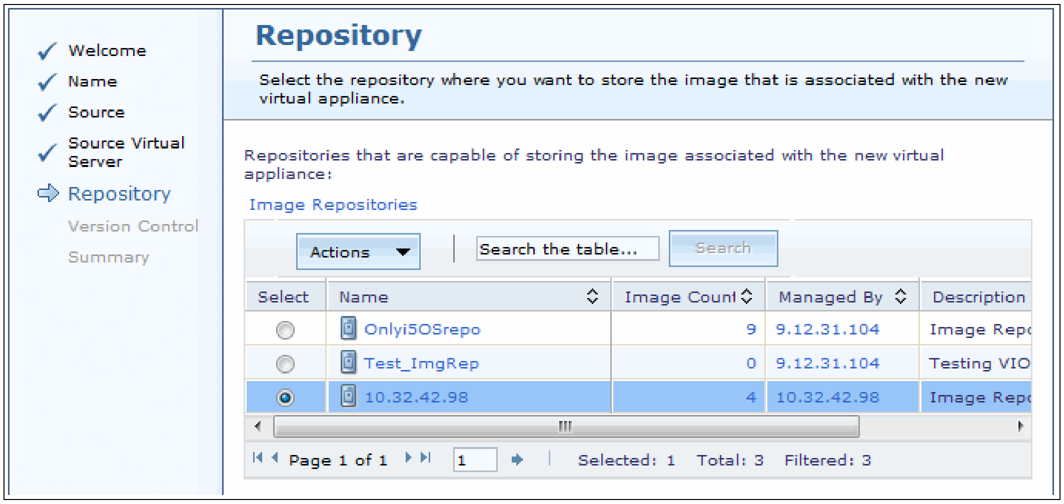

Repository

Select the repository where you want to store the image that is associated with the new virtual appliance.

Repositories that are capable of storing the image associated with the new virtual appliance:

Select	Name	Image Count	Managed By	Description
○	Onlyi5OSrepo	9	9.12.31.104	Image Repo
○	Test_ImgRep	0	9.12.31.104	Testing VIO
●	10.32.42.98	4	10.32.42.98	Image Repo

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

Figure 5-96 VMControl: Virtual appliance Repository configuration

6. Figure 5-97 shows the Network Mapping configuration. Specify a description to use for each virtual network.

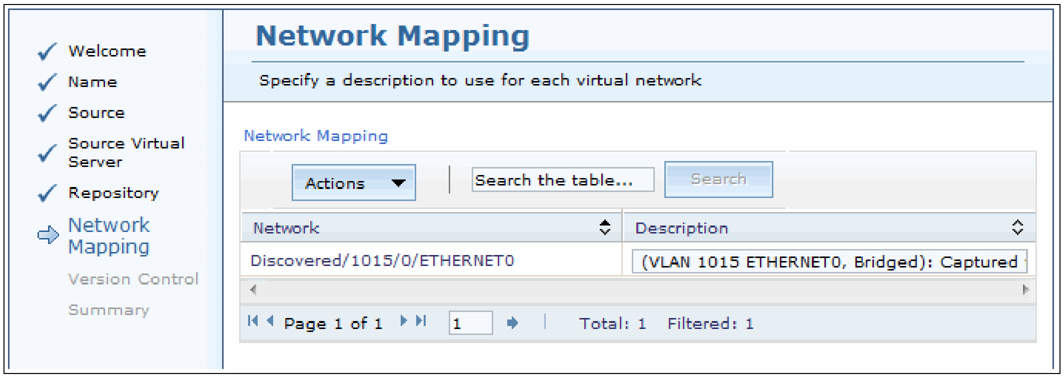

Network Mapping

Specify a description to use for each virtual network

Network Mapping

Network	Description
Discovered/1015/0/ETHERNET0	(VLAN 1015 ETHERNET0, Bridged): Captured

Figure 5-97 VMControl: Virtual appliance Network Mapping configuration

7. Select the version information for **Version Control** (Figure 5-98 on page 339). During any tasks that involve virtual appliances (for example, capture, deploy, import), IBM Systems Director VMControl automatically generates, maintains, and manages version information for the virtual appliances.

   For more information about virtual appliance versions, see the following site:

   http://pic.dhe.ibm.com/infocenter/director/pubs/index.jsp?topic=%2Fcom.ibm.director.vim.helps.doc%2Ffsd0_vim_r_revisions.html

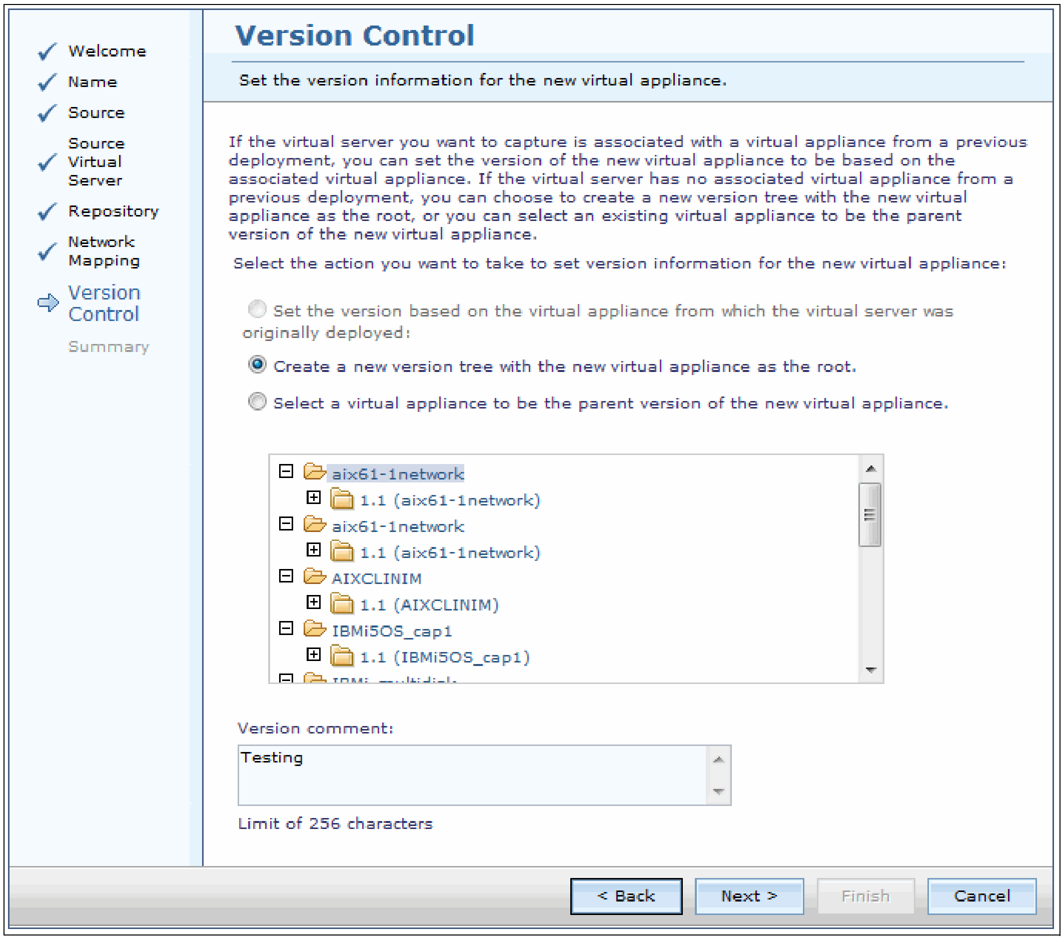

Version Control

Set the version information for the new virtual appliance.

If the virtual server you want to capture is associated with a virtual appliance from a previous deployment, you can set the version of the new virtual appliance to be based on the associated virtual appliance. If the virtual server has no associated virtual appliance from a previous deployment, you can choose to create a new version tree with the new virtual appliance as the root, or you can select an existing virtual appliance to be the parent version of the new virtual appliance.

Select the action you want to take to set version information for the new virtual appliance:

○ Set the version based on the virtual appliance from which the virtual server was originally deployed:

● Create a new version tree with the new virtual appliance as the root.

○ Select a virtual appliance to be the parent version of the new virtual appliance.

Version comment: Testing

Limit of 256 characters

Figure 5-98 VMControl: Virtual Appliance version control configuration

8. Submit the job and check the results. See the example in Figure 5-99.

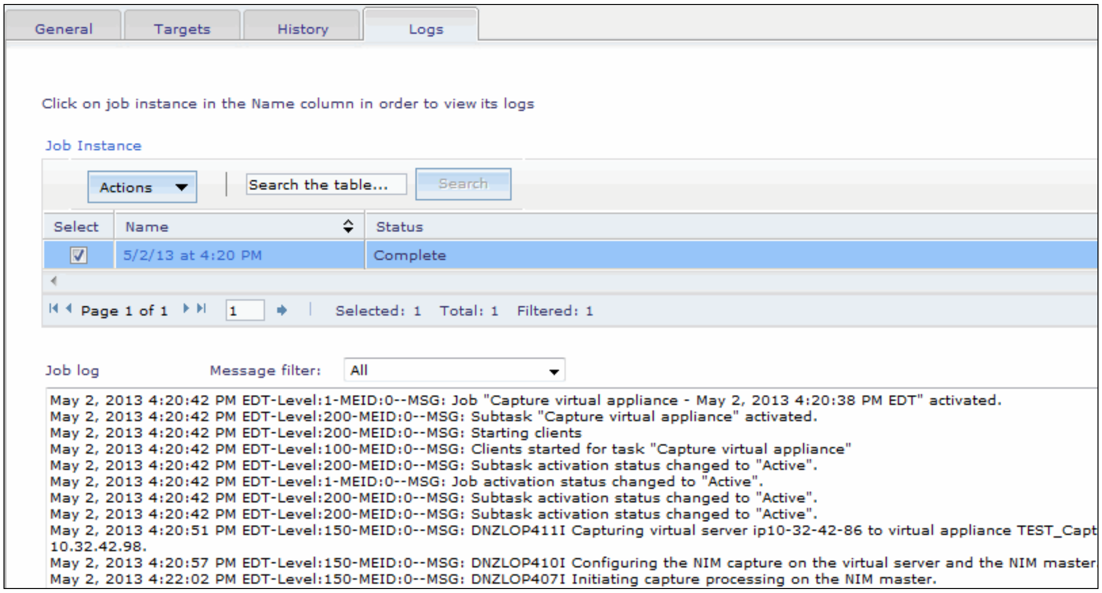

Click on job instance in the Name column in order to view its logs

Job Instance

Select	Name	Status
☑	5/2/13 at 4:20 PM	Complete

Selected: 1 Total: 1 Filtered: 1

Job log Message filter: All

```
May 2, 2013 4:20:42 PM EDT-Level:1-MEID:0--MSG: Job "Capture virtual appliance - May 2, 2013 4:20:38 PM EDT" activated.
May 2, 2013 4:20:42 PM EDT-Level:200-MEID:0--MSG: Subtask "Capture virtual appliance" activated.
May 2, 2013 4:20:42 PM EDT-Level:200-MEID:0--MSG: Starting clients
May 2, 2013 4:20:42 PM EDT-Level:100-MEID:0--MSG: Clients started for task "Capture virtual appliance"
May 2, 2013 4:20:42 PM EDT-Level:200-MEID:0--MSG: Subtask activation status changed to "Active".
May 2, 2013 4:20:42 PM EDT-Level:1-MEID:0--MSG: Job activation status changed to "Active".
May 2, 2013 4:20:42 PM EDT-Level:200-MEID:0--MSG: Subtask activation status changed to "Active".
May 2, 2013 4:20:42 PM EDT-Level:200-MEID:0--MSG: Subtask activation status changed to "Active".
May 2, 2013 4:20:51 PM EDT-Level:150-MEID:0--MSG: DNZLOP411I Capturing virtual server ip10-32-42-86 to virtual appliance TEST_Capt
10.32.42.98.
May 2, 2013 4:20:57 PM EDT-Level:150-MEID:0--MSG: DNZLOP410I Configuring the NIM capture on the virtual server and the NIM master
May 2, 2013 4:22:02 PM EDT-Level:150-MEID:0--MSG: DNZLOP407I Initiating capture processing on the NIM master.
```

Figure 5-99 Capturing virtual appliance job logs

Capture a virtual appliance using Storage Copy Services

You can create a virtual appliance by capturing any of the following sources:

- An IBM Power Systems virtual server (logical partition) or workload that contains a virtual server that is running AIX 5.3 or newer.
- An IBM Power Systems virtual server (logical partition) or workload that contains a virtual server that is running IBM i v7.1 TR3 or newer.
- An IBM Power Systems virtual server (logical partition) or workload that contains a virtual server that is running SUSE Linux Enterprise Server (SLES) 10 SP3 or newer, or Red Hat Enterprise Linux (RHEL) 5.4 or newer.

> **Note:** When capturing a virtual appliance using SCS, the virtual server must be powered off.

Check the requirements for capturing a virtual appliance using SCS by referring to the information center at the following link:

`http://pic.dhe.ibm.com/infocenter/director/pubs/topic/com.ibm.director.vim.helps.doc/fsd0_vim_r_sb_aix_on_power_capture_reqs.html`

Follow the same steps that are explained in the section, “Capture a virtual appliance using NIM” on page 335, to capture a virtual appliance using SCS.

Deploy a virtual appliance using workloads

When you create your image repositories and capture a virtual appliance, you can start working with virtual appliances and workloads in your environment.

1. Go to the VMControl Workloads page and click **Deploy virtual appliance** (Figure 5-100).

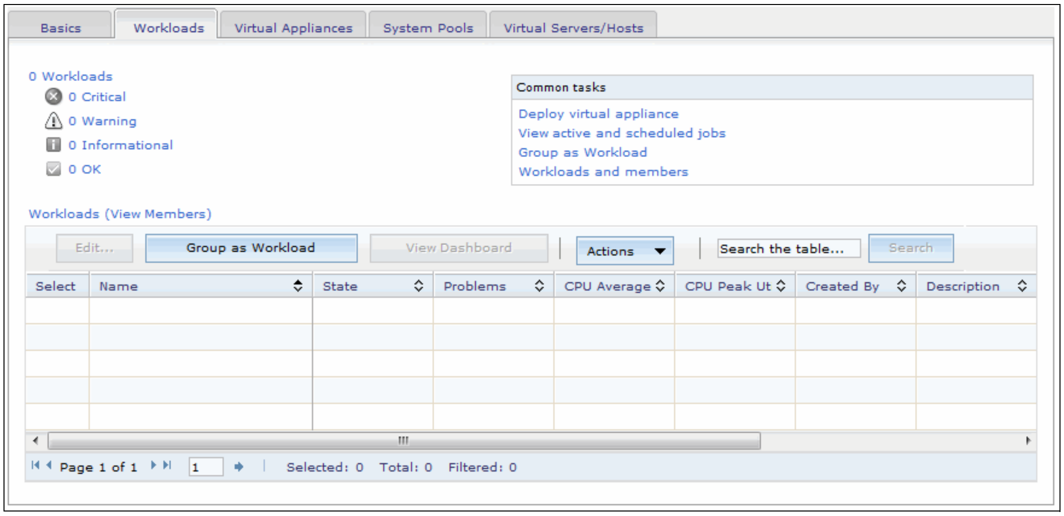

Select	Name	State	Problems	CPU Average	CPU Peak Ut	Created By	Description

Figure 5-100 VMControl Workloads page

2. Select the virtual appliance that you want to deploy (Figure 5-101 on page 341).

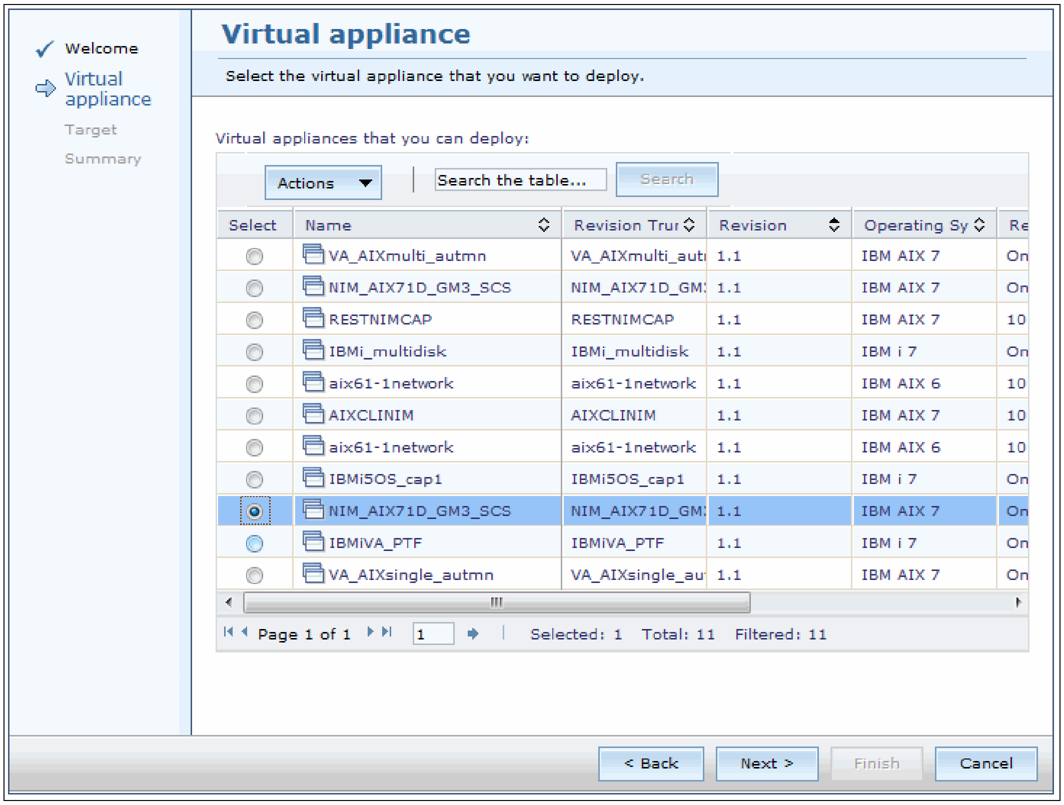

✓ Welcome
⇨ Virtual appliance
Target
Summary

Virtual appliance

Select the virtual appliance that you want to deploy.

Virtual appliances that you can deploy:

Actions ▼ | Search the table... | Search

Select	Name	Revision Trui	Revision	Operating Sy	Re
○	VA_AIXmulti_autmn	VA_AIXmulti_auti	1.1	IBM AIX 7	On
○	NIM_AIX71D_GM3_SCS	NIM_AIX71D_GM:	1.1	IBM AIX 7	On
○	RESTNIMCAP	RESTNIMCAP	1.1	IBM AIX 7	10
○	IBMi_multidisk	IBMi_multidisk	1.1	IBM i 7	On
○	aix61-1network	aix61-1network	1.1	IBM AIX 6	10
○	AIXCLINIM	AIXCLINIM	1.1	IBM AIX 7	10
○	aix61-1network	aix61-1network	1.1	IBM AIX 6	10
○	IBMi5OS_cap1	IBMi5OS_cap1	1.1	IBM i 7	On
●	NIM_AIX71D_GM3_SCS	NIM_AIX71D_GM:	1.1	IBM AIX 7	On
○	IBMiVA_PTF	IBMiVA_PTF	1.1	IBM i 7	On
○	VA_AIXsingle_autmn	VA_AIXsingle_au	1.1	IBM AIX 7	On

Page 1 of 1 | 1 | Selected: 1 Total: 11 Filtered: 11

< Back | Next > | Finish | Cancel

Figure 5-101 VMControl Workloads: Deploy virtual appliance page

3. Select your server system pool (Figure 5-102) where you are planning to deploy your virtual appliance. You can also override an existing virtual server.

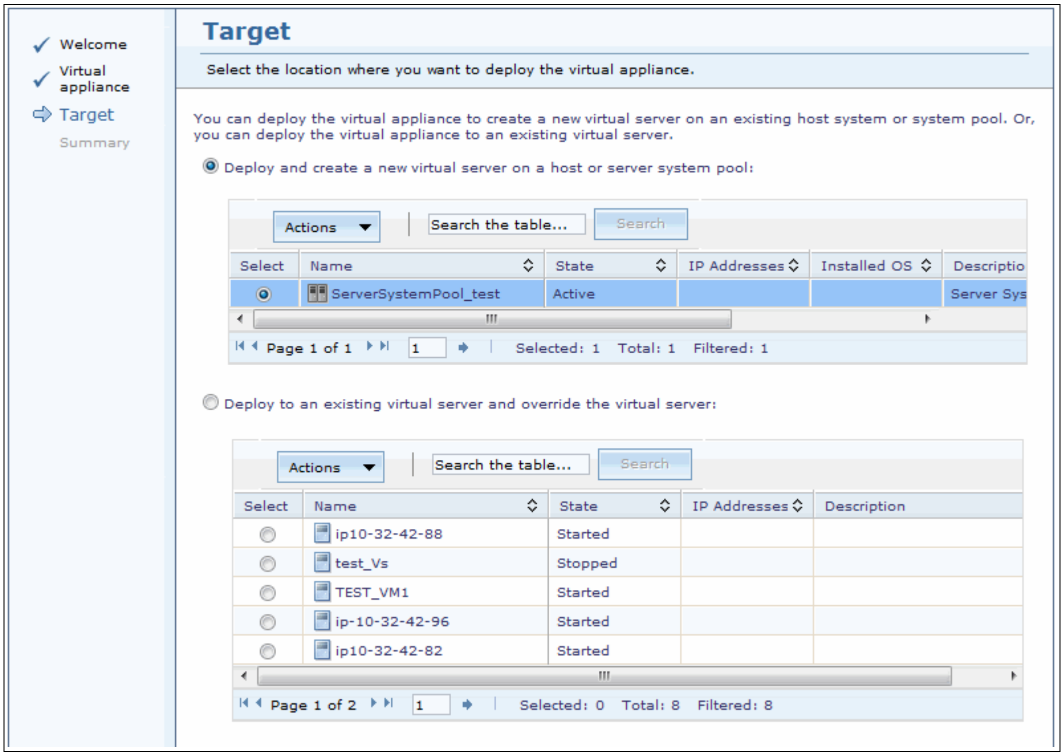

✓ Welcome
✓ Virtual appliance
⇨ Target
Summary

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 and create a new virtual server on a host or server system pool:

Actions ▼ | Search the table... | Search

Select	Name	State	IP Addresses	Installed OS	Descriptio
●	ServerSystemPool_test	Active			Server Sys

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

○ Deploy to an existing virtual server and override the virtual server:

Actions ▼ | Search the table... | Search

Select	Name	State	IP Addresses	Description
○	ip10-32-42-88	Started		
○	test_Vs	Stopped		
○	TEST_VM1	Started		
○	ip-10-32-42-96	Started		
○	ip10-32-42-82	Started		

Page 1 of 2 | 1 | Selected: 0 Total: 8 Filtered: 8

Figure 5-102 VMControl Workloads: Virtual appliance target

4. Specify the setting for the virtual disks as shown in Figure 5-103.

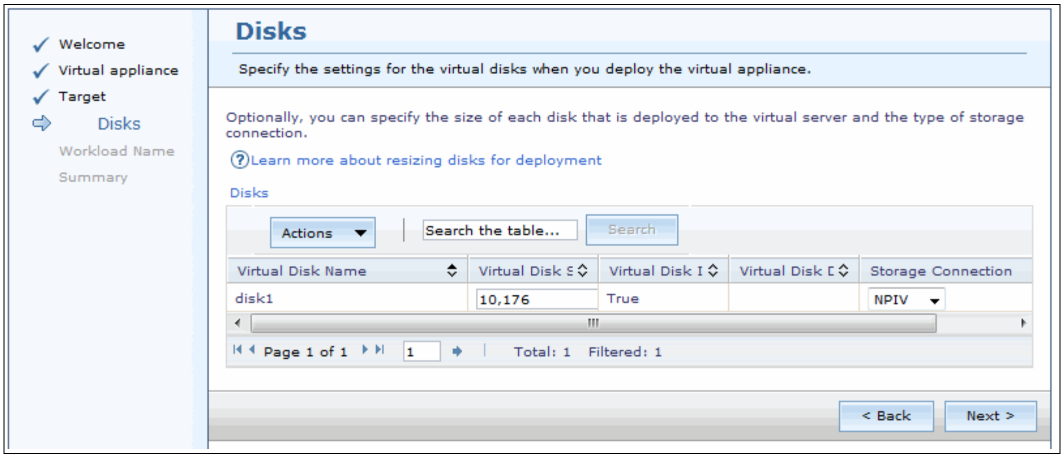

Disks

Specify the settings for the virtual disks when you deploy the virtual appliance.

Optionally, you can specify the size of each disk that is deployed to the virtual server and the type of storage connection.

Learn more about resizing disks for deployment

Disks

Virtual Disk Name	Virtual Disk S	Virtual Disk I	Virtual Disk D	Storage Connection
disk1	10,176	True		NPIV

Total: 1 Filtered: 1

Figure 5-103 VMControl Workloads: Virtual appliance disk configuration

5. Enter the name of your new workload as shown in Figure 5-104.

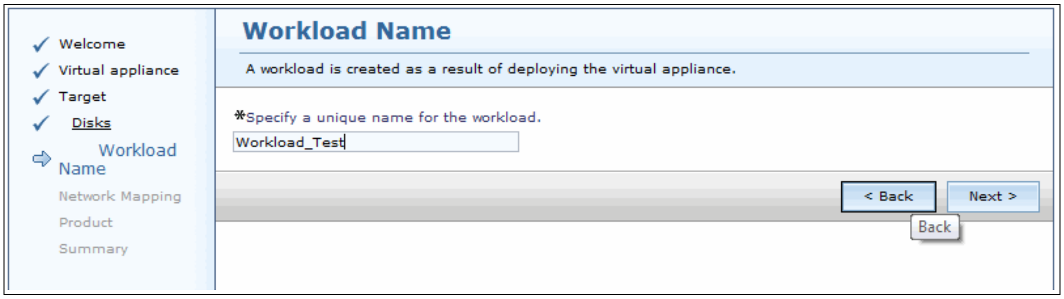

Workload Name

A workload is created as a result of deploying the virtual appliance.

*Specify a unique name for the workload.

Workload_Test

Figure 5-104 VMControl workloads: Workload Name page

6. Figure 5-105 shows the Network Mapping configuration. Select a virtual network for each network that is defined for the appliance.

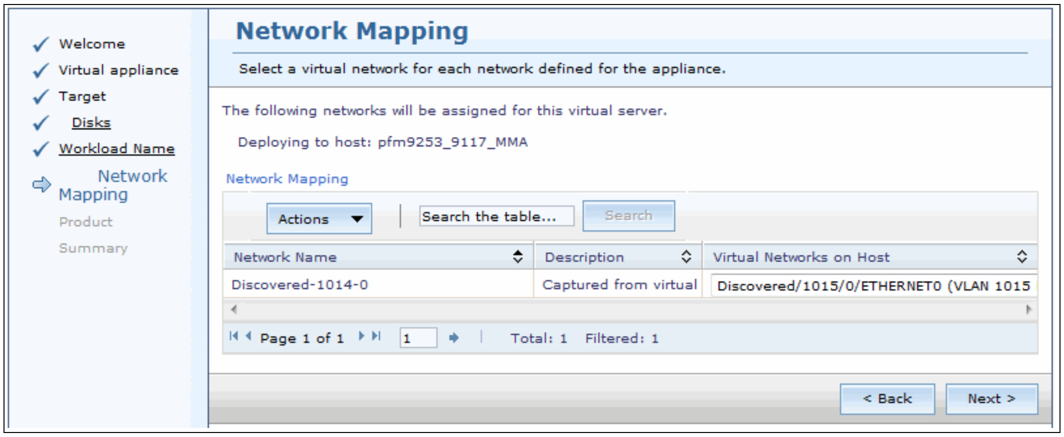

Network Mapping

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

The following networks will be assigned for this virtual server.

Deploying to host: pfm9253_9117_MMA

Network Mapping

Network Name	Description	Virtual Networks on Host
Discovered-1014-0	Captured from virtual	Discovered/1015/0/ETHERNET0 (VLAN 1015

Total: 1 Filtered: 1

Figure 5-105 VMControl workloads: Networking Mapping configuration

7. Configure the new virtual server. Enter the new server name, the DNS domain name, the DNS IP, and the default gateway IP. Also, enter the new IP address and the network mask, as shown in Figure 5-106.

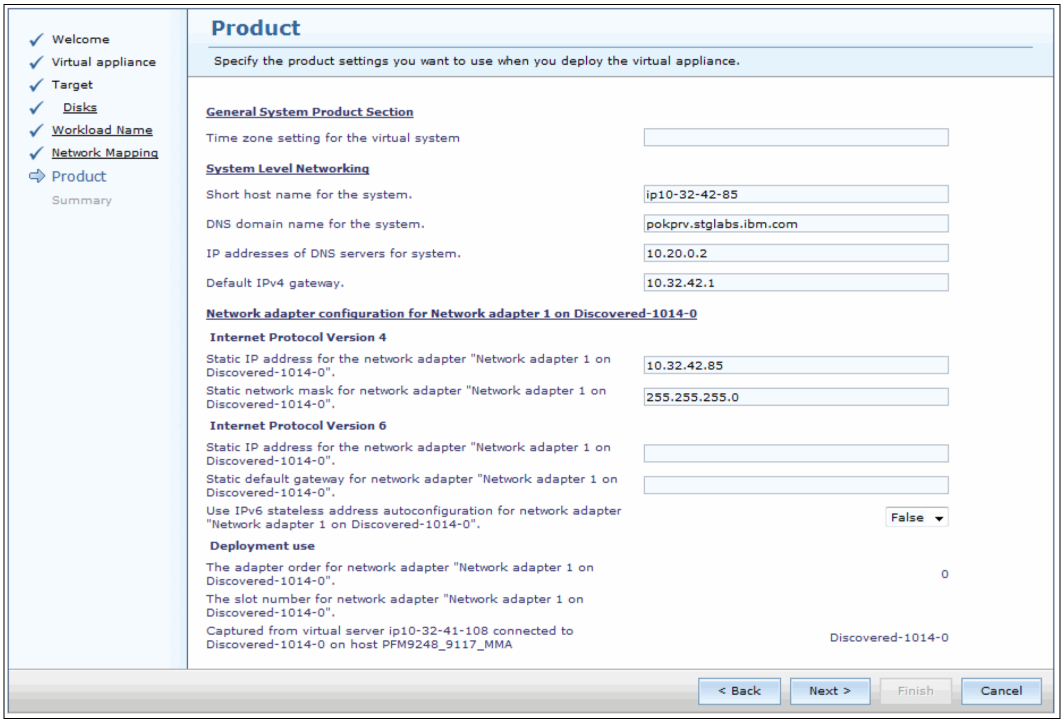

Product

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

General System Product Section

Time zone setting for the virtual system:

System Level Networking

Short host name for the system.: ip10-32-42-85

DNS domain name for the system.: pokprv.stglabs.ibm.com

IP addresses of DNS servers for system.: 10.20.0.2

Default IPv4 gateway.: 10.32.42.1

Network adapter configuration for Network adapter 1 on Discovered-1014-0

Internet Protocol Version 4

Static IP address for the network adapter "Network adapter 1 on Discovered-1014-0".: 10.32.42.85

Static network mask for network adapter "Network adapter 1 on Discovered-1014-0".: 255.255.255.0

Internet Protocol Version 6

Static IP address for the network adapter "Network adapter 1 on Discovered-1014-0".:

Static default gateway for network adapter "Network adapter 1 on Discovered-1014-0".:

Use IPv6 stateless address autoconfiguration for network adapter "Network adapter 1 on Discovered-1014-0".: False

Deployment use

The adapter order for network adapter "Network adapter 1 on Discovered-1014-0".: 0

The slot number for network adapter "Network adapter 1 on Discovered-1014-0".:

Captured from virtual server ip10-32-41-108 connected to Discovered-1014-0 on host PFM9248_9117_MMA: Discovered-1014-0

Figure 5-106 VMControl workloads: Virtual server configuration

8. Submit the job and check the log. See the example in Figure 5-107.

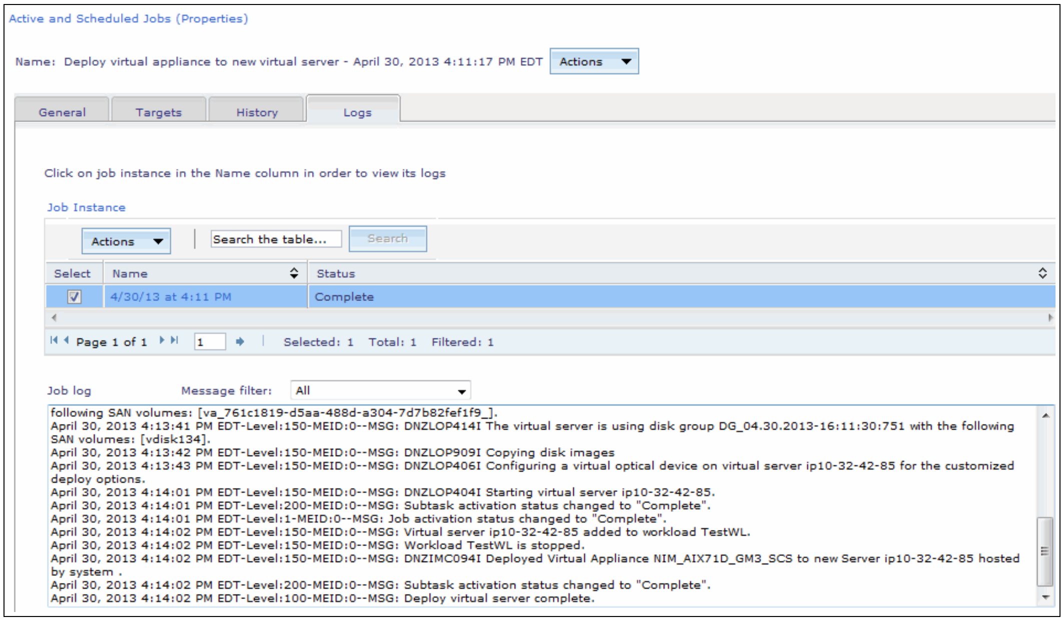

Name: Deploy virtual appliance to new virtual server - April 30, 2013 4:11:17 PM EDT

Click on job instance in the Name column in order to view its logs

Job Instance

Select	Name	Status
☑	4/30/13 at 4:11 PM	Complete

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

Job log Message filter: All

```
following SAN volumes: [va_761c1819-d5aa-488d-a304-7d7b82fef1f9_].
April 30, 2013 4:13:41 PM EDT-Level:150-MEID:0--MSG: DNZLOP414I The virtual server is using disk group DG_04.30.2013-16:11:30:751 with the following
SAN volumes: [vdisk134].
April 30, 2013 4:13:42 PM EDT-Level:150-MEID:0--MSG: DNZLOP909I Copying disk images
April 30, 2013 4:13:43 PM EDT-Level:150-MEID:0--MSG: DNZLOP406I Configuring a virtual optical device on virtual server ip10-32-42-85 for the customized
deploy options.
April 30, 2013 4:14:01 PM EDT-Level:150-MEID:0--MSG: DNZLOP404I Starting virtual server ip10-32-42-85.
April 30, 2013 4:14:01 PM EDT-Level:200-MEID:0--MSG: Subtask activation status changed to "Complete".
April 30, 2013 4:14:01 PM EDT-Level:1-MEID:0--MSG: Job activation status changed to "Complete".
April 30, 2013 4:14:02 PM EDT-Level:150-MEID:0--MSG: Virtual server ip10-32-42-85 added to workload TestWL.
April 30, 2013 4:14:02 PM EDT-Level:150-MEID:0--MSG: Workload TestWL is stopped.
April 30, 2013 4:14:02 PM EDT-Level:150-MEID:0--MSG: DNZIMC094I Deployed Virtual Appliance NIM_AIX71D_GM3_SCS to new Server ip10-32-42-85 hosted
by system .
April 30, 2013 4:14:02 PM EDT-Level:200-MEID:0--MSG: Subtask activation status changed to "Complete".
April 30, 2013 4:14:02 PM EDT-Level:100-MEID:0--MSG: Deploy virtual server complete.
```

Figure 5-107 Deploy virtual appliance job log

9. You can check the progress by opening the terminal from the HMC, as shown in Figure 5-108.

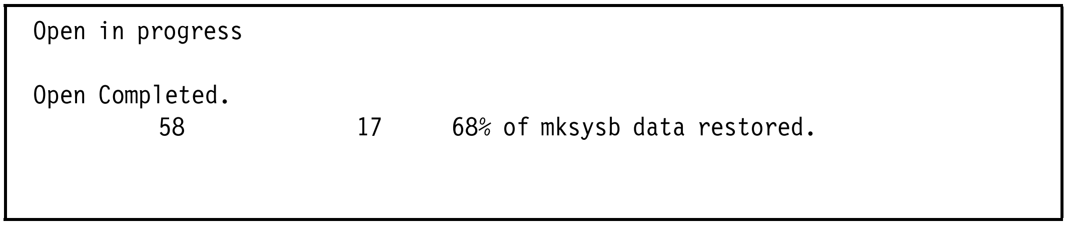

```
Open in progress

Open Completed.
         58            17      68% of mksysb data restored.
```

Figure 5-108 AIX terminal window

10. To run the initial configuration (Name and Network settings), VMControl creates a virtual optical device (client SCSI) to save all the configuration on this device. This initial configuration is applied during the first system startup. After one hour, this device is automatically deleted. You can check this device on the HMC. See the example in Figure 5-109.

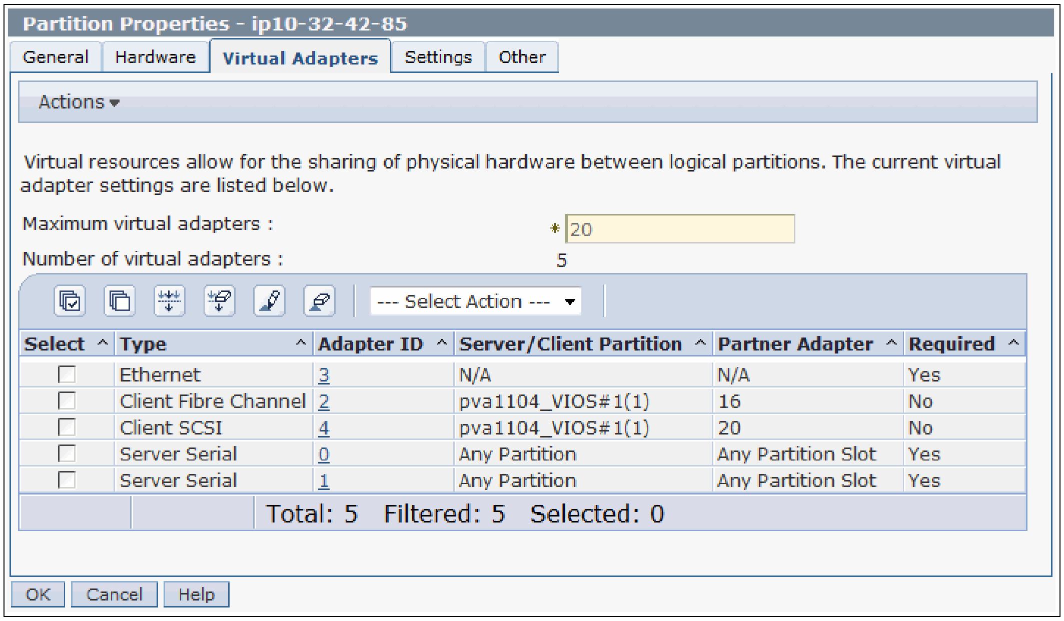

Partition Properties - ip10-32-42-85

Virtual resources allow for the sharing of physical hardware between logical partitions. The current virtual adapter settings are listed below.

Maximum virtual adapters : * 20

Number of virtual adapters : 5

Select	Type	Adapter ID	Server/Client Partition	Partner Adapter	Required
☐	Ethernet	3	N/A	N/A	Yes
☐	Client Fibre Channel	2	pva1104_VIOS#1(1)	16	No
☐	Client SCSI	4	pva1104_VIOS#1(1)	20	No
☐	Server Serial	0	Any Partition	Any Partition Slot	Yes
☐	Server Serial	1	Any Partition	Any Partition Slot	Yes

Total: 5 Filtered: 5 Selected: 0

Figure 5-109 HMC: Virtual server properties

5.7.9 Best practices

Listed in this section are the best practices for IBM Systems Director VMControl:

- Before you start working with VMControl, ensure that your PowerVM environment is ready.
- Before relocating a server, validate from the HMC if that server is able to be relocated.
- When creating virtual farms, ensure that you check *live relocation* to leverage this functionality.
- When creating a server systems pool, select only the hosts that are capable of live virtual server relocation.
- If you are planning to use SCS to capture a virtual appliance, keep in mind that the source must be powered off.

6

Active Energy Manager

This chapter describes the energy monitoring and management features that are offered by IBM Systems Director Active Energy Manager along with the best practice, which needs to be followed when using the IBM Systems Director Active Energy Manager.

This chapter contains the following topics:

6.1 Active Energy Manager

In this section, we briefly introduce the Active Energy Manager (AEM) plug-in for IBM Systems Director along with supported energy monitoring and management features.

The bulk of this chapter focuses on the best practices and the most important questions to ask before using AEM. We also show how you can manage and monitor the energy usage by using the AEM graphical interface and the command-line interface (CLI).

6.1.1 Terms to know

Before starting to read about IBM Systems Director energy management, it is important that you familiarize yourself with some of the widely used AEM-specific terminologies:

- Managed system

  A system that is being controlled by a given system management application, for example, a system managed by IBM Systems Director.

- Metering device

  A resource, such as a power distribution unit (PDU) or sensor that measures things such as power use and thermal values of other objects.

- Metered device

  A resource, such as a server, that is associated with a metering device.

- Metering interval

  The interval, in minutes, that resources are polled for energy information.

- Input power

  The ac (alternating current) power that is supplied to the server power supply from the external sources such as PDUs.

- Output power

  The dc (direct current) power that is consumed by the components inside the server.

6.1.2 Introduction to Active Energy Manager

Active Energy Manager (AEM) is an advanced plug-in to IBM Systems Director that helps users to monitor the power usage of their IT equipment and facility equipment, as well as manage the power usage of the supported managed resources.

Supported energy-related tasks that you can perform using Active Energy Manager, include:

- Monitoring and trending power consumption as well as thermal data
- Managing power, which includes:
  - Setting power savings options
  - Setting power caps
  - Creating and setting power policies
- Monitoring uninterruptible power supply (UPS) units for current status and battery state, and notify associated resources if there are any changes
- Configuring metering devices, such as PDUs and sensors
- Exporting the collected power and thermal trend data

- Viewing events that are related to energy management
- Calculating energy cost
- Calculating estimated energy savings for selected IBM Power Systems servers
- Setting thresholds on power and thermal monitors
- Monitoring of power and cooling equipment that affect the IT resources

IBM Systems Director Active Energy Manager also integrates with the IBM Tivoli Monitoring for Energy Management Agent to provide you with the capability to monitor Active Energy Manager devices from IBM Tivoli Enterprise Portal workspace.

6.1.3 AEM features: No-charge versus charged

The IBM Systems Director Active Energy Manager plug-in offers a few no-charge features and chargeable features. Table 6-1 gives information about what features are available at no charge versus charged.

Table 6-1 No charge-based versus fee-based Active Energy Manager features

Active Energy Manager feature	Is the feature available at no charge?
Monitoring power consumption of IBM servers	Yes
Monitoring inlet and outlet temperatures of IBM servers	
Monitoring power consumption of non-IBM servers and traditional IBM servers using metering devices such as intelligent power distribution units (iPDUs)	
Facility equipment monitoring	
Configuring metering and cooling devices	
Energy-related event monitoring	
Creating custom monitors and thresholds on energy parameters	
Setting power caps and power savings	No
Applying power policies	
Check UPS status and notify associated IBM Power Systems servers	

6.2 Requirements for installing AEM

Starting with version 4.4, the Active Energy Manager full edition with a 90-day evaluation period for power management features is installed along with the installation of IBM Systems Director server.

However, AEM is in a deactivated state by default. You should activate Active Energy Manager before the AEM features can be used. The evaluation period begins after you activate Active Energy Manager.

Note: The 90-day evaluation period begins to count down after activating AEM for the first time. It continues to run even if you deactivate Active Energy Manager thereafter.

Perform the following steps to activate Active Energy Manager:

1. Activate the Active Energy Manager plug-in by using either of the following two methods:
   - Activate by using the IBM Systems Director web interface
     i. Log in to the IBM Systems Director server, go to the home page, and click the **Plug-ins** tab.
     ii. In the “Additional Plug-ins to activate” section, click **Activate now** under the Active Energy Manager heading.
   - From the IBM Systems Director server command-line prompt, run the following command:

     ```
     smcli activatemgrs "Active Energy Manager"
     ```
2. After the activation has completed, restart the IBM Systems Director server by running the following commands, and monitor the status:

   ```
   smstop;smstart;smstatus -r
   ```
   - When the IBM Systems Director server status is “Active”, run the following command to ensure that IBM Systems Director Active Energy Manager is activated:

     ```
     smcli lsmgrs
     ```
   - On the IBM Systems Director web console, go to the **Home** → **Plug-ins** page where you see that the Active Energy Manager is activated with the evaluation period, as shown in Figure 6-1.

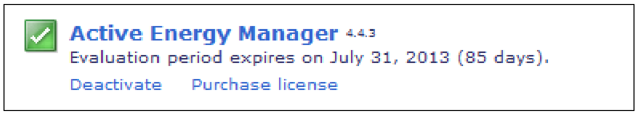
Active Energy Manager 4.4.3

Evaluation period expires on July 31, 2013 (85 days).

Figure 6-1 Active Energy Manager plug-in activated

Visit the following information center link to understand more about the performance and scalability considerations for AEM:

http://pic.dhe.ibm.com/infocenter/director/pubs/index.jsp?topic=%2Fcom.ibm.director.aem.helps.doc%2Ffrb0_t_performance_considerations.html

6.2.1 Learn which devices are supported for monitoring

IBM Systems Director Active Energy Manager can monitor the power usage as well as receive power and cooling-related events on data center equipment:

- IT equipment, includes
  - Servers
  - Network
  - Storage
- Facility equipment, includes
  - IBM power distribution units (PDU+).
  - Non IBM facility equipment such as UPSs, PDUs, and computer room air conditioners, on which are supplied by various facility equipment vendors.

In addition to doing power monitoring, AEM allows you to create logical relationships between the facility equipment such as cooling units, UPSs, PDUs, and the IT servers, which allow the event flow from the facility equipment to the associated IT servers. For more information and directions, see section 6.4.5, “Configuring metering and cooling devices” on page 363.

For more information about supported firmware and software version requirements, see the Supported hardware information center at the following link:

http://pic.dhe.ibm.com/infocenter/director/pubs/index.jsp?topic=%2Fcom.ibm.director.aem.helps.doc%2Ffrb0_r_supported_hardware.html

6.2.2 Learn which devices are supported for power management

IBM Systems Director AEM can be used to perform power management such as power saving and power capping on the following servers:

- IBM System x
- IBM BladeCenter
- IBM Power Systems
- IBM System z

Not all of the supported servers have power save and power cap capabilities. See section 6.5, “Management features that are available” on page 364 to get more information about the management features.

AEM can be used to turn on, off, or reboot individual power outlets on the following supported PDUs:

- Eaton PowerWare Switched PDUs
- Raritan Dominion PX PDUs
- APC Switched Rack PDUs
- Server Tech Switched or POPS PDUs
- Some R-Series Geist PDUs

For more information about supported firmware and software version requirements, see the information center:

http://pic.dhe.ibm.com/infocenter/director/pubs/index.jsp?topic=%2Fcom.ibm.director.aem.helps.doc%2Ffrb0_r_supported_hardware.html

6.3 Learn how to start using Active Energy Manager

The IBM Systems Director AEM communicates to the service processor or the platform managers to retrieve and manage the power and thermal usage. Figure 6-2 on page 350 shows the components with which AEM communicates to retrieve the required data and perform power management.

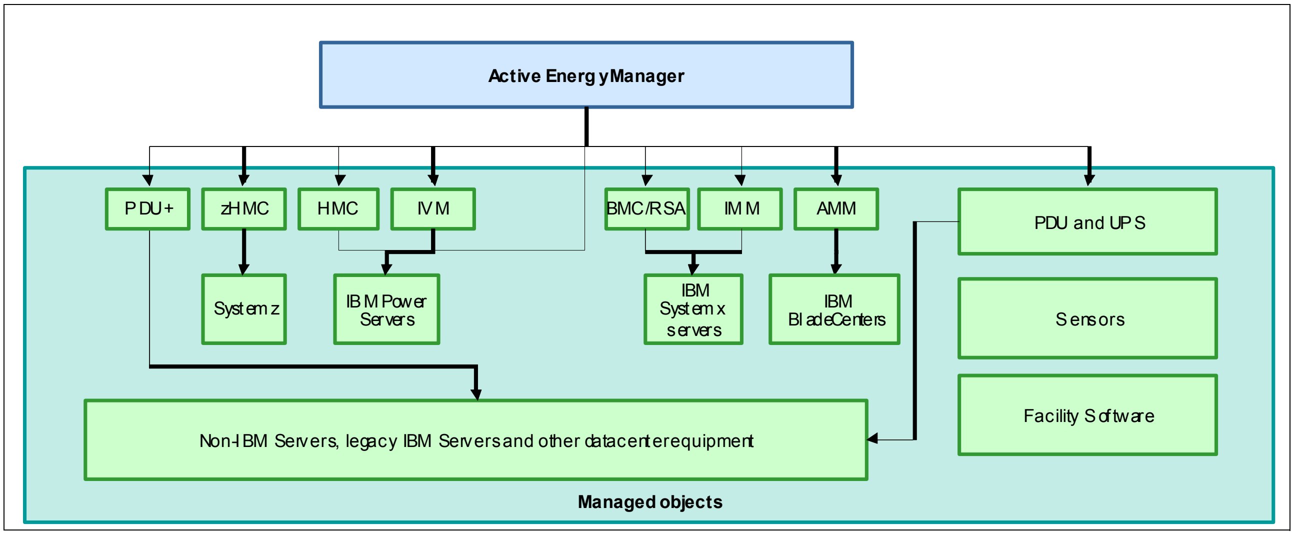

Figure 6-2 IBM Systems Director Active Energy Manager management domain

> **Note:** AEM does not need any agents to be installed on the resources that it needs to power monitor and power manage.

For AEM to start the power monitoring and management on supported endpoints, the endpoints must be first discovered and unlocked in the IBM Systems Director server. Table 6-2 shows the endpoint types that need to be discovered and the discovery method to enable for energy monitoring and management.

Table 6-2 AEM endpoint types and methods of discovery

Managed resource type	Managed by	Resource type to select while discovering	Discovery method [Basic/Advanced]
Power distribution units (PDUs) or uninterruptible power supplies (UPSs)	PDU/UPS itself	Power unit	Either
IBM Power Systems	HMC/IVM/FSP	Server	
IBM System x servers	IMM/BMC/RSA II	Server	
IBM BladeCenter chassis and its components	AMM/MM	IBM BladeCenter chassis	

Managed resource type	Managed by	Resource type to select while discovering	Discovery method [Basic/Advanced]
IBM System z servers and IBM System z with BladeCenter Extensions	zHMC	Server	Advanced
Sensor devices	SynapSense SNMP agent	Generic system	
	► Sensatronics, 1-Wire ► iButtonLink ► Arch Rock sensor networks	Operating system	Either
	Rittal sensors and power units	Power unit	
Facility Software	► Emerson-Liebert SiteScan instance ► Eaton Power Xpert Reporting Database server ► APC InfraStruXure Central server	Operating system	

The following devices can also be power monitored by AEM if they are connected to any of the supported PDU devices:

- Storage devices
- Switch devices
- Non IBM hardware
- Legacy IBM hardware
- Other data center equipment

As of AEM version 4.4.2, energy monitoring and management of IBM Flex nodes are not supported.

For more information about discovering managed resources for energy monitoring and management, see the following site:

http://pic.dhe.ibm.com/infocenter/director/pubs/index.jsp?topic=%2Fcom.ibm.director.aem.helps.doc%2Ffrb0_t_adding_managed_objects.html

IBM Systems Director Active Energy Manager communicates to the managed endpoints during the polling interval, which is set to default values initially, to get the energy-related values. Table 6-3 on page 352 has the default polling intervals for the managed device types.

Table 6-3 Polling intervals

Managed resource type	Metering interval supported by the managed resource	AEM default global polling interval
PDUs or UPSs	N/A	5 minutes
IBM Power Systems	N/A	5 minutes
IBM System x servers	30 seconds (IMMv2)	5 minutes
IBM BladeCenter chassis and its components	10 minutes	10 minutes
SynapSense sensors	5 minutes	5 minutes
Sensatronics	N/A	5 minutes
1-Wire	N/A	5 minutes
iButtonLink	N/A	5 minutes
Arch Rock sensor networks	N/A	5 minutes
Rittal sensors and power units	N/A	5 minutes
Eaton Power Xpert Reporting Database server	60 minutes	15 minutes
Emerson-Liebert SiteScan instance	N/A	5 minutes
APC InfraStruXure Central server	N/A*	5 minutes

> **Best practice:** It is recommended not to set the AEM polling interval to a lesser value than the polling interval supported by the managed resource.

6.3.1 Accessing Active Energy Manager resources

In IBM Systems Director server, there are two ways to access managed resources, which also apply to AEM:

- Steps to access managed resources using the resource explorer are as follows:
  a. On the IBM Systems Director web console, go to **Resource Explorer** where you see the default groups and any custom created groups. When you activate AEM, the Active Energy Manager Groups group gets added to the Resource Explorer page, as shown in Figure 6-3.

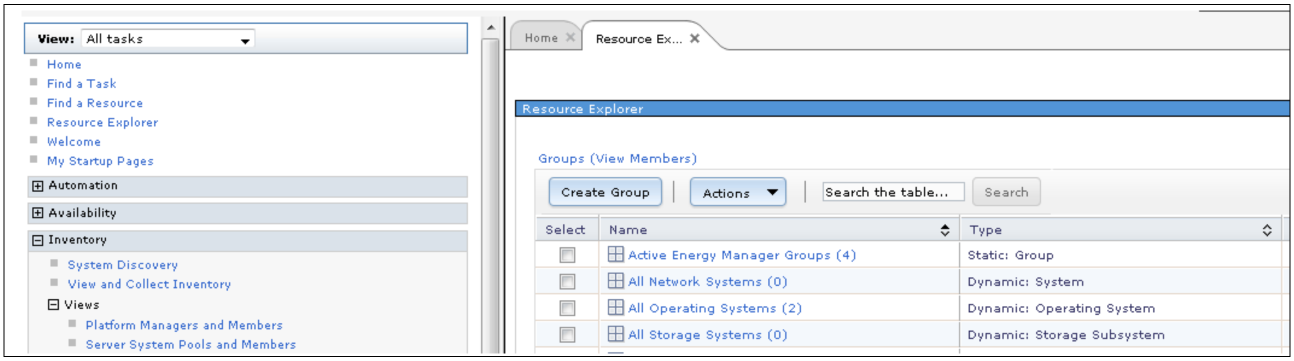

Select	Name	Type
□	Active Energy Manager Groups (4)	Static: Group
□	All Network Systems (0)	Dynamic: System
□	All Operating Systems (2)	Dynamic: Operating System
□	All Storage Systems (0)	Dynamic: Storage Subsystem

Figure 6-3 Resource Explorer page

b. Click the **Active Energy Manager Groups** group. Four default subgroups are displayed, as shown in Figure 6-4.

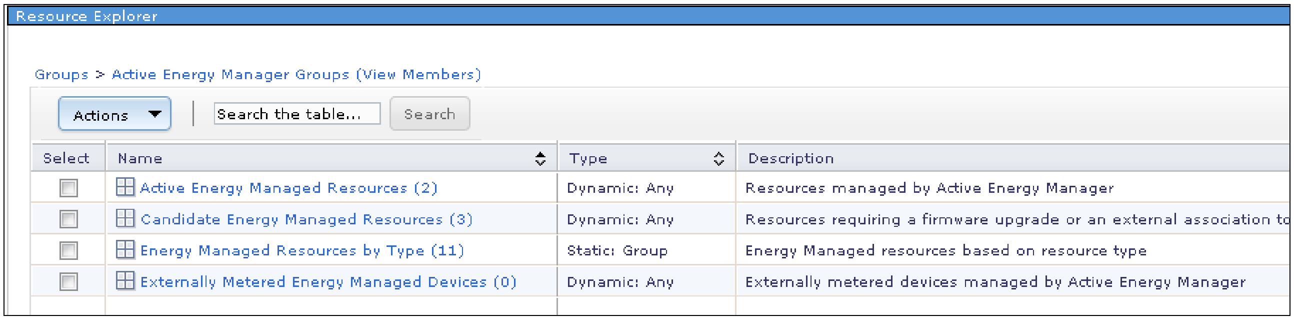

Select	Name	Type	Description
□	Active Energy Managed Resources (2)	Dynamic: Any	Resources managed by Active Energy Manager
□	Candidate Energy Managed Resources (3)	Dynamic: Any	Resources requiring a firmware upgrade or an external association to
□	Energy Managed Resources by Type (11)	Static: Group	Energy Managed resources based on resource type
□	Externally Metered Energy Managed Devices (0)	Dynamic: Any	Externally metered devices managed by Active Energy Manager

Figure 6-4 Active Energy Manager Groups view

The subgroups are defined as follows:

- Active Energy Managed Resources: This subgroup contains the managed resources that are currently being monitored by AEM.
- Candidate Energy Managed Resources: This subgroup contains the managed resources that require a firmware upgrade or association with external metering devices, such has PDUs, to enable AEM to monitor them.
- Energy Managed Resources by Type: This subgroup contains the managed resources by type.
- Externally Metered Energy Managed Devices: This subgroup contains metering devices such as PDUs that are currently being monitored by AEM.

c. Click **Active Energy Managed Resources** to take you to the page where you can access all the energy managed resources, as shown in Figure 6-5.

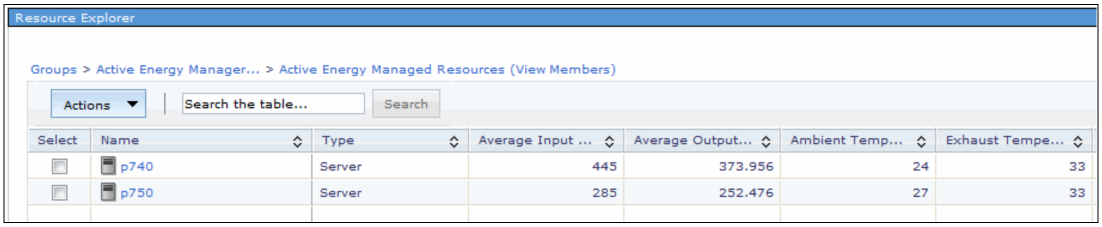
Resource Explorer
Groups > Active Energy Manager... > Active Energy Managed Resources (View Members)
Actions
Search the table...
Search

Select	Name	Type	Average Input ...	Average Output...	Ambient Temp...	Exhaust Tempe...
□	p740	Server	445	373.956	24	33
□	p750	Server	285	252.476	27	33

Figure 6-5 All energy managed resources

d. Click **Energy Managed Resources by Type** and the energy managed resources by type are shown, as in Figure 6-6.

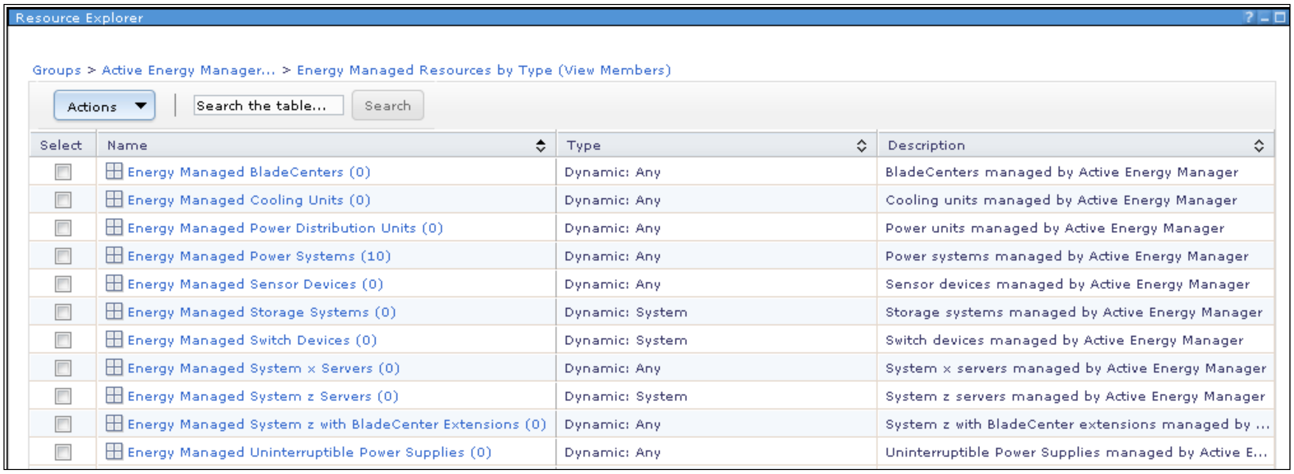

Groups > Active Energy Manager... > Energy Managed Resources by Type (View Members)

Select	Name	Type	Description
□	Energy Managed BladeCenters (0)	Dynamic: Any	BladeCenters managed by Active Energy Manager
□	Energy Managed Cooling Units (0)	Dynamic: Any	Cooling units managed by Active Energy Manager
□	Energy Managed Power Distribution Units (0)	Dynamic: Any	Power units managed by Active Energy Manager
□	Energy Managed Power Systems (10)	Dynamic: Any	Power systems managed by Active Energy Manager
□	Energy Managed Sensor Devices (0)	Dynamic: Any	Sensor devices managed by Active Energy Manager
□	Energy Managed Storage Systems (0)	Dynamic: System	Storage systems managed by Active Energy Manager
□	Energy Managed Switch Devices (0)	Dynamic: System	Switch devices managed by Active Energy Manager
□	Energy Managed System x Servers (0)	Dynamic: Any	System x servers managed by Active Energy Manager
□	Energy Managed System z Servers (0)	Dynamic: System	System z servers managed by Active Energy Manager
□	Energy Managed System z with BladeCenter Extensions (0)	Dynamic: Any	System z with BladeCenter extensions managed by ...
□	Energy Managed Uninterruptible Power Supplies (0)	Dynamic: Any	Uninterruptible Power Supplies managed by Active E...

Figure 6-6 Energy managed resources by type

- The other way in which managed resources can be accessed is through the plug-ins summary page. On the IBM Systems Director web console, go to the **Home** → **Plug-ins** page and click **Active Energy Manager**, which takes you to the AEM summary page, as shown in Figure 6-7 on page 355.

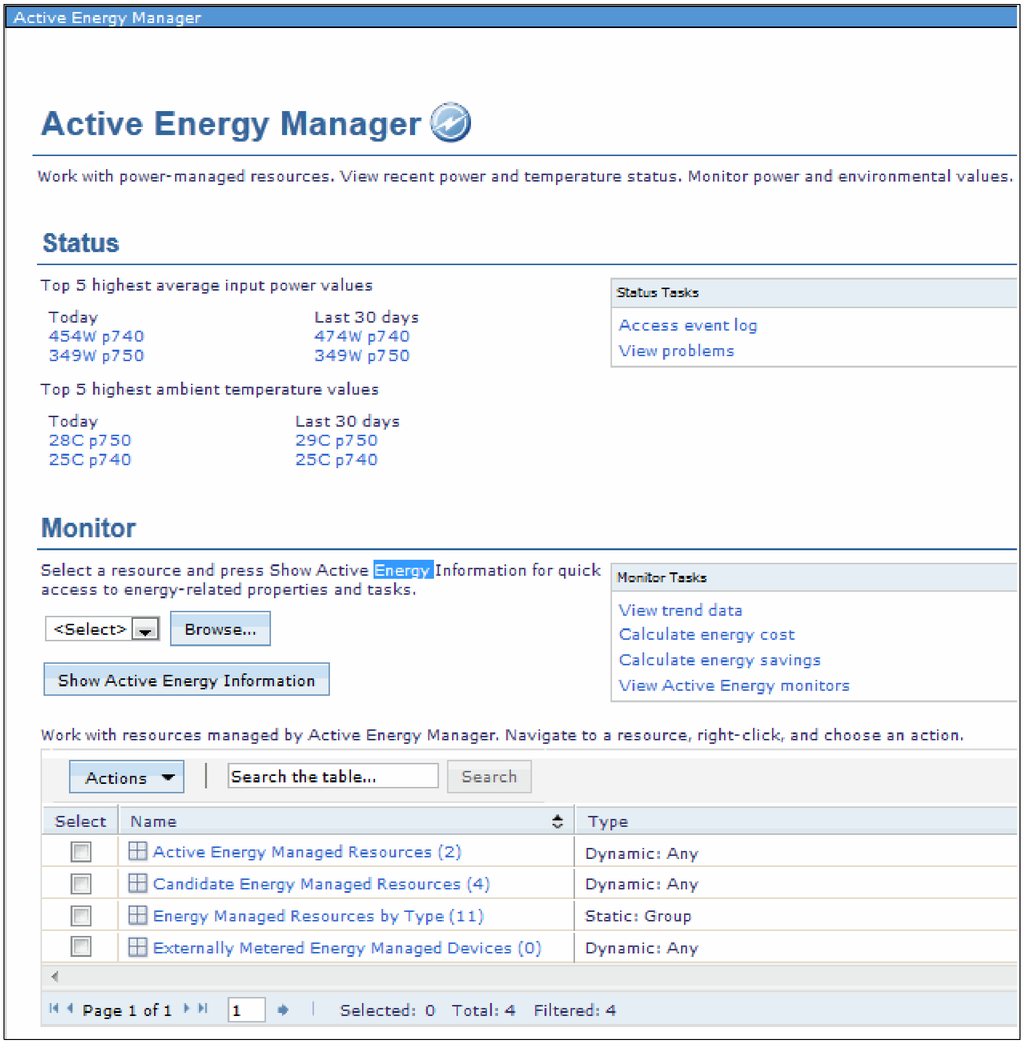

Active Energy Manager

Work with power-managed resources. View recent power and temperature status. Monitor power and environmental values.

Status

Top 5 highest average input power values

Today	Last 30 days
454W p740	474W p740
349W p750	349W p750

Top 5 highest ambient temperature values

Today	Last 30 days
28C p750	29C p750
25C p740	25C p740

Status Tasks

- Access event log
- View problems

Monitor

Select a resource and press Show Active Energy Information for quick access to energy-related properties and tasks.

Monitor Tasks

- View trend data
- Calculate energy cost
- Calculate energy savings
- View Active Energy monitors

Work with resources managed by Active Energy Manager. Navigate to a resource, right-click, and choose an action.

Select	Name	Type
□	Active Energy Managed Resources (2)	Dynamic: Any
□	Candidate Energy Managed Resources (4)	Dynamic: Any
□	Energy Managed Resources by Type (11)	Static: Group
□	Externally Metered Energy Managed Devices (0)	Dynamic: Any

Page 1 of 1 1 Selected: 0 Total: 4 Filtered: 4

Figure 6-7 AEM summary page

Under the "Monitor" section of the summary page, you find the same AEM groups that were present on the Resource Explorer page. Go to the managed resources by clicking the required group.

For more information about the AEM summary page, see the following information center site:

`http://pic.dhe.ibm.com/infocenter/director/pubs/index.jsp?topic=%2Fcom.ibm.director.aem.helps.doc%2Ffrb0_t_aem_summary_page.html`

Further in this chapter, we demonstrate AEM tasks on the managed energy resources by accessing them using the first method that is listed in section 6.3.1, "Accessing Active Energy Manager resources" on page 352.

6.4 Monitoring features that are available

After you discovered the managed resources in AEM, the following operations can be performed that are part of the monitoring feature:

- Viewing AEM-managed resources and trend data
- Calculating energy cost and energy savings
- Viewing energy events and energy properties
- Viewing supported monitors and setting thresholds
- Configuring metering and cooling devices
- Facility Software integration
- Automatic notification upon UPS status change

6.4.1 Viewing trend data

Perform the following steps to view managed resources and trend data:

1. Log in to the IBM Systems Director web console and go to **Resource Explorer** → **Active Energy Manager Groups (View Members)** → **Energy Managed Resources by Type (View Members)**. You are able to see the types of resources that are supported along with the number indicating how many of those resources are already discovered by AEM, as shown in Figure 6-8.

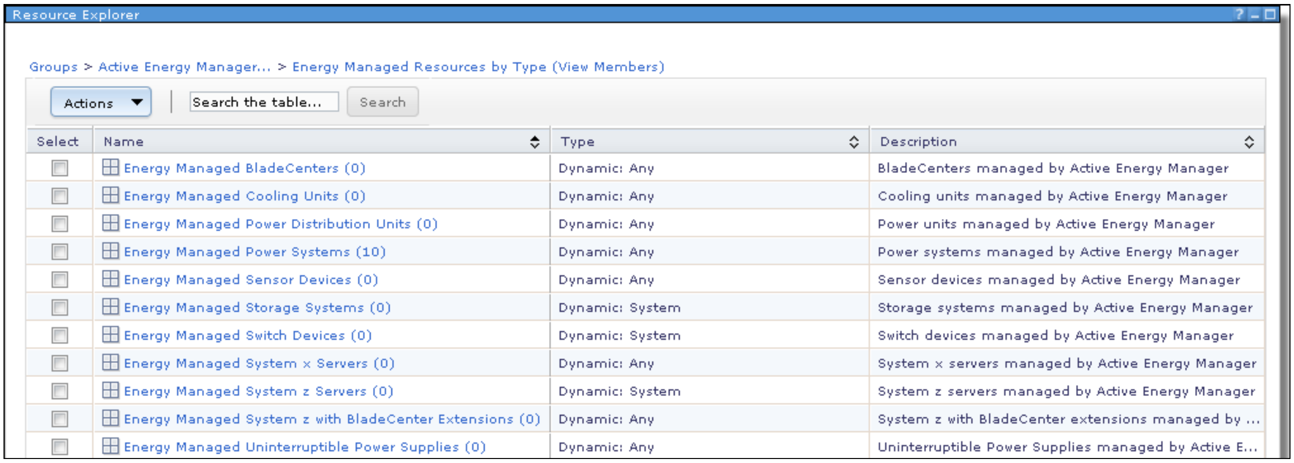
Resource Explorer
Groups > Active Energy Manager... > Energy Managed Resources by Type (View Members)
Actions
Search the table...
Search

Select	Name	Type	Description
□	Energy Managed BladeCenters (0)	Dynamic: Any	BladeCenters managed by Active Energy Manager
□	Energy Managed Cooling Units (0)	Dynamic: Any	Cooling units managed by Active Energy Manager
□	Energy Managed Power Distribution Units (0)	Dynamic: Any	Power units managed by Active Energy Manager
□	Energy Managed Power Systems (10)	Dynamic: Any	Power systems managed by Active Energy Manager
□	Energy Managed Sensor Devices (0)	Dynamic: Any	Sensor devices managed by Active Energy Manager
□	Energy Managed Storage Systems (0)	Dynamic: System	Storage systems managed by Active Energy Manager
□	Energy Managed Switch Devices (0)	Dynamic: System	Switch devices managed by Active Energy Manager
□	Energy Managed System x Servers (0)	Dynamic: Any	System x servers managed by Active Energy Manager
□	Energy Managed System z Servers (0)	Dynamic: System	System z servers managed by Active Energy Manager
□	Energy Managed System z with BladeCenter Extensions (0)	Dynamic: Any	System z with BladeCenter extensions managed by ...
□	Energy Managed Uninterruptible Power Supplies (0)	Dynamic: Any	Uninterruptible Power Supplies managed by Active E...

Figure 6-8 AEM resources grouped by type

2. Click any of the displayed resource type groups to see the group members (Figure 6-9) that show the “Energy Managed Power Systems” group members.

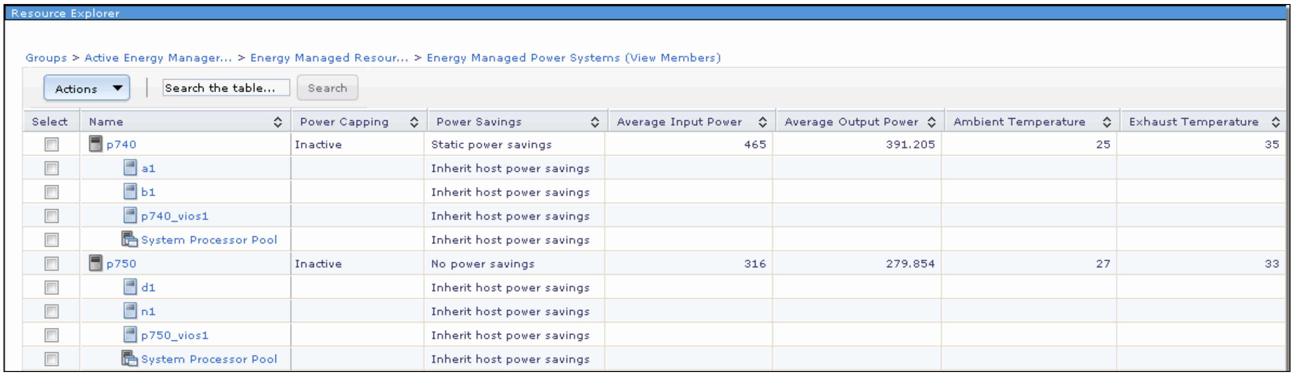
Resource Explorer
Groups > Active Energy Manager... > Energy Managed Resour... > Energy Managed Power Systems (View Members)
Actions
Search the table...
Search

Select	Name	Power Capping	Power Savings	Average Input Power	Average Output Power	Ambient Temperature	Exhaust Temperature
□	p740	Inactive	Static power savings	465	391.205	25	35
□	a1		Inherit host power savings				
□	b1		Inherit host power savings				
□	p740_vios1		Inherit host power savings				
□	System Processor Pool		Inherit host power savings				
□	p750	Inactive	No power savings	316	279.854	27	33
□	d1		Inherit host power savings				
□	n1		Inherit host power savings				
□	p750_vios1		Inherit host power savings				
□	System Processor Pool		Inherit host power savings				

Figure 6-9 AEM-managed IBM Power Systems server resource group

3. To view the collected trend data, right-click the managed resource and select **Energy** → **Trend Data**, as highlighted in Figure 6-10.

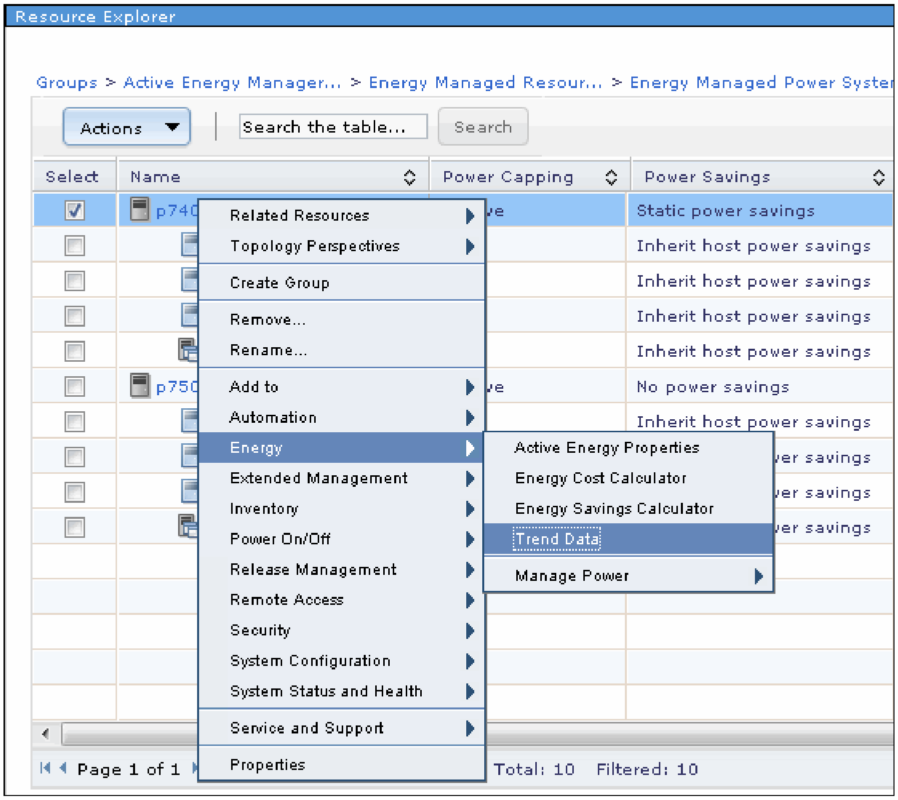

Select	Name	Power Capping	Power Savings
☑	p740	ve	Static power savings
□			Inherit host power savings
□			Inherit host power savings
□			Inherit host power savings
□			Inherit host power savings
□	p750	ve	No power savings
□			Inherit host power savings
□			ver savings
□			ver savings
□			ver savings

Figure 6-10 Viewing trend data

4. The Trend Data page is displayed as shown in Figure 6-11, where you can see the power and thermal trend data for the selected AEM-managed resource that is based on the time range. In this example, we selected an IBM Power 7 processor-based server, which would also display information about the power management capabilities along with the energy events reported for the server resource.

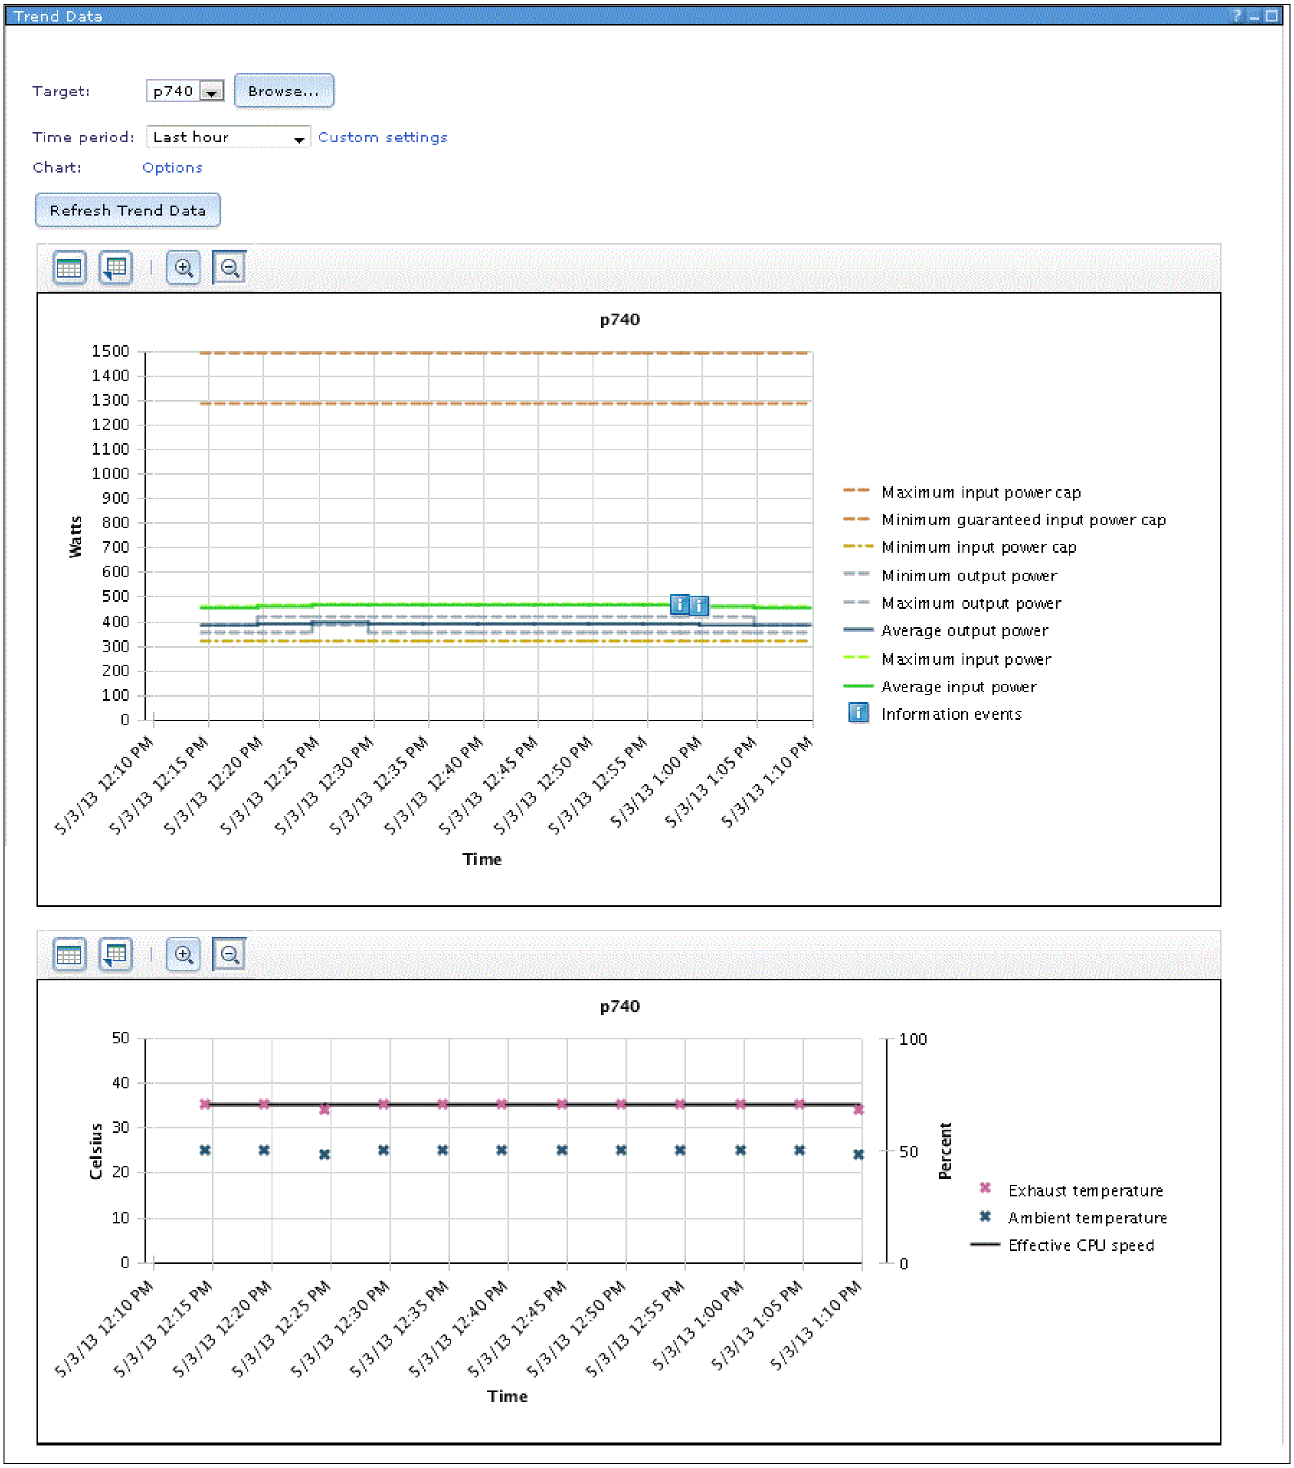

Trend Data

Target: p740 Browse...

Time period: Last hour Custom settings

Chart: Options

Refresh Trend Data

Figure 6-11 AEM Trend Data page

Power and thermal trend data can also be viewed in table format and can be exported to a comma-separated values (CSV) file.

Active Energy Manager runs compression on the trend data that is older than seven days to reduce the amount of storage space occupied. For the collected power and environmental values older than seven days, each night at midnight, this data is compressed to one-hour average values.

6.4.2 Calculating energy cost and energy savings

With AEM, you can calculate the cost that is incurred for running a managed resource over a period of time. To get more accurate cost readings, you need to customize the AEM settings before calculating the energy cost for any managed resources. Perform the following steps to configure the energy cost settings:

1. On the IBM Systems Director web console, go to **Energy** → **Active Energy Manager**. On the Active Energy Manager summary page, click **Settings**, as shown Figure 6-12.

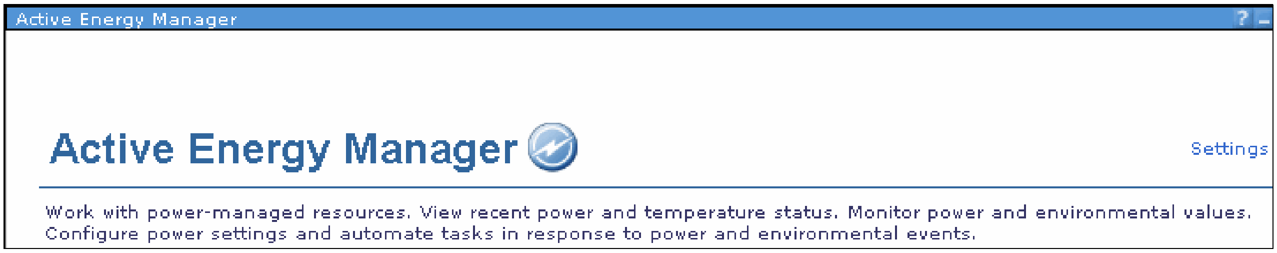

Work with power-managed resources. View recent power and temperature status. Monitor power and environmental values. Configure power settings and automate tasks in response to power and environmental events.

Figure 6-12 AEM global settings

2. Input the custom values for energy cost per kilowatt hour, cooling rate multiplier, and currency type in the AEM settings page, which is shown in Figure 6-13.

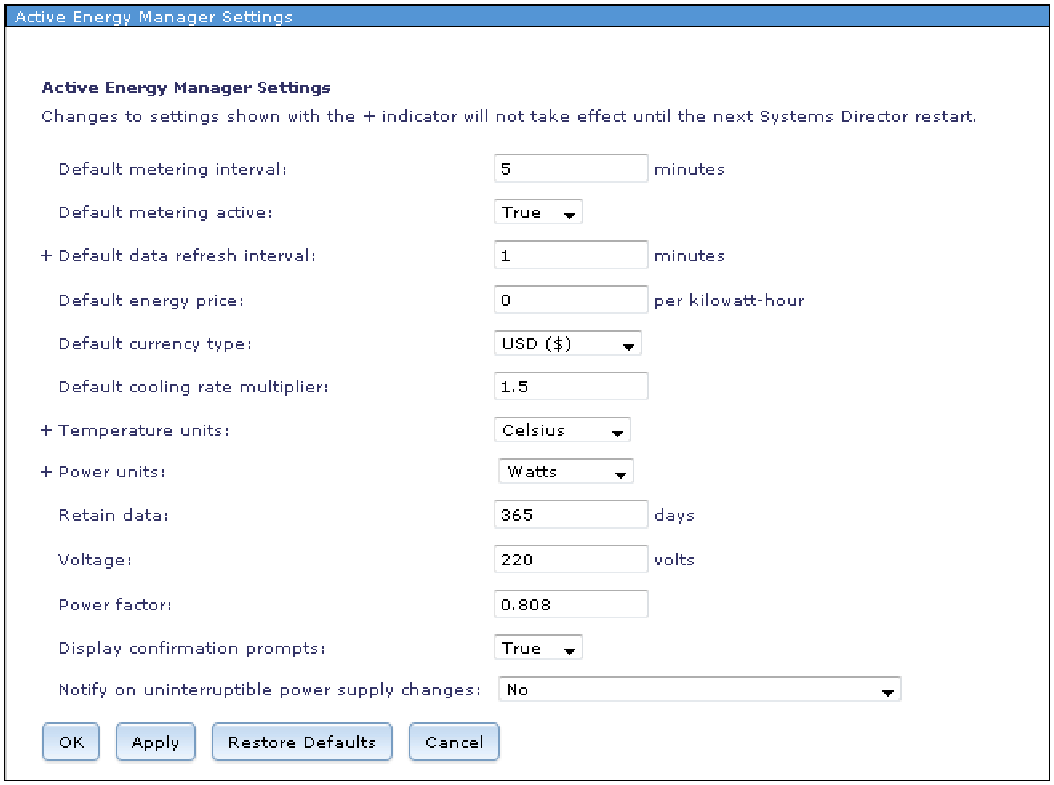

Active Energy Manager Settings

Changes to settings shown with the + indicator will not take effect until the next Systems Director restart.

Default metering interval: 5 minutes

Default metering active: True

+ Default data refresh interval: 1 minutes

Default energy price: 0 per kilowatt-hour

Default currency type: USD ($)

Default cooling rate multiplier: 1.5

+ Temperature units: Celsius

+ Power units: Watts

Retain data: 365 days

Voltage: 220 volts

Power factor: 0.808

Display confirmation prompts: True

Notify on uninterruptible power supply changes: No

Figure 6-13 AEM global settings

3. Click **OK**, which applies the settings. Exit the Settings page.

These settings when saved would be the global settings across the managed resources of AEM. If you are monitoring resources across the geographical regions under the same AEM instance, you can also set energy cost per kilowatt hour, the cooling rate multiplier, and currency type per managed endpoint. The following equation explains how the cost would be calculated by AEM:

```
Total cost = (metered power x time x energy price) + (metered power x time x
energy price x cooling rate multiplier)
```

See the following site to get steps to define cost settings per managed endpoint:

http://pic.dhe.ibm.com/infocenter/director/pubs/index.jsp?topic=%2Fcom.ibm.director.aem.helps.doc%2Ffrb0_t_calculating_energy_cost.html

AEM also lets users calculate the energy savings, both future and past, derived due to the power savings. This feature is currently limited to support these IBM Power Systems servers:

- IBM Power 750 (8233-E8B)
- IBM Power 755 (8236-H8B)

For other supported servers, this feature shows the percentage of time that the Power Savings mode was enabled.

For more information about energy savings calculations, see the following site:

http://pic.dhe.ibm.com/infocenter/director/pubs/index.jsp?topic=%2Fcom.ibm.director.aem.helps.doc%2Ffrb0_t_calculating_energy_savings.html

6.4.3 Viewing energy events and energy properties

When you have the Active Energy Manager plug-in in use, the events that are related to power and thermal monitoring are logged in to the IBM Systems Director event log. You can view the events related to AEM per managed system or across the managed systems. On the IBM Systems Director web console, navigating to **System Status and Health** → **Event Log** and filtering the events by the "Active Energy Manager events" filter enables you to see only the events reported from AEM, as shown in Figure 6-14.

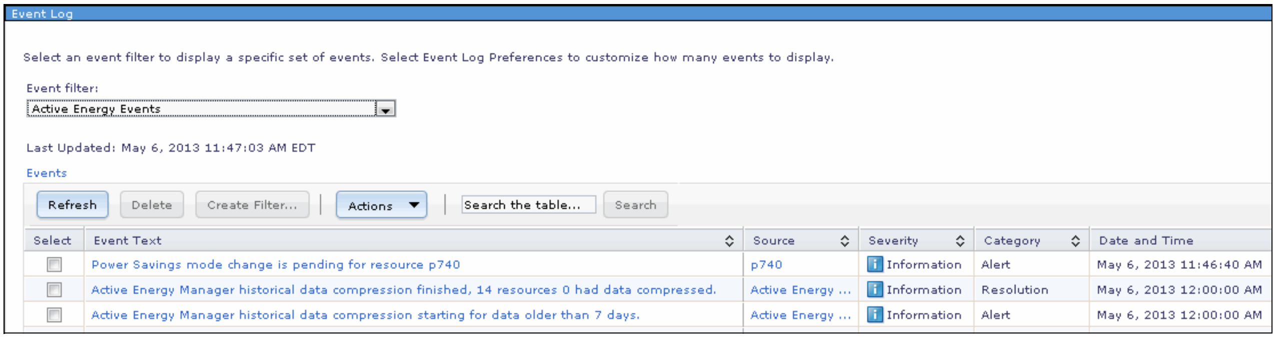

Event Log

Select an event filter to display a specific set of events. Select Event Log Preferences to customize how many events to display.

Event filter: Active Energy Events

Last Updated: May 6, 2013 11:47:03 AM EDT

Events

Select	Event Text	Source	Severity	Category	Date and Time
□	Power Savings mode change is pending for resource p740	p740	Information	Alert	May 6, 2013 11:46:40 AM
□	Active Energy Manager historical data compression finished, 14 resources 0 had data compressed.	Active Energy ...	Information	Resolution	May 6, 2013 12:00:00 AM
□	Active Energy Manager historical data compression starting for data older than 7 days.	Active Energy ...	Information	Alert	May 6, 2013 12:00:00 AM

Figure 6-14 AEM events in the event log of IBM Systems Director server

There are two default event filters created for IBM Systems Director Active Energy Manager, namely:

- Active Energy Critical Events: Displays only those events that are generated by AEM that have a Critical severity
- Active Energy Events: Displays only those events that are generated by AEM

These filters can be used in the event automation plans to carry out any event actions against the reported events. For more information about event automation plans, refer to "Creating an event automation plan" on page 163.

For more information about the type of events that are reported by AEM, see the following information center site:

http://pic.dhe.ibm.com/infocenter/director/pubs/index.jsp?topic=%2Fcom.ibm.director.aem.helps.doc%2Ffrb0_r_events_trend_pane.html

When any devices that are also supported by AEM are discovered in IBM Systems Director, AEM queries the properties that are related to energy parameters and stores them in the device properties under the "Active Energy" section. To access the Active Energy properties on the IBM Systems Director web console, go to **Resource Explorer** → **Groups** → **Active Energy Manager Groups** → **Active Energy Managed Resources** and right-click any managed endpoint and select **Properties** → **Active Energy**.

As an example, Figure 6-15 shows the Active Energy properties of an IBM Power Systems server.

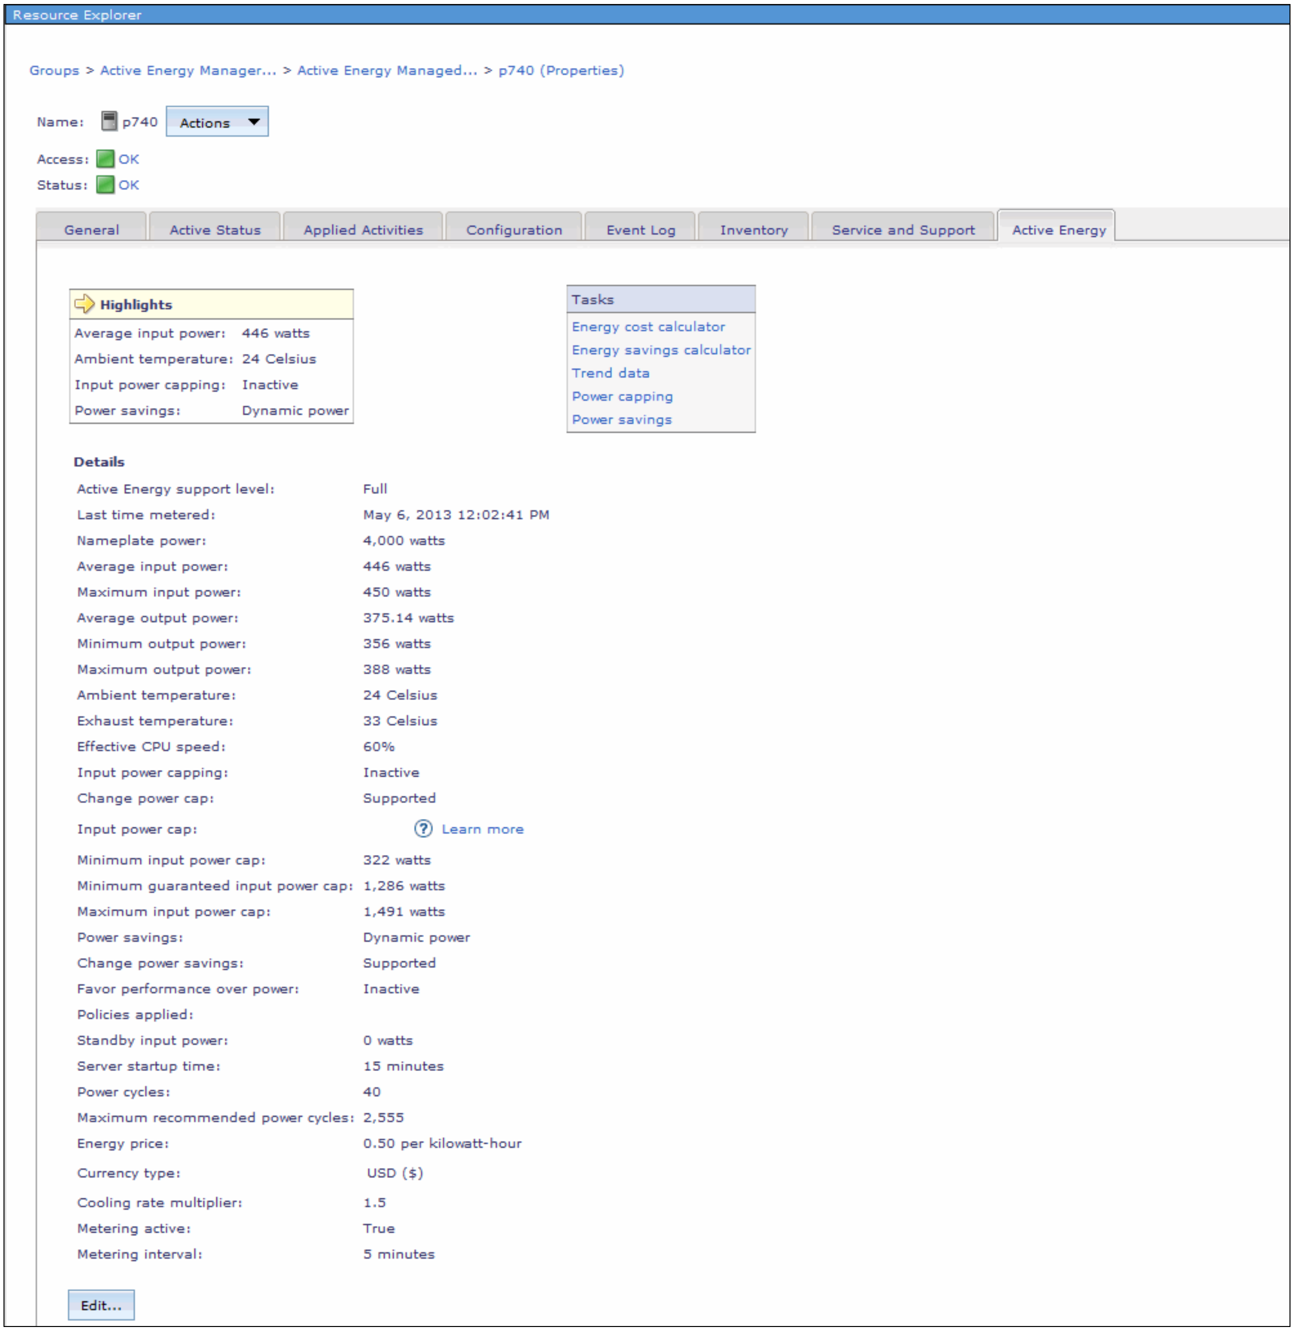

Name: p740

Access: OK

Status: OK

Highlights

Average input power:	446 watts
Ambient temperature:	24 Celsius
Input power capping:	Inactive
Power savings:	Dynamic power

Tasks

- Energy cost calculator
- Energy savings calculator
- Trend data
- Power capping
- Power savings

Details

Active Energy support level:	Full
Last time metered:	May 6, 2013 12:02:41 PM
Nameplate power:	4,000 watts
Average input power:	446 watts
Maximum input power:	450 watts
Average output power:	375.14 watts
Minimum output power:	356 watts
Maximum output power:	388 watts
Ambient temperature:	24 Celsius
Exhaust temperature:	33 Celsius
Effective CPU speed:	60%
Input power capping:	Inactive
Change power cap:	Supported
Input power cap:	? Learn more
Minimum input power cap:	322 watts
Minimum guaranteed input power cap:	1,286 watts
Maximum input power cap:	1,491 watts
Power savings:	Dynamic power
Change power savings:	Supported
Favor performance over power:	Inactive
Policies applied:	
Standby input power:	0 watts
Server startup time:	15 minutes
Power cycles:	40
Maximum recommended power cycles:	2,555
Energy price:	0.50 per kilowatt-hour
Currency type:	USD ($)
Cooling rate multiplier:	1.5
Metering active:	True
Metering interval:	5 minutes

Figure 6-15 AEM properties for an IBM Power Systems server

For more information about the energy properties that are supported for each of the AEM supported managed endpoints, see the following page in the information center:

http://pic.dhe.ibm.com/infocenter/director/pubs/index.jsp?topic=%2Fcom.ibm.director.aem.helps.doc%2Ffrb0_t_viewing_aem_properties.html

6.4.4 Viewing supported monitors and setting thresholds

Active Energy Manager when activated adds “Active Energy” monitors to the set of available monitors under IBM Systems Director server. Following are the two default monitor groups that are created by AEM:

- Active Energy Common Monitors
- Active Energy Monitors

You can set thresholds on the AEM monitors to generate event notifications that can then be filtered to take appropriate actions by using the event automation plan.

6.4.5 Configuring metering and cooling devices

With Active Energy Manager, you can configure metering devices such as PDUs and sensors, and the cooling devices like CRAC devices to associate them with other related resources in the data center.

Configuration of metering and cooling devices enables the following features:

- Monitoring and managing data (power, temperature, humidity, dew point) for resources that are associated with the metering device.
- Events to be generated for the associated resources whenever a severe event is received for the metering device.
- Viewing the data center power flow by using the Active Energy Power perspective.
- Viewing all resources that are cooled by a cooling device and viewing all cooling devices that cool a resource. This also facilitates the generation of events for cooled resources when an associated cooling device experiences a severe event.

See the following information center sites to get steps to configure metering and cooling devices:

- http://pic.dhe.ibm.com/infocenter/director/pubs/index.jsp?topic=%2Fcom.ibm.director.aem.helps.doc%2Ffrb0_t_configure_metering_device.html
- http://pic.dhe.ibm.com/infocenter/director/pubs/index.jsp?topic=%2Fcom.ibm.director.aem.helps.doc%2Ffrb0_t_configure_cooling_device.html

6.4.6 Facility Software integration

Active Energy Manager extends the scope of energy monitoring and management by integrating with various facility endpoints to enable a more complete view of energy consumption within the data center. Integrating facility managers with IBM Active Energy Manager enables IT administrators to be alerted about issues with facilities equipment such as overheating, low battery power-on uninterruptible power supplies, or other conditions that might keep IT equipment in a data center from running properly.

Currently, AEM can integrate and monitor the facility equipment when you have the supported Facility Management Software that is installed in your data center as:

- APC InfraStruXure Central server at version 5.1 or 6.0
- Eaton Power Xpert Reporting versions V1.0 and V2.0
- Emerson Liebert SiteScan (SSWEB) Version 3.0 or Version 4.0

To get step-by-step procedures on the discovery and monitoring of facility equipment managed by the Facility Software Managers listed above, see Chapter 10 of the following IBM Redbooks publication:

http://www.redbooks.ibm.com/abstracts/sg247780.html

6.4.7 Automatic notification upon UPS status change

Active Energy Manager supports metering the current status and battery state of the following uninterruptible power supply models and notifies associated resources of when the status changes. This is a licensed feature:

- IBM branded tower and rack uninterruptible power supplies
- Eaton PowerWare uninterruptible power supplies that utilize a Web/SNMP card

The associated resources include IBM POWER7 models, with firmware level eFW7.4 or greater, which support taking action upon notification of a change in the status of an associated uninterruptible power supply. For example, with IBM POWER7 systems in the case of a utility failure, each of the Power Systems' software partitions starts a timer. If the utility power is not restored before the timer expires, each partition shuts down gracefully. When all partitions are shut down, the server powers off.

The requirements below should be taken care for sending automatic UPS status change notifications to the associated IBM Power 7 servers:

- The Configure Metering Device feature must be used to create an association between the uninterruptible power supply and the resource to notify. This can be a direct association between the uninterruptible power supply and the resource or it can be an indirect association via a PDU.
- The Notify on uninterruptible power supply changes property, for the resource to notify, must indicate that notification should be sent.
- The servers to be notified of the uninterruptible power supply status changes must be discovered in IBM Systems Director by using a connection to the Flexible Server Processor (FSP) of the server.

Best practices: When you are working with UPS state notification, consider the following factors:

- For the most timely notification, configure SNMP traps that are generated by the UPS to be sent to the IBM Systems Director server.
- The metering interval for the UPS should be set as short as possible to detect changes sooner. To change a metering interval, go to the properties of the UPS resource and edit the "Metering interval" property value.

6.5 Management features that are available

The following AEM operations are part of the management features:

- Enabling and disabling power cap
- Enabling and disabling power savings
- Work with power policies

6.5.1 Working with power cap

In a normal and traditional scenario, the power budget for the data centers are calculated based on the UL rating or the nameplate power of the equipment that the data center would host in the future.

In the case of servers, the nameplate power or the UL rating is the amount of power calculated considering the following factors:

- The maximum configuration of that server
- All the components, such as CPU, memory, disk, PCI cards, running the highest possible workload under extreme environmental conditions
- Redundancy of the power supply

If the nameplate power is taken as-is for calculating the amount of cooling that is required for the data center, you end up over-sizing the cooling requirements, which results in energy wastage as well as higher costs.

Actual power usage of a server will be far below the nameplate power rating. Trending power usage of a server in Active Energy Manager for a period of time allows you to understand the trend of power consumption of your servers.

When you understand the power trend, the power capping feature allows you to set a cap on the power that is allocated for the servers, thereby releasing the extra amount of power that you had allocated to the servers considering its name plate power value. This can help save on data center infrastructure costs, and then potentially allow more servers to be put into an existing infrastructure with the same power budget.

Depending on the existing configuration, the system firmware calculates the power cap range [minimum input power cap, minimum guaranteed input power cap, and maximum input power cap] between which the actual power cap can be set. AEM retrieves this range from the server firmware and lets users activate the power cap on the managed servers.

The power cap feature is supported on the IBM x86, IBM Power Systems, and IBM System z servers, but not all managed servers support the power cap feature. Refer to section 6.2.2, “Learn which devices are supported for power management” on page 349 to determine the support matrix.

There are two modes in which a power cap can be enforced from AEM, namely:

- Hard power cap

  This is a guaranteed power cap that is a value that is set between the “minimum guaranteed power cap” and “maximum input power cap”. When this option is selected, if the server tries to use more power than the power cap value, the processor speed is throttled to retain the power usage below the power cap. The hard power cap can be set between the minimum guaranteed power cap and maximum power cap value.

- Soft power cap

  In many scenarios, it was observed that the typical power consumption by the servers was below the “minimum guaranteed input power cap” value. In order to give more flexibility to users in allocating the power budget, the soft power cap was introduced, which allowed users to set the power cap value below the “minimum guaranteed input power cap”. The soft power cap can be set between the minimum input power cap and the maximum input power cap. If the servers start to exceed the set soft power cap value, best efforts are made to hold the power consumption under the soft power cap value that is set by throttling the processor speed. If other components such as storage devices and memory continue to consume more power, the soft power cap is lifted and the server is allowed to consume the power that it needs.

To set the power cap on supported managed servers, perform the following steps:

1. On the IBM Systems Director server, go to **Resource explorer** → **Active Energy Manager Groups** → **Active Energy Managed Resources**.
2. Right-click any managed system that supports power savings, and select **Energy** → **Manage Power** → **Power Capping**. Figure 6-16 shows an example where the power capping mode can be selected and is then enforced on the selected managed resources when you click **Save**.

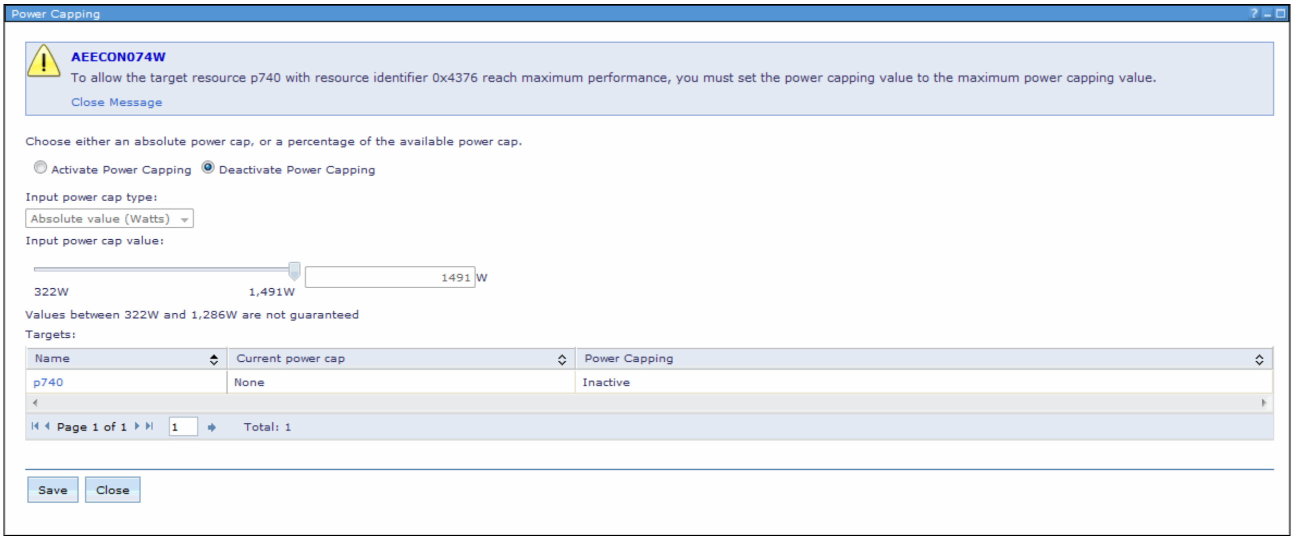

AEECON074W
To allow the target resource p740 with resource identifier 0x4376 reach maximum performance, you must set the power capping value to the maximum power capping value.
Close Message

Choose either an absolute power cap, or a percentage of the available power cap.

○ Activate Power Capping ● Deactivate Power Capping

Input power cap type:
Absolute value (Watts)

Input power cap value:
1491 W
322W 1,491W

Values between 322W and 1,286W are not guaranteed

Targets:

Name	Current power cap	Power Capping
p740	None	Inactive

Page 1 of 1 1 Total: 1

Figure 6-16 Power cap settings

In Figure 6-16, the power cap values can be read as follows:

- Minimum input power cap value: 322W
- Minimum guaranteed input power cap value: 1286W
- Maximum input power cap value: 1491W

It is important to understand that no power capping or processor speed throttling occurs until the power consumption of the server is below the power capping value set. It is a best practice to enable hard power capping unless you have a need for accommodating more servers within the given power budget that is not met by the [Minimum guaranteed input power value - Maximum input power value] range.

Active Energy Manager allows the users to set power cap values on individual servers and also groups of servers.

> **Note:** If hot-swappable components are added to or removed from an IBM System x server while it is powered on, the minimum input power cap and maximum input cap values that are displayed by Active Energy Manager will not be updated until the next reboot of the System x server.

See the following link for steps to enable and disable the power cap on managed resources:

http://pic.dhe.ibm.com/infocenter/director/pubs/index.jsp?topic=%2Fcom.ibm.director.aem.helps.doc%2Ffrb0_t_power_cap.html

6.5.2 Working with power save

Running a server with full capacity, such as processor and memory at 100% nominal frequency all the time independent of the workload utilization, results in an inefficient usage of the power.

Currently, the major component that can be power controlled within a server is the processor and it is said to be one of the major power consumers inside a server. Running a processor at a fixed frequency all the time independent of how many of the processor cycles are really being utilized to run the business workload, results in inefficient power usage.

The power save feature of AEM allows you to use either of the following modes:

- Static Power Save (SPS): Put the servers into low-power mode when the workload utilization is less. In this mode, the jobs executed on the processor can take more time to complete compared to executing the job on the processor running at nominal frequency.
- Dynamic Power Save (DPS): Put the servers into a mode where processor frequency gets varied dynamically depending on the workload utilization. This mode has two submodes, namely:
  - DPS favor performance: Where the performance is not reduced at any time, but at the same time, attempts to save the power are done whenever possible.
  - DPS favor power savings: Where the processor frequency gets reduced by a predetermined value from its nominal frequency, depending on the model of the server, to get constant power savings all the time. Within the remaining processor frequency range, depending on the workload utilization, additional power savings can be achieved.

The power save feature is supported on IBM Power Systems and IBM System z servers, but not all of these servers support all power save modes. Refer to section 6.2.2, “Learn which devices are supported for power management” on page 349 to determine the support matrix.

The power save feature does not apply to System x servers because power savings on System x are provided by DBS on Intel and PowerNow on AMD, which is controlled by the BIOS and controlled thereafter by the operating system.

On IBM Power Systems, AEM can be used to enable and disable the power save mode on the following components:

- Entire Power Systems server
- Physical system processor pool
- Individual partitions

IBM Power Systems servers, certain Power 7 processor-based servers, support enabling and disabling the power save option per logical partition (LPAR) and physical system processor pool. Following are the minimum requirements for performing power savings at the logical partition level and system processor pool level:

- IBM Power Systems server should support partition level and processor pool power savings.
- IBM Systems Director VMControl express edition, at the minimum, should be in activated state.
- Inventory should be collected on the platform managers, such as Hardware Management Console (HMC), and the physical Power servers.

Limitations: As of AEM version 4.4.2:

- AEM shows set power save modes on the LPAR, but not the mode actually in effect. The actual mode in effect can depend on the following states:
  - Host power save state
  - System processor pool power save state
- Power policies do not support partition-level power save option.
- Power monitoring is not supported per logical partition.

Following are the states of power save in which a managed resource can be put in to at any time:

- No power savings: No power savings. The processor runs at high speed.
- Static power savings: Reduces power usage by lowering processor speed.
- Inherit host setting: In this state, the managed resource uses the same power savings as that of the resource that hosts it. For example, when you put a partition or virtual server into this state, it inherits the power save state from the Power Host server on which it is residing.
- Dynamic power savings: Automatically balances power usage and processor performance:
  - Favor power

    or
  - Favor performance
- Component Level Control: Enables the power savings mode of component resources to be set individually. This state is currently available only on the IBM System z server BladeCenter extension.

To access and enable or disable the power savings options that are available for a supported managed system, perform the following steps:

1. On the IBM Systems Director server, go to **Resource explorer** → **Active Energy Manager Groups** → **Active Energy Managed Resources**.
2. Right-click any managed system that supports power savings, and select **Energy** → **Manage Power** → **Power Savings**. Figure 6-17 shows an example of a list of options to choose the power savings mode that you need, where you can enforce on the selected managed resources that are displayed.

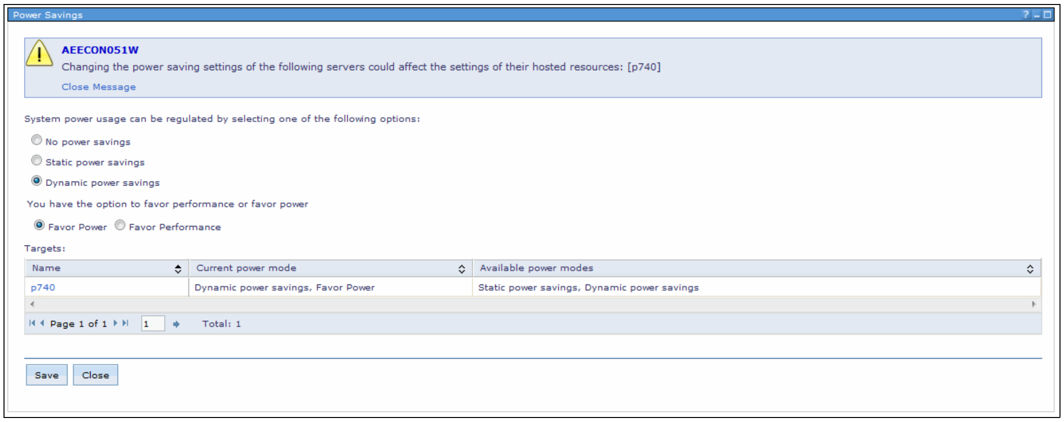

AEECON051W
Changing the power saving settings of the following servers could affect the settings of their hosted resources: [p740]
Close Message

System power usage can be regulated by selecting one of the following options:

○ No power savings
○ Static power savings
● Dynamic power savings

You have the option to favor performance or favor power

● Favor Power ○ Favor Performance

Targets:

Name	Current power mode	Available power modes
p740	Dynamic power savings, Favor Power	Static power savings, Dynamic power savings

Figure 6-17 Power save settings

3. You can choose the right power savings mode dependent on the workload utilization on that managed system and click ‘Save’.

> **Note:** Although the power save enable/disable operation completes immediately after you click ‘Save’, the state is not reflected in Active Energy Manager until the next polling interval.

To get more information about the latest IBM EnergyScale™ features and power saving modes that are supported on IBM Power 7 processor-based servers, see the following link:

`ftp://public.dhe.ibm.com/common/ssi/ecm/en/pow03039usen/POW03039USEN.PDF`

6.5.3 Working with power policies

The power management settings are set using steps that are provided in section 6.5.1, “Working with power cap” on page 365 and 6.5.2, “Working with power save” on page 367. The settings will not be monitored always by AEM, which means that AEM enforces the power management settings once and will not guarantee those settings over a period of time. It can be modified by other users.

Power policies provide a way to continuously monitor and manage the power settings on a managed system or a group of systems. When you have power policies that are applied on managed systems, at every polling interval, AEM checks if the power settings applied by the policy hold good or not. If not, AEM re-enforces the power settings present in the policy on the managed servers/group.

Some resource types like IBM zEnterprise® 196 and zEnterprise BladeCenter Extension support setting a power cap and the power savings mode, but a policy cannot be used to hold the resource in that state.

Two types of power policies exist, namely:

- System policy

  This policy can be either a power save or a power cap setting that can be applied to individual servers or a group of managed servers. When you apply this policy to group of managed servers, the policy is applied individually to each server in that group.

- Group policy

  This policy applies only to a power cap setting, which can be applied only on a managed systems group.

If a policy is in effect at the time the Active Energy Manager license expires, the policy is deactivated and the power management function is deactivated from the system. When you install the Active Energy Manager license, the power savings and power capping functions become active again and the existing power values are used.

Best practices: When applying policies, consider the following factors:

- Do not set system power capping or system power savings policies on both a slot and on the blade in that slot.
- Do not set group capping policies on two groups, where one group contains a slot and the other group contains a blade in that slot.
- Ensure that the resource against which you are applying power cap policy supports the same power cap type (ac or dc).

  To determine if a server supports input or output power capping, view the properties of the server and select the Active Energy tab.
- You can use the IBM Systems Director schedule job feature to schedule policy application and de-application as wanted.
- You can apply policy as an event action in response to an event using event automation plans.

To get step-by-step details about creating, deleting, and editing power policies, see the following information center link:

http://pic.dhe.ibm.com/infocenter/director/pubs/index.jsp?topic=%2Fcom.ibm.director.aem.helps.doc%2Ffrb0_t_work_with_power_policies.html

6.6 Active Energy Manager: smcli references

This section describes the smcli commands for most commonly used energy monitoring and management scenarios:

- List the power and thermal trend data

  Syntax:

  ```
  smcli lstrenddata [-v] [-d delimiter] [-t objectType] [-r timeinterval] [-o
  option] {-a |-f file |-N group_list | -n object_list}
  ```

  Example:

  – Retrieving the past hour power and thermal trend data:

  ```
  smcli lstrenddata -n "p740"
  ```

  Where: “p740” is the display name of the power system

  – Retrieving power and thermal trend data from the specified date:

  ```
  smcli lstrenddata -n “p740” -r "5/6/13"
  ```

  Where:

  - “p740” is the display name of the power system
  - “5/6/13” is the start date for AEM to retrieve the power and thermal data

- Start and stop metering and list all devices that are currently being metered

  Syntax:

  ```
  smcli startcollect -n <display name of the resource>
  smcli stopollect -n <display name of the resource>
  smcli lscollect -a | -n <display name of the resource>
  ```

Example:

- Enable metering for one of the AEM managed resources:
  ```
  smcli startcollect -n "p740"
  ```
  Where: “p740” is the display name of the power system
- Displaying metering status and interval for all of the AEM managed resource:
  ```
  smcli lscollect -a
  ```

- Enabling and disabling the power cap setting on AEM managed resource

  Syntax:

  ```
  smcli setpcap -p pcap_value [-v] [-d delimiter] [-t objectType] {-f file | -N
  group_list | -n object_list}
  ```

  Example:

  - Setting power cap for one of the AEM managed resource:
    ```
    smcli setpcap -n "p740" -p 645 -T AC
    ```
    Where:
    - “p740” is the display name of the power system
    - “645” is the power cap value
    - “AC”’ is the type of power cap
  - Disabling the power cap on one of the AEM managed resource:
    ```
    smcli setpowercap -n "p740" -p -1
    ```
    Where:
    - “p740” is the display name of the power system
    - “-1” used for disabling power cap feature

- Enabling and disabling power save setting on AEM managed resource

  Syntax:

  ```
  smcli setpsaver -p psaver [-v] [-d delimiter] [-t objectType] {-f file | -N
  group_list | -n object_list}
  ```

  Example:

  - Setting static power save for one of the AEM managed resource:
    ```
    smcli setpsaver -n "p740" -p static
    ```
    Where:
    - “p740” is the display name of the power system
    - “static” is the power save type
  - Setting dynamic power save favor performance for one of the AEM managed resources:
    ```
    smcli setpsaver -n "p740" -p dynamic,favorperformance=1
    ```
    Where:
    - “p740” is the display name of the power system
    - “dynamic,favorperformance=1” is the power save type
  - Disabling power save on one of the AEM managed resource:
    ```
    smcli setpowercap -n "p740" -p 0
    ```
    Where:
    - “p740” is the display name of the power system
    - “0” used for disabling power save feature

- Create, edit, and delete power policies

  Syntax:

  ```
  smcli chpolicy [-v] [-d delimiter] [-r] [-p policy_target]
  ```

  Example:

  - Creating power policy for group capping:
    - smcli chpolicy -p
    - Groupcappolicy,policytype=group,pcap=200,powertype=output

    Where:

    - "Groupcappolicy" is the name provided for the new power policy
    - "group" is the power policy type; other types include "saving" and "capping"
    - "200" is the power cap value to be used
    - "output" is the power cap type, which is dc; other type is "input", which is ac
  - Creating power policy for power saving favoring performance:
    - smcli chpolicy -p
    - powersavepolicy,policytype=saving,psavertype=dynamic,favorperformance=on

    Where:

    - "powersavepolicy" is the name provided for the new power policy
    - "saving" is the power policy type; other types include "group" and "capping"
    - "dynamic" is the power save type, which has "favorperformance" turned on; other types include "off" and "static"
  - To update/edit a power policy to change the power cap value:

    ```
    smcli chpolicy -p Groupcappolicy,pcap=350
    ```

    Where:

    - "Groupcappolicy" is the name of the power policy
    - "350" would be the new power cap value
  - To remove a policy:

    ```
    smcli chpolicy -r -p Groupcappolicy
    ```

    Where: "Groupcappolicy" is the name of the power policy

- Applying power policies to managed AEM resources

  Syntax:

  ```
  smcli setpolicy [-v] [-d delimiter] [-t objectType] [-r] [-p policy_target] {-f
  file | -N group_list | -n object_list}
  ```

  Example:

  - Set power policy on a AEM managed resource:

    ```
    smcli setpolicy -p "powersavepolicy" -n "p750"
    ```

    Where:

    - "powersavepolicy" is the name of the power policy
    - "p750" is the display name of the AEM managed resource

- To remove the power policy from the AEM managed resource:

  ```
  smcli setpolicy -p "powersavepolicy" -r -n "p750"
  ```

  Where:

  - “powersavepolicy” is the name of the power policy
  - “p750” is the display name of the AEM managed resource

See the following link to get all the commands supported by Active Energy Manager:

http://pic.dhe.ibm.com/infocenter/director/pubs/index.jsp?topic=%2Fcom.ibm.director.aem.helps.doc%2Ffrb0_r_ref_commands.html

6.7 Best practices

The following list provides best practices to be followed when using AEM:

- Before activating AEM, plan for the list of things that you need to try and evaluate. The 90-day AEM evaluation period begins to count down after activating AEM for the first time and it continues to run even if you deactivate Active Energy Manager later.
- Plan for power monitoring on IBM Power Systems servers:
  - If you need to monitor I/O drawer power and use a licensed feature of AEM to receive alerts from UPS, discover the FSP of the IBM Power Systems servers in IBM Systems Director server.
  - If you need to power manage partitions of an IBM Power 7 server that is managed by an HMC, IBM Systems Director VMControl must be active and the inventory of the server that contains the partitions must be collected.
- It is recommended not to set the AEM polling interval to a lesser value than the polling interval that is supported by the managed resource.
- For timely notification, configure SNMP traps that are generated by the UPS to be sent to the IBM Systems Director server.
- The metering interval for the UPS should be set as short as possible to detect changes and notify the associated resources as early as possible.
- Do not set system power capping or system power savings policies on both a slot and on the blade in that slot.
- Do not set group capping policies on two groups, where one group contains a slot and the other group contains a blade in that slot.
- Ensure that the resource against which you are applying the power cap policy supports the same power cap type (ac or dc).
- Make use of the base IBM Systems Director features such as “Scheduling jobs”, “Event automation plan”, and “Monitors and thresholds”, along with AEM power monitoring to get more advantages.
- If a resource has both natively metered data and externally metered data, set the metering interval for the metering device and the resource to the same value.
- Use AEM power policies to enforce the power management settings to ensure that the power savings value is continuously enforced.
- Adding the same PDU to multiple IBM Systems Director servers could cause the PDU’s minimum and maximum power averages to be incorrect. Some PDUs such as LinkHub do not support to be managed by multiple AEM instances. The best practice is to add a PDU to a single IBM Systems Director instance.

- For performing outlet switching on supported PDUs, always discover the PDU that is using the SNMPv1 Read/Write community name or SNMPv3.
- Study the power consumption trend and the CPU utilization of servers for a good amount of time. Then, determine the time windows where you can enable the power savings feature.
- Enable hard power capping unless you have a need for accommodating more servers within the given power budget that is not met by the [Minimum guaranteed input power value - Maximum input power value] range.

7

AIX Profile Manager

This chapter provides a full description of AIX Profile Manager (APM). The AIX Profile Manager is a tool that can help implement and monitor the security of all AIX servers in a production environment, but also implement and monitor the system compliance of those AIX servers.

This chapter contains the following topics:

7.1 AIX Profile Manager

The AIX Profile Manager is an excellent tool that can help implement and monitor (from a central location) the security of all AIX servers in a production environment. The tool can also implement and monitor the system compliance of those AIX servers. System compliance is the tuning and system settings, which must be the same on all servers like memory tuning options and system dump settings.

The most important thing to understand is that the security and system compliance are quite different to manage and implement into AIX Profile Manager. The common way is always to create or use an existing profile, create or use an existing template from the profile, then deploy the template on the servers.

Important: You can deploy a template to several servers but a server can have only one template that is deployed.

The first part of this chapter covers the security implementation through examples. The second part covers the system compliance monitoring and settings, which can be customized to improve this compliance.

7.2 AIX Profile Manager: Security

APM is used for implementing and monitoring compliance security rules that are based on predefined standards profiles on AIX and VIO servers.

To access the AIX Profile Manager interface through the IBM Systems Director home page interface, select **Plugins** → **AIX Profile Manager**, as shown in Figure 7-1.

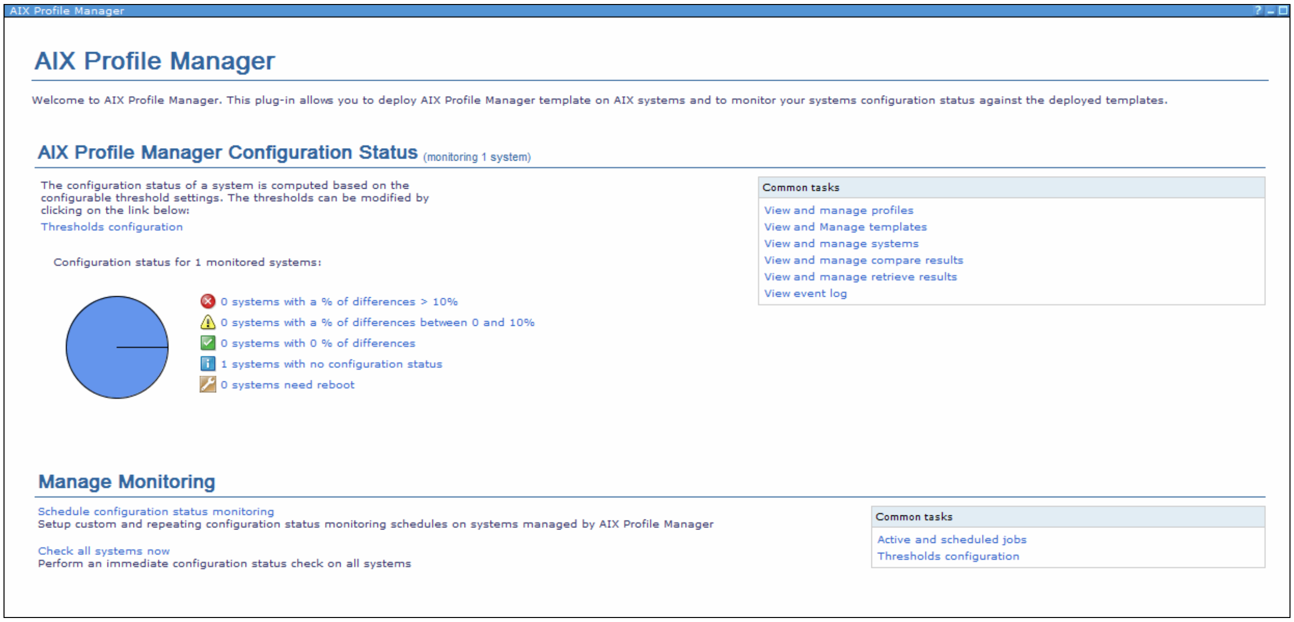

AIX Profile Manager

Welcome to AIX Profile Manager. This plug-in allows you to deploy AIX Profile Manager template on AIX systems and to monitor your systems configuration status against the deployed templates.

AIX Profile Manager Configuration Status (monitoring 1 system)

The configuration status of a system is computed based on the configurable threshold settings. The thresholds can be modified by clicking on the link below:
Thresholds configuration

Configuration status for 1 monitored systems:

- 0 systems with a % of differences > 10%
- 0 systems with a % of differences between 0 and 10%
- 0 systems with 0 % of differences
- 1 systems with no configuration status
- 0 systems need reboot

Common tasks
- View and manage profiles
- View and Manage templates
- View and manage systems
- View and manage compare results
- View and manage retrieve results
- View event log

Manage Monitoring

Schedule configuration status monitoring
Setup custom and repeating configuration status monitoring schedules on systems managed by AIX Profile Manager

Check all systems now
Perform an immediate configuration status check on all systems

Common tasks
- Active and scheduled jobs
- Thresholds configuration

Figure 7-1 AIX Profile Manager home page

The predefined standards that are available with AIX Profile Manager are listed in Table 7-1:

Table 7-1 AIX Profile Manager predefined standards

Security model in APM	Corresponds to:
ice_DLS	Default IBM Security model for AIX
ice_LLS	Low level of security defined by IBM for AIX
ice_MLS	Medium level of security defined by IBM for AIX
ice_HLS	High level of security defined by IBM for AIX
ice_DOD	Department of Defense (DoD) STIG compliance settings
ice_PCI	Payment card industry Data Security Standard compliance
ice_SOX	Sarbanes-Oxley Act and COBIT compliance
ice_Hipaa	Health Insurance Portability and Accountability Act (HIPAA) compliance

7.2.1 Requirements

AIX Profile Manager is supported on every IBM system and operating system that is also supported in IBM Systems Director 6.3.

Following are the basic requirements to install and use AIX Profile Manager with IBM Systems Director:

- IBM PowerSC™ Express Edition must be installed on all servers to enable the use of all policies.
- For AIX 6.1: TL7 or higher is required.
- For AIX 7.1: TL1 or higher is required.
- The file set `bos.ahafs.rte` is required on all systems.
- AIX Profile Manager plug-in must be installed on the IBM Systems Director server.
- AIX servers must be discovered and inventoried before using APM.

To check if prerequisites are available, run the commands that are shown in Example 7-1:

Example 7-1 Commands to check if all prerequisites are installed (AIX 6.1 is used in this example)

```
oslevel -s
6100-07-XX-XXX

lslpp -l powerscExp.licence
Fileset                      Level  State      Description
  ----------------------------------------------------------------------------
Path: /usr/lib/objrepos
  powerscExp.license       6.1.6.15  COMMITTED  PowerSC Express Edition

lslpp -l powerscExp.ice.cmds
  Fileset                      Level  State      Description
  ----------------------------------------------------------------------------
Path: /usr/lib/objrepos
  powerscExp.ice.cmds        1.1.2.0  COMMITTED  ICE Express Security Extension

Path: /etc/objrepos
  powerscExp.ice.cmds        1.1.2.0  COMMITTED  ICE Express Security Extension

lslpp -l bos.ahafs
Fileset                      Level  State      Description
  ----------------------------------------------------------------------------
Path: /usr/lib/objrepos
  bos.ahafs                 6.1.7.15  APPLIED    Aha File System

Path: /etc/objrepos
  bos.ahafs                 6.1.7.15  APPLIED    Aha File System
```

AIX 6.1 is used in the Example 7-1.

Important: The profiles that are available in AIX Profile Manager cannot be customized. If you want your own model of security, either use AIX Security Expert (AIXPert) or create your own profiles and templates from the **artex** command. More information about AIX Runtime Expert (artex) commands is available further in this chapter. As a best practice, do not edit the existing templates (PCI, SOX, DoD, and HIPAA; create your own.

Attention: The XML files from the aixpert commands cannot be imported into AIX Profile Manager; this is not supported.

7.2.2 Before implementing

Before implementing on all servers to the desired security model, consider the following factors.

Choose the right security policy. A security policy that is too high can prevent any application from working correctly and users from logging in. A security model that is too low could be a serious risk. To check all available rules for policies in AIX Profile Manager, check the content of the XML files in the `/etc/security/aixpert/custom` directory for each policy. To generate XML files for DLS, LLS, MLS, and HLS profiles, run the **`aixpert`** command that is shown in Example 7-2.

Example 7-2 Generate XML files from aixpert command

```
cd /etc/security/aixpert/custom
aixpert -l default -n -o security_default.xml
aixpert -l low -n -o security_low.xml
aixpert -l medium -n -o security_medium.xml
aixpert -l high -n -o security_high.xml
ls securit*
security_default.xml security_low.xml security_medium.xml security_high.xml
```

The files for PCI, SOX-COBIT, DoD, and HIPAA are already located in the directory if package `powerscExp.ice.cmds` is installed on the operating system. You can check if it is installed by running the command in Example 7-3.

Example 7-3 Run this command to check the other security XML files

```
lslpp -l powerscExp.ice.cmds
Fileset                      Level  State      Description
  ----------------------------------------------------------------------------
Path: /usr/lib/objrepos
  powerscExp.ice.cmds        1.1.2.0  COMMITTED  ICE Express Security Extension

Path: /etc/objrepos
  powerscExp.ice.cmds        1.1.2.0  COMMITTED  ICE Express Security Extension

cd /etc/security/aixpert/custom
ls PCI.xml DoD.xml Hipaa.xml SOX-COBIT.xml
DoD.xml        Hipaa.xml      PCI.xml        SOX-COBIT.xml
```

The format for each rule in an XML file is shown in Example 7-4.

Example 7-4 Example of a rule in an XML format

```
<AIXPertEntry name="pci_maxage" function="maxage">
    <AIXPertRuleType type="PLS"/>
    <AIXPertDescription>Implements PCI Section 8.5.9 for Maximum age for password:
Specifies the maximum number of weeks    (13 weeks, atleast 90 days) that a
password is
valid</AIXPertDescription>

<AIXPertPrereqList>bos.rte.date,bos.rte.commands,bos.rte.security,bos.rte.shell,bo
s.rte.ILS</AIXPertPrereqList>
    <AIXPertCommand>/etc/security/aixpert/bin/chusrattr</AIXPertCommand>
    <AIXPertArgs>maxage=13 ALL pci_maxage</AIXPertArgs>
    <AIXPertGroup>Password policy rules</AIXPertGroup>
```

```
</AIXPertEntry>
```

In Example 7-4 on page 379, the most important lines are the following rules:

- *AIXPertEntry name*: This is the name of the rule
- *AIXPertDescription*: This is a description of the rule
- *AIXPertPrereqList*: Contains a list of prerequisites to execute the command (list required file sets here)
- *AIXPertCommand*: This is the command that will be executed. In our example, this is the **chusrattr** command
- *AIXPertArgs*: This contains the arguments of the command

In Example 7-4 on page 379, the rule implements a maxage with a value of 13 weeks for all users.

Attention: This rule could lead to password expiration if the maxage setting is not already set in your environment. Remember to check the age of all passwords before implementing this rule; otherwise, critical application users could be locked.

The following steps must be followed before implementing security rules on all systems:

1. Test the rules in a test environment before implementing the selected security model on all systems.
2. When tested, apply the selected model step by step. For example:
   a. Apply the security model on non-critical systems, then wait at least one or two weeks. This could help to identify and solve all problems before implementing the security model on critical systems. If there are too many problems, apply a rollback on the systems or a lower security model.
   b. Implement the model on the production systems and closely monitor those systems.

   Tip: Do not implement the security levels on all productions at one time. Go step by step, server by server.

7.2.3 Implementation

As an example in this subchapter, we work through the implementation procedure for the LLS security model on an AIX server.

This example is simple because profiles and templates are already created in APM. This example works for all security models that are available into APM:

1. From the AIX Profile Manager console, select **View and Manage templates**. In Figure 7-2 on page 381, you can see the default available templates in AIX Profile Manager.

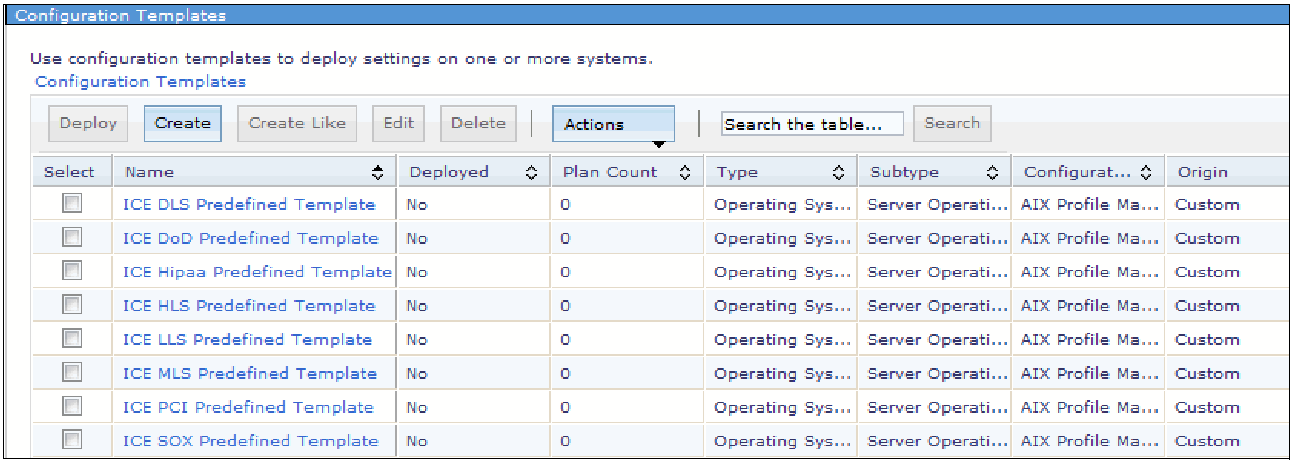

Configuration Templates

Use configuration templates to deploy settings on one or more systems.

Configuration Templates

Select	Name	Deployed	Plan Count	Type	Subtype	Configurat...	Origin
□	ICE DLS Predefined Template	No	0	Operating Sys...	Server Operati...	AIX Profile Ma...	Custom
□	ICE DoD Predefined Template	No	0	Operating Sys...	Server Operati...	AIX Profile Ma...	Custom
□	ICE Hipaa Predefined Template	No	0	Operating Sys...	Server Operati...	AIX Profile Ma...	Custom
□	ICE HLS Predefined Template	No	0	Operating Sys...	Server Operati...	AIX Profile Ma...	Custom
□	ICE LLS Predefined Template	No	0	Operating Sys...	Server Operati...	AIX Profile Ma...	Custom
□	ICE MLS Predefined Template	No	0	Operating Sys...	Server Operati...	AIX Profile Ma...	Custom
□	ICE PCI Predefined Template	No	0	Operating Sys...	Server Operati...	AIX Profile Ma...	Custom
□	ICE SOX Predefined Template	No	0	Operating Sys...	Server Operati...	AIX Profile Ma...	Custom

Figure 7-2 Security templates available in AIX Profile Manager

2. Select **ICE LLS Predefined template** → **Deploy**.

   From this point, there are several options to select the targeted host and tabs for running the deployment job.

   To select the host, there are three options that are available in the Target tab in the Show menu, which are All Targets, Groups, and Recent Targets, as shown in Figure 7-3.

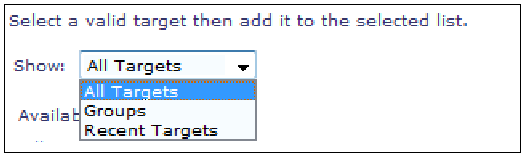

Select a valid target then add it to the selected list.

Show: All Targets

Figure 7-3 Options that are available to find the targeted host

3. Select **All Targets**. Then, enter the name of the client in the Search field (Figure 7-4 on page 382) and then click **Search**.
4. When the targeted host is found, click **Select** → **Add**. For this example, the targeted host is **itso-cb-sys5.itso.ral.ibm.com**.

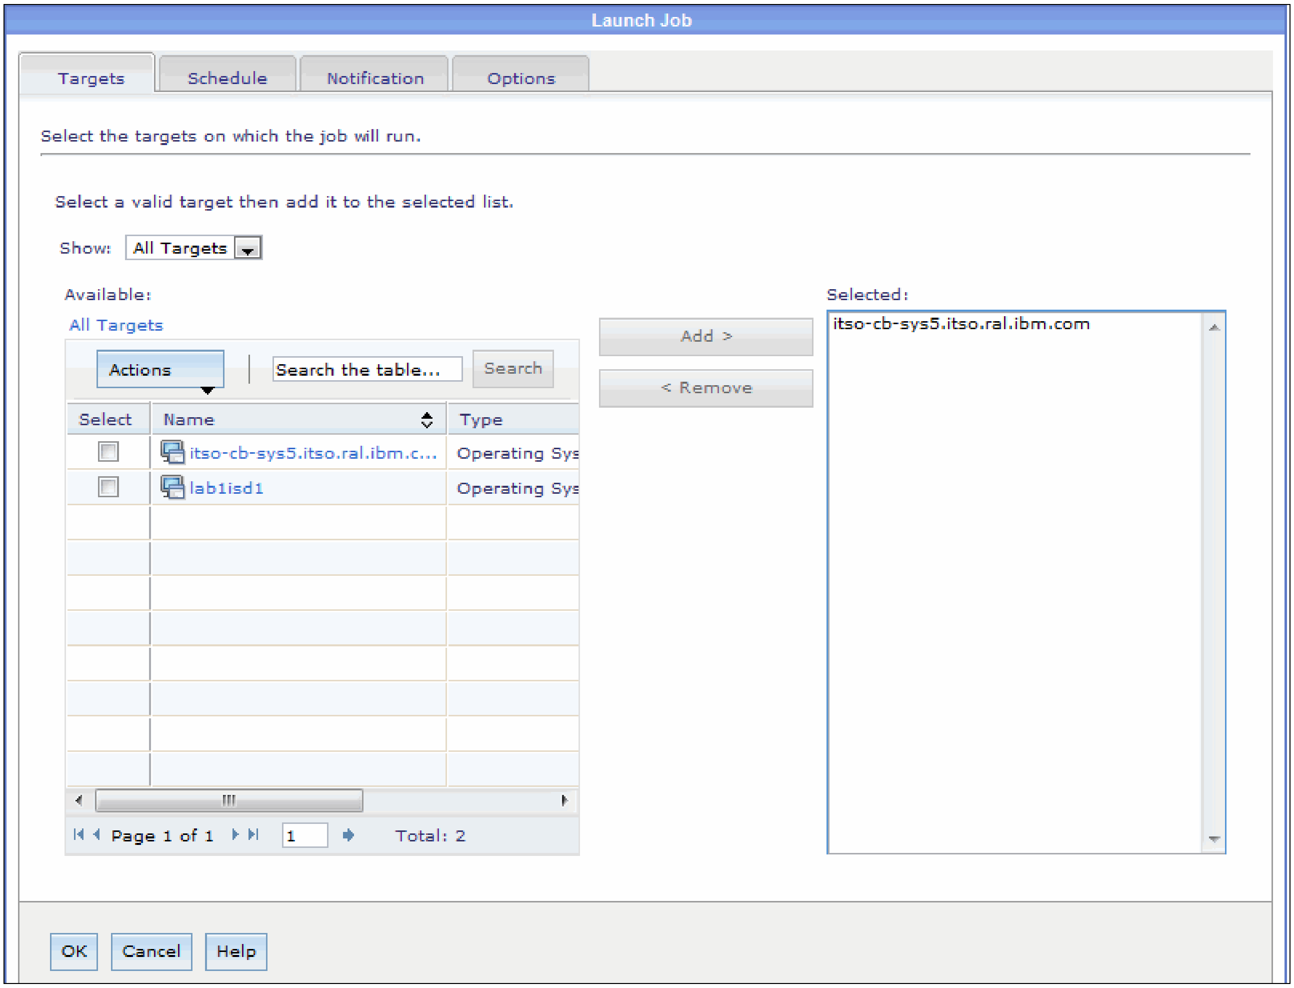

Select the targets on which the job will run.

Select a valid target then add it to the selected list.

Show: All Targets

Available:

All Targets

Search the table...

Select	Name	Type
□	itso-cb-sys5.itso.ral.ibm.c...	Operating Sys
□	lab1isd1	Operating Sys

Page 1 of 1 1 Total: 2

Selected:

itso-cb-sys5.itso.ral.ibm.com

Figure 7-4 Selecting the host where to deploy the security policy

The other tabs are Schedule, Notification, and Options. In theses tabs, there are options for scheduling the jobs, managing notifications to send mails on events during the job, and other options.

> **Note:** In the Schedule tab for the Job Name, it is useful to insert the host name of the targeted host at the beginning of the name of the job. This way, it is easier to find the job in the Task Management of IBM Systems Director.

5. Click **OK** to run the job. To check the status of the job, click **Display Properties** and then go to the Logs tab.

   When the job has run successfully, the security model LLS is deployed on the target host.

7.3 AIX Profile Manager: System compliance

System compliance is quite different to implement and to manage compared with security compliance.

First, no profiles are available by default into AIX Profile Manager. The reason is simple. The system configuration must be taken from one server that is set as a standard and then, as a best practice, use the *artex* commands to generate a new profile, which contains the system settings of the standard server. Finally, import the profile into the AIX Profile Manager, create a template, and deploy to additional servers.

Notice: You can create XML profiles from AIX Profile Manager console, but with AIX Runtime Expert, you can manage the content of the XML files that you created. For example, you can create a profile for the “no” options (network tuning options) from AIX Profile Manager by retrieving values from a selected system. But this profile contains all the “no” settings that are available. If you want to have only a few of these options in AIX Profile Manager (`tcp_sendspace` and `tcp_recvspace`, as an example), it is not possible to manage it with the graphical interface of AIX Profile Manager. Using AIX Runtime Expert gives you the ability to manage your files as you want.

On the same way, if you want to merge several profiles, you will have a thousand settings when creating the profile from AIX Profile Manager whereas you can decide which settings you want in your XML with AIX Runtime Expert commands. In this document, we decided to use AIX Runtime Expert commands as a best practice.

Online Content: At the following site, there is a demo that shows how to create profiles from the AIX Profile Manager graphical interface:

http://youtu.be/h2CT7KOJM7c

You can also scan the QR code that is displayed in the left margin to go directly to the video.

Before doing these steps, we explore AIX Runtime Expert, which is required to create and manage profiles.

Complete documentation about AIX Runtime Expert can be found in the AIX Information Center:

http://pic.dhe.ibm.com/infocenter/aix/v7r1/index.jsp?topic=%2Fcom.ibm.aix.baseadmn%2Fdoc%2Fbaseadmndita%2Fartex_main.htm

7.3.1 Profiles from AIX Runtime Expert

This section shows how to use AIX Runtime Expert commands to create profiles, import them into AIX Profile Manager, and then deploy them on the other systems.

Prerequisites

The AIX Runtime Manager is available for AIX 6.1 and AIX 7.1. Check if the following packages are installed on the host, as shown in Example 7-5.

Example 7-5 Example for an AIX 6.1 server

```
lslpp -l |grep artex
  artex.base.agent          6.1.7.15  APPLIED    AIX Runtime Expert CAS agent
```

```
 artex.base.rte            6.1.7.15  APPLIED    AIX Runtime Expert
 artex.base.samples         6.1.7.0  APPLIED    AIX Runtime Expert sample
 artex.base.agent          6.1.7.15  APPLIED    AIX Runtime Expert CAS agent
 artex.base.rte            6.1.7.15  APPLIED    AIX Runtime Expert
 artex.base.samples         6.1.7.0  APPLIED    AIX Runtime Expert sample
```

Default profiles that are available

When the AIX Runtime Expert is installed on the AIX server, go to the `/etc/security/artex/samples` directory, where you find the default profiles from artex.

> **Important:** Do not modify the profile files. The process of modifying options is explained later in this document. These files are just samples and do not contain any value.

To list the profiles that are available, use the **`artextlist`** command as shown in Example 7-6.

Example 7-6 List profiles that are known by AIX Runtime Expert

```
artexlist
/etc/security/artex/samples/acctctlProfile.xml
/etc/security/artex/samples/aixpertProfile.xml
/etc/security/artex/samples/all.xml
/etc/security/artex/samples/alogProfile.xml
/etc/security/artex/samples/authProfile.xml
/etc/security/artex/samples/authentProfile.xml
/etc/security/artex/samples/chconsProfile.xml
/etc/security/artex/samples/chdevProfile.xml
/etc/security/artex/samples/chlicenseProfile.xml
/etc/security/artex/samples/chservicesProfile.xml
/etc/security/artex/samples/chssysProfile.xml
/etc/security/artex/samples/chsubserverProfile.xml
/etc/security/artex/samples/chuserProfile.xml
/etc/security/artex/samples/classProfile.xml
/etc/security/artex/samples/coreProfile.xml
/etc/security/artex/samples/default.xml
/etc/security/artex/samples/dumpctrlProfile.xml
/etc/security/artex/samples/envProfile.xml
/etc/security/artex/samples/errdemonProfile.xml
/etc/security/artex/samples/ewlmProfile.xml
/etc/security/artex/samples/ffdcProfile.xml
/etc/security/artex/samples/filterProfile.xml
/etc/security/artex/samples/gencopyProfile.xml
/etc/security/artex/samples/iooProfile.xml
/etc/security/artex/samples/krecoveryProfile.xml
/etc/security/artex/samples/login.cfgProfile.xml
/etc/security/artex/samples/lvmoProfile.xml
/etc/security/artex/samples/mktcpipProfile.xml
/etc/security/artex/samples/mkuser.defaultProfile.xml
/etc/security/artex/samples/namerslvProfile.xml
/etc/security/artex/samples/nfsProfile.xml
/etc/security/artex/samples/nfsoProfile.xml
/etc/security/artex/samples/nisProfile.xml
/etc/security/artex/samples/noProfile.xml
/etc/security/artex/samples/probevueProfile.xml
/etc/security/artex/samples/rasoProfile.xml
```

```
/etc/security/artex/samples/roleProfile.xml
/etc/security/artex/samples/ruserProfile.xml
/etc/security/artex/samples/schedoProfile.xml
/etc/security/artex/samples/secattrProfile.xml
/etc/security/artex/samples/shconfProfile.xml
/etc/security/artex/samples/smtctlProfile.xml
/etc/security/artex/samples/syscorepathProfile.xml
/etc/security/artex/samples/sysdumpdevProfile.xml
/etc/security/artex/samples/trcctlProfile.xml
/etc/security/artex/samples/trustchkProfile.xml
/etc/security/artex/samples/tsdProfile.xml
/etc/security/artex/samples/viosdevattrProfile.xml
```

Each profile manages a category of settings for the system.

Create a profile

For example, to create a profile to manage and control the sysdumpdev configuration, follow these steps:

1. On the standard system, run the command that is shown in Example 7-7.

Example 7-7 View configuration settings of the system with the artexget command

```
artexget -r /etc/security/artex/samples/sysdumpdevProfile.xml
<?xml version="1.0" encoding="UTF-8"?>
<Profile origin="get" version="2.0.1" date="2013-04-22T14:23:05Z">
 <Catalog id="sysdumpdevParam" version="2.1">
  <Parameter name="primary" value="/dev/lg_dumplv"/>
  <Parameter name="secondary" value="/dev/sysdumpnull"/>
  <Parameter name="copy_directory" value="/var/adm/ras"/>
  <Parameter name="forced_copy_flag" value="1"/>
  <Parameter name="always_allow_dump" value="0"/>
  <Parameter name="type_of_dump" value="traditional" applyType="nextboot" 
reboot="true"/>
  <Parameter name="full_memory_dump"/>
 </Catalog>
</Profile>
```

   The result of this command is the values from the sysdump configuration of the standard system that is displayed in XML format. This is the profile for the AIX Profile Manager.

2. Now, create the profile XML file by using the command that is shown in Example 7-8:

Example 7-8 Create the profile file in XML format

```
artexget -r /etc/security/artex/samples/sysdumpdevProfile.xml > 
/tmp/sydumpdev_custom_profile.xml
```

   The file that is created is now different from the sample file because it contains the values of the current system configuration.

   In this manner, the file can be fully controlled before being imported into AIX Profile Manager. The file can be modified and customized.

Add the profile into AIX Profile Manager

In this example, the host where the template is deployed is the same location where we created the template because of some limitations within our lab environment.

Perform the following steps to add created profiles into AIX Profile Manager:

1. From the AIX Profile Manager console, click **View and manage profiles**, as shown in Figure 7-5.

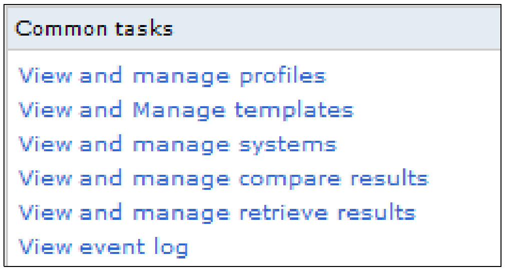

Common tasks

View and manage profiles
View and Manage templates
View and manage systems
View and manage compare results
View and manage retrieve results
View event log

Figure 7-5 Select from menu to manage profiles

2. Then, select **Import** → **Import from a system**, select the standard host where the XML file has been generated, and click **OK**.
3. Enter the name of the directory where the XML file is located and select **View**. Select the file and click **Add**, as shown in Figure 7-6.

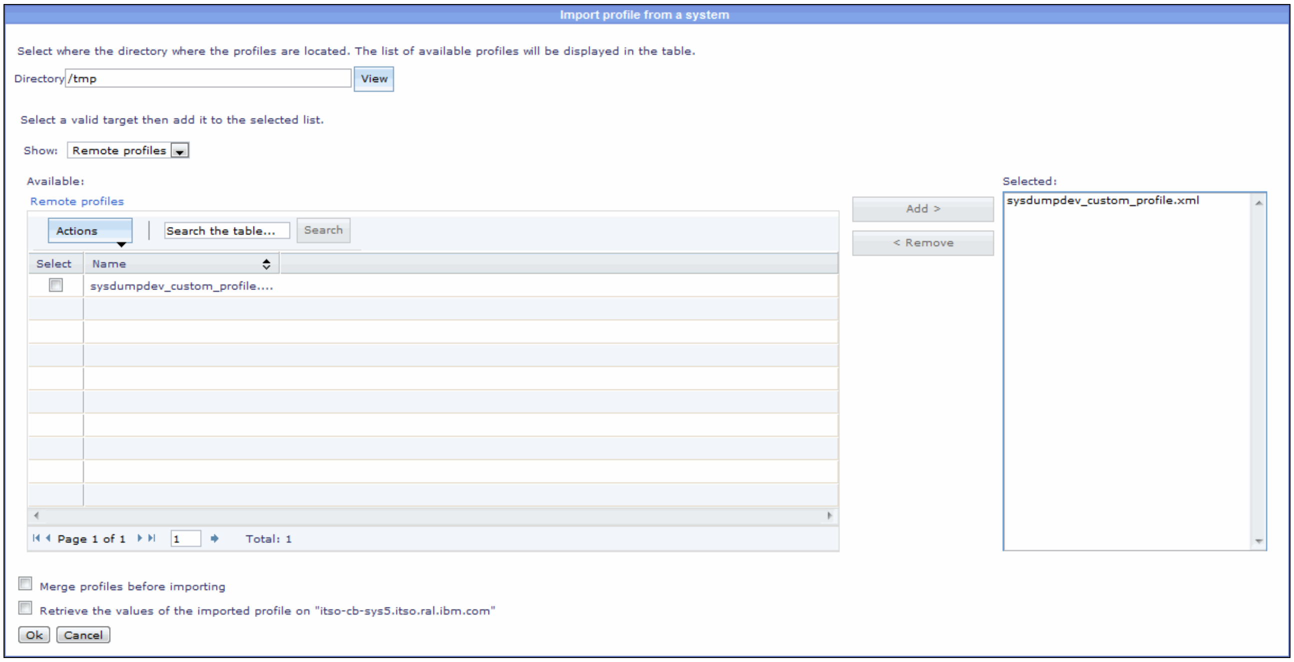

Import profile from a system

Select where the directory where the profiles are located. The list of available profiles will be displayed in the table.

Directory /tmp

Select a valid target then add it to the selected list.

Show: Remote profiles

Available:
Remote profiles

Select	Name
□	sysdumpdev_custom_profile....

Page 1 of 1   1   Total: 1

Selected:
sysdumpdev_custom_profile.xml

□ Merge profiles before importing

□ Retrieve the values of the imported profile on "itso-cb-sys5.itso.ral.ibm.com"

Figure 7-6 Import the XML file to AIX Profile Manager

4. Select **OK** to import the file.

The file is now integrated into the AIX Profile Manager, as shown in Figure 7-7.

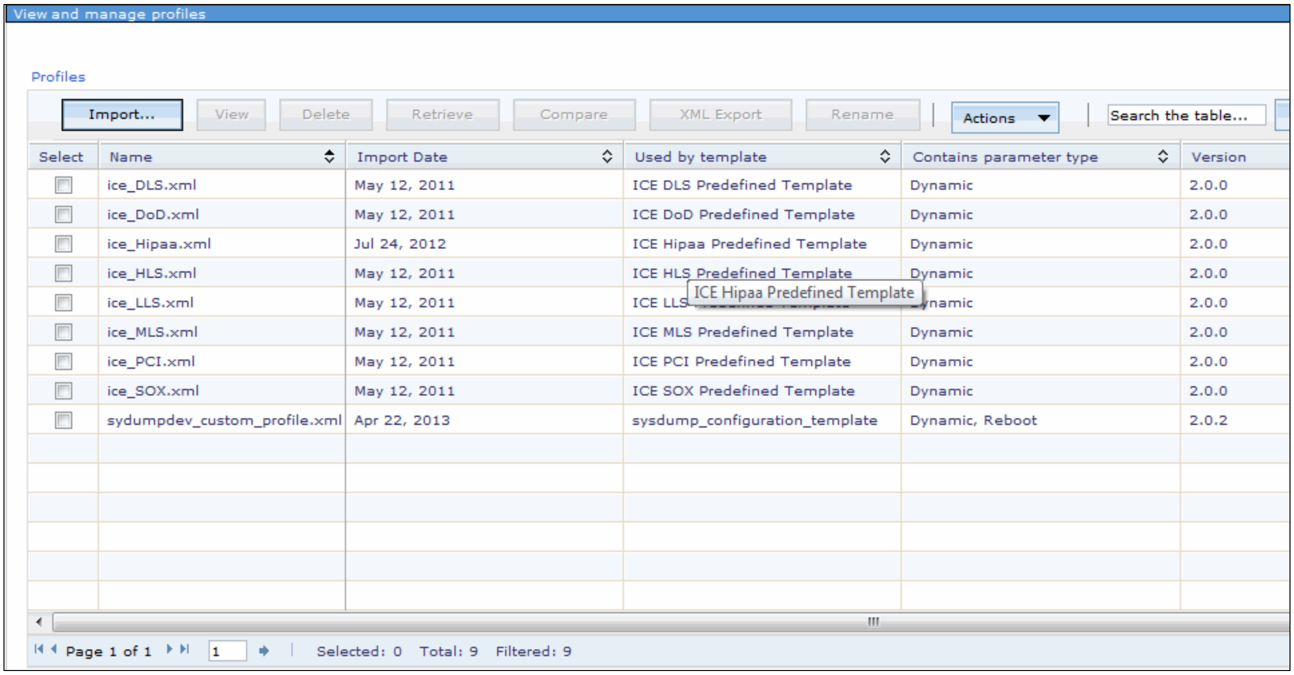

Select	Name	Import Date	Used by template	Contains parameter type	Version
□	ice_DLS.xml	May 12, 2011	ICE DLS Predefined Template	Dynamic	2.0.0
□	ice_DoD.xml	May 12, 2011	ICE DoD Predefined Template	Dynamic	2.0.0
□	ice_Hipaa.xml	Jul 24, 2012	ICE Hipaa Predefined Template	Dynamic	2.0.0
□	ice_HLS.xml	May 12, 2011	ICE HLS Predefined Template	Dynamic	2.0.0
□	ice_LLS.xml	May 12, 2011	ICE LLS Predefined Template	Dynamic	2.0.0
□	ice_MLS.xml	May 12, 2011	ICE MLS Predefined Template	Dynamic	2.0.0
□	ice_PCI.xml	May 12, 2011	ICE PCI Predefined Template	Dynamic	2.0.0
□	ice_SOX.xml	May 12, 2011	ICE SOX Predefined Template	Dynamic	2.0.0
□	sydumpdev_custom_profile.xml	Apr 22, 2013	sysdump_configuration_template	Dynamic, Reboot	2.0.2

Figure 7-7 List of profiles, which now contains the new profile

> **Note:** It is also possible to import XML files by using the AIX Profile Manager menus, but as a best practice, we recommend the usage of artex commands, which allow you to have better control of the XML file content.

Create the template

Create a template that is based on the previously created profile to allow a deployment to other systems:

1. From the AIX Profile Manager console, select **View and Manage templates** → **Create**.

2. Select **Operating System** as the “Template type”. Select **AIX Profile Manager** as the “Configuration to create a template”. Insert a name for the template and description of the template and finish by selecting **Continue**, as shown in Figure 7-8.

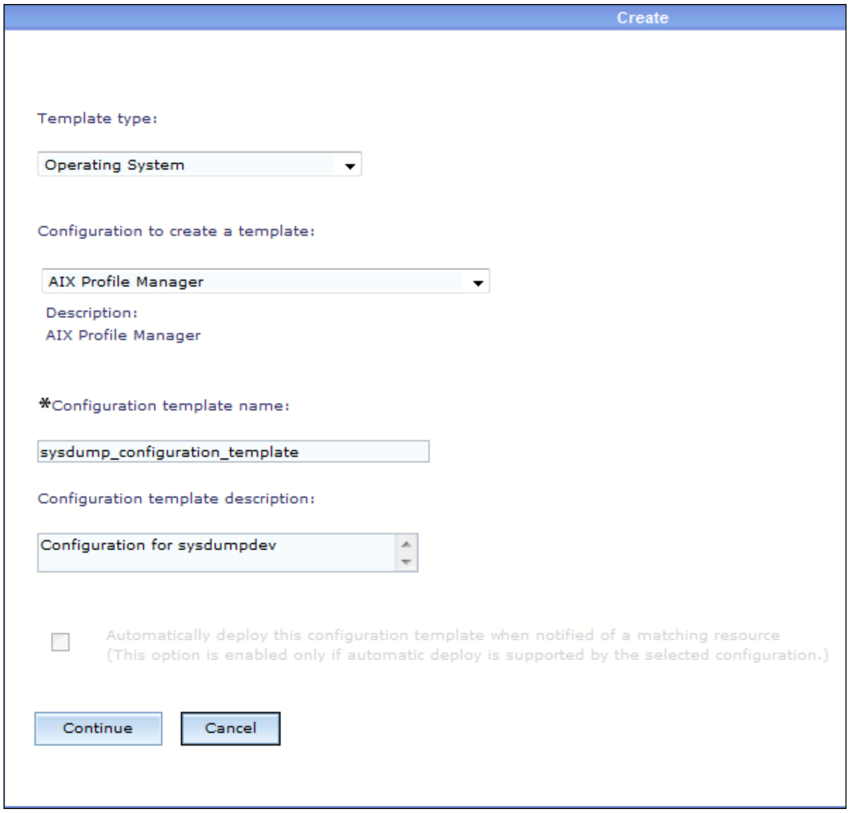

Create

Template type:

Operating System

Configuration to create a template:

AIX Profile Manager

Description:
AIX Profile Manager

*Configuration template name:

sysdump_configuration_template

Configuration template description:

Configuration for sysdumpdev

□ 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.)

Figure 7-8 Create a template

3. Select **Browse** to find the previously created profile. In the selection window, remove the preselected profile name **USMEveryoneRole** and add the **sysdumpdev_custom_profile** profile, as shown in Figure 7-9.

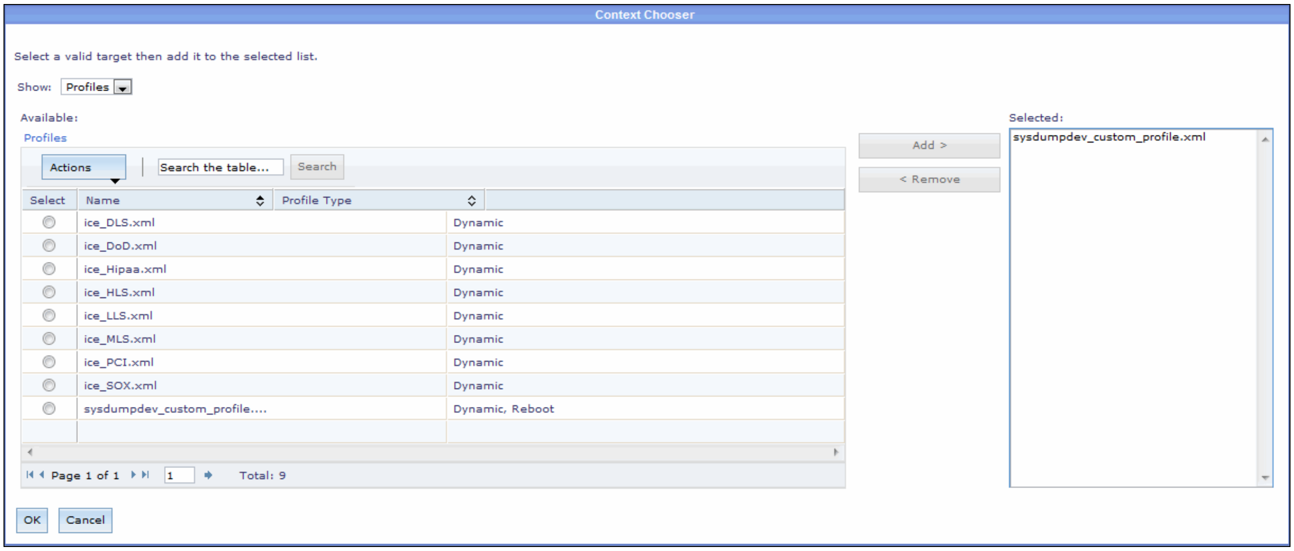

Context Chooser

Select a valid target then add it to the selected list.

Show: Profiles

Available:
Profiles

Select	Name	Profile Type
○	ice_DLS.xml	Dynamic
○	ice_DoD.xml	Dynamic
○	ice_Hipaa.xml	Dynamic
○	ice_HLS.xml	Dynamic
○	ice_LLS.xml	Dynamic
○	ice_MLS.xml	Dynamic
○	ice_PCI.xml	Dynamic
○	ice_SOX.xml	Dynamic
○	sysdumpdev_custom_profile....	Dynamic, Reboot

Page 1 of 1 1 Total: 9

Selected:
sysdumpdev_custom_profile.xml

Figure 7-9 Adding the new profile into the template

4. Select **OK**. In our example, we did not change the default options in the Template Options. See Figure 7-10.

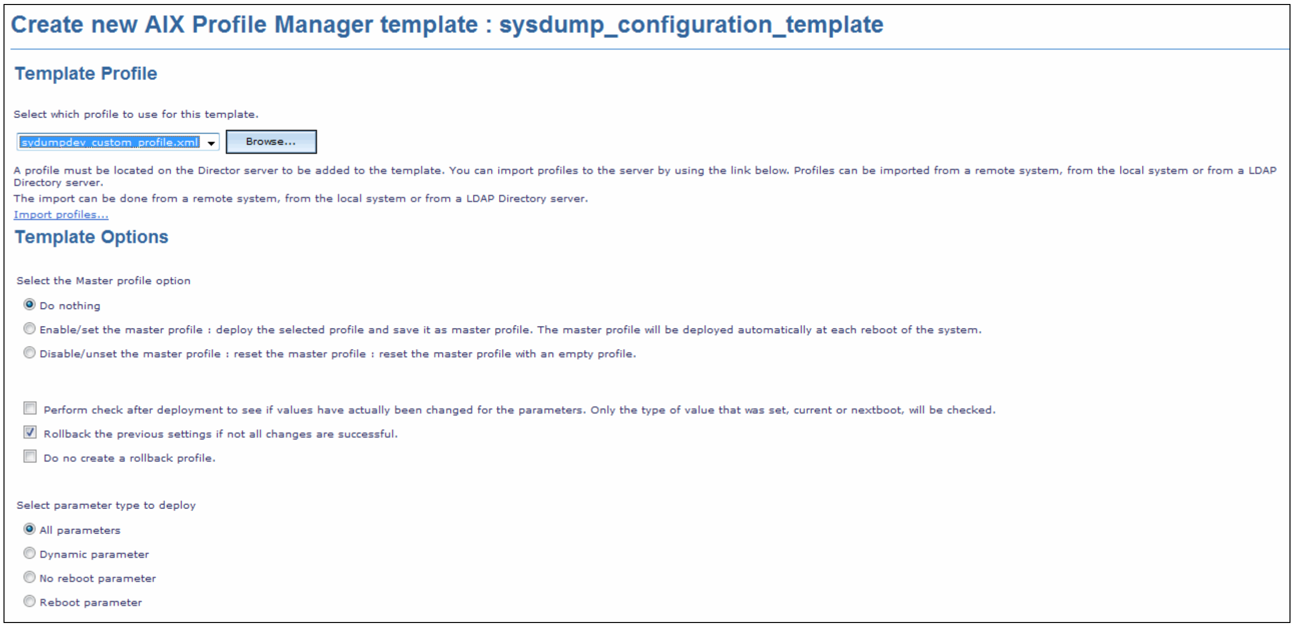

Create new AIX Profile Manager template : sysdump_configuration_template

Template Profile

Select which profile to use for this template.

sysdumpdev_custom_profile.xml ▾ [Browse...]

A profile must be located on the Director server to be added to the template. You can import profiles to the server by using the link below. Profiles can be imported from a remote system, from the local system or from a LDAP Directory server.
The import can be done from a remote system, from the local system or from a LDAP Directory server.
Import profiles...

Template Options

Select the Master profile option

● Do nothing

○ Enable/set the master profile : deploy the selected profile and save it as master profile. The master profile will be deployed automatically at each reboot of the system.

○ Disable/unset the master profile : reset the master profile : reset the master profile with an empty profile.

☐ Perform check after deployment to see if values have actually been changed for the parameters. Only the type of value that was set, current or nextboot, will be checked.

☑ Rollback the previous settings if not all changes are successful.

☐ Do no create a rollback profile.

Select parameter type to deploy

● All parameters

○ Dynamic parameter

○ No reboot parameter

○ Reboot parameter

Figure 7-10 Create new AIX Profile Manager template: sysdump_configuration_template

5. Click **Save**. The template is now ready to be deployed. As shown in Figure 7-11, we can now see the template **sysdump_configuration_template** in the list.

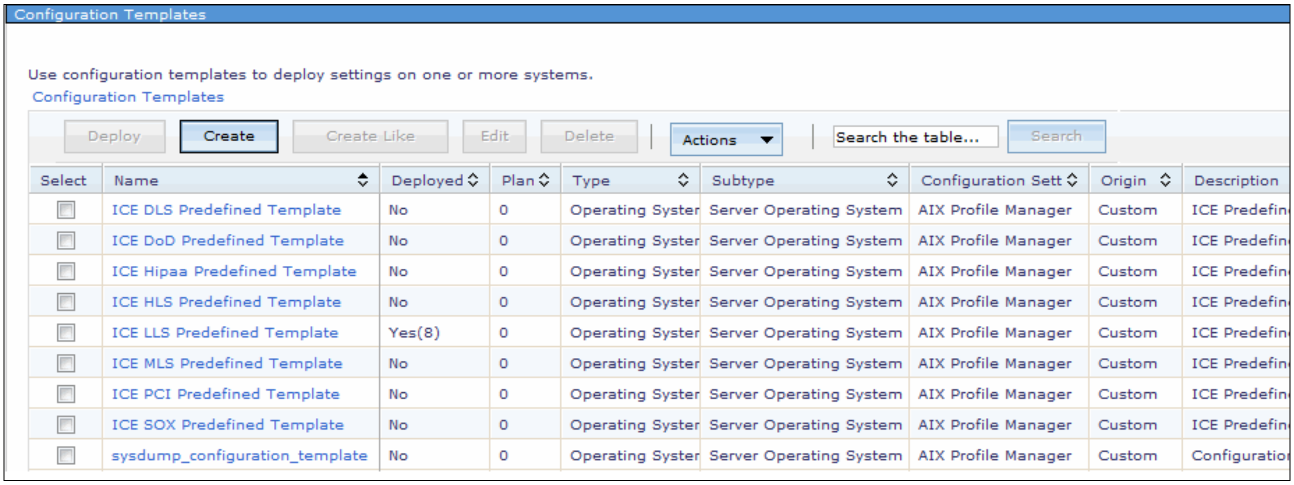

Configuration Templates

Use configuration templates to deploy settings on one or more systems.
Configuration Templates

Select	Name	Deployed	Plan	Type	Subtype	Configuration Sett	Origin	Description
☐	ICE DLS Predefined Template	No	0	Operating Syster	Server Operating System	AIX Profile Manager	Custom	ICE Predefin
☐	ICE DoD Predefined Template	No	0	Operating Syster	Server Operating System	AIX Profile Manager	Custom	ICE Predefin
☐	ICE Hipaa Predefined Template	No	0	Operating Syster	Server Operating System	AIX Profile Manager	Custom	ICE Predefin
☐	ICE HLS Predefined Template	No	0	Operating Syster	Server Operating System	AIX Profile Manager	Custom	ICE Predefin
☐	ICE LLS Predefined Template	Yes(8)	0	Operating Syster	Server Operating System	AIX Profile Manager	Custom	ICE Predefin
☐	ICE MLS Predefined Template	No	0	Operating Syster	Server Operating System	AIX Profile Manager	Custom	ICE Predefin
☐	ICE PCI Predefined Template	No	0	Operating Syster	Server Operating System	AIX Profile Manager	Custom	ICE Predefin
☐	ICE SOX Predefined Template	No	0	Operating Syster	Server Operating System	AIX Profile Manager	Custom	ICE Predefin
☐	sysdump_configuration_template	No	0	Operating Syster	Server Operating System	AIX Profile Manager	Custom	Configuratio

Figure 7-11 List of templates, which now contains the new template

Deploy the template

To deploy the template, perform the following steps:

1. Select the previously created template name **sysdump_configuration_template** in the example, and click **Deploy**.
2. Select the host where the template will be deployed (`itso-cb-sys5.itso.ral.ibm.com` in the example). Then, in the Selected column, click **OK**.
3. Check on the deployed system that the sysdumpdev settings are correct by running the **`sysdumpdev`** command, as shown in Example 7-9 on page 390.

Example 7-9 Result of the sysdumpdev command

```
sysdumpdev -l
primary                /dev/lg_dumplv
secondary              /dev/sysdumpnull
copy directory         /var/adm/ras
forced copy flag       TRUE
always allow dump      FALSE
dump compression       ON
type of dump           traditional
```

Monitor the configuration with AIX Profile Manager

Now, we change the sysdumpdev configuration on one of the deployed systems and run a verification command to see what happens.

1. Run the **sysdumpdev** command as shown in Example 7-10.

Example 7-10 Change the primary device in the sysdumpdev configuration

```
sysdumpdev -p /dev/lv_dump
primary                /dev/lv_dump
secondary              /dev/sysdumpnull
copy directory         /var/adm/ras
forced copy flag       TRUE
always allow dump      FALSE
dump compression       ON
type of dump           traditional
```

2. From the AIX Profile Manager, select **Schedule configuration status monitoring** and select the host that has been modified (for our example, the host is `itso-cb-sys5.itso.ral.ibm.com`). Then, click **OK**. See Figure 7-12.

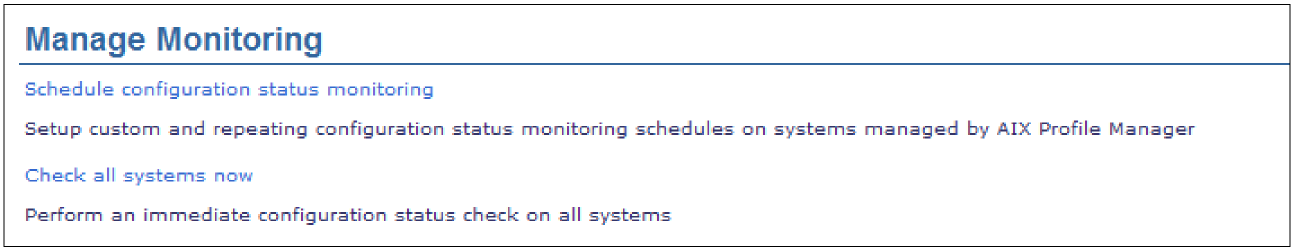

Manage Monitoring

Schedule configuration status monitoring

Setup custom and repeating configuration status monitoring schedules on systems managed by AIX Profile Manager

Check all systems now

Perform an immediate configuration status check on all systems

Figure 7-12 Manage Monitoring panel

3. Again, check the welcome page. The graphic should have changed depending of the configuration threshold. It should look similar to what is shown in Figure 7-13.

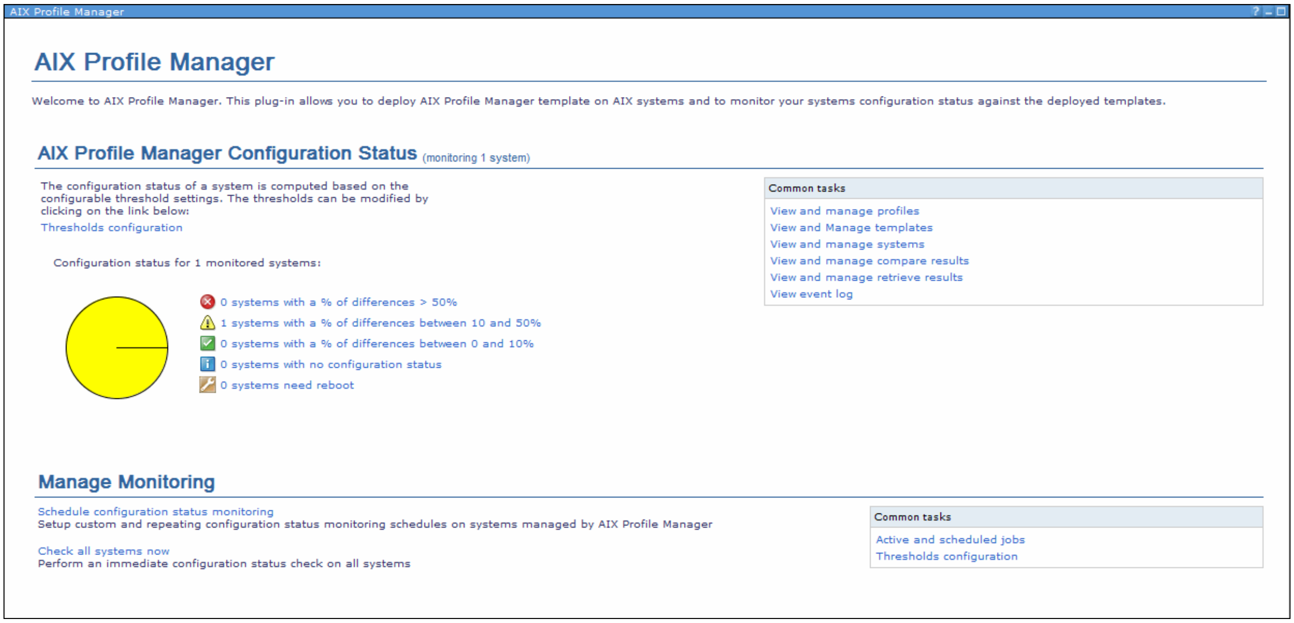

AIX Profile Manager

Welcome to AIX Profile Manager. This plug-in allows you to deploy AIX Profile Manager template on AIX systems and to monitor your systems configuration status against the deployed templates.

AIX Profile Manager Configuration Status (monitoring 1 system)

The configuration status of a system is computed based on the configurable threshold settings. The thresholds can be modified by clicking on the link below:
Thresholds configuration

Configuration status for 1 monitored systems:

- 0 systems with a % of differences > 50%
- 1 systems with a % of differences between 10 and 50%
- 0 systems with a % of differences between 0 and 10%
- 0 systems with no configuration status
- 0 systems need reboot

Common tasks
- View and manage profiles
- View and Manage templates
- View and manage systems
- View and manage compare results
- View and manage retrieve results
- View event log

Manage Monitoring

Schedule configuration status monitoring
Setup custom and repeating configuration status monitoring schedules on systems managed by AIX Profile Manager

Check all systems now
Perform an immediate configuration status check on all systems

Common tasks
- Active and scheduled jobs
- Thresholds configuration

Figure 7-13 The graphic indicates that an operating system is not compliant with our rule

4. In Figure 7-13, we click **1 systems with a % of differences between 10 and 50%** to see all settings that are detected as non-compliant, as shown in Figure 7-14.

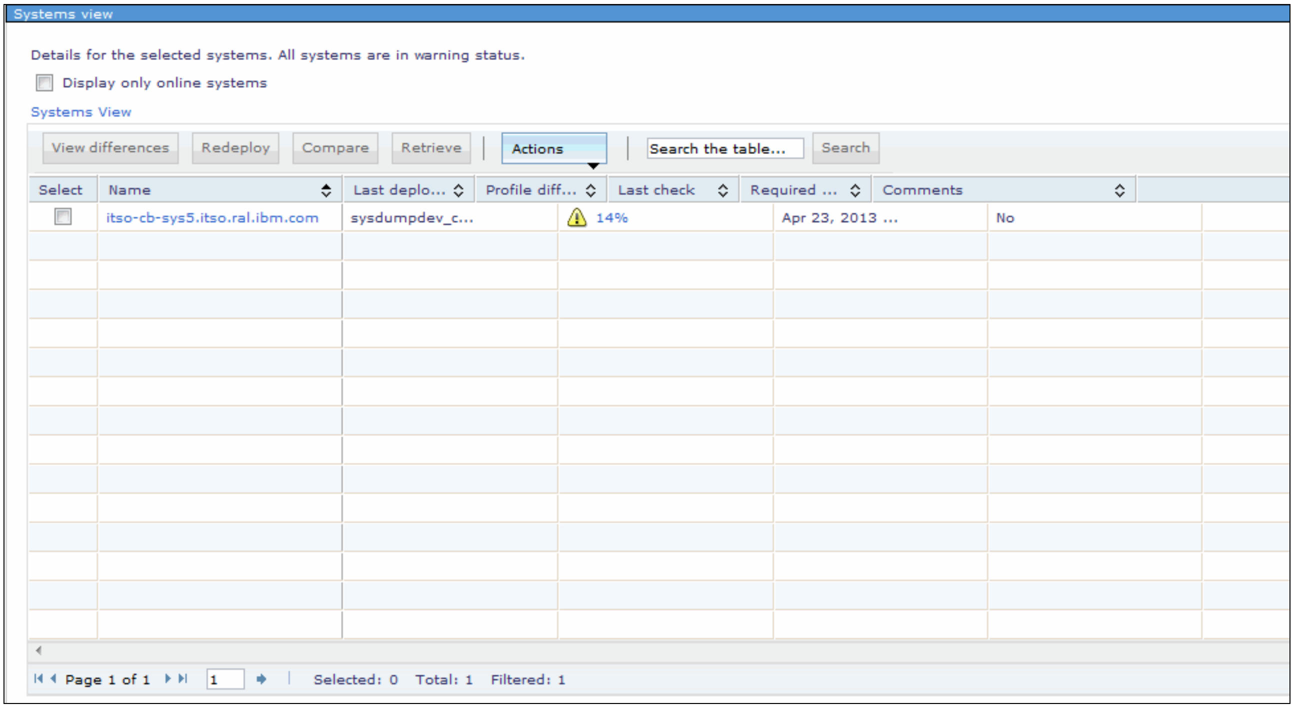

Details for the selected systems. All systems are in warning status.

□ Display only online systems

Systems View

Select	Name	Last deplo...	Profile diff...	Last check	Required ...	Comments
□	itso-cb-sys5.itso.ral.ibm.com	sysdumpdev_c...	⚠ 14%	Apr 23, 2013 ...	No	

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

Figure 7-14 Shows that the system is not compliant with defined rules (sysdumpdev)

5. Select the non-compliant host (`itso-cb-sys5.itso.ral.ibm.com`, shown in Figure 7-14) and click **View differences**. The differences are shown in Figure 7-15 on page 392.

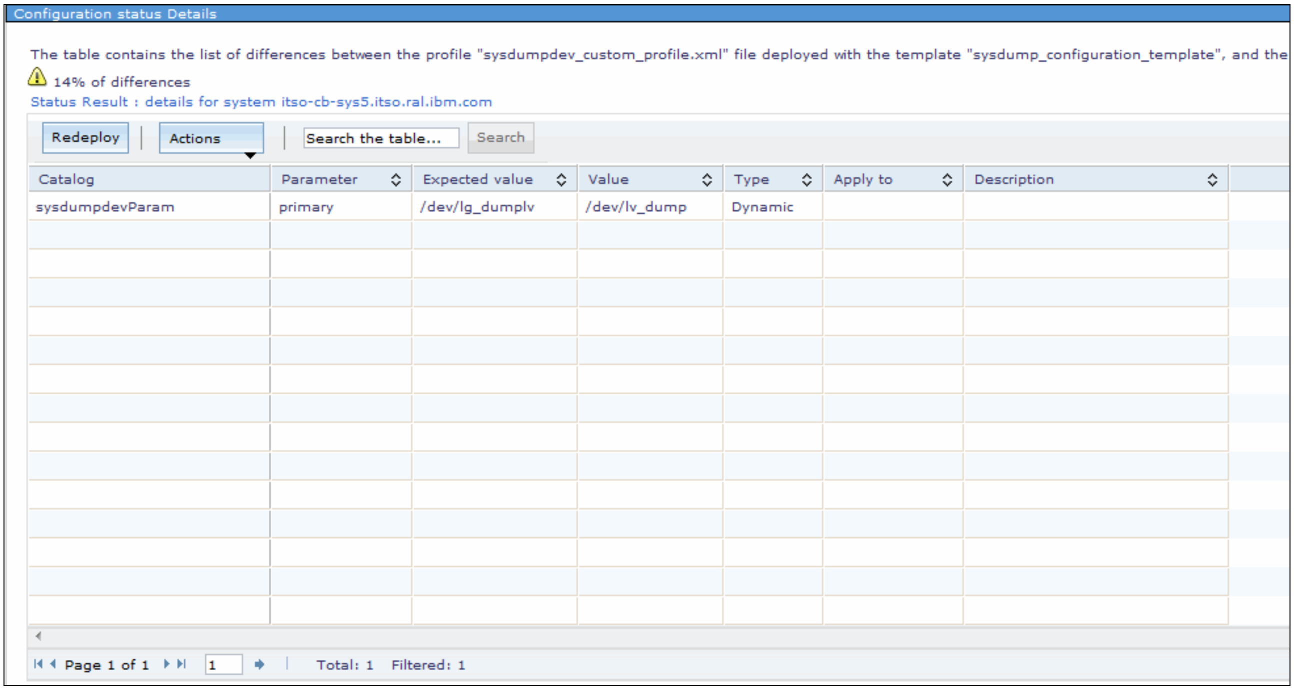

The table contains the list of differences between the profile "sysdumpdev_custom_profile.xml" file deployed with the template "sysdump_configuration_template", and the

⚠ 14% of differences

Status Result : details for system itso-cb-sys5.itso.ral.ibm.com

Catalog	Parameter	Expected value	Value	Type	Apply to	Description
sysdumpdevParam	primary	/dev/lg_dumplv	/dev/lv_dump	Dynamic		

Page 1 of 1 | 1 | Total: 1 Filtered: 1

Figure 7-15 Settings that are not compliant with the defined standard (sysdumpdev)

> **Warning:** If the Type column that is shown in Figure 7-15 is set to Dynamic, the value will be changed online, but some tuning settings require a reboot of the system. The Type column will be changed from Dynamic to Reboot.

6. From this point, it is possible to redeploy the template to correct the problem. Click **Redeploy** → **OK**, as shown in Figure 7-16.

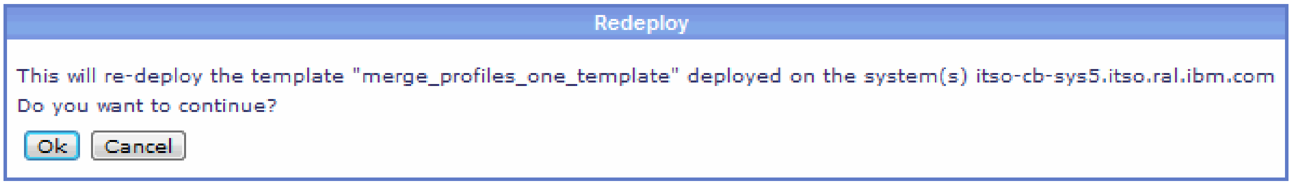

This will re-deploy the template "merge_profiles_one_template" deployed on the system(s) itso-cb-sys5.itso.ral.ibm.com

Do you want to continue?

Figure 7-16 Redeploy

7. After the deployment, the server is back to its normal situation as shown in Figure 7-17 on page 393.

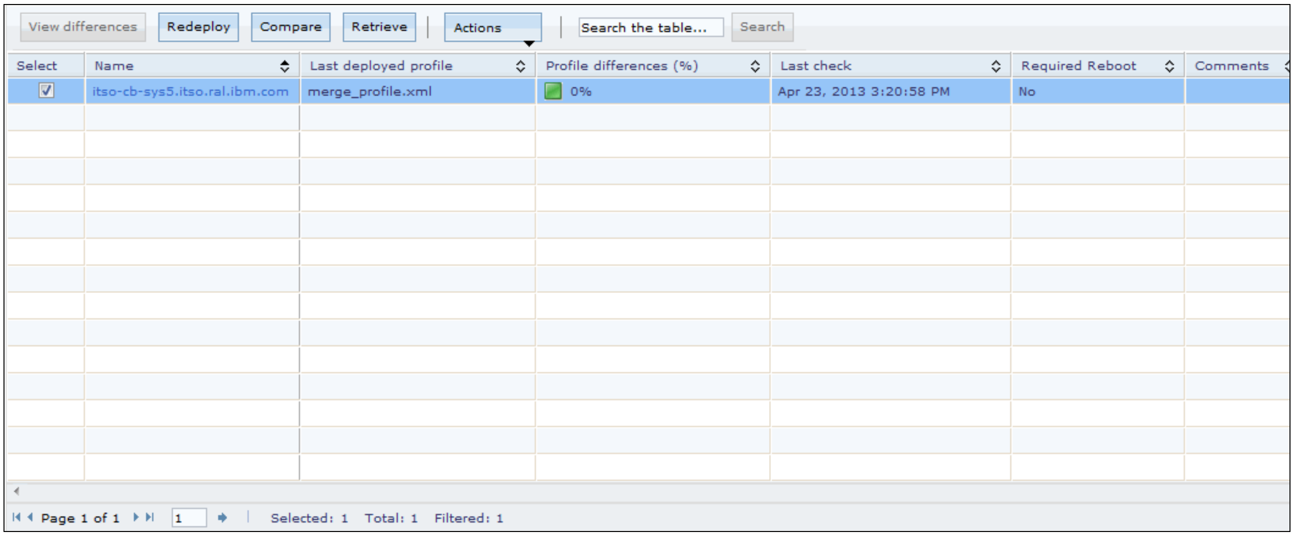

Select	Name	Last deployed profile	Profile differences (%)	Last check	Required Reboot	Comments
☑	itso-cb-sys5.itso.ral.ibm.com	merge_profile.xml	🟩 0%	Apr 23, 2013 3:20:58 PM	No	

Selected: 1 Total: 1 Filtered: 1

Figure 7-17 Profile differences for the host: itso-cb-sys5.itso.ral.ibm.com are shown at 0% compared to the deployed profile: merge_profile.xml

This process could help to massively and quickly correct nonstandard settings on several AIX servers.

> **Online content:** For a video demonstration of the preceding tasks, see the following video:
>
> `http://youtu.be/SnRxIEgeti4`
>
> You can also scan the QR code that is displayed in the left margin to go directly to the video.

Doing more with the AIX Profile Manager and artex commands

In the following example, we show how to merge several policies in only one XML file. We take as an example the values of `sysdumpdev profile` and `no profile` and put them into a new file. That last file is imported into AIX Profile Manager. This allows you to import only one profile, which contains all the desired configuration values into AIX Profile Manager.

> **Important:** At the time of writing, the existing option into AIX Profile Manager that allows you to merge profiles, was not working correctly. That is why you should use the command-line interface (CLI) and AIX Runtime Expert to create merged profiles.

Run the following examples to merge files from CLI on a first host, which is considered as standard in your architecture:

1. We create an XML file that contains the sysdumpdev configuration value and another XML file that contains the network settings that are managed by the **`no`** command, as shown in Example 7-11.

   Example 7-11 Create new XML files

   ```
   artexget -r /etc/security/artex/samples/sysdumpdevProfile.xml >
   /tmp/sysdumpdev.xml
   artexget -r /etc/security/artex/samples/noProfile.xml > /tmp/no.xml
   ```

2. Merge the two files into one and edit the file to delete the non-required lines, as shown in Example 7-12 on page 394.

Example 7-12 Merge the two files into a new one and remove the lines in red

```
cat sysdumpdev.xml no.xml >> /tmp/merge_profile.xml
vi /tmp/merge_profile.xml
<?xml version="1.0" encoding="UTF-8"?>
<Profile origin="get" version="2.0.1" date="2013-04-22T21:57:16Z">
 <Catalog id="noParam" version="2.1">
  <SubCat id="tcp_network">
   <Parameter name="tcp_nagle_limit" value="65535"/>
   <Parameter name="tcp_sendspace" value="16384"/>
  </SubCat>
 </Catalog>
 <Catalog id="sysdumpdevParam" version="2.1">
  <Parameter name="primary" value="/dev/lv_dump"/>
  <Parameter name="secondary" value="/dev/sysdumpnull"/>
  <Parameter name="copy_directory" value="/var/adm/ras"/>
  <Parameter name="forced_copy_flag" value="1"/>
  <Parameter name="always_allow_dump" value="0"/>
  <Parameter name="type_of_dump" value="fw-assisted" applyType="nextboot"
reboot="true"/>
  <Parameter name="full_memory_dump" value="disallow"/>
 </Catalog>
</Profile>
<?xml version="1.0" encoding="UTF-8"?>
<Profile origin="get" version="2.0.1" date="2013-04-22T21:57:09Z">
 <Catalog id="sysdumpdevParam" version="2.1">
  <Parameter name="primary" value="/dev/lv_dump"/>
  <Parameter name="secondary" value="/dev/sysdumpnull"/>
  <Parameter name="copy_directory" value="/var/adm/ras"/>
  <Parameter name="forced_copy_flag" value="1"/>
  <Parameter name="always_allow_dump" value="0"/>
  <Parameter name="type_of_dump" value="fw-assisted" applyType="nextboot"
reboot="true"/>
  <Parameter name="full_memory_dump" value="disallow"/>
 </Catalog>
</Profile>
<?xml version="1.0" encoding="UTF-8"?>
<Profile origin="get" version="2.0.1" date="2013-04-23T15:00:16Z">
...
<SubCat id="restricted">
   <Parameter name="extendednetstats" value="0" applyType="nextboot"
readOnly="true" reboot="true"/>
   <Parameter name="inet_stack_size" value="16" applyType="nextboot"
readOnly="true" reboot="true"/>
   <Parameter name="net_malloc_police" value="0" readOnly="true"/>
   <Parameter name="netm_affinity" value="0" applyType="nextboot"
readOnly="true" reboot="true"/>
   <Parameter name="pseintrstack" value="24576" readOnly="true"/>
   <Parameter name="use_isno" value="1" readOnly="true"/>
  </SubCat>
 </Catalog>
</Profile>
```

3. Save the new file.

4. Check the profile with the **artexset** command. Run the example that is shown in Example 7-13.

Example 7-13 Check the correct format of the new file

```
artexset -t /tmp/merge_profile.xml
Profile correctness check successful.
```

5. Import the file into AIX Profile Manager, create a template, and deploy on a second server.

6. On this second server, change the value of the sysdumpdev primary device, change the value of `tcp_sendspace` and `tcp_nagle_limit`, and then run a verification on the system. The result should be similar to what is shown in Figure 7-18.

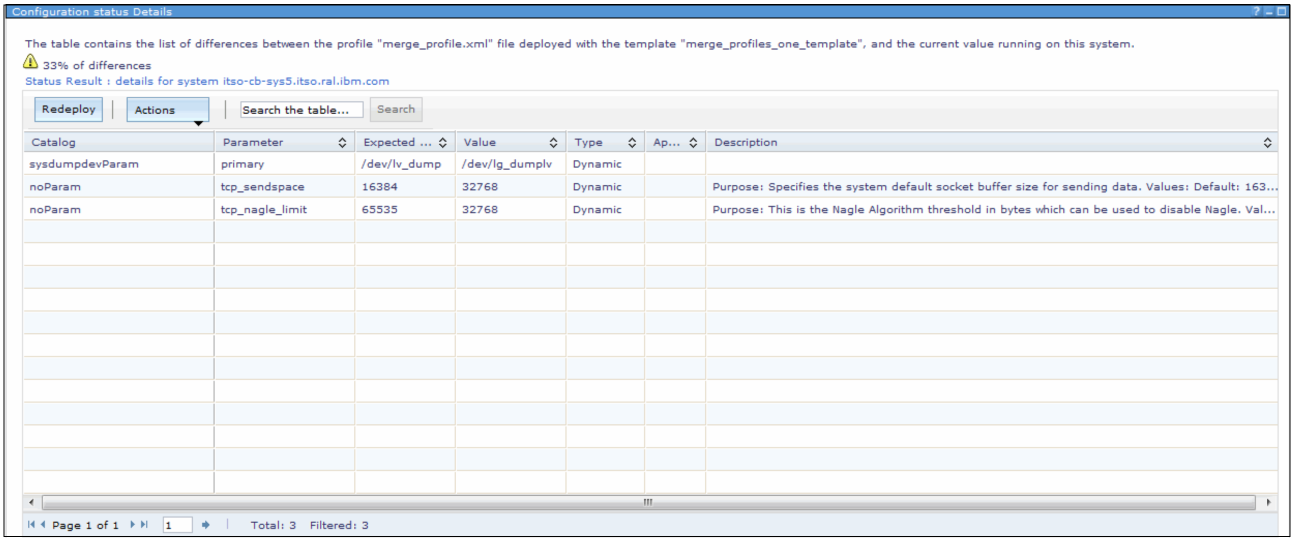

The table contains the list of differences between the profile "merge_profile.xml" file deployed with the template "merge_profiles_one_template", and the current value running on this system.

⚠ 33% of differences

Status Result : details for system itso-cb-sys5.itso.ral.ibm.com

Catalog	Parameter	Expected ...	Value	Type	Ap...	Description
sysdumpdevParam	primary	/dev/lv_dump	/dev/lg_dumplv	Dynamic		
noParam	tcp_sendspace	16384	32768	Dynamic		Purpose: Specifies the system default socket buffer size for sending data. Values: Default: 163...
noParam	tcp_nagle_limit	65535	32768	Dynamic		Purpose: This is the Nagle Algorithm threshold in bytes which can be used to disable Nagle. Val...

Figure 7-18 List of settings that are detected by APM as non-compliant with the default profile

Online Content: For a video demonstration of the preceding tasks, see the following video:

http://youtu.be/g-BD_f36Aaw

You can also scan the QR code that is displayed in the left margin to go directly to the video.

7.4 Event automation plan

In this chapter, we describe how to configure the IBM Systems Director and the AIX Profile Manager to send email in case of a critical event.

To configure the event automation plan feature into IBM Systems Director, perform the following steps:

1. From IBM Systems Director, click **Event Automation Plans** from the Automation menu, as shown in Figure 7-19.

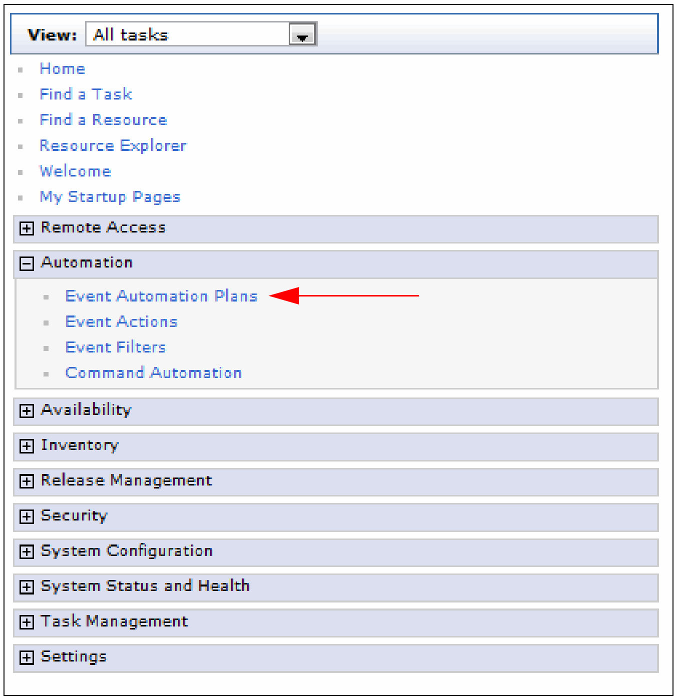

Figure 7-19 Click “Event Automation Plans”

2. Click **Create** → **Next**. Insert a name and description, and click **Next** as shown in Figure 7-20.

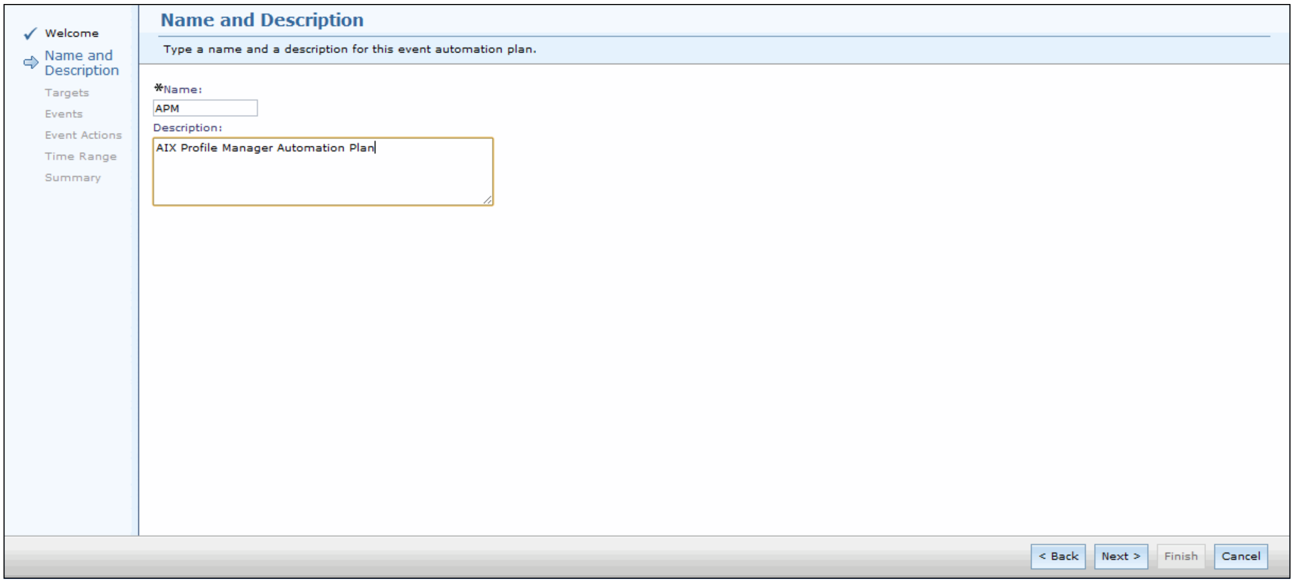

Name and Description

Type a name and a description for this event automation plan.

*Name: APM

Description: AIX Profile Manager Automation Plan

Figure 7-20 Definition and description of the event automation plan

3. Define “All Operating Systems” as target, as shown in Figure 7-21, and click **Next**.

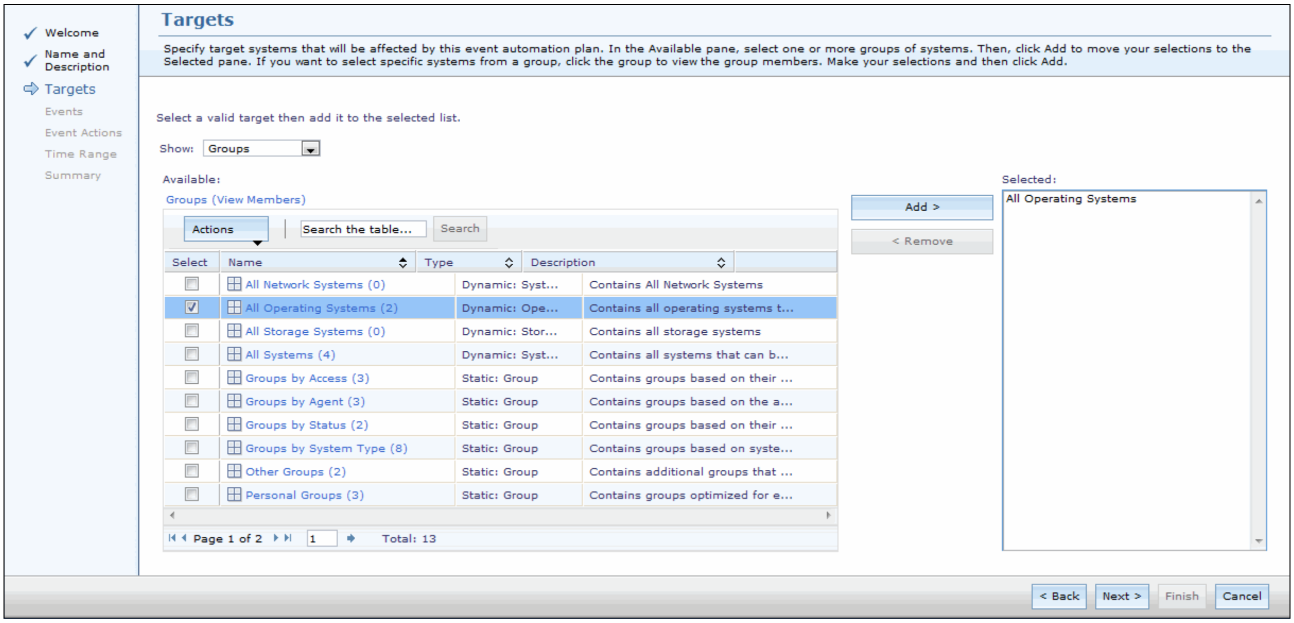

Targets

Specify target systems that will be affected by this event automation plan. In the Available pane, select one or more groups of systems. Then, click Add to move your selections to the Selected pane. If you want to select specific systems from a group, click the group to view the group members. Make your selections and then click Add.

Select a valid target then add it to the selected list.

Show: Groups

Available:

Groups (View Members)

Select	Name	Type	Description
□	All Network Systems (0)	Dynamic: Syst...	Contains All Network Systems
☑	All Operating Systems (2)	Dynamic: Ope...	Contains all operating systems t...
□	All Storage Systems (0)	Dynamic: Stor...	Contains all storage systems
□	All Systems (4)	Dynamic: Syst...	Contains all systems that can b...
□	Groups by Access (3)	Static: Group	Contains groups based on their ...
□	Groups by Agent (3)	Static: Group	Contains groups based on the a...
□	Groups by Status (2)	Static: Group	Contains groups based on their ...
□	Groups by System Type (8)	Static: Group	Contains groups based on syste...
□	Other Groups (2)	Static: Group	Contains additional groups that ...
□	Personal Groups (3)	Static: Group	Contains groups optimized for e...

Page 1 of 2 1 Total: 13

Selected:

All Operating Systems

Figure 7-21 Select group “All Operating Systems” as target

4. Select **Advanced Event Filters** in Events. Select **AIX Profile Manager events** in the list, as shown in Figure 7-22, and then click **Next**.

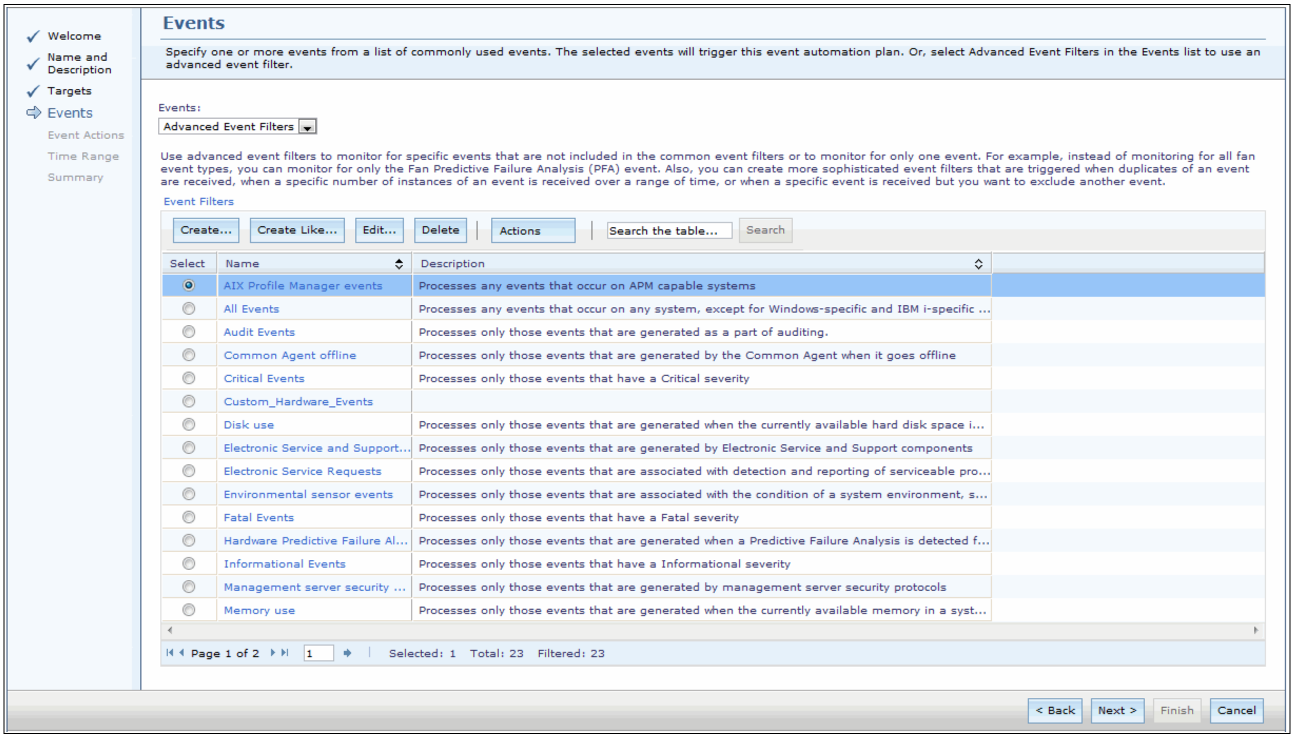

Events

Specify one or more events from a list of commonly used events. The selected events will trigger this event automation plan. Or, select Advanced Event Filters in the Events list to use an advanced event filter.

Events: Advanced Event Filters

Use advanced event filters to monitor for specific events that are not included in the common event filters or to monitor for only one event. For example, instead of monitoring for all fan event types, you can monitor for only the Fan Predictive Failure Analysis (PFA) event. Also, you can create more sophisticated event filters that are triggered when duplicates of an event are received, when a specific number of instances of an event is received over a range of time, or when a specific event is received but you want to exclude another event.

Event Filters

Select	Name	Description
●	AIX Profile Manager events	Processes any events that occur on APM capable systems
○	All Events	Processes any events that occur on any system, except for Windows-specific and IBM i-specific ...
○	Audit Events	Processes only those events that are generated as a part of auditing.
○	Common Agent offline	Processes only those events that are generated by the Common Agent when it goes offline
○	Critical Events	Processes only those events that have a Critical severity
○	Custom_Hardware_Events	
○	Disk use	Processes only those events that are generated when the currently available hard disk space i...
○	Electronic Service and Support...	Processes only those events that are generated by Electronic Service and Support components
○	Electronic Service Requests	Processes only those events that are associated with detection and reporting of serviceable pro...
○	Environmental sensor events	Processes only those events that are associated with the condition of a system environment, s...
○	Fatal Events	Processes only those events that have a Fatal severity
○	Hardware Predictive Failure Al...	Processes only those events that are generated when a Predictive Failure Analysis is detected f...
○	Informational Events	Processes only those events that have a Informational severity
○	Management server security ...	Processes only those events that are generated by management server security protocols
○	Memory use	Processes only those events that are generated when the currently available memory in a syst...

Page 1 of 2 1 Selected: 1 Total: 23 Filtered: 23

Figure 7-22 Select “AIX Profile Manager events” in the list

5. Click **Create**.

6. Select **Send an e-mail (Internet SMTP)**, as shown in Figure 7-23.

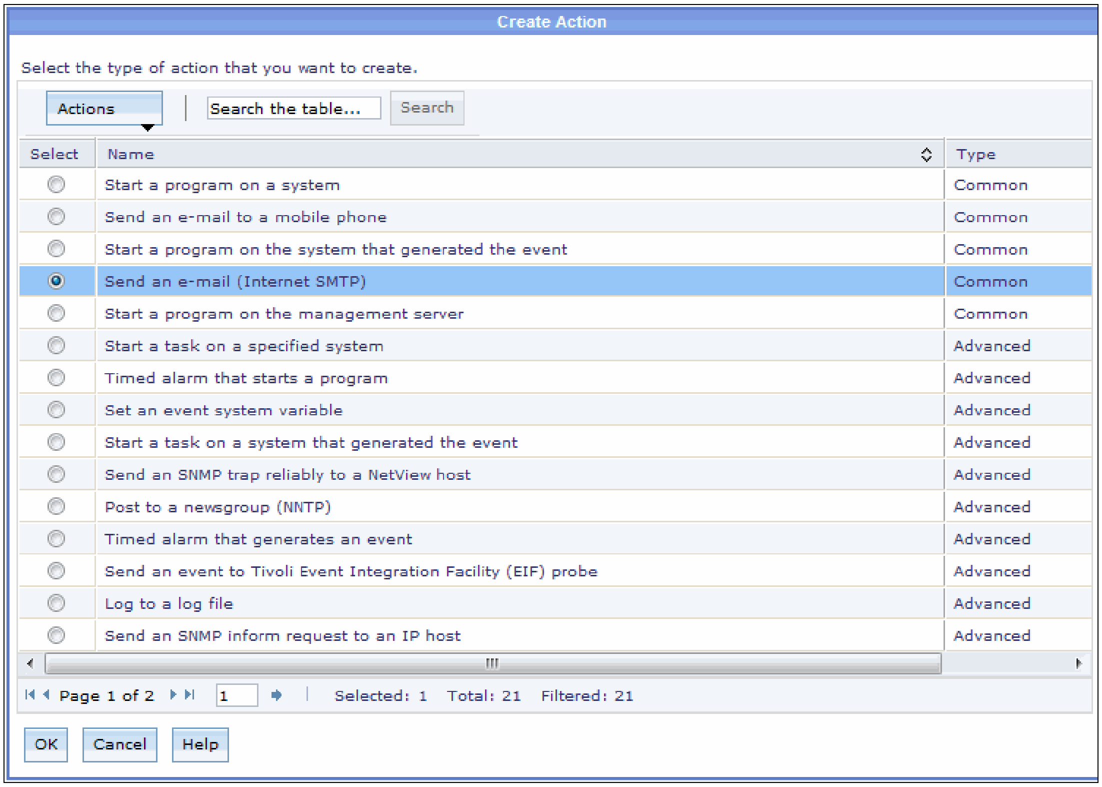

Create Action

Select the type of action that you want to create.

Select	Name	Type
○	Start a program on a system	Common
○	Send an e-mail to a mobile phone	Common
○	Start a program on the system that generated the event	Common
●	Send an e-mail (Internet SMTP)	Common
○	Start a program on the management server	Common
○	Start a task on a specified system	Advanced
○	Timed alarm that starts a program	Advanced
○	Set an event system variable	Advanced
○	Start a task on a system that generated the event	Advanced
○	Send an SNMP trap reliably to a NetView host	Advanced
○	Post to a newsgroup (NNTP)	Advanced
○	Timed alarm that generates an event	Advanced
○	Send an event to Tivoli Event Integration Facility (EIF) probe	Advanced
○	Log to a log file	Advanced
○	Send an SNMP inform request to an IP host	Advanced

Page 1 of 2 1 Selected: 1 Total: 21 Filtered: 21

Figure 7-23 Select "Send an e-mail (Internet SMTP)"

7. Complete the form as shown in Figure 7-24 with your desired settings to configure the email, then click **OK**.

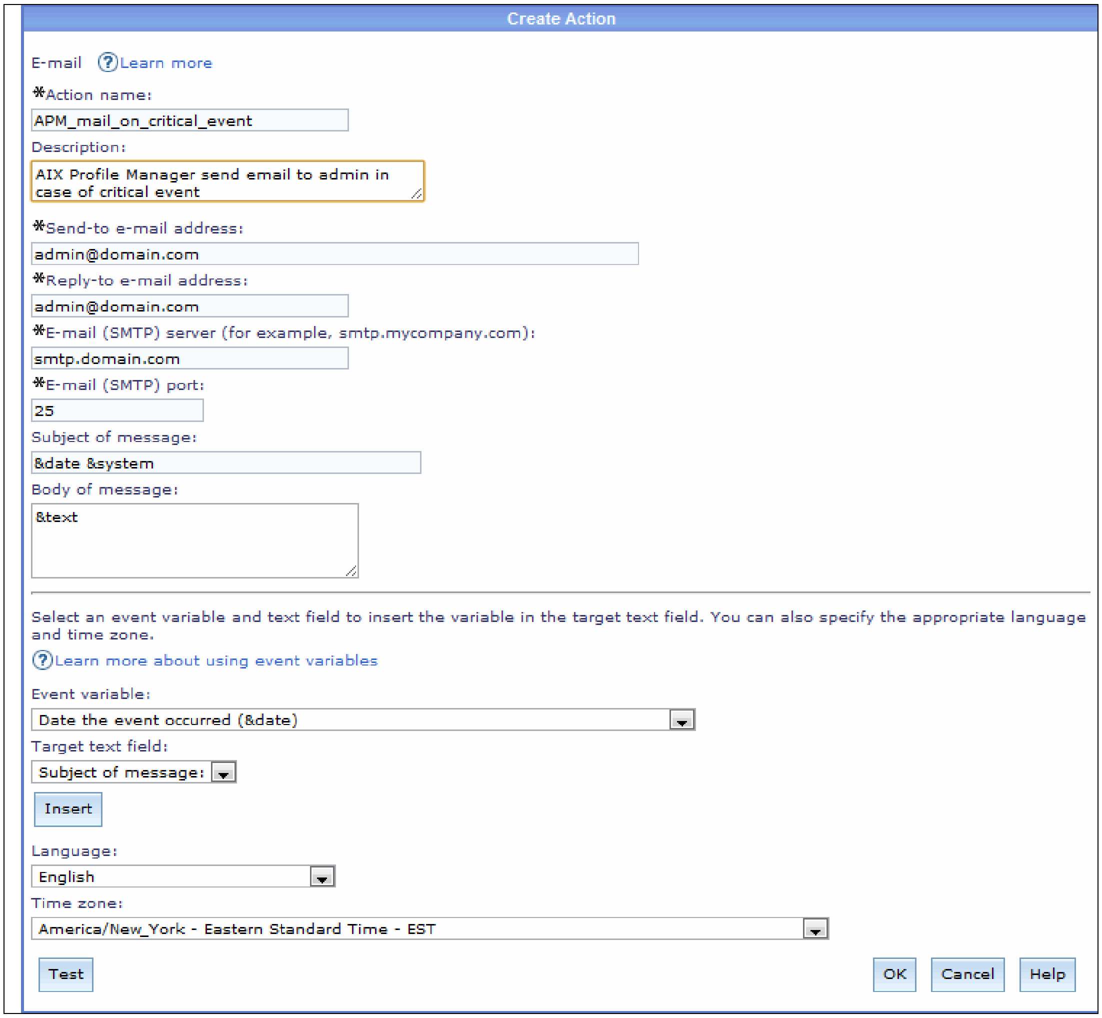

Create Action

E-mail ?Learn more

*Action name:

APM_mail_on_critical_event

Description:

AIX Profile Manager send email to admin in case of critical event

*Send-to e-mail address:

admin@domain.com

*Reply-to e-mail address:

admin@domain.com

*E-mail (SMTP) server (for example, smtp.mycompany.com):

smtp.domain.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:

Language:

English

Time zone:

America/New_York - Eastern Standard Time - EST

Figure 7-24 Complete the form for email configuration

8. Select the event action that has been created (in this example, `APM_mail_on_critical_event`) and click **Next**, as shown in Figure 7-25.

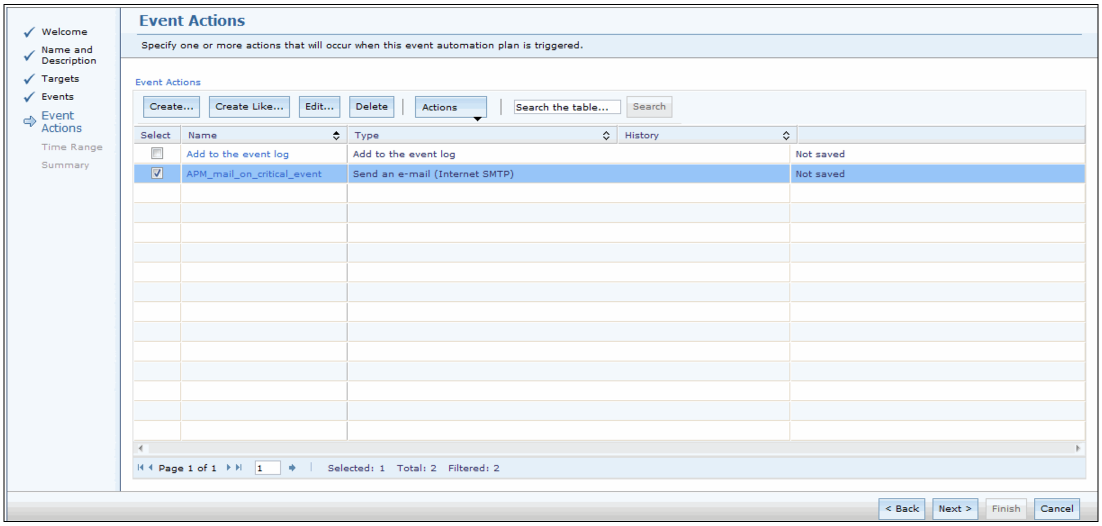

Event Actions

Specify one or more actions that will occur when this event automation plan is triggered.

Event Actions

Select	Name	Type	History
☐	Add to the event log	Add to the event log	Not saved
☑	APM_mail_on_critical_event	Send an e-mail (Internet SMTP)	Not saved

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

Figure 7-25 Event Actions panel

9. Select **All the time (24x7)** and then click **Next**, as shown in Figure 7-26.

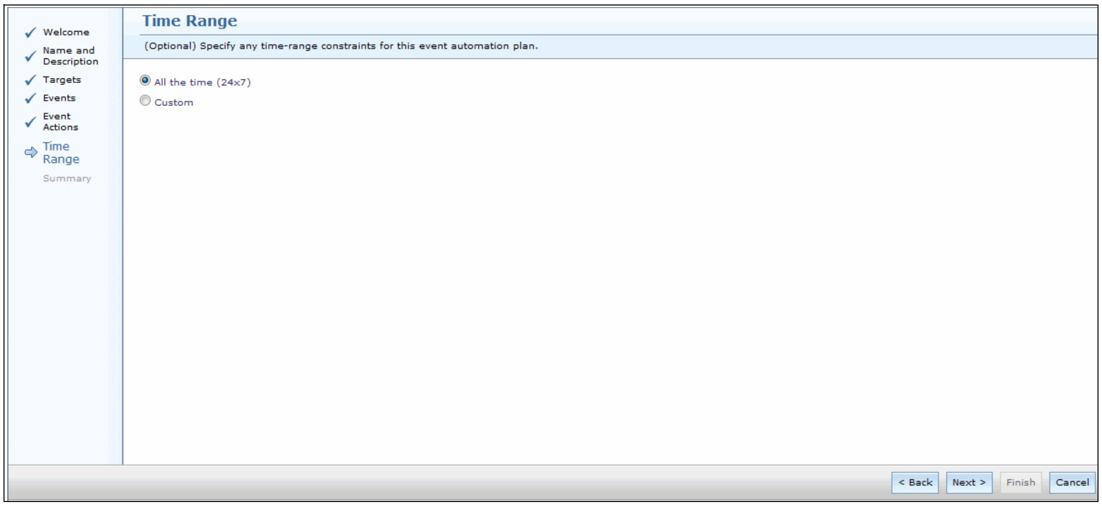

Time Range

(Optional) Specify any time-range constraints for this event automation plan.

● All the time (24x7)

○ Custom

Figure 7-26 Select the time range

10. Check the summary and click **Finish.**

11. The event automation plan is now displayed in the IBM Systems Director console, as shown in Figure 7-27.

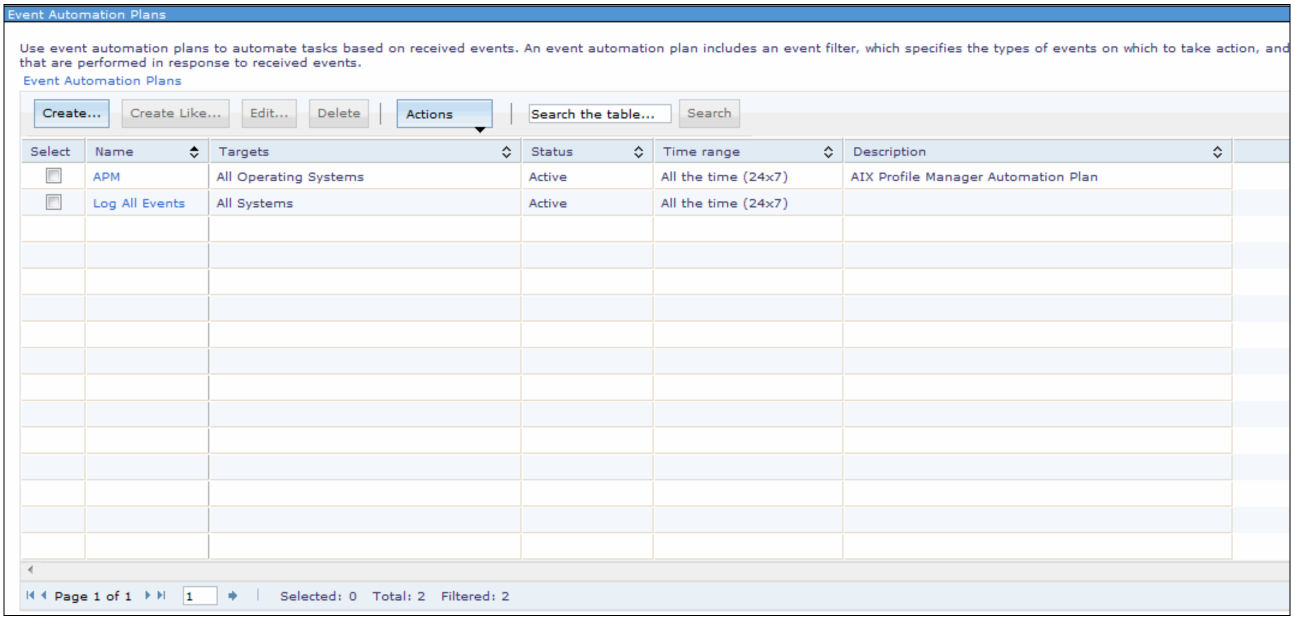

Use event automation plans to automate tasks based on received events. An event automation plan includes an event filter, which specifies the types of events on which to take action, and that are performed in response to received events.

Select	Name	Targets	Status	Time range	Description
□	APM	All Operating Systems	Active	All the time (24x7)	AIX Profile Manager Automation Plan
□	Log All Events	All Systems	Active	All the time (24x7)	

Figure 7-27 Event automation plans

7.5 Best practices

Listed in this section are the best practices for AIX Profile Manager:

- If your security standard is different from the models that are available in AIX Profile Manager, use IBM Security AIXPert. More information about AIXPert can be found here:

  http://pic.dhe.ibm.com/infocenter/aix/v7r1/topic/com.ibm.aix.security/doc/security/security_pdf.pdf

- You cannot import AIXPert XML files into AIX Profile Manager.
- Do not hesitate to make tests before implementing the security model that you want.
- Implement your security and system compliance step by step.
- If using one of the non IBM models (PCI, SOX, HIPPA, or DoD), you must have PowerSC installed on all systems.
- Do not modify the PCI, SOX, Cobit, HIPAA, and DoD security models.
- Use the artex commands to create your own profiles and then import them into AIX Profile Manager.
- Read the artex documentation before you start to use it:

  http://pic.dhe.ibm.com/infocenter/aix/v7r1/topic/com.ibm.aix.baseadmn/doc/baseadmndita/artex_concepts.htm

- Test your XML file with the **`artexset -t`** command to determine if the format is correct.
- Create groups into IBM Systems Director. It is easier to manage AIX Profile Manager with IBM Systems Director groups.
- Create an event automation plan for AIX Profile Manager. You can configure either to send email, to log a file, or to send SNMP messages to the monitoring server.

8

Workload Partition Manager

This chapter covers best practices and common questions to consider before implementing workload partitions (WPARs) with WPAR Manager.

In the first part, we briefly introduce the WPAR Manager plug-in for IBM Systems Director and cover the basics of installation.

The bulk of this chapter focuses on the best practices and the most important questions to ask before creating a WPAR and WPAR Manager infrastructure. We also show how you can manage and relocate WPARs using WPAR Manager graphical interface and the command-line interface (CLI). This chapter does not cover the step-by-step implementation of WPAR.

The end of this chapter includes recommended sites that you can read to get all the detailed procedures to implement WPARs.

The following topics are covered:

8.1 WPAR Manager

WPAR Manager is an IBM Systems Director advanced manager that allows you to manage multiple WPARs from a central location.

From a single access point, you can perform the following functions:

- Discover WPARs
- Create and delete WPARs
- Back up and restore WPARs
- Manage relocation domain groups and relocation policies
- Modify dynamically WPAR CPU and processor provisioning
- Move WPARs (Mobility)
- Clone WPARs
- Synchronize WPARs after an AIX upgrade
- Manage LPARs that host the WPARs, which include event manager, performance reports, compliance reports, and so on.

> **Information:** With IBM Systems Director and WPAR Manager, the LPARs used to host WPARs are called *managed systems*.

All these tasks can be done by using the graphical user interface of IBM Systems Director or the command-line interface. With command-line interface, you can create scripts for added flexibility.

> **Tip:** To see the commands that are available in IBM Systems Director for managing WPARs, use the **`smcli wparmgr help`** command.

With AIX 7.1, the following new features are supported:

- Fibre Channel support with MPIO drivers for disk and *sctape* or *atape* drivers for tapes.
- Trusted kernel extension support.
- Capability to host AIX 5.2 versioned WPARs.

> **Tip:** AIX 5.2 WPARs can be installed on an AIX 7.1 managed system only with an `mksysb` backup file.

8.2 Install IBM WPAR Manager

IBM WPAR Manager can be installed on IBM Systems Director server, which runs on Windows, Linux, or AIX operating systems. IBM WPAR Manager must be installed on the same server as IBM Systems Director.

The binaries can be downloaded from the following link:

http://www-03.ibm.com/systems/software/director/downloads/plugins.html

8.2.1 Prerequisites

The following prerequisites are needed to install WPAR Manager:

- At least 125 MB of free memory
- 5 MB in the root file system (/)

- Minimum of 180 MB in/var
- 15 MB in/opt
- For AIX and Linux, root privileges are required; and for Windows, administrative privileges are required

The following requirements are needed to install the WPAR Manager agent:

- AIX 6.1 or higher
- `mcr.rte` file set is installed (installed by default with AIX 7.1)
- 200 MB free in/var
- IBM Systems Director Common Agent installed

8.2.2 Installation

You can find the installation process of WPAR Manager at the following location:

http://pic.dhe.ibm.com/infocenter/director/pubs/topic/com.ibm.director.wparmgt.helps.doc/wparlpp_pdf.pdf

Attention:

- If you are in a production environment, do not specify an automatic restart of the IBM Systems Director server while installing WPAR Manager. Restart IBM Systems Director manually at the end of the installation by using the **`smstop`** and **`smstart`** commands. Check that no jobs are running in the *Tasks Management* module before restarting it.
- Remember to run discovery and inventory on the managed system after the WPAR manager installation. Discovery must be run from the IBM Systems Director console with the profile, *Extended WPAR Inventory*. Otherwise, IBM Systems Director will not detect the managed system as a WPAR-capable system.

8.3 Accessing WPAR Manager

When installation is done, you can access WPAR Manager in two different ways:

1. You can log in through the graphical interface of IBM Systems Director server and then click the **Plug-ins** tab and select **WPAR Manager**.

Figure 8-1 shows the WPAR Manager welcome page.

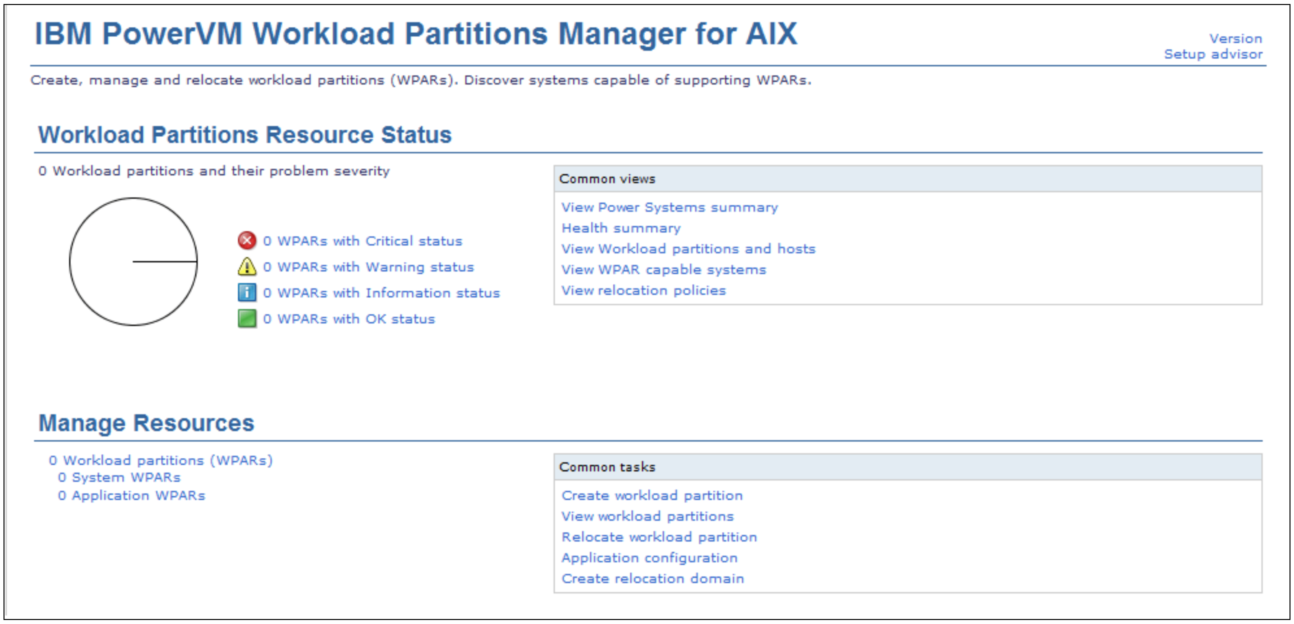

IBM PowerVM Workload Partitions Manager for AIX

Version
Setup advisor

Create, manage and relocate workload partitions (WPARs). Discover systems capable of supporting WPARs.

Workload Partitions Resource Status

0 Workload partitions and their problem severity

0 WPARs with Critical status
0 WPARs with Warning status
0 WPARs with Information status
0 WPARs with OK status

Common views

View Power Systems summary
Health summary
View Workload partitions and hosts
View WPAR capable systems
View relocation policies

Manage Resources

0 Workload partitions (WPARs)
0 System WPARs
0 Application WPARs

Common tasks

Create workload partition
View workload partitions
Relocate workload partition
Application configuration
Create relocation domain

Figure 8-1 Welcome page of WPAR Manager

2. You can use the command-line interface of the IBM Systems Director server. Run the **smcli wparmgr help** command to display all available options, as shown in Example 8-1.

Example 8-1 List of available commands for the CLI

```
smcli wparmgr help

The following WPAR Manager commands are supported:

- help - Display this help information
- lswpar - List all WPARs, or properties of specific WPARs
- mkwpar - Create a WPAR
- rmwpar - Delete a WPAR
- chwpar - Modify a WPAR
- deploywpar - Deploy a WPAR definition
- startwpar - Start a WPAR
- stopwpar - Stop a WPAR
- savewpar - Save a WPAR
- restwpar - Restore a WPAR
- movewpar - Relocate a WPAR
- lswparcapablesys - List WPAR-capable systems, its WPARs, and device
       information
- lswparcompat - View compatibility results for a WPAR
- syncwpar - Synchronize a WPAR
- clonewpar - Clone a WPAR
- lswparpolicy - View WPAR relocation policies
- mkwparpolicy - Create a WPAR relocation policy
- chwparpolicy - Modify a WPAR relocation policy
- rmwparpolicy - Delete a WPAR relocation policy
- lsrelocdomain - View a WPAR relocation domain
- mkrelocdomain - Create a WPAR relocation domain
- chrelocdomain - Modify a WPAR relocation domain
- rmrelocdomain - Delete a WPAR relocation domain
- lswparmgrsettings - Show WPAR Manager settings
- chwparmgrsettings - Modify WPAR Manager settings
- lswparmgrinfo - View WPAR Manager information (such as version and
                  license information)

Type "smcli <command_name> -h" to see the usage statement of the command.
Type "smcli <command_name> --help" for a full description of the command and
       its usage
```

The command-line interface provides you the ability to create scripts to automate numerous operations, such as massive deployment, backup operations, and relocations.

8.4 Managing WPARs' infrastructure with WPAR Manager

All operations for managing WPARs with IBM WPAR Manager are described in the IBM Redbooks document *Exploiting IBM AIX Workload Partitions*, SG24-7955, in section 4.5:

http://www.redbooks.ibm.com/abstracts/sg247955.html

8.4.1 WPARs

In this chapter, we cover different scenarios to show that requirements and the overall architecture can change a lot depending on what you want to achieve.

Creation of WPARs

Although the process to create a WPAR is simple, you still need to consider these questions first:

- Do you need a system WPAR or an application WPAR?
- Do you need a private `/usr` and `/opt`?
- Where will the file systems be located? Local? Network File System (NFS)? Storage area network (SAN)?
- Will you use relocation?
- Will your WPAR have its private storage?
- What about versioned WPARs?
- Can you create a copy of an existing WPAR?

Depending on the answers to those questions, the way to install WPARs varies.

Do you need a system WPAR or an application WPAR?

Answering this question is easy if you consider the following factors:

- Application WPAR: An application WPAR is generally used for one process or one job and is a temporary WPAR. When the jobs or the script ends, the WPAR is automatically removed from the managed system. It has no private file systems, users, or network configuration. This kind of WPAR is used for High Computing Clusters.
- System WPAR: The system WPAR can host your application, have its private file systems and users, have a network configuration, and its own service such as cron. The system WPAR uses more resources compared to the application WPAR.

Do you need private /usr and /opt?

Using read-only /usr and /opt is a good solution because it saves space and reduces the risk of human error if someone tries to remove a file or directory in those file systems from a WPAR. Many applications work without any problem in this case, but some application requirements force you to reconsider this option, such as:

- An application needs to write in /usr or /opt, or both
- An application needs a specific file set that is not installed on the managed system

In those cases, you can perform the following functions:

- Create a local file system in the WPAR for the application, such as /usr/local.
- Grant write access to the file system (this solution is a risk because all other WPARs in the same managed system will see the new files created in /usr).
- Create the WPAR with its own /opt and /usr. This is the recommended solution. It will solve the access problem but will increase the disk size (approximately 2 GB for private /usr and /opt, compared to 450 MB for non-private /usr and /opt) and the WPAR will also use more memory.

Tip: To create a system WPAR from the IBM Systems Director command line with private /usr and */opt*, use the **smcli mkwpar** command with the **-l** option (or **--private_usr_opt**).

Where will the file systems be located?

If you do not want to use relocation options, you can use standard shared file systems. They will be created in the managed system volume group (VG).

If you want to use the relocation feature, select one of the following two options:

- NFS
- SAN

If you choose NFS, you need an NFS server that hosts all the file systems of your WPARs.

To create an NFS WPAR server, follow the instructions in section 8.8.1 of the IBM Redbooks publication, *Exploiting IBM AIX Workload Partitions*, SG24-7955.

As a best practice, configure a dedicated NFS server. It also should not be located on the IBM Systems Director server. This NFS server should have sufficient disk space (if WPARs have private /usr and /opt), CPU, memory, and network resources to cover the performance requirements. The more NFS WPARs that you have, the more hardware resources you will need.

Warning: WPARs that use IPv6 must use NFSv4.

Tip: If you want to relocate NFS WPARs, remember to enable access for all managed servers to the NFS shares.

If you use SAN, there is no need to configure an additional server. The managed systems should have the SAN disks configured with drivers and when creating the new WPARs, you should then specify the disk name where you want to configure the WPAR.

Tip: If you want to relocate SAN WPARs, remember to configure all managed systems to access the SAN disks.

Will you use relocation?

There are a couple of things to consider before using relocation:

- Both managed systems (the current and the target) used for mobility must be in the same network subnet.
- The WPARs must have been created with the checkpointable feature.

> **Tip:** Use the **`smcli lswpar -G -n WPAR_NAME`** command to see if the checkpointable feature is active or not. The option to use when creating the WPAR is **`-c`** when using the command-line interface.

- The WPARs must be moved to a managed system with compatible hardware and software.
- NFS file systems or SAN disks must be accessible on both managed systems.

What about versioned WPARs?

If you want to use versioned WPARs, consider the following advice before implementing it.

Why should you use versioned WPARs?

Listed are some examples where versioned WPARs can be used:

- An application that is installed on an AIX 5.2 or 5.3 cannot be migrated to AIX 6.1 or 7.1
- Eliminate the old hardware and consolidate all old AIX servers into one managed system, which hosts those old AIX in WPARs. This helps reduce the hardware maintenance cost and reduce the electric and cooling costs.
- Boost the performance of your old AIX 5.2 and AIX 5.3 servers by running them on new hardware such as IBM Power7 processors.
- Benefits from the latest Power7 and PowerVM features (AME, SMT4, and so on).

What are the basic requirements for versioned WPARs?

The requirements to host AIX 5.2 and 5.3 WPARs are shown in Table 8-1.

Table 8-1 Requirements for using IBM AIX 5L™ WPARs

	AIX 5.2	**AIX 5.3**
Minimum level of AIX	TL10 SP8	TL 12 or higher
vwpar.images level	1.1.2	1.1.2
Managed system minimum level	AIX 7.1	AIX 7.1 TL1
Hardware required	IBM Power7	IBM Power7

> **Tip:** You need an mksysb backup to create versioned WPARs. The versioned WPAR always owns its `/usr` and `/opt` file systems. You cannot use shared file systems. The option that is used to indicate which mksysb to use in the **`smcli mkwpar`** command is **`-B`** *mksysb_image_name*.

> **Important:** A separate licensed program is required to run versioned WPARs in your environment.

Versioned WPARs and Live Application Mobility

There are some prerequisites before using Live Application Mobility with versioned WPARs:

- The WPARs must have the software vwpar.images 1.1.2 or higher installed.
- On AIX 5.2, the following APARs must be installed:
  - APAR IZ72315
  - APAR IZ90201
- On AIX 5.3, the following APAR must be installed: APAR IZ89583

More information about versioned WPARs requirements and limitations can be found at the following site:

http://pic.dhe.ibm.com/infocenter/aix/v7r1/index.jsp?topic=%2Fcom.ibm.aix.wpar%2Fcreconfig-create-wpar.htm

Will your WPAR have its own private storage?

For different reasons, you might want to configure and manage storage directly on the WPARs. This is possible but there are some questions, limitations, and prerequisites to consider:

> **Attention:** Storage allocation is not supported for application WPARs.

- Will you configure the disks on the managed system and then export them to the WPAR?

  The disks are configured on the managed systems, then exported to WPAR. Using this method allows you to manage several disks for several WPARs while using a minimum of Fibre Channel adapters. It also allows you to manage allocation of the disks directly from the managed systems. You can also easily check the disk configuration of each WPAR by using the **`lswpar -D`** command, which gives you a complete view of devices that are allocated to WPARs.

> **Tip:** Use the **`smcli chwpar`** command from the IBM Systems Director command line to manage devices on WPARs. The **`smcli chwpar --help`** command provides information about available options.

- Will you directly configure the Fibre Channel adapter into the WPAR?

  This method allows the WPAR administrator to manage its own disk configuration. The disks are not viewed by the managed system.

  This feature includes serious limitations:

  - The Fibre Channel adapter is not usable by other WPARs
  - Live Application Mobility is not supported
  - This method is not supported for versioned WPARs

> **Important:** Only fiber-attached devices, FC adapters, and virtual SCSI (vSCSI) devices are supported in WPARs. All disks that are managed by the MPIO subsystem and supported by AIX are also supported on WPARs.

Is it possible to create WPAR from an existing one?

With WPAR Manager, you can create clones of existing WPARs. Data on external devices is not copied (if external devices contain rootvg information, this data is copied to a specified external device).

You can create a clone by using the graphical interface or the command line.

Tips:

- You can clone local WPAR, NFS WPAR, RootVG WPAR, and versioned WPAR.
- For your deployment, you can use the clone of a standard WPAR that you created and then use it to deploy new WPARs.

8.4.2 Relocation

Warning: Live Application Mobility is not supported between hardware that is not from the same processor family. For example, you cannot migrate a WPAR that is running on IBM Power6 to an IBM Power7 server.

Do you want automatic relocation?

If you want to use the features of automatic relocation, WPAR Manager automatically manages the relocation that is based on relocation domain groups that you created and on relocation policies.

Relocation domain groups

You must create relocation domain groups in IBM Systems Director. A relocation domain group restricts the number of servers to which WPARs can be moved automatically by WPAR Manager.

As an example you can create the following groups:

- IBM_PROD_P7, which contains the managed systems where critical WPARs are hosted.
- IBM_DEV_P6, which contains the managed systems where your non-critical production is hosted.

By creating these two groups, you can avoid a relocation of one of your critical WPARs on an older IBM Power Systems server and avoid a performance issue.

You can manage the relocation domain groups with the IBM Systems Director command-line interface for your scripts, as shown in Example 8-2.

Example 8-2 Commands to manage relocation domain groups

```
smcli mkrelocdomain -m itso-cb-sys5.itso.ral.ibm.com -n IBM_PROD_P7
Create relocation domain - IBM_PROD_P7
======================================

smcli lsrelocdomain -D
===================================
Relocation domain name: IBM_PROD_P7
===================================

Description: -
Membership: itso-cb-sys5.itso.ral.ibm.com
Relocation policy: -
```

Warnings:

- A managed system can be in only one relocation domain group.
- If you try to manually relocate a WPAR that is hosted in a managed system that is a part of a relocation domain group, IBM Systems Director can again move the WPAR based on the relocation policies.

Relocation policies

When you create relocation domains, you must configure relocation policies. These relocation policies allow the WPAR Manager to know when it should start using relocation with your different WPARs.

The relocation policies are based on three important parameters:

- Period: In minutes. If you specify 20 minutes, WPAR Manager starts to relocate WPARs if an event happens during at least 20 minutes.
- CPU Threshold: If you set 80%, WPAR Manager starts to relocate WPARs if a managed system uses more than 80% of the CPU.
- Memory Threshold: If you 80%, WPAR Manager starts to relocate WPARs if a managed system uses more than 80% of the available memory.

In this example, WPAR Manager will not relocate WPARs if all managed systems use 80% or more of CPU and memory.

Tip: Do not specify values that are too low or WPAR Manager starts relocating WPARs too often.

You can manage relocation policies with the command-line interface from IBM Systems Director server, as shown in Example 8-3.

Example 8-3 Command to manage relocation policies

```
smcli mkwparpolicy -a -d IBM_PROD_P7 -p 30 -c 80 -m 80 -n IBM_PROD_P7_Reloc_Policy
Create policy - IBM_PROD_P7_Reloc_Policy
========================================
Enable policy - IBM_PROD_P7_Reloc_Policy
========================================

Managed system: itso-cb-sys5.itso.ral.ibm.com

smcli lswparpolicy
Name                     Auto Reloc Avg Period CPU Util Mem Util Reloc Domains
-----------------------------------------------------------------------------
IBM_PROD_P7_Reloc_Policy Yes        30         80.0     80.0     IBM_PROD_P7


smcli rmwparpolicy -n IBM_PROD_P7_Reloc_Policy
Delete policy - IBM_PROD_P7_Reloc_Policy
========================================
Disable policy - itso-cb-sys5.itso.ral.ibm.com
==============================================

Managed system: itso-cb-sys5.itso.ral.ibm.com
```

```
smcli lsrelocdomain -D
===================================
Relocation domain name: IBM_PROD_P7
===================================

Description: -
Membership: itso-cb-sys5.itso.ral.ibm.com
Relocation policy: IBM_PROD_P7_Reloc_Policy
```

Do you want to do use manual relocation?

You can decide to use manual relocation. Manual relocation must be done from WPAR Manager. It can be done from the graphical user interface (GUI) or from the command-line interface with the **smcli movewpar** command. You have to decide several things before running the manual relocation:

- On which managed system do you want to relocate your WPAR?
- Does this managed system meet all the requirements to allow a relocation of your WPAR?
- Do you want to use static or live relocation?

Before trying a manual relocation, check the compatibility between the source system and the target system. This can be done from the WPAR Manager graphical interface with a right-click on the WPAR and selecting **Compatibility**, as shown in Figure 8-2.

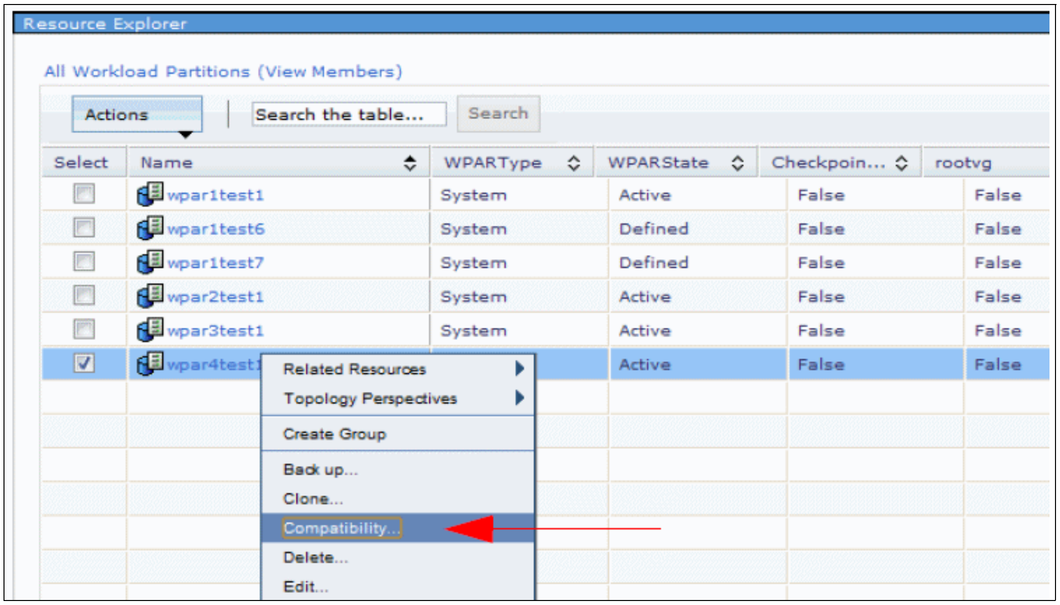

Select	Name	WPARType	WPARState	Checkpoin...	rootvg
□	wpar1test1	System	Active	False	False
□	wpar1test6	System	Defined	False	False
□	wpar1test7	System	Defined	False	False
□	wpar2test1	System	Active	False	False
□	wpar3test1	System	Active	False	False
☑	wpar4test1		Active	False	False

Figure 8-2 Checking the compatibility between managed systems

> **Tip:** You can also use the command-line interface by running **smcli lswparcompat -n** *wpar_name.*

Consider the following compatibility checks:

- Managed systems must have the same level of operating system (technical level and service pack).
- The processors class on the arrival system must be at least as high as the processor class on the departure system.
- The `bos.rte`, `bos.wpars`, and `mcr.rte` file sets must be strictly on the same release on both managed systems.
- The `bos.rte.libc` file set must be the same on both managed systems.
- Allocated and exported devices must be available on the target managed system and must not be used by another WPAR.
- Ensure that versioned WPARs meet all requirements that are mentioned in the section: *Versioned WPARs and Live Application Mobility*.

You can also decide to add some criteria on the WPARs to avoid issues when running relocation. These options can be edited in the WPAR's properties. From the WPAR Manager interface, right-click the WPAR and click **Edit**. Go to the **Advanced settings** tab and select the **Compatibility** tab. You can then choose one of the following options, as shown in Figure 8-3.

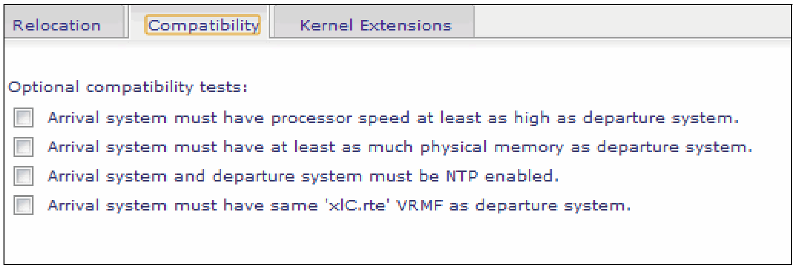

Optional compatibility tests:

☐ Arrival system must have processor speed at least as high as departure system.

☐ Arrival system must have at least as much physical memory as departure system.

☐ Arrival system and departure system must be NTP enabled.

☐ Arrival system must have same 'xlC.rte' VRMF as departure system.

Figure 8-3 List of additional options that are available for relocation compatibility checks

8.4.3 Patches and upgrades

Patches allow you to keep your environment stable and to benefit from new features for your environment. WPAR Manager should be kept up to date every time.

Patch WPAR Manager

Upgrade WPAR Manager each time a new release is available on the IBM website. Use the IBM Systems Director release management feature to stay up to date.

> **Tip:** If WPAR Manager is installed on AIX, do not hesitate to upgrade the operating system.

Patch WPAR Manager agent

It is not necessary to immediately patch the WPAR Manager agent after an upgrade of WPAR Manager, but you should nevertheless try to stay up to date.

Patch managed systems and WPARs

For AIX maintenance, such as technical level and service packs, you must upgrade the managed systems first before any WPARs.

Attention: AIX upgrades cannot be managed by WPAR Manager or IBM Systems Director. You have to run the AIX upgrade directly from the managed system or from a Network Installation Management (NIM) server.

Shared WPARs

When the managed system has been upgraded, synchronize the WPARs. This is necessary to allow the WPARs to use the latest release of the AIX kernel.

Attention: This does not apply to versioned WPARs. The synchronize option is not available for application WPARs.

The next example shows the difference between a classical LPAR architecture and a WPAR architecture when upgrading the AIX, as shown in Table 8-2.

Table 8-2 AIX upgrade comparison between LPARs and WPARs

	40 LPARs	**40 WPARs on one LPAR**
Number of AIX upgrades	40	1
Number of reboots	40	1
Time of maintenance	40 x X minutes (upgrade AIX) + 40 x X minutes (reboot)	1 x X minutes (upgrade AIX) 1 x X minutes (reboot) + 40 x X minutes (WPARs sync) No reboot needed for the sync.

Tip: The synchronization step can be done from the WPAR Manager graphical interface or by the command line with the **`smcli syncwpar`** command. This allows you to run only one command to synchronize all WPARs.

Attention: The synchronization between managed systems and WPARs is not supported for versioned WPARs, application WPARs, and undeployed WPARs.

RootVG WPARs

Because RootVG WPARs have their own `/usr` and `/opt`, they have to be upgraded like the global instance.

Tip: While upgrading AIX on the managed systems, you can decide if you want to upgrade WPARs in the same time with the **`smit install_all`** command.

For more information about patch and file set management for WPARs, read section 10.1 in the IBM Redbooks publication, *Exploiting IBM AIX Workload Partitions*, SG24-7955.

http://www.redbooks.ibm.com/abstracts/sg247955.html

8.4.4 Backup

Using the **`smcli savewpar`** command from IBM Systems Director can help you prepare scripts to back up all your WPARs from a single command and in a single location.

All backups can be stored in an NFS mount point that will be secured by backup software after the **savewpar** command has been run, as shown in Example 8-4.

Example 8-4 Back up all WPARs in directory /backup (NFS share mounted on all managed systems)

```
smcli lswpar  |grep Active |awk '{print $5}' |while read hostname ; do smcli
savewpar -n $hostname -i -M -F /nfs/backup/$hostname.bff ; done
```

You can also use this mount point to restore the images.

8.5 Best practices

Listed in this section are the best practices for WPAR Manager:

- If you want to use versioned WPARs, you must have IBM Power7 servers and AIX 7.1 servers as managed systems.
- If you want to use versioned WPARs, ensure that you meet all technical requirements defined in this chapter and ensure that you have acquired the necessary license programs.
- Relocation is possible only from the WPAR Manager graphical interface or from the IBM Systems Director command-line interface.
- Relocation is not possible on hardware that is not from the same processor family.
- Relocation is not possible between managed systems that do not have the same operating system version.
- Before using automatic relocation, define relocation domain groups and relocation policies on IBM Systems Director server.
- If you want to use NFS shared WPARs with automatic relocation, configure a separate NFS server and configure all managed systems to access it.
- If you want to use SAN WPARs with automatic relocation, configure disks on all managed systems that are in the relocation domain group.
- If you choose to do a manual relocation, do a compatibility check before the relocation.
- For easy deployment, create a clone from an existing WPAR and use it to create other WPARs.
- Define a backup policy and configure a central location where the backups will be stored.
- Check the WPAR logs on managed systems when you have problems with WPARs.
- Read “Additional documentation” to have all the information and technical how-to documentation that you need.

8.6 Additional documentation

IBM Redbooks *Exploiting IBM AIX Workload Partitions*, SG24-7955:

http://www.redbooks.ibm.com/abstracts/sg247955.html

IBM PowerVM Workload Partition Manager for AIX:

http://pic.dhe.ibm.com/infocenter/director/pubs/index.jsp?topic=%2Fcom.ibm.director.wparmgt.helps.doc%2Fwparlpp-kickoff.html

WPAR AIX 6.1 Information center:

http://pic.dhe.ibm.com/infocenter/aix/v6r1/index.jsp?topic=/com.ibm.aix.wpar/wpar-kickoff.htm

WPAR 7.1 information center:

http://pic.dhe.ibm.com/infocenter/aix/v7r1/index.jsp?topic=/com.ibm.aix.wpar/wpar-kickoff.htm

IBM developerWorks: WPAR in AIX 6.1:

http://www.ibm.com/developerworks/aix/library/au-wpar61aix

IBM developerWorks: WPAR in AIX 7.1:

http://www.ibm.com/developerworks/aix/library/au-wparaix7

IBM developerWorks: FC Adapter-based WPAR creation with Oracle database configuration:

http://www.ibm.com/developerworks/aix/library/au-fc_adapter_wpar

IBM developerWorks: Creation and relocation of system WPARs with SAN-based data model:

http://www.ibm.com/developerworks/aix/library/au-wpar_san

Creating a versioned WPAR information center:

http://pic.dhe.ibm.com/infocenter/aix/v7r1/index.jsp?topic=%2Fcom.ibm.aix.wpar%2Fcreconfig-create-wpar.htm

9

Network Control

This chapter covers the basic functionalities of Network Control, how to plan for Network Control deployments, and describes a number of common scenarios with best practices.

The following topics are covered:

9.1 Overview of Network Control

IBM Systems Director Network Control is an optional and chargeable plug-in for IBM Systems Director that extends the basic capabilities available in IBM Systems Director Network Manager. Network Control is a network management tool that can help discover, configure, and automate network components.

Note: Network Control is always included in the Flex System Manager (FSM) appliance that is available in IBM Flex System and IBM PureFlex configurations.

Following are the basic functionalities of Network Control:

- Physical and virtual network switches discovery and monitoring: It uses Simple Network Management Protocol (SNMP) v1, v2, or v3 to get health status from switches. Events generated can then be used in event action plans and forwarded to your event management system.
- Protocol and VLAN configuration of physical and virtual network devices: You can create VLANs on discovered devices and assign those VLANs to specific ports. You can also create port groups on virtual switches.
- Depending on your Ethernet switch type, you might also be able to use Network Control to configure Converged Enhanced Ethernet (CEE) networking and Edge Virtual Bridging (EVB).
- Network topology views for physical and virtual network devices: You can view systems by subnets and VLANs, and you can generate graphical topology views that show how switches are inter-connected and which physical and virtual servers are connected to them. Figure 9-1 on page 421 shows an example of a discovered network topology view that shows the inter-switch link (ISL) between two managed switches.

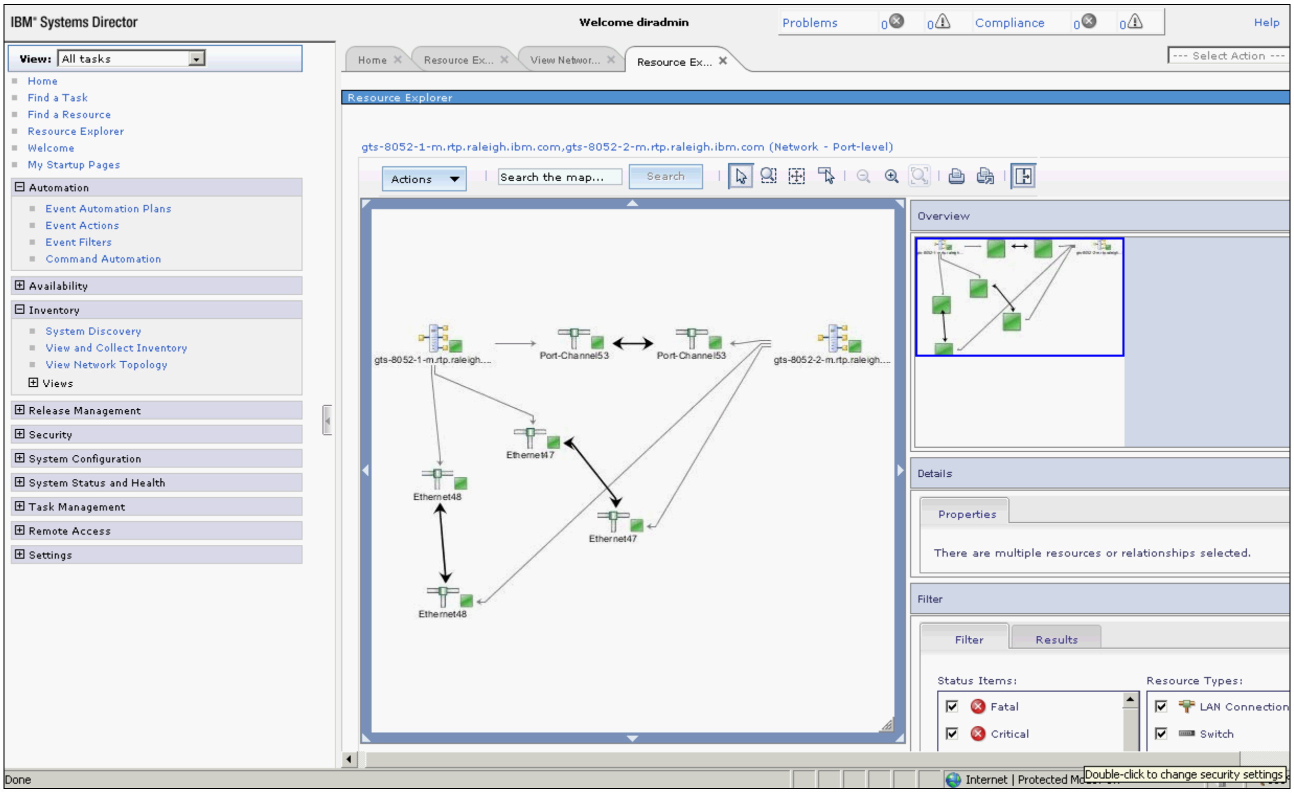

There are multiple resources or relationships selected.

Figure 9-1 Example of a discovered network topology view

Figure 9-2 is an example of the Systems by VLAN and Subnet view.

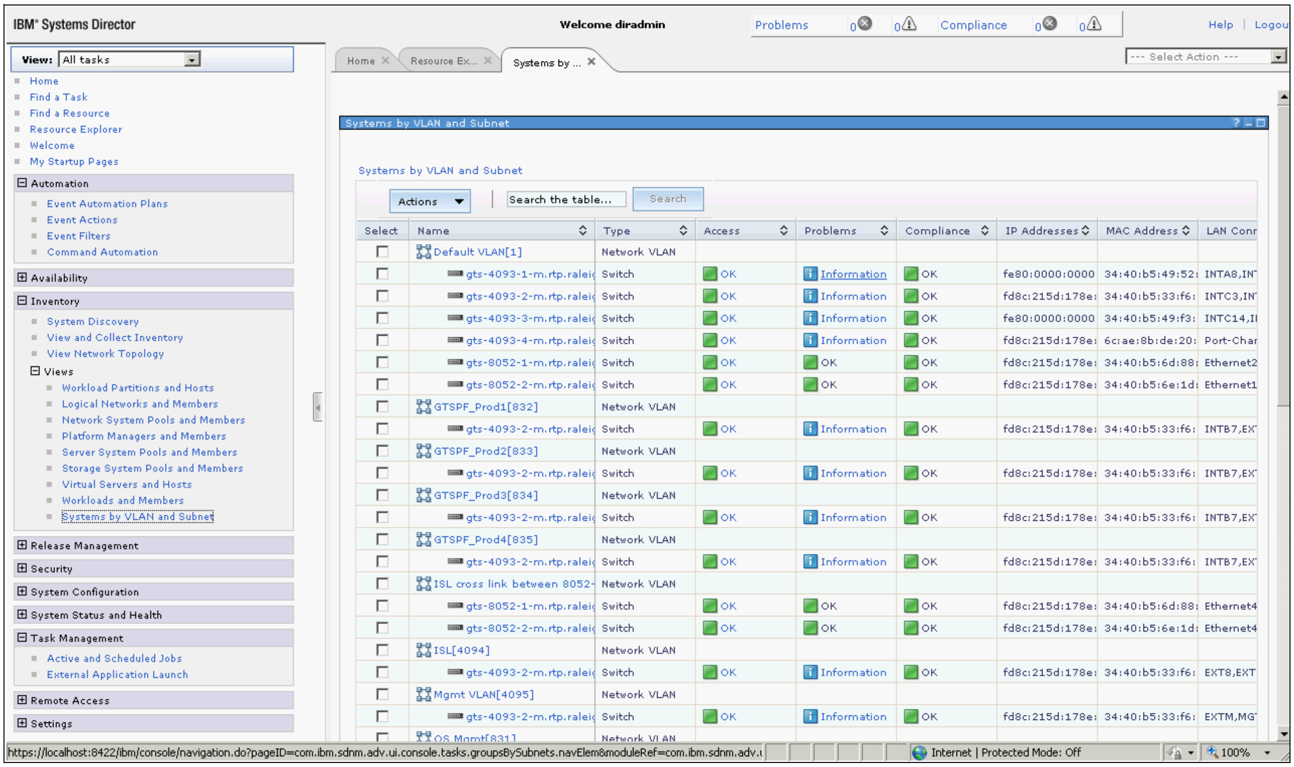
IBM® Systems Director
Welcome diradmin
Problems
Compliance
Help
Logou
View: All tasks
Home
Find a Task
Find a Resource
Resource Explorer
Welcome
My Startup Pages
Automation
Event Automation Plans
Event Actions
Event Filters
Command Automation
Availability
Inventory
System Discovery
View and Collect Inventory
View Network Topology
Views
Workload Partitions and Hosts
Logical Networks and Members
Network System Pools and Members
Platform Managers and Members
Server System Pools and Members
Storage System Pools and Members
Virtual Servers and Hosts
Workloads and Members
Systems by VLAN and Subnet
Release Management
Security
System Configuration
System Status and Health
Task Management
Active and Scheduled Jobs
External Application Launch
Remote Access
Settings
Home
Resource Ex...
Systems by ...
--- Select Action ---
Systems by VLAN and Subnet
Systems by VLAN and Subnet
Actions
Search the table...
Search
https://localhost:8422/ibm/console/navigation.do?pageID=com.ibm.sdnm.adv.ui.console.tasks.groupsBySubnets.navElem&moduleRef=com.ibm.sdnm.adv.u
Internet | Protected Mode: Off
100%

Select	Name	Type	Access	Problems	Compliance	IP Addresses	MAC Address	LAN Conn
□	Default VLAN[1]	Network VLAN						
□	gts-4093-1-m.rtp.raleig	Switch	🟩 OK	Information	🟩 OK	fe80:0000:0000	34:40:b5:49:52:	INTA8,INT
□	gts-4093-2-m.rtp.raleig	Switch	🟩 OK	Information	🟩 OK	fd8c:215d:178e:	34:40:b5:33:f6:	INTC3,INT
□	gts-4093-3-m.rtp.raleig	Switch	🟩 OK	Information	🟩 OK	fe80:0000:0000	34:40:b5:49:f3:	INTC14,II
□	gts-4093-4-m.rtp.raleig	Switch	🟩 OK	Information	🟩 OK	fd8c:215d:178e:	6c:ae:8b:de:20:	Port-Char
□	gts-8052-1-m.rtp.raleig	Switch	🟩 OK	🟩 OK	🟩 OK	fd8c:215d:178e:	34:40:b5:6d:88:	Ethernet2
□	gts-8052-2-m.rtp.raleig	Switch	🟩 OK	🟩 OK	🟩 OK	fd8c:215d:178e:	34:40:b5:6e:1d:	Ethernet1
□	GTSPF_Prod1[832]	Network VLAN						
□	gts-4093-2-m.rtp.raleig	Switch	🟩 OK	Information	🟩 OK	fd8c:215d:178e:	34:40:b5:33:f6:	INTB7,EX
□	GTSPF_Prod2[833]	Network VLAN						
□	gts-4093-2-m.rtp.raleig	Switch	🟩 OK	Information	🟩 OK	fd8c:215d:178e:	34:40:b5:33:f6:	INTB7,EX
□	GTSPF_Prod3[834]	Network VLAN						
□	gts-4093-2-m.rtp.raleig	Switch	🟩 OK	Information	🟩 OK	fd8c:215d:178e:	34:40:b5:33:f6:	INTB7,EX
□	GTSPF_Prod4[835]	Network VLAN						
□	gts-4093-2-m.rtp.raleig	Switch	🟩 OK	Information	🟩 OK	fd8c:215d:178e:	34:40:b5:33:f6:	INTB7,EX
□	ISL cross link between 8052-	Network VLAN						
□	gts-8052-1-m.rtp.raleig	Switch	🟩 OK	🟩 OK	🟩 OK	fd8c:215d:178e:	34:40:b5:6d:88:	Ethernet4
□	gts-8052-2-m.rtp.raleig	Switch	🟩 OK	🟩 OK	🟩 OK	fd8c:215d:178e:	34:40:b5:6e:1d:	Ethernet4
□	ISL[4094]	Network VLAN						
□	gts-4093-2-m.rtp.raleig	Switch	🟩 OK	Information	🟩 OK	fd8c:215d:178e:	34:40:b5:33:f6:	EXT8,EXT
□	Mgmt VLAN[4095]	Network VLAN						
□	gts-4093-2-m.rtp.raleig	Switch	🟩 OK	Information	🟩 OK	fd8c:215d:178e:	34:40:b5:33:f6:	EXTM,MG
□	OS Mgmt[831]	Network VLAN						

Figure 9-2 Systems by VLAN and Subnet view

- In context access to remote command and in context launch of third-party tools: You can open a command prompt to any managed network device right from the IBM Systems Director console, or you can choose to start a third-party management tool from that same console.
- Automation of physical and virtual network configuration with network system pools: You can create network system pools to have Network Control automatically assign VLANs to physical switch ports and reconfigure virtual switches as virtual workloads are moved from one host to another. This is often used with VMControl to manage PowerVM, Kernel-based Virtual Machine (KVM), and VMware environments.

> **Note:** For more information about how to configure network system pools, see the information center:
>
> http://pic.dhe.ibm.com/infocenter/director/pubs/topic/com.ibm.sdnm.adv.helps.doc/fnc0_t_network_ctrl_managing_nsps_and_lnps.html

For more information about Network Control, go to the IBM Systems Director Information Center:

http://pic.dhe.ibm.com/infocenter/director/pubs/topic/com.ibm.sdnm.adv.helps.doc/fnc0_p_network_ctrl.html

To find out what is new in IBM Systems Director Network Control Version 1.4, read the following information center article:

http://pic.dhe.ibm.com/infocenter/director/pubs/topic/com.ibm.sdnm.adv.helps.doc/fnc0_r_whats_new_14.html

For a comparison of basic IBM Systems Director Network Manager versus Network Control, see the following information center page:

http://pic.dhe.ibm.com/infocenter/director/pubs/topic/com.ibm.sdnm.adv.helps.doc/fnc0_r_sdnm_compare_overview.html

9.2 Planning for Network Control

Planning for Network Control is a critical step in the deployment process. Network Control can discover and manage many different types of network devices, but each has different capabilities.

You should not assume that because you can configure VLANs on a type of switch, you will be able to do the same for all your network devices. Management capabilities are exposed to Network Control by using vendor plug-ins that vary greatly in what functionalities they provide.

For a table of supported capabilities per switch device type, see the following information center page:

http://pic.dhe.ibm.com/infocenter/director/pubs/topic/com.ibm.sdnm.adv.helps.doc/fnc0_r_network_ctrl_device_support.html

The tables that are shown at this information center sometimes aggregate some of the management functionalities into categories. For example, it indicates if configuration management is supported or not on the device, but it does not provide the details of which type of configuration is supported (protocol, VLAN, or CEE configuration).

For each switch device type, IBM Systems Director Network Control has one or more specific configuration templates. It is strongly recommended that you review those templates carefully in order to determine if the product capabilities meet your requirements.

To access configuration templates and view what options are configurable for your device type, follow these simple steps:

1. Access the **System Configuration** → **Configuration Templates** menu and click **Create**, as shown in Figure 9-3.

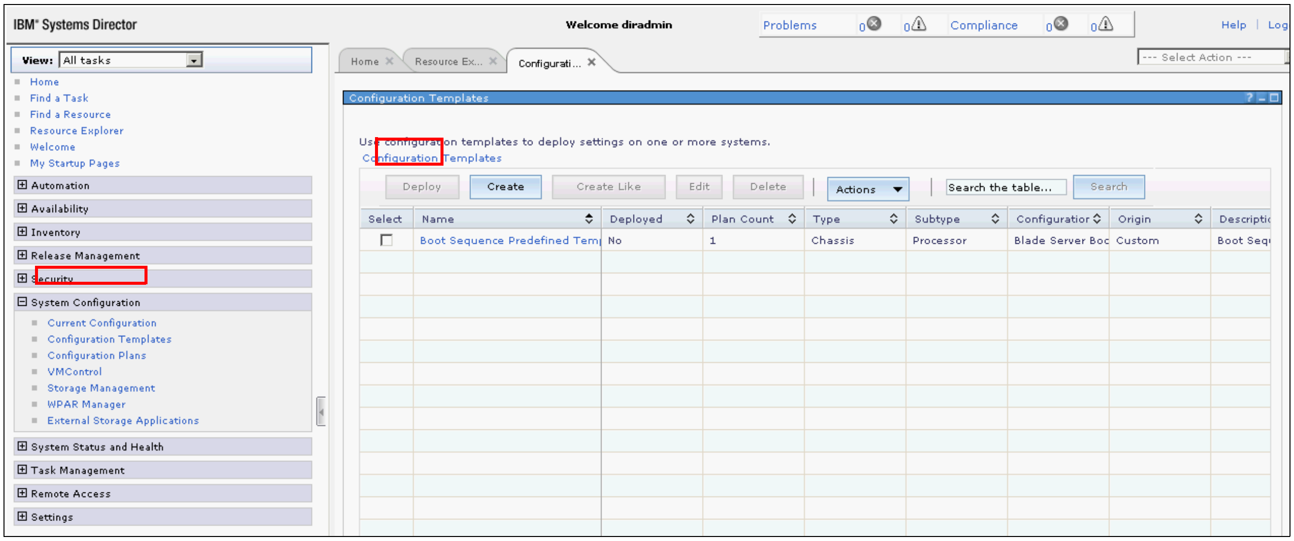

Use configuration templates to deploy settings on one or more systems.
Configuration Templates

Select	Name	Deployed	Plan Count	Type	Subtype	Configuration	Origin	Descriptic
□	Boot Sequence Predefined Temp	No	1	Chassis	Processor	Blade Server Boc	Custom	Boot Seq

Figure 9-3 The configuration template

2. Select **Ethernet Switch** from the Template type drop-down list, and select your device type from the “Configuration to create a template” drop-down list, as shown in Figure 9-4.

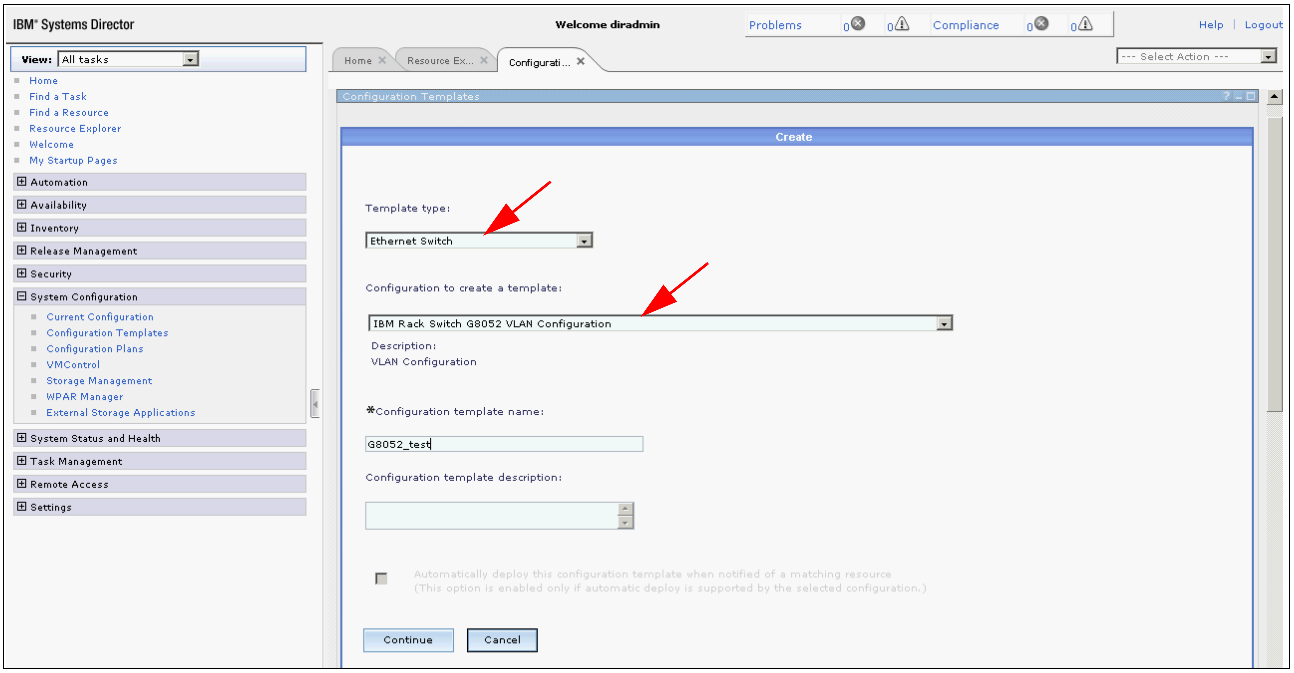

Template type:
Ethernet Switch

Configuration to create a template:
IBM Rack Switch G8052 VLAN Configuration

Description:
VLAN Configuration

*Configuration template name:
G8052_test

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.)

Figure 9-4 Creating a configuration template

3. Finally, navigate through the available options for this specific type of configuration template, as shown in Figure 9-5, and ensure that it has the settings that you require.

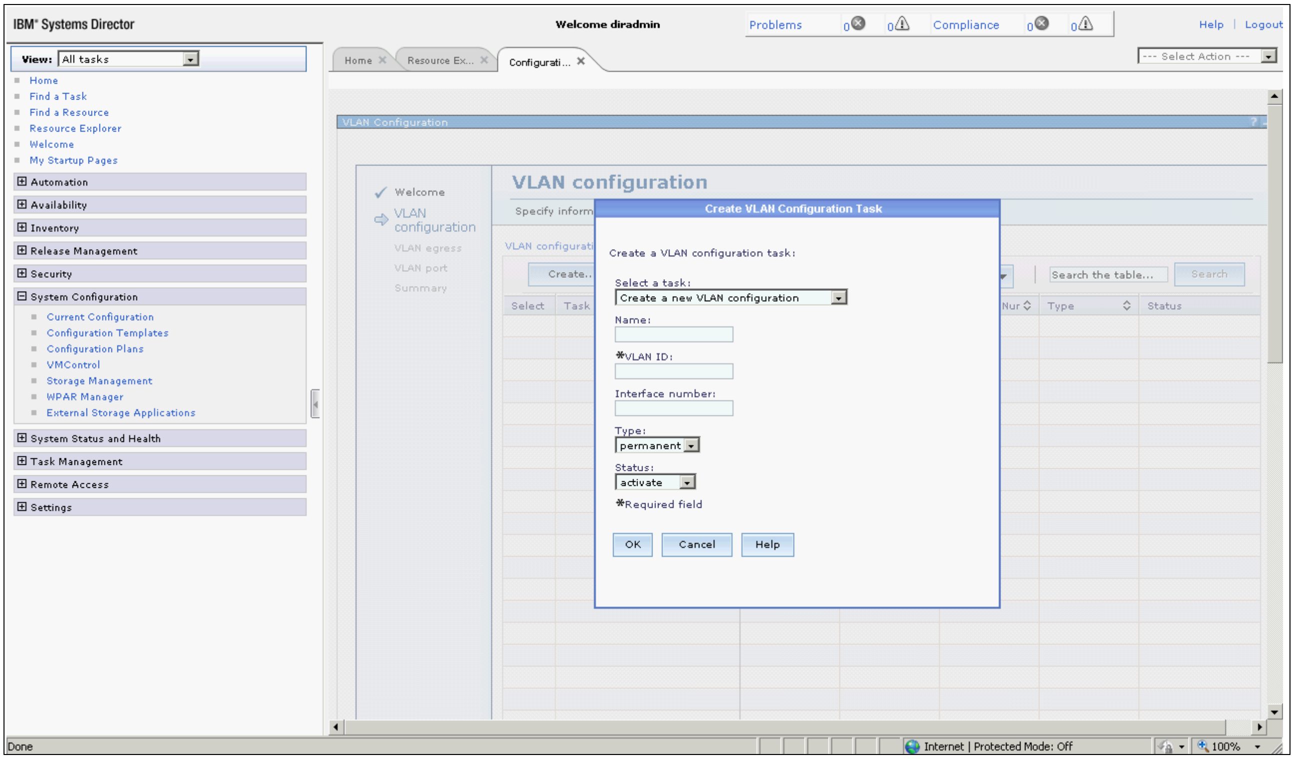

Create VLAN Configuration Task

Create a VLAN configuration task:

Select a task: Create a new VLAN configuration

Name:

*VLAN ID:

Interface number:

Type: permanent

Status: activate

*Required field

Figure 9-5 Examining the available options in a configuration template

When you have determined that the configuration settings that you need are supported by the available templates, perform testing on your specific device types in order to confirm that all requirements are met.

The information center does not clearly document all the requirements for each device type, which is why it is important to test your network devices before implementing your network using IBM Systems Director Network Manager and Network Control. For example, you might discover that certain protocols such as Telnet and HTTP are required for device configuration but are prohibited by your company security policy.

9.3 Discovering and managing virtual networks

Network Control can be used to manage networking in different virtual environments, such as PowerVM, KVM, and VMware vSphere. For more information about how to manage PowerVM and KVM environments using VMControl, Storage Control, and Network Control, refer to Chapter 5, “VMControl” on page 247.

Network Control can also discover virtual switches in VMware vSphere environments. VMware distributed virtual switches are not supported by Network Control. The only type of distributed virtual switch that is supported is the IBM Systems Networking 5000V.

Whereas standard network switches are discovered using their management IP address, VMware vSphere virtual switches do not have management IP addresses and must be discovered by running an inventory on the vCenter Server.

To discover standard VMware vSphere virtual switches, perform the following steps:

1. Discover the vCenter Server operating system by using System Discovery and unlock access to the discovered managed endpoint using credentials that have both local administrative privileges on the operating system and privileged access to the vCenter Server application instance.
2. Perform a full inventory of the vCenter Server. This automatically discovers all associated ESXi physical hosts that are managed by the vCenter Server as well as virtual servers managed by those hosts.
3. Perform a full inventory of those ESXi hosts. This discovers virtual switches.
4. You can then use the **Inventory** → **View Network Topology** view to see how virtual machines are connected to the virtual network. For example, in Figure 9-6, you can see how the virtual machine is connected to the host via the virtual switch, and what the network interface MAC addresses are.

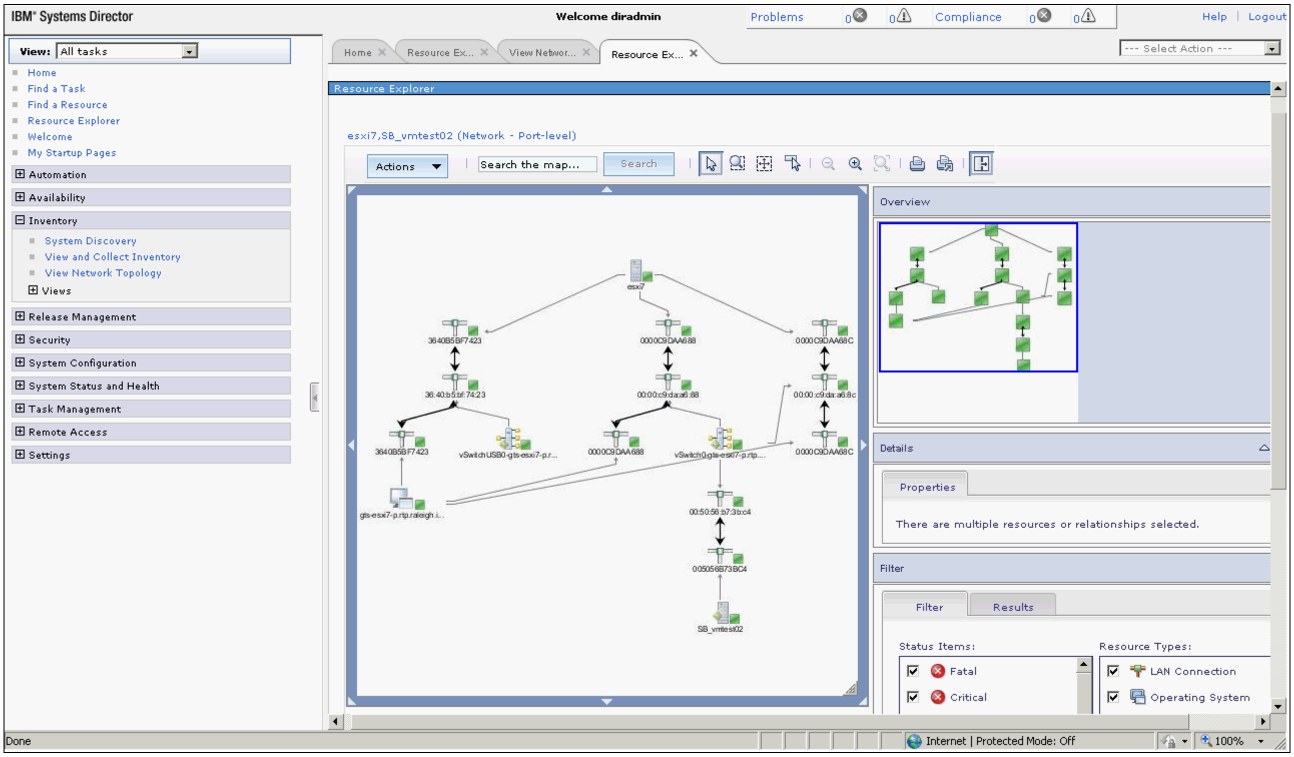

Figure 9-6 Example of a virtual network topology

After you successfully discover your virtual switches, you will also be able to deploy new port groups by editing the current configuration of the virtual switch managed endpoint. To do so, select the vSwitch in the **Resource Explorer** → **Groups** → **All Network Systems** view, then select **Actions** → **System Configuration** → **Current Configuration**.

You can then click **Virtual Switch Module VLAN Configuration** and create a new VLAN.

Alternatively, if you want to create or remove multiple port groups on multiple vSwitches, you can use configuration templates:

1. From the IBM Systems Director console left pane, select **System Configuration** → **Configuration Templates**.
2. Click **Create** to create a new configuration template.

3. Select **Virtual Switch** from the “Template type” drop-down list and select **Virtual Switch Module VLAN Configuration** from the “Configuration to create a template” drop-down list.
4. Enter a name for the configuration template, such as “`My VLAN configuration`”.
5. Click **Create**.
6. To add a port group, leave the default **Create a new VLAN configuration** from the “Select a task” drop-down list and enter a VLAN ID. To remove a port group, select **Delete an existing VLAN configuration** from the “Select a task” drop-down list and enter the VLAN ID of the port group that you want to remove.
7. Add or remove as many port groups as you want and save your configuration when you are done. You can then click **Deploy** and select vSwitches as targets.

> **Note:** When you create port groups using Network Control, they are automatically named `DirectorVLAN`*ID_vSwitch name*. You cannot change that default naming convention.
>
> In addition, Network Control does not create the port group if another port group with the same VLAN ID already exists, even if the names are different.
>
> Finally, when you remove port groups, Network Control selects the port group that is based on the VLAN ID, regardless of the name that it has.

This can be useful if you have many ESXi hosts and VLANs and you are using standard virtual switches because the vSphere client graphical user interface does not let you deploy port groups to multiple virtual switches in a single step. If your virtual infrastructure administrators do not know Powershell, they can use Network Control to automate that task.

9.4 Discovering and managing IBM Systems Networking switches

IBM Systems Networking switches are best discovered using a custom discovery profile and using at least SNMP v2. When discovering this type of switch, consider the following factors:

1. The Telnet protocol might be required for configuration of your switch, including configuration of the initial SNMP trap server address that occurs automatically after the switch MEP has been fully unlocked. If you get a warning after discovery that the SNMP trap address could not be configured, rediscover the switch after you enable Telnet.
2. Ensure that the switch is in ISCLI or prompt mode. You can use the **`boot cli-mode iscli`** or **`boot cli-mode prompt`** ISCLI commands from a Secure Shell (SSH) prompt on your switch.
3. Ensure that the switch can resolve the fully qualified domain name of the IBM Systems Director server using DNS. This is because IBM Systems Director uses its fully qualified domain name (FQDN) as the SNMP trap server address. You can use the **`show ip dns`** ISCLI command to view the current DNS configuration on your switch, and the ISCLI **`ip dns primary-address <ip address>`** and **`ip dns secondary-address <ip address>`** to configure a primary and a secondary DNS server for name resolution.
4. Ensure that you have correctly configured the SNMP trap source on the switch or the IBM Systems Director server might not receive events. Use the **`snmp-server trap-source 128`** command to configure SNMP trap events to be sent over the management interface of the switch.

5. If you are not receiving events in a format that is appropriate, load all the Management Information Base (MIB) files available from the IBM Systems Networking image update package compressed file, which you can get from IBM FixCentral. Go to **Settings** → **Manage MIBs** and manually import all RFC MIB files included in that package in alphabetical order before you are able to successfully import the switch main MIB file. In addition, you might have to edit the map file, which defines severity and event text based on the SNMP trap OID. For more information about this topic, see the following information center topic:

   `http://pic.dhe.ibm.com/infocenter/director/pubs/index.jsp?topic=%2Fcom.ibm.director.automation.helps.doc%2Ffqm0_c_ea_mapping_snmp_traps_events.html`

   > **Note:** At the time of writing, the G8052 and G8264 MIB files could not be compiled and loaded successfully in IBM Systems Director 6.3.2.1. Only the EN4093 MIB file could be tested successfully.
   >
   > The G8052 and G8264 files contain a syntax error in the `dhcpSnoopingBindingInfoExpiry` section. To correct that error and successfully load the MIB files in IBM Systems Director, open the MIB files in a text editor, search for `dhcpSnoopingBindingInfoExpiry`, and replace the line, `SYNTAX Counter32` with `SYNTAX Unsigned32`. Save the file and try to load it again in IBM Systems Director.

6. Perform discovery on the SNMP write community, not the SNMP read community.
7. Always completely unlock the switch. If it remains in partial access, the SNMP trap server address will not be configured and your IBM Systems Director server will not receive SNMP events.

   > **Note:** A simple way to test if you are correctly receiving SNMP traps from your switches is to simply log in and out of your switch using either the web user interface (also known as the *BBI*), Telnet, or SSH. Those events should generate traps of *Unknown* or *Information* severity on your IBM Systems Director server.

When the switches have been discovered and inventoried, perform the following tasks:

1. Configure additional SNMP trap server addresses using the protocol configuration template: Go to **Resource Explorer** → **Groups** → **All Network Systems**, select your switch and then select **Actions** → **System Configuration** → **Current Configuration** and finally, select the **Protocol Configuration** template.
2. Create VLANs and assign VLANs to switch ports using the VLAN configuration template: Go to **Resource Explorer** → **Groups** → **All Network Systems**, select your switch and then select **Actions** → **System Configuration** → **Current Configuration** and finally, select the **VLAN Configuration** template.
3. Create and deploy CEE networking and EVB configurations: Go to **System Configuration** → **Configuration Templates** and create a new template of type, **Ethernet Switch**. Then, select **CEE Configuration for `<your switch type>`** or **EVB Configuration for `<your switch type>`**.
4. Start a remote command line on the switch: Go to **Resource Explorer** → **Groups** → **All Network Systems**, select your switch, and then select **Actions** → **Remote Access** → **Remote Command Line**.
5. View a network topology for your switch: Go to **Resource Explorer** → **Groups** → **All Network Systems**, select your switch, and then select **Actions** → **Topology Perspectives** → **Network**. Then, select either the **Basic**, **Port-level**, **Subnet**, or **System-level** view.

Note: For more advanced management about capabilities of IBM Systems Networking switches, consider using the IBM Systems Networking Switch Center product.

Switch Center replaces what was formerly known as IBM Systems Networking Element Manager (SNEM).

For more information about the IBM System Networking Switch Center, go to the following website:

http://www-03.ibm.com/systems/networking/software/snsc/index.html

You can also find the IBM System Networking Switch Center 7.1.1 User Guide here:

http://www-01.ibm.com/support/docview.wss?uid=isg3T7000624

9.5 Discovering and managing the IBM Systems Networking 5000V

Unlike traditional virtual switches, you must discover the IBM System Networking Distributed Virtual Switch 5000V directly by using its management IP address and a custom discovery profile, as documented in the information center:

http://pic.dhe.ibm.com/infocenter/director/pubs/topic/com.ibm.sdnm.adv.helps.doc/fqm0_t_network_ctrl_distributed_virtual_switches_discover.html

If you have not yet discovered the VMware vCenter server that is connected to the 5000V, the operating system MEP for the vCenter Server will be automatically added to the IBM Systems Director inventory when inventory is collected on the 5000V MEP. Unlock the vCenter Server operating system MEP and collect inventory from it to enable all management functionalities of the 5000V from Network Control.

Note: Enabling the Telnet protocol is not required for the 5000V.

Basic management capability is limited to receiving SNMP events, viewing network topology views (as shown in Figure 9-7 on page 429), and accessing the remote command line on the 5000V.

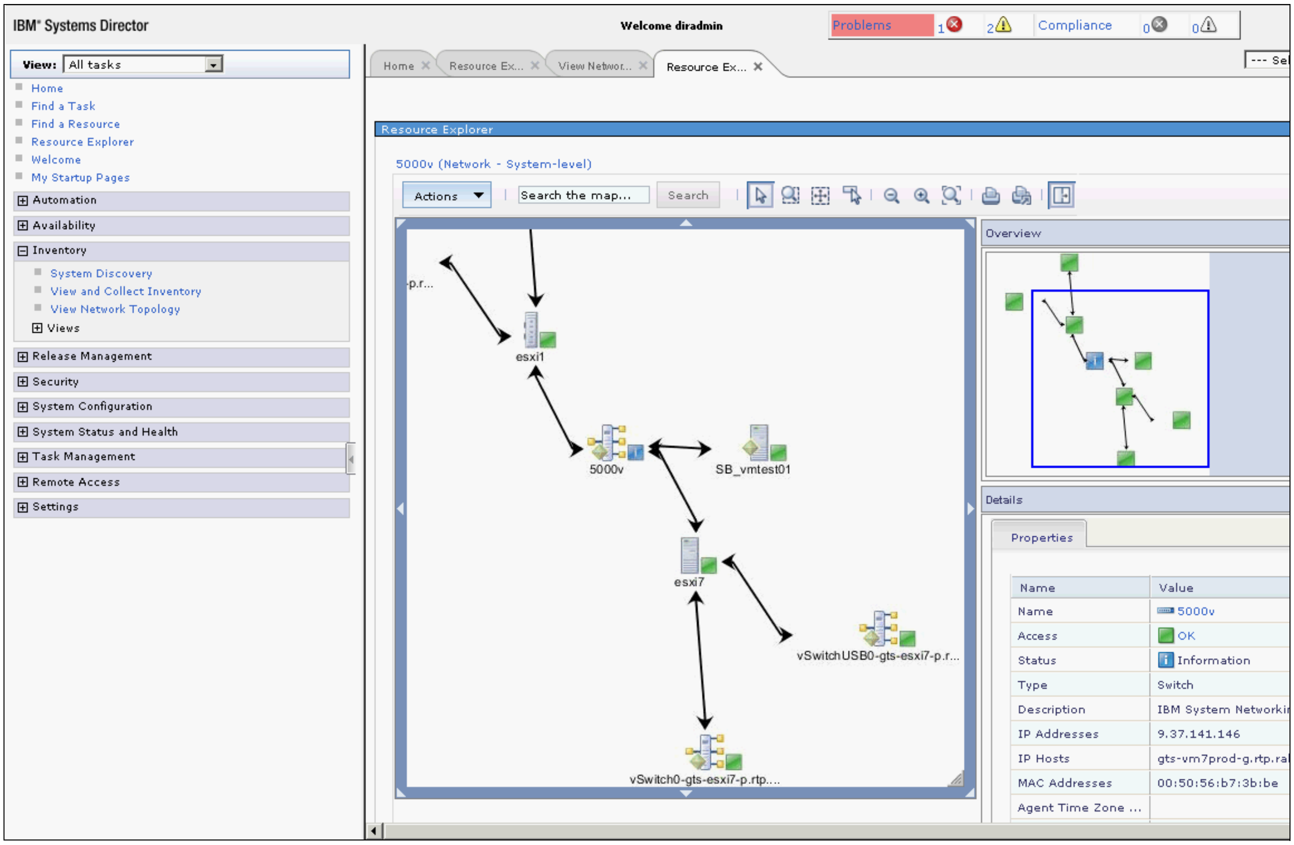

Name	Value
Name	5000v
Access	OK
Status	Information
Type	Switch
Description	IBM System Networki
IP Addresses	9.37.141.146
IP Hosts	gts-vm7prod-g.rtp.ra
MAC Addresses	00:50:56:b7:3b:be
Agent Time Zone ...	

Figure 9-7 Network topology view for the 5000V

The VLAN and EVB configuration of the 5000V by using Network Control is not supported.

In order to receive events from your 5000V, perform the following tasks:

1. Manually configure the SNMP system name on your 5000V by using the **`snmp-server name`** command in the configuration menu.
2. Manually configure the SNMP manager target address (your IBM Systems Director server address) by using the **snmp-server target-address `<a number>` address `<ip address of your IBM Systems Director server>` name `<a name for this entry>`** command.
3. Manually upload the `iswitch.mib` file that is provided with the 5000V into your IBM Systems Director server by going to **Settings** → **Manage MIBs**.

Example 9-1 shows the list of SNMP events that are defined in the `iswitch.mib` MIB file.

Example 9-1 List of SNMP events that are defined in iswitch.mib

```
VM disconnect
VM connect
create vnic profile
delete vnic profile
add ports to vnic profile
delete ports from vnic profile
change pvid settings of vnic profile
change tagging settings of vnic profile
change tagpvid settings of vnic profile
change designated-uplinks settings of vnic profile
change vepa settings of vnic profile
change dps settings of vnic profile
add vlans to vnic profile
remove vlans from vnic profile
change vsi settings of vnic profile
change dot1p settings of vnic profille
add RX ACL for vnic profile
remove RX ACL from vnic profile
change RX/TX policy map settings of vnic profile
change RX/TX rate limit settings of vnic profile
```

> **Note:** At the time of writing, the `iswitch.mib` file did not enable IBM Systems Director version 6.3.2.1 to interpret all SNMP events from the 5000V correctly. As a result, some events were coming in to the server with a severity of type, *unknown*.
>
> This is due to missing entries in the `BNTTraps.map`. You can learn how to customize the map file in the information center:
>
> http://pic.dhe.ibm.com/infocenter/director/pubs/index.jsp?topic=%2Fcom.ibm.director.automation.helps.doc%2Ffqm0_c_ea_mapping_snmp_traps_events.html

If you want to configure 802.1Qbg (also known as EVB) and you are planning on using your IBM Systems Director server as the VSI manager (aka the VSI database), configure the VSI manager IP address on your 5000V manually by using the **iswitch myvsidb `<ip address of your IBM Systems Director server>`** command.

For more information about how to use your IBM Systems Director server as a VSI manager, read the following information center document:

http://pic.dhe.ibm.com/infocenter/director/pubs/topic/com.ibm.sdnm.adv.helps.doc/fnc0_r_learnmore_VSI_config.html

For more information about how to configure 802.1Qbg, see the following IBM Redbooks publication:

http://www.redbooks.ibm.com/abstracts/sg247985.html?Open

> **Note:** IBM Systems Director 6.3.3 with Network Control 1.4.1 also enables EVB support for Flex Power compute nodes that are managed by an FSMI.
>
> With IBM Systems Director 6.3.2, you can use the IBM Systems Director server as a VSI manager to enable EVB support for KVM environments. You can then create Logical Network Profiles (LNPs), which use the VSI manager ID, VSI type ID, and VSI type version that is specified in your VSI database and assign those LNPs to your network system pools.

9.6 Best practices

When implementing your network using IBM Systems Director Network Manager and Network Control, keep the following best practices in mind:

- *Carefully plan your Network Control deployment* by ensuring that your network devices are supported and that available configuration templates include the options that you require.
- *Always test Network Control discovery and management* on your specific network device type before deployment in order to ensure that device requirements can be met.
- *Use custom discovery profiles* so that you can use non-default SNMP community names.
- Fully *unlock access to your network device managed endpoints,* or the SNMP trap server address will not be automatically configured and your devices will not be correctly monitored.
- Ensure that your *network devices can resolve the IBM Systems Director server fully qualified domain name*.
- *Import MIB files* for all your network devices when the default system behavior does not meet your monitoring requirements. In addition, *edit the map file* to handle event text and severity appropriately.
- Always *inventory your network devices* after discovering them or topology views might be inaccurate.
- For configuration management, your *IBM Systems Networking switches will need to be configured in ISCLI mode*.
- *Use configuration templates* to deploy or remove multiple port groups to multiple VMware vSphere virtual switches.
- *Consider using EVB* in the Logical Network Profiles of your network system pools in VMControl because this is the emerging standard for automated network configuration.

10

Service and Support Manager

This chapter provides information about the IBM Systems Director Service and Support Manager (SSM). It describes how the Service and Support Manager must be set up and how it handles serviceable events.

The following topics are covered:

10.1 Service and Support Manager

IBM Systems Director Service and Support Manager (SSM) manages serviceable problems and reports the events to IBM. A *serviceable problem* is a problem to which IBM service typically responds, such as a failure of a hardware component that is under warranty.

IBM Systems Director Service and Support Manager is documented in the information center:

http://publib.boulder.ibm.com/infocenter/director/pubs/topic/com.ibm.esa.director.help/esa_kickoff.html

Service and Support Manager is viewed as an advanced manager within Systems Director, although it is installed with the base Systems Director server installation. The core objective of Service and Support Manager is to work with service information:

- Supported systems monitoring
- Serviceable event processing
- Support file management
- CLI support
- Collection of performance management data to send to IBM

Service and Support Manager support different system types in the environment. See Figure 10-1 on page 435. Following are the different system types:

- In-band communication:
  - Windows on x86 systems
  - Linux on x86 system
  - BladeCenter servers (no JS and QS blades)
  - AIX and Linux on IBM Power Systems with Common Agent installed (not managed by Integrated Virtualization Manager (IVM) or Hardware Management Console (HMC))
  - IBM i partition on Power Systems managed by HMC
  - IBM Flex System V7000 Storage Node
- Out-of-band communication:
  - Advanced management module (AMM)
  - Integrated management module (IMM and IMM v2)
  - Chassis management module (CMM)
  - Remote Supervisor Adapter (RSA and RSA II) and baseboard management controller (BMC)

Figure 10-1 shows systems that are supported by SSM.

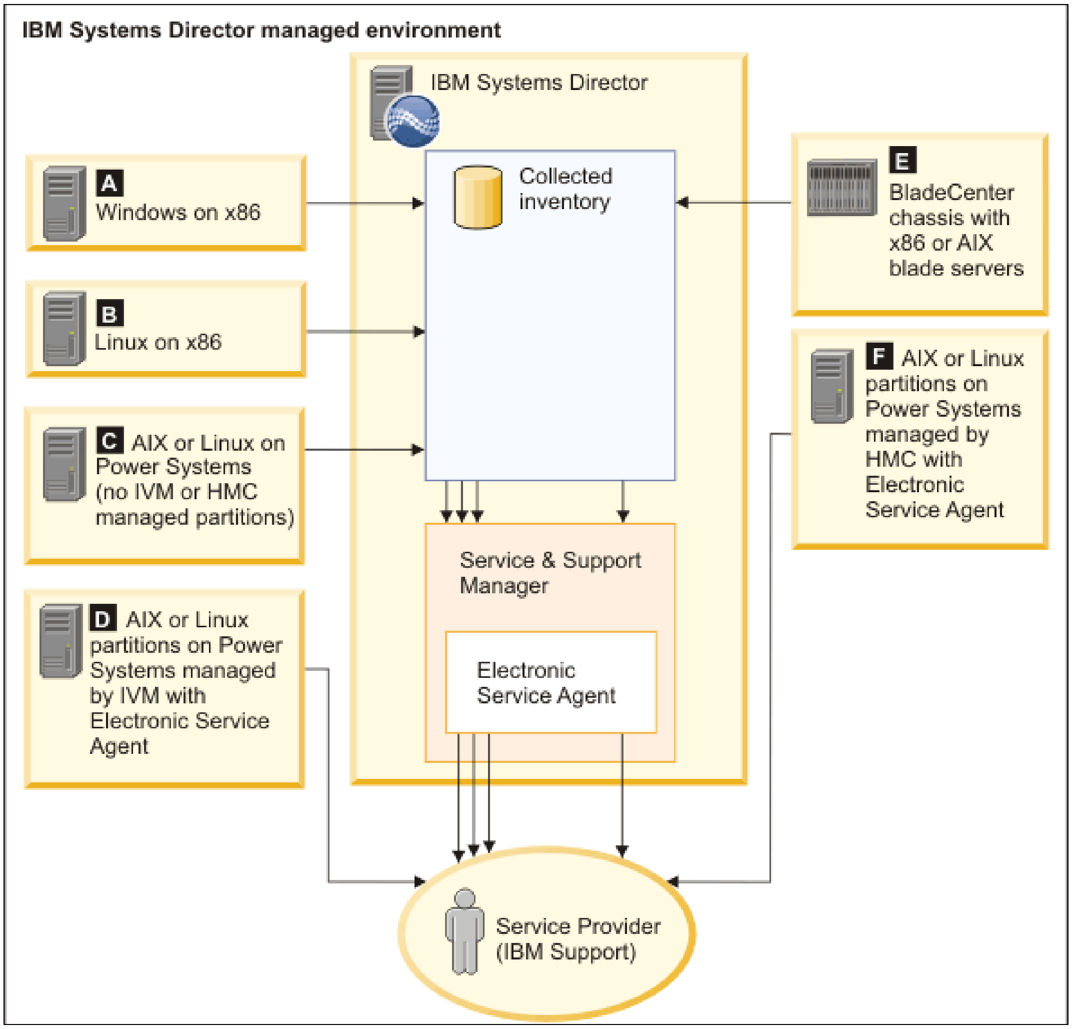
IBM Systems Director managed environment
IBM Systems Director
Collected inventory
A Windows on x86
B Linux on x86
C AIX or Linux on Power Systems (no IVM or HMC managed partitions)
D AIX or Linux partitions on Power Systems managed by IVM with Electronic Service Agent
E BladeCenter chassis with x86 or AIX blade servers
F AIX or Linux partitions on Power Systems managed by HMC with Electronic Service Agent
Service & Support Manager
Electronic Service Agent
Service Provider (IBM Support)

Figure 10-1 Systems supported by Service and Support Manager

A complete list of systems and resources that are eligible for monitoring by Service and Support Manager are listed in the information center:

http://pic.dhe.ibm.com/infocenter/director/pubs/index.jsp?topic=%2Fcom.ibm.esa.director.help%2Fesa_eligibility.html

To enable a system for the Service and Support Manager, complete the following steps:

- System must be discovered and unlocked
- A full inventory of the system must be collected

After finishing the inventory collection and when the Service and Support Manager is configured, the monitoring of the systems begins automatically. You can check if a system is eligible by right-clicking the system and selecting **Service and Support** → **Check Eligibility**. When eligible for the Service and Support Manager, it is automatically added to the **Group** → **Service and Support Groups** → **Monitored Systems** group. Systems that are not eligible for the Service and Support Manager can be found in the **Excluded Systems** group. Systems where the IBM Systems Director has no access or the access state in unknown can be found in the **Unknown Systems** group.

Service and Support Manager subscribes to Systems Director events and filters out unserviceable events. When a serviceable event is received by Service and Support Manager, it submits a service request for the applicable event to IBM. Service and Support Manager runs data collectors on managed endpoints by using the **snap** command for AIX and Linux on Power Systems. Service and Support Manager runs data collectors on managed endpoints by using Dynamic System Analysis (DSA) for Linux and Windows on IBM x86 systems.

The following data will be collected by the different systems and tools:

- System x
  - Dynamic Systems Analysis (DSA)

    SSM uses DSA to collect data from System x endpoints. SSM remotely uploads the DSA tool to the endpoint system. DSA creates a compressed XML report, which will be sent back to the SSM. This file will be submitted to IBM for the serviceable problem.

    DSA collects the following data:

    - System configuration
    - Installed packages
    - Kernel modules
    - Network interface and settings
    - Performance data and details for running processes
    - Hardware inventory information
    - IBM Lightpad status
    - Service processor status and configuration
    - Vital product data (VPD) data, firmware, and basic input/output system (BIOS) information
    - ServeRAID and LSI Redundant Array of Independent Disks (RAID) configuration
    - Event logs from operating system, RAID controller, and service processors
  - sosreport

    System x systems that are running RedHat Enterprise Linux with sosreport version 1.7 or newer are eligible for sosdata collection. SSM runs the **/usr/sbin/sosreport** command and creates a compressed tar file. The IBM Electronic Service Agent™ (ESA) transmits this file and sosreport to the SSM.
  - IMM service log

    The IMM service log contains service information that is collected by the IMM service processor. These files are collected from x86 ITE within a Flex Enterprise Chassis and from System x systems that have an IMM v2 service processor.
- IBM BladeCenter chassis
  - Service log

    Collects data from the BladeCenter chassis by collection service information from the AMM and stores this data in a support file
- AIX and Linux on Power Systems
  - snap

    SSM uses the **snap** command to collect system data on AIX and Linux on Power Systems. The **snap** command gathers system configuration information and compresses this information into a pax file. This file is sent to IBM support.

Detailed information about the **snap** command can be found at the following link:

http://publib.boulder.ibm.com/infocenter/pseries/v5r3/index.jsp?topic=/com.ibm.aix.cmds/doc/aixcmds5/snap.htm

  - sos report

    Linux on Power Systems that are running RedHat Enterprise Linux with sosreport version 1.7 or newer are eligible for sosdata collection. SSM runs the **/usr/sbin/sosreport** command and creates a compressed tar file. The ESA transmits this file as a report file to the SSM.

- IBM i partition on a Power System managed by HMC
  - APAR Library

    Authorized program analysis report (APAR) data is saved in a unique APAR library. The **aparlib** command collects partition local problem error data in the APAR library.

    When automatically collected, the default options collect primary APAR save files and manifest files from the IBM i partition. Additional collection is available by manual collection.

- IBM Flex Enterprise Chassis

  SSM collects several different types of data from monitored x86 compute nodes, IBM Flex System V7000 Storage Nodes, and from the IBM Flex Enterprise Chassis itself.

  - CMM service log

    The CMM service log contains service information from the CMM

  - IMM service log

    Contains information that is collected from IMM on x86 compute nodes within a Flex Enterprise Chassis

  - SVC snap

    Collects storage data from an eligible IBM Flex System V7000 Storage Node

10.2 Launch Service and Support Manager

To launch Service and Support Manager from the Systems Director home page, click **Plug-ins** and scroll down to Service and Support Manager.

Before starting the Service and Support Manager the first time, it needs to be configured. To configure the Service and Support Manager, click **Getting Started with Electronic Service Agent** (see Figure 10-2).

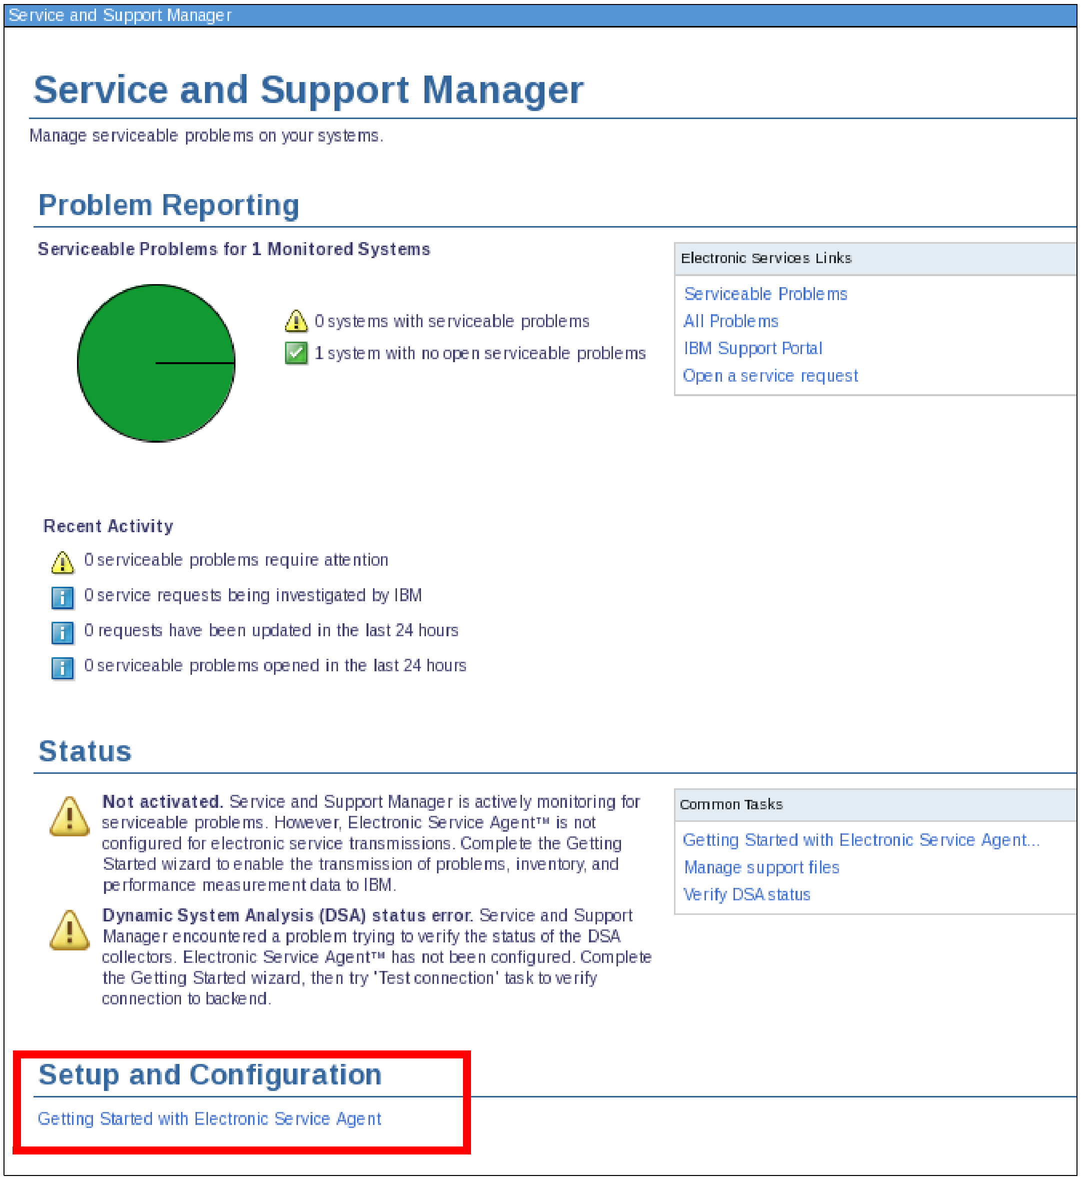

Service and Support Manager

Service and Support Manager

Manage serviceable problems on your systems.

Problem Reporting

Serviceable Problems for 1 Monitored Systems

⚠ 0 systems with serviceable problems

✓ 1 system with no open serviceable problems

Electronic Services Links

- Serviceable Problems
- All Problems
- IBM Support Portal
- Open a service request

Recent Activity

⚠ 0 serviceable problems require attention

ℹ 0 service requests being investigated by IBM

ℹ 0 requests have been updated in the last 24 hours

ℹ 0 serviceable problems opened in the last 24 hours

Status

⚠ **Not activated.** Service and Support Manager is actively monitoring for serviceable problems. However, Electronic Service Agent™ is not configured for electronic service transmissions. Complete the Getting Started wizard to enable the transmission of problems, inventory, and performance measurement data to IBM.

⚠ **Dynamic System Analysis (DSA) status error.** Service and Support Manager encountered a problem trying to verify the status of the DSA collectors. Electronic Service Agent™ has not been configured. Complete the Getting Started wizard, then try 'Test connection' task to verify connection to backend.

Common Tasks

- Getting Started with Electronic Service Agent...
- Manage support files
- Verify DSA status

Setup and Configuration

Getting Started with Electronic Service Agent

Figure 10-2 Service and Support Manager

A wizard opens. Put in the necessary data and information and then test the connectivity.

This necessary data includes the contact data (see Figure 10-3). The minimum required data is marked with an asterisk. After finishing the setup wizard, you can add additional contacts for using IBM Systems Director Service and Support Manager.

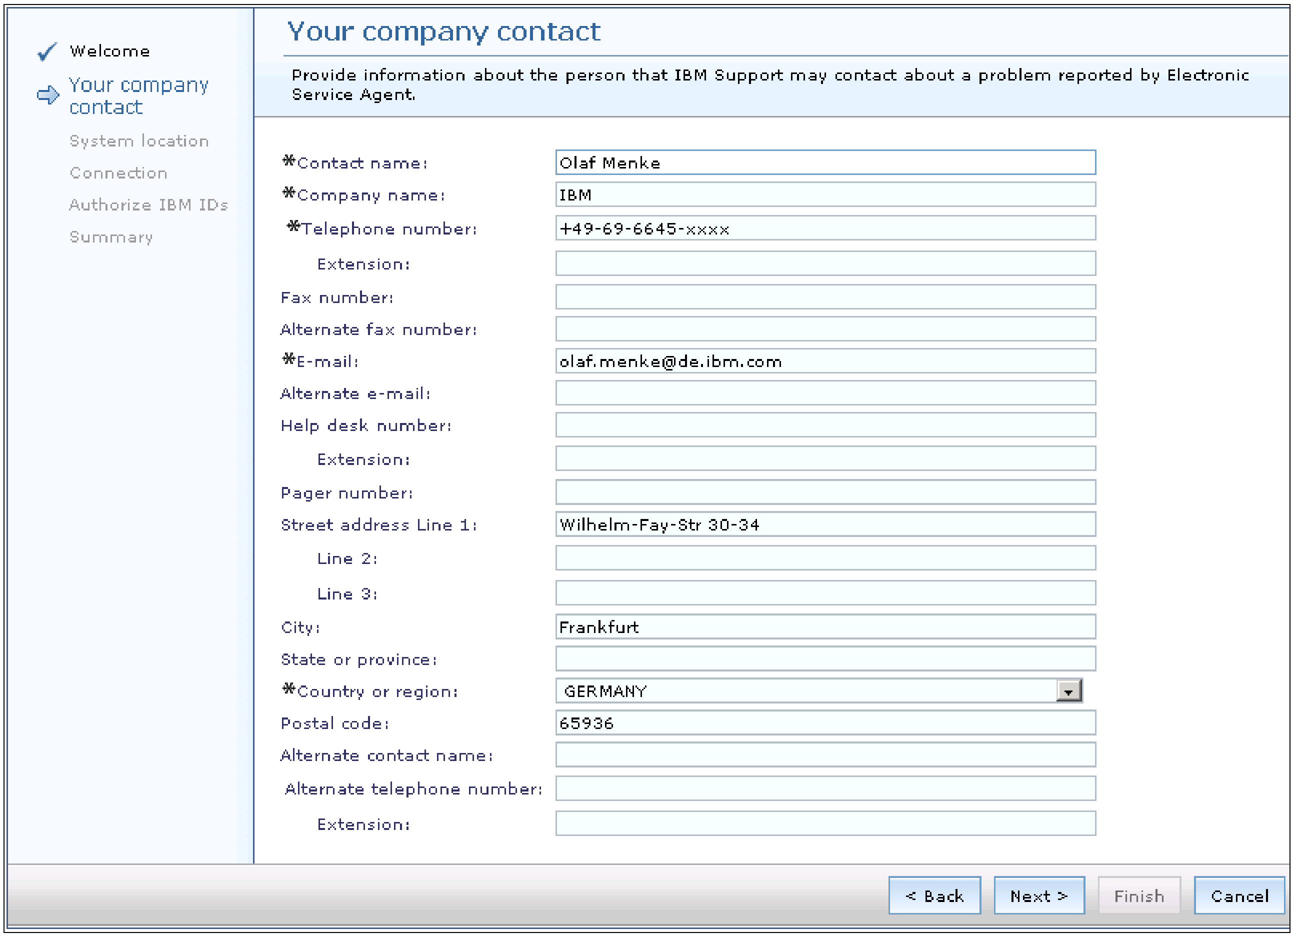

Your company contact

Provide information about the person that IBM Support may contact about a problem reported by Electronic Service Agent.

*Contact name:	Olaf Menke
*Company name:	IBM
*Telephone number:	+49-69-6645-xxxx
Extension:	
Fax number:	
Alternate fax number:	
*E-mail:	olaf.menke@de.ibm.com
Alternate e-mail:	
Help desk number:	
Extension:	
Pager number:	
Street address Line 1:	Wilhelm-Fay-Str 30-34
Line 2:	
Line 3:	
City:	Frankfurt
State or province:	
*Country or region:	GERMANY
Postal code:	65936
Alternate contact name:	
Alternate telephone number:	
Extension:	

Figure 10-3 Configure Service and Support manager: Contact data

On the next page, put in the location for your systems. If there are different locations, put in data for the main location. Also, the minimum required data is marked with an asterisk (see Figure 10-4).

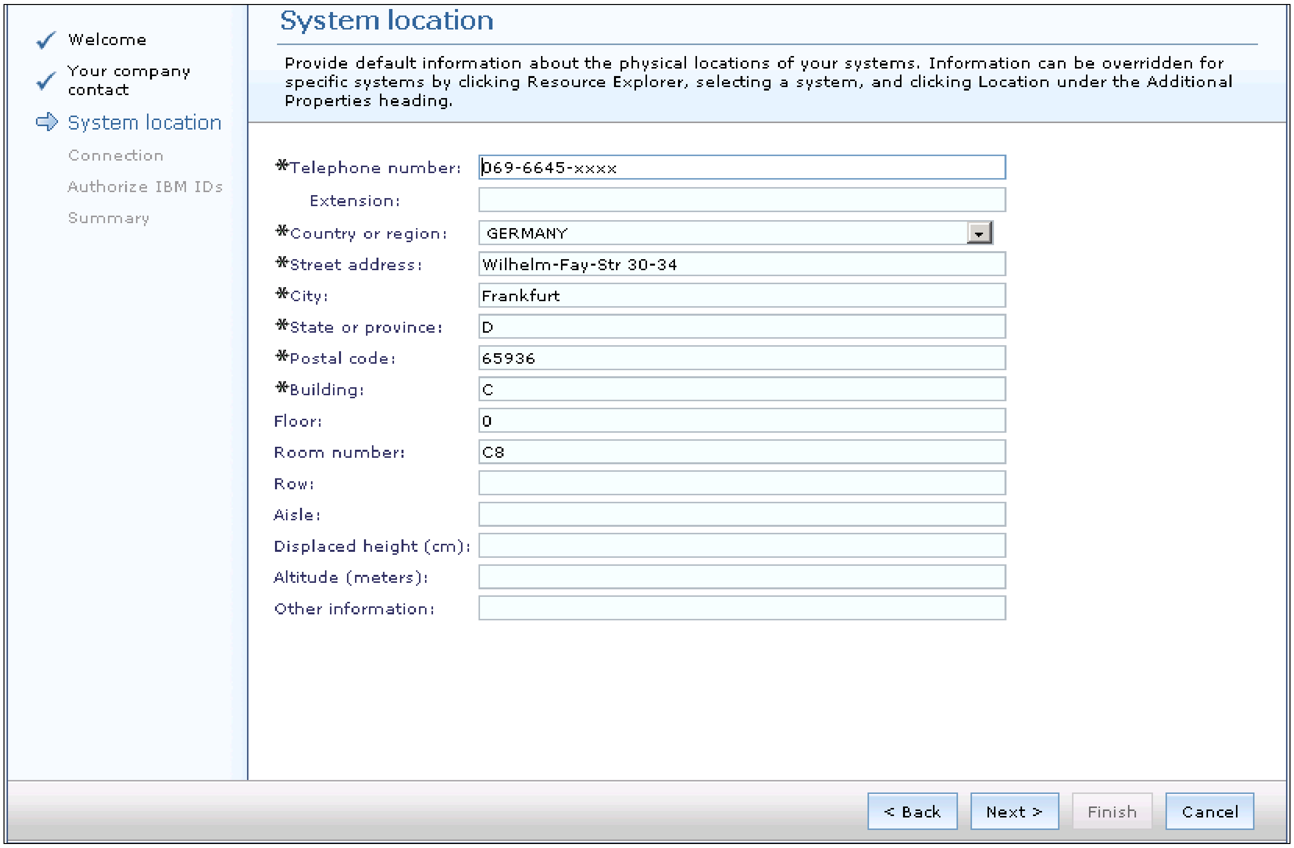

System location

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.

*Telephone number:	069-6645-xxxx
Extension:	
*Country or region:	GERMANY
*Street address:	Wilhelm-Fay-Str 30-34
*City:	Frankfurt
*State or province:	D
*Postal code:	65936
*Building:	C
Floor:	0
Room number:	C8
Row:	
Aisle:	
Displaced height (cm):	
Altitude (meters):	
Other information:	

Figure 10-4 Configure Service and Support manager: System location

On the next page, you can configure the Internet connection. To run and automatically create hardware calls at IBM, an Internet connection is necessary. Figure 10-5 shows the two possible methods:

- Direct Internet connection
- Connection over an HTTP proxy

Select the method that fits best for your environment.

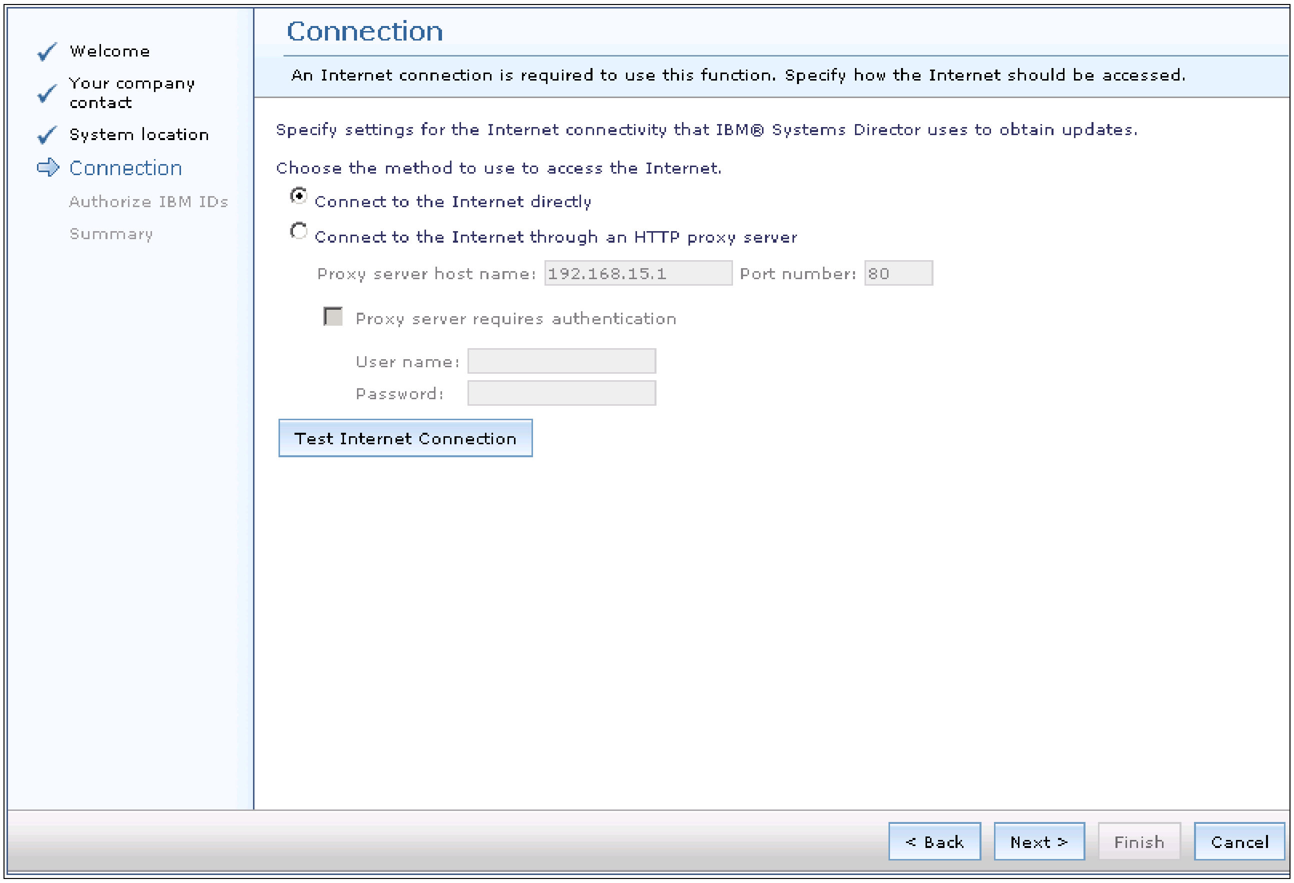

Connection

An Internet connection is required to use this function. Specify how the Internet should be accessed.

Specify settings for the Internet connectivity that IBM® Systems Director uses to obtain updates.

Choose the method to use to access the Internet.

● Connect to the Internet directly

○ Connect to the Internet through an HTTP proxy server

Proxy server host name: 192.168.15.1 Port number: 80

□ Proxy server requires authentication

User name:

Password:

Figure 10-5 Configure Service and Support Manager: Internet connection

On the next page, you can put in the IBM ID. Providing your IBM ID enables you to access the service information that is transmitted to IBM. See Figure 10-6.

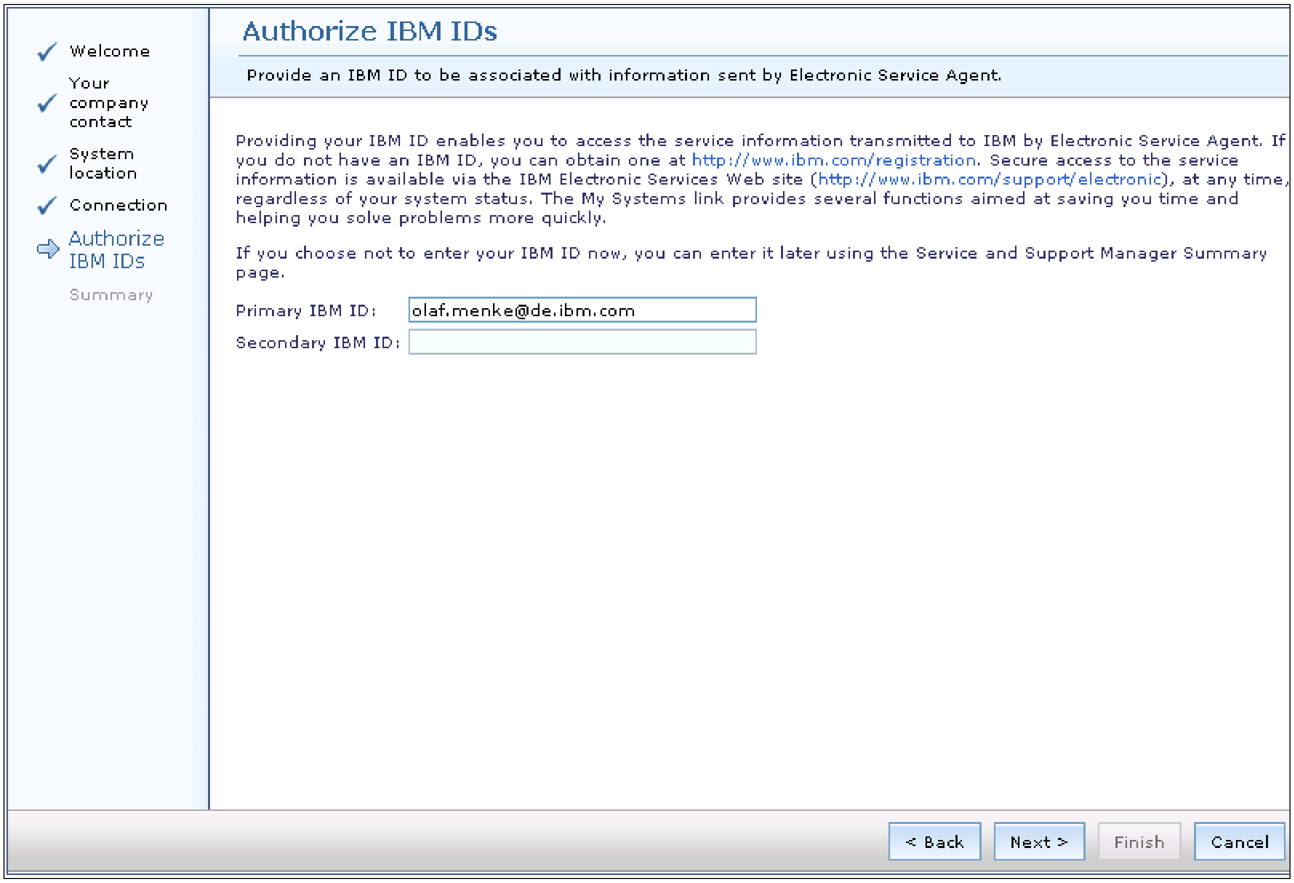

Authorize IBM IDs

Provide an IBM ID to be associated with information sent by Electronic Service Agent.

Providing your IBM ID enables you to access the service information transmitted to IBM by Electronic Service Agent. If you do not have an IBM ID, you can obtain one at http://www.ibm.com/registration. Secure access to the service information is available via the IBM Electronic Services Web site (http://www.ibm.com/support/electronic), at any time, regardless of your system status. The My Systems link provides several functions aimed at saving you time and helping you solve problems more quickly.

If you choose not to enter your IBM ID now, you can enter it later using the Service and Support Manager Summary page.

Primary IBM ID: olaf.menke@de.ibm.com

Secondary IBM ID:

Figure 10-6 Configure Service and Support Manager: IBM ID

At the end, you see a summary page with all the information and data that you submitted. Click **Finish** and the data is saved.

After you enter the required information, the Service and Support Manager is ready, as shown in Figure 10-7.

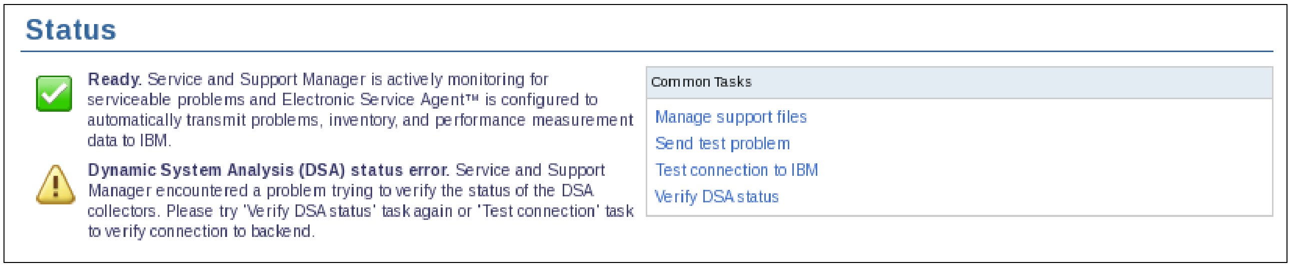

Status

✓ **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.

⚠ **Dynamic System Analysis (DSA) status error**. Service and Support Manager encountered a problem trying to verify the status of the DSA collectors. Please try 'Verify DSA status' task again or 'Test connection' task to verify connection to backend.

Figure 10-7 Service and Support Manager: Status

The Service and Support Manager creates default groups within Systems Director. The groups are under **Resource Explorer → Groups → Service and Support Groups**, as shown in Figure 10-8. The groups are dynamic and automatically populated.

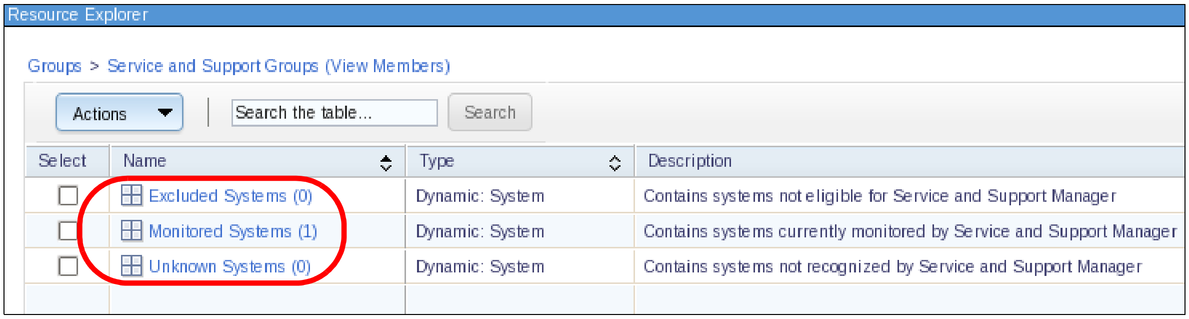

Select	Name	Type	Description
□	Excluded Systems (0)	Dynamic: System	Contains systems not eligible for Service and Support Manager
□	Monitored Systems (1)	Dynamic: System	Contains systems currently monitored by Service and Support Manager
□	Unknown Systems (0)	Dynamic: System	Contains systems not recognized by Service and Support Manager

Figure 10-8 Service and Support Manager groups

When the Service and Support manager is configured, you see additional links at the bottom of the page:

- Manage settings
- Manage your system contacts
- Getting started with Electronic Service Agent (which reopens the configuration wizard)

When you select **Manage your system contacts**, a new window opens where you can see the actual defined contacts. You can also put in new contacts and assign systems to each of the contacts. Select which contact will be the default (this is the person that will be contacted by IBM Service in case of an event. See Figure 10-9.

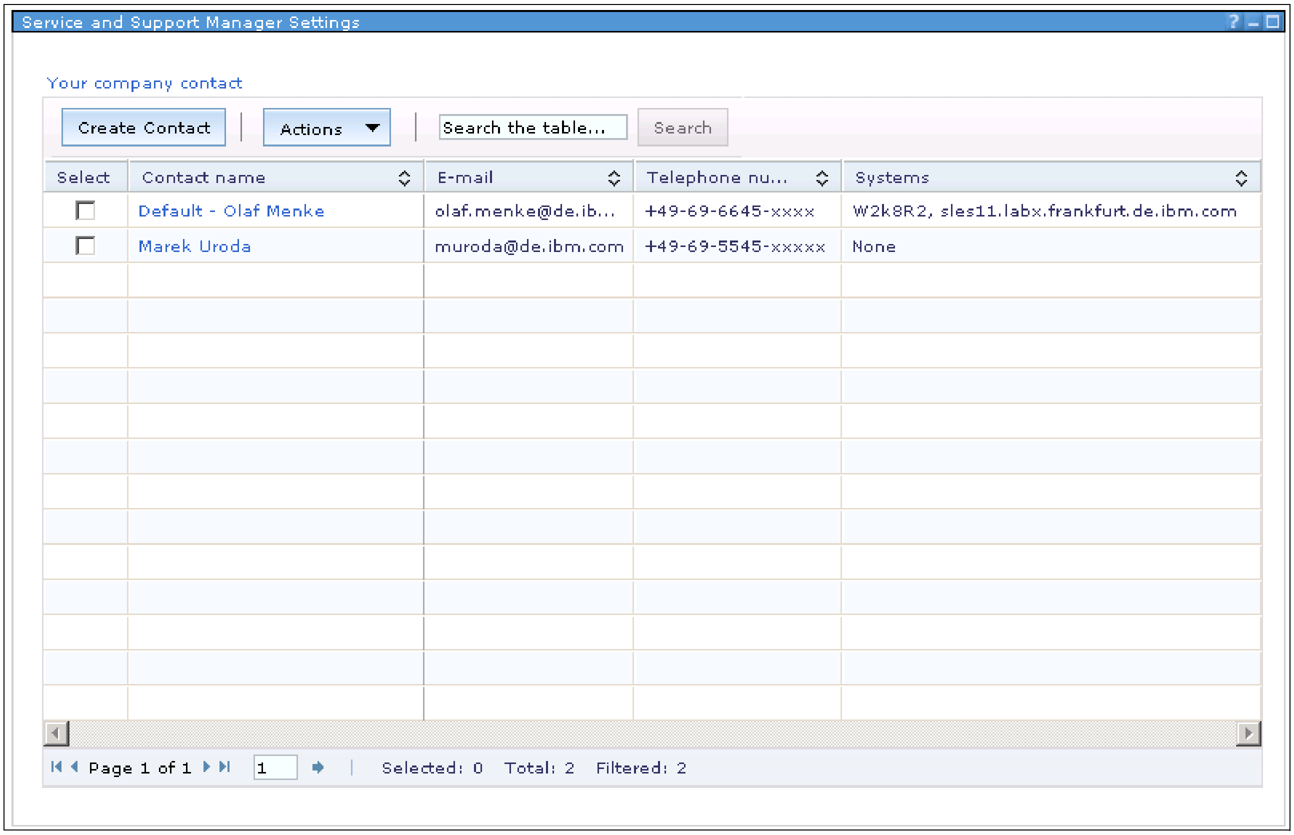

Select	Contact name	E-mail	Telephone nu...	Systems
□	Default - Olaf Menke	olaf.menke@de.ib...	+49-69-6645-xxxx	W2k8R2, sles11.labx.frankfurt.de.ibm.com
□	Marek Uroda	muroda@de.ibm.com	+49-69-5545-xxxxx	None

Page 1 of 1 1 Selected: 0 Total: 2 Filtered: 2

Figure 10-9 Managing your system contacts

When selecting Manage Settings, a window opens where you can perform the following functions:

- Set or change the IBM ID (like in the setup wizard before)
- Test the connectivity (like in the setup wizard before)
- Enable or disable the problem reporting and define settings for the reporting (Figure 10-10)

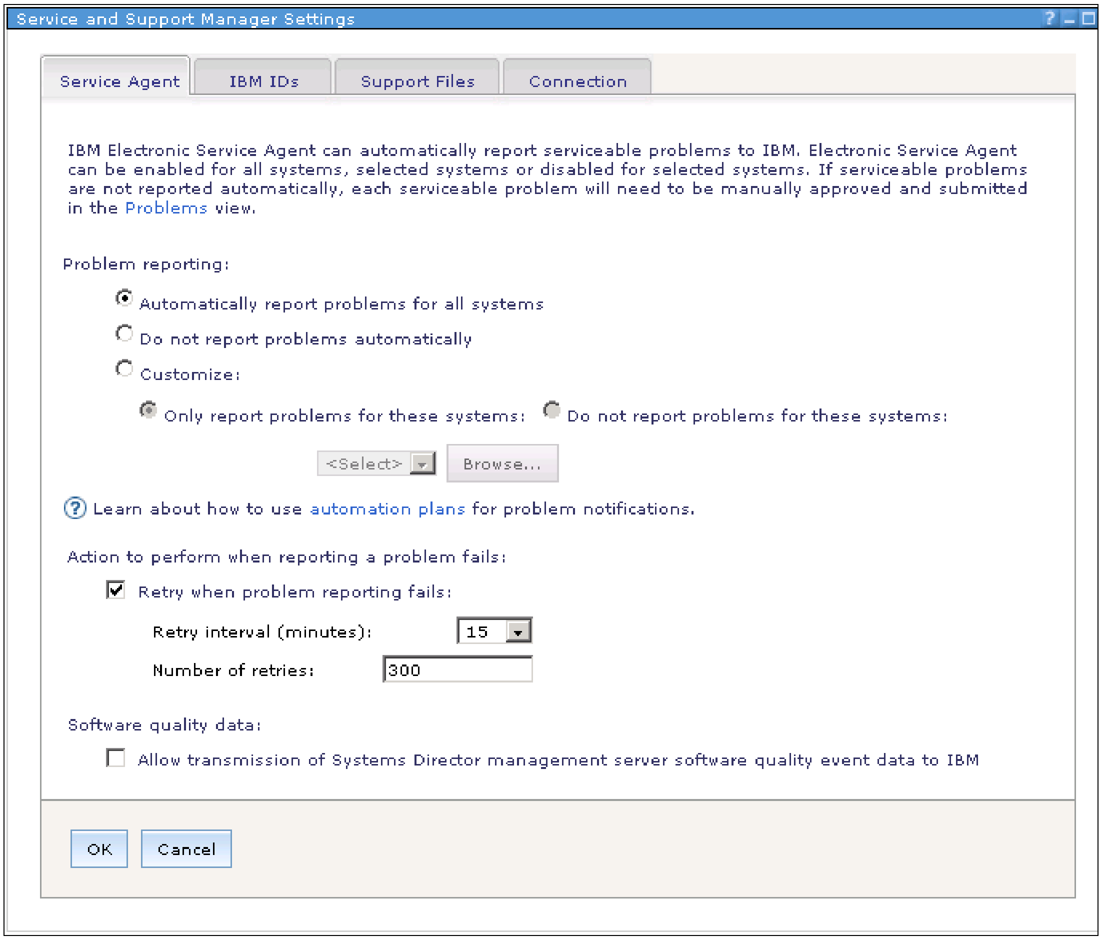

IBM Electronic Service Agent can automatically report serviceable problems to IBM. Electronic Service Agent can be enabled for all systems, selected systems or disabled for selected systems. If serviceable problems are not reported automatically, each serviceable problem will need to be manually approved and submitted in the Problems view.

Problem reporting:

● Automatically report problems for all systems

○ Do not report problems automatically

○ Customize:

● Only report problems for these systems: ○ Do not report problems for these systems:

Learn about how to use automation plans for problem notifications.

Action to perform when reporting a problem fails:

☑ Retry when problem reporting fails:

Retry interval (minutes): 15

Number of retries: 300

Software quality data:

□ Allow transmission of Systems Director management server software quality event data to IBM

Figure 10-10 Service Agent settings

- Define the cache settings for the support files (default is 500 MB), the lifetime for the support files (default is seven days), and the support file collector application settings (Figure 10-11 on page 445).

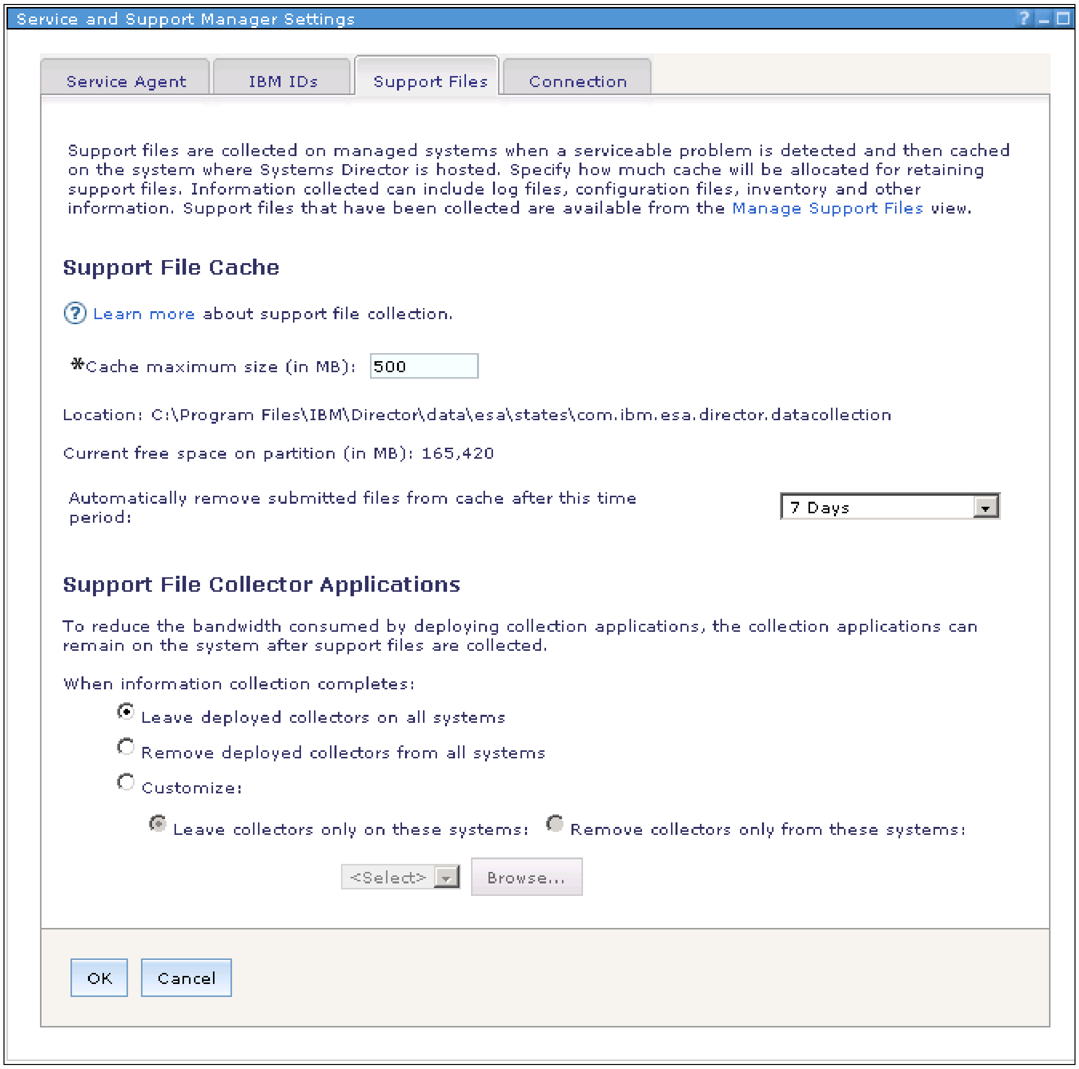

Support files are collected on managed systems when a serviceable problem is detected and then cached on the system where Systems Director is hosted. Specify how much cache will be allocated for retaining support files. Information collected can include log files, configuration files, inventory and other information. Support files that have been collected are available from the Manage Support Files view.

Support File Cache

Learn more about support file collection.

*Cache maximum size (in MB): 500

Location: C:\Program Files\IBM\Director\data\esa\states\com.ibm.esa.director.datacollection

Current free space on partition (in MB): 165,420

Automatically remove submitted files from cache after this time period: 7 Days

Support File Collector Applications

To reduce the bandwidth consumed by deploying collection applications, the collection applications can remain on the system after support files are collected.

When information collection completes:

- ● Leave deployed collectors on all systems
- ○ Remove deployed collectors from all systems
- ○ Customize:
  - ● Leave collectors only on these systems: ○ Remove collectors only from these systems:

Figure 10-11 Support file settings

10.3 Connectivity to IBM

For the Service and Support Manager to function, enable connectivity through your firewall to IBM. The addresses and ports that are used are listed in Table 10-1.

Table 10-1 SSM proxy

Host name	IP address	Port
www6.software.ibm.com	207.25.253.41	443
	192.109.81.20	443
download2.boulder.ibm.com	207.25.253.8	80
download3.boulder.ibm.com	207.25.253.76	80
eccgw01.boulder.ibm.com	207.25.252.197	443
eccgw02.rochester.ibm.com	129.42.160.51	443

Host name	IP address	Port
www-945.ibm.com	129.42.26.224 129.42.34.224 129.42.42.224	443
www.ibm.com	129.42.56.216 129.42.58.216 129.42.60.216	443 or 80
www-03.ibm.com	204.146.30.17	80

If you encounter problems, see the following information center link:

http://pic.dhe.ibm.com/infocenter/director/pubs/topic/com.ibm.director.tbs.helps.doc/fqm0_r_tbs_um_proxy_issues.html

10.4 Enabling systems for service and support

Perform the following tasks to enable a system for monitoring:

- The system is discovered.
- The system is unlocked.
- An inventory is collected.

After the tasks are complete, check whether the ESA agent is running on your endpoints. For guidance, see the following information center page:

http://pic.dhe.ibm.com/infocenter/director/pubs/topic/com.ibm.esa.director.help/esa_problem_optimize.html

To enable reporting to IBM, follow the steps in the following information center page:

http://pic.dhe.ibm.com/infocenter/director/pubs/topic/com.ibm.esa.director.help/esa_enable_disable_problem.html

To check the functionality, send a test problem by using the **Send test problem** link in the Common Tasks section, as shown in Figure 10-12. This should be done only if there might be a problem with sending events to IBM and when recommended by IBM service personal. Also, this sends a test problem and will create a hardware call at IBM.

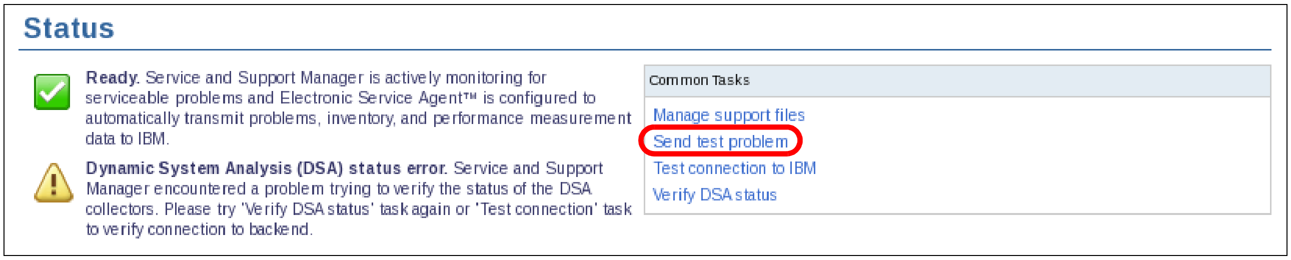

Status

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.

Dynamic System Analysis (DSA) status error. Service and Support Manager encountered a problem trying to verify the status of the DSA collectors. Please try 'Verify DSA status' task again or 'Test connection' task to verify connection to backend.

Figure 10-12 Electronic Service Agent Status panel

A reminder appears that advises you that this action sends an actual report to IBM support (Figure 10-13).

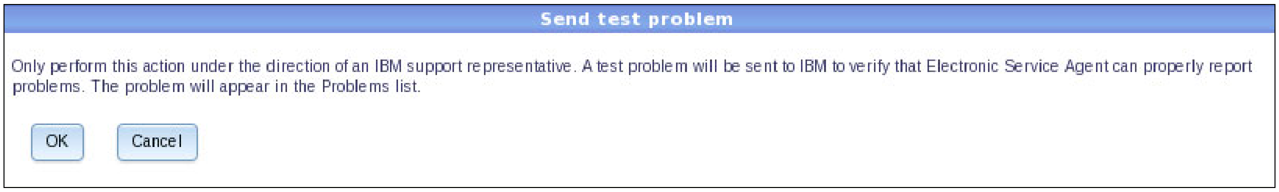

Send test problem

Only perform this action under the direction of an IBM support representative. A test problem will be sent to IBM to verify that Electronic Service Agent can properly report problems. The problem will appear in the Problems list.

Figure 10-13 Send test problem confirmation

After you click **OK**, the test problem report is submitted to IBM support. This report is visible under the dashboard and the Problem Reporting view (Figure 10-14).

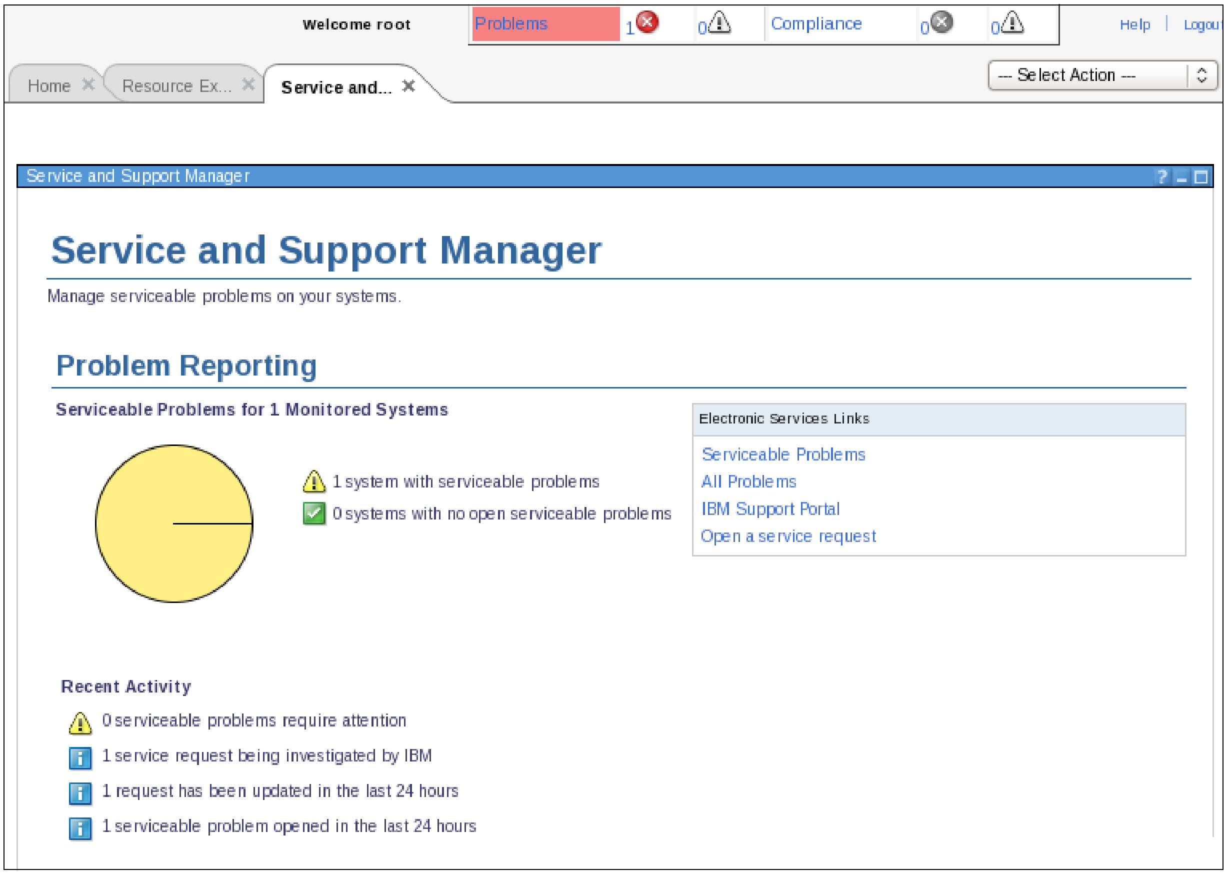

Service and Support Manager

Manage serviceable problems on your systems.

Problem Reporting

Serviceable Problems for 1 Monitored Systems

1 system with serviceable problems

0 systems with no open serviceable problems

Recent Activity

0 serviceable problems require attention

1 service request being investigated by IBM

1 request has been updated in the last 24 hours

1 serviceable problem opened in the last 24 hours

Figure 10-14 Service and Support Manager: Problem Reporting view

10.5 Serviceable event processing

Serviceable event processing is the management and transmission of serviceable events for hardware problems to IBM.

Serviceable events are determined by IBM and cannot be altered. The analysis component of Service and Support Manager determines whether the event warrants the creation of a serviceable event. Serviceable events are viewable in the Service and Support Manager plug-in, as shown in Figure 10-14.

If an event is serviceable and the Service and Support Manager is fully configured, the service request is transferred automatically to IBM unless otherwise configured.

Duplicate event processing is supported. Any duplicate event that is generated within a 24-hour window does not generate a new ticket with IBM.

10.6 Managing support files

After the problem is submitted to IBM support, additional data that is associated with the problem can be uploaded to IBM to help diagnose the problem. The data includes detailed system information, dump files, and event logs.

Take the following steps to view the support files and then submit them to IBM:

1. Under Common Tasks in the Service and Support Manager home page (Figure 10-2 on page 438), select **Manage Support Files**. The window that is shown in Figure 10-15 opens.

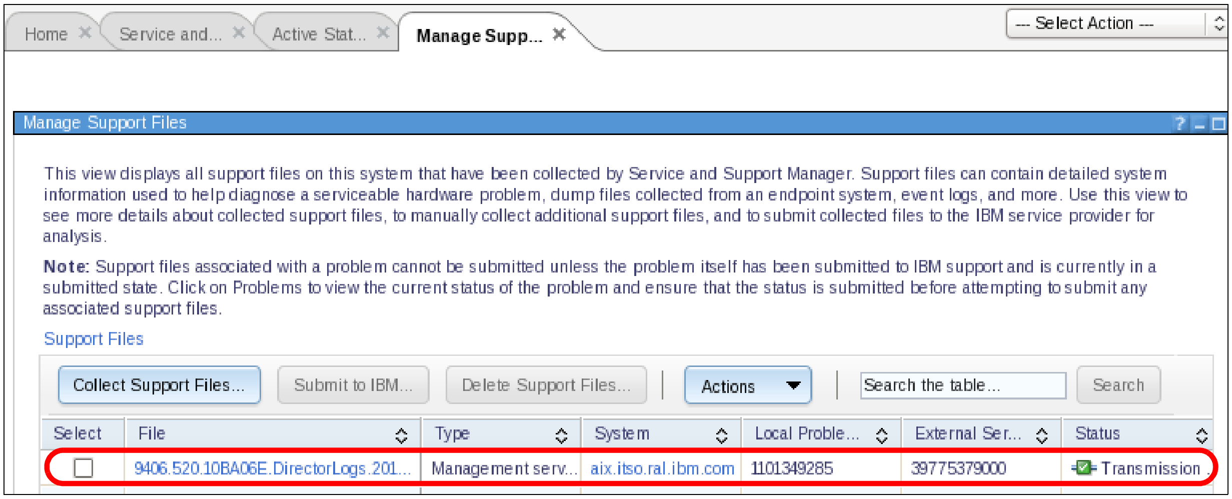

This view displays all support files on this system that have been collected by Service and Support Manager. Support files can contain detailed system information used to help diagnose a serviceable hardware problem, dump files collected from an endpoint system, event logs, and more. Use this view to see more details about collected support files, to manually collect additional support files, and to submit collected files to the IBM service provider for analysis.

Note: Support files associated with a problem cannot be submitted unless the problem itself has been submitted to IBM support and is currently in a submitted state. Click on Problems to view the current status of the problem and ensure that the status is submitted before attempting to submit any associated support files.

Support Files

Select	File	Type	System	Local Proble...	External Ser...	Status
□	9406.520.10BA06E.DirectorLogs.201...	Management serv...	aix.itso.ral.ibm.com	1101349285	39775379000	✓ Transmission ...

Figure 10-15 Manage Support Files panel

2. When you collect support files, select the *monitored system* where you want to collect the support files and click **Collect Support Files**. Use the predefined groups that are listed in Figure 10-8 on page 443.

3. Choose the target system as shown in Figure 10-16.

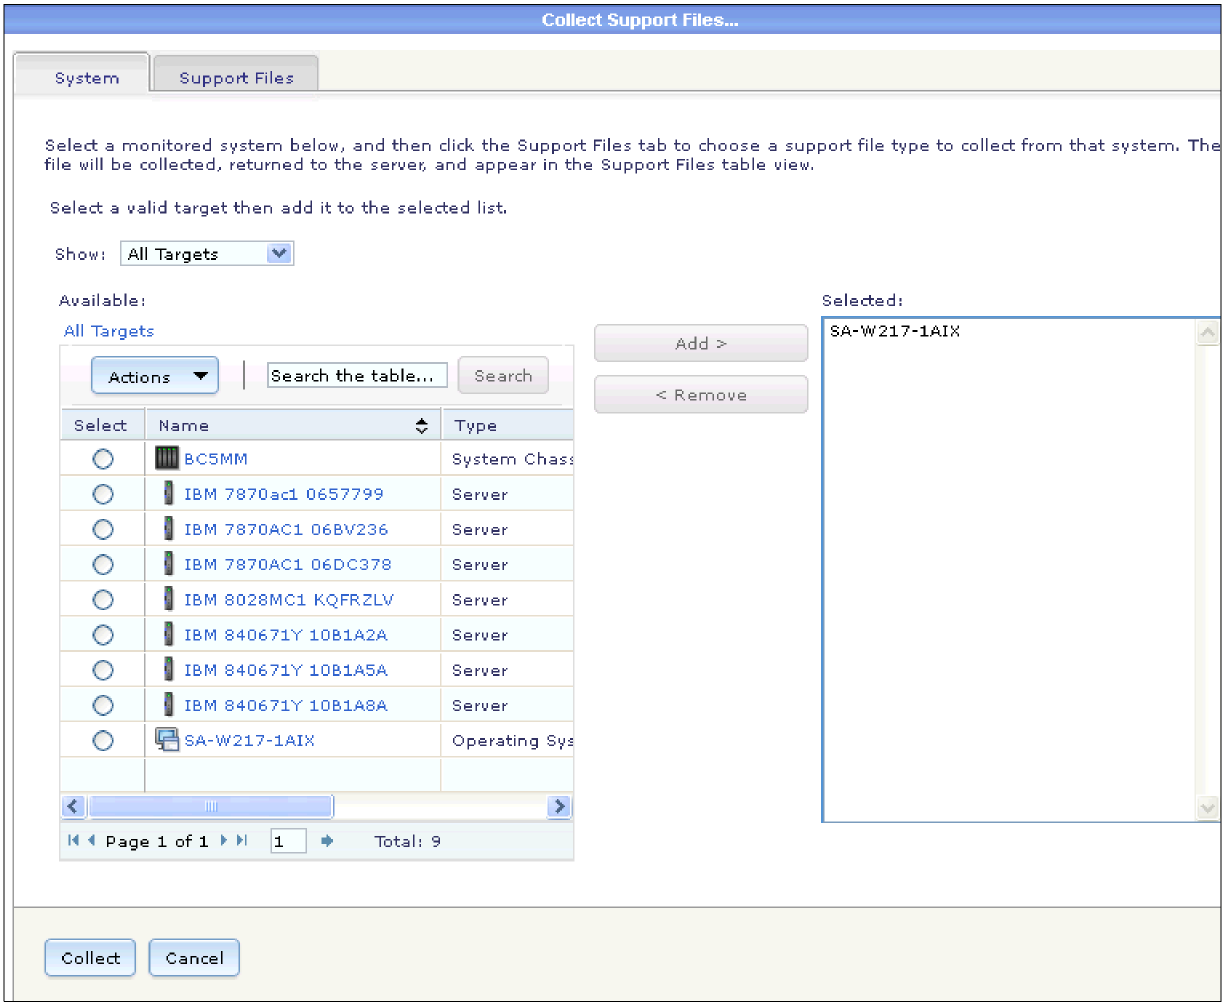

Select a monitored system below, and then click the Support Files tab to choose a support file type to collect from that system. The file will be collected, returned to the server, and appear in the Support Files table view.

Select a valid target then add it to the selected list.

Show: All Targets

Available:

All Targets

Select	Name	Type
○	BC5MM	System Chass
○	IBM 7870ac1 0657799	Server
○	IBM 7870AC1 06BV236	Server
○	IBM 7870AC1 06DC378	Server
○	IBM 8028MC1 KQFRZLV	Server
○	IBM 840671Y 10B1A2A	Server
○	IBM 840671Y 10B1A5A	Server
○	IBM 840671Y 10B1A8A	Server
○	SA-W217-1AIX	Operating Sys

Page 1 of 1 1 Total: 9

Selected:

SA-W217-1AIX

Figure 10-16 Monitored systems

4. As shown in Figure 10-17, click the **Support Files** tab to select the support files.

5. Figure 10-17 displays the types of support files to collect. Select the required option and click **Collect**.

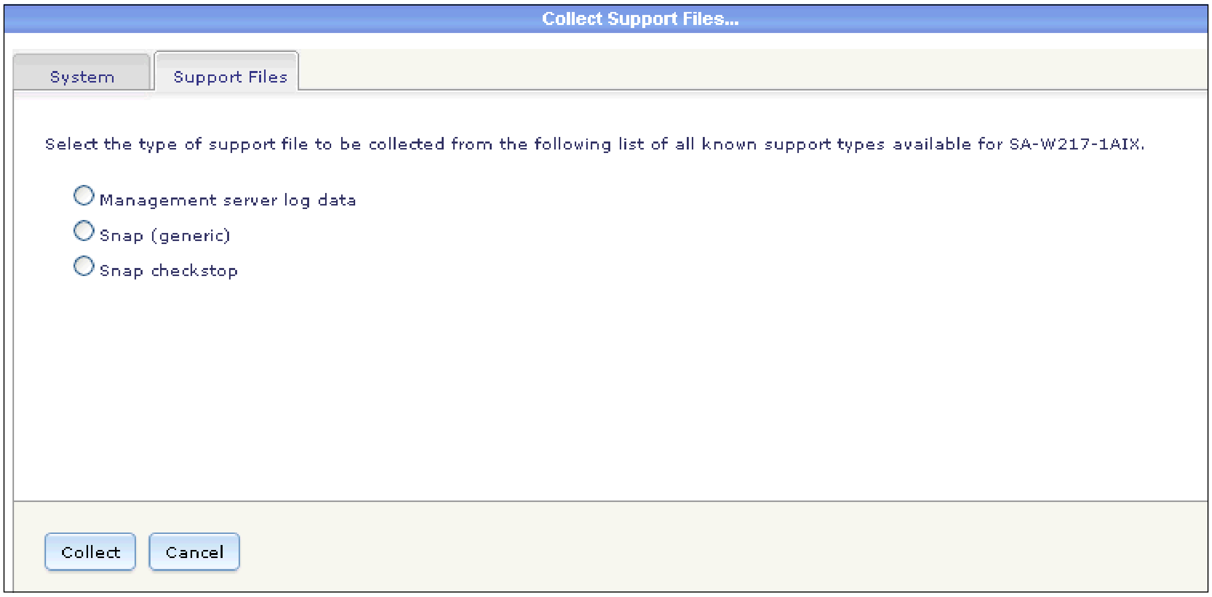

Select the type of support file to be collected from the following list of all known support types available for SA-W217-1AIX.

○ Management server log data
○ Snap (generic)
○ Snap checkstop

Figure 10-17 Collect Support Files panel

6. On the window that is shown in Figure 10-18, you can select the support files to send to IBM. After submission, you can delete files manually. However, Service and Support Manager removes support files after seven days after the successful file transmission of data to IBM.

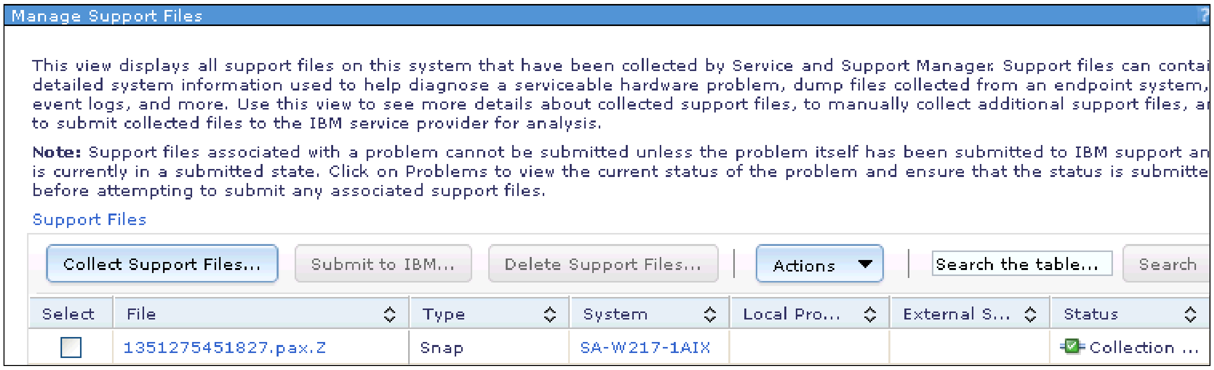

This view displays all support files on this system that have been collected by Service and Support Manager. Support files can contai detailed system information used to help diagnose a serviceable hardware problem, dump files collected from an endpoint system, event logs, and more. Use this view to see more details about collected support files, to manually collect additional support files, a to submit collected files to the IBM service provider for analysis.

Note: Support files associated with a problem cannot be submitted unless the problem itself has been submitted to IBM support an is currently in a submitted state. Click on Problems to view the current status of the problem and ensure that the status is submitte before attempting to submit any associated support files.

Select	File	Type	System	Local Pro...	External S...	Status
□	1351275451827.pax.Z	Snap	SA-W217-1AIX			Collection ...

Figure 10-18 Snap files

10.7 Best practices

To use the Service and Support Manager, apply the following best practices:

- *Enable* Service and Support Manager by clicking **Getting Started with Electronic Service Agent** and completing the necessary information.
- *Test connectivity* to the Internet to ensure that the Service and Support Manager can send events to IBM.
- *Prepare your systems*: Discover, unlock, and run inventory on your systems.

11

Storage Management solutions and Storage Control

This chapter describes the storage monitoring and management capabilities that are offered by IBM Systems Director server along with best practices.

The following topics are covered:

- 11.1, “Storage management solutions in IBM Systems Director” on page 452
- 11.2, “Base edition: IBM Systems Director Storage Manager” on page 457
- 11.3, “Advanced edition: IBM Systems Director Storage Control” on page 469

11.1 Storage management solutions in IBM Systems Director

IBM Systems Director supports monitoring and management (includes storage volume creation/deletion, zone creation/deletion, logical volume management) of a wide variety of storage devices, including disks, switches, internal Redundant Array of Independent Disks (RAID) controllers, and storage subsystems.

The IBM Systems Director portfolio offers storage monitoring and management capabilities in two editions:

- Base edition: IBM Systems Director Storage Manager, which is part of IBM Systems Director base edition.
- Advanced edition: IBM Systems Director Storage Control, which is an advanced paid plug-in of IBM Systems Director base edition. It comes with a 90-day trial license for users to evaluate the features.

This section explains both of the storage management editions, including introduction, support, and configuration, along with best practice tips for the supported tasks.

11.1.1 Terms to know

Before starting to read about IBM Systems Director storage management, it is important that you familiarize yourself with some of the storage terminologies that are used by IBM Systems Director:

- Storage volume

  A storage volume is the basic unit of storage, such as allocated space on a disk, or a single tape cartridge. *Logical unit number* (LUN) can be considered as a synonym to the term, *storage volume*.

- Storage pool

  A storage pool is a collection of storage volumes. Array can be considered as a synonym to the term, *storage pool*.

- SMI-S provider

  An *SMI-S provider* is a vendor-specific module that is used so that independent management software, such as IBM Systems Director, can manage a vendor device by using a standard interface that is based on the Common Information Model (CIM) protocol.

11.1.2 Storage device support and management

IBM Systems Director offers monitoring and management of the following storage types:

- Dedicated Local Storage, which is accessed with Integrated RAID Controllers (IRC).
- IBM BladeCenter integrated storage, which is accessed by using IBM BladeCenter S SAS RAID Controller Modules.
- Network storage, which is an external storage system that is accessed with storage switches, adapters, and protocols such as Fibre Channel, serial-attached SCSI (SAS), or Internet Small Computer System Interface (iSCSI).

The diagram that is shown in Figure 11-1 depicts the required and optional components that are involved in managing the storage devices from IBM Systems Director.

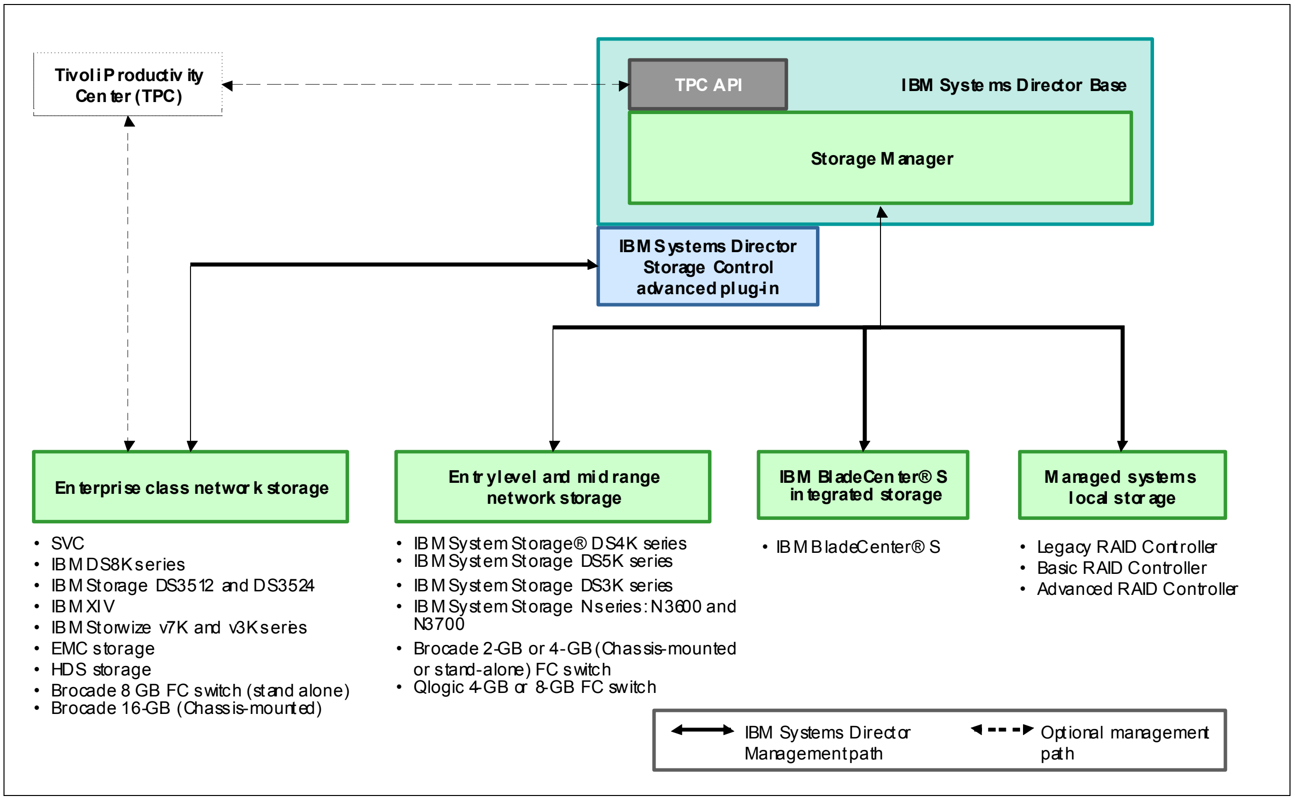

Tivoli Productivity Center (TPC)
TPC API
IBM Systems Director Base
Storage Manager
IBM Systems Director Storage Control advanced plug-in
Enterprise class network storage
- SVC
- IBM DS8K series
- IBM Storage DS3512 and DS3524
- IBM XIV
- IBM Storwize v7K and v3K series
- EMC storage
- HDS storage
- Brocade 8 GB FC switch (stand alone)
- Brocade 16-GB (Chassis-mounted)
Entry level and mid range network storage
- IBM System Storage® DS4K series
- IBM System Storage DS5K series
- IBM System Storage DS3K series
- IBM System Storage N series: N3600 and N3700
- Brocade 2-GB or 4-GB (Chassis-mounted or stand-alone) FC switch
- Qlogic 4-GB or 8-GB FC switch
IBM BladeCenter® S integrated storage
- IBM BladeCenter® S
Managed systems local storage
- Legacy RAID Controller
- Basic RAID Controller
- Advanced RAID Controller
IBM Systems Director Management path
Optional management path

Figure 11-1 IBM Systems Director storage device management

As shown in Figure 11-1, there are a few storage devices that are supported for management via IBM Systems Director Storage Manager, whereas some other storage devices require either the IBM Systems Director Storage Control or the IBM Tivoli Storage Productivity Center.

You can choose to use the optional management path that is shown in Figure 11-1 if your environment already has the IBM Tivoli Storage Productivity Center installed. In this case, the IBM Systems Director communicates with Tivoli Storage Productivity Center by using Tivoli Storage Productivity Center application programming interfaces (APIs) to monitor and manage the enterprise class network storage and Fibre Channel (FC) switches. IBM Systems Director can integrate with Tivoli Storage Productivity Center basic edition 4.1 or later releases.

Best Practice: Although integration with IBM Tivoli Productivity Center version 4.1 or later is supported, to get the additional storage device management support, it is recommended to upgrade IBM Tivoli Productivity Center to at least version 4.2.2 FP1+ before integrating it with the IBM Systems Director.

How does communication with storage devices happen?

IBM Systems Director storage manager and Storage Control plug-in communicates to the managed storage devices using the Storage Management Initiative Specification (SMI-S) protocol.

Note: Storage Management Initiative Specification is a storage standard developed and maintained by the Storage Networking Industry Association (SNIA). For more information, see the following website:

http://www.snia.org/tech_activities/standards/curr_standards/smi

Following are types of SMI-S providers with which the IBM Systems Director Storage Manager or Storage Control communicates for monitoring and managing the end-storage devices:

- External SMI-S providers, which are vendor-provided.
- Native SMI-S providers, which are built into the storage subsystems.
- IBM Systems Director Platform Agent: Has native interfaces to monitor and manage the local storage on the IBM System x servers.

Which storage devices are supported?

Table 11-1 depicts the list of supported storage devices by IBM Systems Director 6.3.2 along with the required providers and the management components.

Table 11-1 IBM Systems Director v6.3.2 storage support matrix

Hardware type	Model	SMI-S provider component	IBM Systems Director storage management component required
Integrated RAID Controller (IRC) attached to System x	LSI IRC 1064/1064e/1068/1078	00.32.05.xx/SMI-S 1.2 IBM Systems Director Platform Agent 6.3.2	IBM Systems Director Storage Manager
	LSI IRC 2208		
	LSI IRC M5014/ M5015/M5025		
	LSI IRC M5110		
	Adaptec ServeRAID 4/5/6/7/8/9		
IBM BladeCenter S SAS RAID Controller	RSSM (System x)	7.89, SMI-S 1.1 IBM Systems Director Platform Agent 6.3.0/6.3.1	

FC Switch	QLogic 4-GB	QLogic SMI-S 1.1 [embedded]	
	QLogic 8-GB		
	IBM Flex System Fabric CN4093 10-GB Converged Scalable Switch	IBM SMI-S 1.1 [embedded]	
	Brocade 2-GB or 4-GB (Chassis-mounted or stand-alone)	Brocade 120.10.0 (SMI-S 1.2)	
	Brocade 8-GB (stand-alone)	Brocade SMI Agent 120.11.0 or Brocade Network Advisor 11.1.x (with Integrated SMI Agent)	IBM Systems Director Storage Control
	Brocade 16-GB (Chassis-mounted)		
NetApp Storage	DS3300, DS3400, DS4100, DS4200, DS4300, DS4400, DS4500, DS4700, DS4800, DS5020, DS5100, and DS5300	Eagle 10.19.GG.xx, SMI-S 1.4	IBM Systems Director Storage Manager
	DS3512, DS3524	Native	IBM Systems Director Storage Control
IBM Storage	DS8300, 8700, SAN Volume Controller 4.3 and 5.1		
	DS8300; SAN Volume Controller 6.1, 6.2, and 6.3; IBM Storwize V7000		
	DS8300; SAN Volume Controller 6.1, 6.2, and 6.3; IBM Flex System V7000 Storage Node, IBM Storwize V7000		
	Storwize V3700 and Storwize V3500		
	IBM XIV® 2810-A14, 2812-A14		
	IBM XIV Gen 3 2810-114, 2812-114		
NAS storage	N-series (N3600 or N3700) - NFS only	NetApp 3.0.2 SMI-S 1.2	IBM Systems Director Storage Manager
EMC storage	CLARiiON AX Series, AX4 Series, CX Series, CX3 Series, CX4 Series, Symmetrix DMX Series, and V-Max Series3	EMC SMI-S Provider v4.2.0/Solutions Enabler v7.2-1108 0.0	IBM Systems Director Storage Control

Hitachi Data Systems	Various models mentioned at the following link: http://www-304.ibm.com/support/docview.wss?rs=0&uid=swg27019305#hds	Hitachi Device Manager 7.0	

For more information about where to install the required SMI-S providers, details about supported firmware versions, and restrictions that apply on storage management functionality as of IBM Systems Director Version 6.3.2, see the following site:

http://pic.dhe.ibm.com/infocenter/director/pubs/index.jsp?topic=%2Fcom.ibm.director.plan.helps.doc%2Ffqm0_r_hardware_compatibility_storage_devices.html

Interoperation with IBM Systems Director VMControl plug-in

IBM Systems Director VMControl is an advanced plug-in to IBM Systems Director, which is responsible for virtualization management including virtual machine lifecycle management, OS deployment, and server and storage pool management. For more information about VMControl, see Chapter 5, “VMControl” on page 247.

IBM Systems Director VMControl communicates with IBM Storage Manager or IBM Storage Control, depending on the storage device that is involved, for serving the operations involving creation and assignment of storage volumes from the network storage devices to the virtual machines. Following are the operations that need VMControl interactions with the Storage Manager or Storage Control plug-in:

- Creation and deletion of storage volumes in the network storage for virtual server deployments
- Creation and deletion of zones on the Fibre Channel switch for N-Port ID Virtualization (NPIV)-based virtual server deployments

Best Practice: Before performing any operations that involve dealing with assignment of new and existing volumes on the supported network storage using IBM Systems Director VMControl, ensure that IBM Systems Director Storage Manager or Storage Control is configured with the necessary providers (in Table 11-1 on page 454) to communicate with the network and storage devices involved.

IBM Systems Director VMControl supports an environment with NPIV on only Power Systems servers and only the following storage devices:

- SAN Volume Controller
- IBM Storwize V7000, Storwize V3700, and Storwize V3500
- IBM Flex System V7000 Storage Node
- Brocade switches:
  - Brocade 4-GB FC storage switches (chassis-mounted or stand-alone)
  - Brocade 8-GB FC storage switches (stand-alone)
  - Brocade 16-GB FC storage switches (chassis-mounted)

Table 11-2 shows the task that is supported by IBM Systems Director Storage Control working with IBM Systems Director VMControl plug-in.

Table 11-2 Storage volume allocation to the virtual servers by the Storage Control plug-in

Hypervisor platform	Operation	Supported by VMControl Express Edition	Supported by VMControl Standard or Enterprise Edition
PowerVM and RHEL Kernel-based Virtual Machine (KVM)	Server to Storage Mapping view	Yes	Yes
	Allocate new storage volume to the host server	Yes	Yes
	Allocate storage disks while deploying virtual servers from IBM Systems Director server	No	Yes
	Creating and working with Server System Pools and Storage System Pools	No	Yes
PowerVM	Create new virtual servers with: ▶ New and existing SAN disks with NPIV ▶ New and existing vSCSI-based VIOS local and SAN disks	Yes	Yes
	Edit virtual server and add: ▶ New and existing SAN disks with NPIV ▶ New and existing vSCSI-based VIOS local and SAN disks	Yes	Yes
Red Hat KVM	Create new virtual servers with new and existing SAN disks	Yes	Yes
	Edit existing virtual servers and add new and existing SAN disks	Yes	Yes

11.2 Base edition: IBM Systems Director Storage Manager

IBM Systems Director Storage Manager is part of the base IBM Systems Director server. There is no need for a separate installation of Storage Manager. This section describes the basic needs and best practices to be followed when using IBM Systems Director Storage Manager.

IBM Systems Director Storage Manager communicates to providers that are external to the IBM Systems Director server to manage the storage devices.

Following are the two components to which the Systems Director Storage Manager communicates to monitor and manage the storage devices:

- *IBM Systems Director Platform Agent*: To manage internal storage of the managed systems.

> **Note:** IBM Systems Director Platform Agent is also included as part of the IBM Systems Director Common Agent package.

- *The SMI-S Provider Agent*, which is released by the supported storage vendors: To monitor and manage the specific storage devices. This management software has many different names: SMI-S CIM provider, storage proxy, storage provider, or storage agent. It might also be referred to as the provider or agent as well. In this book, we refer to this software as SMI-S Provider Agent.

If your environment does not have any of the enterprise storage devices that you are planning to monitor and manage by using IBM Systems Director server, IBM Systems Director Storage Manager, working with SMI-S providers, can serve the need of monitoring and management of the existing storage devices.

11.2.1 Planning for storage management

To manage most storage systems with IBM Systems Director, you need to install an SMI-S provider released by the respective storage device vendor. This is required for all storage systems except IBM DS8000, XIV, SAN Volume Controller, and V7000 storage systems, which can be managed directly.

If your storage system requires an SMI-S Provider Agent, print Table 11-3 and complete the information and then use it when installing the SMI-S Provider Agent.

Table 11-3 Information to gather before installing an SMI-S Provider Agent

Information needed	Your information
Operating systems that the SMI-S CIM provider can run on, such as AIX or Windows Review the storage system product information to determine this.	
Server that is running the supported operating system For example, this would be the server that has supported an operating system for SMI-S Provider Agent installation.	
Ports for the SMI-S Provider Agent to use Determine whether the default CIM ports are in use. The default ports are 5988 (HTTP) and 5989 (HTTPS). ► If the default ports are available, use default ports when installing the SMI-S Provider Agent. ► If the default ports are not available, use non-default ports when you install the SMI-S Provider Agent. A common practice is to add 1000 to the port numbers. For example, you could specify ports 6988 (HTTP) and 6989 (HTTPS).	

11.2.2 SMI-S provider installation and configuration

For communicating with certain storage devices, IBM Systems Director Storage Manager needs an SMI-S provider. Table 11-1 on page 454 provides information about the storage devices and SMI-S provider version that is needed for IBM Systems Director Storage Manager.

Figure 11-2 shows the SMI-S Provider Agent connectivity to IBM Systems Director server.

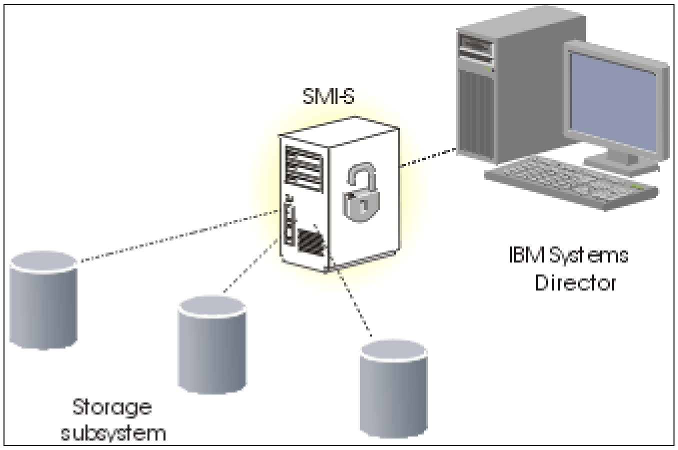

Figure 11-2 SMI-S provider connection illustration

For detailed information about where to download and how to configure these providers for IBM Systems Director V6.3.2, see the following site:

http://pic.dhe.ibm.com/infocenter/director/pubs/index.jsp?topic=%2Fcom.ibm.director.storage.helps.doc%2Ffqm0_t_sm_managing_smis_providers.html

Following are some of the main best practices to be followed when using the storage providers in your environment:

- Choose a system other than IBM Systems Director server or a systems hosting platform agent for installation of the SMI-S provider.
- Try to avoid installing more than one SMI-S provider on the same system. If there is a need to install more than one provider on the same server, ensure that the ports on which the provider agent CIM server listens are unique.
- Uninstalling some of the SMI-S providers does not remove the CIM component. If the system chosen for SMI-S provider installation was hosting any other SMI-S providers previously, ensure that the CIM components are removed completely before going ahead with a new installation of an SMI-S provider.
- There is no hard limit on how many storage systems can be supported by a single provider instance. The recommended maximum is 10 storage systems per provider.
- After SMI-S provider installation, ensure that both CIM object manager (CIMOM) and Service Location Protocol (SLP) services are active:
  - On AIX/Linux, run the following command to verify the status of CIM and SLP services:

    ```
    ps -ef | grep cim --> look for cimserver and provider processes
    ps-ef | grep slp --> look for slp related processes
    ```

  - On Microsoft Windows, look at “Services” to determine the CIM server and SLP service status.

- If you are reinstalling any SMI-S provider, it is recommended that you copy files that contain the IP addresses of the systems under management and the ports that are in use to a safe location.

11.2.3 IBM Systems Director Platform Agent installation

IBM Systems Director Storage Manager can monitor and manage the following devices by communicating to the IBM Systems Director Platform Agent on the managed endpoint:

- Integrated RAID Controller (IRC) attached to a System x system
- IBM BladeCenter S SAS RAID Controller: RSSM (System x)

For more information about the IBM Systems Director v6.3.2 supported IRC cards and the operating systems on which the platform agent installation is required, see the following information center link:

http://pic.dhe.ibm.com/infocenter/director/pubs/index.jsp?topic=%2Fcom.ibm.director.plan.helps.doc%2Ffqm0_r_hardware_compatibility_storage_devices.html

Information on installing IBM Systems Director Platform Agent can be found at the following information center link:

http://pic.dhe.ibm.com/infocenter/director/pubs/index.jsp?topic=%2Fcom.ibm.director.install.helps.doc%2Ffqm0_t_installing_platform_agent.html

Following are some of the best practices to be followed while installing and using IBM Systems Director Platform Agent:

1. Ensure that the system clocks on systems that contain IBM Systems Director server and Platform Agent remain synchronized.
2. On the system, where the agent installation is planned, run the IBM Systems Director pre-Installation utility to ensure that your system meets all the applicable requirements. If not, take appropriate actions to meet the requirements before proceeding with the installation.

   For instructions to run Pre-Installation Utility on Windows, see the following link:

   http://pic.dhe.ibm.com/infocenter/director/pubs/index.jsp?topic=%2Fcom.ibm.director.install.helps.doc%2Ffqm0_t_running_piu_win_agents.html

   For instructions to run Pre-Installation Utility on Linux, see the following link:

   http://pic.dhe.ibm.com/infocenter/director/pubs/index.jsp?topic=%2Fcom.ibm.director.install.helps.doc%2Ffqm0_t_running_piu_aix_agents.html
3. If you are installing IBM Systems Director Platform Agent on a Linux system, ensure that the Linux system has all prerequisite RPMs. See the following link to get the list of all prerequisite RPMs:

   http://pic.dhe.ibm.com/infocenter/director/pubs/index.jsp?topic=%2Fcom.ibm.director.install.helps.doc%2Ffqm0_t_preparing_to_install_core_on_xseries.html
4. After you install IBM Systems Director Platform Agent, before you discover the platform agent in the IBM Systems Director server, run the following SLP query from the IBM Systems Director server to ensure that the platform agent service is advertised properly and that there is no firewall blocking between the platform agent and the IBM Systems Director server.

You can run either the **slptool** command or **slp_query** tool on the IBM Systems Director server based on the operating system where you installed the IBM Systems Director server.

On Linux, run the following command:

```
/opt/ibm/icc/bin/slptool -u <ip_of_managed_resource> findsrvs
service:management-software.IBM:platform-agent
```

On Windows, run the following command:

```
slp_query --type=service:management-software.IBM:platform-agent
--address=<ip_of_managed_resource>
```

11.2.4 IBM Systems Director Storage Manager: Discovery

Similar to managing other type of endpoints, discover the storage devices as a first step before monitoring and managing them through IBM Systems Director server.

We use an example of discovering an SMI-S Provider Agent for IBM DS4000 series SAN storage in IBM Systems Director:

1. As a first step, install and configure the SMI-S Provider Agent for the storage device. Each SAN component that is managed by Storage Manager requires an SMI-S Provider Agent to be installed and configured. SMI-S Provider Agent can be installed on any system. For more information about installing SMI-S Provider Agent, see “SMI-S provider installation and configuration” on page 459.
2. On the IBM Systems Director server console, discover the SMI-S provider using the following steps:
   a. Log on to the IBM Systems Director server console, go to **Inventory** → **System Discovery**. Click **Create new profile** under Advanced Tasks, as shown in Figure 11-3.

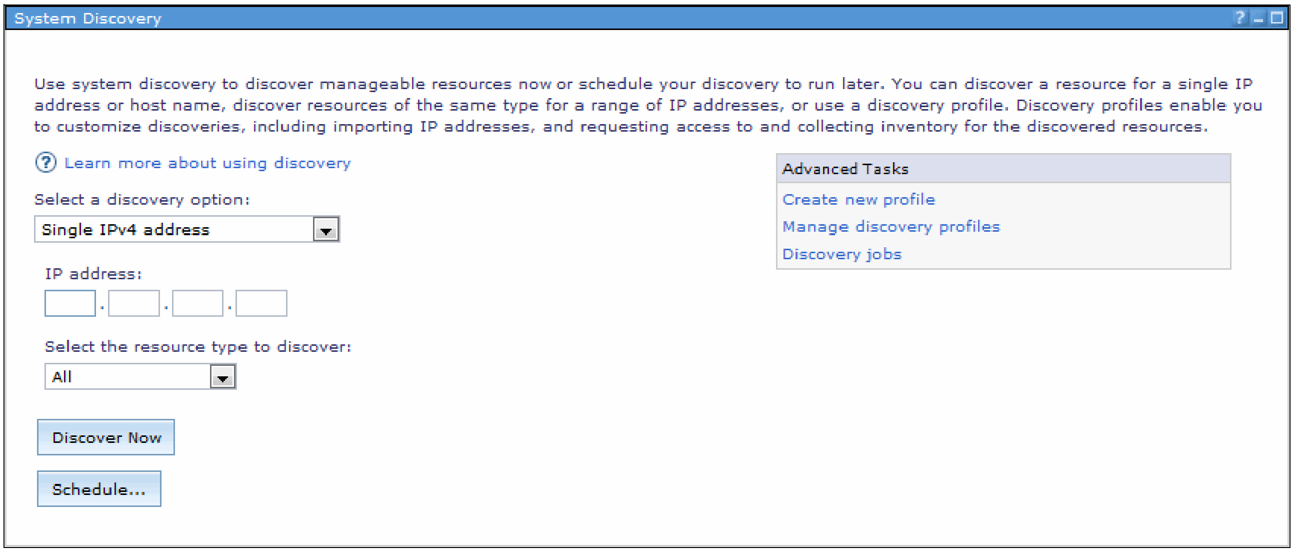

System Discovery

Use system discovery to discover manageable resources now or schedule your discovery to run later. You can discover a resource for a single 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 inventory for the discovered resources.

Select a discovery option: Single IPv4 address

IP address:

Select the resource type to discover: All

Figure 11-3 System Discovery: Create new profile

   b. On the Discovery Profile Wizard, click **Next** on the Welcome page.

c. On the Profile Properties panel, as shown in Figure 11-4, enter the profile name and select the resource type, **Operating System**. Select **All** for the resource subtype. Click **Next**.

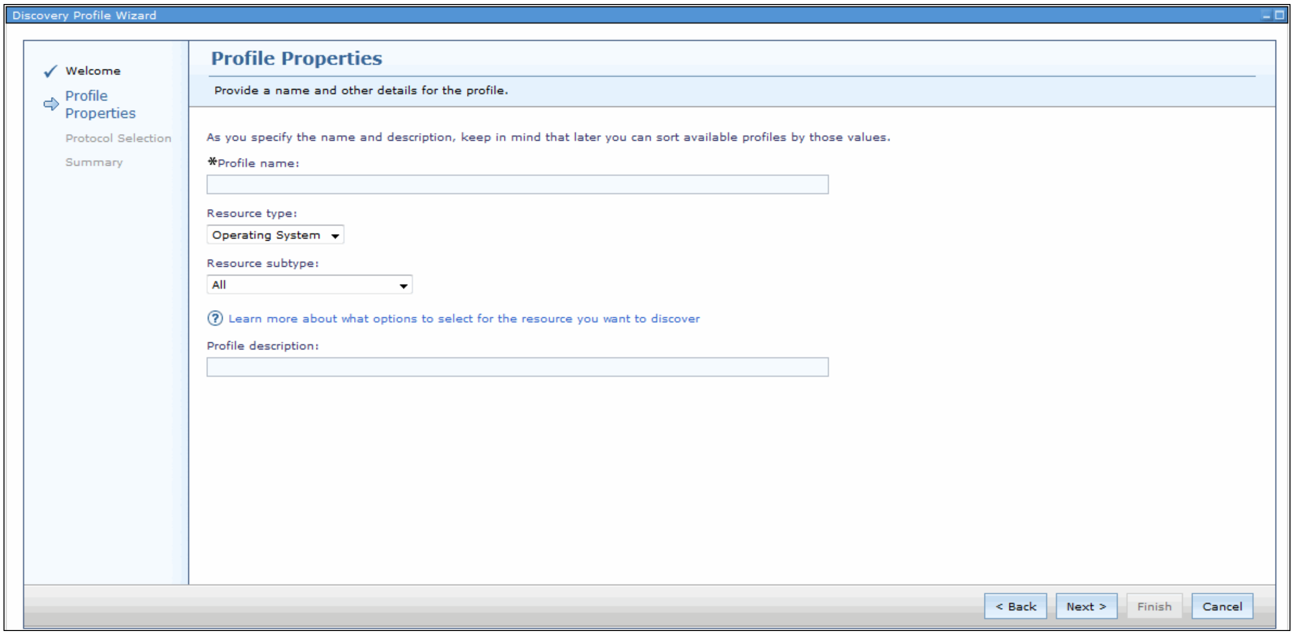

Profile Properties

Provide a name and other details for the profile.

As you specify the name and description, keep in mind that later you can sort available profiles by those values.

*Profile name:

Resource type: Operating System

Resource subtype: All

Learn more about what options to select for the resource you want to discover

Profile description:

Figure 11-4 Discovery Profile Wizard: Profile Properties page

d. On the Protocol Selection page, select **Storage Management Initiative Specification (SMI-S) Discovery** from the list of protocols that are displayed, as shown in Figure 11-5. Click **Next**.

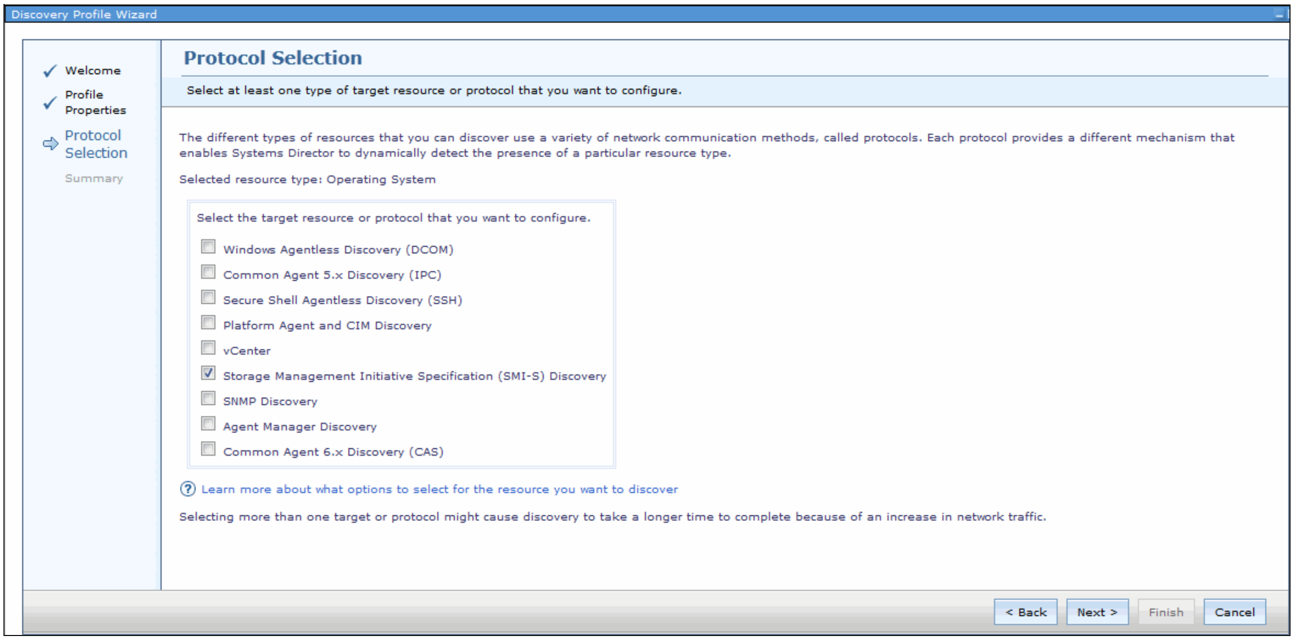

Protocol Selection

Select at least one type of target resource or protocol that you want to configure.

The different types of resources that you can discover use a variety of network communication methods, called protocols. Each protocol provides a different mechanism that enables Systems Director to dynamically detect the presence of a particular resource type.

Selected resource type: Operating System

Select the target resource or protocol that you want to configure.

- ☐ Windows Agentless Discovery (DCOM)
- ☐ Common Agent 5.x Discovery (IPC)
- ☐ Secure Shell Agentless Discovery (SSH)
- ☐ Platform Agent and CIM Discovery
- ☐ vCenter
- ☑ Storage Management Initiative Specification (SMI-S) Discovery
- ☐ SNMP Discovery
- ☐ Agent Manager Discovery
- ☐ Common Agent 6.x Discovery (CAS)

Learn more about what options to select for the resource you want to discover

Selecting more than one target or protocol might cause discovery to take a longer time to complete because of an increase in network traffic.

Figure 11-5 Discovery Profile Wizard: Protocol Selection page

e. On the SMI-S Configuration page, perform the steps that are shown in Figure 11-6:

   i. Choose SMI-S **Direct connection** discovery.

   ii. Choose the hardware type depending on the type of storage device that you have.

   iii. Select the type of protocol that is configured on the SMI-S Provider Agent.

   iv. Enter the IP address and port details of the SMI-S Provider Agent.

   v. Click **Next**.

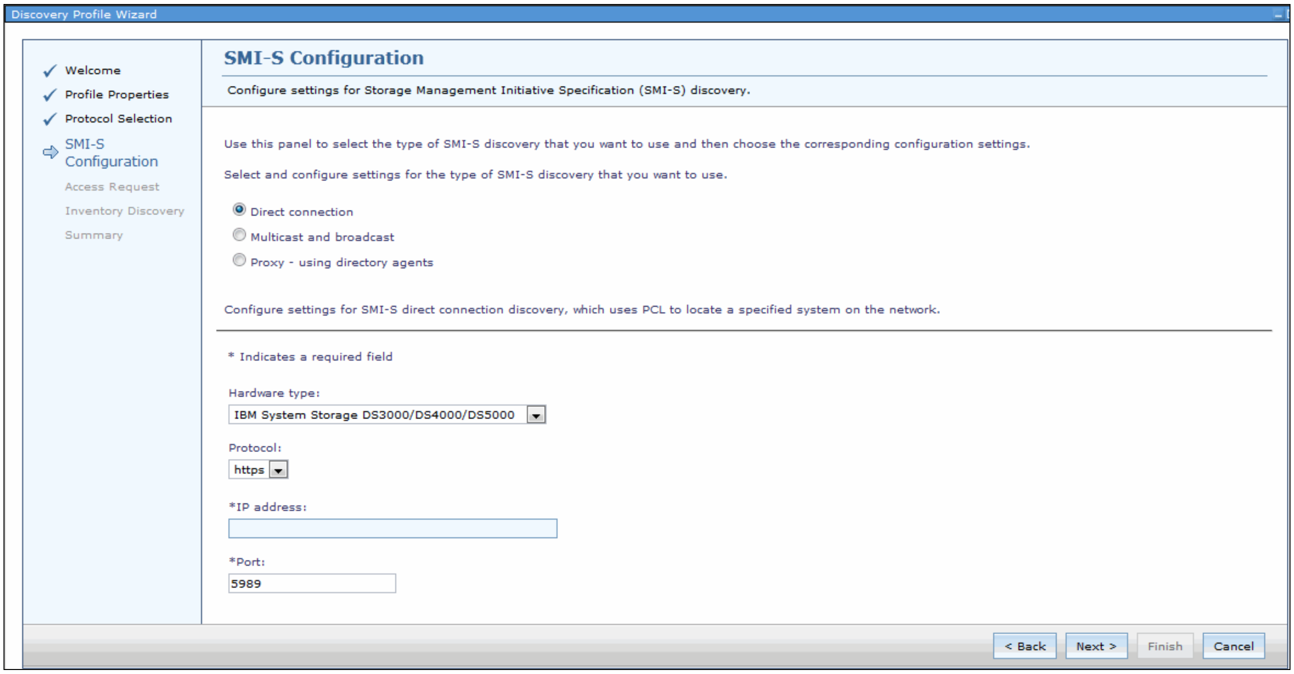

SMI-S Configuration

Configure settings for Storage Management Initiative Specification (SMI-S) discovery.

Use this panel to select the type of SMI-S discovery that you want to use and then choose the corresponding configuration settings.

Select and configure settings for the type of SMI-S discovery that you want to use.

● Direct connection
○ Multicast and broadcast
○ Proxy - using directory agents

Configure settings for SMI-S direct connection discovery, which uses PCL to locate a specified system on the network.

* Indicates a required field

Hardware type: IBM System Storage DS3000/DS4000/DS5000

Protocol: https

*IP address:

*Port: 5989

Figure 11-6 Discovery Profile Wizard: SMI-S Configuration page

f. *Optionally*, on the Access Request page you can provide the access details of the operating system where SMI-S Provider Agent has been installed. Click **Next**.

g. *Optionally*, on the Inventory Discovery page you can configure settings to automatically discover (collect) inventory. Click **Next**.

h. On the Summary page, review the summary of your profile selections. If there are no modifications required, click **Finish** to create a new profile.

3. On the System Discovery page that is shown in Figure 11-7 on page 464, perform the following steps:

   a. Under “Select a discovery option:”, choose **Select a discovery profile to run**.
   b. Select newly created discovery profile.
   c. Click **Discover Now**.

   When the discovery job is completed, you are able to see the SMI-S Provider Agent system and the managed storage devices, if you provided the access details in the discovery profile already.

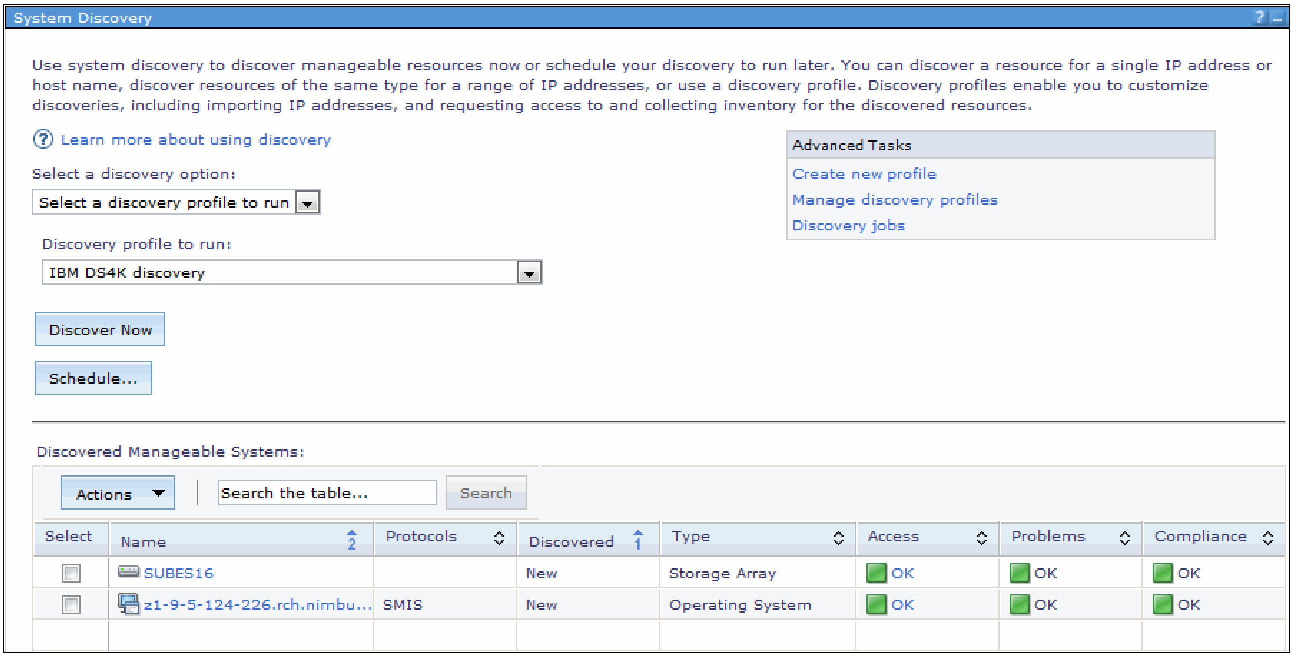

Use system discovery to discover manageable resources now or schedule your discovery to run later. You can discover a resource for a single IP address or host name, discover resources of the same type for a range of IP addresses, or use a discovery profile. Discovery profiles enable you to customize discoveries, including importing IP addresses, and requesting access to and collecting inventory for the discovered resources.

Select a discovery option: Select a discovery profile to run

Discovery profile to run: IBM DS4K discovery

Advanced Tasks
Create new profile
Manage discovery profiles
Discovery jobs

Discovered Manageable Systems:

Select	Name	Protocols	Discovered	Type	Access	Problems	Compliance
□	SUBES16		New	Storage Array	🟩 OK	🟩 OK	🟩 OK
□	z1-9-5-124-226.rch.nimbu...	SMIS	New	Operating System	🟩 OK	🟩 OK	🟩 OK

Figure 11-7 Discover SMI-S Provider Agent using the customized discovery profile

> **Note:** If you have IBM Systems Director Storage Control activated in your environment, following the steps mentioned above to discover a storage device creates a data source in the Storage Control [e-Tivoli Storage Productivity Center basic edition] database.

11.2.5 IBM Systems Director Storage Manager: Inventory collection

The inventory collection operation should be the first step performed after discovering any storage devices in the IBM Systems Director server.

> **Best Practice:** Users should ensure that the inventory collection is done for each endpoint before running any IBM Systems Director or plug-in tasks. This ensures that the IBM Systems Director server tasks when run would be starting from a known state. In other words, IBM Systems Director server and its plug-ins would be aware of the environment on which they were installed and activated. The inventories on the endpoints can be done in any order, or even at the same time.

If there are any changes or tasks, such as volume addition or deletion done on the storage devices using native management consoles, an inventory rerun should be performed on those devices for IBM Systems Director server to reflect the changes made.

To collect inventory on storage devices, perform the following steps:

1. On the IBM Systems Director console, go to the Resource Explorer page.
2. Click the **All Storage Systems** group, right-click the required storage device, and select **Inventory** → **Collect Inventory**. Run the inventory job.

3. When the inventory collection succeeds, you are able to see the amount of raw storage that is available, usable capacity of the storage, and available capacity, as shown in Figure 11-8.

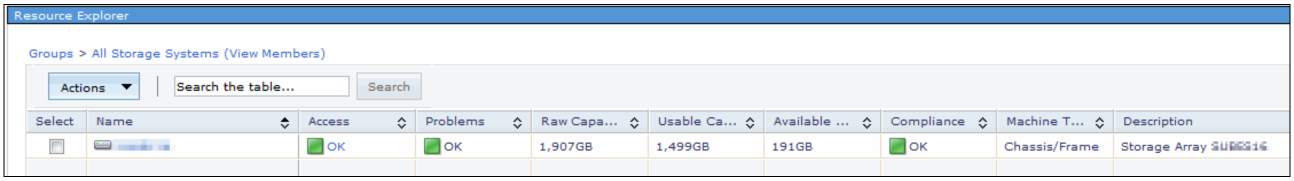

Select	Name	Access	Problems	Raw Capa...	Usable Ca...	Available ...	Compliance	Machine T...	Description
☐	[redacted]	🟩 OK	🟩 OK	1,907GB	1,499GB	191GB	🟩 OK	Chassis/Frame	Storage Array [redacted]

Figure 11-8 Resource Explorer → All Storage Systems (View Members)

4. Right-clicking the storage device and selecting **Related Resources**, as shown in Figure 11-9, allows you to see related disk drives, IP interfaces, storage volumes, and storage pools information.

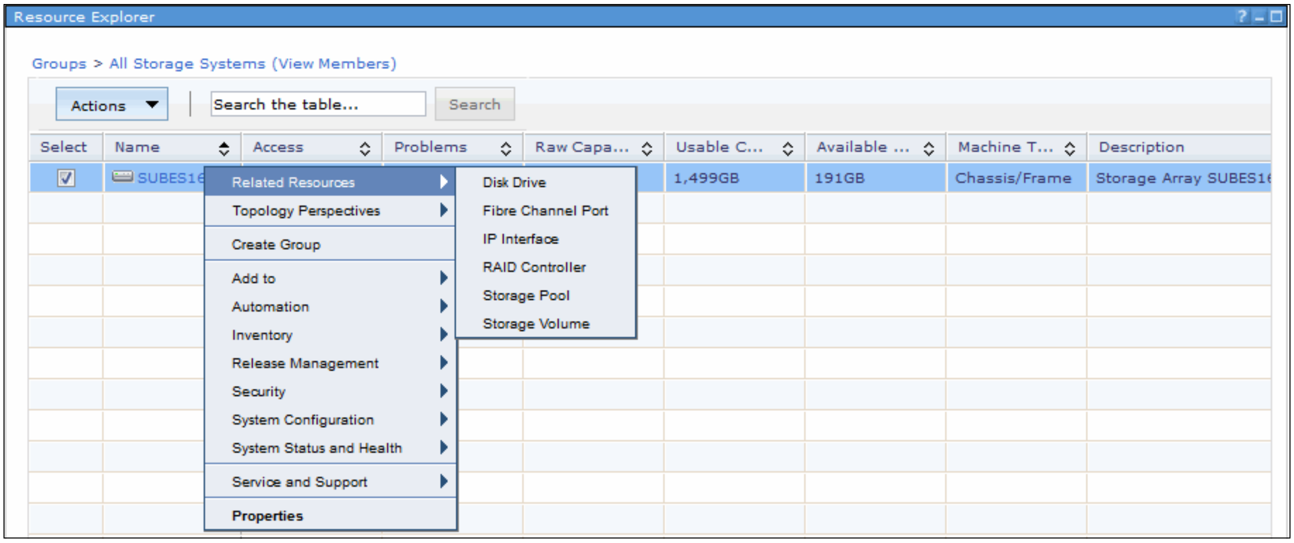

Select	Name	Access	Problems	Raw Capa...	Usable C...	Available ...	Machine T...	Description
☑	SUBES16				1,499GB	191GB	Chassis/Frame	Storage Array SUBES16

Figure 11-9 Storage device Related Resources page

11.2.6 IBM Systems Director Storage Manager: Managing storage

You can perform real-time management of the storage systems that are attached to IBM BladeCenter or System x systems.

IBM Systems Director Storage Manager supports management of these storage systems that are attached to or integrated with IBM BladeCenter and System x systems:

- Integrated RAID Controllers
- ServeRAID MR Controllers
- BladeCenter SAS Connectivity Modules
- IBM BladeCenter S SAS RAID Controller Modules

For more information about how to configure and manage the preceding devices, see the following link:

http://pic.dhe.ibm.com/infocenter/director/pubs/index.jsp?topic=%2Fcom.ibm.director.storage.helps.doc%2Ffqm0_t_sm_managing_storage_in_real_time.html

In this book, we use an example of discovering an IBM System x86 server that has IRC running IBM Systems Director Platform Agent v6.3.2 IBM Systems Director:

1. Install the platform agent on the IBM System x86 that has supported IRC. Refer to “IBM Systems Director Platform Agent installation” on page 460 for more information about installation and best practices.
2. On the IBM Systems Director server web console, go to **Inventory** → **System Discovery**.
3. Input the IP address and host name of the server that needs to be discovered. Select **Operating System** as the resource type. Click **Discover Now**, as shown in Figure 11-10.

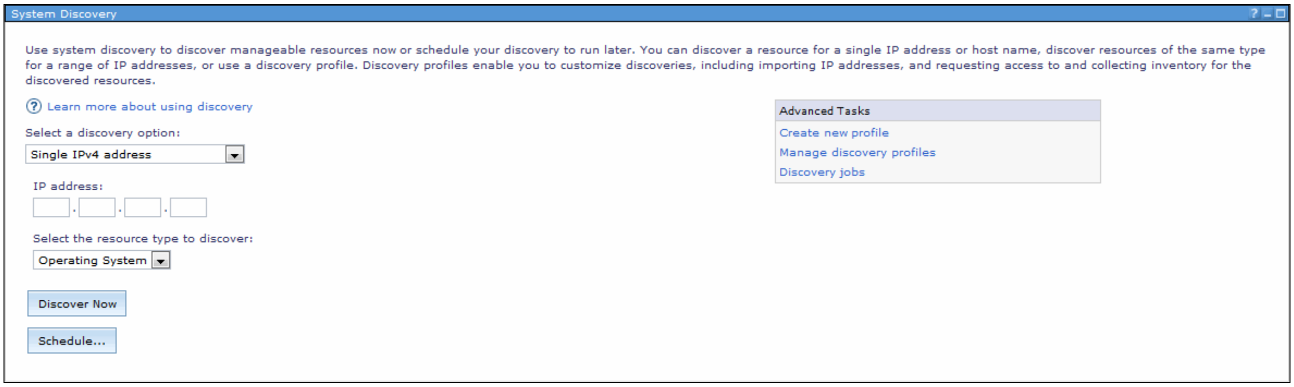

System Discovery

Use system discovery to discover manageable resources now or schedule your discovery to run later. You can discover a resource for a single 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 inventory for the discovered resources.

Learn more about using discovery

Select a discovery option: Single IPv4 address

IP address: ___.___.___.___

Select the resource type to discover: Operating System

Figure 11-10   System Discovery page

4. When the discovery is complete, the operating system (OS) instance is listed under Discovered Manageable Systems, as shown in Figure 11-11.

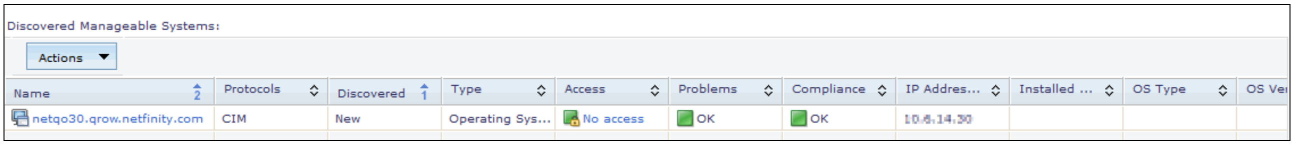

Discovered Manageable Systems:

Name	Protocols	Discovered	Type	Access	Problems	Compliance	IP Addres...	Installed ...	OS Type	OS Ve
netqo30.qrow.netfinity.com	CIM	New	Operating Sys...	No access	🟩 OK	🟩 OK	10.6.14.30			

Figure 11-11   Discovered Manageable Systems panel

5. Request access to the discovered endpoint by right-clicking and select **Security** → **Request Access**. Provide the user name and credentials that have root privileges to unlock the OS instance endpoint. When the request access is successful, close the Request Access page and return to the System Discovery page.
6. If you have the management controller/service processor on the server hosting the OS instance that was discovered, on the network reachable by IBM Systems Director server, you see a server instance that is created. The instance is in a “locked” mode, as shown in Figure 11-12 on page 467.

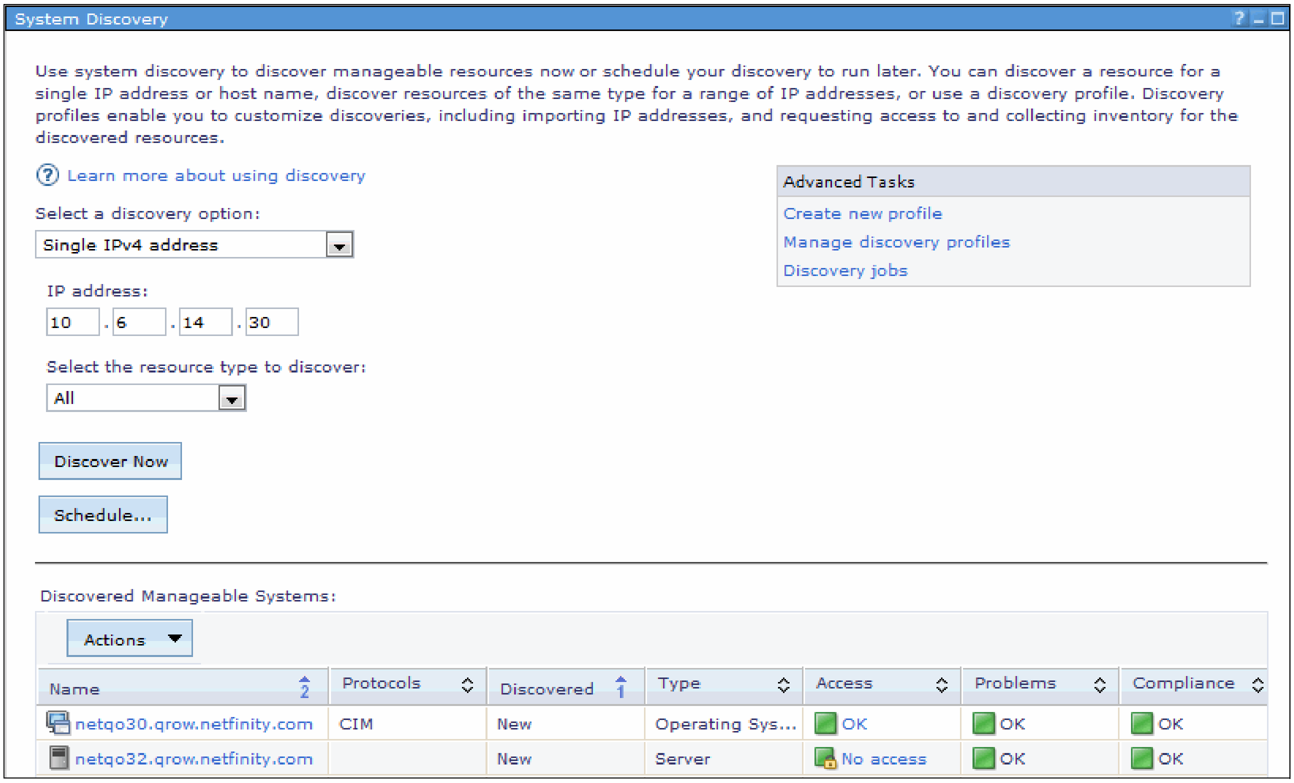
System Discovery

Use system discovery to discover manageable resources now or schedule your discovery to run later. You can discover a resource for a single 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 inventory for the discovered resources.

Learn more about using discovery

Select a discovery option: Single IPv4 address

IP address: 10 . 6 . 14 . 30

Select the resource type to discover: All

Advanced Tasks
- Create new profile
- Manage discovery profiles
- Discovery jobs

Discovered Manageable Systems:

Name	Protocols	Discovered	Type	Access	Problems	Compliance
netqo30.qrow.netfinity.com	CIM	New	Operating Sys...	🟩 OK	🟩 OK	🟩 OK
netqo32.qrow.netfinity.com		New	Server	No access	🟩 OK	🟩 OK

Figure 11-12 After getting access to the OS instance

7. Request access to the server instance by right-clicking and select **Security** → **Request Access**. Input the management controller/service processor user name and credentials to unlock the server instance endpoint.
8. Go to **Resource Explorer** → **Groups** → **All Systems (View Members)**. Right-click the OS instance and select **Inventory** → **Collect Inventory**. Run the inventory collection and wait until the job is completed.

> **Note:** Running inventory collection is a necessary step to get the storage management options for the system having supported-IRC.

9. When inventory collection is complete, go to **Resource Explorer** → **Groups** → **All Systems (View Members)**. Right-click the OS instance. Select **System Configuration** → **Storage**, as shown in Figure 11-13 on page 468. This option can be used to view and manage the storage that is associated with the managed system.

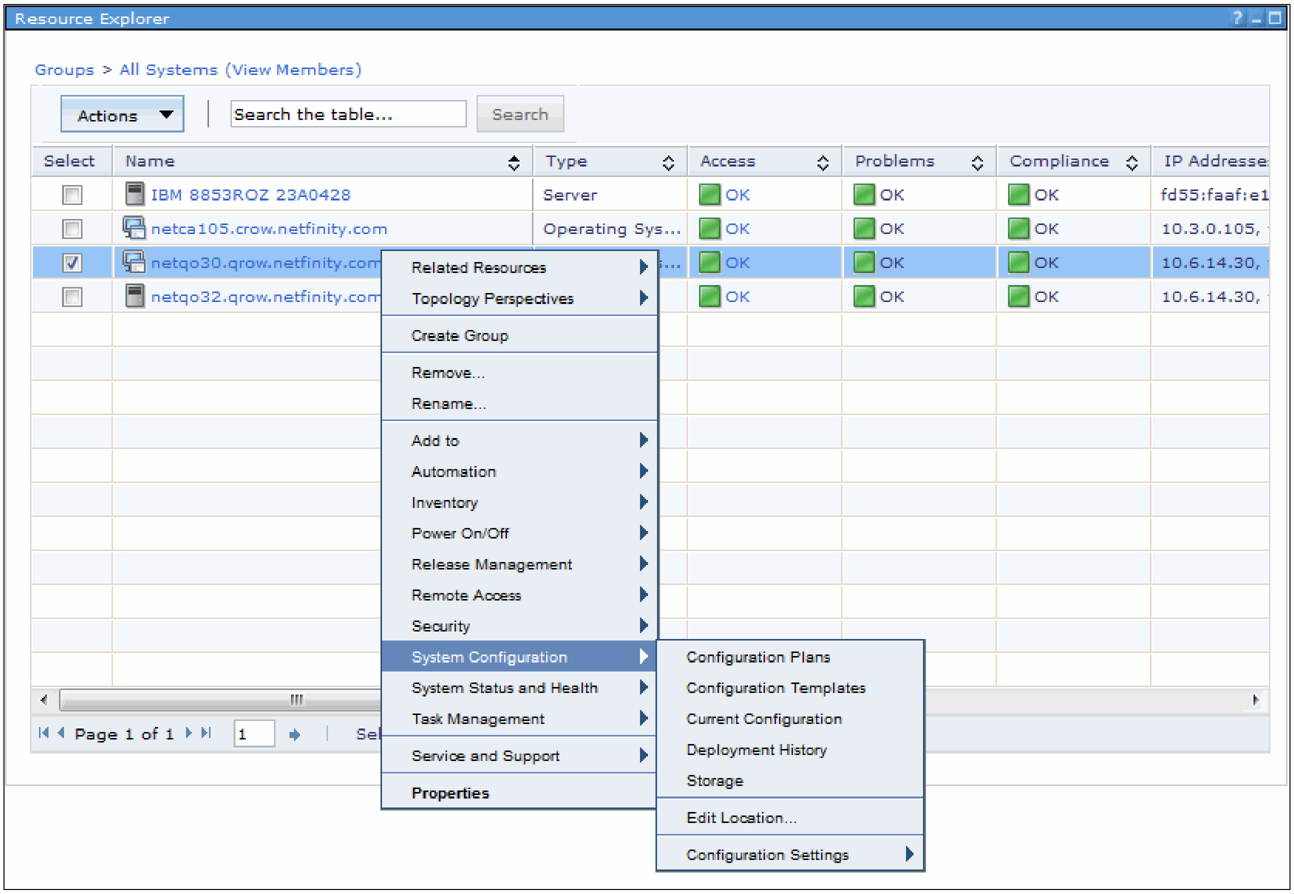

Select	Name	Type	Access	Problems	Compliance	IP Addresse
☐	IBM 8853ROZ 23A0428	Server	🟩 OK	🟩 OK	🟩 OK	fd55:faaf:e1
☐	netca105.crow.netfinity.com	Operating Sys...	🟩 OK	🟩 OK	🟩 OK	10.3.0.105,
☑	netqo30.qrow.netfinity.com	...s...	🟩 OK	🟩 OK	🟩 OK	10.6.14.30,
☐	netqo32.qrow.netfinity.com		🟩 OK	🟩 OK	🟩 OK	10.6.14.30,

Figure 11-13 Storage monitor and manage option for IRC on IBM System x86 server

10. Clicking the **System Configuration** → **Storage** option leads you to the Storage page where, as shown in Figure 11-14, you can view the local storage devices and volumes that are associated with the system and perform any necessary tasks that are associated with storage volumes, such as creating and deleting volumes.

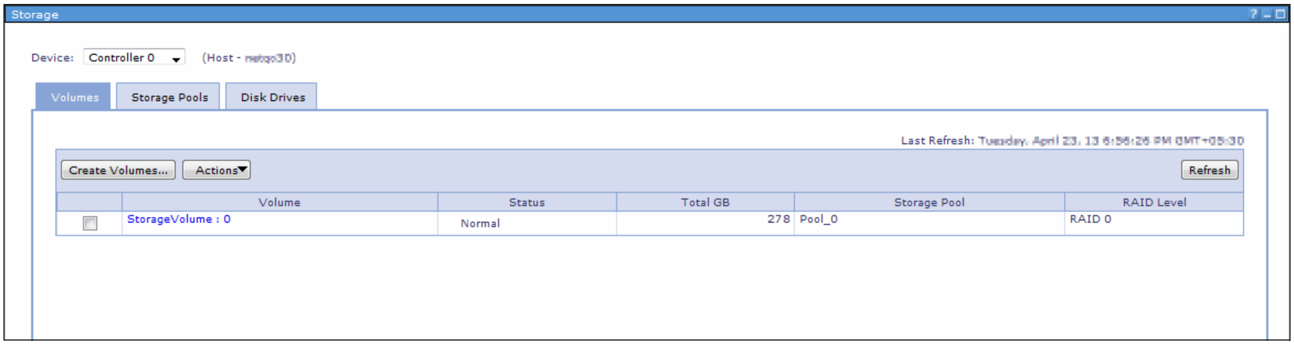

Last Refresh: Tuesday, April 23, 13 6:56:26 PM GMT+05:30

	Volume	Status	Total GB	Storage Pool	RAID Level
☐	StorageVolume : 0	Normal	278	Pool_0	RAID 0

Figure 11-14 Monitor and Management panel for IRC-managed storage devices

> **Note:** The Volumes and Storage Pools pages do not support all the Action menu choices. The Volumes page only supports Create Volumes and Delete Volumes. The Storage Pools page only supports Create Storage Pool and Delete Storage Pool. All other Action menu choices for those pages are not supported.

Removing managed devices

If the storage devices are managed through an SMI-S Provider Agent, it is always recommended to remove the SMI-S Provider Agent followed by removal of the managed storage devices. For more information about removal of managed objects from IBM Systems Director server, see the following link:

http://pic.dhe.ibm.com/infocenter/director/pubs/index.jsp?topic=%2Fcom.ibm.director.console.helps.doc%2Ffqm0_t_rn_removing_a_resource.html

11.3 Advanced edition: IBM Systems Director Storage Control

IBM Systems Director Storage Control is the advanced plug-in for IBM Systems Director server, which provides the users with flexibility to monitor and manage the enterprise-level storage devices from the same console that is used for server and network device monitoring.

This section outlines the tasks that are supported in IBM Systems Director Storage Control along with best practices, tips, and troubleshooting topics.

Need for Storage Control plug-in

Below are the scenarios depending on which the requirement to buy IBM Systems Director Storage Control plug-in can be decided:

- Storage Control plug-in is required if your environment has storage devices or the Fibre Channel (FC) switches are not supported by the base IBM Systems Director Storage Manager and there is no Tivoli Storage Productivity Center management server in your environment.
- When you already have a Tivoli Storage Productivity Center management server in your environment, a Storage Control plug-in is not a mandatory requirement. IBM Systems Director can communicate with an external Tivoli Storage Productivity Center to manage the enterprise storage devices and the supported FC switches.
- As of IBM Systems Director server version 6.3.2, if your environment has Cisco MDS 9100, 9200, and 9500 family switches, it can be managed only by an external Tivoli Storage Productivity Center server. IBM Systems Director server can communicate with external Tivoli Storage Productivity Center to manage these switches.
- When using VMControl plug-in, IBM Systems Director manages and communicates with the virtual SCSI storage devices through the SMI-S providers, the HMC, and optionally, IBM Tivoli Storage Productivity Center. For N-Port ID Virtualization (NPIV) storage devices, the devices must be managed through either IBM Systems Director Storage Control or IBM Tivoli Storage Productivity Center. You cannot manage the storage devices by using the SMI-S provider through IBM Systems Director Storage Manager.

11.3.1 Planning for Storage Control

Listed below are the requirements for using IBM Systems Director Storage Control:

- The virtual/physical server that is chosen for the installation of IBM Systems Director server should have at least the minimum hardware configuration. See Table 11-4 on page 470.

Table 11-4 Hardware requirements

Storage Control minimum hardware requirements		
Processor	Installed memory	Disk storage
3 GHz	► 3 GB (Total requirement for Systems Director, DB2, and Storage Control) ► 0.8 GB RAM (Requirement for Storage Control only)	5.15 GB For AIX or Linux installations, the 5.15 GB must be in these directories: ► /tmp - 2.15 GB ► /opt - 2.6 GB ► /home - 400 MB ► /etc - 10 KB

- The Storage Control plug-in supports most of the Red Hat Enterprise Linux, Microsoft Windows, and AIX operating systems that Systems Director also supports, with these exceptions: Red Hat Enterprise Linux on Power Systems, Red Hat Enterprise Linux on System z systems, and SUSE Linux Enterprise.
- IBM Systems Director 6.2.1 or higher is required in order to use Storage Control.
- As a best practice, plan to use a local IBM DB2 database that is managed by Systems Director server. If not, you must use one of these versions of IBM DB2 as the local database application for Systems Director to use Storage Control:
  - IBM DB2 Enterprise Edition v. 9.7
  - IBM DB2 Enterprise Edition v. 9.7 with fix pack 4

For more information about the IBM Systems Director Storage Control requirements, see the following site:

http://pic.dhe.ibm.com/infocenter/director/pubs/index.jsp?topic=%2Fcom.ibm.director.storagectrl.helps.doc%2Ffqm0_c_sc_planning.html

11.3.2 Installation of Storage Control

IBM Systems Director Storage Control is a separately downloadable and installable plug-in that is available on the following platforms:

- Red Hat Enterprise Linux (RHEL) running on IBM x86 servers
- Microsoft Windows running on IBM x86 servers
- AIX running on IBM Power Servers

Before installing Storage Control, ensure that the supported version of the managed DB2 server is running on IBM Systems Director server:

- On Linux/AIX operating system:

  Command to show the version, fix pack, and installation information of the managed DB2 installed:

  ```
  <IBM Systems Director installation directory>/director/db2/adm/db2level
  ```

- On Microsoft Windows operating system:

  Command to show the version, fix pack, and installation information of the managed DB2 installed:

  ```
  <IBM Systems Director installation directory>\director\db2\adm\db2level
  ```

If the managed DB2 version is other than the supported version mentioned in 11.3.1, “Planning for Storage Control” on page 469, ensure that the supported managed DB2 version is installed before moving ahead with IBM Systems Director Storage Control plug-in installation. The following site provides the instruction for performing fix pack installation for DB2 version 9.7:

http://pic.dhe.ibm.com/infocenter/db2luw/v9r7/index.jsp?topic=%2Fcom.ibm.db2.luw.qb.server.doc%2Fdoc%2Fc0025016.html

IBM Systems Director Storage Control 4.2.1.1 must be installed before you apply updates, including Storage Control 4.2.3.1.

You can use IBM Systems Director update manager to acquire and install IBM Systems Director Storage Control versions that are compatible with IBM Systems Director 6.3.1 and 6.3.2. Table 11-5 lists the installation log file locations for the IBM Systems Director Storage Control plug-in.

Table 11-5 Installation log location

Operating system platform	Installation log file location
IBM AIX/RHEL	► /opt/IBM/TPC/TPC.log ► /var/log/SCinst.log ► /opt/IBM/TPC/log/install
Microsoft Windows	► C:\Program Files\IBM\TPC\TPC.log ► C:\Program Files\IBM\TPC\log\install

For more information about installing Storage Control 4.2.1.1, see the following information center link:

http://pic.dhe.ibm.com/infocenter/director/pubs/index.jsp?topic=%2Fcom.ibm.director.storagectrl.helps.doc%2Ffqm0_t_sc_installing_storage_control.html

After you install IBM Systems Director Storage Control, activate the plug-in by using the following steps:

1. Run the following command to get the advanced plug-in status:

```
# smcli lsmgrs
Network Control : Activated
Active Energy Manager : Deactivated
VMControl : Activated
Storage Control : Deactivated
```

2. If you see Storage Control in the deactivated state, run the following command to activate:

```
#smcli activatemgrs “Storage Control”
```

3. When Storage Control is activated, run the following commands to restart IBM Systems Director server:

```
On Linux/AIX operating system,
<IBM Systems Director installation location>/bin/smstop
<IBM Systems Director installation location>/bin/smstart

On Microsoft Windows operating system,
net stop dirserver
net start dirserver
```

When there are Storage Control *software quality* or *feature function* updates, an update is made available.

Storage Control updates must be applied to an existing Storage Control installation. The procedure for updating Storage Control depends on the version that is installed. To determine the version of Storage Control that is installed, open the Plug-ins tab of the IBM Systems Director Home page and scroll to the Storage Management section. The Storage Control version is listed in the Storage Management section. For steps to update storage control, see the following site:

http://pic.dhe.ibm.com/infocenter/director/pubs/index.jsp?topic=%2Fcom.ibm.director.storagectrl.helps.doc%2Ffqm0_t_sc_updating_storage_control.html

11.3.3 IBM Systems Director Storage Control lifecycle operations

There could be situations where users need to stop, start, or rediscover the Storage Control Farm object which could be a part of troubleshooting issues as well. Listed below are a few of those instances:

- After performing an upgrade from IBM Systems Director 6.2 to Systems Director 6.3.2, the access type for the remote service access point (RSAP) of a switch or storage device does not have the correct value of Tivoli Storage Productivity Center.
- If you changed the password that IBM Systems Director uses to access IBM DB2 and Storage Control can no longer access IBM DB2.
- If you changed the IP address of the management server and you can no longer communicate with Storage Control.
- After you switch to using the managed IBM DB2 database for IBM Systems Director.

Stop Storage Control plug-in

Stop Storage Control by running the appropriate command on the management server, where DIRHOME is the root directory of your IBM Systems Director installation:

On Windows: From a command prompt, run the following command:

```
DIRHOME\StorageControl\bin\stopStorageControl.bat
```

On AIX or Linux: From the command line, run the following command:

```
DIRHOME/StorageControl/bin/stopStorageControl.sh
```

Start Storage Control plug-in

Start Storage Control by running the appropriate command on the management server, where `DIRHOME` is the root directory of your IBM Systems Director installation:

On Windows: From a command prompt, run the following command:

```
DIRHOME\StorageControl\bin\startStorageControl.bat
```

On AIX or Linux: From the command line, run the following command:

```
DIRHOME/StorageControl/bin/startStorageControl.sh
```

Remove and rediscovering the Storage Control Farm object

Removing Storage Control Farm object is similar to removing any other resource from the IBM Systems Director database.

For more information about removing a resource, see the following link:

http://pic.dhe.ibm.com/infocenter/director/pubs/index.jsp?topic=%2Fcom.ibm.director.console.helps.doc%2Ffqm0_t_rn_removing_a_resource.html

Follow the steps below to rediscover and request access to the Storage Control Farm.

For Windows, any of the following two methods can be used:

- Create and run a discovery profile to discover and request access to the farm. When you create the profile, be sure to specify the following settings:
  - Tivoli Storage Productivity Center database configuration:
    - IP Address: The management server address
    - Port: 50010
    - Database Name: TPCDB
    - Database User ID: The Windows administrator user name
    - Database Password: The password for the specified Windows administrator user name
  - Tivoli Storage Productivity Center server configuration:
    - TPC User ID: The Windows administrator user name
    - TPC Password: The password for the specified Windows administrator user name
  - TPC Storage Resource Group Selection: Select **ALL**
  - Access Request: Specify the Windows administrator user name and password
- From the Director Server command line, run the following command, where `DIRHOME` is the path of the directory of your IBM Systems Director installation:

  ```
  DIRHOME\StorageControl\bin\SCDiscoverUnlock.sh
  ```

On AIX or Linux: On the management server from a command line, run the following command, where `DIRHOME` is the root directory of your IBM Systems Director installation:

```
DIRHOME/StorageControl/bin/SCDiscoverUnlock.sh
```

Note: This command discovers and automatically requests access to the Storage Control Farm.

11.3.4 Discovering storage devices

To start managing storage devices with Storage Control, the storage device should be discovered in the IBM Systems Director server as a first step. Below are the type of storage devices that you can discover and manage from IBM Systems Director Control:

- FC switch
- IBM Systems Storage: includes DS4000/3000/5000 series
- IBM DS8000 series
- XIV storage systems
- IBM SAN Volume Controller
- IBM Flex System V7000 and Storwize V7000

Starting with IBM Systems Director Storage Control 4.2.3, use the Discover Storage task that is present on the IBM Systems Director console. Using Discover Storage makes it easy and straightforward to discover the data sources for your storage devices and subsystems so that Storage Control can communicate with them. You also can automatically collect inventory on the discovered devices.

Note: The term *data source* describes how the storage device is managed. Some data sources are the location of the SMI-S provider for the device. For example, the data source for a Brocade Fibre Channel switch is the location of the SMI-S provider for the switch. Other data sources are the device itself. For example, for the current firmware levels of the SAN Volume Controller and IBM Storwize V7000, the device is the data source.

Before starting to work with Storage Control, run the following command to ensure that the Storage Control plug-in is activated:

```
# smcli lsmgrs
Network Control : Activated
Active Energy Manager : Deactivated
VMControl : Activated
Storage Control : Activated
```

If Storage Control is in the deactivated state, follow the steps that are listed in section 11.3.2, "Installation of Storage Control" on page 470 to activate the IBM Systems Director Storage Control.

The following sections provide demonstrations about how to discover storage devices in IBM Systems Director Storage Control, using IBM Storwize V7000 and the Brocade 16 GB Fibre Channel switch as examples.

Discovering an SVC storage device in IBM Systems Director

To discover an IBM Storwize V7000 SVC node, perform the following steps:

1. Ensure that you have a Secure Shell (SSH) key pair generated and associated with an IBM Storwize V7000 user. Following are the steps to generate an SSH key pair and associating it with a user on v7000:

   a. On any host, perform the following step to set up an RSA key pair:

      - On an AIX or Linux host, create an RSA key pair by issuing a command on the host that is similar to the following command. Issue the command from the `$HOME/.ssh` directory:

        ```
        #ssh-keygen -t rsa
        ```

        This process generates two user named files. The files are named <key name> and <key name>.pub. Where "key name" is the name of the private key and 'keyname'.pub is the name of the public key.

      - If you need to perform the preceding step from a server running Microsoft Windows operating system, tools such as *puttygen* can be used.

   b. Download the public key that is generated to the workstation from where you are accessing Storwize V7000. Associate the public key that is generated in the previous step with a user on the Storwize V7000 or Storwize V7000 system, using the following steps:

      i. Log in to the Storwize V7000 web console.

      ii. Go to **Access** → **Users**. Click **New User**, as shown in Figure 11-15 on page 475, to create a new user.

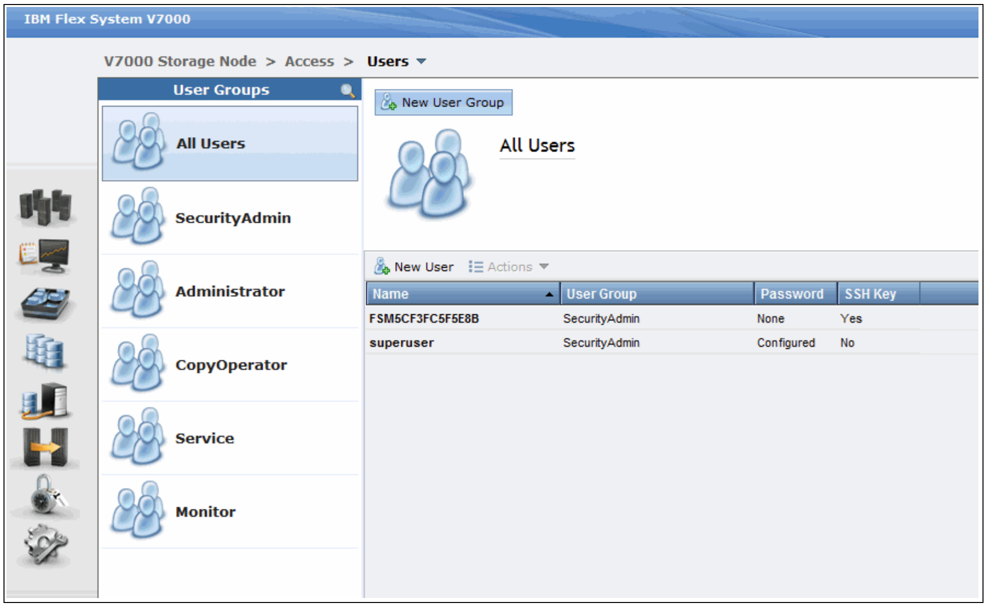

Name	User Group	Password	SSH Key
FSM5CF3FC5F5E8B	SecurityAdmin	None	Yes
superuser	SecurityAdmin	Configured	No

Figure 11-15 IBM Storwize V7000 user management console

iii. Complete the information about the New User panel and upload the public key that is generated in step (a). When done, the new user is displayed under the All Users section, as shown in Figure 11-16.

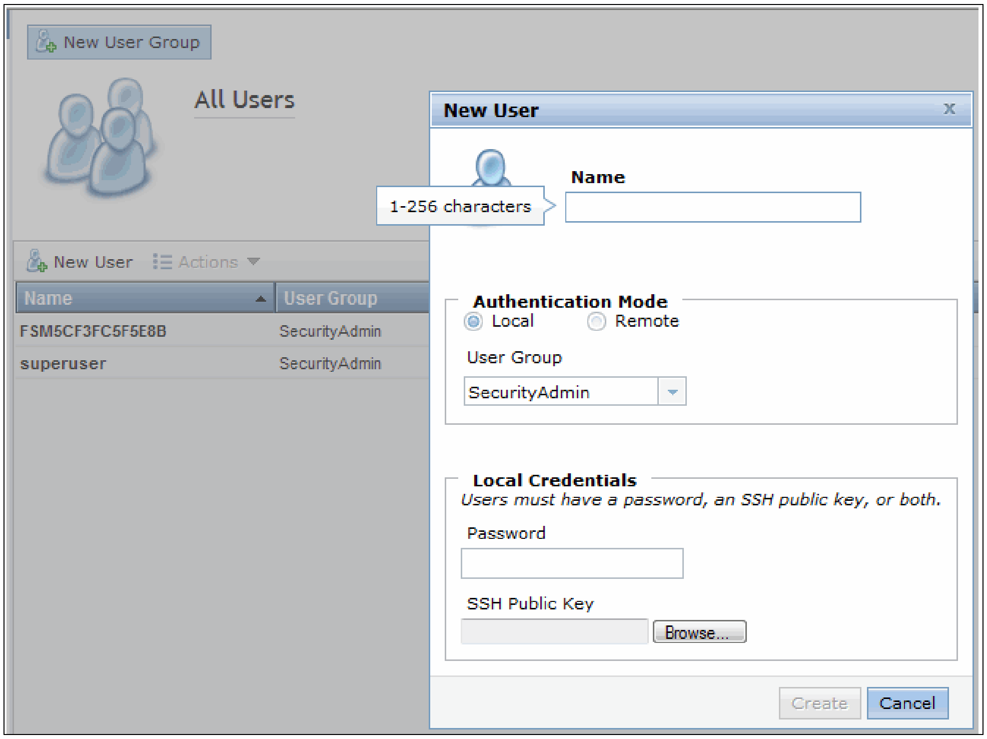

Name	User Group
FSM5CF3FC5F5E8B	SecurityAdmin
superuser	SecurityAdmin

Name

Authentication Mode
● Local ○ Remote

User Group: SecurityAdmin

Local Credentials
Users must have a password, an SSH public key, or both.

Password

SSH Public Key

Figure 11-16 IBM Storwize V7000 add user panel

2. Copy the generated SSH private key to IBM Systems Director server or the workstation where you will access the IBM Systems Director web console depending on the method, either graphical user interface (GUI) or command line, that you are using for discovery of the storage device. Both methods are explained below.

Discovery using the graphical user interface

Perform the following steps to discover by using the GUI:

1. Navigate to the Home window of the IBM Systems Director server console. Click **Discover Storage** under the Storage Management section, as shown in Figure 11-17.

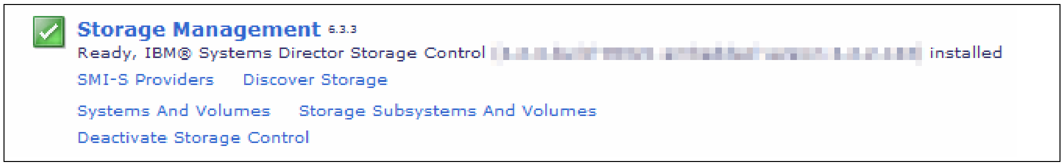

Storage Management 6.3.3
Ready, IBM® Systems Director Storage Control [redacted] installed

Figure 11-17 IBM Systems Director: Discover storage

2. In the Discover Storage page, select **IBM Flex System V7000 and Storwize V7000** for the storage device type.
3. Specific settings for discovering “IBM Flex System V7000 and Storwize V7000” are displayed, as shown in Figure 11-18.

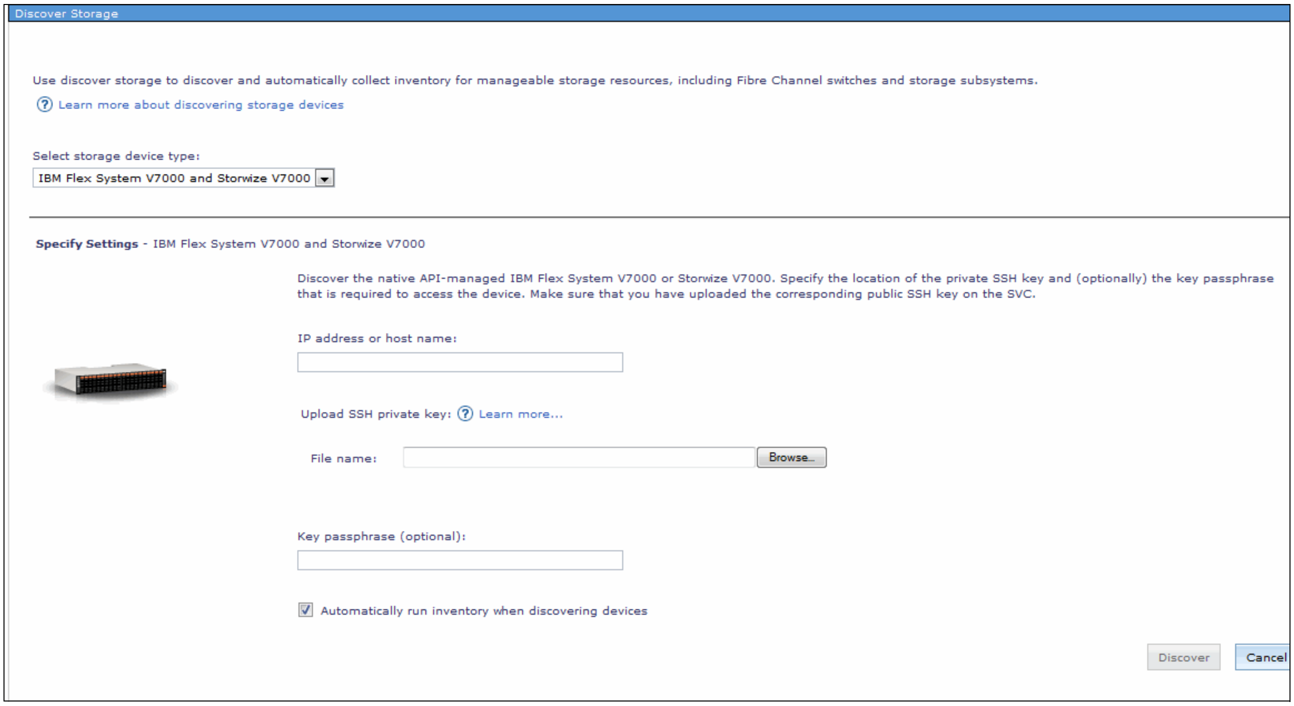

Discover Storage

Use discover storage to discover and automatically collect inventory for manageable storage resources, including Fibre Channel switches and storage subsystems.

Select storage device type: IBM Flex System V7000 and Storwize V7000

Specify Settings - IBM Flex System V7000 and Storwize V7000

Discover the native API-managed IBM Flex System V7000 or Storwize V7000. Specify the location of the private SSH key and (optionally) the key passphrase that is required to access the device. Make sure that you have uploaded the corresponding public SSH key on the SVC.

IP address or host name:

Upload SSH private key: Learn more...

File name:

Key passphrase (optional):

☑ Automatically run inventory when discovering devices

Figure 11-18 Discover Storage page

Where:

a. “IP address or the host name” field: Enter the IP address of IBM Storwize v7000.

b. “Upload SSH private key” field: Specify the location where you have the private SSH key located on the machine from where you are accessing the IBM Systems Director server console.

c. “Key passphrase (optional)” field: Enter the SSH key passphrase, if any.

d. “Automatically run inventory when discovering devices” field: This check box is selected by default. It is recommended to collect inventory on the Storage Farm when you add new storage devices.

e. Click **Discover** to start the discovery job.

4. When the discovery of the IBM Storwize V7000 is done, inventory collection starts as shown in Figure 11-19. Click **Close** to continue.

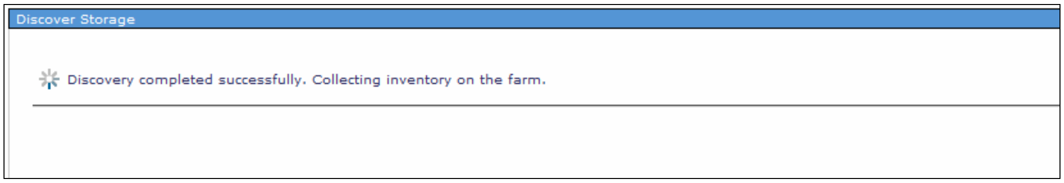

Discovery completed successfully. Collecting inventory on the farm.

Figure 11-19 Discover Storage page: Completed discovery of the storage device

5. Both discovery and inventory collection can be seen on the “Active and Scheduled Jobs” list, as shown in Figure 11-20.

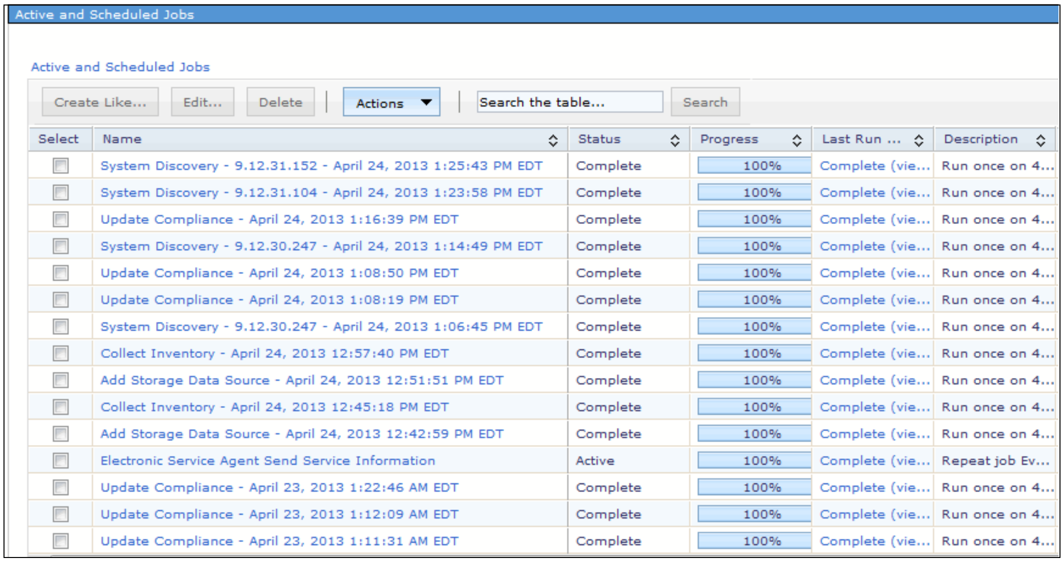

Select	Name	Status	Progress	Last Run ...	Description
□	System Discovery - 9.12.31.152 - April 24, 2013 1:25:43 PM EDT	Complete	100%	Complete (vie...	Run once on 4...
□	System Discovery - 9.12.31.104 - April 24, 2013 1:23:58 PM EDT	Complete	100%	Complete (vie...	Run once on 4...
□	Update Compliance - April 24, 2013 1:16:39 PM EDT	Complete	100%	Complete (vie...	Run once on 4...
□	System Discovery - 9.12.30.247 - April 24, 2013 1:14:49 PM EDT	Complete	100%	Complete (vie...	Run once on 4...
□	Update Compliance - April 24, 2013 1:08:50 PM EDT	Complete	100%	Complete (vie...	Run once on 4...
□	Update Compliance - April 24, 2013 1:08:19 PM EDT	Complete	100%	Complete (vie...	Run once on 4...
□	System Discovery - 9.12.30.247 - April 24, 2013 1:06:45 PM EDT	Complete	100%	Complete (vie...	Run once on 4...
□	Collect Inventory - April 24, 2013 12:57:40 PM EDT	Complete	100%	Complete (vie...	Run once on 4...
□	Add Storage Data Source - April 24, 2013 12:51:51 PM EDT	Complete	100%	Complete (vie...	Run once on 4...
□	Collect Inventory - April 24, 2013 12:45:18 PM EDT	Complete	100%	Complete (vie...	Run once on 4...
□	Add Storage Data Source - April 24, 2013 12:42:59 PM EDT	Complete	100%	Complete (vie...	Run once on 4...
□	Electronic Service Agent Send Service Information	Active	100%	Complete (vie...	Repeat job Ev...
□	Update Compliance - April 23, 2013 1:22:46 AM EDT	Complete	100%	Complete (vie...	Run once on 4...
□	Update Compliance - April 23, 2013 1:12:09 AM EDT	Complete	100%	Complete (vie...	Run once on 4...
□	Update Compliance - April 23, 2013 1:11:31 AM EDT	Complete	100%	Complete (vie...	Run once on 4...

Figure 11-20 Active and Scheduled Jobs page

6. When both discovery and inventory jobs are completed, the discovered storage device is displayed under the Discover Storage page, as shown in Figure 11-21.

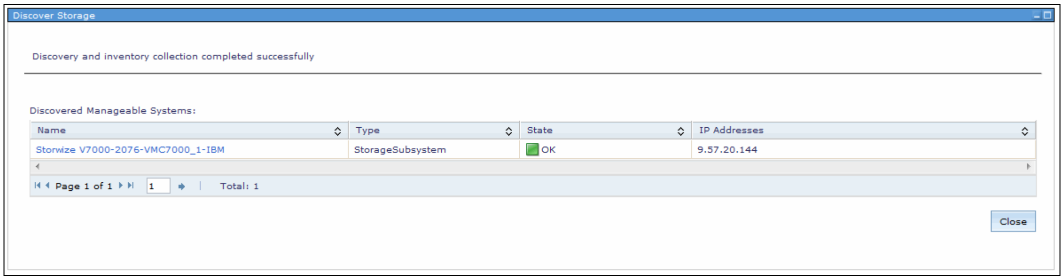

Discovery and inventory collection completed successfully

Discovered Manageable Systems:

Name	Type	State	IP Addresses
Storwize V7000-2076-VMC7000_1-IBM	StorageSubsystem	🟩 OK	9.57.20.144

Figure 11-21 Discovered IBM Storwize V7000

7. To further manage and monitor the IBM Storwize V7000, go to **Resource Explorer** → **Groups** → **All Storage Systems (View Members)**. The discovered V7000 is displayed along with additional information, such as Raw Capacity, Usable Capacity, Available Capacity, as shown in Figure 11-22.

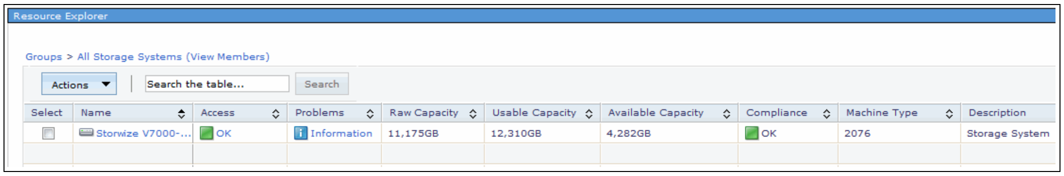
Resource Explorer

Groups > All Storage Systems (View Members)

Select	Name	Access	Problems	Raw Capacity	Usable Capacity	Available Capacity	Compliance	Machine Type	Description
□	Storwize V7000-...	🟩 OK	Information	11,175GB	12,310GB	4,282GB	🟩 OK	2076	Storage System

Figure 11-22 Resource Explorer panel

Discovery using the command line

Perform the following steps to discover by using the command line:

1. Log in to the IBM Systems Director server command line.
2. Run the following command to discover the IBM Storwize V7000:

```
smcli mkdatasource -c svc -i <IP address of the v7000> -f <path to generated
private SSH Key on systems director server> -v V7000
```

When the v7000 discovery is done, inventory collection is done automatically on the Storage Control Farm object.

11.3.5 Discovering Fibre Channel fabric in IBM Systems Director

Discovery of the Fibre Channel (FC) switch is supported in both GUI and command line, similar to discovery of storage devices. Both methods are documented below.

Discovery using the GUI

To discover a Brocade FC switch, perform the following steps:

1. On a dedicated physical/virtual server, install and configure the “Brocade SMI Agent” or “Brocade Network Advisor (with Integrated SMI Agent)” with the Brocade FC switch. For more information about installation and configuration, see the following link:

   http://www.brocade.com/services-support/drivers-downloads/smi-agent/index.page

2. Navigate to the Home page of the IBM Systems Director server console. Click **Discover Storage** under the Storage Management section, as shown in Figure 11-23.

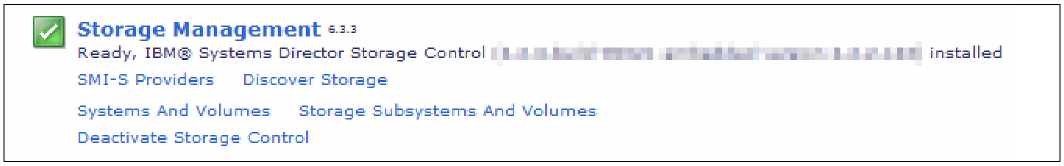
Storage Management 6.3.3

Ready, IBM® Systems Director Storage Control [redacted] installed

Figure 11-23 IBM Systems Director: Discover Storage page

3. In the Discover Storage page, select **FC switch** for the storage device type.
4. Specific settings for discovering FC switch are displayed, as shown in Figure 11-24 on page 479.

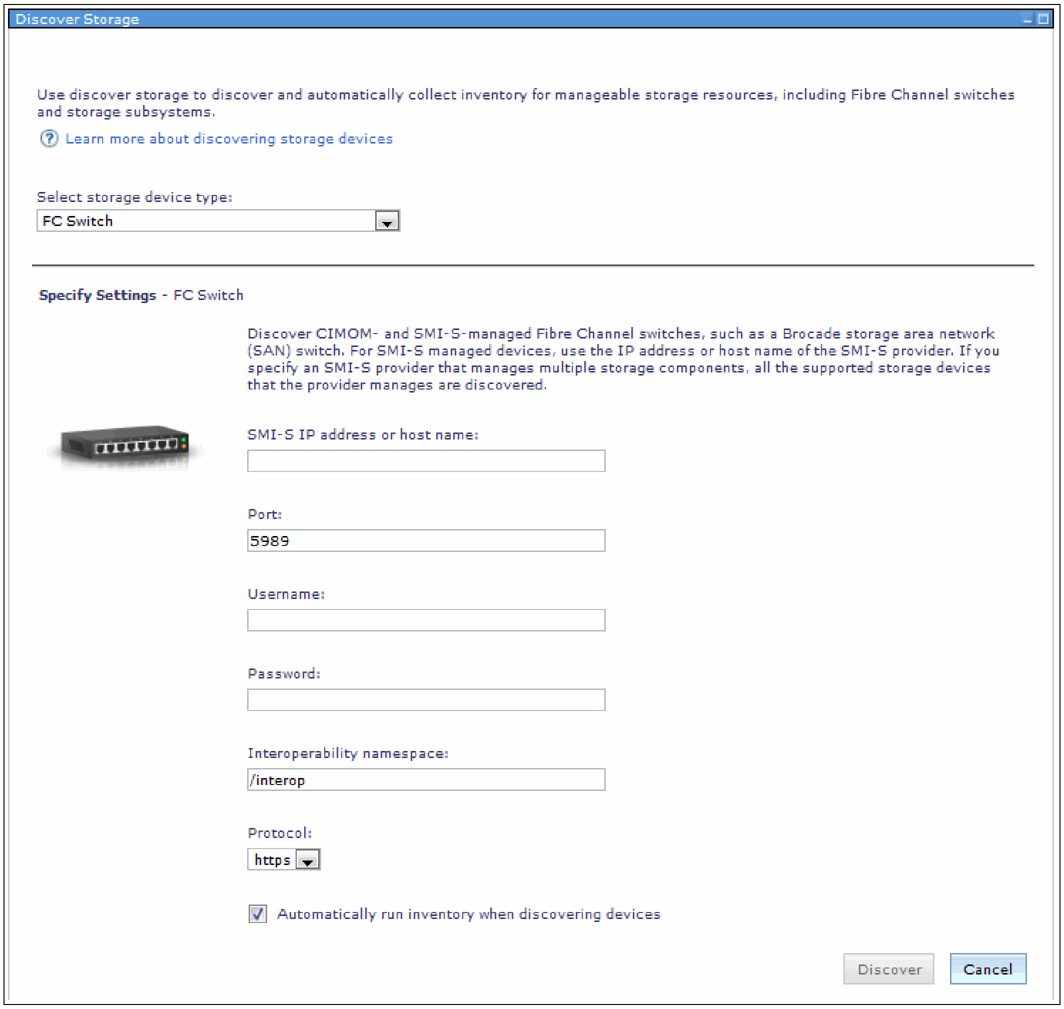
Discover Storage

Use discover storage to discover and automatically collect inventory for manageable storage resources, including Fibre Channel switches and storage subsystems.

Learn more about discovering storage devices

Select storage device type: FC Switch

Specify Settings - FC Switch

Discover CIMOM- and SMI-S-managed Fibre Channel switches, such as a Brocade storage area network (SAN) switch. For SMI-S managed devices, use the IP address or host name of the SMI-S provider. If you specify an SMI-S provider that manages multiple storage components, all the supported storage devices that the provider manages are discovered.

SMI-S IP address or host name:

Port: 5989

Username:

Password:

Interoperability namespace: /interop

Protocol: https

☑ Automatically run inventory when discovering devices

Figure 11-24 Discover storage: FC switch

Where:

i. "SMI-S IP address or host name" field: Enter the IP address of the server on which SMI-S Provider Agent has been installed.

ii. "Port" field: Leave the default value if you have not chosen a different port while installing the SMI-S Provider Agent. If there is a different port selected, enter that port number in this text box.

iii. "Username" and "Password" fields: Enter the user name and password of the server on which SMI-S Provider Agent has been installed.

iv. "Interoperability namespace" field: Leave the default value if you have not chosen a different name for the namespace while installing the SMI-S Provider Agent. If there is a different name, enter that name in this text box.

v. "Protocol" field: Choose the protocol using which IBM Systems Director Storage Control communicates with the SMI-S Provider Agent.

vi. "Automatically run inventory when discovering devices" option: This check box is selected by default. It is recommended to collect inventory on the Storage Farm when you add new storage devices.

Click **Discover** to start the discovery job.

vii. When the discovery of the FC switch is done, inventory collection is started. Both the discovery and the inventory collection tasks can be seen on the “Active and Scheduled Jobs” list of IBM Systems Director server.

viii.When both discovery and inventory jobs are completed, the discovered FC switches managed by the SMI-S provider are displayed under the Discover Storage page, as shown in Figure 11-25.

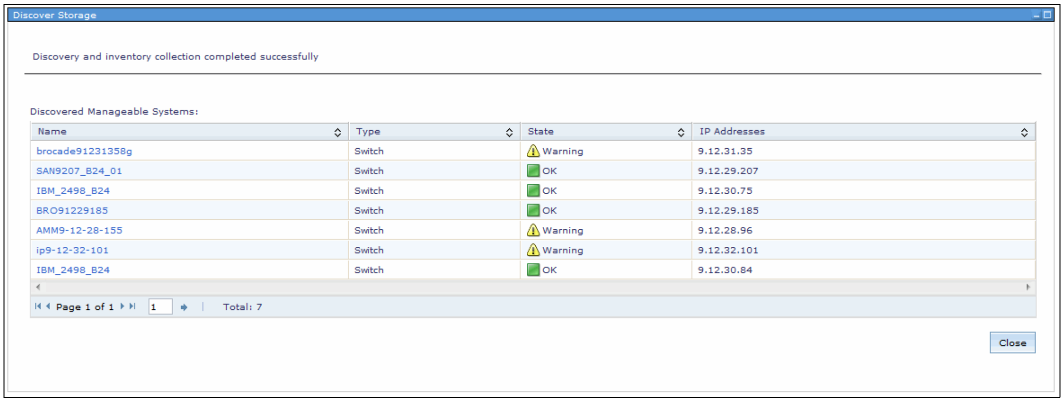

Discovery and inventory collection completed successfully

Discovered Manageable Systems:

Name	Type	State	IP Addresses
brocade91231358g	Switch	⚠ Warning	9.12.31.35
SAN9207_B24_01	Switch	🟩 OK	9.12.29.207
IBM_2498_B24	Switch	🟩 OK	9.12.30.75
BRO91229185	Switch	🟩 OK	9.12.29.185
AMM9-12-28-155	Switch	⚠ Warning	9.12.28.96
ip9-12-32-101	Switch	⚠ Warning	9.12.32.101
IBM_2498_B24	Switch	🟩 OK	9.12.30.84

Page 1 of 1 | 1 | Total: 7

Figure 11-25 Discover Storage page: FC switches discovered

Discovery using the command line

From the IBM Systems Director command line, run the following command to discover the FC switch:

```
smcli mkdatasource -c fabric -i <IP address of the SMI-S provider agent> -t
<https/http> -p <https/http port> -u <Username> -w <password> -n <name of the
interoperability namespace>
```

When the discovery FC switch is done, inventory collection is done automatically on the Storage Control Farm object.

For more information about the commands used to discover other types of Storage Control-supported storage devices, see the following link:

http://pic.dhe.ibm.com/infocenter/director/pubs/index.jsp?topic=%2Fcom.ibm.director.cli.helps.doc%2Ffqm0_r_cli_mkdatasource.html

11.3.6 Inventory collection on Storage Control-managed devices

Before using any tasks that are related to IBM Systems Director Storage Control or the dependent plug-ins such as VMControl, inventory collection has to be done on the managed storage devices.

IBM Systems Director Storage Control uses the embedded version of the IBM Tivoli Storage Productivity Center server.

Discovering and running inventory against the Storage Farm that represents the embedded IBM Tivoli Storage Productivity Center server causes Systems Director to fully populate all of the devices that are being managed by the embedded IBM Tivoli Storage Productivity Center, including Fibre Channel switches and storage subsystems.

Note: You can collect inventory for only one device in a Storage Farm at a time. Collecting inventory at the same time for multiple devices that are managed by the embedded IBM Tivoli Storage Productivity Center is not supported.

Below are a couple of the scenarios where a user needs to run inventory collection on the storage devices:

- As a first step after discovering storage devices in IBM Systems Director server.
- When the configuration changes for a storage device, that is managed by IBM Systems Director Storage Control, using external management such as the native device management tools. For IBM Systems Director Storage Control to update its database, an inventory collection on the specific storage device or FC switch should be done.

There is no specific sequence to be followed in an inventory collection operation for FC switches, Storage Control Farm, and storage devices.

11.3.7 IBM Systems Director Storage Control: Allocating storage volumes to virtual servers

When IBM Systems Director Storage Control is working in conjunction with IBM Systems Director VMControl [express/standard/enterprise edition], you manage storage allocation and deallocation on the Power VM/KVM virtual servers.

This book provides an example of allocating and viewing storage volume relationships of existing PowerVM virtual servers to demonstrate the capability of IBM Systems Director Storage Control:

1. As a best practice, ensure that you have collected inventory on the following instances, as applicable to your configuration:

   a. The IBM HMC/IVM instances
   b. The IBM Power Servers instances
   c. VIOS Virtual Server instances

   Note: It is not necessary to collect inventory every time before doing VMControl or Storage Control operations. See section 11.3.6, "Inventory collection on Storage Control-managed devices" on page 480 for more information about when collecting inventory is necessary.

2. Go to **Inventory** → **Views** → **Virtual Servers and Hosts** on the IBM Systems Director server web console. Right-click the virtual server to which you need to allocate the storage disk and select **System Configuration** → **Create Storage Volumes**, as shown in Figure 11-26.

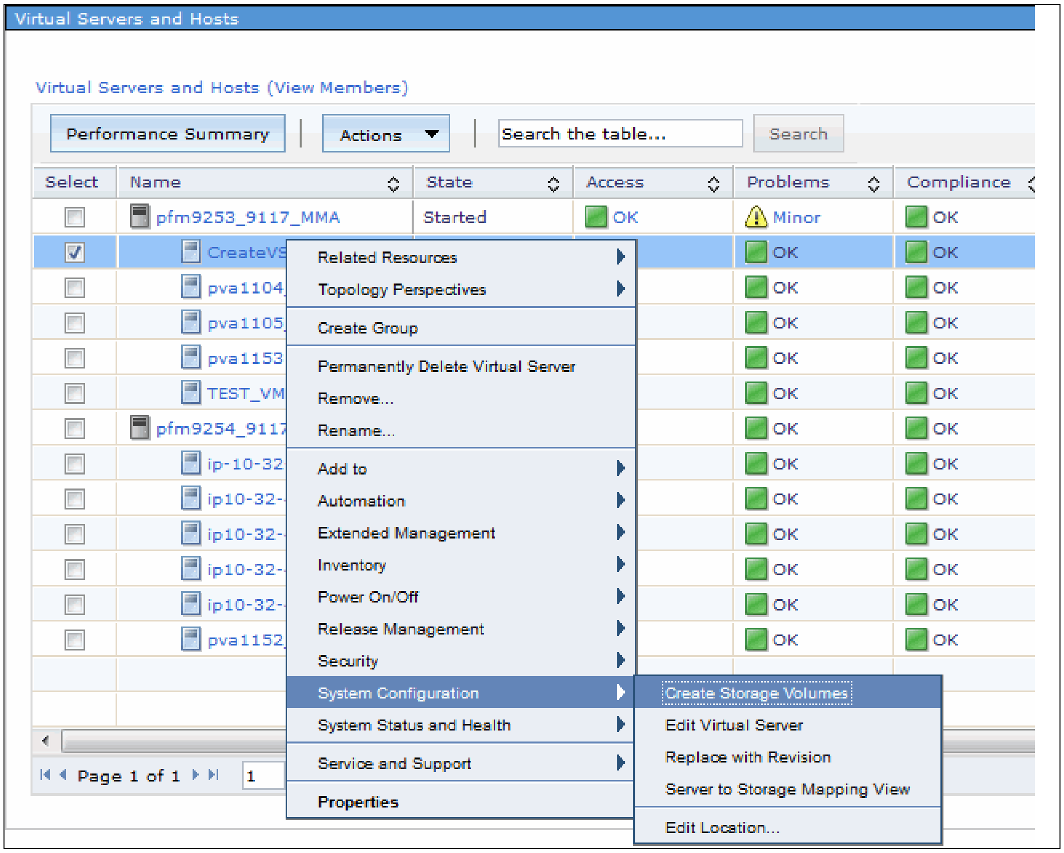

Select	Name	State	Access	Problems	Compliance
☐	pfm9253_9117_MMA	Started	🟩 OK	⚠ Minor	🟩 OK
☑	CreateVS			🟩 OK	🟩 OK
☐	pva1104			🟩 OK	🟩 OK
☐	pva1105			🟩 OK	🟩 OK
☐	pva1153			🟩 OK	🟩 OK
☐	TEST_VM			🟩 OK	🟩 OK
☐	pfm9254_9117			🟩 OK	🟩 OK
☐	ip-10-32			🟩 OK	🟩 OK
☐	ip10-32-			🟩 OK	🟩 OK
☐	ip10-32-			🟩 OK	🟩 OK
☐	ip10-32-			🟩 OK	🟩 OK
☐	ip10-32-			🟩 OK	🟩 OK
☐	pva1152			🟩 OK	🟩 OK

Figure 11-26 Create Storage Volumes option

3. On the Create Storage Volumes wizard, click **Next** to get past the Welcome page.

4. On the Storage Pool page, you see all the storage pools to which the virtual server has access to, as shown in Figure 11-27. Select the appropriate storage pool and click **Next**.

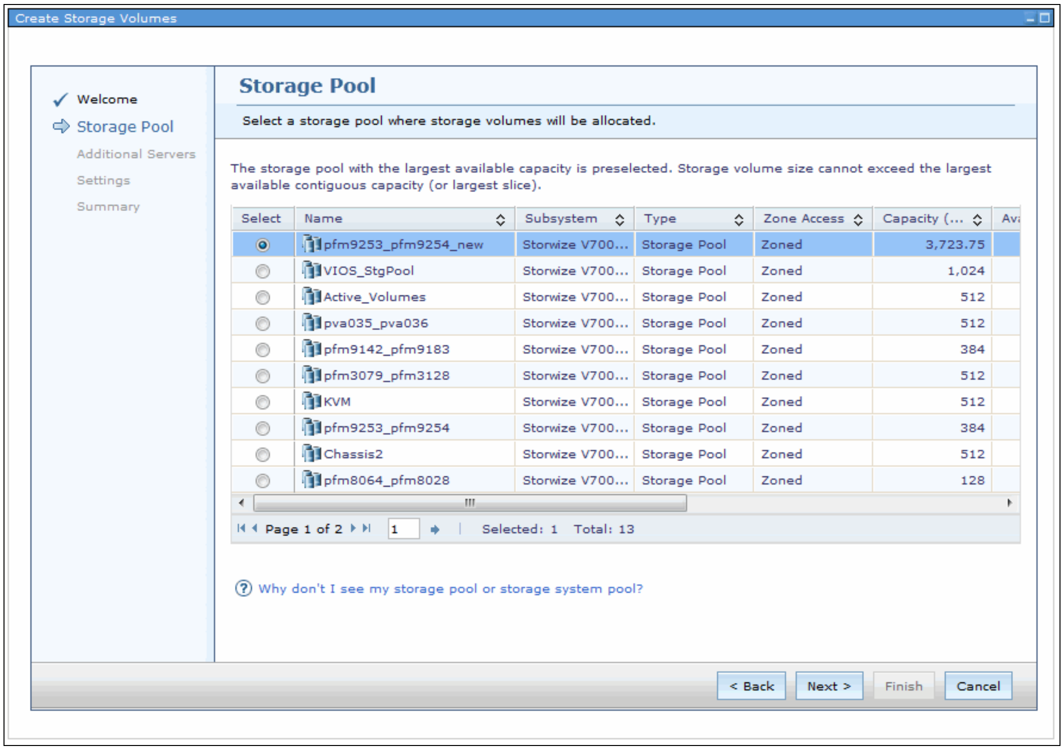

Storage Pool

Select a storage pool where storage volumes will be allocated.

The storage pool with the largest available capacity is preselected. Storage volume size cannot exceed the largest available contiguous capacity (or largest slice).

Select	Name	Subsystem	Type	Zone Access	Capacity (...	Ava
●	pfm9253_pfm9254_new	Storwize V700...	Storage Pool	Zoned	3,723.75	
○	VIOS_StgPool	Storwize V700...	Storage Pool	Zoned	1,024	
○	Active_Volumes	Storwize V700...	Storage Pool	Zoned	512	
○	pva035_pva036	Storwize V700...	Storage Pool	Zoned	512	
○	pfm9142_pfm9183	Storwize V700...	Storage Pool	Zoned	384	
○	pfm3079_pfm3128	Storwize V700...	Storage Pool	Zoned	512	
○	KVM	Storwize V700...	Storage Pool	Zoned	512	
○	pfm9253_pfm9254	Storwize V700...	Storage Pool	Zoned	384	
○	Chassis2	Storwize V700...	Storage Pool	Zoned	512	
○	pfm8064_pfm8028	Storwize V700...	Storage Pool	Zoned	128	

Page 1 of 2 | 1 | Selected: 1 Total: 13

Why don't I see my storage pool or storage system pool?

Figure 11-27 Storage pools that are accessible to the selected virtual server

Note applies only to PowerVM: In this example, we chose a virtual server that has vFC adapters and is configured to use NPIV. When you are trying to assign storage to existing virtual servers, the following disk types are supported to be added from the IBM Systems Director GUI:

- New/existing NPIV disks
- Existing vSCSI disks

Creating a new vSCSI disk and assigning it to an existing virtual server is supported from the command line and REST interface. Following is an example of the **smcli** command that is used to allocate a new vSCSI disk to an existing PowerVM virtual server:

```
# smcli lsvrtcap -c chvs -n <name of virtual server>

Note the storage pool key from the above command output

# smcli chvs -s "assigneddisks=diskname:<new disk name>;disksize:<disk size
in MB>;adddisklocation:storagepools[<Pool Key>]" -n <name of virtual server>
```

5. On the Additional Servers page, you can select any other virtual servers that you want to have access to the storage volume being created using the current task. In this example, we selected virtual server **ip10-32-42-88**, as shown in Figure 11-28. Click **Next** to continue.

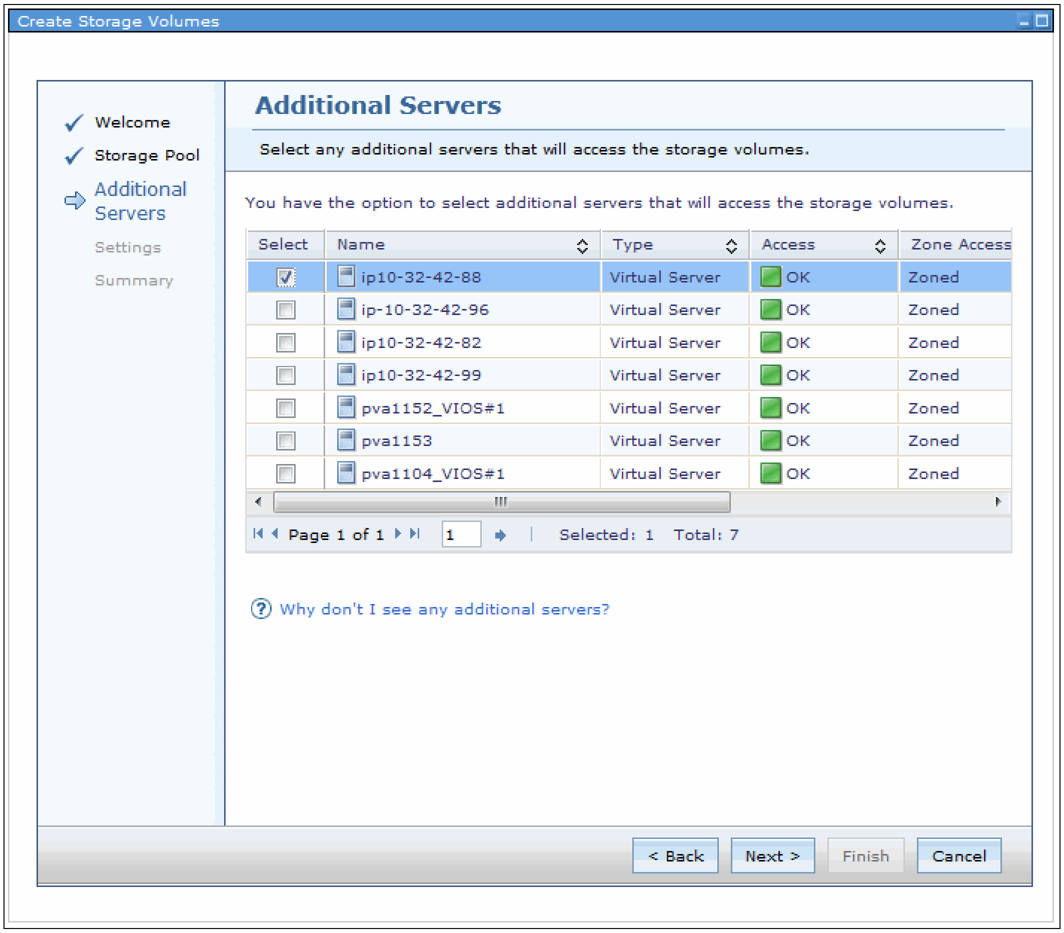

Additional Servers

Select any additional servers that will access the storage volumes.

You have the option to select additional servers that will access the storage volumes.

Select	Name	Type	Access	Zone Access
☑	ip10-32-42-88	Virtual Server	🟩 OK	Zoned
□	ip-10-32-42-96	Virtual Server	🟩 OK	Zoned
□	ip10-32-42-82	Virtual Server	🟩 OK	Zoned
□	ip10-32-42-99	Virtual Server	🟩 OK	Zoned
□	pva1152_VIOS#1	Virtual Server	🟩 OK	Zoned
□	pva1153	Virtual Server	🟩 OK	Zoned
□	pva1104_VIOS#1	Virtual Server	🟩 OK	Zoned

Page 1 of 1 | 1 | Selected: 1 Total: 7

Why don't I see any additional servers?

Figure 11-28 Additional servers that need to have access to the new storage volume

6. On the Settings page, enter the volume name and size details, as shown in Figure 11-29.

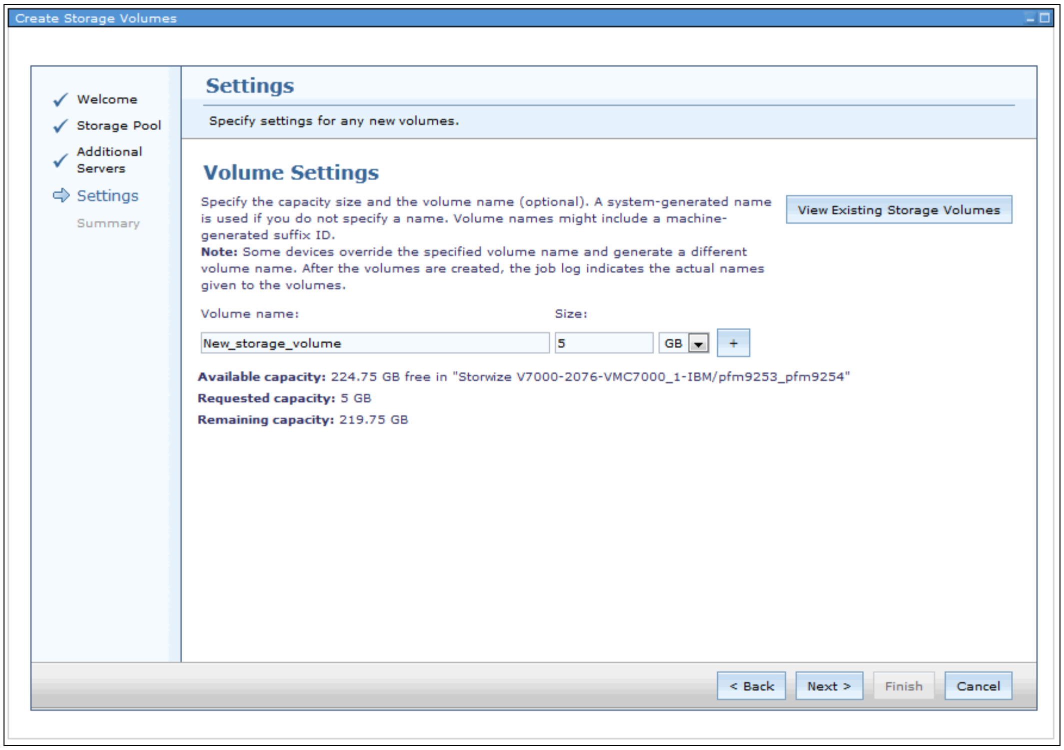

Settings

Specify settings for any new volumes.

Volume Settings

Specify the capacity size and the volume name (optional). A system-generated name is used if you do not specify a name. Volume names might include a machine-generated suffix ID.
Note: Some devices override the specified volume name and generate a different volume name. After the volumes are created, the job log indicates the actual names given to the volumes.

Volume name: New_storage_volume

Size: 5 GB

Available capacity: 224.75 GB free in "Storwize V7000-2076-VMC7000_1-IBM/pfm9253_pfm9254"
Requested capacity: 5 GB
Remaining capacity: 219.75 GB

Figure 11-29 Volume settings

7. If you need to look at the existing storage volumes on the selected storage pool, click **View Existing Storage Volumes**. Existing storage volumes are displayed as shown in Figure 11-30. Click **Close** to exit.

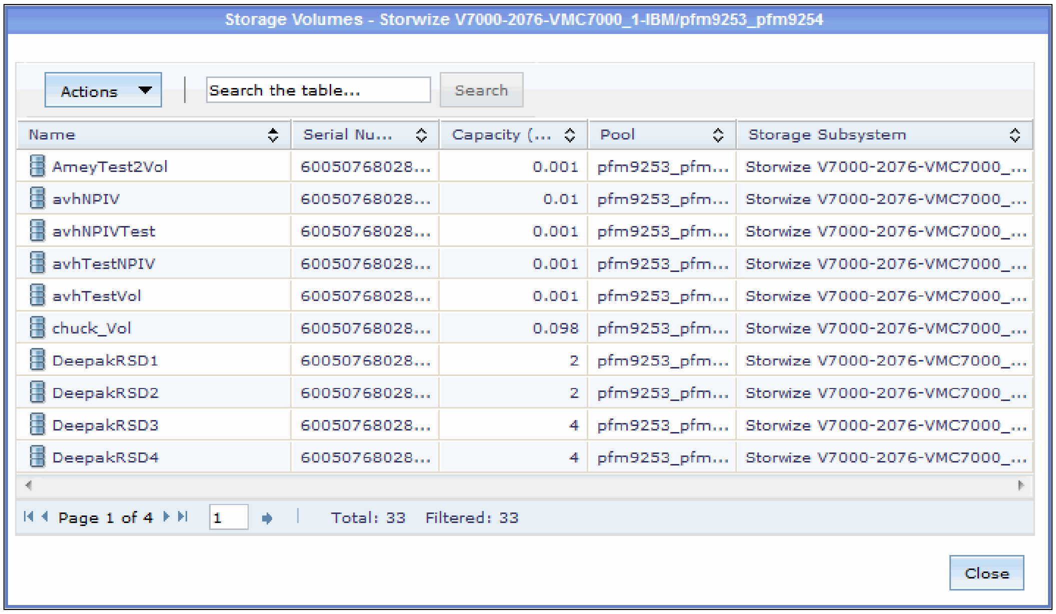

Storage Volumes - Storwize V7000-2076-VMC7000_1-IBM/pfm9253_pfm9254

Name	Serial Nu...	Capacity (...	Pool	Storage Subsystem
AmeyTest2Vol	60050768028...	0.001	pfm9253_pfm...	Storwize V7000-2076-VMC7000_...
avhNPIV	60050768028...	0.01	pfm9253_pfm...	Storwize V7000-2076-VMC7000_...
avhNPIVTest	60050768028...	0.001	pfm9253_pfm...	Storwize V7000-2076-VMC7000_...
avhTestNPIV	60050768028...	0.001	pfm9253_pfm...	Storwize V7000-2076-VMC7000_...
avhTestVol	60050768028...	0.001	pfm9253_pfm...	Storwize V7000-2076-VMC7000_...
chuck_Vol	60050768028...	0.098	pfm9253_pfm...	Storwize V7000-2076-VMC7000_...
DeepakRSD1	60050768028...	2	pfm9253_pfm...	Storwize V7000-2076-VMC7000_...
DeepakRSD2	60050768028...	2	pfm9253_pfm...	Storwize V7000-2076-VMC7000_...
DeepakRSD3	60050768028...	4	pfm9253_pfm...	Storwize V7000-2076-VMC7000_...
DeepakRSD4	60050768028...	4	pfm9253_pfm...	Storwize V7000-2076-VMC7000_...

Page 1 of 4 | 1 | Total: 33 Filtered: 33

Figure 11-30 Existing storage volumes on the storage pool selected

8. Clicking **Next** on the Settings page brings you to the Summary page, as shown in Figure 11-31. Review the settings and click **Finish**.

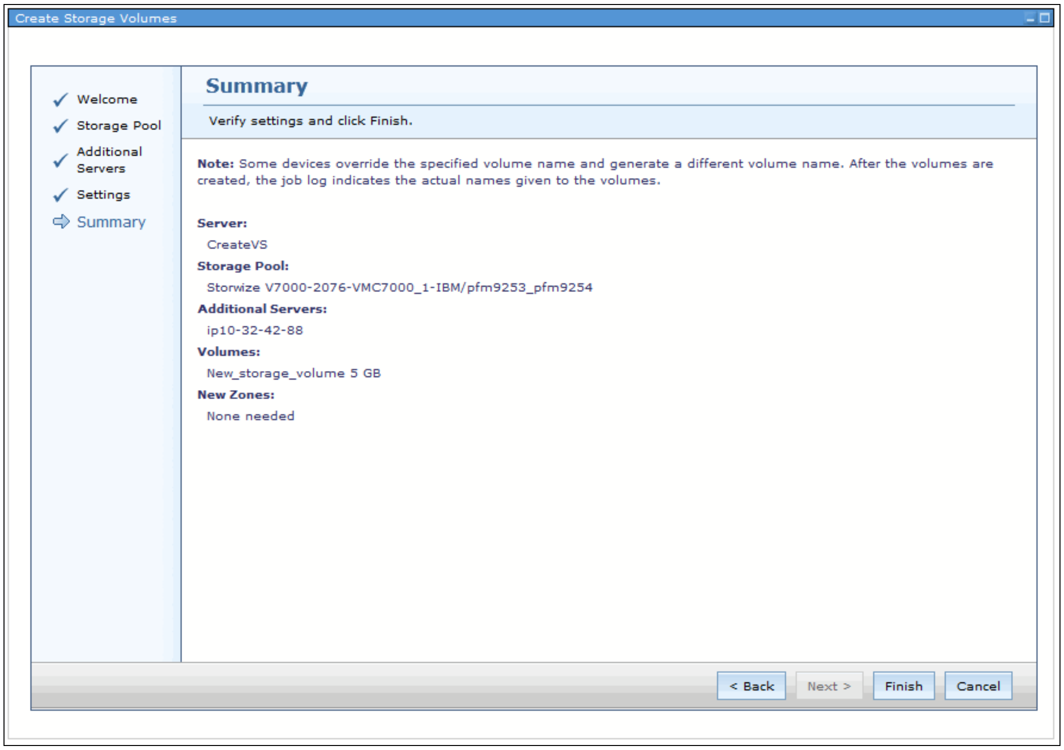

Summary

Verify settings and click Finish.

Note: Some devices override the specified volume name and generate a different volume name. After the volumes are created, the job log indicates the actual names given to the volumes.

Server:
CreateVS
Storage Pool:
Storwize V7000-2076-VMC7000_1-IBM/pfm9253_pfm9254
Additional Servers:
ip10-32-42-88
Volumes:
New_storage_volume 5 GB
New Zones:
None needed

Figure 11-31 Summary page

9. Monitor the task in the Active and Scheduled Jobs list until it is complete.

10. To view the storage volume that was allocated in previous steps, go to **Inventory** → **Views** → **Virtual Servers and Hosts** on the IBM Systems Director server web console. Right-click the virtual server to which you just allocated the storage disk and select **Related Resources** → **Storage Volumes**.

The resulting page shows you the storage volumes that are related to the virtual server. Both “CreateVS” and “ip10-32-42-88”, as shown in Figure 11-32 and Figure 11-33 on page 487, are allocated with the newly created volume.

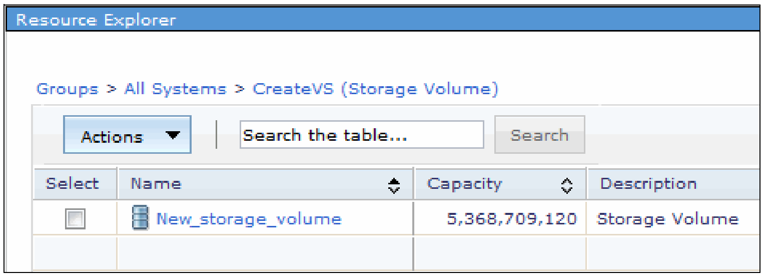

Select	Name	Capacity	Description
□	New_storage_volume	5,368,709,120	Storage Volume

Figure 11-32 CreateVS: Virtual server: storage volume relationship

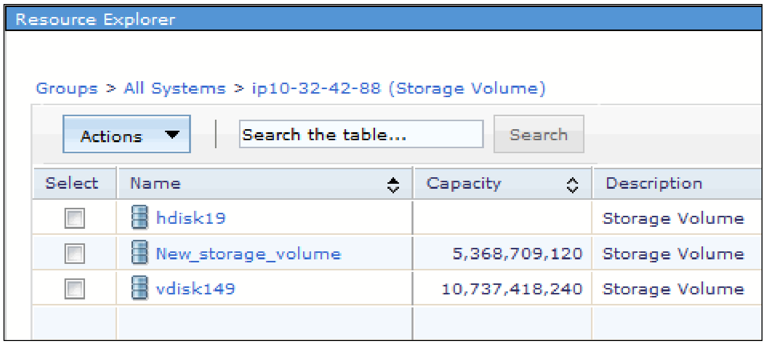

Resource Explorer

Groups > All Systems > ip10-32-42-88 (Storage Volume)

Select	Name	Capacity	Description
□	hdisk19		Storage Volume
□	New_storage_volume	5,368,709,120	Storage Volume
□	vdisk149	10,737,418,240	Storage Volume

Figure 11-33 ip10-32-42-88: Virtual server: storage volume relationship

11. To see the “Server to Storage Mapping View”, go to **Inventory** → **Views** → **Virtual Servers and Hosts** on the IBM Systems Director server web console. Right-click the virtual server to which you allocated the storage disk and select **System Configuration** → **Server to Storage Mapping View**.

    The resulting page shows you an end-to-end storage configuration, as shown in Figure 11-34.

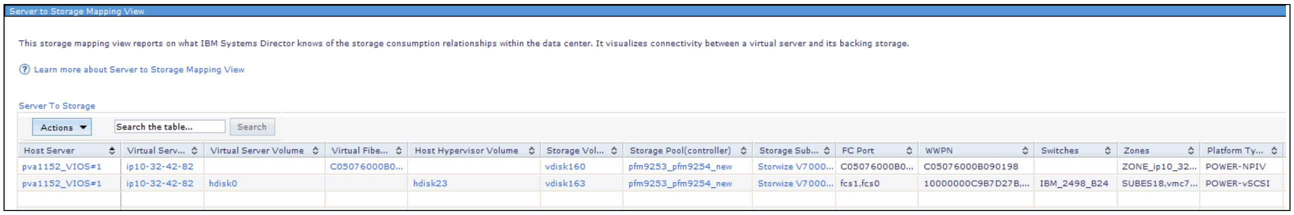

Server to Storage Mapping View

This storage mapping view reports on what IBM Systems Director knows of the storage consumption relationships within the data center. It visualizes connectivity between a virtual server and its backing storage.

Host Server	Virtual Serv...	Virtual Server Volume	Virtual Fibe...	Host Hypervisor Volume	Storage Vol...	Storage Pool(controller)	Storage Sub...	FC Port	WWPN	Switches	Zones	Platform Ty...
pva1152_VIOS#1	ip10-32-42-82		C05076000B0...		vdisk160	pfm9253_pfm9254_new	Storwize V7000...	C05076000B0...	C05076000B090198		ZONE_ip10_32...	POWER-NPIV
pva1152_VIOS#1	ip10-32-42-82	hdisk0		hdisk23	vdisk163	pfm9253_pfm9254_new	Storwize V7000...	fcs1,fcs0	10000000C9B7D27B,...	IBM_2498_B24	SUBES18,vmc7...	POWER-vSCSI

Figure 11-34 Server to Storage Mapping View page

11.3.8 Removing Storage Control-managed devices

When you discover the storage devices that are going to be managed by IBM Systems Director Storage Control, two types of objects are created in the IBM Systems Director database. One is a data source that is defined by IBM Systems Director Storage Control and the other is a corresponding managed endpoint instance that is created by IBM Systems Director server:

1. Running the **`smcli lsdatasource`** command retrieves all the data sources that are defined by IBM Systems Director Storage Control. Below is an example command output that lists all the data sources that are defined by Storage Control in the IBM Systems Director server database:

    ```
    #smcli lsdatasource

    Data Source: 1
    Management IP Addresses: 9.x.x.x
    Category: svc
    User ID:
    Managed Devices OID: [0x7a6c]


    Data Source: 2
    Management IP Addresses: 9.x.x.x
    Category: svc
    User ID:
    ```

```
Managed Devices OID: [0xaab1]


Data Source: 3
Management IP Addresses: 9.x.x.x
Category: fabric
User ID: Administrator
Managed Devices OID: [0x8150, 0x7f92, 0x7fe8, 0x81a4, 0x8068, 0x80f2, 0x821e]
Port: 5989
Protocol: https
Namespace: /interop
```

2. Running the **smcli lsmeps** command retrieves all the managed endpoints that are defined by IBM Systems Director server, which also includes the FC switches and storage devices. Included below is an example command output that lists the extracts pointing to the managed endpoint instances defined by IBM Systems Director server in its database for the storage devices and FC switches. You are able to see the following instances in the IBM Systems Director GUI interface:

```
smcli lsmeps
..
..
MEP: Storwize V7000-2076-VMC7000_1-IBM (resource name: 00000200A0203BD9+0)
        OID: 31340
        GUID: 988DB4C7FA483DB29BF84208B9D2AE3A
        ResourceType StorageSubsystem
MEP: Flex V7000-4939-V7000 Storage Node-IBM (resource name: 0000000020600062+0)
        OID: 43697
        GUID: E617B60819D7307895789421691B499B
        ResourceType StorageSubsystem
..
MEP: IBM_2498_B24 (resource name: IBM_2498_B24)
        OID: 32872
        GUID: A09E353C97E7392287997A1CBC5555B2
        ResourceType Switch
```

When you need to remove the storage devices or FC switches that are managed by IBM Systems Director Storage Control, both of following endpoint types should be removed:

- Farm objects that are defined by IBM Systems Director Storage Control. See the following link for more information about the removal of farm objects:

  http://pic.dhe.ibm.com/infocenter/director/pubs/index.jsp?topic=%2Fcom.ibm.director.cli.helps.doc%2Ffqm0_r_cli_rmdatasource.html

- Managed endpoint instances that are defined by IBM Systems Director server. See the following link for more information about the removal of managed objects from IBM Systems Director server:

  http://pic.dhe.ibm.com/infocenter/director/pubs/index.jsp?topic=%2Fcom.ibm.director.console.helps.doc%2Ffqm0_t_rn_removing_a_resource.html

11.4 Storage management: smcli references

This section describes the *smcli* commands for the most commonly used storage management scenarios:

- Listing all related storage volumes for a resource

  Syntax:

  ```
  smcli lsstvol -n <resource name>
  ```

  Example:

  ```
  #smcli lsstvol -n "Storwize V7000-2076-VMC7000_1-IBM"
  ```

  Where, resource selected is an IBM Storwize V7000 Storage

  ```
  #smcli lsstvol -n ip10-32-42-88
  ```

  Where, resource selected is an IBM PowerVM LPAR having vFC adapters

  > **Tip:** This command should be used to list related storage volumes for a server that has Fibre Channel ports. If your server has Fibre Channel ports and has storage volume allocated using external tools other than storage control, ensure that you have collected inventory on the server and farm object at least once.

- Creating a new storage volume in the network storage

  Syntax:

  ```
  smcli mkstvol -n "<virtual servers name separated by comma, if more than one>"
  -P "<Host name>" -p <name of the new storage volume> -s <size> -u <unit>
  ```

  Example:

  ```
  smcli mkstvol -n "CreateVS, ip10-32-42-88" -P "pfm9253_pfm9254" -p
  New_storage_volume -s 5.0 -u GB
  ```

  – Where:

  - "New_storage_volume" is the name of the new storage volume
  - "pfm9253_pfm9254" is the name of the existing storage pool on which the storage volume is created
  - "CreateVS & ip10-32-42-88" are the two virtual servers that have access to the newly created storage volume

  Tip:

  Before running this command, ensure that all server, switch, and storage devices are discovered and inventory collection is done at least once.

- Detach and delete storage volumes

  Syntax:

  ```
  smcli rmstvol -n volume_name
  ```

  Example:

  ```
  smcli rmstvol -n New_storage_volume
  ```

  – Where:

  "New_storage_volume" is the storage volume name that will be detached from the servers that it is currently attached and then deleted from the storage sub system.

Tip:

- Before running this command, you can verify that all servers are using the storage volume by using the following command:

  ```
  smcli svsrelationships --src server --rel uses --tgt storagevolume
  ```

- In a Power Server environment, you can find physical volume names present on the VIOS. To find the storage volumes backing them, run the following command:

  ```
  smcli svsrelationships -src physicalvolume -tgt storagevolume
  ```

For more storage management commands, see the following information center link:

http://pic.dhe.ibm.com/infocenter/director/pubs/index.jsp?topic=%2Fcom.ibm.director.cli.helps.doc%2Ffqm0_r_cli_storage_cmds.html

11.5 Best practices

Listed below are the best practices to be followed in using storage management features efficiently:

- If you are planning to use IBM Systems Director server for monitoring and managing the storage devices, before starting to use Storage Manager or Storage Control, perform the following functions:
  - List the storage devices that you have
  - Determine the following for each of these storage devices:
    - The type of SMI-S Provider Agents that are required.
    - The edition of storage management that is required.
    - Verify the existing firmware versions on the storage devices and ensure that they comply with IBM Systems Director server-supported firmware versions.
- Before installing IBM Systems Director Storage Control, ensure that you have a managed DB2 version that meets storage control requirements.
- For the SMI-S provider, the best practice is to consider the following factors:
  - Determine the servers where you can install the SMI-S Provider Agents, if required, for monitoring and managing the storage devices.
  - It is required to choose a system other than IBM Systems Director server or systems hosting Platform Agent for installation of the SMI-S provider (because of port conflicts).
  - Download the required SMI-S Provider Agents followed by installation and configuration of the SMI-S Provider Agent on the predetermined servers.
  - Do not install more than one SMI-S provider on the same system. If there is a need to install more than one provider on the same server, ensure that the ports on which the provider agent CIM server listens are unique.
  - After SMI-S provider installation, ensure that both CIMOM and SLP services are active. These protocols are needed for discovery and communication between IBM Systems Director server and the SMI-S provider.
  - Uninstalling some of the SMI-S providers does not remove the CIM component. If the system chosen for SMI-S provider installation was hosting any other SMI-S provider previously, ensure that the CIM components are removed completely before going ahead with the new installation of an SMI-S provider.

  - If you are reinstalling any SMI-S provider, it is recommended that you copy files that contain the IP addresses of the systems under management and the ports in use to a safe location.
  - It is always recommended to remove the SMI-S Provider Agent followed by removal of the managed storage devices.
  - There is no hard limit on how many storage systems can be supported by a single provider instance. The recommended maximum is 10 storage systems per provider.
- If your environment already has IBM Tivoli Storage Productivity Center, ensure that it is at version level 4.2.2 FP1+ before integrating it with IBM Systems Director.
- Before performing any operations that involve dealing with the assignment of new and existing volumes on the supported network storage using IBM Systems Director VMControl, ensure that the IBM Systems Director Storage Manager or Storage Control plug-in is configured with the necessary providers to communicate with the required network and storage devices.
- Ensure that the inventory collection is done for each endpoint at least once before running any IBM Systems Director or plug-in tasks. The inventory collection on the endpoints can be done in any order, or even at the same time.
- If there are any changes in the configuration of the devices managed by IBM Systems Director Storage Manager or Storage Control using the native management tools, an inventory should be collected on those devices in IBM Systems Director for it to detect the change and update the database.
- While collecting inventory on the devices managed by Storage Control or IBM Tivoli Storage Productivity Center, ensure that you collect inventory for only one device in a Storage Farm at a time.

If the storage devices are managed through Storage Control and you want to remove this storage device, it is recommended to remove both the farm objects as well as managed endpoint instances defined by IBM Systems Director server representing this storage device.

12

Troubleshooting

This chapter covers troubleshooting techniques for the different components of IBM Systems Director.

The following topics are covered:

12.1 Troubleshooting the installation of IBM Systems Director components

This section covers the information for checking and troubleshooting the installation of IBM Systems Director.

12.1.1 Pre-installation check

Since IBM Systems Director 6.3, there is a tool that is available that checks the installation requirement for the IBM Systems Director server. The **`checkds`** program is located in the `checkds` folder of the installation medium (DVD or installation package).

There are two ways to use the **`checkds`** program:

- Run as a separate program before starting the installation
- Integrated into the installation process

When running checkds as a separate program, the result of the check is shown as a web page html format on Windows and Linux (with graphical interface) or as a text file on AIX and Linux. Examples for the output can be found in the installation chapters for Windows and Linux above.

The result files can be found in the following directories:

- Windows

  `c:Users/%Username%/AppData/Temp/checkds/reports`

- Linux/AIX

  `/tmp/checkds/reports`

12.1.2 Post Installation Validator

Post Installation Validator (PIV) is a new tool that comes with IBM Systems Director version 6.3.2. It is intended for use by service personnel, not the normal user.

PIV is located on the installation media in the `PIV` folder. It must be run manually after finishing the installation of the IBM Systems Director.

The PIV analyzes the system to determine the state of the IBM Systems Director installation. The PIV tool generates a report that contains a list of errors (if there are some) extracted from the log files. The tool also analyzes the status of the IBM Systems Director and shows used ports.

To start the PIV, run the following command from the command line:

- Windows: **`PostInstallValidator_win.exe`**
- Linux on x86: **`PostInstallValidator_xLin`**
- Linux on Power: **`PostInstallValidator_pLin`**
- Linux on z: **`PostInstallValidator_zLin`**
- AIX: **`PostInstallValidator_AIX`**

Per default, a text (PostInstallationReport.txt) and HTML (PostInstallationReport.html) report is created. Following is the default location for these reports:

- Windows:

  `%temp%`

- Linux and AIX:

  `/tmp`

There are different options that are available. With the **-h** option, a list of all available options is shown (Figure 12-1).

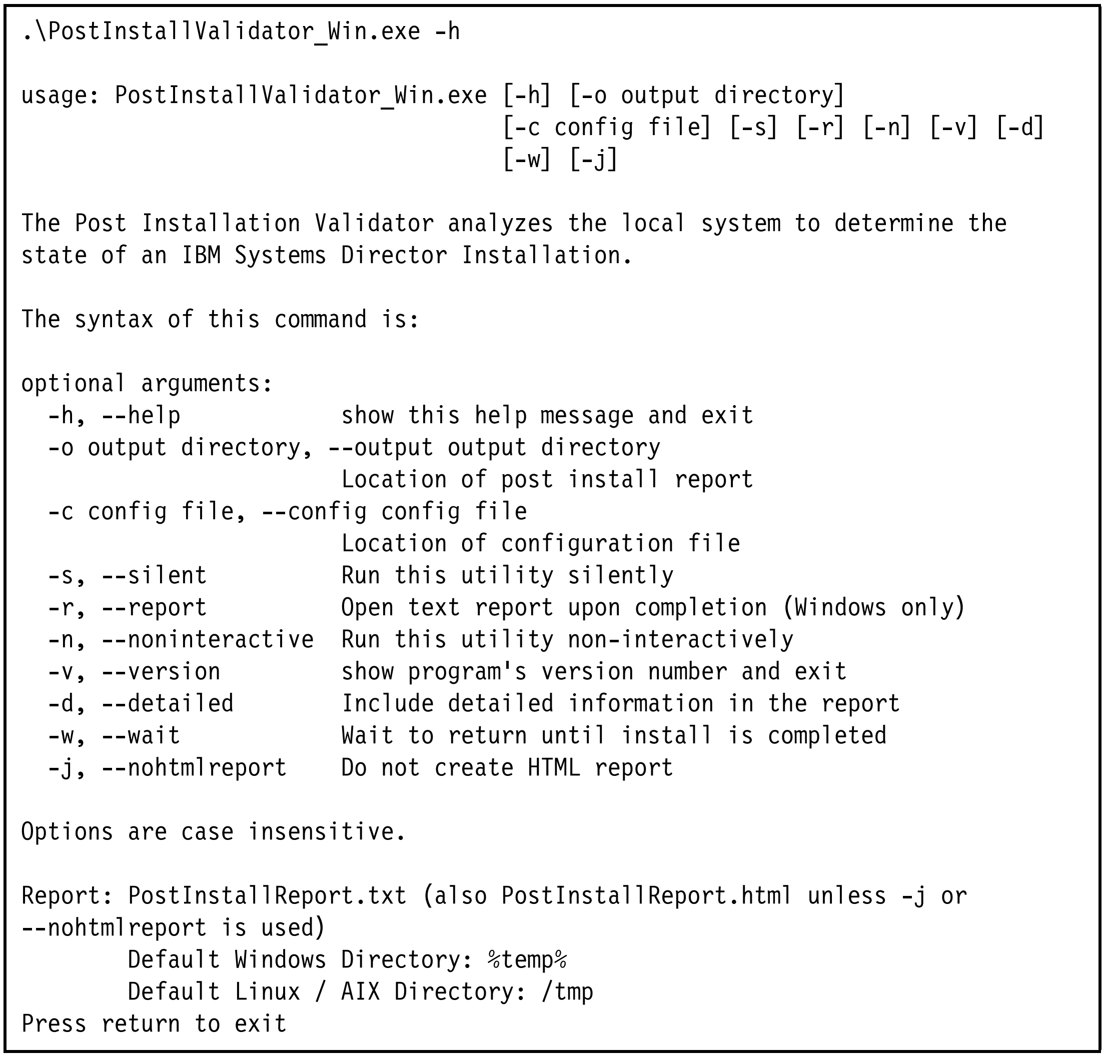

```
.\PostInstallValidator_Win.exe -h

usage: PostInstallValidator_Win.exe [-h] [-o output directory]
                                    [-c config file] [-s] [-r] [-n] [-v] [-d]
                                    [-w] [-j]

The Post Installation Validator analyzes the local system to determine the
state of an IBM Systems Director Installation.

The syntax of this command is:

optional arguments:
  -h, --help            show this help message and exit
  -o output directory, --output output directory
                        Location of post install report
  -c config file, --config config file
                        Location of configuration file
  -s, --silent          Run this utility silently
  -r, --report          Open text report upon completion (Windows only)
  -n, --noninteractive  Run this utility non-interactively
  -v, --version         show program's version number and exit
  -d, --detailed        Include detailed information in the report
  -w, --wait            Wait to return until install is completed
  -j, --nohtmlreport    Do not create HTML report

Options are case insensitive.

Report: PostInstallReport.txt (also PostInstallReport.html unless -j or
--nohtmlreport is used)
        Default Windows Directory: %temp%
        Default Linux / AIX Directory: /tmp
Press return to exit
```

Figure 12-1 Post Installation Validator options

When running the PIV tool in a Windows environment, for example, you see the different steps and the result for each one of the steps (Figure 12-2).

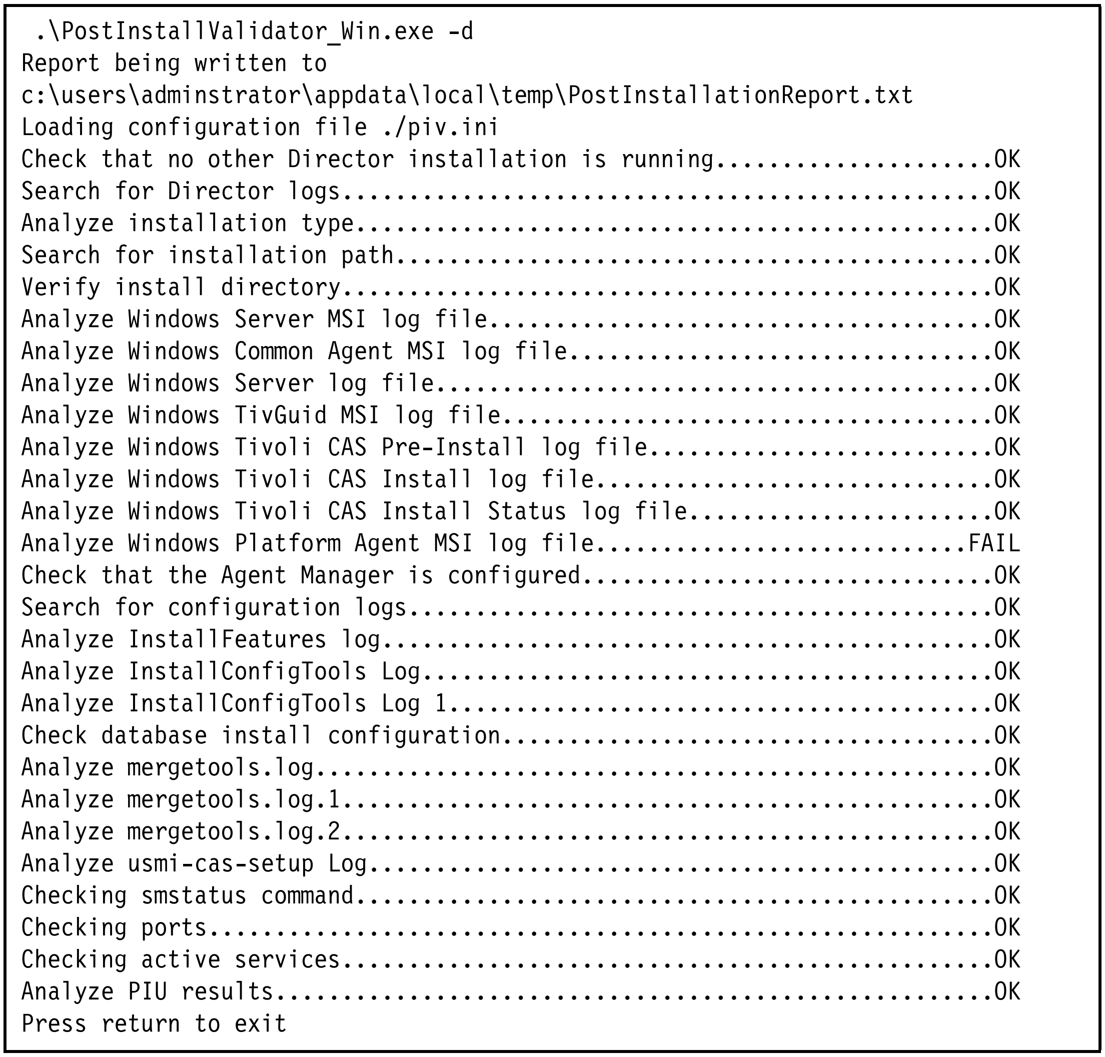

```
 .\PostInstallValidator_Win.exe -d
Report being written to
c:\users\adminstrator\appdata\local\temp\PostInstallationReport.txt
Loading configuration file ./piv.ini
Check that no other Director installation is running.....................OK
Search for Director logs.................................................OK
Analyze installation type................................................OK
Search for installation path.............................................OK
Verify install directory.................................................OK
Analyze Windows Server MSI log file......................................OK
Analyze Windows Common Agent MSI log file................................OK
Analyze Windows Server log file..........................................OK
Analyze Windows TivGuid MSI log file.....................................OK
Analyze Windows Tivoli CAS Pre-Install log file..........................OK
Analyze Windows Tivoli CAS Install log file..............................OK
Analyze Windows Tivoli CAS Install Status log file.......................OK
Analyze Windows Platform Agent MSI log file............................FAIL
Check that the Agent Manager is configured...............................OK
Search for configuration logs............................................OK
Analyze InstallFeatures log..............................................OK
Analyze InstallConfigTools Log...........................................OK
Analyze InstallConfigTools Log 1.........................................OK
Check database install configuration.....................................OK
Analyze mergetools.log...................................................OK
Analyze mergetools.log.1.................................................OK
Analyze mergetools.log.2.................................................OK
Analyze usmi-cas-setup Log...............................................OK
Checking smstatus command................................................OK
Checking ports...........................................................OK
Checking active services.................................................OK
Analyze PIU results......................................................OK
Press return to exit
```

Figure 12-2 Run of PostInstallValidator_Win.exe

In our example, you see a failed result for the Platform Agent MSI log. Now you can check the `PostInstallationReport` file to find the problem.

An excerpt from the report file that shows the failure can be seen in Figure 12-3.

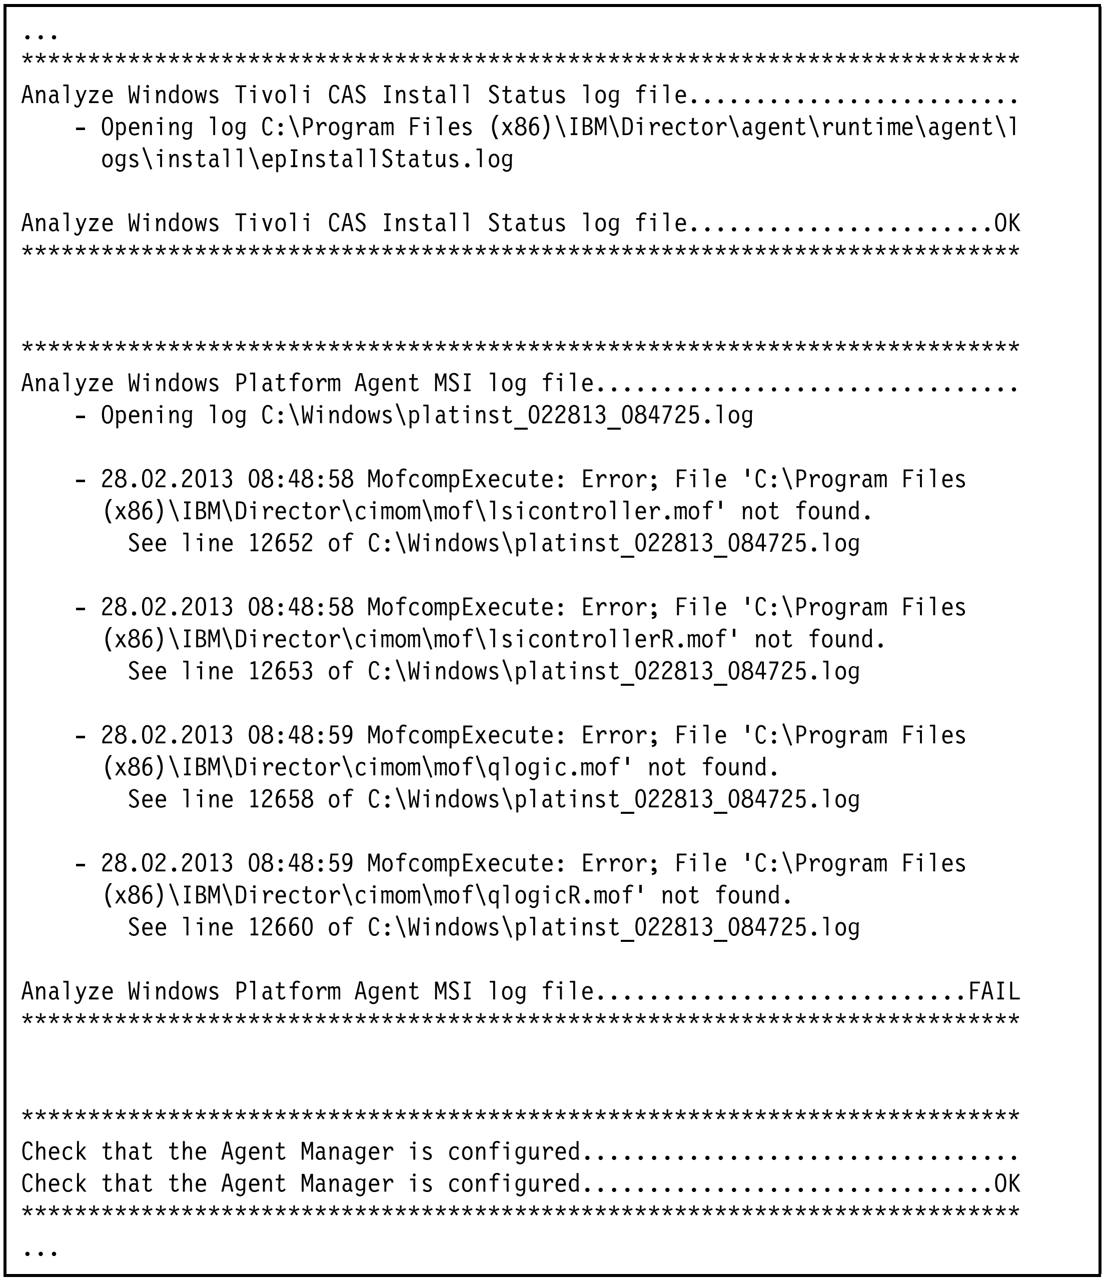

```
...
***************************************************************************
Analyze Windows Tivoli CAS Install Status log file.........................
    - Opening log C:\Program Files (x86)\IBM\Director\agent\runtime\agent\l
      ogs\install\epInstallStatus.log

Analyze Windows Tivoli CAS Install Status log file.......................OK
***************************************************************************



***************************************************************************
Analyze Windows Platform Agent MSI log file................................
    - Opening log C:\Windows\platinst_022813_084725.log

    - 28.02.2013 08:48:58 MofcompExecute: Error; File 'C:\Program Files
      (x86)\IBM\Director\cimom\mof\lsicontroller.mof' not found.
        See line 12652 of C:\Windows\platinst_022813_084725.log

    - 28.02.2013 08:48:58 MofcompExecute: Error; File 'C:\Program Files
      (x86)\IBM\Director\cimom\mof\lsicontrollerR.mof' not found.
        See line 12653 of C:\Windows\platinst_022813_084725.log

    - 28.02.2013 08:48:59 MofcompExecute: Error; File 'C:\Program Files
      (x86)\IBM\Director\cimom\mof\qlogic.mof' not found.
        See line 12658 of C:\Windows\platinst_022813_084725.log

    - 28.02.2013 08:48:59 MofcompExecute: Error; File 'C:\Program Files
      (x86)\IBM\Director\cimom\mof\qlogicR.mof' not found.
        See line 12660 of C:\Windows\platinst_022813_084725.log

Analyze Windows Platform Agent MSI log file............................FAIL
***************************************************************************



***************************************************************************
Check that the Agent Manager is configured.................................
Check that the Agent Manager is configured...............................OK
***************************************************************************
...
```

Figure 12-3 Excerpt from the PostInstallationReport.txt file

In our example, the IBM Systems Director ran in a virtual machine (VM). Therefore, no LSI adapter and no QLogic adapter are installed. That is why no Managed Object Format (MOF) files for these adapters were compiled.

12.1.3 IBM Systems Director messages

When running a task or getting a message box in the IBM Systems Director console, you see a message code in front of the message.

This code has a specific format, a unique message identifier (example: ATKSRV614E).

The format **XXXYYY######Z** is defined in the following way:

- XXX identifies the component prefix that sends the message
- YYY identifies the subsystem from which the message comes (this is optional)
- The next 3..6 digits ###### identify the message number
- Z, the last character, identifies the severity code:
  - E: Error message
  - I: Information message
  - W: Warning message

There is no list with all the messages for download available, but you can search for the message in the IBM Systems Director Information Center at the following link:

`http://pic.dhe.ibm.com/infocenter/director/pubs/index.jsp`

Put the message in the search field on top of the left column and then the result shows with the explanation of the message (Figure 12-4). You can also put in a part of the message key, for example, ATK*, to get all messages that start with ATK. Then, you can go through the list and search the message that you want to get information from.

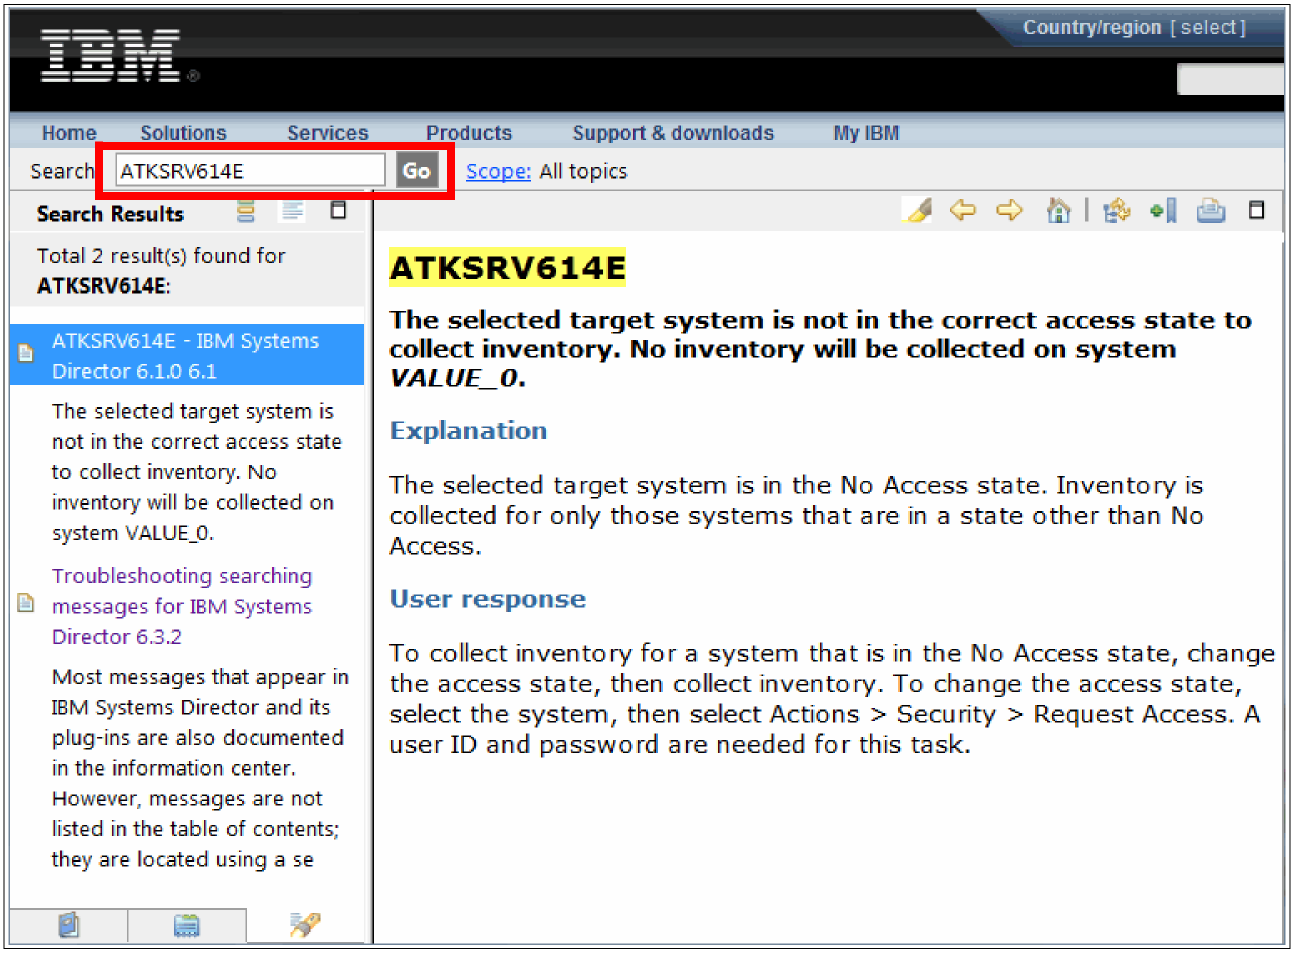

Total 2 result(s) found for **ATKSRV614E**:

ATKSRV614E - IBM Systems Director 6.1.0 6.1

The selected target system is not in the correct access state to collect inventory. No inventory will be collected on system VALUE_0.

Troubleshooting searching messages for IBM Systems Director 6.3.2

Most messages that appear in IBM Systems Director and its plug-ins are also documented in the information center. However, messages are not listed in the table of contents; they are located using a se

ATKSRV614E

The selected target system is not in the correct access state to collect inventory. No inventory will be collected on system *VALUE_0*.

Explanation

The selected target system is in the No Access state. Inventory is collected for only those systems that are in a state other than No Access.

User response

To collect inventory for a system that is in the No Access state, change the access state, then collect inventory. To change the access state, select the system, then select Actions > Security > Request Access. A user ID and password are needed for this task.

Figure 12-4 Search for messages in IBM Systems Director Information Center

12.1.4 IBM Systems Director log files

IBM Systems Director creates some log files during the installation and also during the work. These log files can help to find problems and solve them. In Table 12-1, you can see the location of the most important log files for IBM Systems Director and agents.

Table 12-1 Location of log files by OS

Log file type	Windows	Linux	AIX	IBM i
IBM Systems Director server installation logs	c:\Windows\ ► dirserverinst_*time*.log ► dirserversetup.log ► db2prereqcheck.log ► prereqcheckdb2.log ► installdb2.log ► tivguidinst64.log %inst_dir%\IBM\Director\log\recovery.txt	► /var/log/dirinst.log ► /opt/ibm/director/log/*.log ► /opt/ibm/director/log/recovery.txt	► /var/log/dirinst.log ► /var/log/director/* ► /opt/ibm/director/log/recovery.txt	-
Common Agent installation logs	c:\Windows\ ► agent_install.log ► CasInst.log ► certutil.log ► diragnetinst_*time*.log ► diragnetsetup.log ► tivguidinst.log	/var/log/dirinst.log	/var/log/dirinst.log	► /tmp/director/installAgent.log ► /tmp/director/installDirAgent.log ► www/cas/lwi/runtime/agent/logs/install
Platform Agent installation logs	c:Windows\platinst_time.log	/var/opt/ibm/platform/log/install.log	/var/log/dirinst.log	-
Common Agent uninstall logs			-	/tmp/director/uninstallDirAgent.log
Platform agent uninstall logs	c:\Windows\ ► platuninst.logs ► sysdiremoval.log	/var/opt/ibm/platform/log/uninstall.log		-
Configuration logs	%inst_dir%\IBM\Director\log\*.log	/opt/ibm/director/log/*.log	/var/log/director/*	-

Logcollector tool

To help collect all necessary log_files and information for a support case, a tool named **`logcollector`** is available. The tool is located in the `%install_dir%/IBM/Director/bin` directory. It is a command line tool and creates a packed report file in the following format `dirlogs-`*`systename-date-time.zip`*. This report per default is located in the user home directory from the user that runs the **`logcollector`** tool.

The **`logcollector`** tool collects information from and runs checks with the following parts:

- Common Agent on the system
- IBM Systems Director database
- Network Control
- Platform Agent information
- IBM Systems Director server Core (actual status and configuration of the server)
- Server JavaCoreDump
- smcli collection
- Storage Control collection
- Local System collection (network configuration, systeminfo, firewall)

The size of the report file depends on the number of installed advanced managers and the number of endpoints managed by the IBM Systems Director server.

The report file has a data structure like what is shown in Figure 12-5. There can be more information in the report when, for example, Network Control and Storage Control are installed on the server.

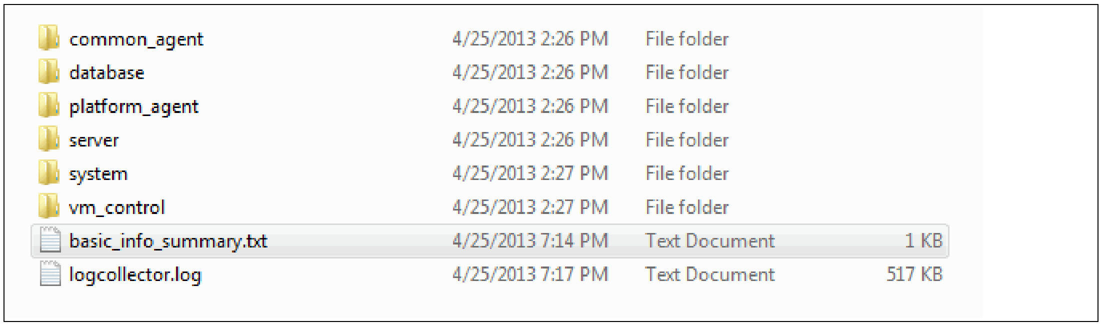

common_agent	4/25/2013 2:26 PM	File folder	
database	4/25/2013 2:26 PM	File folder	
platform_agent	4/25/2013 2:26 PM	File folder	
server	4/25/2013 2:26 PM	File folder	
system	4/25/2013 2:27 PM	File folder	
vm_control	4/25/2013 2:27 PM	File folder	
basic_info_summary.txt	4/25/2013 7:14 PM	Text Document	1 KB
logcollector.log	4/25/2013 7:17 PM	Text Document	517 KB

Figure 12-5 Logcollector report file: data structure

This report file is often needed by the support if there are problems with IBM Systems Director and a defect call is opened.

In some rare cases, especially in large environments, the log collector can fail to collect all information because of a timeout problem. The client might not see this because a report file is created that contains all log information until the time of failure.

To avoid this problem, increase the time that the program waits for commands to complete. Run the following command:

```
logcollector -t 600
```

Where: The value of `-t 600` defines a 5-minute (set in seconds) time to wait for completion. The default value is 120 seconds (2 minutes).

12.2 Troubleshooting security-related issues

If you are unable to log on to the Systems Director server or if the server fails to start, review the logs to determine the error. You can increase the logging with Systems Director in two ways:

- **`lwilog.sh`** script
- `logging.properties` file

If the server is active, use the **`lwilog.sh`** script as shown in Figure 12-6.

```
-bash-3.2# /opt/ibm/director/lwi/bin/lwilog.sh -addlogger -name
com.ibm.lwi.security.rolemanagers.ldap -level FINEST
ALR0299I: Logger successfully added for package com.ibm.lwi.security.rolemanagers.ldap.
SUCCESS
-bash-3.2
```

Figure 12-6 lwilog.sh script

If the server is not active, edit the `logging.properties` file as shown in Figure 12-7.

```
-bash-3.2#echo -e "#additional LDAP
logging\ncom.ibm.lwi.security.rolemanagers.ldap.level=FINEST" >>logging.properties
```

Figure 12-7 Increase the logging level

The following additional logging attributes are available:

- For the **`lwilog.sh`** script:

  ```
  com.ibm.usmi.kernel.security -level FINEST
  com.ibm.usmi.console.security -level FINEST
  ```

- For the `logging.properties` file:

  ```
  com.ibm.usmi.kernel.security.level=ALL
  com.ibm.usmi.console.security.level=ALL
  ```

After the logging threshold is changed by using the **`lwilog.sh`** script, refresh the logs by using the **`-refresh`** parameter:

```
/opt/ibm/director/lwi/bin/lwilog.sh -refresh
```

If you edit the `logging.properties` file, restart the Systems Director server. Log files are stored in the `/opt/ibm/director/lwi/logs/error-log-0.html` directory.

12.2.1 Restoring local OS authentication

To restore Systems Director to use local OS authentication, use the **`cfguserreg.sh`** script (Figure 12-8). Before you use the script, restore the original `security.properties` file. Remove the `securityLDAP.properties` file from the `/opt/ibm/director/lwi/conf/overrides/` directory.

```
-bash-3.2# cfguserreg.sh -os
/opt/ibm/director/bin
Security settings have been set to use operating system registry.
Restart IBM Systems Director Server to complete configuration.

-bash-3.2#smstop;smstart;smstatus -r
Shutting down IBM Director...
Starting IBM Director...
The starting process may take a while. Please use smstatus to check if the server is
active.
Starting
Active
```

Figure 12-8 cfguserreg.sh script

12.2.2 Additional information

For OpenLDAP support with Systems Director, see the IBM Support Portal:

http://ibm.com/support/search.wss?q1=openldap&tc=SGZ2Z3

For common troubleshooting steps with LDAP, see this site:

http://ibm.com/support/docview.wss?uid=nas7917752a664b2c71a8625768e0001ab13

For additional common troubleshooting steps with LDAP, see the following site:

http://ibm.com/support/docview.wss?uid=nas7cf1a05b97228ef0d86257749007b7025

12.3 Troubleshooting VMControl

This section covers troubleshooting steps for the VMControl advanced manager plug-in.

12.3.1 KVM troubleshooting

This section describes basic operations for common troubleshooting tasks with Kernel-based Virtual Machine (KVM) on Red Hat Enterprise Linux and IBM Systems Director VMControl.

Handling request access failures with Common Agent

The management software uses an agent manager to communicate with Common Agent after it is installed on a managed system.

The agent manager provides authentication and authorization services for installed common agents and the management software. It also maintains a registry of configuration information about Common Agent-managed systems.

There are several actions that you can take to investigate and resolve Request Access failures for a system that is managed by IBM Systems Director and has Common Agent installed:

- Ensure that the system is not being managed by more than one agent manager

  An IBM Systems Director-managed resource can be managed by only one agent manager at a time. If the managed resource is an operating system and IBM Systems Director stops managing it, its agent is not unregistered from the agent manager. To manage the agent with IBM Systems Director unregisters the agent manually.

- Query current manager

  On the operating system endpoint, query the usma service detail information to determine the current agent manager. Issue the following command that is applicable for the managed compute node operating system. See Example 12-1.

> **Note:** To run the following command on a system with Virtual I/O Server (VIOS), you must first exit the restricted shell. To get out of the restricted shell, run the **oem_setup_env** command before you attempt to query the usma service detail information.

AIX or Windows:

```
slp_query --type=* --address=<system_IP_address>
```

See Example 12-1 for the querying manager in Linux.

Example 12-1 Querying manager in Linux

```
[root@node11229 ~]# /opt/ibm/icc/bin/slptool -u 127.0.0.1 findattrs
service:management-software.IBM:usma
(ip-address=192.168.122.1@10.31.21.203@10.31.21.148),(mac-address=3440B5BE9178),(t
ivguid=CE32DB6061CD11E2BFC63440B5BE9178),(uid=31322f0d5b16c1ea),(uuid=C4E2DA1E-206
D-11E2-A23E-3440B5BE9178),(vendor=IBM),(System-Name=node11229),(timezone-offset=12
0),(version=6.3.2),(port=9510),(manager=10.31.54.199)
```

- Reconnect host to manager:

  On the operating system endpoint, issue the following command that is applicable for the managed compute node operating system:

  AIX or Linux

  ```
  /opt/ibm/director/agent/runtime/agent/toolkit/bin/configure.sh -unmanaged
  -force
  ```

  Windows

  ```
  <install_root>\agent\runtime\agent\toolkit\bin\configure.bat -unmanaged -force
  ```

  Remove the operating system endpoint from the IBM Systems Director and rediscover the operating system endpoint, then request access again.

- Ensure that you are using the correct user ID and password

  When requesting access to a secured system from the IBM Systems Director Request Access page, you cannot enter a password longer than 45 characters. The Request Access task has a password length limit of 45 characters.

  Additional characters are ignored. You cannot gain access to a Common Agent-managed system by using a user credential that has a password greater than 45 characters in length. To correct this problem, shorten the password on the managed system to gain access.

  Tip: Check by manually logging in to the managed system via SSH, Telnet (tn), or mstsc.exe to confirm that the password is correct.

  Many browsers pre-fill the password field. Empty it and ensure that it is the correct password.

- Check the firewall

  Request Access requires a reliable connection between the IBM Systems Director and Common Agent. You need a connection in both directions: agent to server and server to agent. Ensure that you can use SSH, Telnet (tn), or mstsc.exe from either direction.

  Agent and server-side example checks follow (Example 12-2):

  Windows agent side:

  ```
  telnet ISD_IPADDR 9513
  ```

  AIX or Linux agent side:

  ```
  /opt/ibm/director/agent/runtime/agent/toolkit/bin/checkconn.sh -host AM_IP
  -password AM_password
  ```

  Example 12-2 Connection check in Linux

  ```
  [root@node11229 ~]#
  /opt/ibm/director/agent/runtime/agent/toolkit/bin/checkconn.sh -host
  10.31.54.199 -password PASSWORD

  BTC8614I The connection to the agent manager has been established successfully.
  BTC8619I The common agent would register successfully with the given password.
  ```

  Windows:

  ```
  <install_root>\agent\runtime\agent\toolkit\bin\checkconn.bat -host AM_IP
  -password AM_password
  ```

Server side:

Test from IBM Systems Director server that network connectivity to CAS agent is valid:

```
telnet Agent_IP 9510
```

- Check to see if the IBM Systems Director IP address has changed:

  If there is an IBM Systems Director IP address change, restart the server before issuing a Request Access. A correct server IP address is a prerequisite for Request Access.

  You can run the following commands to check the IP address. See Example 12-3.

Example 12-3 Check that Systems Director is in the correct IP address

```
[root@node11229 ~]# wget -O AgentMgr.Info http://10.31.54.199:9513/AgentMgr/Info
--2013-04-18 22:08:17--  http://10.31.54.199:9513/AgentMgr/Info
Connecting to 10.31.54.199:9513... connected.
HTTP request sent, awaiting response... 200 OK
Length: unspecified [text/html]
Saving to: âAgentMgr.Infoâ

    [ <=>                                                                          ]
15,374      --.-K/s   in 0s

2013-04-18 22:08:17 (68.9 MB/s) - âAgentMgr.Infoâ
```

- Agent Health check. See Example 12-4.

  AIX or Linux:

  ```
  /opt/ibm/director/lwi/runtime/agentmanager/toolkit/bin/HealthCheck.sh
  -toolkitPassword AM_password
  ```

  Windows:

  ```
  <install_root>\lwi\runtime\agentmanager\toolkit\bin\HealthCheck.bat
  -toolkitPassword AM_password
  ```

Example 12-4 Agent Health check

```
C:\Program Files\IBM\Director\lwi\runtime\agentmanager\toolkit\bin>HealthCheck.bat -toolkitPassword Salasana1
CTGEM2470I It is the agent manager instance ID: 7e61ca6dd647361cbcd94863afc27cd8

AGENT_CONFIG_SERVICE                        = https://127.0.0.1:9512/AgentMgr/AgentConfiguration
CRL_REQUEST                                 = http://127.0.0.1:9513/AgentMgr/CRLRequest
MULTI_SCHEDULER_SYNCHRO_SERVICE             = https://127.0.0.1:9512/AgentMgr/MultiScheduleSynchronizer
FILE_DOWNLOAD_SERVICE                       = http://127.0.0.1:9513/AgentMgr/Patches
UPGRADE_SERVICE                             = https://127.0.0.1:9512/AgentMgr/UpgradeService
REGISTRATION_SERVICE                        = https://127.0.0.1:9511/AgentMgr/Registration
SERVICE_CATALOGUE_REQUEST                   = http://127.0.0.1:9513/AgentMgr/ServiceCatalogueRequest
AGENT_MANAGER_QUERY                         = http://127.0.0.1:9513/AgentMgr/AgentManagerQuery
AGENT_QUERY                                 = https://127.0.0.1:9512/AgentMgr/AgentQuery
MIGRATION_SERVICE                           = https://127.0.0.1:9512/AgentMgr/MigrationService
MANAGER_CREDENTIALS_SERVICE                 = https://127.0.0.1:9512/AgentMgr/AuthAdmin
TRUSTSTORE_REQUEST                          = http://127.0.0.1:9513/AgentMgr/TrustedCertificateQuery
PATCH_SERVICE                               = https://127.0.0.1:9512/AgentMgr/PatchService
COMMON_AGENT_QUERY                          = https://127.0.0.1:9512/AgentMgr/CommonAgentQuery
CERT_REVOCATION_SERVICE                     = https://127.0.0.1:9512/AgentMgr/CertificateRevocation
SCHEDULER_ADMIN_SERVICE                     = https://127.0.0.1:9512/AgentMgr/ScheduleManager
DELETE_AGENTS_SERVICE                       = https://127.0.0.1:9512/AgentMgr/DeleteAgents
SCHEDULER_SYNCHRO_SERVICE                   = https://127.0.0.1:9512/AgentMgr/ScheduleSynchronizer
VERSION_SERVICE                             = http://127.0.0.1:9513/AgentMgr/Version
DEREGISTER_SERVICE                          = https://127.0.0.1:9512/AgentMgr/DeregistrationService
CONFIG_UPDATE_SERVICE                       = https://127.0.0.1:9512/AgentMgr/ConfigurationUpdate
CERT_RENEWAL_SERVICE                        = https://127.0.0.1:9512/AgentMgr/CertificateRenewal
FILE_ADMIN_SERVICE                          = https://127.0.0.1:9512/AgentMgr/PatchAdmin
INFO_PAGE                                   = http://127.0.0.1:9513/AgentMgr/Info
IPADDRESS_SERVICE                           = http://127.0.0.1:9513/AgentMgr/IPAddress
JOB_MANAGER_SERVICE                         = https://127.0.0.1:9512/AgentMgr/JobManager

CTGEM2450I The Health Check tool passed.
```

- Check the Common Agent status. See Example 12-5.

  Request Access requires that the Common Agent is in good status. You can run the following queries to check agent status:

  AIX or Linux

Example 12-5 Check endpoint status in host

```
[root@node11229 director]#
/opt/ibm/director/agent/runtime/agent/bin/endpoint.sh status
Running.
[root@node11229 director]#
```

Example 12-6 shows the test connector status.

Example 12-6 Test connector

```
[root@node11229 director]#
/opt/ibm/director/agent/runtime/agent/bin/agentcli.sh connector alive
BTC7101I The connector is active.
```

Windows:

- `<install_root>\agent\runtime\agent\bin\endpoint.bat status Running.`
- `<install_root>\agent\runtime\agent\bin\agentcli.bat connector alive BTC7101I The connector is active.`

If the Common Agent status is not active, complete the following steps:

Issue the following command to check to ensure that the SLP is in an operational state. See Example 12-7.

Example 12-7 Check that you have SLP services running

```
[root@node11229 director]# ps -ef|grep slp
daemon   25079     1  0 Jan30 ?        00:08:32 /opt/ibm/icc/bin/slpd
root     25672 27826  0 23:15 pts/3    00:00:00 grep slp
root     25700     1  0 Jan30 ?        00:00:04 /opt/ibm/platform/bin/tier1slp
```

Test from IBM Systems Director server that you can connect the Common Agent host to port 427:

```
telnet AGENT_IPADDRESS 427
```

If you cannot reach the port, the SLP service is not working properly.

In this case, you might need to install with yum openslp package. Restart the server and remove the Common Agent host from the Systems Director inventory and place the host in the unmanaged state.

Common Agent running on AIX or Linux:

```
/opt/ibm/director/agent/runtime/agent/toolkit/bin/configure.sh -unmanaged
-force
```

Common Agent running on Windows:

```
<install_root>\agent\runtime\agent\toolkit\bin\configure.bat -unmanaged -force
```

After that, perform the following command on the Common Agent host.

Common Agent running on AIX or Linux:

```
/opt/ibm/director/agent/runtime/agent/toolkit/bin/configure.sh -amhost
agentmanager_IP_address -passwd agentregistration_password -force
```

Common Agent running on Windows:

```
<install_root>\agent\runtime\agent\toolkit\bin\configure.bat -amhost
active_agent_manager_IP_address -passwd active_agent_registration_password
-force
```

Discover the Common Agent host from the Systems Director user interface and connect it with root credentials.

- Check DNS settings:

  Ensure that the DNS is set correctly on the agent system. If no DNS is configured, check the hosts files on the agent system. The following entry is available:

  ```
  - 127.0.0.1 localhost
  - ::1 localhost
  ```

  > **Note:** -:1 localhost applies only if there are IPv6 addresses.
  >
  > Every global IP address should be mapped to the host name of the endpoint.

12.3.2 VMware vSphere troubleshooting

This topic describes common problems when working with IBM Systems Director and VMware environments.

Launching VMware management software

From VMControl, you can open VMware management software such as the VMware vSphere Client or the VMware Infrastructure Client. To launch a VMware client from IBM Systems Director VMControl, the VMware client must be installed in the default location:

```
C:\\Program Files (x86)\\VMware\\Infrastructure\\Virtual Infrastructure Client\\
Launcher\\VpxClient.exe
```

However, if you installed the VMware client in a different location, you can use a text editor to create a properties file to override the default location. Create the file named `vmcontrolvmware.properties` in the following directory on your IBM Systems Director server system:

- AIX or Linux:

  ```
  install_path/ibm/director/lwi/conf/overrides/vmcontrolvmware.properties
  ```

- Windows:

  ```
  install_path\IBM\Director\lwi\conf\overrides\vmcontrolvmware.properties
  ```

The “install_path” field is the location where IBM Systems Director is installed.

You can copy the following sample and paste it into your `vmcontrolvmware.properties` file. Revise it to specify where the VMware client is installed in your environment. Ensure that the installation path is all on one line in the `vmcontrolvmware.properties` file that you create:

```
RequiredPathToClientVC4=C:\\Program Files\\VMware\\Infrastructure\\Virtual
Infrastructure Client\\Launcher\\VpxClient.exe
RequiredPathToClientVC5=C:\\Program Files\\VMware\\Infrastructure\\Virtual
Infrastructure Client\\Launcher\\VpxClient.exe
```

The defined variable must be formatted as a Windows path with file separators specified as `\\`.

The following variables map to versions of VMware software:

- `RequiredPathToClientVC4 = VMware vCenter 4`
- `RequiredPathToClientVC5 = VMware vCenter 5`

> **Note:** After the `vmcontrolvmware.properties` file has been created to set the defined VMware client installation path, every Windows system that is used to launch the VMware client must have the VMware Infrastructure Client installed in the specified location.

After the `vmcontrolvmware.properties` file is created, you must collect inventory on the VMware vCenter systems.

To launch a VMware client, complete the following instructions.

In the IBM Systems Director navigation pane, click **Resource Explorer** to locate the host or virtual server from which you want to start the VMware client. Right-click the host or virtual server, and select **VMware Client**. The VMware client is started in a new window.

VMware ESXi hosts not discovered in a BladeCenter chassis

This problem affects discovery of VMware ESXi version 4.0 hosts in a BladeCenter chassis:

Problem During discovery of multiple VMware ESXi version 4.0 operating systems, only one server-managed endpoint is associated after the request access task completes. The problem occurs if the BladeCenter chassis has not been discovered and *request access* has not been granted before discovering the operating system with VMware ESXi version 4.0.

Explanation In VMware ESXi version 4.0, the same IP address is commonly surfaced for all of the VMware ESXi hosts. IBM Systems Director interprets the same IP address as the same hardware. This problem is resolved in VMware ESXi version 4.1 or later versions.

Resolution Remove the VMware ESXi host systems and the single server managed endpoint. Discover the BladeCenter chassis and request access to it. Rediscover the VMware ESXi operating systems and request access. The servers are now associated properly to each VMware ESXi operating-system managed endpoint.

12.3.3 PowerVM troubleshooting

This section covers the troubleshooting for PowerVM.

HMC discovery

Ensure that the ID that is used to request access to the HMC has **`hmcsuperadmin`** rights or **`hmcoperator`** rights.

Check the status of the following protocols: SSH and CIM.

VMControl NIM subagent

If the VMControl subagent installation fails, check the following error logs for detailed error messages:

- AIX or Linux: **`director/agent/logs/*.xml`**
- Windows: **`director\agent\logs\*.xml`**

Where: The “director” field is the path where IBM Systems Director is installed.

For more information about how to install the IBM Systems Director VMControl subagent manually, see the following site:

http://pic.dhe.ibm.com/infocenter/director/pubs/topic/com.ibm.director.vim.helps.doc/fsd0_vim_t_installing_agent_manual.html

For additional configurations after the installation, see the following information center link:

http://pic.dhe.ibm.com/infocenter/director/pubs/topic/com.ibm.director.vim.helps.doc/fqm0_t_configuring_vmc.html

Check if the SLP query lists required services advertisement:

```
# slp_query --type=* --address=10.32.42.98
...
URL: service:wbem:http://10.32.42.98:5988
URL: service:wbem:https://10.32.42.98:5989
ATTR: (template-url-syntax=https://10.32.42.98:5989) (https://10.32.42.98:5989%29)
URL: service:management-software.IBM:platform-agent://10.32.42.98
ATTR:
(vendor=IBM),(version=6.3.3),(uid=7F836092D21E574B),(ip-address=10.32.42.98),(host
name=ip10-32-42-98.pokprv.stglabs.ibm.com)
URL: service:management-software.IBM:usma://ip10-32-42-98.pokprv.stglabs.ibm.com
ATTR:
(ip-address=10.32.42.98),(mac-address=8e.15.19.c7.f3.3),(tivguid=A0F2F5A6AC5311E2B
A4A8E1519C7F303),(uid=7f836092d21e574b),(vendor=IBM),(System-Name=ip10-32-42-98.po
kprv.stglabs.ibm.com),(timezone-offset=-240),(version=6.3.3),(port=9510),(manager=
unmanaged)
URL: service:TivoliCommonAgent://ip10-32-42-98.pokprv.stglabs.ibm.com:9510
ATTR:
(ca-uid=file:///var/opt/tivoli/ep/runtime/agent),(am-host=null),(ca-ips=10.32.42.9
8),(ca-basic-port=9510),(ca-cert-port=9510),(ca-version=1.4.2.4),(os-uid=A0F2F5A6A
C5311E2BA4A8E1519C7F303)
URL: service:service-agent://10.32.42.98
ATTR: (service-type=service:management-software.IBM:usma,service:service-agent)
```

Virtual server relocation validation

If you are not able to perform a relocation, check from your HMC if the logical partition (LPAR) is valid to be relocated:

1. Go to Servers and then select the host where your LPAR is located.
2. Select the server and open the pop-up menu.
3. Select **Operations** → **Mobility** → **Validate** (Figure 12-9 on page 509).

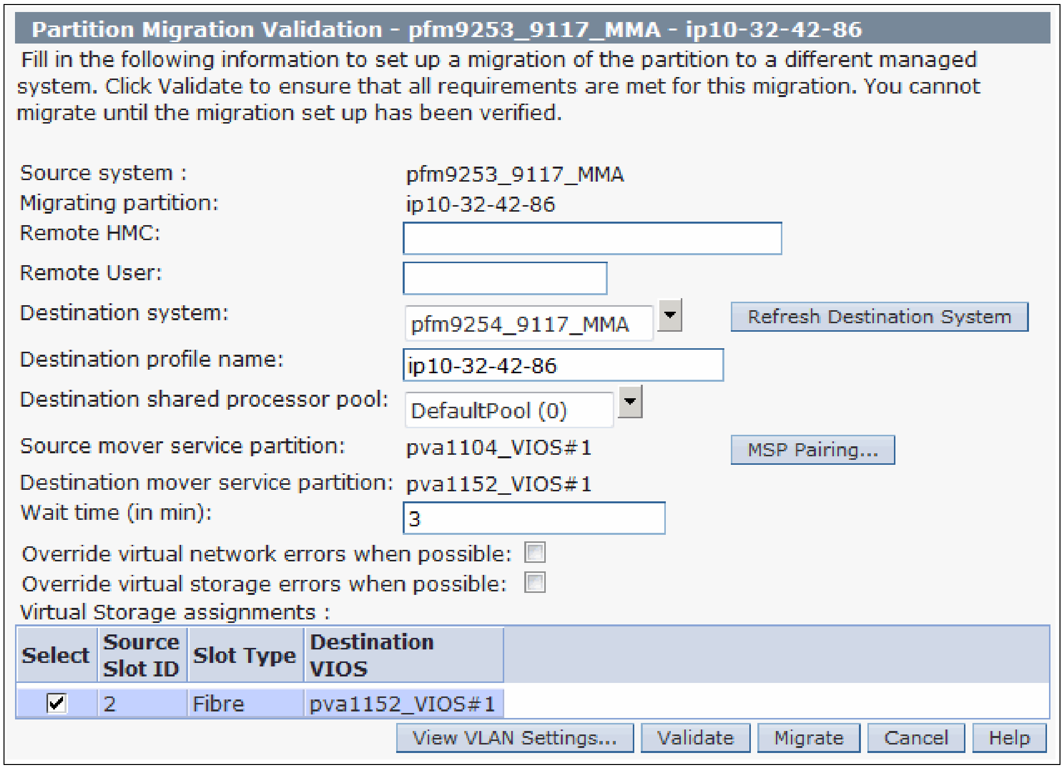

Partition Migration Validation - pfm9253_9117_MMA - ip10-32-42-86

Fill in the following information to set up a migration of the partition to a different managed system. Click Validate to ensure that all requirements are met for this migration. You cannot migrate until the migration set up has been verified.

Source system :	pfm9253_9117_MMA
Migrating partition:	ip10-32-42-86
Remote HMC:	
Remote User:	
Destination system:	pfm9254_9117_MMA
Destination profile name:	ip10-32-42-86
Destination shared processor pool:	DefaultPool (0)
Source mover service partition:	pva1104_VIOS#1
Destination mover service partition:	pva1152_VIOS#1
Wait time (in min):	3

Override virtual network errors when possible: □

Override virtual storage errors when possible: □

Virtual Storage assignments :

Select	Source Slot ID	Slot Type	Destination VIOS
☑	2	Fibre	pva1152_VIOS#1

Figure 12-9 HMC Partition Migration Validation page

12.4 Troubleshooting AIX Profile Manager

When having problems with AIX Profile Manager, consider the following troubleshooting tips:

- When deploying a template, all logs are available in the Task Management tab. Locate the job, open it, and go to the Logs tab. Everything is logged here.
- When using IBM security AIXPert, all logs are available in the `/etc/security/aixpert/log` directory.
- With artex commands, you can specify a debug level to have more information when a command is not working properly. The debug level must be specified in the `/etc/security/artex/artex.conf` file.

12.5 Troubleshooting Workload Partition Manager

This section shows some examples that can help in case you have issues in your environment with the WPAR Manager.

12.5.1 WPAR-capable system does not appear in WPAR Manager console

Check the basic requirements:

1. Does the managed system have the IBM Systems Director Common Agent installed?

Use the **lslpp** command as shown in Example 12-8 to check if IBM Systems Director Common Agent is installed on the server.

Example 12-8 Checking if IBM Systems Director Common Agent is installed on the managed system

```
lslpp -l DirectorCommonAgent
  Fileset                      Level  State      Description
  ----------------------------------------------------------------------------
Path: /usr/lib/objrepos
  DirectorCommonAgent        6.2.1.3  COMMITTED  All required files of Director
                                                 Common Agent, including JRE,
                                                 LWI

Path: /etc/objrepos
  DirectorCommonAgent        6.2.1.3  COMMITTED  All required files of Director
                                                 Common Agent, including JRE,
                                                 LWI
```

2. Is the managed system discovered and fully inventoried in IBM Systems Director?
3. Is the WPAR Manager Agent installed on the managed systems?

Use the **lslpp** command as shown in Example 12-9 to check if WPAR Manager Agent is installed on the server.

Example 12-9 Check if WPAR Manager Agent is installed on the managed system

```
lslpp -l wparmgt.agent.rte
  Fileset                      Level  State      Description
  ----------------------------------------------------------------------------
Path: /usr/lib/objrepos
  wparmgt.agent.rte          2.3.1.1  COMMITTED  Workload Partitions Manager
                                                 Agent

Path: /etc/objrepos
  wparmgt.agent.rte          2.3.1.1  COMMITTED  Workload Partitions Manager
                                                 Agent
```

4. Have you run an inventory with the **Extended WPAR Inventory** profile enabled on the managed system, as shown in Figure 12-10.

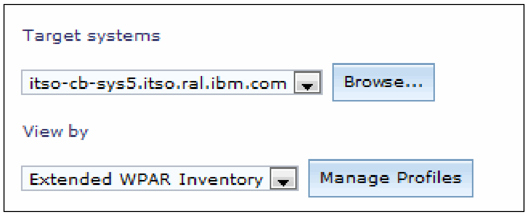
Target systems

itso-cb-sys5.itso.ral.ibm.com

View by

Extended WPAR Inventory

Figure 12-10 Run the inventory with Extended WPAR Inventory profile

12.5.2 Remove WPAR from Director

In some cases, you might still see a WPAR on a managed system in IBM Systems Director after it has been removed by using the **smcli rmwpar** command. To definitely remove the WPAR from IBM Systems Director server, go to the IBM Systems Director console in WPAR Manager Plug-ins and then click **View workload partitions**, as shown in Figure 12-11 on page 511.

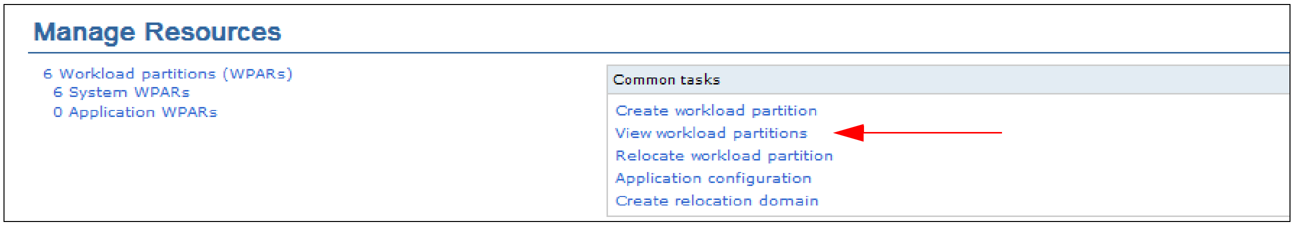

Figure 12-11 Remove WPAR from IBM Systems Director

Select the WPAR. Then, click **Actions** → **Remove** and click **OK**.

Go back to the command line and run the **smcli lswpar** command. It should produce the following result, as shown in Example 12-10.

Example 12-10 List WPAR with CLI

```
smcli lswpar
DNZWML806E No valid WPAR target(s) found.
```

12.5.3 Creating WPARs

Creating a WPAR by using an IP that is already used on another WPAR in the same managed system is not possible. The creation task will fail. If the creation is done on another managed system, the task will be successful but you end up with two WPARs sharing the same IP.

12.5.4 WPAR logs

All WPAR operation logs can be found on managed systems in the `/var/adm/wpars/event.log` file, as shown in Example 12-11.

Example 12-11 Event.log

```
cat /var/adm/wpars/event.log
V 2013-04-26 09:38:52 7667912 startwpar - COMMAND START, ARGS: wpar3test1
V 2013-04-26 09:38:53 9961476 ckwpar wpar3test1 ckwpar: Passed.
I 2013-04-26 09:38:53 7667912 startwpar wpar3test1 Lock acquired.
I 2013-04-26 09:38:53 7667912 startwpar wpar3test1 Starting workload partition
'wpar3test1'.
I 2013-04-26 09:38:53 7667912 startwpar wpar3test1 Mounting all workload partition
file systems.
I 2013-04-26 09:38:54 7667912 startwpar wpar3test1 Loading workload partition.
I 2013-04-26 09:38:54 7667912 startwpar wpar3test1 Exporting workload partition
devices.
I 2013-04-26 09:38:55 7667912 startwpar wpar3test1 Exporting device /dev/null, 1
I 2013-04-26 09:38:55 7667912 startwpar wpar3test1 Exporting device /dev/tty, 1
I 2013-04-26 09:38:55 7667912 startwpar wpar3test1 Exporting device /dev/console,
1
I 2013-04-26 09:38:55 7667912 startwpar wpar3test1 Exporting device /dev/zero, 1
I 2013-04-26 09:38:55 7667912 startwpar wpar3test1 Exporting device /dev/clone, 1
I 2013-04-26 09:38:55 7667912 startwpar wpar3test1 Exporting device /dev/sad, 3
I 2013-04-26 09:38:55 7667912 startwpar wpar3test1 Exporting device /dev/xti/tcp,
3
I 2013-04-26 09:38:55 7667912 startwpar wpar3test1 Exporting device /dev/xti/tcp6,
3
```

```
I 2013-04-26 09:38:55 7667912 startwpar wpar3test1 Exporting device /dev/xti/udp,
3
I 2013-04-26 09:38:55 7667912 startwpar wpar3test1 Exporting device /dev/xti/udp6,
3
I 2013-04-26 09:38:55 7667912 startwpar wpar3test1 Exporting device
/dev/xti/unixdg, 3
I 2013-04-26 09:38:55 7667912 startwpar wpar3test1 Exporting device
/dev/xti/unixst, 3
I 2013-04-26 09:38:55 7667912 startwpar wpar3test1 Exporting device /dev/error, 1
I 2013-04-26 09:38:55 7667912 startwpar wpar3test1 Exporting device /dev/errorctl,
1
I 2013-04-26 09:38:55 7667912 startwpar wpar3test1 Exporting device /dev/audit, 1
I 2013-04-26 09:38:55 7667912 startwpar wpar3test1 Exporting device /dev/nvram, 1
I 2013-04-26 09:38:55 7667912 startwpar wpar3test1 Exporting device /dev/kmem, 1
I 2013-04-26 09:38:55 7667912 startwpar wpar3test1 Exporting workload partition
kernel extensions.
W 2013-04-26 09:38:55 9765080 lswpar - lswpar: 0960-679 wpar3test1 has no kernel
extension configuration.
I 2013-04-26 09:38:55 7667912 startwpar wpar3test1 Starting workload partition
subsystem 'cor_wpar3test1'.
I 2013-04-26 09:38:55 7667912 startwpar wpar3test1 Verifying workload partition
startup.
V 2013-04-26 09:38:56 12124344 runwpar - COMMAND START, ARGS: wpar3test1
E 2013-04-26 09:38:56 12124344 runwpar wpar3test1 runwpar: 0960-246 Workload
partition 'wpar3test1' is currently locked by startwpar (PID = 7667912).
I 2013-04-26 09:38:56 12124344 runwpar wpar3test1 Starting workload partition
init.
I 2013-04-26 09:38:56 12124344 runwpar wpar3test1 Removing work directory
/tmp/.workdir.9764952.12124344_1
I 2013-04-26 09:38:56 7667912 startwpar wpar3test1 Lock released.
I 2013-04-26 09:38:56 7667912 startwpar wpar3test1 Removing work directory
/tmp/.workdir.9961680.7667912_1
V 2013-04-26 09:38:56 7667912 startwpar wpar3test1 Return Status = SUCCESS.
```

12.6 Troubleshooting Storage Control

This section includes the basic troubleshooting steps that should be followed for IBM Systems Director Storage Management.

12.6.1 Runtime log locations

The IBM Systems Director runtime storage management logs can be found in the following specified locations:

- On Microsoft Windows:
  - `<IBM Systems Director installation location>\director\lwi\logs\`
  - `C:\Program Files\IBM\TPC\device\log\`
- On Linux/AIX:
  - `/opt/ibm/director/lwi/logs/`
  - `/opt/IBM/TPC/device/log/`

IBM Systems Director Storage Control logs are also included as part of log collection when a user runs the following command:

- On Microsoft Windows:

  ```
  <IBM Systems Director installation location>\Director\bin\logcollector.bat
  ```

- On Linux/AIX:

  ```
  <IBM Systems Director installation location>/director/bin/logcollector.sh
  ```

12.6.2 Troubleshooting

Troubleshooting steps for various storage management-related issues are documented in the following IBM Systems Director information center link:

http://pic.dhe.ibm.com/infocenter/director/pubs/index.jsp?topic=%2Fcom.ibm.director.tbs.helps.doc%2Ffqm0_r_tbs_sm_storage_troubleshooting.html

Troubleshooting the basic tasks speeds up the configuration time that is taken. Following are some of the storage management issues related to discovery and inventory tasks:

- Unable to remove storage device by using the **rmdatasource** command

  Description:

  There is a rare case where a user might try to remove a data source from their Storage Control server using the **rmdatasource** command or the remove data source function in the GUI, but the data source will not be removed. The commands return with a success message even though the data source is still there if the client lists them using the **smcli lsdatasource** command.

  Resolution:

  The reason this is happening is due to a timeout issue in the TPC database when it is attempting to remove the device from the database. The TPC server has a default timeout value but it might not be sufficient in some cases. There is not a fix. The default timeout value is a permanent restriction.

- Unable to discover Storage Management Initiative Specification (SMI-S) Provider Agent that is installed on a system with multiple IP addresses

  Description:

  There might be an issue in discovering the SMI-S Provider Agent system with dual IPs configured.

  Resolution:

  – If the network of the system is configured with dual IPs and the SMI-S provider is installed on the dual IP systems, the provider will be listening to both of the dual IPs. Under this situation, running the discovery profile to discover the SMI-S protocol for one of the IPs will not discover the DSxxxx storages.

  – A user must run the Discovery profile for both of the dual IPs so that SMI-S protocol will be discovered for both the IPs to discover the DSxxxx storages in IBM Systems Director.

- Inventory on a V7000/SAN Volume Controller (SVC) might fail after performing a Tier 3 or 4 recovery procedure on the storage

  Description:

  After the V7000/SVC gets managed in IBM Systems Director, the inventory on V7000/SVC fails after a Tier 3 or 4 recovery procedure is performed on the storage.

Resolution:

i. Run the following command to remove the storage from the IBM Systems Director server command line:

```
smcli rmdatasource -c svc -i <storageIP>
```

ii. Run the following command to add the storage back into IBM Systems Director server:

```
smcli mkdatasource -c svc -i <storageIP> -f <sshKeyPath> -r
<keyPassphrase>
```

- Cannot rediscover a Storage Farm object after the farm was removed manually

  Description:

  If the Storage Control Farm object is not removed completely or if the process of removing it is interrupted, the farm resource still exists in the IBM Systems Director internal database and the rediscovery task cannot create a new Storage Control Farm object.

  Resolution:

  The resolution is to remove the Farm resource instance and then rediscover the Farm MEP.

  On the IBM Systems Director server command line, run the following commands:

  i. Locate the Globally Unique Identifier (GUID) value from the output of the following command:
  `smcli lsresource farm`

  ii. Run the following command to remove the Storage Control Farm object:

  ```
  smcli rmresource (GUID of the farm obtained from Step 3 above)
  ```

  iii. Change directory to: <IBM Systems Director installation path>`/StorageControl/bin` and run the following command, which discovers the Storage Control Farm object:

  ```
  SCDiscoverUnlock.sh
  ```

- Inventory collection on Storage Control Farm object fails

  Description:

  When a user performs an inventory collection on the Storage Control Farm object, it might complete with errors.

  Resolution:

  The inventory collection issue can be due to one of the following reasons:

  - Caused by the wrong data stored in DB2 by Storage Control.
  - Caused due to the wrong credentials. This can result in the inappropriate access state occurring to the Storage Control Farm object.

  To fix this problem, complete the following steps:

  i. From the IBM Systems Director command line, change the directory to: <IBM Systems Director installation path>`/StorageControl/bin` and run the following command:

  ```
  SCDiscoverUnlock.sh
  ```

  ii. Collect the inventory on the associated Storage Control Farm object from the IBM Systems Director web console or by running the following command:

  ```
  smcli collectinv -n <name of the farm object> -p “All Inventory”
  ```

- Inventory collection on Storage Control Farm object fails

  Description:

  If the Storage Control Farm object is managing an IBM Storwise V7000 and some of the disk drives in V7000 show as offline and have enclosure IDs and slot IDs missing, inventory on the Storage Control Farm object fails.

  Resolution:

  To resolve the inventory failed issue, perform the following steps:

  i. Recover from the offline physical disk from V7000 (follow the fix procedure on V7000 to fix the disk offline error).

  ii. Remove the v7000 storage MEP from the IBM Systems Director server. See "Removing Storage Control-managed devices" on page 487 to get more information about how to remove storage devices.

  iii. Add the v7000 storage by using the following command:

  ```
  smcli mkdatasource -c svc -i <IP address of the v7000> -f <path to
  generated private SSH Key on systems director server> -v V7000
  ```

  iv. Collect the inventory on the associated Storage Control Farm object from the IBM Systems Director web console or by running the following command:

  ```
  smcli collectinv -n <name of the farm object> -p “All Inventory”
  ```

A

Technical articles

This appendix provides information about sources of information for the IBM Systems Director. This includes the IBM Systems Director Information Center, social media references such as YouTube, and information about the IBM Systems Director forum and wiki.

There are also links provided for downloading IBM Systems Director components and other additional, useful links.

A.1 Information center

One of the first sources to go to for information is the IBM Systems Director Information Center. The Systems Director Information Center is an online and regularly updated version of the Systems Director product publications. For the latest version of Systems Director, see the following website:

http://pic.dhe.ibm.com/infocenter/director/pubs/index.jsp

If you use older versions of Systems Director, see the URLs listed in Table A-1.

Table A-1 Information center links

IBM Systems Director version	Information center
6.2.*x*	Systems Director version 6.2.*x* Information Center: http://publib.boulder.ibm.com/infocenter/director/v6r2x/index.jsp
6.1.*x*	Systems Director version 6.1.*x* Information Center: http://publib.boulder.ibm.com/infocenter/director/v6r1x/index.jsp
5.2.*x*	IBM Director version 5.20.*x* Information Center: http://publib.boulder.ibm.com/infocenter/eserver/v1r2/topic/diricinfo_5.20/fqm0_main.html

In the information centers, you can search for information about the installation, configuration, management, and problem determination. You can also download the PDF versions of the Installation Guide, Planning Guide, Troubleshooting Guide, and Systems Management Guide.

A.2 Social media and support

Other users can be a great source of help with IBM Systems Director. You can connect to them over the IBM forum, the wiki, or see information and comments on YouTube and Facebook. There are also information messages available through the My Notification function.

A.3 Forum

The IBM Systems Director Forum is available at the following link:

http://www.ibm.com/developerworks/forums/forum.jspa?forumID=759

This forum provides a place for all IBM Systems Director topics. You can post your questions and comments and share your thoughts, ideas, and solutions with other users.

The forum also offers an RSS feed. Click the orange RSS icon to access the RSS subscription page. You can then select the method to use to subscribe to the feed and click **Subscribe Now** (Figure A-1).

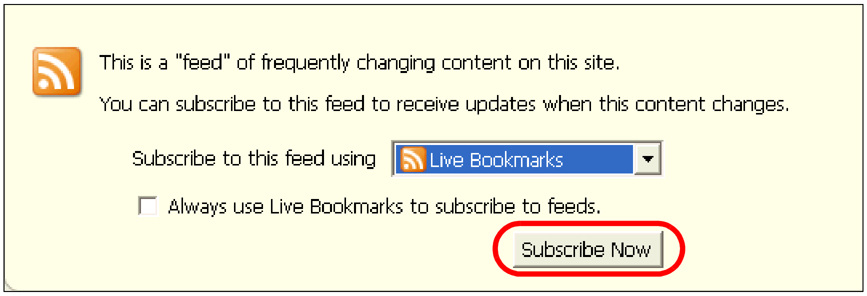

This is a "feed" of frequently changing content on this site.

You can subscribe to this feed to receive updates when this content changes.

Subscribe to this feed using: Live Bookmarks

□ Always use Live Bookmarks to subscribe to feeds.

Figure A-1 Systems Director Forum RSS feeds

A.4 Wiki

There are two wikis for IBM Systems Director. Both wikis are focused on IBM Systems Director running on a Power Systems platform. However, some of the information applies to all platforms:

- IBM Systems Director Wiki:

  `https://www.ibm.com/developerworks/mydeveloperworks/wikis/home?lang=en#/wiki/Power%20Systems/page/IBM%20Systems%20Director`

- IBM Systems Director Best Practices Wiki:

  `https://www.ibm.com/developerworks/community/wikis/home?lang=en#!/wiki/W3e8d1c956c32_416f_a604_4633cd375569/page/Best%20Practices`

A.5 YouTube channel

To subscribe to this channel, go to `http://www.youtube.com/user/IBMSystemsDirector/feed` (Figure A-2) and click **Subscribe**.

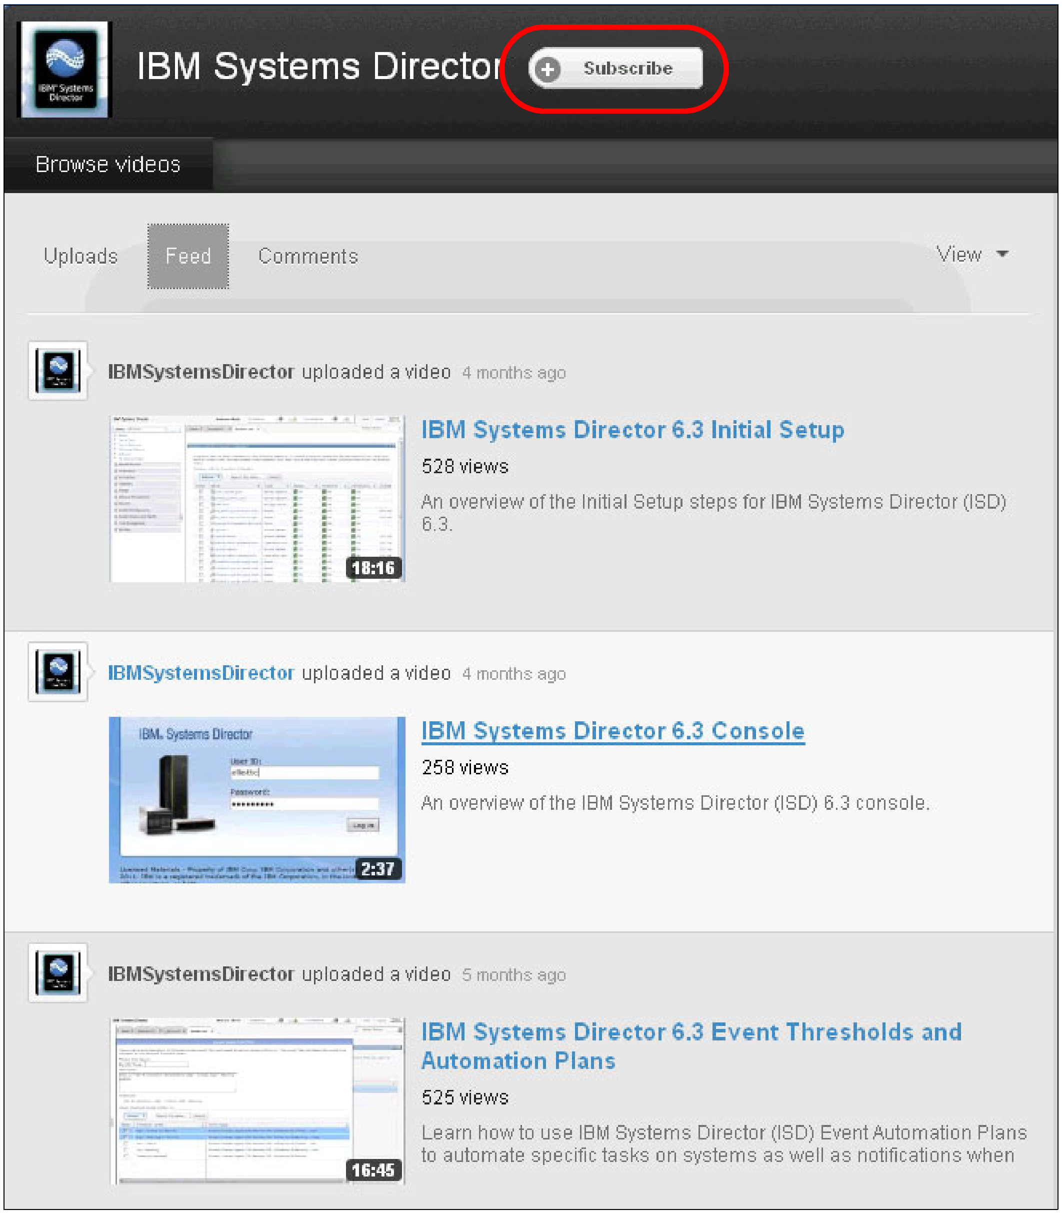

IBMSystemsDirector uploaded a video 4 months ago

IBM Systems Director 6.3 Initial Setup

528 views

An overview of the Initial Setup steps for IBM Systems Director (ISD) 6.3.

IBMSystemsDirector uploaded a video 4 months ago

IBM Systems Director 6.3 Console

258 views

An overview of the IBM Systems Director (ISD) 6.3 console.

IBMSystemsDirector uploaded a video 5 months ago

IBM Systems Director 6.3 Event Thresholds and Automation Plans

525 views

Learn how to use IBM Systems Director (ISD) Event Automation Plans to automate specific tasks on systems as well as notifications when

Figure A-2 Systems Director at YouTube

A.6 My Notifications email announcements

With My Notifications, you can subscribe to support updates for any IBM product.

Tip: The My Notifications tool replaces My Support, a similar tool.

With My Notifications, you can specify that you want to receive daily or weekly email announcements. You can specify the type of information that you want to receive:

- Publications
- Hints and tips
- Product flashes (also known as *alerts*)
- Downloads
- Drivers

With My Notifications, you can customize and categorize the products about which you want to be informed and the delivery methods that best suit your needs.

Complete the following steps to subscribe to My Notifications:

1. Go to http://www.ibm.com/support/mynotifications.
2. Enter your IBM ID and password and click **Submit**.
3. Identify the updates that you want to receive and the method through which you want to receive them:
   a. Click the **Subscribe** tab.
   b. Select **IBM Systems Director**.
   c. Specify or select your notifications and other preferences.
   d. Click **Submit**.

A.7 Education and training

IBM offers various educational offerings, including instructor-led online (ILO) courses, classroom courses, virtual learning courses, private/on-site training, and web-based learning videos. Links are available in the IBM Systems Director console for online training.

A.8 Integrated education modules in IBM Systems Director

The Systems Director console provides links for online training. The links are on the Systems Director Home page under the **Learn** tab, as shown in Figure A-3.

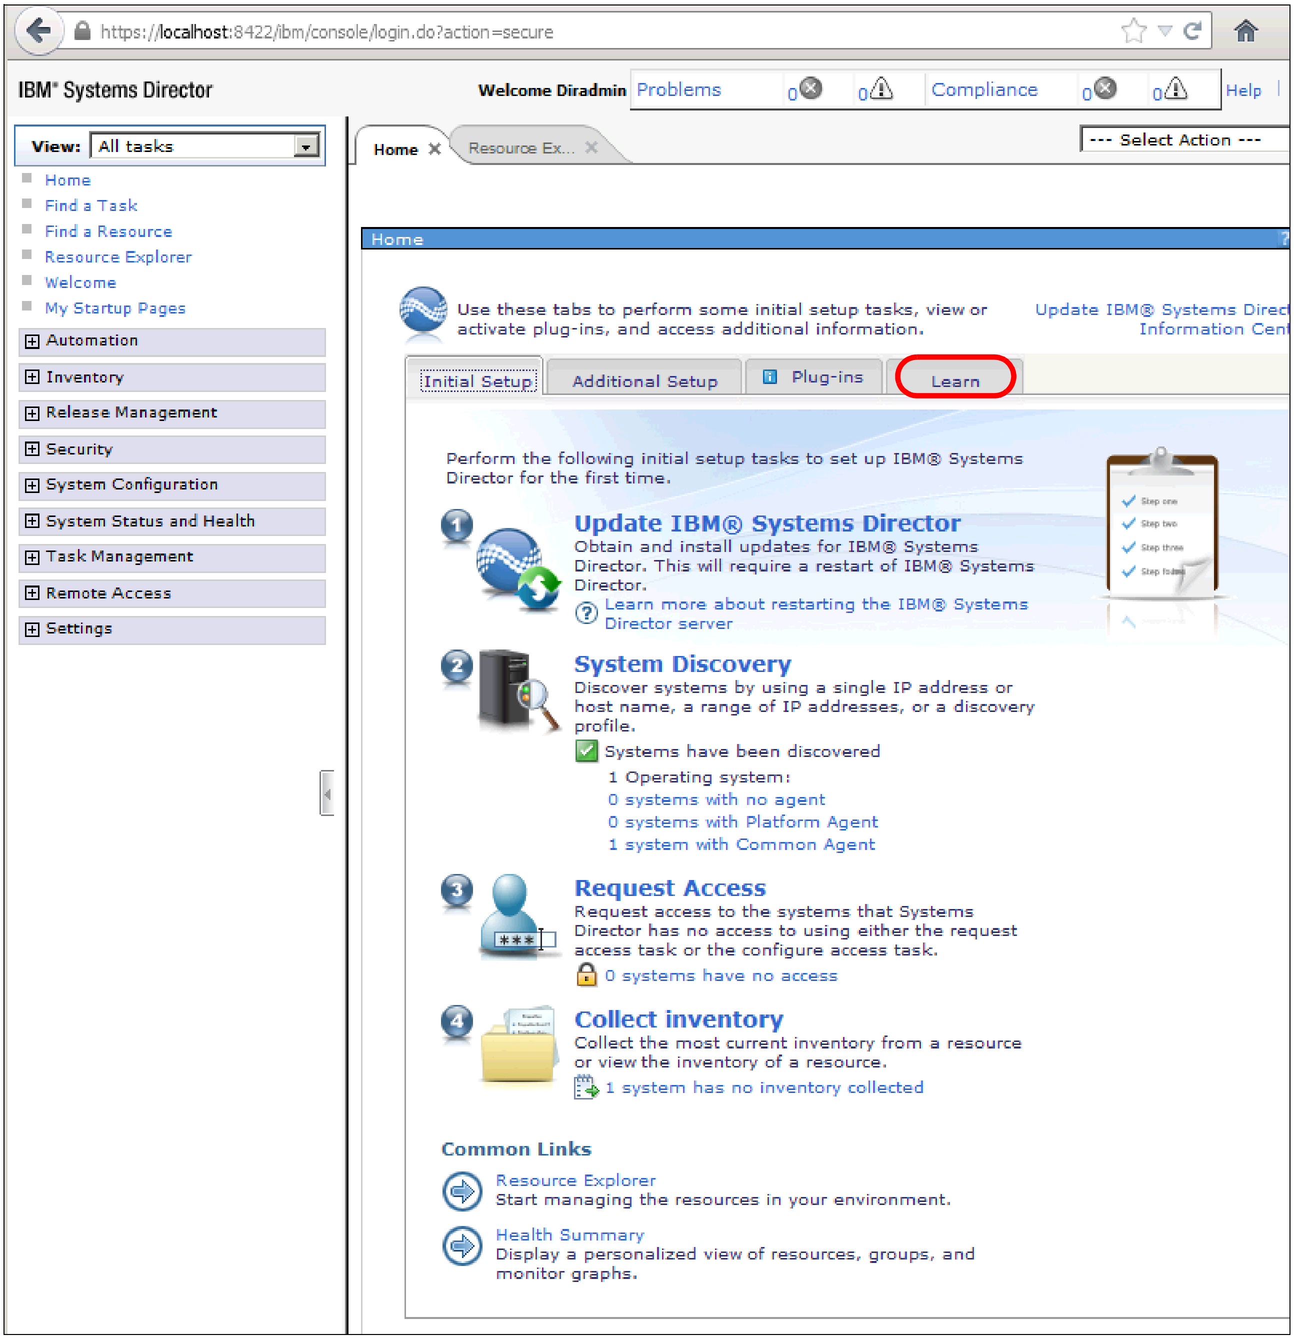

Use these tabs to perform some initial setup tasks, view or activate plug-ins, and access additional information.

Update IBM® Systems Direc
Information Cent

Perform the following initial setup tasks to set up IBM® Systems Director for the first time.

1 **Update IBM® Systems Director**
Obtain and install updates for IBM® Systems Director. This will require a restart of IBM® Systems Director.
Learn more about restarting the IBM® Systems Director server

2 **System Discovery**
Discover systems by using a single IP address or host name, a range of IP addresses, or a discovery profile.
✓ Systems have been discovered
1 Operating system:
0 systems with no agent
0 systems with Platform Agent
1 system with Common Agent

3 **Request Access**
Request access to the systems that Systems Director has no access to using either the request access task or the configure access task.
0 systems have no access

4 **Collect inventory**
Collect the most current inventory from a resource or view the inventory of a resource.
1 system has no inventory collected

Common Links

Resource Explorer
Start managing the resources in your environment.

Health Summary
Display a personalized view of resources, groups, and monitor graphs.

Figure A-3 IBM Systems Director console Home page

On the Learn page, you can see the available learning modules for Systems Director (Figure A-4).

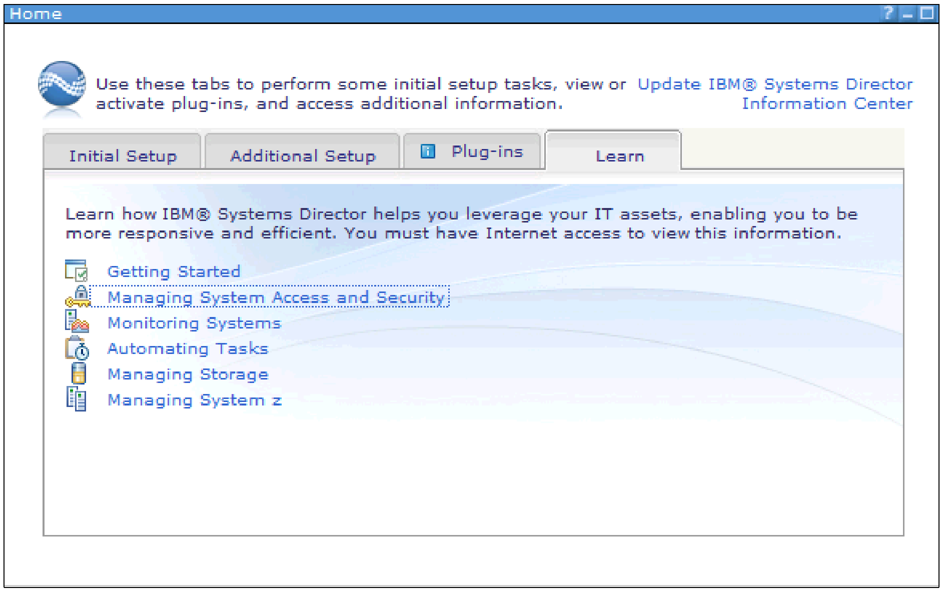

Use these tabs to perform some initial setup tasks, view or activate plug-ins, and access additional information.

Update IBM® Systems Director Information Center

Learn how IBM® Systems Director helps you leverage your IT assets, enabling you to be more responsive and efficient. You must have Internet access to view this information.

- Getting Started
- Managing System Access and Security
- Monitoring Systems
- Automating Tasks
- Managing Storage
- Managing System z

Figure A-4 Learn tab

When you click one of the topics, you link to the Systems Director Information Center website. Videos for the selected topic are displayed. In our example, we select the “Managing system access and security” topic (Figure A-5).

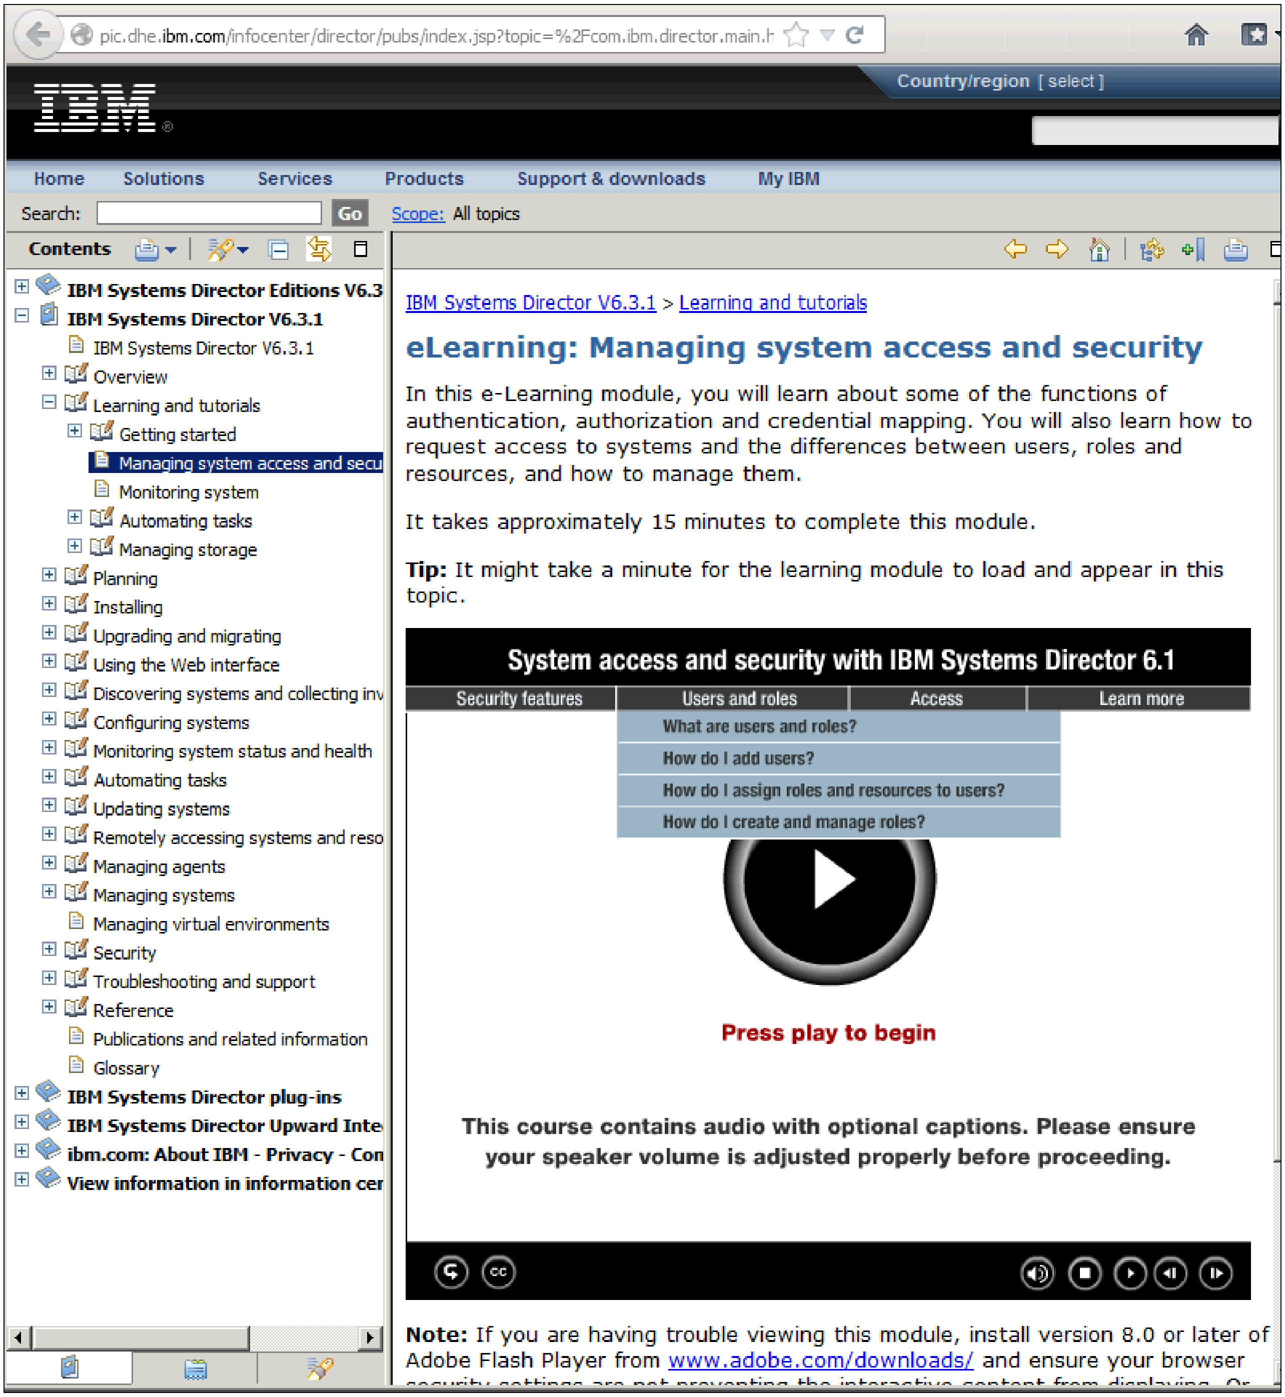
pic.dhe.ibm.com/infocenter/director/pubs/index.jsp?topic=%2Fcom.ibm.director.main.h
IBM
Country/region [select]
Home
Solutions
Services
Products
Support & downloads
My IBM
Search:
Go
Scope: All topics
Contents
IBM Systems Director Editions V6.3
IBM Systems Director V6.3.1
IBM Systems Director V6.3.1
Overview
Learning and tutorials
Getting started
Managing system access and secu
Monitoring system
Automating tasks
Managing storage
Planning
Installing
Upgrading and migrating
Using the Web interface
Discovering systems and collecting inv
Configuring systems
Monitoring system status and health
Automating tasks
Updating systems
Remotely accessing systems and reso
Managing agents
Managing systems
Managing virtual environments
Security
Troubleshooting and support
Reference
Publications and related information
Glossary
IBM Systems Director plug-ins
IBM Systems Director Upward Inte
ibm.com: About IBM - Privacy - Con
View information in information cer
IBM Systems Director V6.3.1 > Learning and tutorials

eLearning: Managing system access and security

In this e-Learning module, you will learn about some of the functions of authentication, authorization and credential mapping. You will also learn how to request access to systems and the differences between users, roles and resources, and how to manage them.

It takes approximately 15 minutes to complete this module.

Tip: It might take a minute for the learning module to load and appear in this topic.

System access and security with IBM Systems Director 6.1
Security features
Users and roles
Access
Learn more
What are users and roles?
How do I add users?
How do I assign roles and resources to users?
How do I create and manage roles?
Press play to begin

This course contains audio with optional captions. Please ensure your speaker volume is adjusted properly before proceeding.

Note: If you are having trouble viewing this module, install version 8.0 or later of Adobe Flash Player from www.adobe.com/downloads/ and ensure your browser

Figure A-5 Managing system access and security module

You can also access the training modules at the Systems Director Information Center through the following link:

http://pic.dhe.ibm.com/infocenter/director/pubs/index.jsp?topic=%2Fcom.ibm.director.main.helps.doc%2Ffqm0_c_elearning.html

A.9 Education courses

The following classroom (XTR) courses and ILO (XTV) courses for IBM Systems Director 6.3 are available in the US at the time of writing:

- XTR/XTV 42: IBM Systems Director 6.3 Hands-on Workshop
- XTR/XTV 46: IBM Systems Director 6.3 - Introduction
- XTR/XTV 47: IBM Systems Director 6.3 for IBM System x and BladeCenter Servers - Base
- XTR/XTV 48: IBM Systems Director 6.3 for IBM System x and BladeCenter Servers - Advanced
- XTRD1+2/XTVD1+2: IBM Systems Director 6.3 for IBM System x and BladeCenter Servers
- AN940/AX940: IBM Systems Director 6.3 for Power Systems I: Installation and Management

A self-paced virtual class (SPVC) is available:

- AN0D0/XTRD0: IBM Systems Director 6.3 - Power and System x - Planning and Installation

The availability of onsite or classroom training depends on your country. For detailed information about the offerings in your country, see the following link:

http://www-304.ibm.com/jct03001c/services/learning/ites.wss/zz/en?pageType=page&c=a0011023

Complete the following steps to check the availability of learning courses in your country:

1. Go to the following link:

   http://www-304.ibm.com/jct03001c/services/learning/ites.wss/zz/en?pageType=page&c=a0011023

2. Select your country and click the arrow.
3. Scroll to the bottom of the page and click **Search for training courses**.
4. Type `IBM Systems Director 6.3` in the search field and click **Search**.
5. The available courses are displayed.
6. Select your course.

A.10 Downloads

All downloads for Systems Director, including server, agents, and plug-ins or advanced managers, are at the following link:

http://ibm.com/systems/software/director/downloads

On the Downloads overview page, you can view information about the recent product updates that you can download (Figure A-6 on page 526). By selecting the tabs for Management servers, Agents, Plug-ins, and Partner integration, you can access the pages for download.

You need an IBM ID to download the codes. If you do not have an IBM ID, you can request one on the website where the ID is required. The IBM ID is available at no charge.

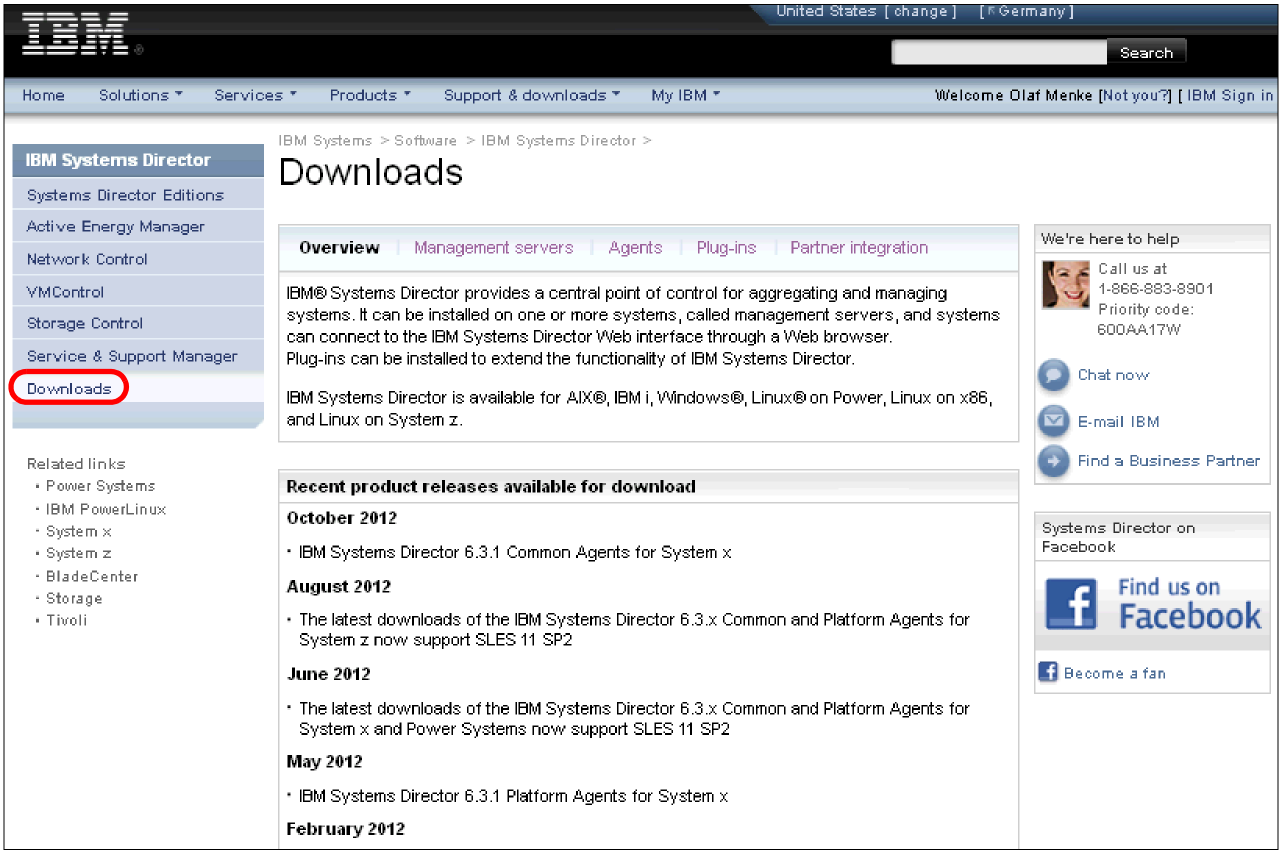

Downloads

IBM® Systems Director provides a central point of control for aggregating and managing systems. It can be installed on one or more systems, called management servers, and systems can connect to the IBM Systems Director Web interface through a Web browser. Plug-ins can be installed to extend the functionality of IBM Systems Director.

IBM Systems Director is available for AIX®, IBM i, Windows®, Linux® on Power, Linux on x86, and Linux on System z.

Recent product releases available for download

October 2012

- IBM Systems Director 6.3.1 Common Agents for System x

August 2012

- The latest downloads of the IBM Systems Director 6.3.x Common and Platform Agents for System z now support SLES 11 SP2

June 2012

- The latest downloads of the IBM Systems Director 6.3.x Common and Platform Agents for System x and Power Systems now support SLES 11 SP2

May 2012

- IBM Systems Director 6.3.1 Platform Agents for System x

February 2012

Figure A-6 Downloads for Systems Director

A.11 Other useful links

The following references are useful:

IBM websites about System x, BladeCenter, and Flex Systems:

- xREF: IBM x86 Server Reference:

  http://www.redbooks.ibm.com/xref
- IBM Configuration and Options Guide (COG):

  http://ibm.com/support/entry/portal/docdisplay?lndocid=SCOD-3ZVQ5W
- BladeCenter Interoperability Guide (BIG):

  http://ibm.com/support/entry/portal/docdisplay?lndocid=MIGR-5073016
- IBM Flex System Interoperability Guide:

  http://www.redbooks.ibm.com/fsig
- IBM ToolsCenter (ServerGuide, Bootable Media Creator, Advanced Settings Utility):

  http://ibm.com/support/entry/portal/docdisplay?lndocid=TOOL-CENTER

IBM Information Center:

- IBM Systems Director Information Center:

  http://pic.dhe.ibm.com/infocenter/director/pubs/index.jsp

- IBM BladeCenter Information Center:

  http://publib.boulder.ibm.com/infocenter/bladectr/documentation/index.jsp

- IBM PureFlex System Information Center:

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

- IBM Power Systems Hardware Information Center (including IBM System p®, IBM System i, and Hardware Management Console information):

  http://pic.dhe.ibm.com/infocenter/powersys/v3r1m5/index.jsp?topic=/ipha8/hwicwelcome.htm

External website:

- IBM Quicklinks (link collection for IBM products):

  http://www.ibmquicklinks.com

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 Systems Director 6.3 Best Practices: Installation and Configuration,* REDP-4932-00
- *IBM Systems Director Management Console: Introduction and Overview,* SG24-7860-00
- *IBM CSM to IBM Systems Director Transformation Guide,* SG24-8002-00

You can search for, view, download or order these documents and other Redbooks, Redpapers, Web Docs, drafts, and additional materials, at the following website:

ibm.com/redbooks

Online resources

These websites are also relevant as further information sources:

- IBM Systems Workload Estimator for Systems Director 6.3:

  http://pic.dhe.ibm.com/infocenter/director/pubs/index.jsp?topic=%2Fcom.ibm.director.plan.helps.doc%2Fwle.html

- Hardware requirements for IBM Systems Director:

  http://pic.dhe.ibm.com/infocenter/director/pubs/topic/com.ibm.director.plan.helps.doc/fqm0_r_hardware_requirements.html

- Supported operating systems:

  http://pic.dhe.ibm.com/infocenter/director/pubs/index.jsp?topic=%2Fcom.ibm.director.plan.helps.doc%2Ffqm0_r_supported_operating_systems.html

- Security features and considerations are documented in the information center at:

  http://pic.dhe.ibm.com/infocenter/director/pubs/index.jsp?topic=%2Fcom.ibm.director.security.helps.doc%2Ffqm0_c_security.html

- File system requirements that are needed for the installation are documented at this site:

  http://pic.dhe.ibm.com/infocenter/director/pubs/index.jsp?topic=%2Fcom.ibm.director.plan.helps.doc%2Ffqm0_r_hardware_requirements_servers_running_aix.html

- A list of all TCP/IP ports that are used by IBM Systems Director is provided at this site:

  http://pic.dhe.ibm.com/infocenter/director/pubs/index.jsp?topic=%2Fcom.ibm.director.plan.helps.doc%2Ffqm0_r_all_available_ports.html

- You can download lsof as part of the AIX Expansion Pack from this link:

  http://www-03.ibm.com/systems/power/software/aix/expansionpack/index.html

- Information about ports used by the server, managed systems, and important port considerations:

  http://pic.dhe.ibm.com/infocenter/director/pubs/index.jsp?topic=%2Fcom.ibm.director.plan.helps.doc%2Ffqm0_r_all_available_ports.html

- IBM Tivoli Provisioning Manager for OS Deployment and IBM Tivoli Provisioning Manager for Images, Version 7.1.1 information:

  http://pic.dhe.ibm.com/infocenter/tivihelp/v3r1/index.jsp?topic=/com.ibm.tivoli.tpm.osd.doc/welcome/osdhome.html

- Preparing agentless managed systems:

  http://pic.dhe.ibm.com/infocenter/director/pubs/index.jsp?topic=%2Fcom.ibm.director.install.helps.doc%2Ffqm0_t_preparing_agentless_managed_systems.html

- Configuring access to agentless managed systems:

  http://pic.dhe.ibm.com/infocenter/director/pubs/topic/com.ibm.director.install.helps.doc/fqm0_t_setting_up_access_to_agentless_systems.html

- Platform Agent systems:

  http://pic.dhe.ibm.com/infocenter/director/pubs/index.jsp?topic=%2Fcom.ibm.director.main.helps.doc%2Ffqm0_c_platform_agent.html

- Common Agent systems:

  http://pic.dhe.ibm.com/infocenter/director/pubs/index.jsp?topic=%2Fcom.ibm.director.main.helps.doc%2Ffqm0_c_common_agent.html

- Coexistence of Director V6 CAS Agent with other Tivoli CAS Agents:

  https://www.ibm.com/developerworks/community/wikis/home?lang=en#!/wiki/W3e8d1c956c32_416f_a604_4633cd375569/page/Coexistence+of+Director+V6+CAS+Agent+with+other+Tivoli+CAS+Agents

- Obtaining licenses for Common Agent:

  http://pic.dhe.ibm.com/infocenter/director/pubs/topic/com.ibm.director.install.helps.doc/fqm0_t_obtaining_licenses.html

- Choosing the level of agent capabilities to deploy on managed systems:

  http://pic.dhe.ibm.com/infocenter/director/pubs/index.jsp?topic=%2Fcom.ibm.director.plan.helps.doc%2Ffqm0_t_agent_tiers.html

- Discovering systems that use a mirrored image:

  http://pic.dhe.ibm.com/infocenter/director/pubs/index.jsp?topic=%2Fcom.ibm.director.discovery.helps.doc%2Ffqm0_t_discovering_systems_mirrored_image.html

- Choosing the IBM Systems Director database application:

  http://pic.dhe.ibm.com/infocenter/director/pubs/topic/com.ibm.director.plan.helps.doc/fqm0_t_selecting_the_ibm_director_database_application.html

- Binding Platform Agent to specific IP addresses:

  http://pic.dhe.ibm.com/infocenter/director/pubs/topic/com.ibm.director.agent.helps.doc/fqm0_t_binding_pa_specific_ip_addresses.html

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services

IBM Systems Director 6.3 Best Practices

(1.0” spine)
0.875”<->1.498”
460 <-> 788 pages

IBM Systems Director 6.3 Best Practices

Provides additional guidance beyond standard instructions

Covers basic and advanced features

Learn from the experts in the field

This IBM Redbooks publication describes the positioning of the IBM Systems Director in the complete management range. It also compares the IBM Systems Director with the IBM Flex Systems Manager (FSM) and describes the environments for which each tool is best suited.

This publication helps you plan, install, tailor, and configure the IBM Systems Director on different platforms. It contains information about required system resources and which network ports are used. It shows how to use the Workload Estimator to select the appropriate hardware for IBM Systems Director server and provides information about the IBM Systems Director Editions.

Best practices are covered for the basic management tasks that are available in IBM Systems Director, including how to perform discovery; how to collect inventory on discovered resources; how to deploy agent, driver, and firmware updates; how to manage hardware events; and other miscellaneous tasks.

An overview of best practices is provided for using IBM Systems Director VMControl. Systems Director VMControl is a cross-platform product that assists you in rapidly deploying virtual appliances to create virtual servers that are configured with the operating system and software applications that you want. It also enables you to group resources into system pools, which enable you to centrally manage and control the different workloads in your environment.

The following plug-in offerings are described:

- Energy monitoring and management features offered by IBM Systems Director Active Energy Manager along with the best practice, which needs to be followed in using the IBM Systems Director Active Energy Manager.
- The IBM AIX Profile Manager is a tool that can help implement and monitor the security of all AIX servers in a production environment but also implement and monitor the system compliance of those AIX servers.
- Best practices and the most important questions to ask before creating Workload Partition Manager (WPAR) and WPAR Manager infrastructure. In addition, how you can manage and relocate WPARs using WPAR Manager graphical interface and the command-line interface.
- Network Control basic functionalities and how to plan for Network Control deployments and also a number of common scenarios with best practices.
- The IBM Systems Director Service and Support Manager describes how to set up and how to handle serviceable events.
- Best practices for the Storage Monitoring and Management capabilities offered by IBM Systems Director server.

This publication is for IBM IT specialists and IT architects, IBM Business Partners, and clients, who are utilizing or considering implementing the IBM Systems Director.

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

SG24-8141-00 ISBN 0738438650