


IBM CSM to IBM Systems Director Transformation Guide

Provides planning guidelines

Includes features comparison

Discusses transformation scenarios



Dino Quintero
Glen Corneau
Antonio Garcia
Vinicius Goncalves
Tejaswini Kaujalgi Athanikar
Piotr Mazur
Laszlo Niesz
Srinivasa Raghavan
Christian Schmidt
Gheorghe Tudor

Redbooks



International Technical Support Organization

**IBM CSM to IBM Systems Director Transformation
Guide**

August 2012

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

First Edition (August 2012)

This edition applies to AIX 5.3.12 or 5.3 TL12, AIX 6.1.6 or 6.1 TL6, AIX 7.1.0, CSM 1.7.1, IBM Systems Director 6.2.1, VMControl V2.3.1, HMC V7.7.2.

© Copyright International Business Machines Corporation 2012. All rights reserved.

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

Contents

Notices	vii
Trademarks	viii
 Preface	 ix
The team who wrote this book	ix
Now you can become a published author, too!	xi
Comments welcome	xi
Stay connected to IBM Redbooks	xi
 Chapter 1. Abstract	 1
1.1 Introduction to CSM	2
1.2 Introduction to IBM Systems Director	2
1.3 Scope of this book	3
1.3.1 What we cover	3
1.3.2 What we do not cover	4
1.4 Why transform from CSM to IBM Systems Director	4
 Chapter 2. Planning and preparation	 7
2.1 Comparison	8
2.2 Terminology	8
2.2.1 Components	8
2.2.2 Manageable system types	9
2.2.3 Base functions	9
2.2.4 Mapping terminology	10
2.3 Agent Manager considerations	11
2.4 Product prerequisites	12
2.4.1 Software	12
2.4.2 Hardware and performance	13
2.5 Licensing	15
2.5.1 Licensing requirement for IBM Cluster Systems Management	15
2.5.2 Licensing requirements for IBM Systems Director	15
2.6 Extendibility	16
2.6.1 Application programming interfaces	16
2.6.2 IBM Systems Director Advanced Managers (Plug-ins)	17
2.6.3 Command line	20
2.7 Features not directly available in IBM Systems Director	21
2.7.1 Configuration File Manager	21
2.7.2 Cluster-Ready Hardware Server (CRHS)	22
2.7.3 Managing node status information with csmstat	22
2.7.4 High availability management server	26
2.8 Features and functions not available in CSM	26
2.8.1 LPAR lifecycle management	26
2.8.2 Update Manager	28
2.9 Database	29
2.9.1 RSCT System Registry versus an actual DB product	29
2.10 Transformation as a project and future impacts	32
2.10.1 Transformation project	32
2.10.2 BAU after transformation to IBM Systems Director	36

Chapter 3. Transformation scenarios	39
3.1 Overview of the scenarios	40
3.1.1 What we do not cover	40
3.1.2 What we cover	41
3.1.3 Prerequisites	41
3.1.4 Recommendations	41
3.1.5 Considerations and time estimates	41
3.2 Coexistence	42
3.2.1 The cluster environment	42
3.2.2 Cluster verification	44
3.2.3 Transformation to IBM Systems Director	46
3.2.4 Testing the coexistence scenario	54
3.3 Migration scenario	58
3.3.1 Initial environment	59
3.3.2 The target environment	61
3.3.3 Preparation	62
3.3.4 Effective migration	67
3.4 Decommissioning CSM	70
Chapter 4. Functional comparison	73
4.1 Monitoring of resources	74
4.2 Resource Monitoring and Control (RMC) versus IBM Systems Director Agent Services	74
4.2.1 CSM using Resource Monitoring and Control (RMC)	74
4.2.2 IBM Systems Director Common Agent and related resource managers	76
4.2.3 Comparison between various system resource monitors' output	79
4.2.4 Monitor migration considerations	85
4.3 Hardware and software management	85
4.3.1 Remote commands	85
4.3.2 System firmware updates	91
4.4 Security	95
4.4.1 Authentication and Authorization when using CSM	96
4.4.2 Authentication and authorization when using IBM Systems Director	98
4.4.3 Security consideration when defining users and groups	101
4.4.4 Roles	104
4.5 OS deployment	105
4.5.1 Comparison of installation services	106
4.5.2 Post-install customization	118
4.5.3 Deployment in work	122
4.6 Agents	124
4.6.1 Daemons, filesets, and protocols	127
4.6.2 Port conflicts - ITM?	128
Chapter 5. Special topics	131
5.1 SDMC	132
5.1.1 Overview of SDMC	132
5.1.2 Parallel management with the HMC and the SDMC	134
5.1.3 High availability of SDMC	135
5.1.4 Command Line for SDMC as compared to HMC	136
5.2 High Performance Computing	136
5.3 High availability IBM Systems Director management server	137
5.3.1 Additional considerations	138
5.3.2 Highly available IBM Systems Director management servers	138
5.3.3 Redundant IBM System Director management server options	138

5.3.4 Conclusion	139
5.4 Backup and restore of the management server	139
5.4.1 IBM Cluster Systems Management (CSM) backup and restore	139
5.4.2 IBM Systems Director backup and restore	142
5.4.3 smsave	143
5.4.4 smrestore	143
Appendix A. isdstat script	145
isdstat script sample	146
Related publications	149
IBM Redbooks	149
Other publications	149
Online resources	149
Help from IBM	150
Index	151

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 give 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.

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:

AIX 5L™	LoadLeveler®	Storwize®
AIX®	Power Systems™	System i®
BladeCenter®	POWER6®	System p5®
DB2®	POWER7®	System p®
DS8000®	PowerHA®	System Storage®
Electronic Service Agent™	PowerVM®	System x®
Global Technology Services®	POWER®	System z®
GPFS™	Redbooks®	Systems Director VMControl™
HACMP™	Redbooks (logo)  ®	Tivoli®
IBM®	RS/6000®	

The following terms are trademarks of other companies:

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.

Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium 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.

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

Preface

For many years, IBM® Cluster Systems Management (CSM) provided a single point of management for IBM Power Systems™ servers running the AIX® operating system. Now you can transform your environment to IBM Systems Director®, which provides CSM clients with the next generation of Cluster Systems Management for their Power Systems servers.

The target audience for this IBM Redbooks® publication includes technical professionals (IT consultants, technical support staff, IT Architects, and IT Specialists) responsible for planning and implementing the Cluster Systems Management software transformation from CSM to IBM Systems Director.

The team who wrote this book

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

Dino Quintero is a Project Leader and IT generalist with the ITSO in Poughkeepsie, NY. His areas of knowledge include enterprise continuous availability planning and implementation, enterprise systems management, virtualization, and clustering solutions. He is currently an Open Group Master Certified IT Specialist. Dino holds a Master of Computing Information Systems degree and a Bachelor of Science degree in Computer Science from Marist College.

Glen Corneau is a Level 2 Certified Consulting IT Specialist in the United States. He has 20 years of experience in AIX and Power Systems. He holds a degree in Computer Science from Texas A&M University in College Station, Texas. His areas of expertise include systems management on AIX clusters (with Cluster Systems Management and IBM Systems Director) and IBM General Parallel File System on AIX. He has presented extensively to clients in one-on-one, group, and IBM conferences on Power Systems clustering and management technologies.

Antonio Garcia is a Software Engineer in Austin, Texas, and has been with IBM for 12 years. He has experience in AIX performance and Linux performance tools development. Antonio currently works in the AIX Development Support Organization. He holds a Bachelor of Arts degree in Political Science and Psychology, a Bachelor of Science degree in Computer Science, and a Master of Science degree in Computer Science.

Vinicius Goncalves is an IT Specialist in Brazil. He has seven years of experience in Power Systems and AIX, including VIO/HACMP/CSM and GPFS™. Currently he focuses on providing cloud solutions and support for several Atos Brasil clients. His areas of expertise include virtualization, high availability, and networking. Vini holds a degree in Computer Science, as well as several senior level industry certifications.

Tejaswini Kaujalgi Athanikar is a System Software Engineer in Bangalore, India. She has four years of experience in AIX and Power Systems. Her areas of expertise include High Availability, Virtualization, and Security components such as LDAP and RBAC. Tejaswini has published an article in the Developer Works Forum, and has also contributed to IBM Redbooks. She holds a Bachelor's degree in Electronics and Communication from VTU, Belgaum, India, and is an IBM Certified System p® Administrator.


Piotr Mazur is an IT Specialist at IBM Global Technology Services® in Warsaw, Poland. He has six years of IBM experience in designing, implementing, and supporting solutions around IBM POWER® technology. He is a Technical Account Advocate supporting enterprise IT environments in Poland. Piotr holds a Master of Science degree in Information Technology from the Polish-Japanese Institute of Information Technology in Warsaw, as well as many technical expert certifications.

Laszlo Niesz is an IT specialist at IBM Hungary. He has 14 years of experience in AIX and clustering solutions on Power Systems. His areas of expertise include systems management with IBM cluster software products, virtualization on POWER systems, and IBM General Parallel File System. Laszlo is an architect on an international server consolidation project implementing and running several CSM clusters EMEA-wide. He holds a degree in Computer Programming from the University of Szeged, and has written extensively about operating system deployment.

Srinivasa Raghavan is an IBM Certified Consulting IT Specialist for the Global Delivery Center, Fishkill, NY. He has more than 16 years of experience in IBM support. He began his career in IBM India prior to moving to IBM US. His areas of expertise include UNIX operating systems, High Availability clusters, Virtualization, and project management. For the last three years, Srinivasa has worked as an Account Focal in the Global Delivery Center, supporting IBM internal and commercial clients. He holds a Bachelor's degree in Computer Science and Engineering, and is an IBM Certified Advanced Technical Expert in Power Systems.

Christian Schmidt is a Consulting IT Specialist at IBM GTS Services Delivery Nordic in Copenhagen, Denmark, where he works as an Infrastructure Architect for the Nordic Design Authority Group. He has more than 18 years of experience at IBM, primarily working with complex system infrastructure, systems design and implementation strategies, tuning and performance and system management. His areas of expertise include “anything that can run on IBM Power and AIX or Linux,” and keeping clients happy. He is especially interested in optimization, monitoring, infrastructure security and automation in virtualized environments. For the last three years he has written and spoken extensively about system management best practices and performance issues on large MetaSAN-attached virtualization farms. Christian has coauthored several IBM Redbooks on PSSP, CSM, PowerHA® (HACMP™) and GPFS. He holds a degree in Computer Science from CBS in Copenhagen, has various IBM Product Certifications, and is an IBM Certified IT Specialist.

Gheorghe Tudor is an IT Specialist and currently works for IBM Global Technologies Services in Romania. He has experience in designing and implementing solutions based on AIX, PowerHA, PowerVM®, General Parallel File System (GPFS), IBM System Storage®, VMware, Linux, Cisco & Brocade SAN and Cloud Computing. Gheorghe holds a degree in Computer Science and has various IBM Product Certifications.



From left to right: Srinivasa Raghavan, Christian Schmidt, Glen Corneau, Gheorghe Tudor, Vinicius Goncalves, Dino Quintero (project leader), Laszlo Niesz, Piotr Mazur, and Tejaswini Kaujalgi

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

David Bennin, Richard Conway, Ella Buslovich, Alfred Schwab, Elsie Ramos
International Technical Support Organization, Poughkeepsie Center

Marty Fullam, Sean Safron, Bruce Potter
IBM Poughkeepsie

Alan Wilcox, Khanh Le, Prasad Potluri, Sandy Amin
IBM Austin

Craig Elliott
IBM Dallas

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>



Abstract

This Redbooks publication contains helpful information for the IBM Cluster Systems Management (CSM) for AIX administrator who is looking to utilize IBM Systems Director for AIX. In 2009, IBM announced that CSM would not be actively enhanced and existing CSM customers should evaluate to moving to one of two products in the future: xCAT (eXtreme Cloud Administration Toolkit used for High Performance Computing customers) and IBM Systems Director with VMControl (for commercial customers). Since that time there has been no guide from IBM for the commercial AIX customer. This Redbook fills the gap and assists these customers who are looking to transform their CSM environment to one managed by IBM Systems Director.

The topics covered in this Redbook include:

- ▶ Comparison of features and functions between IBM Cluster Systems Management for AIX and IBM Systems Director for AIX.
- ▶ Transformational methodologies with real examples of planning, tasks and instructions.
- ▶ References to various informational sources on IBM Systems Director and related products.

The following topics are covered in this chapter:

- ▶ “Introduction to CSM”.
- ▶ “Introduction to IBM Systems Director”.
- ▶ “Scope of this book”.
- ▶ “Why transform from CSM to IBM Systems Director”.

1.1 Introduction to CSM

IBM Cluster Systems Management (CSM) for AIX is designed for the administration of distributed and clustered IBM Power Systems environments. CSM considerably simplifies management of a cluster by providing supervision from a single point-of-control.

CSM provides the following system management capabilities for IBM Power Systems nodes in a managed cluster:

- Remote command execution across multiple nodes in the cluster.
- Configuration file manager to dispense and synchronize files across nodes in the cluster.
- Comprehensive monitoring and automated response of system parameters and status.
- Security configuration across the cluster to eliminate the need to manually distribute encryption keys.
- Centralized management of remote hardware control and console access to increase administrative efficiency.
- Command line interface support to script complete tasks.
- Software diagnostic tools to analyze components and servers to find the root cause of problems.
- Web-based interface to effectively manage and control the CSM cluster.

For detailed planning and installation procedures, refer to the *CSM Planning and Installation Guide* at:

<http://publib.boulder.ibm.com/epubs/pdf/a2313445.pdf>

For detailed CSM Administration procedures, refer to the CSM Administration Guide at:

<http://publib.boulder.ibm.com/epubs/pdf/a2313445.pdf>

1.2 Introduction to IBM Systems Director

IBM Systems Director is an enhanced system management tool to drive operational management of heterogeneous IT environments. IBM Systems Director significantly reduces operational complexity, improves the efficiency of IT staff and systems, and manages costs while meeting challenging business requirements for service delivery.

IBM Systems Director is designed to incorporate discovery, monitoring, optimization and updates from a single point of contact. IBM Systems Director not only manages IBM power-based servers but also x86-based servers, VMware, PowerVM, and storage resources.

The enterprise IBM Systems Director provides the following system management capabilities:

- ▶ System administration
 - Discovery of servers and virtual resources
 - Monitoring and reporting system health
 - Configuration and update of the server operating systems and firmware
- ▶ Virtualization management
 - Creation, deployment and management of virtual machines

- Maintain virtual images in a repository
- Manage virtual server workload
- ▶ Energy management
 - Collection and reporting of system energy usage
 - Analyzing historical energy usage
 - Advanced energy management and reporting
- ▶ Network management
 - Discover, monitor, and manage multi-vendor network devices
 - Automate network management tasks
- ▶ Service and Support
 - Automatic problem reporting
 - Performance and capacity analysis
 - AIX enterprise system reports
 - Enterprise service monitoring

For detailed step-by-step procedures for using IBM Systems Director, refer to the following IBM Redbooks publication, and IBM publications located at the following sites:

<http://www.redbooks.ibm.com/redbooks/pdfs/sg247694.pdf>

http://publib.boulder.ibm.com/infocenter/director/v6r2x/topic/com.ibm.director.main.helps.doc/fqp0_bk_install_gde_aix.pdf

1.3 Scope of this book

With products such as the IBM Cluster Systems Management and the IBM Systems Director, there are many features, functions and capabilities that can be used by customers. This chapter provides additional details about what we cover, and what we do not cover in this book.

1.3.1 What we cover

Both IBM Systems Director and IBM Cluster Systems Management support a wide variety of environments. At a high level, the scope of this book is as follows:

AIX operating system only

While both CSM and IBM Systems Director support multiple operating systems, the majority of commercial CSM cluster customers are using AIX as their primary CSM operating system. This operating system choice implies that Power (and not x86) is the architecture we cover in this book.

CSM functions in IBM Systems Director

After reviewing 1.2, “Introduction to IBM Systems Director” on page 2, it is clear that IBM Systems Director and its associated plug-ins add extended capabilities that do not exist in IBM Cluster Systems Management. Therefore, this book does not delve into many of these functions because they are already covered in the IBM Systems Director product documentation and many books listed in the “Related publications” on page 149.

1.3.2 What we do not cover

The following topics are not covered in this publication.

Detailed installation steps

The installation and configuration of IBM Systems Director, VMControl Standard Edition, AIX and the Network Installation Manager (NIM) are well documented in other publications. Therefore, these topics are not covered in this book except for pertinent information encountered during the cluster systems management transformation. For detailed information on these products, refer to the InfoCenter pages as follows:

- ▶ IBM Systems Director and VMControl

http://publib.boulder.ibm.com/infocenter/director/v6r2x/topic/com.ibm.director.main.helps.doc/fqm0_main.html

- ▶ AIX and NIM (V7.1)

<http://publib.boulder.ibm.com/infocenter/aix/v7r1/index.jsp>

High performance computing

IBM Cluster Systems Management has been utilized in both the commercial and high performance computing (HPC) customer environments. This book focuses on the commercial customers and the transformation to IBM Systems Director. For the HPC customer, refer to *Configuring and Managing AIX Clusters Using xCAT 2*, SG24-7766.

Non-Power Systems

Both IBM Cluster Systems Management and IBM Systems Director support multiple hardware architectures and multiple operating systems. This book only covers Power Systems running AIX.

Application integration

Customer utilization of both IBM Cluster Systems Management and IBM Systems Director includes multiple interfaces: traditional command line; graphical and even advanced programming interfaces (APIs). Detailed comparisons of the programming APIs are beyond the scope of this book.

xCSM graphical user interface

While IBM Systems Director's main interface is graphical, this is not the case with IBM Cluster Systems Management. IBM provided, on an as-is basis, a web-based graphical front-end for CSM via the xCSM package. This was not part of the official licensed product and was supported ad hoc via the CSM forum only. A detailed comparison of the xCSM interface and the IBM Systems Director web interface is not discussed in this book.

1.4 Why transform from CSM to IBM Systems Director

This section discusses the reasons and benefits of transitioning from Cluster Systems Management (CSM) to IBM Systems Director.

Why transform

Cluster Systems Management (CSM) has been used with success in many IBM Power System environments making it easy to manage and deploy multiple systems. However,

introduction of new technologies in POWER and AIX products created a necessity to develop a management product that is aware of the latest features created by new technologies.

IBM Systems Director is a multiplatform management foundation for achieving a single-point-of-control over the entire IT environment. It enables integration with Tivoli® and third party management platforms. A complete redesign of Systems Director allows an administrator to easily understand, manage, and deploy complex systems built on the newest virtualization technologies. With Systems Director we can:

- ▶ Automate data center operations
- ▶ Unify the management of IBM servers, storage and network devices
- ▶ Simplify the management of physical and virtual platform resources
- ▶ Obtain a view of systems status, health, utilization and power usage
- ▶ Understand and map the complex virtualization topologies

Because of these reasons, Cluster Systems Management (CSM) is no longer shipped or supported with the AIX 7.1 operating system, making Systems Director its successor.

Benefits of Systems Director

Transforming your systems management from Cluster Systems Management to Systems Director provides many additional capabilities of managing your IBM Power Systems environment. We consider the following abilities most important:

- ▶ Virtualization management - The ability to view, configure, and deploy complex virtual environments
- ▶ Multivendor, crossplatform hardware support - The support for Power Systems, System i®, System x®, BladeCenter® and many more, including network and storage devices
- ▶ Active energy management - The measuring, monitoring, and managing of the energy components built into IBM systems



Planning and preparation

In this chapter we discuss many important topics that provide the background to any transformation of an existing IBM Cluster Systems Management for AIX cluster to IBM Systems Director. As with any major project, planning is key and the information provided in this chapter should be a good reference for the administrator.

This chapter contains the following sections:

- ▶ Comparison
- ▶ Terminology
- ▶ Agent Manager considerations
- ▶ Product prerequisites
- ▶ Licensing
- ▶ Extendibility
- ▶ Features not directly available in IBM Systems Director
- ▶ Features and functions not available in CSM
- ▶ Database
- ▶ Transformation as a project and future impacts

2.1 Comparison

While this entire book is a reference comparing and contrasting IBM Cluster Systems Management for AIX and IBM Systems Director, a bit of background can be useful to understand some of the major differences between the products.

IBM Cluster Systems Management for AIX is, at its core, a command line-driven set of programs to assist the AIX Administrator managing multiple operating systems and Logical Partitions (LPARs) with as few steps as possible. CSM evolved from the predecessor product IBM Parallel Systems Support Programs (PSSP), which was developed in concert with the RS/6000® SP hardware product in the early 1990s. The tight integration of both hardware and software control was a key component used by many customers. With the spread of High Performance Computing clusters in the Linux and x86 arena, CSM (unlike PSSP) adapted to support those technologies.

IBM Systems Director for AIX developed from a predecessor product, IBM Director, with a background in the Microsoft Windows x86 world focused on monitoring and managing large numbers of disparate desktop and server systems. As management of all the IBM server hardware platforms from a consolidated pane-of-glass became a priority, IBM Director multiplatform was the first step to bring it all together. IBM Systems Director Version 6 was a major rewrite of the previous products using the cross-platform Eclipse-based development environment setting a new, higher baseline for the single pane-of-glass concept.

2.2 Terminology

As with many transformations, we must understand new terminology introduced in the new environment. Naming conventions for some familiar components in IBM Systems Director (Systems Director) are very different from the ones used in IBM Cluster Systems Management (CSM). Here, we will briefly discuss the terminology important for the administrator transforming from CSM to Systems Director. Fully detailed definitions are available in *IBM Systems Director for AIX Planning, Installation and Configuration Guide Version 6.2.1*, GI11-8709-06.

2.2.1 Components

IBM Systems Director is a server-client based solution consisting of IBM Systems Director Server and two client-side operating system agents: Common Agent and Platform Agent. IBM Systems Director can be installed on various IBM-supported operating systems.

- ▶ Agent Manager

Serves as authentication and authorization component for Common Agents. Multiple Management Servers can utilize a single, common Agent Manager.

- ▶ Management Server

The management server is an operating system instance running the IBM Systems Director Server code. It provides a central point of control for aggregating and managing discovered systems based on a service-oriented architecture.

- ▶ Common Agent

IBM-provided software installed on an operating system required to enable a rich set of security, deployment, and management functions from a Systems Director perspective.

- ▶ Platform Agent

Also IBM-provided software that provides a subset of Common Agent functions. Used to communicate with and administer managed systems, including hardware alerts and status information.

- ▶ Managed System

Systems discovered, managed and monitored by the Management Server. Can refer to a virtual operating system instance, a standalone server, a physical server that hosts virtual servers, and more.

- ▶ Agentless managed systems

Systems managed by the Management Server without Common Agent or Platform Agent installed. Communication with these systems uses native protocols such as Secure Shell (SSH), Simple Network Management Protocol (SNMP), Common Information Model (CIM), and more.

2.2.2 Manageable system types

Systems managed by IBM Systems Director can be hardware or software based. All of the entities that have an IP address and can be discovered by IBM Systems Director are considered systems. Here are some important ones when it comes to Power Systems:

- ▶ Operating System

Operating system software instance installed and running on a physical or virtual server. In Power Systems these are Logical Partitions or Full-System Partitions running AIX, IBM i, Virtual I/O Server (VIOS) or Linux on Power operating systems.

- ▶ Server or Host

A physical computer system. In Power Systems, this is a machine able to host multiple LPARs or a single partition system.

- ▶ Platform Manager

A component that provides access to virtualized resources.

- ▶ Hardware Management Console

A special type of Platform Manager that provides access, via the Flexible Service Processor (FSP), to Power System servers.

- ▶ Integrated Virtualization Manager

Another type of Platform Manager that provides access to virtualized Power Blades and small- to medium-class Power Systems servers.

- ▶ System Chassis

An enclosure for one or more blades, both Power and x86.

2.2.3 Base functions

Base functions, also called plug-ins, provide basic functions in IBM Systems Director allowing it to perform management of IBM servers, virtual servers, storage and network resources.

- ▶ Discovery Manager

Provides predefined discovery protocols for virtual and physical systems on TCP/IP networks. Collects information (inventory) about systems and displays system relationships.

- ▶ **Status Manager**
Monitors hardware, software and power status of discovered systems.
- ▶ **Update Manager**
Acquires (downloads or imports) required updates, including operating system software, drivers, and firmware. Creates and monitors update policies.
- ▶ **Automation Manager**
Runs predefined and/or user-defined tasks when designated events occur.
- ▶ **Configuration Manager**
Configures attributes and options on discovered systems. Can automatically configure a newly discovered system if desired.
- ▶ **Remote Access**
Provides remote control over systems either by remote console, Virtual Network Computing (VNC), or web-based remote control.
- ▶ **Storage Management**
Provides monitoring and management of storage resources.
- ▶ **Network Management**
Manages and monitors discovered network devices.
- ▶ **BladeCenter and System x Management**
Provides lifecycle management for BladeCenter, System x and related resources including specific functions and capabilities.
- ▶ **Power Systems Management**
Provides lifecycle management for Power Systems and related resources including Hardware Management Console (HMC) and Integrated Virtualization Manager (IVM).
- ▶ **VMControl**
Allows management of virtual resources. Capabilities here are determined by the Edition of VMControl installed.

2.2.4 Mapping terminology

To make the transition easier, we decided to list some basic terms used in Power Systems and provide their counterparts in Cluster Systems Management (CSM) and IBM Systems Director. Due to the different nature of these products, this table should only be used to familiarize yourself with new terminology because there may not be an exact one-to-one mapping, as shown in Table 2-1.

Table 2-1 Mapping Cluster Systems Management terminology to IBM Systems Director

Term	Cluster Systems Management	IBM Systems Director
Logical Partition (LPAR)	Node	Virtual Server
Operating System (AIX, Linux, Virtual I/O Server)	Node	Operating System
Physical server	Node	Host or Server
	Hardware control	Power Systems Management

Term	Cluster Systems Management	IBM Systems Director
	Monitoring	Status Manager
	rconsole	Remote Access, dconsole

2.3 Agent Manager considerations

A key internal component of IBM Systems Director is the Agent Manager. The Agent Manager is responsible for credentials and authentication between the IBM Systems Director Server and the Common Agent-managed systems.

Notes:

- ▶ The Agent Manager is only available to IBM Systems Director Server 6.1 or later. A Systems Director Server must be configured against one and only one Agent Manager at a given time.
- ▶ A Common Agent-managed system can be registered with one and only one Agent Manager. If you try to request access to a Common Agent-managed system that is already registered with an Agent Manager, it will fail.

As part of the IBM Systems Director implementation you may plan to configure a single Agent Manager for the enterprise. By implementing this configuration you will be able to manage Common Agent systems from multiple management servers, decreasing endpoint-management complexity.

To configure the Agent Manager (during initial implementation) that the server will use to manage Common Agent resources, the following command is used:

```
/opt/ibm/director/bin/configAgtMgr.sh
```

Note: Only the IBM Systems Director Server and Common Agent use the Agent Manager. There is no relation between the Agent Manager and Platform Agents, Agentless-managed systems, or other resource types.

For detailed information about how to configure the Agent Manager, refer to the *IBM Systems Director for AIX Planning, Installation and Configuration Guide*. From a security perspective you must have the ports utilized by the Agent Manager properly configured in your firewalled environment to allow communication to your Systems Director Servers.

Figure 2-1 on page 12 shows a simple diagram displaying the Agent Manager and connected IBM Systems Director Servers.

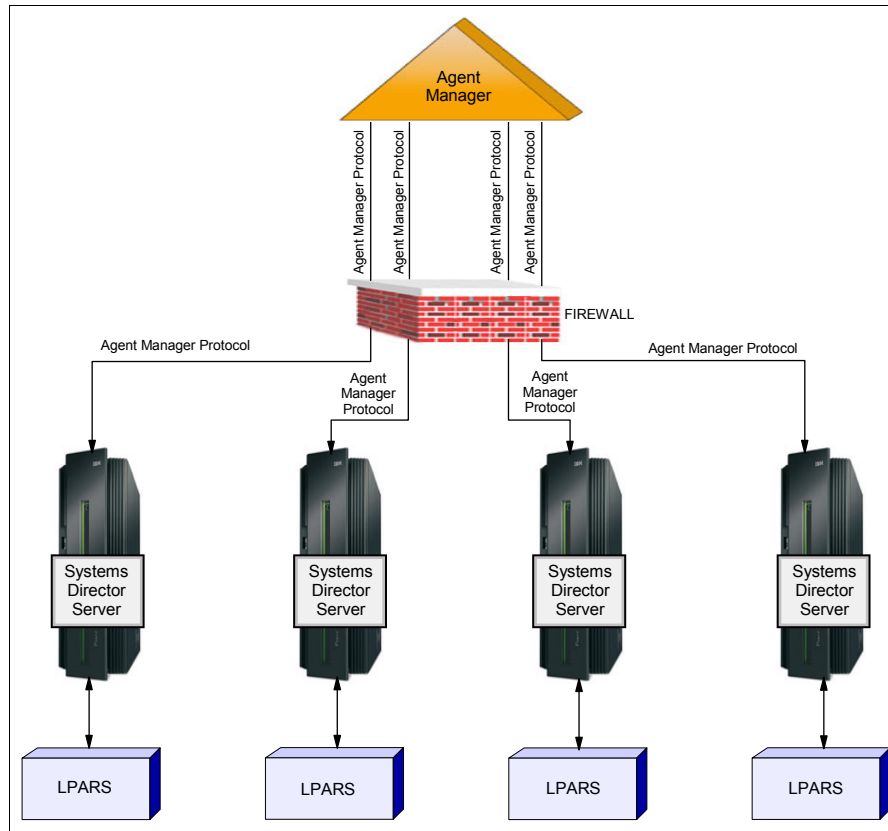


Figure 2-1 Agent Manager

2.4 Product prerequisites

Any successful transformation from CSM to IBM Systems Director requires verification that your current environment will be supported. If you are planning to implement the IBM Systems Director Server on the same operating system instance as your CSM server (as described in 3.1, “Overview of the scenarios” on page 40), you need to bring the environment to a state where both CSM and IBM System Director are supported. This section discusses the requirements for IBM Systems Director and CSM and shows the overlaps between them.

2.4.1 Software

Table 2-2 shows the supported versions for CSM and IBM Systems Director. Check whether you have the supported version levels.

Note: The highest supported level for CSM 1.7.1.12 refers to the highest tested versions of software at the time of writing.

Table 2-2 IBM Systems Director and CSM support for IBM AIX

Software	AIX 5.3	AIX 6.1	AIX 7.1
CSM 1.7.1.12	TL7 up to TL11	TL00 up to TL06	not supported
IBM Systems Director v6.2.x	TL6 SP09 or later SP TL07 SP06 or later SP TL08 SP04 or later SP TL09 or later TL	TL00 SP07 or later SP TL01 SP03 or later SP TL02 or later TL	all

Note: IBM Systems Director recommends the use of the Common Agent on AIX clients. A new AIX base system installation with current versions of AIX V5.3 or V6.1 will include the Common Agent by default.

Table 2-3 shows the firmware levels supported by both products.

Table 2-3 IBM Systems Director and CSM support for firmware

Software	POWER5	POWER6®	POWER7®
CSM 1.7.1.12	All	Up to 350	Up to 720
IBM Systems Director v6.2.x	SF240_338 or later	All	All

Note: It is recommended that you always use the latest possible firmware level.

Both the Hardware Management Console (HMC) and the Integrated Virtualization Manager (IVM) can be used by IBM Systems Director and CSM. Table 2-4 shows the version levels supported.

Table 2-4 IBM Systems Director and CSM support for HMC, IVM and VIOS

Software	HMC	IVM	VIOS
CSM 1.7.1.12	Up to v7 R7.2	All	N/A
IBM Systems Director v6.2.x	v7 R3.3.0 SP2 or later v7 R3.4.0 or later v7 R3.5.0 SP1 or later v7 R7.1.0 or later v7 R7.1.1 or later v7 R7.2.0 or later	v1.5.2.1 or later v2.1.0.10 or later v2.1.1 or later v2.2 or later	v1.5.2.1 or later v2.1.0.10 or later v2.1.1 or later v2.2 or later

Note: VMControl requires IVM to be at a minimum version of 2.1.0.10.

The Common Agent installed with Virtual I/O Server 2.1.1 or later is not started by default on system boot. See the VIOS documentation for details about starting the Common Agent.

2.4.2 Hardware and performance

Since this book covers transformation from CSM, we focus only on hardware resource requisites and performance recommendations for IBM Systems Director.

IBM Systems Director Server for AIX requires installation on a supported IBM Power Systems server running AIX, including IBM Power JS21 and later blade servers. The minimum architecture is POWER5, but POWER7 is recommended for best performance. The hardware resource requirements vary depending on the workload.

Server hardware requirements for running IBM Systems Director for AIX can be obtained by using the IBM Systems Workload Estimator, which is a web tool for estimating hardware requirements for various workloads. You can access the IBM Systems Workload Estimator workload for IBM Systems Director at:

<http://www-01.ibm.com/support/docview.wss?uid=nas7cd6a96f49d05f608862577420075ca9a&aid=3>.

The IBM Systems Workload Estimator asks several questions about your intended usage, such as the number of systems you are managing and whether you are using a local or remote database. After the questions are answered, IBM Systems Workload Estimator will generate a drop-down list of system model and features that meet the resource requirements of Systems Director including an estimate of the number of cores, memory, disk capacity, and the number of disk drives required. It also shows a “Growth Solution” that can handle further growth for managing larger environments.

When reviewing these recommendations, consider the following:

- ▶ Installation and startup times improve with faster disk access times. High-speed disk adapters and faster drives provide the best performance.
- ▶ Disk sizes are arbitrary and indicative of disk requirements.
- ▶ System performance depends on the nature of your requirements and system workload.
- ▶ The IBM DB2® database software sizing should be comparable for Oracle.
- ▶ The Estimator only provides disk space requirements that are not on a file system base. For AIX, consider that space requirements to run IBM Systems Director relate to the following file systems:
 - /opt - for server information not kept in a database unless using the default Derby database. Refer to your database configuration for the file system hosting the non-Derby database.
 - /var - for Common Agent information and other log files.
- ▶ Paging space should be at least 3 GB regardless of the memory size.
- ▶ Using a logical partition in a larger server than displayed is also an option.

It might be necessary to change the maximum Java heap size in order for IBM Systems Director to take advantage of additional memory. On AIX, the Java heap size setting is the following path: /opt/ibm/director/lwi/conf/overrides/director.javaopt. On 64-bit operating systems, the optimal heap size value is typically in the range of 2 GB (-Xmx2g) to 8 GB (-Xmx8g), depending on the systems being managed and the amount of physical memory available on the IBM Systems Director Server.

Table 2-5 on page 15 provides information about disk storage requirements for installing IBM Systems Director Server. These requirements include Common Agent and Platform Agent, which are installed with IBM Systems Director Server. These requirements do not include disk storage requirements for running IBM Systems Director Server and the database, updating IBM Systems Director Server or for downloading and staging updates (such as AIX, HMC, or system firmware updates) with Update Manager.

Table 2-5 Disk storage requirements for the IBM Systems Director Server

/	50 MB
/usr	256 MB
/var	512 MB
/tmp	2 GB ^a
/opt	5 GB

a. This space is required only during installation and is freed after the installation is complete. If /tmp is a JFS2 file system, you can shrink it after the installation.

2.5 Licensing

In this section we provide the licensing requirements.

2.5.1 Licensing requirement for IBM Cluster Systems Management

CSM for AIX is a licensed program charged based on the number of processors (cores) in the system where CSM for AIX will be run, or to which hardware control commands will be used. To use CSM for production purposes, you must purchase a license that will provide the appropriate license key, which will be provided on the product media. During the CSM Management Server installation, you must accept the IBM Program License Agreements (IPLA) license agreement. The license key is a file named `csmlum.full`. The key file is only used on the Management Server.

Note: CSM must be licensed for all components including management server and managed nodes.

2.5.2 Licensing requirements for IBM Systems Director

Licensing requirements for the management server and managed systems are an important consideration when planning to deploy IBM Systems Director, which is licensed on a per-active core basis for each server running IBM Systems Director (both management server and managed endpoint).

IBM Systems Director has three licensed editions: Express, Standard and Enterprise.

- ▶ All three editions are available for Power Systems.
- ▶ Edition availability on the other platforms (x86 and System z®) may be different. Check the latest IBM Announcement Letters for availability.

Each edition includes a license. Edition capabilities include:

- ▶ IBM Systems Director Express Edition
 - Discover and inventory servers and operating systems
 - Monitor, configure and update resources
 - Find and resolve problems faster

- ▶ IBM Systems Director Standard Edition
 - Monitor and manage energy within capacity.
 - Reduce time to deploy workloads.
 - Get a single view and status of network and server systems.
- ▶ IBM Systems Director Enterprise Edition
 - Deploys workloads faster with improved reliability in system pools.
 - Increases productivity with prioritized information and context, real-time and historical health status.
 - Uses predictive capabilities to help with capacity estimation.

The following chargeable plug-ins (*not* the Editions) are provided with a 60 day trial license available:

- ▶ VMControl Standard and Enterprise Edition
- ▶ Active Energy Manager
- ▶ Network Control
- ▶ Storage Control

Note: IBM Systems Director requires a proper license for all platforms, including the management server and server endpoints.

2.6 Extensibility

One of the features of both IBM Cluster Systems Management and IBM Systems Director is the ability to add to the base functionality as shipped either via IBM-defined or administrator-defined programs.

2.6.1 Application programming interfaces

In this section we provide information about the application programming interfaces.

IBM Systems Director REST API

In addition to the graphical user interface and command line interfaces, IBM Systems Director also has a Software Development Kit that utilizes the representational state transfer (REST) API to provide a consistent mechanism for interacting with the Systems Director Server. The REST API allows access to the base Systems Director functions (discover, inventory, configuration, health and monitoring, event notification and automation) and plug-in functions for VMControl and Active Energy Manager. The programming model allows for asynchronous notification via Java Message Service for processing real-time events.

For documentation about the SDK, see:

http://publib.boulder.ibm.com/infocenter/director/sdk/topic/com.ibm.usmi.toolkit.doc/systems_management_landing.html

RSCT RMC API

IBM Cluster Systems Management does not, itself, have a direct programming interface. However, the technology that underpins the software, namely Reliable Scalable Cluster

Technology (RSCT), does contain programming interfaces. RSCT has multiple components, but the one utilized by CSM is 99 Resource Monitoring and Control (RMC). Applications written to the RMC API can:

- ▶ List the resources of a resource class.
- ▶ Monitor changes in attribute values for events of interest.
- ▶ Query dynamic or persistent attributes of resources or resource classes.
- ▶ Change the persistent attributes of resources or resource classes.
- ▶ Define and undefine resources.
- ▶ Bring resources online and take them offline.

For details about monitoring in both CSM and Systems Director, see 4.1, “Monitoring of resources” on page 74. For details about the RSCT RMC API, see *RSCT RMC Programming Guide*, SA23-1346. or online at:

http://publib.boulder.ibm.com/infocenter/clresctr/vxrx/topic/com.ibm.cluster.related_libraries.doc/related.htm?path=3_6#rsct_link

2.6.2 IBM Systems Director Advanced Managers (Plug-ins)

IBM Systems Director provides systems management personnel with a single-point-of-control, which helps reduce the management cost and complexity. Systems Director assists with optimizing the compute and the network resources, is greatly flexible and attains higher levels of service management with streamlined management of physical, virtual, storage and network resources.

IBM Systems Director is delivered with a variety of plug-ins. In addition, there are installable plug-ins (some available with Systems Director Editions) to manage energy, virtualization, operating system installation, high availability, storage, network, workload partitions and provide hardware service call-home. This section provides a brief introduction to the additional plug-ins available.

Active Energy Manager

Active Energy Manager (AEM) is an IBM Systems Director plug-in that offers power and thermal monitoring. It comes with many management capabilities and can be used to better understand the power usage of the data center which, in turn, helps with the optimal utilization of power. AEM can also be used to plan for future energy needs.

AEM can directly monitor IBM Power Systems, System z, System x and Blade Center servers, and extends the scope of energy management to include non-IBM servers and facility providers to enable a more complete view of energy consumption within the data center.

VMControl

The VMControl (VMC) plug-in is designed to manage a cross-platform virtualized environment, including Power Systems servers. It provides rapid deployments of virtual appliances to create virtual servers preconfigured with an operating system and software applications that you need. Virtual Server lifecycle management allows you to create, edit, delete and even relocate virtual servers between physical hosts.

VMC also enables you to group resources into system pools to centrally manage and control the workloads in your environment.

This diversity of capabilities allows you to select the optimal level of VMControl functionality (Express, Standard, or Enterprise Edition) for your virtualization infrastructure and to seamlessly upgrade as it evolves.

PowerHA SystemMirror Director

The PowerHA SystemMirror Director plug-in provides management interfaces for configuring and managing PowerHA SystemMirror high availability clusters. This interface provides an easy step-by-step wizard to create the clusters and define or modify Resource Groups and Run-Time Policies.

The plug-in also provides various functionalities that include:

- ▶ Monitor and manage the status of the clusters, nodes
- ▶ Synchronization of the cluster, cluster recovery, reports
- ▶ Management of the cluster services
- ▶ Manage and monitor the resource groups
- ▶ Network and storage management
- ▶ Snapshots
- ▶ Command line interface for the configuration and the management of the cluster.

Note: The IBM Systems Director for PowerHA SystemMirror plug-in requires, and is shipped with, the licensed program product IBM PowerHA SystemMirror 7.1 for AIX.

AIX Profile Manager

AIX Profile Manager (APM) for System Director is a plug-in that can be used to manage the configuration settings of a group of AIX systems from a single point-of-control. It is able to collect configuration settings from each Common Agent-managed AIX operating system defined in System Director. The collected configuration is a set of properties from multiple areas of an AIX instance such as reliability, availability, serviceability, security, or kernel.

The AIX Profile Manager plug-in manages AIX system configurations through profiles. A profile is an XML-formatted file that contains a set of runtime properties for a given domain of activity, such as a user, TCP/IP, kernel_heap_size, /etc/inetd.conf. A profile contains a set of parameters with optional flags and values.

Profiles are used to deploy configurations on your managed systems in order to align their configuration, check your managed systems compliance with a reference profile, or to retrieve the values of a profile on a system's running configuration. Manual creation of the profiles is not possible because they rely on catalogs residing on the AIX systems you are monitoring. Instead they can be imported on your IBM Systems Director plug-in from the local system or from a managed system.

More details about the installation, configuration and administration of this plug-in can be found in *AIX Profile Manager*, SC23-6766-00, located online at:

http://publib.boulder.ibm.com/infocenter/aix/v7r1/topic/com.ibm.aix.apmgr/apmgr_pdf.pdf

Notes:

- ▶ AIX Profile Manager is shipped with the AIX 7.1 Standard and Enterprise Edition media set on a separate CD.
- ▶ AIX-managed systems require the DirectorCommonAgent fileset and the artex.base.agent and artex.base.rte filesets for APM to function.

Service and Support Manager

IBM Systems Director Service and Support Manager is a no-charge plug-in to IBM Systems Director that automatically reports hardware problems, collects system service information for monitored systems and provides the interface back to IBM for call-home reporting. With these functions, the Electronic Service Agent™ (which is the underpinning of Server and Support

Manager) can monitor, track, and capture system hardware errors and service information and report them directly to IBM support for you.

The benefits of this plug-in include:

- ▶ Automatic problem reporting
- ▶ 24/7 direct routing of reported problems to IBM technical support
- ▶ Reduced personnel time required for gathering and reporting service information
- ▶ Higher availability and shorter downtime
- ▶ Custom IT management tools enabled
- ▶ Secure Internet access
- ▶ Accurate solutions with reduced human error in gathering and reporting service information
- ▶ Secure web access to your service information
- ▶ Consistent IBM worldwide service and support process

Note: The current release of Service and Support Manager does not support HMC- or IVM-managed Power System servers. Standalone servers are supported.

WPAR Manager

This plug-in enables IBM Systems Director to manage AIX Workload Partitions (WPARs) and provides functions to relocate the WPARs from one server to another without the need to restart the application. The Live Application Mobility function can be automated by the policy engine.

The WPAR Manager allows the administrator to create, remove, clone, or start and stop the WPARs. In short, it provides the complete lifecycle management support for WPARs.

The other features provided by the WPAR manager include:

- ▶ Backup and restore
- ▶ Static relocation of WPARs
- ▶ IP v6 support
- ▶ Synchronize WPAR
- ▶ Storage devices support
- ▶ Support for WPAR-owned rootvg
- ▶ Kerberos support

Note: For all the above features, it is required to have AIX Version 6.1 with the 6100-02 Technology Level, or later.

IBM System Director Storage Control

IBM Systems Director Storage Control plug-in facilitates the management and configuration of network-attached storage subsystems, BladeCenter integrated storage, internal RAID storage, and storage switches.

It can discover, collect inventory and monitor the health for various storage subsystems including the IBM DS8000® and DS500 family, SAN Volume Controller, IBM Storwize® V7000, Brocade, and QLogic Fibre Channel switches.

Management of the storage subsystem is done through its respective SMI-S provider, which is a vendor-specific module that must be installed and properly configured.

Note: SMI-S providers and an IBM Systems Director agent should not be installed on the same machine because there can be a port conflict.

IBM Systems Director with Storage Control and VMControl plug-ins provides a single management console for managing servers, storage systems and Fibre Channel switches.

In conjunction with VMControl, there is support for capabilities such as:

- ▶ Storage provisioning for image creation, deployment and cloning
- ▶ Ability to manage storage system pool lifecycle
- ▶ Taking group actions across pool and policy-based storage placement
- ▶ Cloning actions within the pool

Storage Control unifies the management of physical and virtual server and storage resources for integrated end-to-end lifecycle management, including: configuration, discovery, health, capacity, inventory, provisioning, topology, updates, alerts, and retirement.

2.6.3 Command line

This section provides information on the command line interface.

Command line interface

Although IBM Systems Director's web interface is really powerful and can be used to perform most of the tasks, sometimes it is desirable to have a tool that can help you to automate some tasks or processes with the help of scripts.

For these types of activities, System Director comes with a powerful command line interface. IBM System Director has two categories of commands based on the purpose of the command, as follows:

Single-purpose commands

These commands can perform only one function, as opposed to the **smcli** command that can perform many functions.

Examples of functions that can be accomplished by this type of command would be:

- ▶ Initialize IBM Systems Director's database connection: **cfgdbcmd**
- ▶ CIM event management: **cimsubscribe**
- ▶ Reset the IBM Systems Director Server database: **smreset**
- ▶ Backup/restore the Systems Director's persistent data (file system data and database): **smsave/smrestore**
- ▶ Start, stop or get the status of IBM Systems Director Server: **smstart/smstop/smstatus**.

Server-based command-line interface (smcli)

The smcli tool is an alternative way to interact with the IBM Systems Director Server. It is installed by default on the IBM Systems Director Server and must be run as a user with **smadmin** or **smmgr** group authority (or equivalent administratively-defined group). You can run **smcli** commands locally from the management server or remotely by accessing the management server using a remote-access utility, such as Secure Shell (SSH) or Telnet.

Commands in the smcli tool are grouped by function and each function includes sub functions. Examples of command functions are: virtualization commands, storage commands, security commands, discovery commands, resource and process monitor commands, remote access commands and many others.

For an example of the commands in the smcli, see Example 2-1 on page 21. Note that the commands available will vary depending on authority and plug-ins that are installed.

Example 2-1 Sample of smcli commands

```
# /opt/ibm/director/bin/smcli lsbundle
snmp/addsystem
snmp/get
snmp/getbulk
snmp/getnext
...
ActiveEnergyManager/chpolicy
ActiveEnergyManager/chpowerinfo
ActiveEnergyManager/chproperties
...
automation/evtlog
automation/lsevtact
...
event/listeventactionplans
event/listeventactions
event/listevents
...
imagemgrcli/captureva
...
inventory/collectinv
...
sccli/dumpstcfg
...
sysmirror/lcluster
...
user/luser
user/lusergp
...
wparmgr/mkwpar
...
```

2.7 Features not directly available in IBM Systems Director

IBM Systems Director is a lot different from CSM. Not every feature available in CSM has direct equivalence in IBM Systems Director. In this chapter we discuss those features and, when possible, provide recommendations or replacements.

2.7.1 Configuration File Manager

Configuration File Manager (CFM) is a CSM feature that provides a file repository for configuration files that are common to some or all cluster nodes. CFM stores all shared configuration files in one location on the management server and automatically propagates changes to these files throughout the cluster. You must maintain the configuration files stored on the management server. CFM also supports pre- and post-distribution scripts to handle necessary tasks, for example stopping and starting daemons.

No direct replacement for this feature is offered in IBM System Director at the moment of writing this book. You can manage configuration schemes using the AIX Profile Manager plug-in but it does not provide a common file repository feature. The AIX command **rdist** can be used to distribute files in a manual fashion, but does not have any pre- or post-distribution capabilities.

You can also use PowerHA SystemMirror to synchronize files across cluster nodes. More information about this topic can be found in the *PowerHA SystemMirror 7.1 for AIX Standard Edition* documentation available at:

http://publib.boulder.ibm.com/infocenter/aix/v7r1/topic/com.ibm.aix.powerha.adm/ngd/ha_admin_manage_file_coll.htm

2.7.2 Cluster-Ready Hardware Server (CRHS)

Cluster-Ready Hardware Server is a component of CSM that provides advanced management of physical Power Systems servers (POWER5 and POWER6) with regards to their associated Hardware Management Consoles (HMCs). When using CRHS, the CSM Management Server and the HMCs are clustered together in a RSCT Peer Domain. This allows for dynamic assignment of physical servers to their associated HMC. CRHS is required for clusters using the IBM High Performance Switch.

Some features provided by CRHS:

- ▶ Consolidation of service networks into a maximum of two subnets for redundancy
- ▶ Automated association of System p5@ 575, 590 and 595 servers with their frames
- ▶ Ease of movement of System p servers between HMCs
- ▶ Shared database of System p cluster hardware information
- ▶ Reduced number of Dynamic Host Configuration Protocol servers for large clusters

The IBM HPS required CSM for support, and as such is not supported with IBM Systems Director. Therefore, there is no equivalent function in IBM Systems Director.

2.7.3 Managing node status information with csmstat

This section illustrates managing the node status in CSM.

Showing status information in CSM

CSM provides an excellent tool, used by most CSM administrators, to provide an initial view of the status of their cluster nodes: **csmstat**. The command shows the nodes, their primary managing HMC, RSCT status, power status and adapters, which are visible from the CSM management server. Example 2-2 shows the output of the command in our second transformation scenario.

Example 2-2 Sample output of the csmstat command

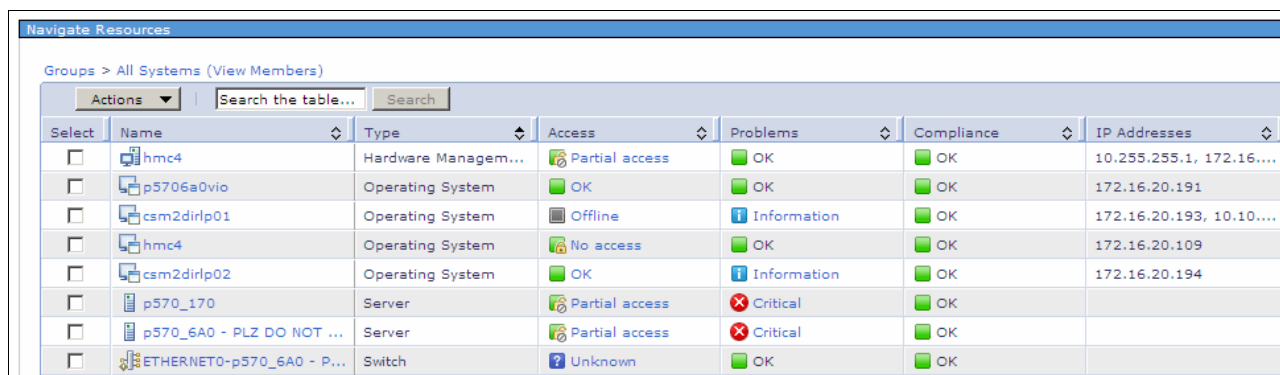
```
[p61p04:/]# csmstat
```

Hostname	HWControlPoint	Status	PowerStatus	Network-Interfaces
p61p02	hmc2	on	on	en0-Online en1-Online
p61p03	hmc2	on	on	en0-Online
p6vio1		unknown	not configured	unknown

In addition to this output, there is an option to the **csmstat** command (-l) to display the individual LPAR LCD display.

Showing status information in the IBM Systems Director GUI

Unfortunately there is no command similar to `csmdat` available in Systems Director at the time of writing this book. Systems Director provides a graphical interface in which the user can configure different status monitors. The default status information is also very useful to show the status of the whole infrastructure. Figure 2-2 shows the status of our third test environment in the Navigate Resources panel after choosing the **All Systems** predefined group. For this example, we stopped the Common Agent on `csm2dirp01` and revoked the access to HMC named `hmc4`.



Select	Name	Type	Access	Problems	Compliance	IP Addresses
<input type="checkbox"/>	hmc4	Hardware Managem...	Partial access	OK	OK	10.255.255.1, 172.16...
<input type="checkbox"/>	p5706a0vio	Operating System	OK	OK	OK	172.16.20.191
<input type="checkbox"/>	csm2dirp01	Operating System	Offline	Information	OK	172.16.20.193, 10.10...
<input type="checkbox"/>	hmc4	Operating System	No access	OK	OK	172.16.20.109
<input type="checkbox"/>	csm2dirp02	Operating System	OK	Information	OK	172.16.20.194
<input type="checkbox"/>	p570_170	Server	Partial access	Critical	OK	
<input type="checkbox"/>	p570_6A0 - PLZ DO NOT ...	Server	Partial access	Critical	OK	
<input type="checkbox"/>	ETHERNET0-p570_6A0 - P...	Switch	Unknown	OK	OK	

Figure 2-2 Navigate Resources - All Systems

Showing status information in the IBM Systems Director CLI

There will be many system administrators who want to use tools and commands available in the command line, so we provide some sample commands and the background information to use those commands.

Note: Some of the parameters and commands we show in the following examples are not documented and can be changed or removed without any notification in future versions or updates of IBM Systems Director and VMControl. Check the availability of these options in your environment and after all updates before integrating them in your production environments.

Use the `smcli lssys -I` command to display information about all types of managed objects.

In case of managed Power Systems with AIX the most important types are:

- HardwareManagementConsole
- Server

We can filter the servers using the HypervisorPlatformRole in the option of the `lssys` command. Use:

- VirtualContainer to list LPARs
- HostPlatform to list Power systems

- OperatingSystem

Each of the managed objects have varying relationships to each other depending on object type.

An OperatingSystem type object's DisplayName attribute relates to a server with VirtualContainer object's role (that is, an LPAR) InstalledOSDisplayName attribute, as shown in Example 2-3.

The first column shows the Server name, the second column shows the OperatingSystem object's name.

Example 2-3 Listing installed operating system on a server

```
# smcli lssys -A\
InstalledOSDisplayName,SerialNumber,PowerState,OperatingState,CommunicationState,AccessState \
-t Server -w "HypervisorPlatformRole=VirtualContainer" -d" " | grep -v Unsupported|sed "s/ //"
csm2dirlp01:csm2dirlp01 100F6A0 1 8 2 Unlocked
csm2dirlp02:csm2dirlp02 100F6A0 1 8 2 Unlocked
VIO_Server1:p5706a0vio 100F6A0 1 8 2 Unlocked
```

A server with the VirtualContainer role is connected to a server with the HostPlatform role (Power Systems) using the SerialNumber attribute, which is common for both objects. Example 2-4 shows how to query the hosting machine.

The managing HMC of a host is listed by querying the TWGMOParentSet attribute of the server.

Example 2-4 Listing the host machine of a VirtualContainer

```
# export SERIAL=100F6A0
# smcli lssys -A TWGMOParentSet,PowerState,OperatingState,CommunicationState,AccessState \
-t Server -w "HypervisorPlatformRole=HostPlatform AND SerialNumber=$SERIAL" -d" "
p570_6A0: { 11843 } 1 8 2 Unlocked
```

The attribute is a list of the managing HMCs. So if there are two HMCs connected to the physical machine using both HMC ports, we will see two object IDs.

Based on this information we created a sample script to show the attributes and the connections between managed objects. This is useful when an administrator wants to use the CLI to list status information. See Example 2-5 for the output of the script in our scenarios.

Example 2-5 Output of sample status script

```
### Scenario 1:
Hostname      HealthState Communicate AccessState LPAR      PowerState HostSystem PowerState
              HMC      State
-----
p5570lp01    1          2          Unlocked  p5570lp01 1          p570-CSM  1
              First HMC: hmc1      8          2          Unlocked
p5570lp03    1          2          Unlocked  p5570lp03 1          p570-CSM  1
              First HMC: hmc1      8          2          Unlocked
p5570lp04    Unsupported 2          Unlocked  p5570lp04 1          p570-CSM  1
              First HMC: hmc1      8          2          Unlocked
p5570lp05    Unsupported 2          Unlocked  p5570lp05 1          p570-CSM  1
              First HMC: hmc1      8          2          Unlocked

### Scenario 2:
Hostname      HealthState Communicate AccessState LPAR      PowerState HostSystem PowerState
              HMC      State
-----
isd           1          2          Unlocked  isd        1          p570_170  1
              First HMC: hmc4      8          2          Unlocked
nim           1          2          Unlocked  nim        1          p570_170  1
```

p6client1	First	HMC:	hmc4	8	2	Unlocked	
	Unsupported	2		Unlocked	p6client1	1	p570_170
	First	HMC:	hmc4	8	2	Unlocked	
p6client2	1	2		Unlocked	p6client2	1	p570_170
	First	HMC:	hmc4	8	2	Unlocked	
p6viol	1	2		Unlocked	p6viol	1	p570_170
	First	HMC:	hmc4	8	2	Unlocked	

More details about the possible attributes can be found in the IBM Systems Director documentation.

The script can be found in Appendix A, “isdstat script” on page 145.

For more monitoring possibilities, see 4.1, “Monitoring of resources” on page 74.

Grouping managed objects

The IBM Systems Director administrator can create groups of managed objects using static or dynamic selection methods. It is possible to assign objects to groups based on relations to other objects, as shown in Figure 2-3.

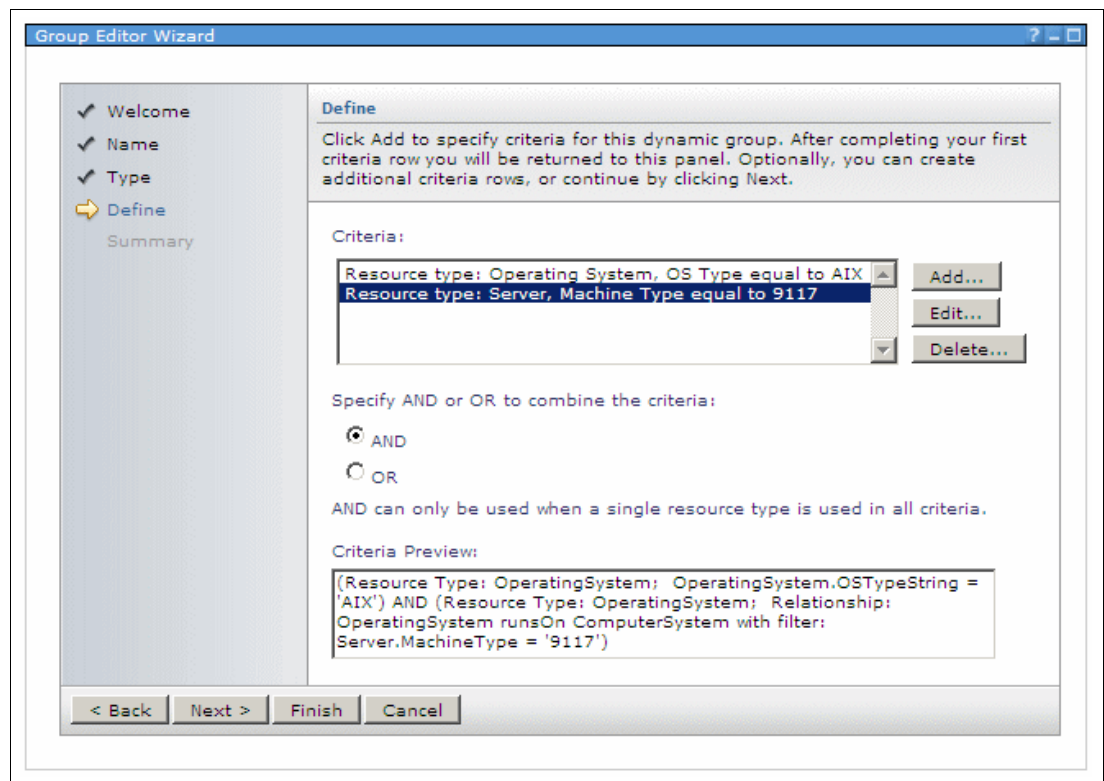


Figure 2-3 Specify criteria for a dynamic group

Using customized dynamic groups and the status information as displayed in Figure 2-2 on page 23, this provides an easy way to check the whole infrastructure at a glance.

The groups can be used in monitoring and also in assigning roles for specific users or user groups to enable certain tasks on objects in the groups.

2.7.4 High availability management server

We can separate the high availability based on Systems Director's main components:

- Systems Director Server with Agent Manager and features installed together on the same operating system

At the time of writing this book there is no preconfigured, supported method to set up a complete highly available server configuration. Some considerations are discussed in 2.3, "Agent Manager considerations" on page 11 and 5.3, "High availability IBM Systems Director management server" on page 137.

- Database server

In case of need for high availability it is preferable to set up a remote database server. The database server can use high availability features available for the chosen product. Note the database versions supported with IBM Systems Director referenced in 2.9, "Database" on page 29 and read the appropriate documentation.

2.8 Features and functions not available in CSM

While there are many features and capabilities that are unique to IBM Systems Director, we felt it useful to document, at a high level, some of the major functions of import to Power Systems customers.

2.8.1 LPAR lifecycle management

Cluster Systems Management for AIX was designed to manage existing AIX or Linux operating system images in both single-system (that is, one operating system per physical server) and logically partitioned (multiple operating systems per physical server) modes with Power hardware. However, it was not designed to manage the entire lifecycle of a logical partition (LPAR, or virtual server). This lifecycle includes:

- Creation of new LPARs from available resources
- Modification of assigned resources to the LPAR (that is, dynamic LPAR operations)
- Relocation of LPARs between servers that are capable of Live Partition Mobility
- Removal (destruction) of LPARs

IBM Systems Director with the VMControl Express Edition plug-in provides all these capabilities in an integrated fashion with both graphical and command-line interfaces. While CSM's remote command capability can execute tasks on the appropriate Hardware Management Console (HMC) or Integrated Virtualization Manager (IVM), it requires an exacting syntax for the commands. Systems Director's knowledge of the entire virtualized environment allows for easier management of complex configurations from the single pane of glass. Components that are inventoried and managed include:

- The platform manager (HMC or IVM) that provides the connection and the interface into the remainder of the virtualized resources. This not only includes the physical Power Systems servers and all the associated subcomponents, but also out-of-band hardware problem detection and reporting through the Service Agent component running on the HMC/IVM. For each Power System server the hardware configuration (drawers, processor, memory, disk, and so on) is gathered along with the on demand resources that may not be currently available to client LPARs. Through this information, Systems Director can display the used and unused resources for a given host.

- ▶ The utility virtual server (Virtual I/O Server) that provides the virtual network and virtual SCSI resources to the various VIO client LPARs. This includes the vscsi and vhost resource definitions and the client LPARs to which those resources are assigned.
- ▶ The end client LPARs where the operating systems are installed. Available knowledge includes the virtual resources (SCSI and Ethernet, physical and virtual processors, memory) and physical resources (assigned slots and adapters, optical and virtual optical devices). For the processor and memory resources, Systems Director knows about the profile (minimum, maximum) values.

Knowledge of the physical server (host)

CSM understands the physical server (host) for functions such as firmware (see 2.8.2, “Update Manager” on page 28 for more details). However, there is little relational data linking the host and the virtual servers it contains. Systems Director understands this relational data and much more. In Example 2-6, the Systems Director command line is used to display the knowledge stored in the database for a physical server and the LPARs it contains. This same information can be found in the graphical interface by right-clicking the Host representation and choosing **Related Resources** → **Server** (or just left-click, because this is the default action).

The **bc** command is used to convert between hexadecimal and decimal for the various commands.

Example 2-6 Listing the hosting relationship between a physical server and a virtual server

```
p5570lp01(root)/> smcli lssys -o p570-CSM
p570-CSM, 0x53e1
p5570lp01(root)/> smcli lssys -o p570-CSM
p5570lp01(root)/> echo "ibase=16;53E1" | bc
21473
p5570lp01(root)/> smcli lssystem 21473 | grep Relationships
Relationships found: 66
p5570lp01(root)/>

p5570lp01(root)/> smcli lssystem 21473 | grep "hosts Server"
Server(21473) hosts Server(21479)
Server(21473) hosts Server(21477)
Server(21473) hosts Server(21474)
Server(21473) hosts Server(21478)
Server(21473) hosts Server(21475)
Server(21473) hosts Server(21476)
Server(21473) hosts Server(2776)
p5570lp01(root)/> echo "obase=16;21479" | bc
53E7
p5570lp01(root)/> smcli lssys -o -n 0x53e7
p5570vio1, 0x53e7
```

Knowledge of the Virtual I/O Server (VIOS)

The previous example (Example 2-6) ends with the display of a Virtual I/O Server LPAR. IBM Systems Director can not only gather standard inventory information from the operating system running inside the VIOS, but also virtualized resources. This information comes from the communication with the Hardware Management Console and the **viosvr cmd** command. The virtual resources and the client LPARs are stored in the Systems Director database. There is nothing to be configured on the VIOS for this data gathering.

For communication with the Common Agent on the VIOS, the agent needs to be started with the `startsvc director_agent` command before being successfully discovered.

2.8.2 Update Manager

The Update Manager available with IBM Systems Director helps to acquire, install and manage the updates on the systems. It can automatically receive new update information, provide the update documentation and determine the appropriate updates for managed systems. In addition, the compliance policies defined can notify the administrator when requested updates are missing.

The Update Manager can be used to update the following:

- ▶ Many hardware components, including system firmware, for IBM System i, IBM System P, IBM Power Systems, IBM System x, and BladeCenter servers
- ▶ The operating systems of IBM i (PTF groups and cumulative PTF packages), IBM AIX (Technology Levels and Service Packs), and Linux systems
- ▶ The IBM Systems Director software itself
- ▶ Hardware Management Console fix packs, service packs and efixes
- ▶ Virtual I/O Server fix packs, service packs and interim fixes
- ▶ Power I/O Firmware for Power Systems components such as adapters, disk drivers, and so on
- ▶ Device Driver and Update Express System Pack updates for System x
- ▶ IBM BladeCenter modules' firmware

The Update Manager can be found in the IBM Systems Director interface on the left-hand navigation bar through **Release Management** → **Updates** as shown in Figure 2-4.

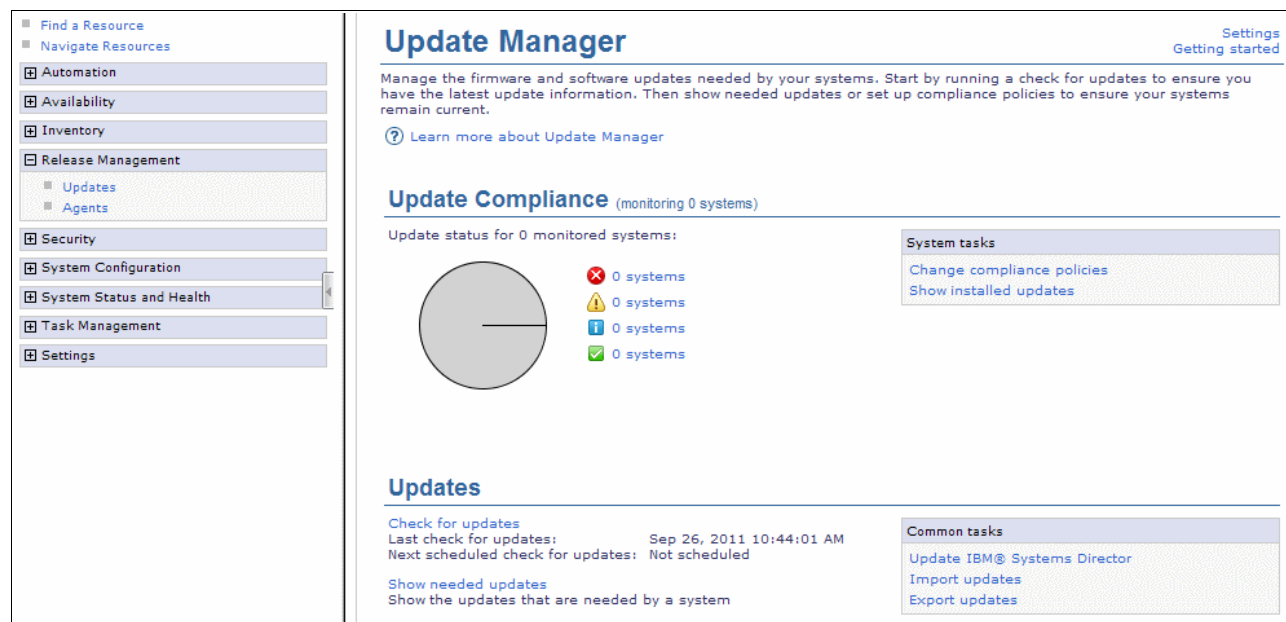


Figure 2-4 Update Manager

One important feature of the Update Manager is providing the notification of which systems require firmware and operating system updates. Using the Update Manager, a compliance

policy can be defined for a single system or a designated group of systems. Compliance policies notify the system administrators when a system or a group of systems need updating.

The Update Manager compliance policies can contain a single update or multiple updates to a single system, group of systems, or multiple groups of systems.

The general steps followed during an update installation include:

- ▶ Download - The IBM Systems Director Server connects to the appropriate Fix Service Provider to retrieve the installable files for an update (stored on the management server).
- ▶ Installation staging - Copies the installable files for an update to an appropriate location for later installation.
- ▶ Installation - Installs an update.
- ▶ Uninstall - Removes an update. Not all updates support the uninstall task.

The Update Manager is not capable of installing new software products or AIX filesets, IBM Systems Director agents, or migration of IBM Systems Director from one release to another.

Note: For more information about Update Manager, refer to:

http://publib.boulder.ibm.com/infocenter/director/v6r2x/topic/com.ibm.director.updates.helps.doc/fqm0_t_um_updating_systems.html

2.9 Database

The centralized management of software and hardware objects requires storage to handle the inventory information about those objects. It also needs a safe method to update that information when an appropriate change occurs with the managed entities. This information is stored in a database.

The database stores both static and dynamic information. Some of this information is changed as result of manual or automated operations. Some information may be monitored to warn the administrators of a possible error.

Both CSM and IBM Systems Director use a database to store and handle information about the managed environment, but they have different approaches.

2.9.1 RSCT System Registry versus an actual DB product

This section provides information on the RSCT system registry in comparison to an actual database product.

RSCT System Registry for CSM

CSM depends on the Reliable Scalable Cluster Technology (RSCT) component to store configuration information. It also uses the RSCT security infrastructure (CtSec) for authentication and access control.

RSCT contains a system registry (SR) that supports both persistent data (kept on disk) and dynamic data (only held in memory). The system registry is enclosed in the RSCT Resource Management and Control (RMC) framework, which is available to other IBM software components.

The `rsct.core` filesets are installed by default during the base AIX installation. Therefore, each AIX machine provides the full RSCT Resource Management and Control infrastructure

as well as the RSCT client commands. Each AIX machine can serve resources (including data stored in its local system registry), and can be a client of any other machine resources.

The data that CSM stores in the system registry is managed through a CSM-specific resource manager that resides on the management server. The system registry persistent data resides in subdirectories underneath the /vary/ct/ directory. Some registry data can be modified by CSM commands.

Notes:

- ▶ Further details about RSCT can be found in the IBM Cluster Library at:
http://publib.boulder.ibm.com/infocenter/clresctr/vxrx/topic/com.ibm.cluster.related_libraries.doc/related.htm
- ▶ For general discussion of RSCT and CSM RSCT data backups, refer to:
<http://www.redbooks.ibm.com/abstracts/tips0090.html?Open>
<http://www.redbooks.ibm.com/abstracts/tips0262.html?Open>

Product databases for IBM Systems Director

Unlike CSM, the IBM Systems Director Server primary uses a relational database to store inventory information in a central location. By default, the Apache Derby database is installed during server installation if other options are not selected.

Advantages of the default Apache Derby database:

- ▶ Configuration is easy.
- ▶ No additional licenses or software acquisition is required.
- ▶ A separate server for the database is not required.

Disadvantages of the default Apache Derby database:

- ▶ Limited to 500 managed objects or less.
- ▶ Database cannot be remote.
- ▶ Cannot be used if the Storage Control plug-in is required.

Note: If you plan to manage a large environment from a single IBM Systems Director Server using an external database, consider this limitation in your planning. There is no path for migrating data from the Apache Derby database to another database. Rediscovery of managed objects is required along with additional recustomization.

IBM Systems Director can use other databases to store inventory and update information. The three databases supported are:

- ▶ IBM DB2
- ▶ Oracle
- ▶ Microsoft SQL Server

The advantages of using a non-default database:

- ▶ The database can be local or remote to the IBM Systems Director management server.
- ▶ If the database is remote, you will reduce the disk space requirements on the management server.
- ▶ Performance and scalability (up to the maximum number of objects supported by a single IBM Systems Director Server) can be greater.

The disadvantages of using a non-default database:

- ▶ If the database server is remote, a separate server is required. Any network connectivity problems with the database server will affect the IBM Systems Director.
- ▶ If not already purchased or licensed, the database software must be acquired.

Database versions supported by IBM Systems Director Server

Table 2-6 lists the database versions supported by IBM Systems Director Server on different management servers and provides information about whether the database server is embedded or can be installed locally or remotely.

Note: The database versions that are listed represent both the database server and the database client where applicable.

Table 2-6 Database versions supported by IBM Systems Director

Database	Supported database versions	AIX	Linux	Windows
Apache Derby	V10.5.3.1 (included with IBM Systems Director Server on AIX, Linux, and Windows)	Embedded	Embedded	Embedded
IBM DB2	<ul style="list-style-type: none"> ▶ Express version 9 ▶ Version 9.1 with Fix Pack 4 or later ▶ Version 9.5 with Fix Pack 1 or later ▶ Version 9.7 with Fix Pack 1 or later <p>Notes:</p> <ol style="list-style-type: none"> 1. An IBM DB2 Version 9.5 or later client is required. 2. A 64-bit client is required to access an IBM DB2 database from 64-bit IBM Systems Director Server. 	Local or remote	Local or remote	Local or remote
Microsoft SQL Server	<ul style="list-style-type: none"> ▶ Microsoft SQL Server 2005 with Service Pack 3 ▶ Microsoft SQL Server 2008 with Service Pack 1 <p>Note: Requires Microsoft SQL Server JDBC Driver 2.0 or later.</p>			Local or remote
	<ul style="list-style-type: none"> ▶ Microsoft SQL Server 2005 Express Edition with Service Pack 2 ▶ Microsoft SQL Server 2008 Express <p>Note: Requires Microsoft SQL Server JDBC Driver 2.0 or later.</p>			Local

Database	Supported database versions	AIX	Linux	Windows
Oracle Database	<ul style="list-style-type: none"> ▶ Version 9.2 ▶ Version 10g release 1 (10.1.0.3 or later) ▶ Version 10g release 2 ▶ Version 11g release 2 <p>Notes:</p> <ul style="list-style-type: none"> ▶ Oracle Cluster configuration is not currently supported. ▶ A 64-bit client is required to access an Oracle Database database from 64-bit IBM Systems Director Server. 	Local or remote	Local or remote	Local or remote

2.10 Transformation as a project and future impacts

The transformation from CSM to IBM Systems Director will result in major changes in the way you manage your infrastructure. We cannot avoid thinking about it as a project, which requires proper planning activities not only for the system administration personnel but for other teams who have the responsibility to keep important applications running.

Having a CSM cluster is not only about installing and monitoring the operating system. The way you manage that operating system has serious implications on the applications that are running on those systems. Changing the core management infrastructure, when you transform it from CSM to IBM Systems Director, may result in planned outages when you are performing the transformation and most likely will result in changes in the way the systems are administered after transformation is complete.

In the following section, we discuss some of the nontechnical aspects of such a major change in your infrastructure. We make appropriate references when we see a technical issue that is documented in this or other published materials.

The transformation itself can be seen as a project, but it will result in changes how you can do your daily work after that is finished, so we split this section into two parts.

2.10.1 Transformation project

We provide two scenarios with detailed guidance for the technical steps and provide further technical explanations for the most important steps of the transformation process. Refer to Chapter 3, “Transformation scenarios” on page 39.

In this section, we concentrate on the broader view.

Preparation

Planning and preparation can take longer than the actual implementation steps of a project and is the most important aspect of any successful project.

Preparation involves not only technical requirements, but non-technical issues as well:

- ▶ Review your management infrastructure together with IBM Systems Director documentation to understand all aspects of your infrastructure, and the parts that will be replaced by the new management infrastructure.
 - Choose what you can replace easily and what kind of development activities you will need in case something is not available or is provided differently in IBM Systems Director.
 - Choose the required plug-ins of IBM Systems Director at the onset so the implementation considerations can be made during initial installation.

We suggest a minimal implementation at first for simplicity and then add features and plug-ins later when you are used to the new environment.

For example, immediate cost saving results can be achieved using Active Energy Manager, but this could require additional investments which were not in the initial responsibility of the server management team.

- Choose the database application that will be used by IBM Systems Director.

If we start with the embedded Apache Derby with a small infrastructure, it could be hard to move to IBM DB2 later, when Systems Director is already managing hundreds of systems. For more details, see “Product databases for IBM Systems Director” on page 30.
- ▶ Prepare for the fact that you may need additional system resources while you are preparing and doing the transformation compared to daily operations under CSM. Because IBM Systems Director can provide more features, using those features can raise the resource requirements even after transformation is complete.

- ▶ Provide education of new technologies implemented with IBM Systems Director.

This is probably the most important step, because the differences between CSM and Systems Director are substantial. These differences are not only in the technical details, but administrators and other involved parties have to think differently when they are working with IBM Systems Director compared to CSM.

IBM Systems Director has many new features and is a cross-platform management tool that is developed to manage today's cloud-based, virtualized infrastructures.

The database used by IBM Systems Director is simply managed and does not require heavy maintenance activities. But it is there, so it is better to know something about it.

Education should provide not only theoretical classroom knowledge, but hands-on workshops and test systems available before, during and after the transformation project provide the greatest benefit.

High-level education on IBM Systems Director features and requirements has to be provided for non-technical personnel also, to understand why some changes in processes are needed.

- ▶ Prior to production, test everything you want to implement; not only the systems management side of the infrastructure. Think about the application- and process-related consequences also.
- ▶ Choose your path of transformation, which can be any combination of the following:
 - CSM and IBM Systems Director on the same management infrastructure.
 - Build a new management infrastructure for IBM Systems Director.
 - Set up one or more NIM servers on separate machines or one on the same server where the IBM Systems Director Server is installed.

- Implement a hierarchical management infrastructure or have only one management server for all managed systems; this includes NIM server placement as well.
- Move all CSM nodes into one maintenance window, move in groups or move one by one.

► Prepare for the worst

The management infrastructure should prevent service loss and be able to restore the service to normal operation in the allotted time based on the Service Level Agreement (SLA).

- Any utilized monitoring capabilities provided by CSM and other related applications require a smooth transformation to minimize the time when an important service can go down without notification.
- A backup infrastructure has to be available in case an unforeseen restoration is required.

This is not only about the capability to restore the recently changed components of your management infrastructure, but in many cases CSM is connected to NIM servers that provide the AIX operating system restoration capability in your managed infrastructure. See the section about OS deployment for more details in 4.5, “OS deployment” on page 105.

► Review security rules appropriate for systems management

- Managing remote servers requires authentication and authorization, which are often strictly determined by corporate rules and standards.

IBM Systems Director, like CSM, provides tools that can cause problems if used by inexperienced, or even malicious, personnel for infrastructure changes and/or data acquisition.

Start with minimal authorization for new administrators in IBM Systems Director and extend in a controlled manner.

See security sections in this publication and other documentation for more technical information and incorporate the requirements into the processes. Refer to 4.4, “Security” on page 95.

- Prepare for integrating the existing post install customization procedures into IBM Systems Director.

Probably the most important of these are the security hardening steps which have to be accomplished after each server installation. Using cloning of already hardened operating systems is one method, for which IBM Systems Director and VMControl can provide assistance. See more details about Storage Copy Services capabilities in IBM System Storage DS Storage Manager Copy Services Guide SG24-7822.

The IBM Systems Director plug-in AIX Profile Manager can help with post install customization and security hardening.

► Review processes related to systems management and start detailed planning early in the transformation because changes such as this can take time to incorporate in a large enterprise.

- IBM Systems Director works with virtualized environments, which may need changes in hardware and operating system requirements handling and lifecycle management.
- Centralized user ID management and security handling of shared user IDs should be reviewed.
- IBM Systems Director provides centralized update notification, and possibly installation, for AIX. This has to be incorporated into the existing processes in a company.

- Status of all managed systems can be displayed in the GUI and additional monitoring can be set up as explained in 4.1, “Monitoring of resources” on page 74, but this can overlap with existing monitoring tools in a company.

Typically any newly managed system or operating system has to be registered for monitoring as part of the activation process. Because IBM Systems Director can provide status and health monitoring by default, this process could require changes.

Transformation from CSM to Systems Director

If you did a good job in the planning, then the actual implementation part of the project should be easy. Just do what was documented and prepared for.

Of course, there are always some things that can go wrong, so listed below are some of those risks:

- Removal of old management software and installation of new management software on managed systems

After removal of the management software from a single CSM node, that node is no longer managed, at least not managed by CSM. This sounds obvious but can cause problems if the planning and the schedule of the project were incorrect.

Without management software you could lose the following functions:

- Restoration of a managed server’s operating system

To avoid this, be sure that at least one NIM server and the backups are still available, so you can use manual restore procedures.

- Monitoring

Inform the monitoring team, if they are separate from the implementors, that some of the monitored features and processes will be stopped or completely removed, so they should not open problem tickets for those.

Without monitoring you cannot see if there is a non-obvious failure that can result in missing the SLA, so use other tools that are still available (such as IBM Tivoli Monitoring or the HMC GUI) to keep an eye on the systems.

- Remote commands and access to managed nodes or devices

Use SSH or HMC options for manual tasks. If the time without a new remote command facility is too long, then modification of automated scripts could be necessary.

When CSM functions are the only path used by administrators to remotely access HMCs or managed nodes, then this will not be available until IBM Systems Director is configured to manage the server. Set up individual authentication and authorization to avoid a situation where an administrator cannot log in and perform any urgent activities after removing an HMC from CSM.

- Hardware control

The physical servers are still managed by the HMCs, so you can use them directly to control the hardware if needed.

- Security

Prepare the security environment first to avoid accidental or malicious modifications of the managed systems.

Create groups of users and assign only the minimal roles that are required. Users with unlimited rights can view and control much more of the IT infrastructure than was possible with CSM because of the extended capabilities of IBM Systems Director.

- ▶ **Monitoring**

Be careful when removing CSM, because it is possible that application-related monitors are used. Stopping these monitors or monitored CSM processes can cause undesired effects at the application level.

A server that was managed by CSM is fully transformed to IBM Systems Director if all of the hosted applications are running and monitored in the manner supported by the required SLA.

- ▶ **Refresh the backups of managed systems**

New system backups should be created post-transformation and integrated into the new OS deployment infrastructure, otherwise a restoration of an old backup image can cause problems.

2.10.2 BAU after transformation to IBM Systems Director

After a successful transformation project you will have a changed management infrastructure along with changed processes as well.

A summarization of some of the most important BAU (business as usual) activities performed by system administrators that can change as the result of transforming from CSM to IBM Systems Director includes:

- ▶ **Hardware provisioning**

Hardware provisioning can be integrated into OS deployment processes if you utilize the virtualization features of Power Systems with VMControl Standard or Enterprise Edition. Virtual I/O Servers and the storage environment has to be prepared for this change, which typically also requires reviewing the processes used by server and storage administration teams.

- ▶ **Operating System (OS) deployment**

The most widely used OS provisioning tool for AIX is the built-in Network Installation Manager. IBM Systems Director VMControl™ Standard and Enterprise Edition provides end-to-end OS deployment in conjunction with NIM. NIM resources are created and removed as necessary for the deployment. New mksysb images need to be integrated not only on the NIM server but with VMControl in virtual appliances format.

NIM resources can be created manually for post-install customization and assigned at OS deployment time. This is likely the best method for initial security hardening, installing necessary agents and configuring application prerequisites on the installed operating system.

For more details, see 4.5, “OS deployment” on page 105.

- ▶ **Operating System updates**

Operating System updates can be automated via IBM Systems Director's Update Manager, which will also utilize NIM. See 2.8.2, “Update Manager” on page 28 for more information.

- ▶ **Firmware updates**

CSM provided the capability to manage the installation of system firmware updates, and this capability was enhanced with IBM Systems Director. Any system firmware update (disruptive or concurrent) can cause problems if the update is not planned very carefully after reading the documentation and understanding any issues and requisites.

See more information about the firmware update capabilities in 4.3.2, “System firmware updates” on page 91.

Note: Adapter firmware updates can have requirements outside of the machines where the adapter itself is installed because they can have prerequisites against the SAN and storage infrastructure. Neither CSM nor IBM Systems Director today can completely guard against a mismatch in component levels.

► Monitoring

Take advantage of dynamic groups in IBM Systems Director based on system attributes. These attributes can determine which monitors are activated on which systems and can speed up installation processes or changes.

► Handling hardware errors

Most IBM hardware products have the capability to notify IBM about any failure that happens on the devices, either natively or using call-home reporting tools such as Enterprise Service Agent. In addition, you can create notifications to administrators or they can see it on the HMC GUI if they are logged in to the specific HMC where the failing device is connected.

IBM Systems Director is designed to provide a single starting point for management of the entire infrastructure and has enhanced grouping features. This makes it much easier to see a problem just watching the high-level view of the whole infrastructure and drill down for more information when you see any notification.

Utilizing Power System's virtualization features and providing a broad view of the whole infrastructure makes it much easier for the administrators to find the necessary replacement resource and configure it when it is necessary to resolve a problem.

► Automating day-to-day activities

IBM Systems Director provides many features that ease the administrator's burden:

- The graphical interface with the capability for filtering system lists using groups, monitors and status information, provides an overview of the whole environment at a glance.
- Operating system, HMC and Virtual I/O Server control functions are easy to access.
- There are multiple methods to run remote commands and operations, and to access operating system prompts.
- Operating System and management software updates can be automated.
- Full inventory available all the time, which can be a huge help dealing with other departments within the company.

► User IDs to improve the efficiency of managing User IDs for admin team members.

This can be a frequent activity because teams managing certain parts of the infrastructure can change rapidly.

Creating roles and groups in IBM Systems Director can provide dynamic authorization of user IDs to perform necessary activities with minimal configuration activities based on external IDs, even if managed by the operating system or by an external directory (LDAP or Active Directory).

Use the auditing feature in IBM Systems Director from the beginning, but be aware that the audit log is circular.

See 4.4, "Security" on page 95 for more information.



Transformation scenarios

In this chapter we provide a few sample scenarios you can follow to transform your existing IBM Cluster Systems Management (CSM) cluster to IBM Systems Director. Specifically, we present two scenarios built in the laboratory based on actual customer configurations including basic transformation steps and the related results. We also provide descriptions of some of the main features that are configured and a proposed list of test and verification operations.

We do not provide all possible variables that may be part of the cluster transformation from CSM to IBM Systems Director. The intention is to get you started on preparing your organization for CSM transformation, or at least understand what it takes to move your cluster from CSM to IBM Systems Director.

The term transformation has been chosen because it reflects the fact that you are changing the form and structure of your cluster and not just performing a simple migration to a new cluster. However, we have labeled one of our scenarios a migration because it also involves migrating to new software and hardware.

3.1 Overview of the scenarios

The two scenarios have two different approaches for transformation. The first scenario is more conservative and relies on reusing the old infrastructure. The second scenario uses a new physical infrastructure for both the IBM Systems Director Server and NIM Server and focuses on keeping existing capabilities in place as well as implementing new functionalities only available in IBM Systems Director.

Here is a brief description of the two scenarios:

- ▶ Coexistence Scenario - CSM Management Server and the IBM Systems Director Server in the same operating system on the same server.
- ▶ Migration Scenario - CSM Management Server and IBM Systems Director Server on different servers.

In order to provide familiarity with both management interfaces of IBM Systems Director we have focused on the Graphical User Interface for the first scenario and the command line interface for the second scenario. No matter what path you choose when transforming your CSM cluster, we recommend that you read both scenarios and familiarize yourself with both interfaces because some tasks may only be available through one interface.

The transformation scenarios show the movement of one or more nodes at a time from CSM to IBM Systems Director while CSM is up and running.

Notes:

- ▶ We strongly recommend having both CSM and IBM Systems Director active at the same time because this should be a temporary solution. Once the transformation is complete, CSM should be decommissioned.
- ▶ Before transforming your environment we strongly recommend to have full backups of all machines and CSM. For information about CSM backup refer to 5.4, "Backup and restore of the management server" on page 139.

3.1.1 What we do not cover

The following topics are not covered in this publication:

- Installation and Administration of CSM

For detailed installation and administration steps, refer to the CSM Installation Guide at:

<http://publib.boulder.ibm.com/epubs/pdf/a2313435.pdf>

- Installation and configuration of IBM Systems Director

For detailed installation and configuration of IBM Systems Director for AIX, refer to:

http://publib.boulder.ibm.com/infocenter/director/v6r2x/topic/com.ibm.director.main.helps.doc/fqp0_bk_install_gde_aix.pdf

- NIM installation and configuration

For information related to NIM, refer to:

http://publib.boulder.ibm.com/infocenter/aix/v7r1/topic/com.ibm.aix.install/doc/insgdrf/nim_intro.htm

- IBM Systems Director VMControl Standard Edition

For information about VMControl, refer to:

http://publib.boulder.ibm.com/infocenter/director/v6r2x/topic/com.ibm.director.vim.helps.doc/fsd0_vim_pdf.pdf

3.1.2 What we cover

The following topics are covered:

- ▶ Transformation of nodes (endpoints) to IBM Systems Director
- ▶ Transformation of hardware devices (HMC and Virtual I/O Servers in IBM Systems Director)
- ▶ CSM and IBM Systems Director common functionality testing
- ▶ CSM decommissioning
- ▶ Migration of AIX 5.3 to AIX 7.1 using IBM Systems Director

Note: AIX 7.1 is not supported by CSM, but is supported by IBM Systems Director V6.2.1.

3.1.3 Prerequisites

For IBM Systems Director-related prerequisites, refer to 2.4, “Product prerequisites” on page 12.

3.1.4 Recommendations

We recommend switching from the default Apache Derby database used in IBM Systems Director to IBM DB2. For more details about databases used and how to change them, refer to 2.9, “Database” on page 29.

For even more information, refer to Chapter 9 of *IBM Systems Director for AIX Installation Guide*:

http://publib.boulder.ibm.com/infocenter/director/v6r2x/topic/com.ibm.director.main.helps.doc/fqp0_bk_install_gde_aix.pdf

3.1.5 Considerations and time estimates

Depending on the customer environment, there may be requirements to migrate AIX, CSM and/or IBM Systems Director to the latest level prior to or during the transformation. Any form of software upgrade has an associated risk, and therefore upgrading software during this process adds risk to the transformation itself.

The transformation steps need to be performed and completed in a given order, as can be seen in Figure 3-1 on page 42. More accurate time estimates will be based on your own environment size and available resources.

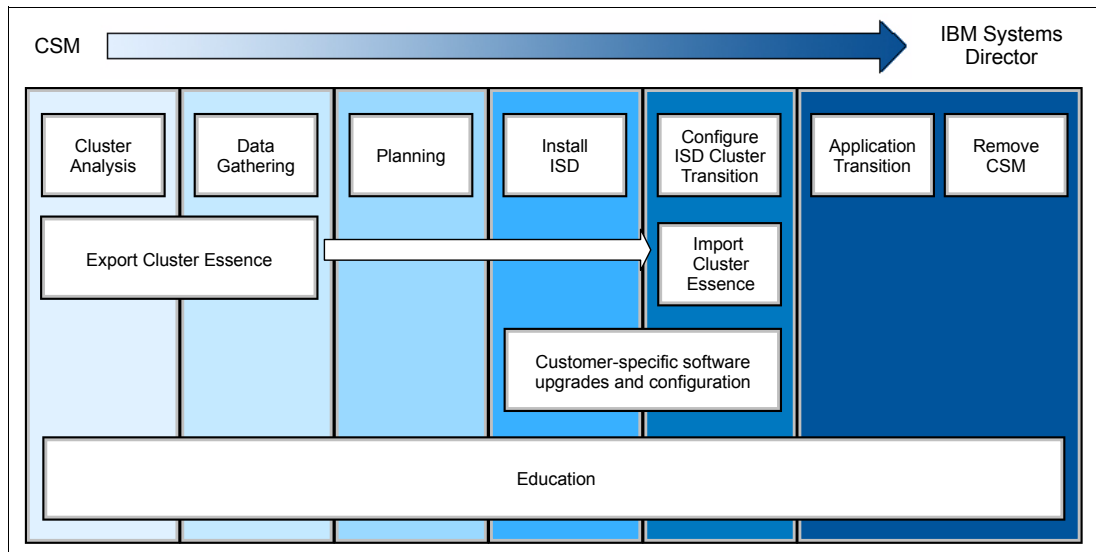


Figure 3-1 Transformation timeline

Note: It is imperative to collect CSM cluster information and import that information to IBM Systems Director during the transformation, as shown in the Figure 3-1.

3.2 Coexistence

In this scenario we lead you through the transformation from IBM CSM to IBM Systems Director where both the CSM Management Server and IBM Systems Director Server coexist on the same server. This scenario is more GUI driven and you can find many screen shots to help you get familiar with the look of IBM Systems Director.

3.2.1 The cluster environment

Figure 3-2 on page 44 represents the current hardware and software environment configured for the coexistence scenario.

The hardware and software levels used for our environment include:

- ▶ The CSM management server is installed on an AIX LPAR.
- ▶ Three additional client LPARs running on the POWER5 server.
- ▶ One client LPAR running on POWER6 server.
- ▶ The Hardware Management Console (HMC) is installed, operational, and managing all systems.
- ▶ The required network infrastructure is in place and functional.
- ▶ The Network Install Manager (NIM) server is installed on the CSM Management Server LPAR.

Table 3-1 on page 43 lists the details of the servers including function, software levels, and hostname.

Table 3-1 Components in the coexistence scenario

	System	CSM/SD Level	OS Level	Hostname
1	CSM/NIM/ISD	CSM-1.7.1.12 SD - 6.2.1.2	AIX 6.1	p5570lp01
2	HMC-1	N/A	7.7.2.0	hmc1
3	Client1	6.2.1.2	AIX 5.3	p5570lp02
4	Client2	CSM-1.7.1.12 SD - 6.2.1.2	AIX 6.1	p5570lp03
5	Client3	CSM-1.7.1.12 SD - 6.2.1.2	AIX 6.1	p5570lp04
6	HMC-2		7.7.2.0	hmc2
7	Client4	CSM-1.7.1.12 SD - 6.2.1.2	AIX 6.1	p6lp01
8	VIO	NA	2.2.0.11	p5570vio1

Figure 3-2 on page 44 shows a graphical representation of the environment. The environment consists of a CSM server that manages nodes from POWER5 and POWER6 servers. The hardware management console for the POWER5 server is HMC1 and for the POWER6 server it is HMC2. HMC2 and VIO1 are added as hardware devices to the CSM server.

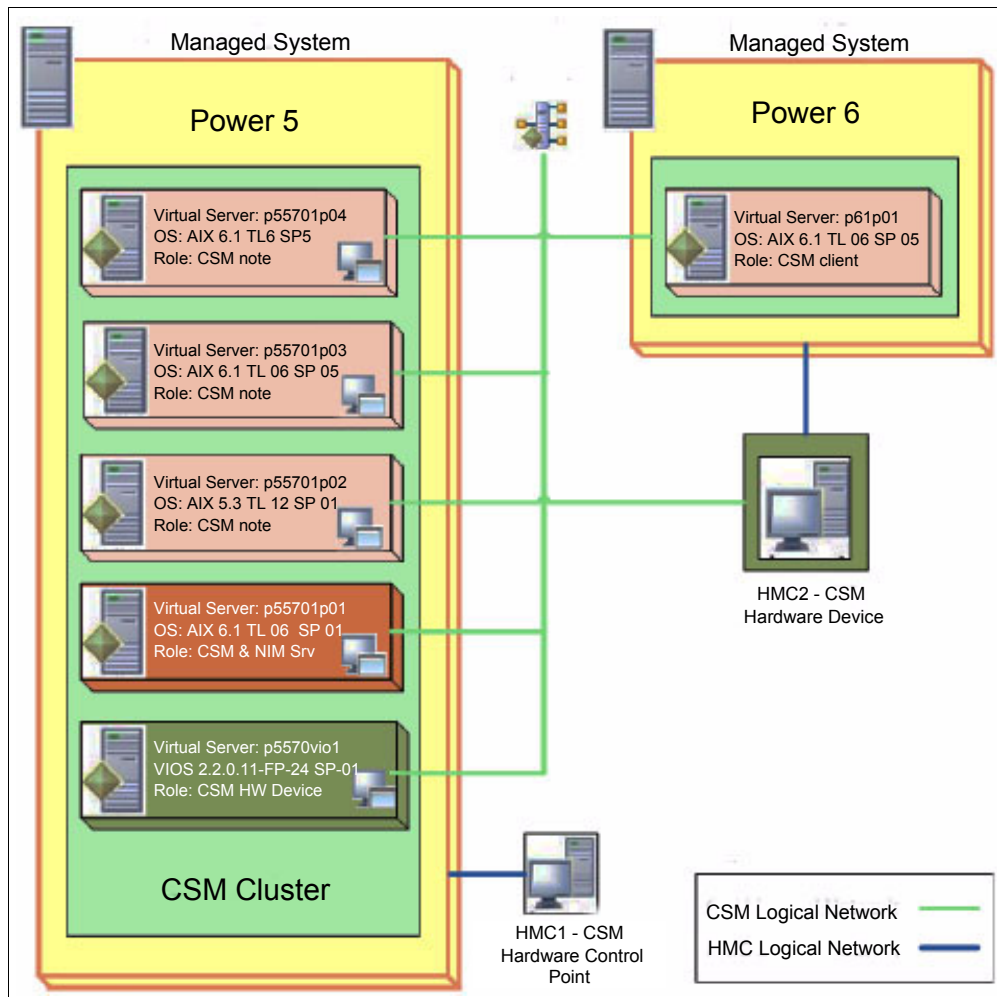


Figure 3-2 CSM environment for the coexistence scenario

3.2.2 Cluster verification

Before starting the transformation, it is important to perform functional verification to ensure your cluster is working correctly.

Note: CSM executable commands can be found in the `/opt/csm/bin` directory. That directory is in the `PATH` in the following examples, as typically recommended during CSM implementation. Therefore, full path names are not used.

- List the nodes that are managed by CSM, as shown in Example 3-1.

Example 3-1 List the CSM-managed nodes

```
p55701p01(root)/> lsnode
p55701p02
p55701p03
p55701p04
p61p01
p55701p01(root)/>
```

The 3 nodes p5570lp02, p5570lp03, p5570lp04 are the POWER5 LPARs and node p6lp01 is the POWER6 LPAR.

- List the hardware devices that are managed by CSM as shown in Example 3-2.

Example 3-2 List all the CSM-managed hardware devices

```
p5570lp01(root)/> lshwdev
hmc2
p5570vio1
p5570lp01(root)/>
```

The HMC that manages the POWER6 server is hmc2 and p5570vio1 is the VIO Server for the POWER5 server.

- The node group *CSM1_nodegroup* contains all the CSM-managed nodes. See Example 3-3.

Example 3-3 Show the members of the CSM node group

```
p5570lp01(root)/> nodegrp CSM1_nodegroup
p5570lp02
p5570lp03
p5570lp04
p6lp01
p5570lp01(root)/>
```

- The **csmdat** command displays the status of each node, as shown in Example 3-4.

Example 3-4 Show the status of the CSM-managed nodes

```
p5570lp01(root)/> csmdat
```

Hostname	HWControlPoint	Status	PowerStatus	Network-Interfaces
p5570lp02	hmc1	on	on	en0-Online en1-Online
p5570lp03	hmc1	on	on	en0-Online en1-Online
p5570lp04	hmc1	on	on	en0-Online
p6lp01	hmc2	on	on	en0-Online

```
p5570lp01(root)/>
```

In many CSM environments, customer utilize the distributed condition-response capabilities. We have created a condition and a response to demonstrate this.

- The condition *Filesystem_Tmp_Monitor1* is defined for the node p5570lp02. This condition sends an email to the root user when the /tmp file system utilization exceeds 90%. The condition gets rearmed when the percentage utilization drops below 75%. Refer to Example 3-5.

Example 3-5 List the details of the distributed RMC condition

```
p5570lp01(root)/> lscondition Filesystem_Tmp_Monitor1
Displaying condition information:
```

```
condition 1:
    Name           = "Filesystem_Tmp_Monitor1"
    Node           = "p5570lp01"
    MonitorStatus  = "Monitored"
    ResourceClass  = "IBM.FileSystem"
    EventExpression = "PercentTotUsed>90"
```

```

EventDescription      = ""
RearmExpression       = "PercentTotUsed<75"
RearmDescription      = ""
SelectionString       = "Name==\"/tmp\"/"
Severity              = "i"
NodeNames             = {"p55701p03"}
MgtScope              = "m"
Toggle                = "Yes"
EventBatchingInterval = 0
EventBatchingMaxEvents = 0
BatchedEventRetentionPeriod = 0
BatchedEventMaxTotalSize = 0
RecordAuditLog        = "ALL"
p55701p01(root)/>

```

- The response E-mail root anytime for the above condition is listed in Example 3-6.

Example 3-6 Show the details of the distributed RMC response

```

p55701p01(root)/> lsresponse "E-mail root anytime"
Displaying response information:

```

```

ResponseName      = "E-mail root anytime"
Node              = "p55701p01"
Action            = "E-mail root"
DaysOfWeek        = 1-7
TimeOfDay         = 0000-2400
ActionScript      = "/usr/sbin/rsct/bin/notifyscript root"
ReturnCode        = -1
CheckReturnCode   = "n"
EventType         = "b"
StandardOut       = "n"
EnvironmentVars   = ""
UndefRes          = "n"
EventBatching     = "n"
p55701p01(root)/>

```

- The **lscondresp** command shows the configured relationship between the condition and the response. Refer to Example 3-7.

Example 3-7 Show the configured response to the condition in distributed RMC

```

p55701p01(root)/> lscondresp
Displaying condition with response information:

```

Condition	Response	Node	State
"Filesystem_Tmp_Monitor1"	"E-mail root anytime"	"p55701p01"	"Active"

- The `/cfmroot` director contains configuration files such as `/etc/hosts`, `/etc/passwd`, and `/etc/group`. These configuration files are synchronized across the cluster nodes using the Cluster File Management (CFM) capability in CSM. The update can be done using the command **cfmupdatenode**.

3.2.3 Transformation to IBM Systems Director

In this scenario the IBM Systems Director Server is installed on the same LPAR as the CSM Management Server. The following section demonstrates the step-by-step procedure. As

mentioned before, we do not document the details of installing IBM Systems Director Server on AIX.

1. Update IBM Systems Director to the latest level.

Note: The latest updates can be downloaded from the IBM Support Portal at:

<http://www-933.ibm.com/support/fixcentral/>

2. Start the IBM Systems Director if not already started, as shown in Example 3-8.

Example 3-8 Command to start the IBM Systems Director Server

```
#/opt/ibm/director/bin/smstart
```

3. Check the status of the IBM Systems Director Server, as shown in Example 3-9.

Example 3-9 Command to check the status of IBM Systems Director Server

```
p55701p01(root)/> smstatus  
Active  
p55701p01(root)/>
```

4. Log in to the web interface as the root user and verify the version. Use the browser to open the IBM Systems Director using the format shown in:

<https://<ipaddress>:8422/ibm/console>

Note: “8422” is the default secure port used by IBM Systems Director. Use the appropriate port if it has been changed in your environment.

Figure 3-3 shows the screen shot of the IBM Systems Director Welcome page. The version of the Systems Director Server code can be seen to the right (here it is 6.2.1.2).

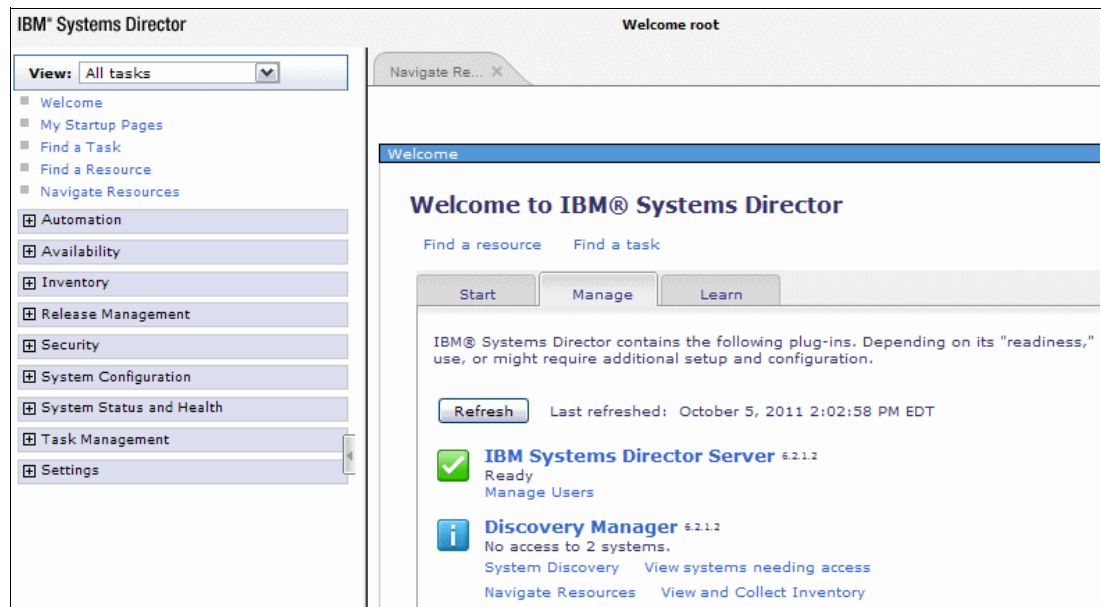


Figure 3-3 IBM Systems Director Welcome page

5. To view the systems managed by IBM Systems Director, use the left-hand navigation area to select **Navigate Resources** → **All Systems**.

In Figure 3-4, only the LPAR installed with IBM Systems Director is listed because no other LPARs have been discovered yet.

It shows two objects that represent the same LPAR, Virtual Server and Operating System. Virtual Server represents the hardware container (LPAR), and Operating System represents the operating system software instance.

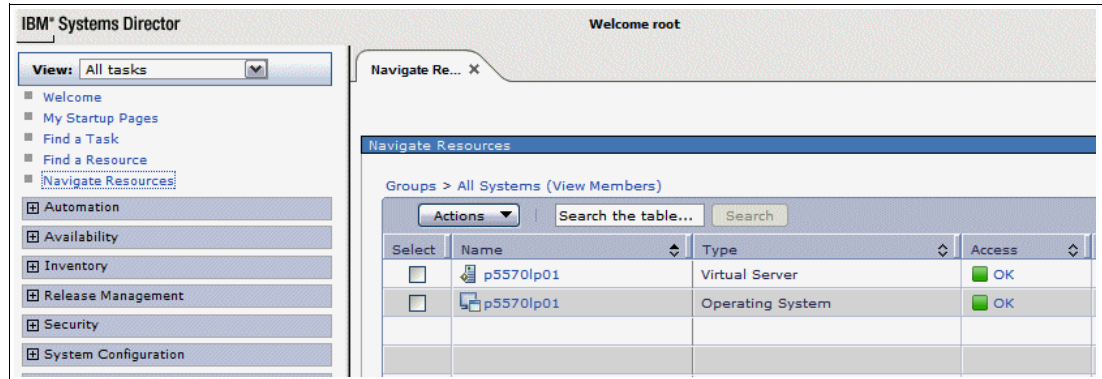


Figure 3-4 Hardware and OS instance in IBM Systems Director

- CSM uses NIM to install the AIX operating system and to update software on client LPARs. Similarly, IBM Systems Director can also use a NIM server to install and update the AIX operating system. Another way to install and update the OS is by using Storage Copy Services with VMControl Standard Edition (also called Power Storage Based Provisioning).

Notes:

- High-level information about Storage Copy Service can be found in “VMControl Storage Copy Services” on page 117.
- The IBM Systems Director Common Agent is responsible for communication between the clients and the IBM Systems Director Server. The Common Agent on the server should be at the highest level of code across the client LPARs.

- Without VMControl Standard Edition, IBM Systems Director cannot interact with a NIM Server for AIX operating system installation. Therefore, the next step is to install the VMControl plug-in.

Note: The VMControl installation steps can be found at:

http://publib.boulder.ibm.com/infocenter/director/v6r2x/topic/com.ibm.director.vim.helps.doc/fsd0_vim_pdf.pdf.

The plug-in can be found at:

<http://www-03.ibm.com/systems/software/director/vmcontrol/>

Once the plug-in is installed, we need to install the VMControl subagent for NIM on the NIM Server. Navigate through **Release Management** → **Agents task** on the left-hand navigation bar. Then click **All Agent Packages** as shown in Figure 3-5 on page 49.

Select the **CommonAgentSubagent_VMControl_NIM-2.3.1**, click **Install Agent** and follow the wizard.

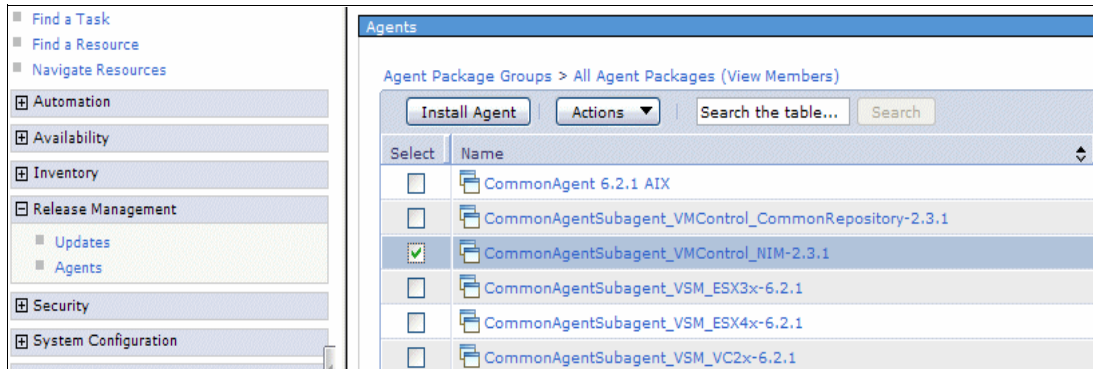


Figure 3-5 Installation of the VMControl subagent for NIM

8. Discover the Hardware Management Console and client LPARs in IBM Systems Director.

Before discovering the operating system client endpoints, we recommend that you discover the HMC.

To discover the HMC, navigate using **Inventory** → **System Discovery** as shown in Figure 3-6.

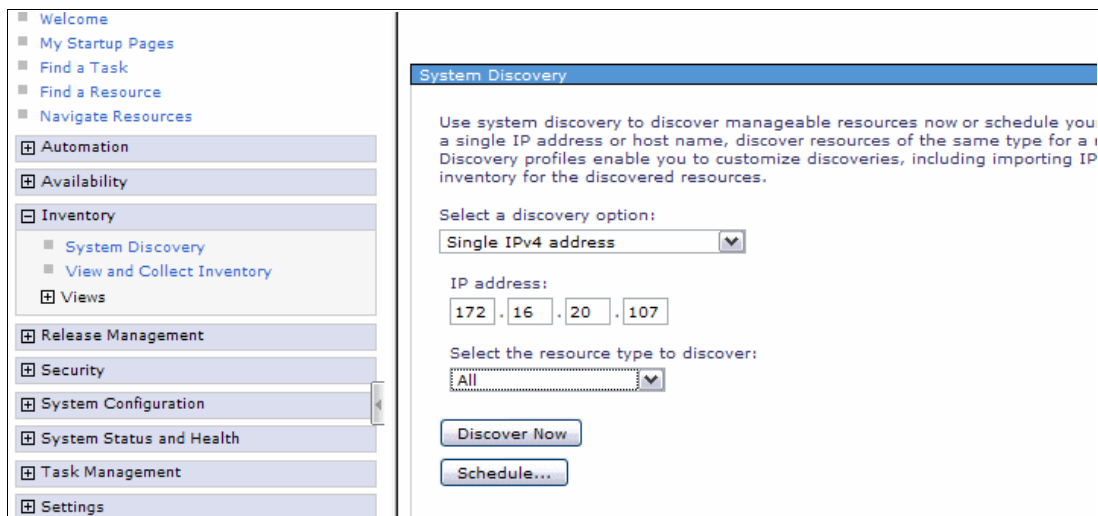


Figure 3-6 HMC discovery

Note: Discovery does not cause any changes on the target systems.

After System Discovery, provide access to the discovered HMC. First, use the left-hand navigation bar to **Navigate Resources** → **All Systems**. Page through the table to find the appropriate HMC. IBM Systems Director creates two objects for the HMC:

- Hardware Management Console
- Operating System

9. Provide access to the Hardware Management Console by clicking **No access** as shown in Figure 3-7 on page 50. Enter the user ID (hscroot or other similar authority user) and password.

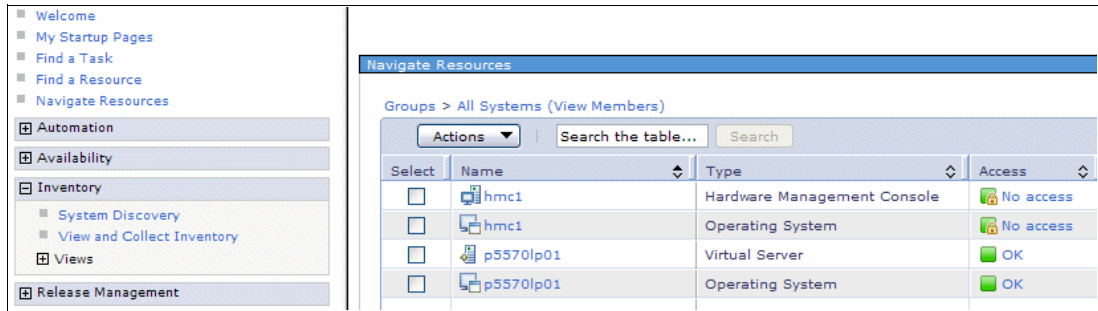


Figure 3-7 Access to the HMC

Discovering the HMC also discovers all the virtual servers connected to HMC, as shown in Figure 3-8 on page 50.

Notes:

- ▶ If you have an HMC that manages many Power Servers, and you do not want to discover all of them in IBM Systems Director, then you can add a user with a specific resource role containing only the servers you choose. This allows your user to manage only the systems you have granted access to.
- ▶ A Virtual Server is a hardware container instance for an operating system. In Power Systems terms, an LPAR. Refer to Figure 3-8.

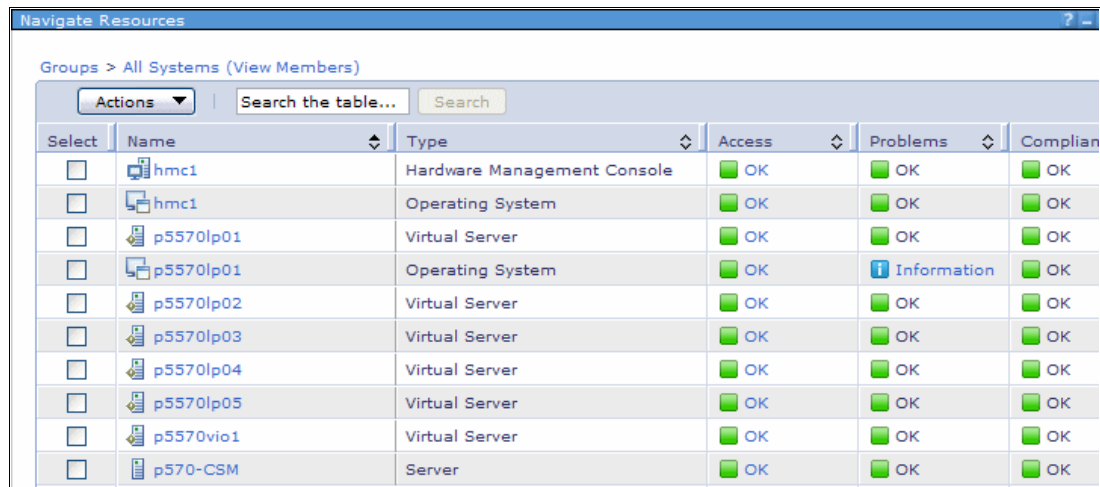


Figure 3-8 Virtual servers discovered after HMC discovery

10. After providing authentication to the HMC, we need to collect inventory,, which can be accomplished via right-clicking the HMC representation and choosing the task **Inventory** → **View and Collect Inventory** as shown in Figure 3-9 on page 51.

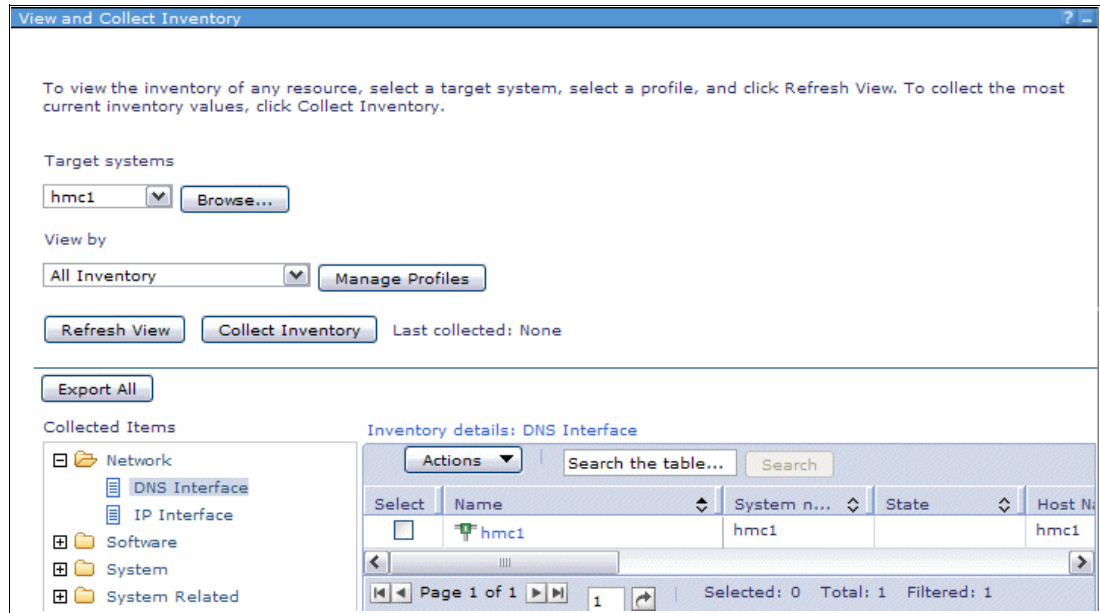


Figure 3-9 Inventory collection of the HMC

11. The next step is to discover the operating systems running inside the virtual servers. In our example we start with endpoint “p5570lp02”.

Using the left-hand navigation bar, expand **Inventory** and click **System Discovery**. From here, enter the IP address of the system that needs to be discovered, as shown in Figure 3-10.

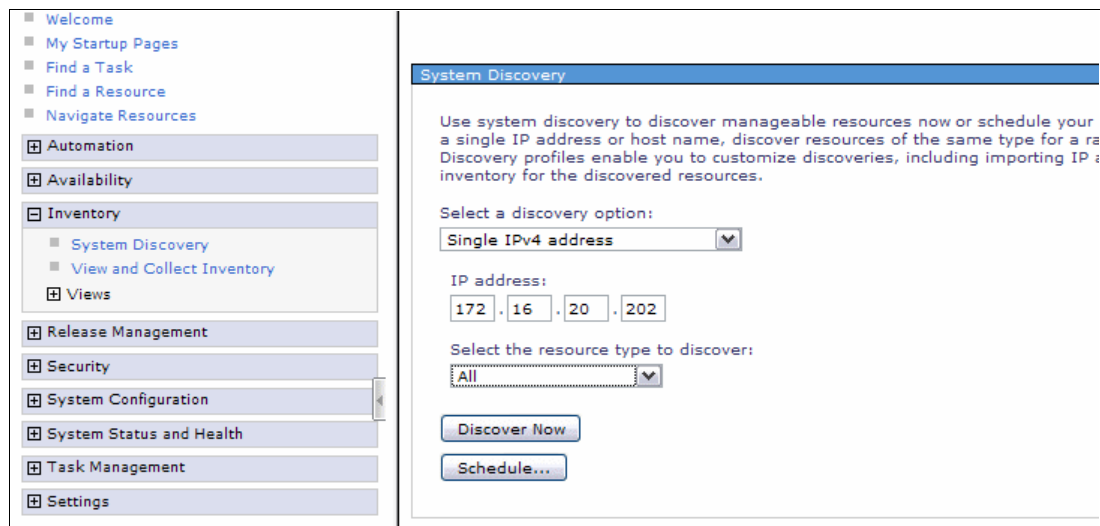


Figure 3-10 Operating system discovery

12. After you have discovered the operating systems, provide the authentication credentials. As shown in Figure 3-11 on page 52, the Access status appears OK for Operating System and Virtual Server instances of client LPARp5570lp2.


Navigate Resources			
Groups > All Systems (View Members)			
<div> <div>Actions ▾</div> <div>Search the table...</div> <div>Search</div> </div>			
Select	Name	Type	Access
<input type="checkbox"/>	 hmc1	Hardware Management Console	 OK
<input type="checkbox"/>	 hmc1	Operating System	 OK
<input type="checkbox"/>	 p5570lp01	Virtual Server	 OK
<input type="checkbox"/>	 p5570lp01	Operating System	 OK
<input type="checkbox"/>	 p5570lp02	Virtual Server	 OK
<input type="checkbox"/>	 p5570lp02	Operating System	 OK

Figure 3-11 :Client LPAR access status

13. The next step is to “View and collect inventory” of the operating system. Inventory collection collects information related to the operating system such as OS level, software filesets installed, network information, and more, and adds it to the database. Some of this information can be viewed on the object’s Properties page.client LPAR Properties.
14. Discover the operating system instances of all the client LPARs, provide access and collect inventory information for each.
15. Next we added hmc2, which manages client LPAR “p6lp01” to IBM Systems Director, granted access and collected inventory as shown in list item 8 on page 49. After you discover the Operating System instance for client LPAR “p6lp01”, provide access and collect inventory as shown in list item 11 on page 51.

At this point you have added all the nodes that were managed by CSM to IBM Systems Director. All the client endpoints are managed by both CSM and the IBM Systems Director, as shown in Figure 3-12 on page 53.




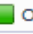



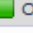



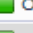



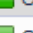











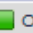



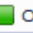



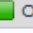



















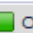



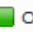

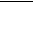

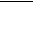




Navigate Resources						
Groups > All Systems (View Members)						
<div> <div>Actions</div> <div>Search the table...</div> <div>Search</div> </div>						
Select	Name	Type	Access	Problems	Compliance	
<input type="checkbox"/>	 hmc1	Hardware Management Console	 OK	 OK	 OK	
<input type="checkbox"/>	 hmc1	Operating System	 OK	 OK	 OK	
<input type="checkbox"/>	 hmc4	Hardware Management Console	 OK	 OK	 OK	
<input type="checkbox"/>	 hmc4	Operating System	 OK	 OK	 OK	
<input type="checkbox"/>	 p5570lp01	Virtual Server	 OK	 OK	 OK	
<input type="checkbox"/>	 p5570lp01	Operating System	 OK	 Information	 OK	
<input type="checkbox"/>	 p5570lp02	Operating System	 OK	 OK	 OK	
<input type="checkbox"/>	 p5570lp02	Virtual Server	 OK	 OK	 OK	
<input type="checkbox"/>	 p5570lp03	Virtual Server	 OK	 OK	 OK	
<input type="checkbox"/>	 p5570lp03	Operating System	 OK	 Information	 OK	
<input type="checkbox"/>	 p5570lp04	Operating System	 OK	 Information	 OK	
<input type="checkbox"/>	 p5570lp04	Virtual Server	 OK	 OK	 OK	
<input type="checkbox"/>	 p5570vio1	Operating System	 OK	 OK	 OK	
<input type="checkbox"/>	 p5570vio1	Virtual Server	 OK	 OK	 OK	
<input type="checkbox"/>	 p570_170	Server	 OK	 OK	 OK	
<input type="checkbox"/>	 p6lp01	Virtual Server	 OK	 OK	 OK	
<input type="checkbox"/>	 p6lp01	Operating System	 OK	 Information	 OK	

Figure 3-12 Endpoints managed by IBM Systems Director

Figure 3-13 on page 54 shows that all the nodes and client endpoints are managed by both CSM and IBM Systems Director.

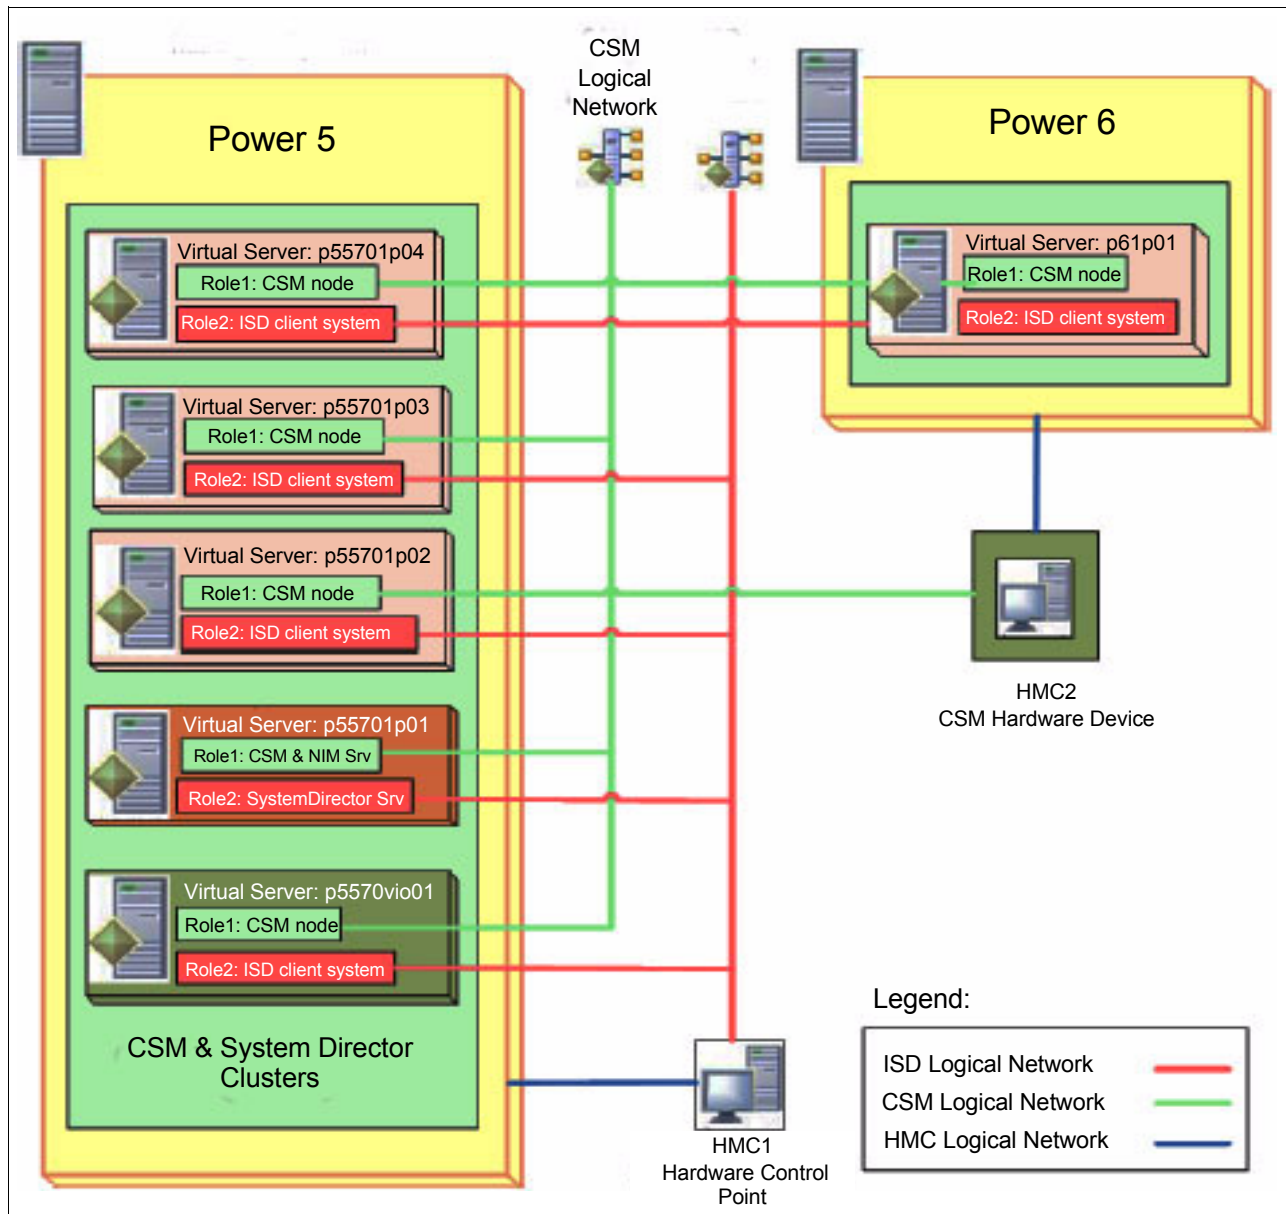


Figure 3-13 Client LPARs managed by CSM and IBM Systems Director

3.2.4 Testing the coexistence scenario

Now that the nodes and client LPARs are visible from both CSM and IBM Systems Director, it is time to test some of the commonly available features on both CSM and IBM Systems Director.

In this section, we test three of the features:

- ▶ Power control
- ▶ Serial console
- ▶ Distributed command

Power Control

In Example 3-10 we power off and on a node using CSM and check the node and client LPAR status on both CSM and IBM Systems Director.

Example 3-10 shows the power status of all the nodes from the CSM server.

Example 3-10 Query the power status of nodes

```
p5570lp01(root)/> rpower -a query
p5570lp02 on
p6lp01 on
p5570lp03 on
p5570lp04 on
p5570lp01(root)/>
```

The command shown in Example 3-11 powers off the node p5570lp03.

Example 3-11 Power off the node

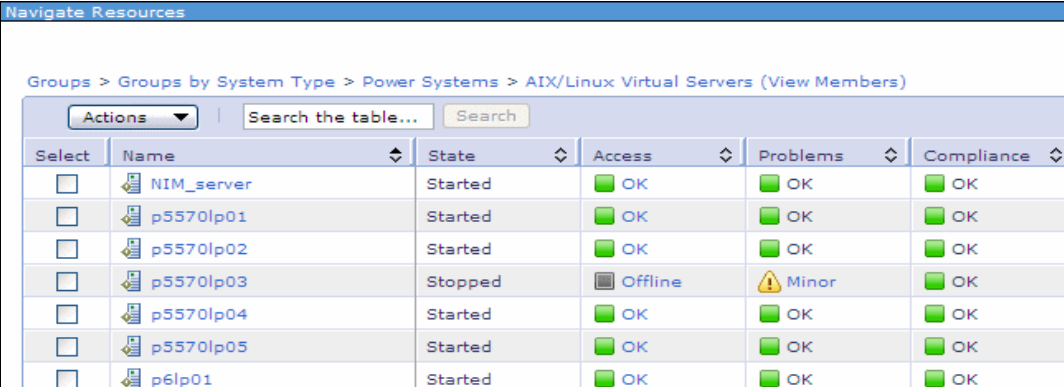
```
p5570lp01(root)/> rpower -n p5570lp03 off
p5570lp03 off complete rc=0
p5570lp01(root)/>
```

To see the status of the nodes, run the query again. As you can see, the power status of the node p5570lp03 is off, as shown in Example 3-12.

Example 3-12 Query the node status after power off

```
p5570lp01(root)/> rpower -a query
p6lp01 on
p5570lp02 on
p5570lp03 off
p5570lp04 on
p5570lp01(root)/>
```

The status of client LPAR p5570lp3 from IBM Systems Director is shown in Figure 3-14.



Navigate Resources						
Groups > Groups by System Type > Power Systems > AIX/Linux Virtual Servers (View Members)						
Actions Search the table... Search						
Select	Name	State	Access	Problems	Compliance	
<input type="checkbox"/>	NIM_server	Started	OK	OK	OK	
<input type="checkbox"/>	p5570lp01	Started	OK	OK	OK	
<input type="checkbox"/>	p5570lp02	Started	OK	OK	OK	
<input type="checkbox"/>	p5570lp03	Stopped	Offline	Minor	OK	
<input type="checkbox"/>	p5570lp04	Started	OK	OK	OK	
<input type="checkbox"/>	p5570lp05	Started	OK	OK	OK	
<input type="checkbox"/>	p6lp01	Started	OK	OK	OK	

Figure 3-14 Power control status of endpoints

Note: It takes a few minutes to reflect the power status of the client LPAR on IBM Systems Director.

Now initiate the Power On task from the IBM Systems Director Server, as shown in Figure 3-15.

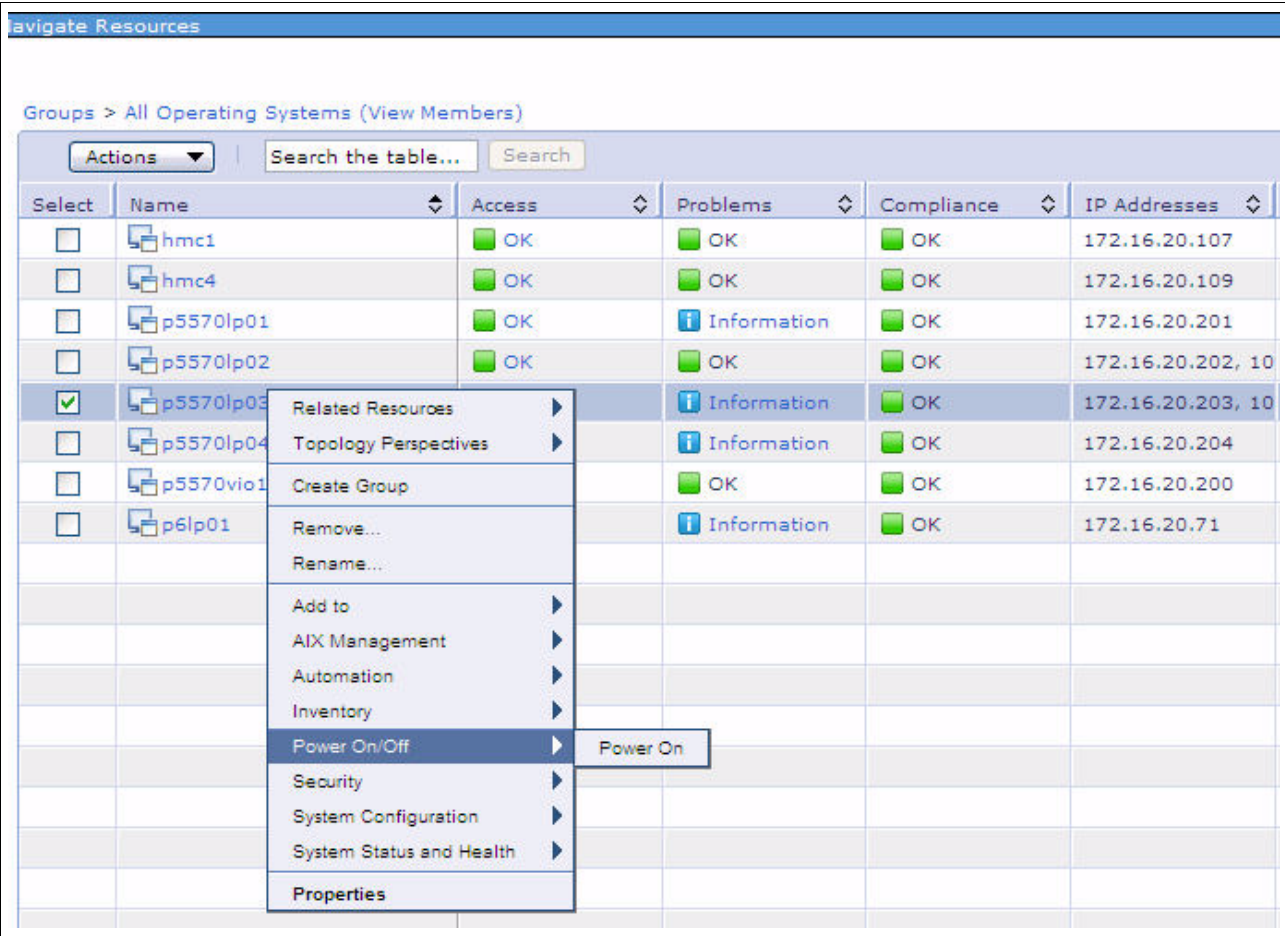


Figure 3-15 Power on from IBM Systems Director

The status of the power change is reflected on both CSM and IBM Systems Director, as shown in Example 3-13.

Example 3-13 Node status after power on from IBM Systems Director

```
p5570lp01(root)/> rpower -a query
p6lp01 on
p5570lp02 on
p5570lp03 on
p5570lp04 on
p5570lp01(root)/>
```

We also tested Power Off from IBM Systems Director (see Figure 3-16 on page 56), and Power On from the CSM server for the same node and client.

<input type="checkbox"/>	p5570lp02	Operating System	OK	OK	OK
<input type="checkbox"/>	p5570lp03	Operating System	OK	OK	OK
<input type="checkbox"/>	p5570lp03	Virtual Server	OK	OK	OK

Figure 3-16 Power Status on IBM Systems Director

Serial console

Next, we verified that we can open the remote console from both CSM and the IBM Systems Director Server.

Open the remote console from the CSM server, as shown in Example 3-14.

Example 3-14 Command to open the serial console

```
rconsole -t -n p6lp01
```

To open the console from IBM Systems Director Server, right-click the mouse on the system, and navigate through **System Configuration** → **Remote Access** → **Serial Console**, as shown in Figure 3-17.

Note: The ability to open a remote serial console via the HMC/IVM in IBM Systems Director is only available on an AIX Systems Director Server with the `dsm.core` fileset installed.

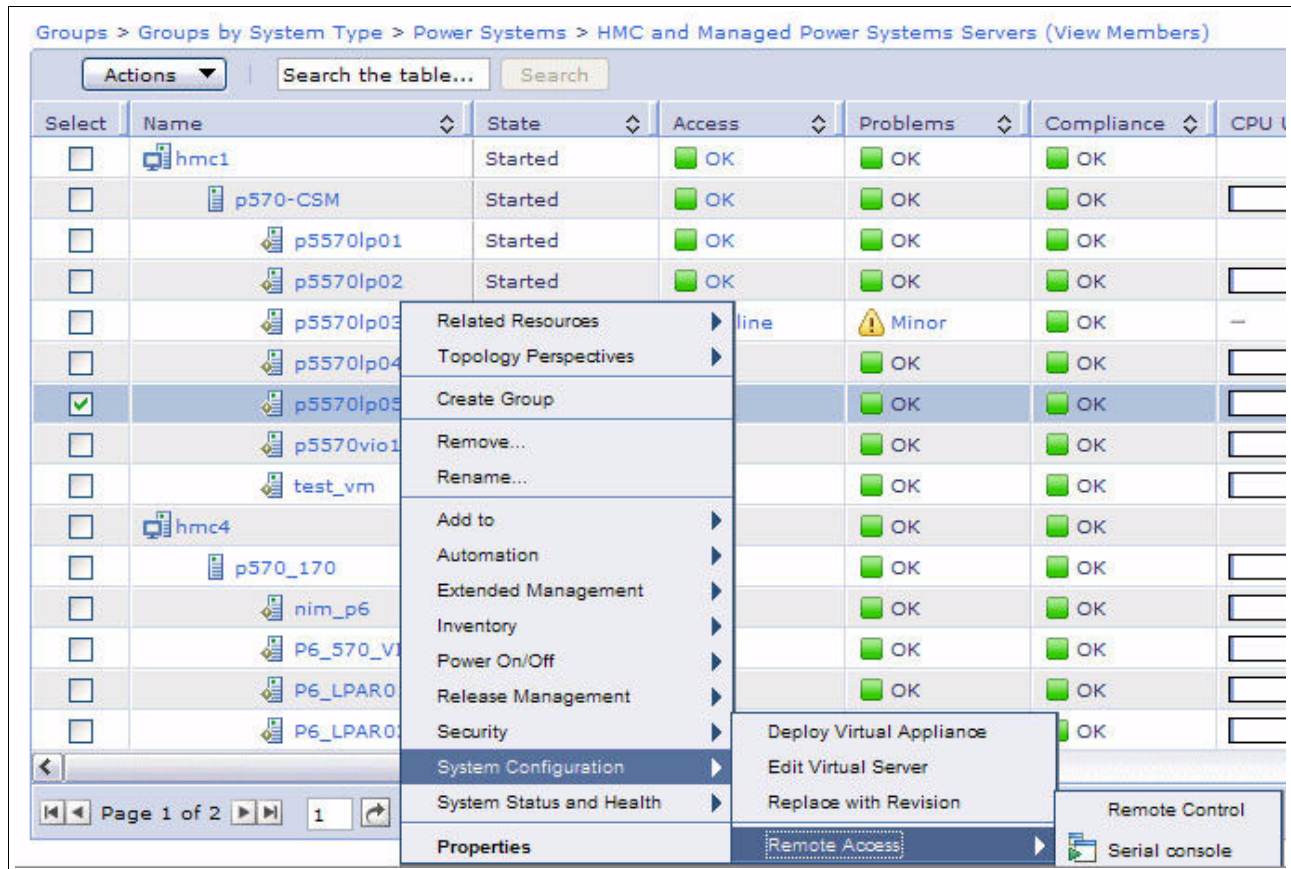


Figure 3-17 Remote console from IBM Systems Director

Distributed commands

Example 3-15 shows testing distributed commands on both CSM and IBM Systems Director.

To run **dsh** from the AIX SMIT menu, use `smitty csm` → Additional CSM commands → CSM Cluster Command (**dsh**), input command to run, select all the node names and press Enter. In our example we run data commands on all four nodes.

Example 3-15 dsh distributed command output

COMMAND STATUS

Command: OK stdout: yes stderr: no

Before command completion, additional instructions may appear below.

```
p5570lp03: Fri Oct 7 10:33:23 EDT 2011
p6lp01: Fri Oct 7 10:33:22 EDT 2011
p5570lp02: Fri Oct 7 10:33:31 EDT 2011
p5570lp04: Fri Oct 7 10:33:41 EDT 2011
Press Enter to continue.
```

To utilize the remote command capability from IBM Systems Director, navigate through **System Configuration** → **Remote Access** → **Distributed command**. Select the targets (client LPARs) and input the commands. The output of the command is shown in Figure 3-18.




Figure 3-18 Output of dsh in IBM Systems Director

3.3 Migration scenario

IBM Systems Director is not only a distributed systems management and monitoring solution like IBM Cluster Systems Management for AIX. It has many other functions that will play an important role in your infrastructure because of the many plug-ins. At the time of writing this book, there are nine plug-ins available; additional plug-ins might be available in later releases.

The base resource requirements for IBM Systems Director in comparison to CSM are higher and can become larger as you add additional plug-ins. In many cases, the CSM Management Server is placed together with other LPARs on an old POWER5 system. That class of server may not have enough memory and processor resources to effectively run the IBM Systems Director Server, so it is advisable to install it on a new LPAR on a more powerful system.

In addition to a separate LPAR for the Systems Director Server, we are installing and configuring a new NIM Server on a new LPAR distinct from the IBM Systems Director Server LPAR. We have chosen this because we are going to use this NIM server for deploying and

updating AIX 7.1 in our environment. For these operations, the NIM Server may require frequent operating system updates because it has to be running, at least, the highest level of AIX it is deploying or updating. These possible frequent changes make it better to keep it on a different LPAR.

Note: There is no restriction on installing the NIM Server on the same LPAR as IBM Systems Director, but installing on a different LPAR provides greater flexibility.

This scenario has two main subsections: preparation and effective migration, as shown in Figure 3-19.

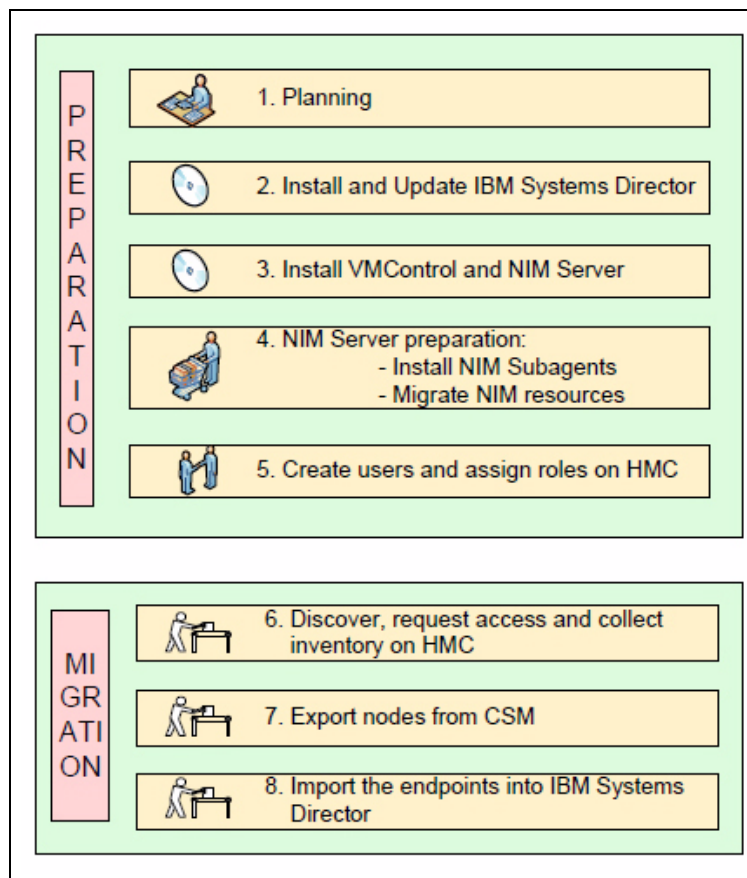


Figure 3-19 Scenario steps - preparation and migration

3.3.1 Initial environment

The laboratory environment where this scenario was tested is comprised of the following components:

- ▶ CSM management server version 1.7.1, which manages three AIX nodes located as follows: two LPARs on a POWER6 server and one LPAR on a POWER7 server.
- ▶ Hardware Management Console (HMC1) version 7.7.2 SP2, used as a Hardware Control Point for POWER6 nodes in CSM. Also, this HMC is added as a Hardware Device in CSM.
- ▶ Hardware Management Console (HMC2) version 7.7.2 was not defined to CSM and will not be defined in IBM Systems Director. We are going to migrate only one LPAR managed by this HMC and nothing more.

- ▶ Old NIM Server, co-located on the same LPAR as the CSM Management Server. It has already defined NIM resources such as nodes, mksysb, lpp_source, and so on.
- ▶ The New NIM Server, located on a new LPAR.
- ▶ IBM Systems Director server version 6.2. This system will be updated to a later version of IBM Systems Director in the next sections and have the IBM Systems Director VMControl plug-in installed.
- ▶ The required network infrastructure (see Figure 3-20) is in place and functional.

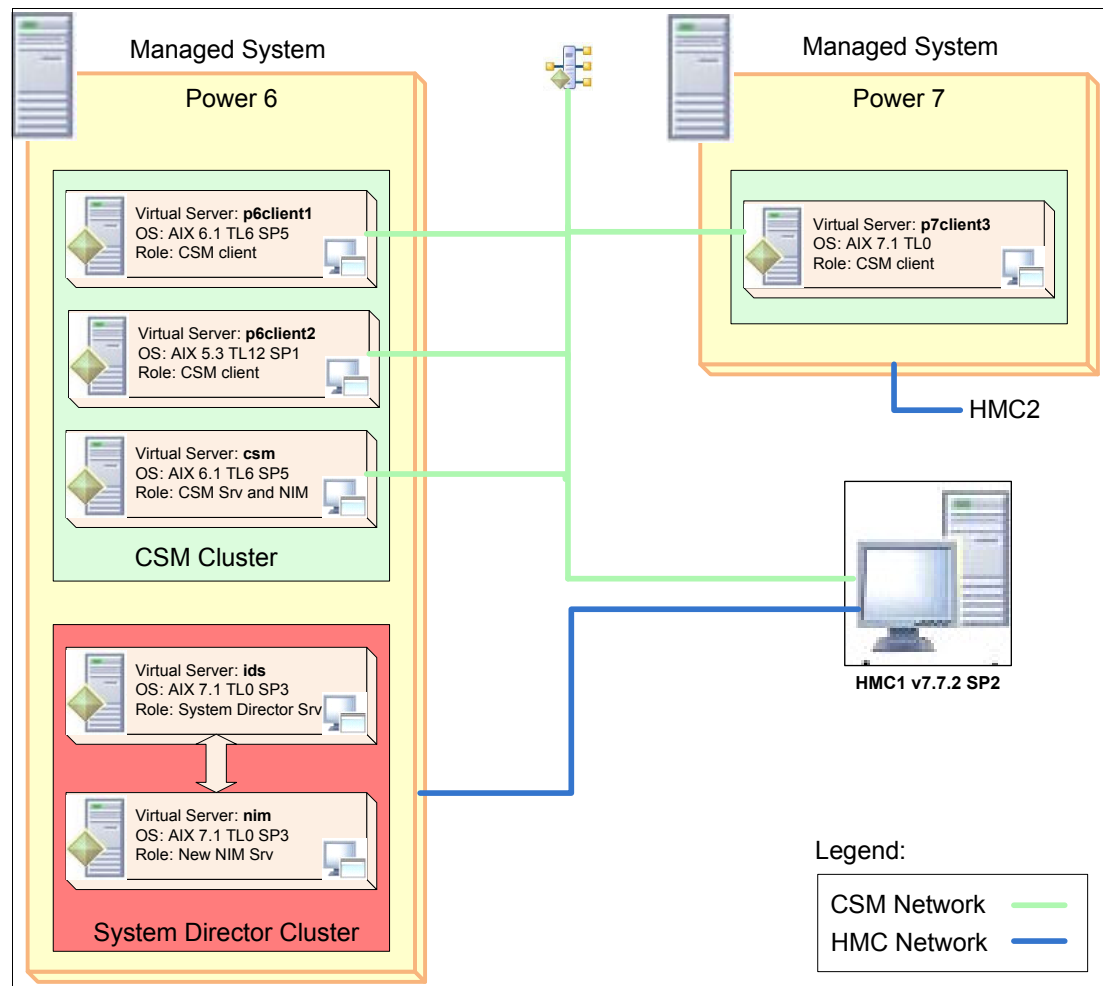


Figure 3-20 Initial setup of the migration scenario

The nodes that are managed by the CSM cluster can be seen in the output shown in Example 3-16 on page 60.

Example 3-16 Output of the csmstat command on the CSM cluster

```
[csm:/]# csmstat
```

Hostname	HWControlPoint	Status	PowerStatus	Network-Interfaces
p6client1	hmc1	on	on	en0-Online
p6client2	hmc1	on	on	en0-Online

3.3.2 The target environment

After going through the preparation and executing the documented migration steps, our environment is shown in Figure 3-21 on page 61. All the nodes and hardware devices defined in the initial setup in CSM are now defined as endpoints to IBM Systems Director. The CSM Management Server has no connection to its old nodes or hardware devices and can be decommissioned.

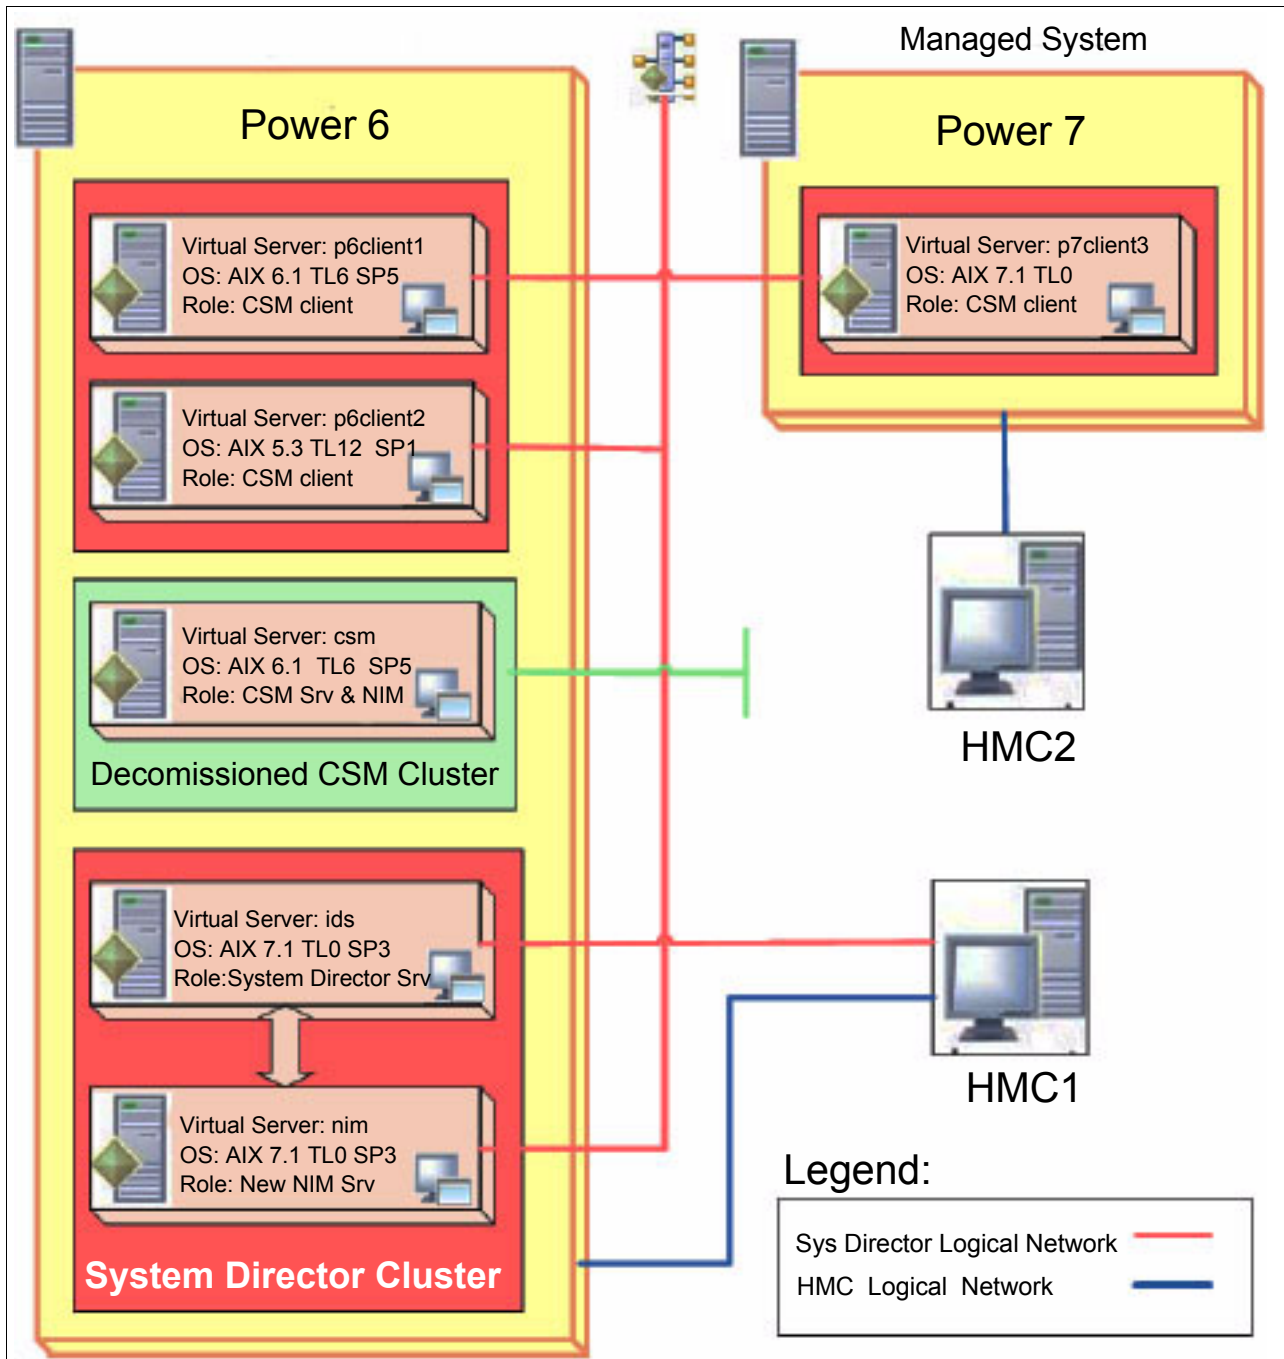


Figure 3-21 Target environment of the migration scenario

3.3.3 Preparation

This section explains the preparation steps for the transformation of the clusters from a CSM management environment to an IBM Systems Director managed clustered environment.

Planning for transforming the CSM cluster

One of the most important aspects of transforming the CSM cluster to IBM Systems Director is the planning. For a clustered environment with diverse hardware, different versions of operating systems and possibly many networks, planning is even more important. For more details about planning considerations, refer to Chapter 2, “Planning and preparation” on page 7. Additional information can be found in the InfoCenter at:

http://publib.boulder.ibm.com/infocenter/director/v6r2x/index.jsp?topic=%2Fcom.ibm.director.plan.helps.doc%2Ffqm0_t_selecting_how_to_install_ibm_director.html

Install and update IBM Systems Director

In this scenario we are not covering a step-by-step installation of IBM Systems Director. We are just documenting the steps you might use to transform your CSM cluster. For installation details, see the Installation Manual located in the InfoCenter at:

http://publib.boulder.ibm.com/infocenter/director/v6r2x/index.jsp?topic=%2Fcom.ibm.director.plan.helps.doc%2Ffqm0_t_selecting_how_to_install_ibm_director.html

After the installation of IBM Systems Director, we recommend updating it to the latest version. The latest fixes can be found on the IBM Support Portal at:

<http://www.ibm.com/support/fixcentral>

Note: Because of the way IBM Systems Director is updated, version information displayed by the `ls1pp` command may not always be accurate. We recommend using the `smcli lsver` command to determine the currently installed version of IBM Systems Director.

If you do not have the latest updates for IBM Systems Director, click the **Update IBM Systems Director** link on the Welcome page, or follow these steps:

- ▶ Download the update package to the management server from FixCentral.
- ▶ Unpack the package file to a local directory.
- ▶ Check the /opt file system for at least 2 GB of free space.
- ▶ Remove existing updates for IBM Systems Director:
`smcli cleanupd -mFv -P Platform=Director`
- ▶ Run the following command to install the updates:
`smcli installneeded -v -F <full path to extracted updates directory>`
- ▶ Restart IBM Systems Director:
`smstop; smstart; smstatus -r`
- ▶ Check to see if you have the latest version installed:
`smcli lsver`

Install the IBM Systems Director VMControl plug-in and NIM Server

In order for IBM Systems Director to update the operating system on managed endpoints or to deploy virtual appliances to new virtual servers (LPARs), it requires two capabilities: the IBM Systems Director VMControl plug-in (Standard or Enterprise Edition), and an AIX NIM Server.

The NIM Server is an image repository capable of importing, capturing, deploying, and deleting AIX virtual appliances. A NIM-based virtual appliance can be either an `mksysb` image or `lpp_source` resource.

IBM Systems Director VMControl plug-in is a cross-platform suite of products that assists with rapid deployment of virtual appliances to create virtual servers that are configured with an operating system and software applications. It also enables you to group resources into server system pools, which enables you to centrally manage and control the different workloads in your environment.

A step-by-step installation of the AIX NIM Server and VMControl plug-in is beyond the scope of this book.

For instructions about installation of an AIX NIM Server, see the InfoCenter at:

http://publib.boulder.ibm.com/infocenter/aix/v6r1/topic/com.ibm.aix.install/doc/insgdrf/nim_basic.htm

For instructions about installation of the VMControl plug-in, see the InfoCenter at:

http://publib.boulder.ibm.com/infocenter/director/v6r2x/topic/com.ibm.director.vim.helps.doc/fsd0_vim_t_installing_uninstalling.html

Note: The VMControl plug-in has to be installed on the IBM Systems Director Server.

NIM Server preparation

Once we have the required software installed, we can proceed to discover the NIM Server in IBM Systems Director. As already mentioned in the introduction to this chapter, this scenario is more command line oriented and we use the IBM Systems Director command-line utility, **smcli**, to do all the necessary operations in IBM Systems Director. The **smcli** command should be executed by the root user, or another user assigned to the smadmin group.

NIM discovery in IBM Systems Director

The command to discover an AIX operating system is shown as follows:

```
smcli discover -i <NIM_ip_address> -t OperatingSystem
```

Assuming the system is successfully discovered, we can list the nodes. For that, we use the **smcli lssys** command together with a filter that displays only the nodes for which we have not configured proper authentication. In this scenario, this is the NIM Server.

```
smcli lssys -l -w "AccessState=Locked"
```

A sample stanza for a correctly discovered operating system is shown in Example 3-17.

Example 3-17 lssys output showing discovered operating system

```
nim:
  DisplayName: nim
  Description: null
  SystemBoardUUID
  CurrentTimeZone: -1
  IPv4Address: { '172.16.20.133' }
  HostName: { 'nim' }
  AccessState: Locked
  CommunicationState: 2
  OperatingState
  DisplayPingTime: 3
  DisplayOperationalStatusTime: 3
  SuspendEventActions: false
...
  Protocols: { 'CAS', 'CIM', 'SSH' }
```

URL
ManagementSoftware: { 'Unknown-IBM Director Agent-v6.2.0',
'IBM-IBM Director Core Services-v6.2.0.1' }

Notes:

- ▶ If the Protocols section does not include CAS, you might not have the Director Common Agent installed on your system. You should install it manually or use the agent push feature once you configure the access in IBM Systems Director. For information about this topic, see the Information Center at:
http://publib.boulder.ibm.com/infocenter/director/v6r2x/topic/com.ibm.director.install.helps.doc/fqm0_t_installing_common_agent.html
- ▶ The base installation of AIX starting with 6.1 TL03 and VIOS starting with version 2.1 includes the Common Agent by default. This applies to base installations, not updates.

Accessing the AIX operating system from IBM Systems Director

After the discovery, the system is in a Locked state. That means that IBM Systems Director does not have any control over the NIM Server component. You need to provide the credentials to access the system.

To request access through the command line interface, you have two options:

```
smcli accesssys <system_name>  
smcli accesssys -u root -p <root_password> <system_name>
```

The command prompts you to enter the user ID and password for every system.

Note: This second method of accessing a system presents a security risk, since the credentials might be recorded in the shell or other operating-system areas.

Tip: The second method can be useful only when you want to configure access to one or more systems at a time from a script.

In our environment, we used the following commands to discover and request access to the NIM server. See Example 3-18.

Example 3-18 Discovering and accessing a new NIM server

```
[isd:/]# smcli discover -i 172.16.20.133 -t OperatingSystem  
Discovery completion percentage 69%  
Discovery completion percentage 100%  
Discovery completed:  
100%  
[isd:/]# smcli accesssys nim  
Type the user ID:root  
Type the password:  
DNZCLI0727I : Waiting for request access to complete on... nim  
Result Value: DNZCLI0734I : Request access was successful.: nim  
DNZCLI0728I : Request access process for system completed.
```

VMControl NIM subagent installation

In order for VMControl to automate the use of NIM Server for importing, capturing, deploying, and deleting virtual appliances, it requires the Common Agent Subagent on the NIM Server.

After you set up and discover your NIM master, you generally do not need to log on to the NIM master to create NIM resources or to run NIM commands. VMControl creates the NIM resources and calls the NIM and AIX commands required to accomplish its import, capture, deploy, and delete tasks.

Note: IBM Systems Director VMControl update manager also relies on a NIM master for updating AIX systems. VMControl and update manager can use the same NIM master, or separate ones.

You can install a VMControl NIM subagent by using the installation wizard or you can manually install the subagent. To use the wizard, go to the “Navigate Resources” section of IBM Systems Director, and select your NIM Server, then from the menu choose **Release Management** → **Install Agent**.

If you need more details for manual or assisted installation, refer to:

http://publib.boulder.ibm.com/infocenter/director/v6r2x/topic/com.ibm.director.vim.helps.doc/fsd0_vim_t_installing_uninstalling.html

Note: The prerequisite for CommonAgentSubagent_VMControl_NIM is dsm.core. Remember to install it on your NIM server before deploying the subagent.

One way to check the installation status of the VMControl NIM Common Agent Subagent is to list the plug-ins installed in the Common Agent with the following command:

```
/opt/ibm/director/agent/bin/lwiplugin.sh -status | grep com.ibm.director.im \  
| cut -f5 -d:
```

This command should output four bundles, as shown in Example 3-19.

Example 3-19 Bundles indicating proper installation of the NIM subagent

```
com.ibm.director.im.rf.nim.imaster.n11  
com.ibm.director.im.rf.nim.imaster  
com.ibm.director.im.rf.nim.master  
com.ibm.director.im.rf.events
```

Note: For more details about VMControl and NIM, refer to the InfoCenter at:

http://publib.boulder.ibm.com/infocenter/director/v6r2x/topic/com.ibm.director.vim.helps.doc/fsd0_vim_r_power_virtualization_using_NIM.html

Migrate NIM resources from the old NIM Server to the new NIM Server

IBM Systems Director together with VMControl Standard Edition has the capability to back up and restore the AIX operating system in a slightly different manner than just using the traditional NIM method.

As mentioned above, IBM Systems Director can capture a running AIX operating system or an existing mksysb and transform it in a virtual appliance with the help of VMControl and a NIM server. This virtual appliance can be used afterwards either to deploy a new system from the image or to restore the image to the original system from which it was captured.

Existing mksysb and lpp_source resources on the old NIM Server need to be available on the new NIM Server in case you have to restore a system or install a fileset from an existing lpp_source.

For more details regarding Virtual Appliances, see 4.5, “OS deployment” on page 105.

We now show the necessary steps to transform your mksysb resources (located on the old NIM Server) into a virtual appliance on IBM Systems Director's image repository (new NIM Server):

1. Export the necessary mksysb resource information into a file, as shown in Example 3-20.

Example 3-20 Command used to export mksysb information

```
[root@csm:/]# for i in `lsnim -t mksysb|awk '{print $1}'`;
do
lsnim -l $i|egrep "location|source_image"|awk '{print $3}'|paste -d, - - >> /nimrepo/mksb.list;
done
```

2. NFS export the location where you store the mksysb images and mount it on the IBM Systems Director Server, as shown in Example 3-21.

Example 3-21 Commands used for exporting and mounting NIM resources

```
[root@csm:/]# exportfs /nimrepo
[root@isd:/]# mkdir /nimrepo; mount csm:/nimrepo /nimrepo
```

Note: Remember to replicate the same path on the target system. In Example 3-21, we export and mount the filesystem as /nimrepo.

3. Display the object ID of your NIM image repository in the IBM Systems Director, as shown in Example 3-22.

Example 3-22 Listing repositories with their object IDs

```
[root@isd:/]# smcli lsrepos -o
nim, 5032
```

4. Using VMControl, move (capture) the mksysb file onto the new NIM server. See Example 3-23.

Example 3-23 Loop used to migrate NIM data

```
[root@isd:/]# for i in `cat /nimrepo/mksb.list`;
do
importname=`echo $i|cut -d, -f2`
importpath=`echo $i|cut -d, -f1`
importfile=`echo $importpath|awk -F '/' '{print $NF}'`
importcpu=`smcli lsvrtsys -n $importname -A "Assigned Processing Units"|awk '{print $2}'`
importmemory=`smcli lsvrtsys -n $importname -A "Assigned Memory Size (MB)"|awk '{print $2}'`
smcli captureva -r 5032 -n ${importname}_${importfile} -F ${importpath} \
-A "cpushare=${importcpu},memsize=${importmemory}"
done
```

Note1: Remember to correctly set the fsize parameter in /etc/security/limits to avoid problems with transferred files being too large. After setting the fsize limit, you need to restart the Common Agent on the new NIM Server.

Note2: Commands and scripts included worked for our lab environment. Remember to test and modify the commands before you proceed with migration of your production environment.

Discover virtual appliances

The last step in the new NIM Server setup is to collect inventory from IBM System Director on this NIM Server in order to discover the virtual appliances imported earlier:

```
smcli collectinv -n <nim_server_name> -p "All Inventory"
```

After successfully collecting the inventory, the Image repository should show up in the VMControl Virtual Appliances tab (**System Configuration** → **VMControl** → **Virtual Appliances**). The name of your image repository will be the name of your new NIM Server.

3.3.4 Effective migration

This section provides migration guidance.

Discovery, Request Access and collect inventory in HMC

The next logic step in our scenario is to discover the Hardware Management Console on the IBM Systems Director Server. When you discover HMC devices with IBM Systems Director, it is a good idea for security and audit reasons to create a special user on the HMC, which IBM Systems Director uses for connectivity.

For example, if you have an HMC that manages many POWER servers, and you do not want to discover all of them into IBM Systems Director, then you can add a user and assign roles to it. These roles will allow your user to manage only the systems you granted access to.

For help on how to create roles for managed resources, you can use section 5.2.3 in *Hardware Management Console V7 Handbook*, SG24-7491.

To create users on the HMC through which Systems Director will manage the resources, again see section 5.2.3 in *Hardware Management Console V7 Handbook*, SG24-7491.

Discovery

The discovery process is not disruptive. It just interrogates the system using various methods and protocols and as a result it adds the HMC in IBM Systems Director's database.

To discover the HMC into Systems Director, use the following command:

```
smcli discover -i <HMC_ip_address> -t OperatingSystem
```

Request access

After the discovery completes successfully, Systems Director cannot access this HMC because the credentials to authenticate have not been provided. The next step is to request access using the user you created in the previous section.

To request access through the command line interface, use the following command:

```
smcli accesssys -u hscdir -p <hscdir_password> <hmc_system_name>
```

Note1: This method of accessing a system presents a security risk, because the credentials might be recorded in the shell or other operating system areas.

Note2: After discovering and providing access to the HMC, you will see only the logical partitions. Operating systems must be discovered and accessed separately.

Collect Inventory

The Collect Inventory operation is done to gather information such as LPARs, Virtual Switches, Virtual Networks and other virtual resources. Use the following command to collect the inventory from Hardware Management Control:

```
smcli collectinv -n <hmc_server_name> -p "All Inventory"
```

After completion of these steps, you have an IBM Systems Director environment ready for migrating the nodes.

Exporting the CSM nodes and discovering them in Systems Director

Because you may have many systems defined into CSM that you want to discover into IBM Systems Director, we have provided a few commands to automate the export and the discovery process.

We could have used the automated discovery of IBM Systems Director based on ranges of IP addresses but this might not fit everyone's needs. We want to discover only the nodes we already had defined in CSM.

After we export necessary node and hardware device data from CSM, we are able to start the first nondisruptive step of operating systems migration into IBM Systems Director.

Exporting the nodes and hardware devices from CSM

To export the hardware devices and nodes, we used the commands shown in Example 3-24.

Example 3-24 Commands used to export data from CSM

```
[csm:/]# lshwdev -i | awk '{print $2}' > /tmp/csmhwdevs  
[csm:/]# lsnodes -i -w "InstallOSName like 'AIX'" | awk '{print $2}' > /tmp/csmnodes
```

These commands created two files, which are then transferred to the IBM Systems Director Server. These files contain only IP addresses of the nodes and hardware devices.

Discovering the nodes and hardware devices into IBM Systems Director

In this example, we transferred the files to the */tmp/* directory. We used the commands shown in Example 3-25 to import the hardware devices and nodes.

Example 3-25 Import CSM hardware devices and nodes in IBM Systems Director

```
[isd:/]# for i in `cat /tmp/csmhwdevs`; do  
> smcli discover -i $i;  
> done  
  
[isd:/]# for i in `cat /tmp/csmnodes`; do  
> smcli discover -i $i -t OperatingSystem;  
> done
```

Note: Discovery does not make any changes on the target system.

Import of CSM nodes to IBM Systems Director endpoints

We recommend to migrate the nodes one at a time to avoid confusion and introduce problems. It is advised to check that every endpoint is in the correct state before beginning with migration of the next node.

By this time, you should have a list of operating system endpoints without access. In the GUI, you can see this list in the Resource Navigator.

In the smcli, we use the `lssys` command to display the endpoints without access. Refer to Example 3-26.

Example 3-26 Partial output of the smcli lssys command

```
[isd:/]# smcli lssys -l -w "AccessState=Locked"
p6client1:
  DisplayName: p6client1
  Description: null
  IPv4Address: { '172.16.20.72' }
  HostName: { 'p6client1' }
  AccessState: Locked
  Protocols: { 'CAS', 'CIM', 'SSH' }
  URL
  ManagementSoftware: { 'Unknown-IBM Director Agent-v6.2.0',
                        'IBM-IBM Director Core Services-v6.2.0.1' }
```

For each endpoint, we recommend the following steps:

1. Remove the node from CSM by running the following command on the CSM Management Server:

```
rmnode <node_name>
```

2. Grant access to the operating system in IBM Systems Director by running the following command on the IBM Systems Director Server:

```
smcli accesssys -u root -p <root_password> <system_name>
```

Note: This method of accessing a system presents a security risk, because the credentials might be recorded in the shell or other operating system areas. An interactive alternative is to leave out the password and be prompted.

3. Collect inventory

Collection of the inventory is needed to gather additional information about the endpoint. It is strongly recommended after accessing the endpoint for the first time. It is also a good idea to perform it once a major change in the configuration happens. Collect inventory using the following command:

```
smcli collectinv -n <node_name> -p "All Inventory"
```

4. Verify that the endpoint is correctly migrated.

A properly migrated endpoint is shown in Example 3-27. The output has been truncated, leaving out irrelevant information.

Example 3-27 Partial output from smcli lssys

```
[isd:/]# smcli lssys -l -n p6client1
p6client1:
  DisplayName: p6client1
  SystemBoardUUID: null
  IPv4Address: { '172.16.20.72' }
  HostName: { 'p6client1' }
  AccessState: Unlocked
  MachineType: 9117
```

```
SerialNumber: 101F170
MACAddress: { '6A8886E46211' }
VMID: 3
Model: MMA
Architecture: ppc64
Virtual: true
p6client1:
  DisplayName: p6client1
  SystemBoardUUID
  CurrentTimeZone: -240
  IPv4Address: { '172.16.20.72' }
  HostName: { 'p6client1' }
  AccessState: Unlocked
  MACAddress: { '6A8886E46211' }
  OSVersion: 6.1
  OSTypeString: AIX
  MachineType: 9117
  Model: MMA
  SerialNumber: 101F170
  VMID: 3
  Architecture: ppc64
  BuildNumber: 6100-06-05-1115
  Protocols: { 'CAS', 'CIM', 'SSH' }
  ManagementSoftware: { 'IBM-IBM Director Core Services-v6.2',
                        'IBM-IBM Director Agent-v6.2.0' }
```

After each system is migrated to the new IBM Systems Director Server, you may want to configure other functions required for that system such as monitoring, automated events, and commands.

3.4 Decommissioning CSM

After a successful migration of all your systems to the new IBM Systems Director Server, or after you have solved all the problems encountered during the transformation, it is time to decommission the CSM Cluster.

Before decommissioning we advise you to make a fresh backup of the CSM configuration data using the **csmbackup** command. This command provides the latest copy of your environment. 5.4, “Backup and restore of the management server” on page 139 provides more details about backup and restore.

Note: Refer to *IBM Cluster Systems Management for AIX 5L™ Command and Technical Reference Version 1.7.1*, SA22-7934 for information about the **csmbackup** command.

We also recommend to save a copy of this backup on external media. This provides the ability to restore your CSM data in event that is necessary.

Note: At this point we assume you have migrated all the required CSM functions available in the IBM Systems Director environment.

To uninstall the CSM Management Server, execute the following steps:

- Stop the `rsct_rm` agent before the unistallation:

```
stopsrc -s rsct_rm
```

- Remove the CSM Server filesets using:

```
uninstallms
```

For more details, including optional flags for `uninstallms`, see the CSM Command Reference, SA22-7934.

Note: Particularly in the coexistence scenario, the CSM uninstallation removes the `csm.dsh` fileset. For a proper functioning of **dsh** features in IBM Systems Director, this fileset has to be installed on the IBM Systems Director Server. This fileset is available starting with AIX 6.1 TL 06 or later.

For the coexistence scenario where both products were running on the same system, after the CSM removal we tested some IBM Systems Director features and determined that this removal did not affect the IBM Systems Director capabilities.



Functional comparison

This chapter provides the Cluster Systems Management administrator with details on comparing major features and functions between their existing environment and one using IBM Systems Director.

We discuss the following high-level topics:

- ▶ The monitoring infrastructure and high-level design between CSM and IBM Systems Director. Examples of how to implement some basic monitoring tasks in both environments are provided.
- ▶ One of the key capabilities in Cluster Systems Management is the management of the state of the logical partitions and physical servers and the ability to execute remote commands on the target nodes. In this chapter we discuss the features in IBM Systems Director that provide these capabilities, as well as managing system firmware.
- ▶ Security is a topic on every IT practitioner's mind in today's world. The security and authentication models in CSM and IBM Systems Director are compared in addition to some information on managing a firewalled environment.
- ▶ The integration of Cluster Systems Management and AIX's Network Installation Manager was a key component in operating system deployments under CSM. IBM Systems Director with VMControl Standard Edition provides similar functions with the differences and some details discussed.
- ▶ Lastly a comparison of the various Agent components in IBM Systems Director versus Cluster Systems Management on AIX.

4.1 Monitoring of resources

Monitoring provides the means to retrieve and visually observe real time changes in system resources and should consist of at least the basic elements such as events, thresholds and system health. Over the years monitoring seems to have shifted lanes from strict server monitoring to application monitoring, because the signal to noise ration is often very poor on server monitoring. Having defined end-to-end or inside-out monitoring for your application will grant a much more granular information flow, as well as list the current state of the application.

More importantly, it can help you track changes made to the system and provide you with a baseline for future comparison.

Both Cluster Systems Management (CSM) and IBM Systems Director (ISD) provide a wide range of monitoring capabilities, including support for customizations and scripting to meet the exact needs.

4.2 Resource Monitoring and Control (RMC) versus IBM Systems Director Agent Services

In this section, we briefly compare the differences between a CSM-managed RMC cluster versus the IBM Systems Director Agent Service approach. Here we strictly focus on monitoring and no other components are discussed.

► CSM

The cluster relies on the historic Reliable Scalable Cluster Technology (RSCT) framework which traces its roots back to the RS/6000 SP and is considered a well respected product that provide a robust cluster management capability. It has been a part of the AIX code base since AIX 5L, and provides support for Linux, Solaris and Windows. CSM is only available for AIX and Linux clusters.

► IBM Systems Director

This is a fairly new way of providing platform management and considered the corner stone of the IBM Smarter System portfolio. It is based on the Common Information Model (CIM) and supports multiple operation systems as well as virtualization technologies across IBM and non-IBM platforms. It also provides an easy-to-use web management interface, and utilizes Common Agent Services (CAS), which is known from the Tivoli Provisioning Manager (TPM) suite and offers a shared infrastructure for managing systems.

IBM Systems Director relies heavily on a web-based console where Cluster Systems Management (CSM) offers a more command line-based approach. WebSM, the web-based tool for system managementdoes, however, provide plug-in options for CSM. but it comes with some limitations when compared to Systems Director. We do not cover the usage of webSM in this book.

4.2.1 CSM using Resource Monitoring and Control (RMC)

The resource monitoring and control (RMC) subsystem is what is known as a generalized framework for managing, monitoring and manipulating resources, those can be both physical and logical. Aresource is the fundamental concept of RMC architecture. It refers to an instance of a physical or logical entity that provides services to some other component of the

system. However, the term resource is used very broadly to refer to software as well as hardware entities. We do not go into the details and structure of RSCT RMC because there are already many books providing this type of information.

For more detailed information about RSCT and CSM, visit the “Libraries Related to Cluster Products” section at the IBM Cluster Products information center website or the following books:

SA23-1343-05 - *CSM for AIX and Linux V1.7.1: Administration Guide*

SA23-1344-05 - *CSM for AIX and Linux V1.7.1: Planning and Installation Guide*

SA22-7889-20 - *RSCT: Administration Guide*

SA22-7890-20 - *RSCT for AIX: Technical Reference*

RMC runs as a daemon process and can be configured to manage and monitor resources of a single machine or what is also known as a management domain, which is what CSM is using, or by using a cluster peer domain which is utilized by applications such as IBM PowerHA, IBM DB2 PureScale, and IBM DB2 HADR, and others.

CSM uses the management server as a single point of control in the management domain by using a subset of RSCT Resource Monitoring and Control (RMC). It lets you configure scripts and response actions to manage general system conditions with little or no involvement from the system administrators. It is important to understand that even though RMC provides the basic abstractions for representing physical or logical entities, it does not represent any actual entities itself—a resource manager simply connects the actual entity to the relevant RMC abstractions. Refer to Figure 4-1.

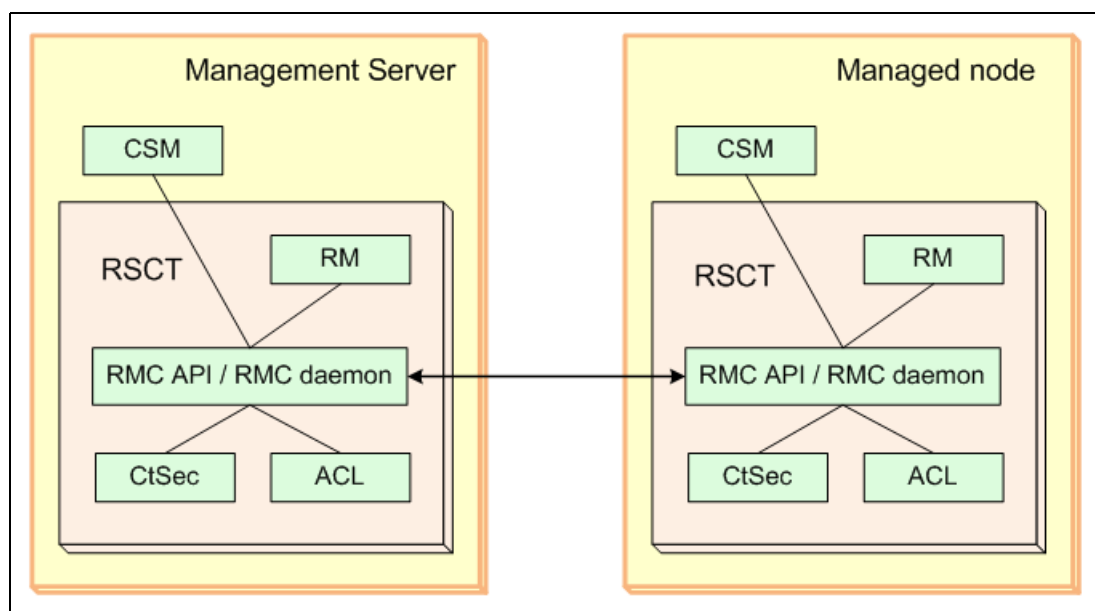


Figure 4-1 CSM RSCT using the RMC communication flow

CSM-related resource managers

RMC provides a modular architecture typical of resource managers and resource classes, each providing a comprehensive cluster management stack. It does not provide a plug-in infrastructure, nor does it provide any browser system interface. In order to simplify the differences, we will briefly describe them.

- ▶ Hardware control resource manager
The hardware control resource manager is a component of CSM that runs on the server designated as the management server (MS).
- ▶ Domain management server resource manager
The domain management server resource manager runs on the server designated as the management server, and controls a number of resources.
- ▶ Management domain configuration
The resource manager implements the necessary function to automatically add the operating system image, as a managed node, to the management domains created automatically on each HMC.

How it all works from a system monitoring perspective

The resource manager maps system resources into resource-class abstractions, which could be physical volume or file system. Each resource class includes individual system resources such as /tmp, /var and so on. Each resource can then be monitored as *Monitored Resource*. Based on the monitoring scheme we can define various conditions, such as PercentToUsed. The state of the conditions is defined by threshold values, expressions such as PercentToUsed >90 or <90 as well as the severity: Informational, Warning or Critical. All in all CSM provides comprehensive monitoring using RMC with predefined monitors for key system metrics, and support for customizing of complex system parameters, functions, and monitors.

4.2.2 IBM Systems Director Common Agent and related resource managers

IBM Systems Director uses what is known as common agent services (CAS) to communicate with its endpoints. It is built on the Common Information Model (CIM). The Common Agent is a superset that also includes the Platform Agent and is part of the IBM Tivoli infrastructure suite. It includes sub agents that provide specific management capabilities such as resource managers, agent managers, and common agent elements. The purpose of the Common Agent architecture is to reduce the infrastructure complexity by providing a simplified method that could be utilized by multiple products while being transparent to the operating system (OS) instead of separate agents that each provide the same functionality. Refer to Figure 4-2 on page 77.

IBM Systems Director is defined as a resource manager and uses its agent manager service to communicate with the manageable endpoint system agents.

Visit the IBM Systems Director Information Center website for more extensive information and education, or refer to the following IBM Systems Director books:

GI11-8709-06 - *Planning, Installation and Configuration Guide for AIX*

GC30-4176-06 - *IBM Systems Director - System Management Guide*

GC30-4170-06 - *IBM Systems Director - Command Reference*

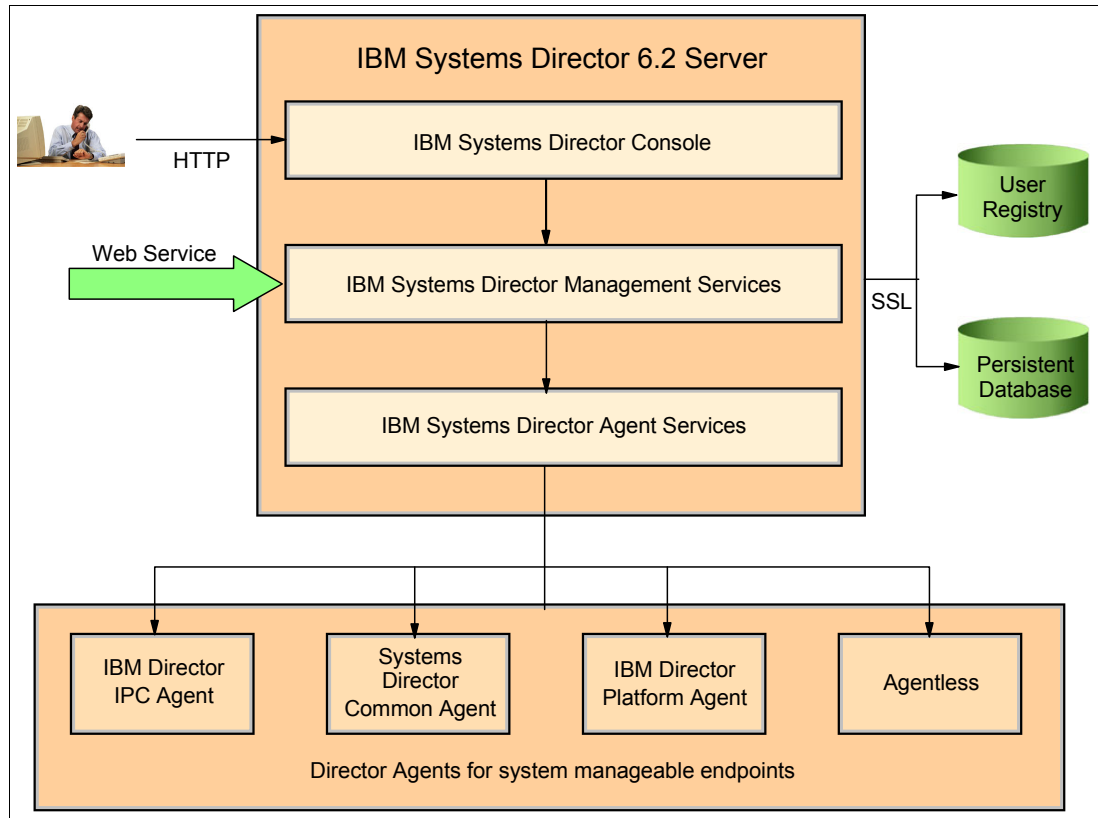


Figure 4-2 IBM Systems Director agent-based communication

The IBM Systems Director uses the common agent services (CAS) architecture, and it provides a shared infrastructure for managing systems. Resource managers (such as IBM Systems Director Server) use an agent manager to communicate with the common agents that are installed on the managed resources.

- **Resource Manager**

The Resource Manager is installed on the management server (MS), which provides the management application used by the agent manager for security and credential management. The IBM Systems Director Server is the resource manager for IBM Systems Director.

- **Agent Manager**

The Agent Manager is the server component of the Common Agent that provides the functionality that allows clients to get information about agents and resource managers. It also provides the registry and agent recovery services, enables secure connections between managed endpoints, and maintains the database information.

- **Common Agent**

The Common Agent is the common container for all the subagents that run within. It provides a rich set of security, deployment, and management functions. The functionality provided varies based on the operating system and hardware in the environment. It is installed on managed systems and reports back to the resource manager, as well as providing resource manager tasks. It is available for all IBM Power Systems, System X, IBM BladeCenter, System z, and some non-IBM systems. AIX systems utilize the Common Agent.

- Platform Agent

The Platform Agent provides a smaller subset compared to the Common Agent. However, it does not sacrifice the many abilities to manage and monitor your system. It is available for all IBM Power Systems, System X, IBM BladeCenter, System z, and some non-IBM systems.

Common Agents and how to monitor resources

The Common Agent gets the information by communication with the CIM Agent. The monitoring definitions are using what is called CIMObjectPath classes or CIM Instance, which consist of different or multiple layers.

- Agent Provider - name of the agent also known as higher-level or top-level monitor
- NameSpace - the namespace also known as-Monitor path or lower-level submonitor
- Metric - name of the object class also known as Monitor path or lower-level submonitor
- KeyBindings - key/value pairs which uniquely identify an instance; for monitor purposes these are called Monitor data.

For example, resource monitors are made up of levels:

[Director Agent][CPUMonitors][Process Count]

You can specify any top-level resource monitors such as [Director Agent] or lower-level submonitors [CPU Monitors][Process Count]. If you specify a top-level monitor, all lower-level submonitors are also targeted.

Development and manipulation monitors were fairly straightforward in RMC. Years of usage had contributed to extensive information and high-quality documentation. Migrating to IBM Systems Director requires a somewhat different mind set, but do not panic—the IBM Systems Director web navigation GUI provides some of the much needed guidance.

So lets get started. You should already be somewhat familiar with the IBM Systems Director content navigation panel, and how to use or navigate within it, when dealing with monitors. Select **System Status and Health** → **Monitors**, and select one of the predefined monitor groups such as the AIX Monitors, then select **CPU Utilization**.

The Display Path, for the selected monitor, is shown at the bottom of Figure 4-3 on page 79.

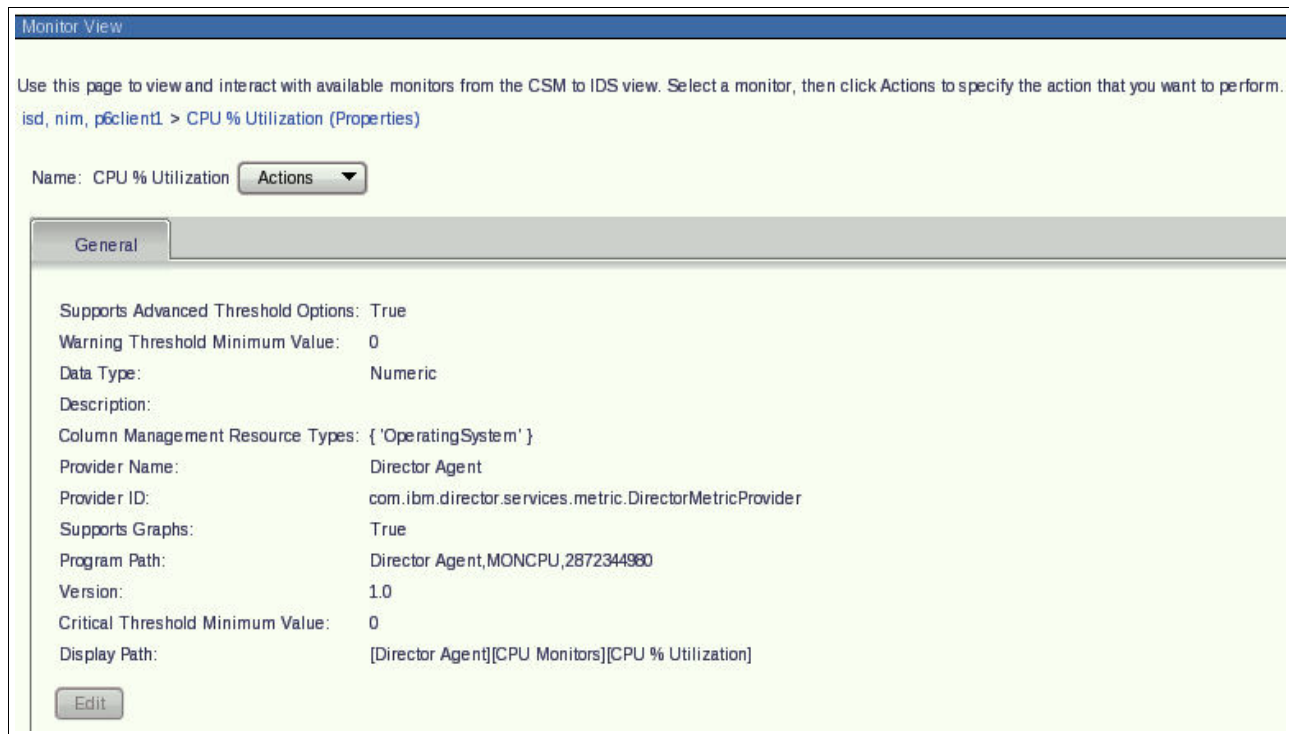


Figure 4-3 Using the Monitor View to list the monitored Object Classes

A similar option is available by using the systems management command line interface (**smcli**), as shown in Example 4-1.

Example 4-1 Object class for AIX monitoring of /tmp listed through smcli

```
[CIM Agent] [root/CIMv2] [AIX_MetricVal] [InstanceID="fs:AvailableSpace:/dev/hd3"] /
[Available Space of Filesystem /tmp
(Megabytes)]
```

4.2.3 Comparison between various system resource monitors' output

Since the architectures of CSM and IBM Systems Director are completely different, the option of reusing existing definitions is considered very limited. We did not find any easy way of migrating the current CSM monitors to IBM Systems Director.

Where CSM lacked a simplified manageable GUI, Systems Director provides scoreboards, dashboards, health summarizes and more, easy to manage through a web console. However, a large portion of the monitoring capabilities is also manageable using **smcli**. In Table 4-1 we have mapped the frequently used monitor commands from CSM to the IBM Systems Director counterpart.

Table 4-1 Mapping between CSM and SDS resource monitoring commands

CSM Commands	IBM SDS commands	Description
mkconditions	runresmon, mkresmonthresh	Create and define a resource to be monitored together with any threshold values.

CSM Commands	IBM SDS commands	Description
lsconditions, lscondres, lsresponse	lsresmon, lsresmonthresh	List defined monitors and their current threshold values.
chcondition	chresmonthresh	Change the settings of a resource being monitored.
rmcondition	rmresmonthresh	Remove a resource monitor and its thresholds.
mkresponse, mkcondresp	mkresmonthresh	Create a link between the resource monitored and the condition monitored.
startcondresp	chresmonthresh -A	Start monitoring of a current resource and/or its threshold.
stopcondresp	chresmonthresh -D	Stop the monitoring of a current resource and/or its thresholds.

In our CSM-managed cluster we have enabled one of the predefined conditions that create an event if the total space available on our /tmp file system becomes critical; the output of the command is shown in Example 4-2.

Example 4-2 Listing a simple file system condition on a CSM managed node

```
lscondition "/tmp space used"
```

Displaying condition information:

condition 1:

```

Name                = "/tmp space used"
Node                = "p55701p01"
MonitorStatus       = "Monitored"
ResourceClass       = "IBM.FileSystem"
EventExpression      = "PercentTotUsed > 90"
EventDescription     = "An event will be generated when more than i
                        90 percent of the total space in the /tmp
                        directory is in use."
RearmExpression      = "PercentTotUsed < 75"
RearmDescription     = "The event will be rearmed when the percent
                        of the space used in the /tmp directory
                        falls below 75 percent."
SelectionString      = "Name == \"/tmp\"
Severity            = "i"
NodeNames           = {}
MgtScope            = "1"
Toggle              = "Yes"
EventBatchingInterval = 0
EventBatchingMaxEvents = 0
BatchedEventRetentionPeriod = 0
BatchedEventMaxTotalSize = 0
RecordAuditLog      = "ALL"
```

Creating a similar condition in IBM Systems Director is straightforward. In Figure 4-4 on page 81, we have created a group view called CSM to IDS, which consists of a small

selection of already discovered systems. In the monitor selection view we selected some of the predefined AIX Monitors that come with IBM Systems Director.

Monitor View

Use this page to view and interact with available monitors from the CSM to IDS view. Select a monitor, then click Actions to specify the action that you want to perform.

isd, nim, p6client1

Select	Name	Monitor Name	Current	Warning	Critical	Threshold Status
<input type="checkbox"/>	isd	CPU % Utilization	20.4%	>= 90.0	>= 100.0	✖ Critical
<input type="checkbox"/>	isd	Memory Usage	4072 Megabytes			
<input type="checkbox"/>	isd	Paging Space Remaining (MB)	1587			
<input type="checkbox"/>	isd	Total Space of Filesystem /tmp (Megabytes)	1696 Megabytes			
<input type="checkbox"/>	isd	Used Space of Filesystem / (Megabytes)	210 Megabytes			
<input type="checkbox"/>	nim	CPU % Utilization	6.97%	>= 75.0	>= 90.0	✔ Activated
<input type="checkbox"/>	nim	Memory Usage	6001 Megabytes			
<input type="checkbox"/>	nim	Paging Space Remaining (MB)	500			
<input type="checkbox"/>	nim	Total Space of Filesystem /tmp (Megabytes)	No Data Available			
<input type="checkbox"/>	nim	Used Space of Filesystem / (Megabytes)	No Data Available			
<input type="checkbox"/>	p6client1	CPU % Utilization	78%			
<input type="checkbox"/>	p6client1	Memory Usage	3346 Megabytes			
<input type="checkbox"/>	p6client1	Paging Space Remaining (MB)	499			
<input type="checkbox"/>	p6client1	Total Space of Filesystem /tmp (Megabytes)	160 Megabytes	<= 100.0	<= 25.0	✔ Activated
<input type="checkbox"/>	p6client1	Used Space of Filesystem / (Megabytes)	204 Megabytes			

Page 1 of 1 | Selected: 0 Total: 15 Filtered: 15

Figure 4-4 Using the Systems Director web interface to define threshold values

By selecting the Action Menu we can define threshold values, and manage processes and system properties. In our example we configured CPU Utilization on our NIM server and total space of file system /tmp on p6client1.

IBM Systems Director also supports the abilities to list monitoring settings and conditions by using **smcli**. Hence the format is different from the information provided by RSCT; refer to Example 4-3.

Example 4-3 ~~Using the smcli lsresmon command to list active monitored thresholds on managed~~ *Using the smcli lsresmonthres command to list the active monitored thresholds* **rewrote in blue text**

```
$ smcli lsresmonthres -l -m "[Director Agent][CPU Monitors][CPU Utilization]" -n nim
```

```
Name: Individual threshold for CPU % Utilization on system nim
IsActive: true
Type: Individual
Target: nim
Attribute:[Director Agent][CPU Monitors][CPU Utilization]
Description:
Double Data      HighError :90.0 HighWarning :75.0
```

Tip: Use the **lssys** command without any options to list all system names, or the **lssys -o** command to list all system IDs.

IBM Systems Director provides another monitor option called process monitoring. In Figure 4-5 on page 82, we have selected the sshd daemon for monitoring, where the default process event options issue an event on the following condition criteria, where one or more apply:

- ▶ Start
- ▶ Terminate
- ▶ Fail to start after reboot

Viewing processes for target: p6client1.

Select a target:

p6client1

Select a view:

Applications

Manage Processes (p6client1)

Select	Name	Command Line	Process ID	Parent Process ID	User	Priority	Memory Usage	CPU Time	Monitored
<input type="radio"/>	wait	wait	983070	0	root	41	448K	00:00:06	False
<input type="radio"/>	random	random	1704002	1	root	20	448K	00:00:04	False
<input type="radio"/>	ldmp_process	ldmp_process	1835106	1	root	20	512K	00:00:00	False
<input type="radio"/>	psmd	psmd	393228	0	root	41	512K	00:00:00	False
<input type="radio"/>	init	/etc/init	1	0	root	20	720K	00:00:00	False
<input type="radio"/>	tier1slp	/opt/ibm/director/cimom/bin/tier1slp	2818214	1	root	20	5220K	00:00:00	False
<input type="radio"/>	cimserver	cimserver	4522152	1	root	20	37216K	00:01:05	False
<input type="radio"/>	IBM_SensorRMd	/usr/sbin/rsc/bin/IBM_SensorRMd	6095052	3014818	root	20	1364K	00:00:00	False
<input type="radio"/>	dirsnmpd	/opt/ibm/icc/cimom/bin/dirsnmpd	4587666	1	root	20	7048K	00:00:00	False
<input checked="" type="radio"/>	sshd	/usr/sbin/sshd	9044076	3014818	root	20	1208K	00:00:00	True
<input type="radio"/>	getty	/usr/sbin/getty	8192040	1	root	20	620K	00:00:00	False
<input type="radio"/>	inetd	/usr/sbin/inetd	3408006	3014818	root	20	572K	00:00:00	False
<input type="radio"/>	java	/var/opt/tivoli/epf_jvm/jre/bin/java -Xmx384m -Xmin10.01 -Xma...	7078114	7143672	root	20	139136K	00:03:43	False
<input type="radio"/>	nfsSM	nfsSM	2621520	0	root	20	512K	00:00:00	False
<input type="radio"/>	aioPpool	aioPpool	2031712	1	root	20	448K	00:00:00	False

Page 2 of 6 2 Selected: 1 Total: 80 Filtered: 80

Figure 4-5 Using the Systems Director web interface to define process monitors

Depending on the condition selected the reporting may vary. Hence every time a condition is met the corresponding result is reported to the event monitor and usually to the active status monitor.

So what would happen if our SSH process was terminated. Clearly any users would lose access to the server and the new process condition would generate an event informing that the process is dead. One obvious way is to restart it by hand or script. Another is to automate it. Based on this new condition we can continue to configure event-based conditions because IBM Systems Director provides something called Event Automation Plans (EAP). Using this we can automate actions based on certain events. EAP supports various actions from post to a news group, and email a cell phone to start a task on a remote system. Figure 4-6 on page 83 shows a subset of the available action selections.

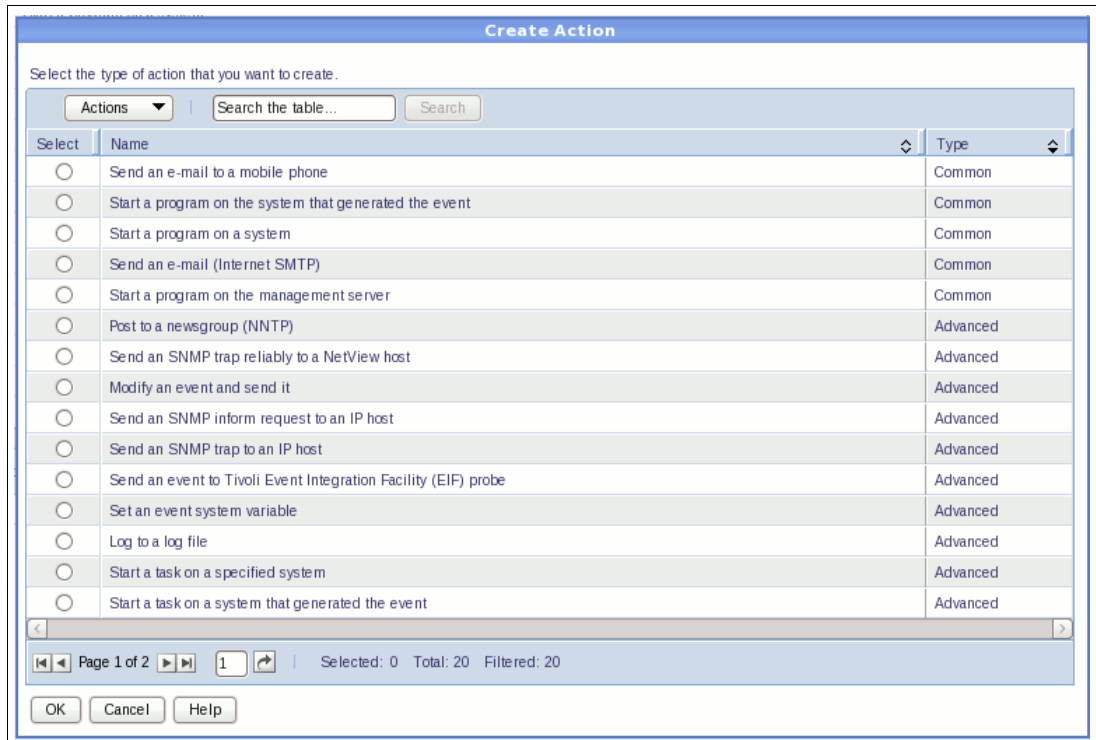


Figure 4-6 Action types available for Event Action Plan (EAP)

Going back to our dead process event, in order for the EAP to assist us we would need to create one first. This is done by selecting **Automation** in the content navigation panel, where we need to define:

- ▶ An Event Filter monitoring our SSH daemon
- ▶ An Event Action that starts our SSH daemon
- ▶ An Event Automation Plan that combines the defined Event Filter and Action into a plan.

In Figure 4-7 on page 84 we have done so. Once applied it will monitor the target defined.

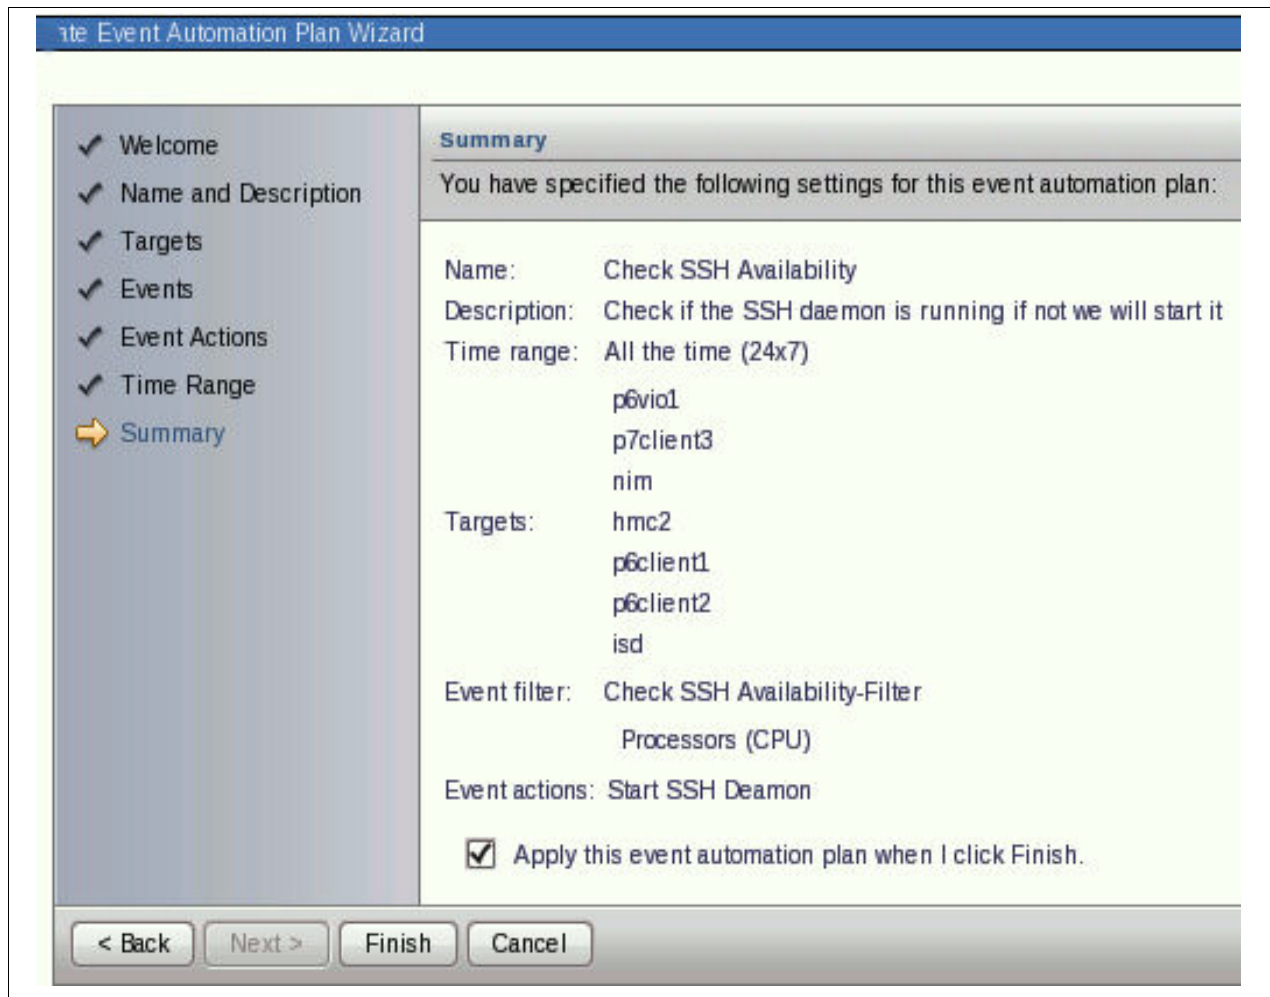


Figure 4-7 Creating an Event Automation plan

The same functionality is available using the **smcli** command to create an EAP. You would be using the **mkevtautopln** parameter. The syntax is rather complex, so unless you are really familiar with the command structure we suggest you stick with the GUI.

In Example 4-4 we use the **smcli lsevtautopln** command to list the EAP that we previously defined.

Example 4-4 Using lsevtautopln to list event-automation plans

```
#smcli lsevtautopln -l 0x12
Name: SSH Availability
Description: Check if the ssh daemon is active, if not start it
Status: Active
Event Filter: SSH Not Running
Time Ranges:
    All the time (24x7)
Actions:
    Start SSH Deamon
Targets:
    System Name: p6client1
```

You can also use the `mkevtautopln` command to create an EAP, or `chevtautopln` to make changes to an existing EAP.

4.2.4 Monitor migration considerations

Since CSM uses RSCT RMC and IBM Systems Director, it is agent protocol-based. We did not find any easy way to transform existing CSM monitor or condition data into an XML format usable by IBM Systems Director, because data is kept in two very different formats.

One could consider whether making such a transformation tool would really be worth it because this is often costly. Our conclusion is that it depends on the environment and would most likely require a cost benefit analysis (CBA), the result of which is a total picture of costs, benefits and risks involved, making it much easier for management to make the final call.

We found it much easier to document the existing CSM monitors by using commands such as `lscondition` and `lscondresp` and then using this information to define new monitors on IBM Systems Director, but because each environment is different, other solutions might be more appropriate.

4.3 Hardware and software management

This section contains information about hardware and software management in the cluster.

4.3.1 Remote commands

Remote management facilities in IBM Systems Director are more comprehensive and include more remote functionalities than CSM. In CSM we got used to performing remote management for more than one platform with the help of distributed shell commands (`dsh`) and distributed copy commands (`dcp`) for remote file transfer. In Systems Director we can still find these functionalities, although the implementation is a bit different and plus we can find many others such as: remote command line, hardware command line, web-based administration, remote control, and remote serial consoles.

Throughout this chapter we compare the implementation of the functionalities available in both products and also describe the new functionalities available only in IBM Systems Director.

Distributed Command Execution

The purpose of `dsh` is to run remote shell commands and their flags on multiple systems at a time.

Cluster Systems Management

The `dsh` used in CSM was based on PSSP's distributed shell commands `dsh` and `dshbak`. When you need to execute a remote command on multiple nodes, `dsh` runs that command for you in parallel on all the nodes you specified, and afterwards `dshbak` formats the output. The two commands were part of the `csm.dsh` fileset.

Because `dsh` could run commands on multiple nodes or groups at a time it was able to populate its targets from three different sources (contexts): CSM database, NIM resource database, or `dsh`-specific methods (node list). Another option was to dynamically construct a list of nodes by specifying them as parameters using the `-n` or `-N` flags.

There were three available methods to use **dsh** for remote shell commands:

- **Command Line Interface**

With this method you had to build the command every time you ran it and you had to know all the flags. The nodes against which you wanted to run the command you had to either enter manually or by setting the `DSH_NODE_LIST` or `WCOLL` variable.

- **Smitty csm_dsh**

This method enabled you to dynamically construct the commands. Smitty menus were able to search for nodes and groups defined in the CSM database and set them as targets for the command you were about to run.

- **Distributed Command Execution Manager**

DCEM was included as a plug-in application in the IBM web-based System Manager. This method was the most comprehensive because you had many options available. You could choose the nodes or devices to run this command against based on DHS, NIM, or CSM's database, and had the option to save and edit the commands you built for later use. DCEM command output and activity were saved in log files so you were able to search through previous runs of the commands. For more details see *IBM Cluster Systems Management for AIX and Linux Planning and Installation Guide Version 1, Release 7.1*, SA23-1344-05.

With CSM it was the system administrator's responsibility to configure and enable remote shell access to the systems, or configure Kerberos. So we can say that CSM was not in charge of security. The distributed shell command **dsh** used the underlying remote shell security protocols, which by default was **rsh**, but you had the option to change it to an even more secure remote command protocol, **ssh**.

Figure 4-8 on page 87 shows the panel of CSM - DCEM.

Command Options Reports

Enter the name of a new command specification or browse for a saved one.

Name:

Description:

Define command

Path:

Command:

Default user:

Specify targets

Include targets from these contexts: ☒ CSM ☒ DSH ☒ NIM

☐ Run on all CSM nodes. ☐ Run on all CSM devices.

Node names:

Node groups:

Device names:

Device groups:

Figure 4-8 Panel of CSM - DCEM

IBM Systems Director

In IBM Systems Director you still have the option to execute remote commands on one or more managed systems in parallel using the distributed shell. The only constraint is that IBM Systems Director has to be installed on an AIX system.

When you run a distributed command it is executed on each remote target and the output from each target is returned to the IBM Systems Director Server. You have the option to format the output in two ways: to compress identical output, and to group output from each target.

In order to run a command on multiple targets, you can use the IBM Systems Director's wizard to select the targets against which you want to run the commands. Another useful option is to save the commands you are running so that you can recall them later on the same target or on different ones. You can choose even the shell in which remote commands should be run and how long the commands should timeout.

To use the wizard go to **System Configuration → Remote Access → Distributed Command** as shown in Figure 4-9.

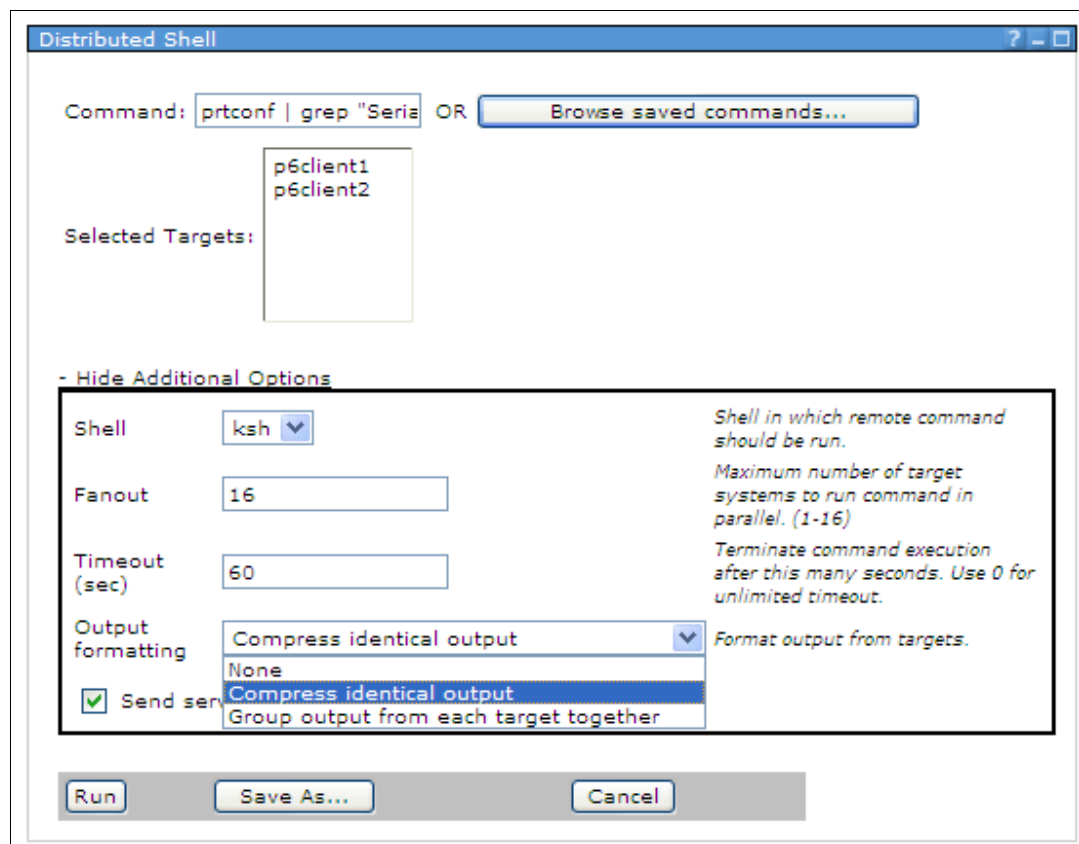


Figure 4-9 Distributed command wizard

After the command executes on the selected targets, the Command Output pane is displayed in the Distributed Shell page to show the output of the command from each target, as well as errors from those targets. For more details about using the Distributed Shell from IBM Systems Director, see:

http://publib.boulder.ibm.com/infocenter/director/v6r1x/topic/director.remote_6.1/fqm0_t_ra_using_distributed_shell.html

The package in charge of **dsh** commands in IBM Systems Director is **dsm.dsh**. If you install IBM Systems Director on a system which has already installed **csm.dsh** then **dsm.dsh** will not be installed. If you need to uninstall **csm.dsh**, as is the case in our coexistence scenario (refer to 3.2, “Coexistence” on page 42), then you have to manually install **dsm.dsh**.

Another option to run remote commands concurrently on remote targets is to use the CLI command **smcli dsh** with different options. The output returned by this command is formatted so that you can easily manage the command results from all targets. For more details, see the manual by typing at your command line interface: **smcli dsh --help**.

Unlike CSM, IBM Systems Director is in charge of the security of the remotely executed commands. For more details in regard to security, refer to 4.4.3, “Security consideration when defining users and groups” on page 101.

Remote management of the configuration files and parameters

The remote files and parameter configuration tools used by CSM and IBM Systems Director cannot perfectly match. They are grouped in the same section but they do not achieve 100% identical functionality.

Cluster Systems Management

In CSM you could maintain the configuration of your nodes in the cluster with the help of the Configuration File Manager. CFM was acting like a file repository for the configuration files that were common among the nodes in a cluster. The file repository was located on the management server and contained all the files that need to be shared among the cluster's nodes.

Thanks to this simple design the system administrator's main responsibility was to maintain the configuration files stored on the management server, and all the changes to these files were automatically propagated throughout the cluster.

Though the files were common, you still had a certain amount of control over your configuration files using different methods, which allowed for variations based on groups, IP address, and host name. You had the flexibility to specify one version of the configuration file for all the nodes of your cluster, while specifying different versions for CSM-defined node groups.

The main components of CFM for remote synchronization of your configuration files were:

- ▶ **Server File Repository:** a directory on the management server called `/cfmroot`, which stored all cluster configuration files.
- ▶ The **`cfmupdatenode`** command and cron job: used on the management server to distribute the configuration files to each node in your cluster every 24 hours via the **`cfmupdatenode`** cron job.

IBM Systems Director

In Systems Director the main tool for remotely managing a system's configuration is through the plug-in named AIX Profile Manager. Using various methods, it is able to collect configurations from systems defined in Systems Director, and based on that it builds its own configuration database. The collected configuration is a set of properties from a few layers of an AIX instance such as reliability, availability, serviceability, security, or kernel.

The AIX Profile Manager plug-in manages AIX system configurations through profiles. A profile is an XML-formatted file that contains a set of runtime properties for a given domain of activity, such as a user, TCP/IP, `kernel_heap_size`, and `/etc/inetd.conf`. A profile contains a set of parameters with optional flags and values.

Profiles are used to deploy configurations on your managed systems in order to align their configuration, to check your managed system's compliance with a reference profile or to retrieve the values of a profile on a system's running configuration.

Beyond the remote synchronization of the configuration files achieved with AIX Profile Manager there may be cases when you want to synchronize files, directories, or drives. This can be achieved using remote access facilities named File Transfer. You can transfer individual files and directories between the following systems: the browser system and the management server, or the browser system and a managed system.

Note: Transferring files directly between a management server and a managed system is not supported. If you need to transfer a file from a management server to a managed system (or the reverse, from a managed system to a management server), you must transfer the file to a browser system first, and then transfer it from the browser system to the managed system or the management server.

You can synchronize a source file, directory, or drive with as many target-system files, directories, or drives as you choose, but you must synchronize the file, directory, or drive on each system individually. You cannot synchronize multiple target systems from a source system at the same time.

From this point on we discuss features available only in IBM Systems Director.

Hardware command line

The hardware command line function, also known as IBM Management Process Command-Line Interface (MPCLI), is run from an established remote session.

Use the hardware command line, with remote systems that have hardware compatible with MPCLI. Hardware command line communicates with Remote Supervisor Adapter (RSA), RSA II, Baseboard Management Controller (BMC), IBM BladeCenter Management Module, or IBM BladeCenter Advanced Management Module.

MPCLI provides system management functions from an easy-to-use command-line interface that connects to a service processor. Using MPCLI, you can access and set a wide range of information about the health, configuration, communication, and state of your system. These functions are immediately available after you make a connection to a service processor.

Remote command line

The remote command line task manages a command-line interface to the remote system, on a management console. You can have multiple remote command-line sessions active at the same time but only one active remote command line session through a management server to a single system at a time.

You can use the remote command line window to establish a fully active command session with a system. A remote command line session is nongraphical, so a command line Java window opens when it is started into the IBM Systems Director main window in a separate tab.

When you connect to a system, remote command line uses the secure shell (ssh) protocol. If the ssh server on the system does not respond, remote command line attempts to connect using the Telnet protocol.

By default, remote command line uses Transmission Control Protocol (TCP). If you disable support for TCP sessions, remote command line uses User Datagram Protocol (UDP).

To access the remote command line from the IBM Systems Director main window into a separate tab, navigate to **System Configuration** → **Remote Access** → **Remote Command Line**.

Remote control tasks

Remote control establishes a full-screen session to the remote system using a remote control application. Examples of remote control applications for different types of endpoints are:

- ▶ BladeCenter and RSA Remote Control - Web-Based Interface
- ▶ AIX endpoint system - Virtual Network Computing (VNC)

To install and configure a remote control application, go to the Remote Access Summary page and click **Set up remote control** in the Common Tasks pane. For remote control tasks through Virtual Network Computing, the server application must be installed first on the endpoint in order to use it with remote control. The steps to install VNC are:

- ▶ Install the RealVNC viewer or another compatible VNC viewer on the web browser client.
- ▶ Ensure that the web browser client has network access to the agent.
- ▶ Install and configure the VNC Server on the agent, using a default port number of 5091.

To install and configure a remote control application, go to the Remote Access Summary page and click **Set up remote control** in the Common Tasks pane.

Remote serial console

The serial console gives you the ability to open console windows to one or more POWER-managed systems. Each window provides access to the system's serial console, accessed out-of-band.

When you use the serial console to open console windows to one or more POWER-managed systems, each window provides access to the system's serial console, accessed out-of-band.

Launching a web browser

Use the Launch Web Browser task to access the default web page for a system that hosts a web server.

4.3.2 System firmware updates

The following section contains information about the system firmware updates.

Firmware updates with CSM

CSM performs hardware inventory scans and maintenance activities on cluster nodes from the CSM management server. For POWER Systems it supports Hardware Management Console (HMC)-attached System p5 and POWER 6 nodes and some limited support for the BladeCenter JS Servers.

The maintenance commands access the controlling HMC using **dsh** to perform tasks like applying service level and release level updates to Licensed Internal Code (LIC) for managed systems or power subsystems. LIC updates are contained within microcode update packages, which could be downloaded from IBM.

Note: In order to use **dsh** commands on HMC you have to run the **updatehwdev -k** command to set up and generate the ssh keys on the CSM management server and transfer the public key to the HMC.

Maintenance commands and their roles are as follows:

► The **mkflashfiles** command

Validates the firmware update packages, determines the hardware type, calculates and stores a checksum and copies the result packages in the `/csminstall/csm/fw` directory. If the `-f` flag is not specified, the **mkflashfiles** command processes all packages in the `/csminstall/csm/fw` directory.

The microcode packages consist of a microcode update file, packaged with an `.rpm` filename extension, and associated XML file, and are stored in subdirectories specific to the hardware and the component being updated.

For different hardware platforms the **mkflashfiles** command copies the microcode update package and associated XML file into automatically created subdirectories specific to the hardware and the component being updated, beneath the base as follows:

/csminstall/csm/fw/8832 <- hardware type for BladeCenter HS20 servers

/csminstall/csm/fw/p5_common <- hardware type for System p5 systems

/csminstall/csm/fw/p6_hv or **/csminstall/csm/fw/p6_ml** <- System POWER 6 servers use different update packages, depending on the model of the server.

► The **rfwflash** command

Runs from the CSM management server to perform LIC updates on the specified HMC-attached System p5 and POWER 6 nodes. It scans the `/csminstall/csm/fw` directory structure for Code Update Packages applicable to the given nodes and components. Output from the **rfwflash** command is sorted by node and written to the `/csminstall/csm/fw/status` directory. This directory contains one file for each node that the **rfwflash** command was run on.

Depending on the Licensed Internal Code (LIC) update that is installed, the affected HMC-attached System p nodes might need to be recycled. The **--activate** flag determines how the affected systems activate the new code. It can be concurrent or disruptive:

- The **concurrent** update option activates code updates that do not require a system recycle.
- The **disruptive** update requires a system recycle, and causes affected systems that are powered on to be powered down before installing and activating the update.

The System p5 and POWER 6 managed system or power subsystem flash chip stores firmware in two locations; the temporary side and the permanent side. By default, most System p5 and POWER 6 systems boot from the temporary side of the flash. When the **rfwflash** command updates code, the contents of the temporary side are written to the permanent side, and the new code is written to the temporary side. The new code is then activated. Therefore, the two sides of the flash contain different levels of code when the update has completed. The two flags that deal with the location of the firmware are:

- The **commit** flag writes the contents of the temporary side of the flash to the permanent side. Use this flag after updating code and verifying correct system operation.
- The **recover** flag writes the permanent side of the flash chip back to the temporary side. Use this flag to recover from a corrupt flash operation, so the previously running code can be restored.

► The **rfwscan** command

Scans the HMC-attached System p5 and POWER 6 nodes to determine currently installed LIC levels. Optionally, the `-w` flag can be used to write selected data to the CSM database.

For HMC-attached System p5 and POWER 6 nodes, a listing of the Managed System and, if applicable, the associated Power Subsystem Licensed Internal Code levels are provided as shown in Example 4-5.

Example 4-5 Listing the managed systems

```
# rfwscan -n p6client01.ibm.com
Nodename = p6client01.ibm.com
Managed System Release Code Level = 01SF240
Active Service Code Level = 320
Installed Service Code Level = 320
Accepted Service Code Level = 320
Power Subsystem MTMS = 9458-100*992056T
Power Subsystem Release Code Level = 02BP240
Active Service Code Level = 214
Installed Service Code Level = 214
Accepted Service Code Level = 214
Managed System Current Release Code Level Primary = 01SF240
Current Service Code Level Primary = 320
Managed System Current Release Code Level Secondary = 01SF240
Current Service Code Level Secondary = 320
Power Subsystem Current Release Code Level A = 02BP240
Current Service Code Level A = 214
Power Subsystem Current Release Code Level B = 02BP240
Current Service Code Level B = 214
```

The **rfwscan** and **rfwflash** commands allow the hostnames or IP addresses of one or more HMCs to be used as the target of the command. When an HMC name is used, the command will perform the requested firmware scan or update on all the hardware controlled by the given HMC. This method may be useful if problems are encountered updating firmware across the cluster. The update may be run only on particular HMCs to help isolate the failure.

A firmware update on Power Systems would look like this:

- ▶ Verify that the **HWType**, **HWModel**, **HWSerialNum**, and **PowerMethod** attributes are defined for each node using:

```
shwinfo -p hmc -c HMC_IP_or_hostname
```
- ▶ Define the target node HMCs as managed devices and set the **RemoteCopyCmd**, **RemoteShell** and **RemoteShellUser** attributes as appropriate.
- ▶ Configure **dsh** and **dcp** for the HMCs defined in the previous step by running the **updatehwdev -k** command.
- ▶ Download the microcode update package and associated XML file from the IBM website and place it into `/csminstall/csm/fw`.
- ▶ Run the **mkflashfiles** command and specify the location of downloaded firmware with the **-f** flag.
- ▶ Run the **rfwflash** command with options like the following:

The **-t** flag to specify the System p5 and POWER 6 components to update: **system** for managed systems or **power** for power subsystems.

Activate the flag to specify the update mode to perform the updates: concurrent or disruptive.

Firmware updates with IBM Systems Director Management Console

IBM Systems Director can handle firmware updates on Power Systems with the help of IBM Systems Director Management Console (SDMC). In order to be able to perform such operations you must have the SDMC properly installed and configured in your environment, as shown in Figure 4-10.

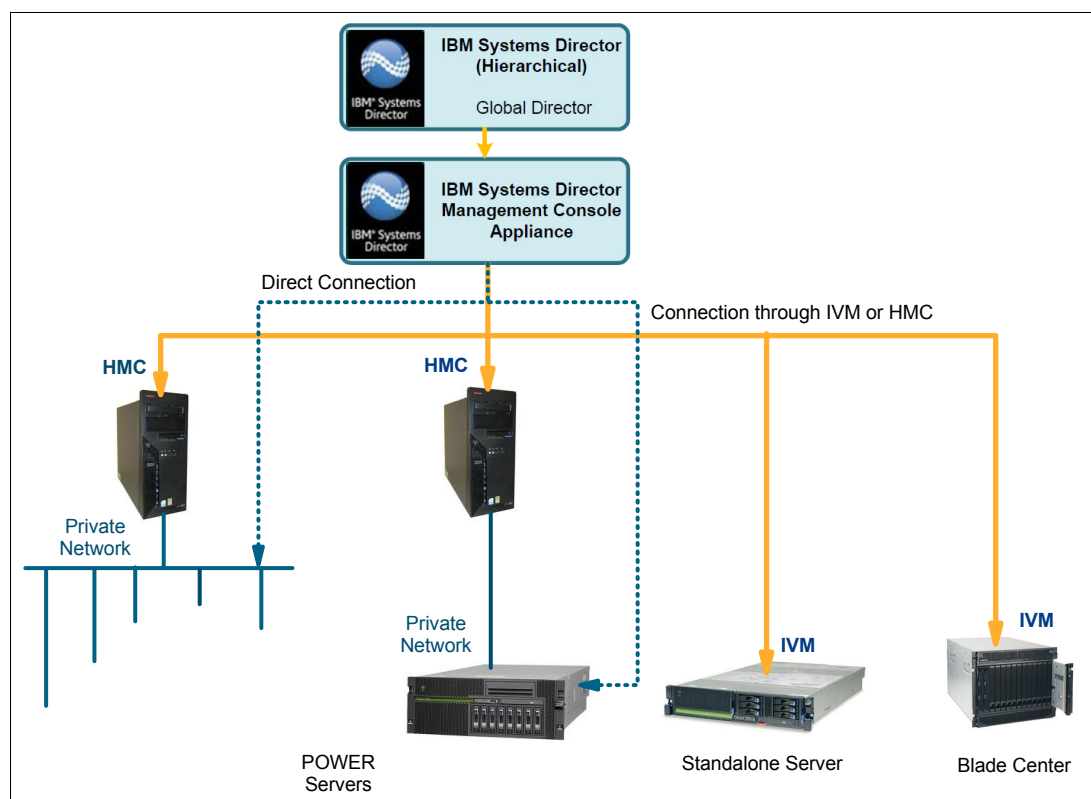


Figure 4-10 IBM Systems Director Management Console clustered environment

SDMC is the successor to the Hardware Management Console (HMC) and the Integrated Virtualization Manager (IVM). It supports only POWER6 and POWER7 processor-based systems (including Power Blades) with the exception of model 575 and is the next generation of management appliances for Power Systems.

IBM Systems Director Management Console is going to replace both Hardware Management Console and Integrated Virtualization Manager in Power Systems administration in the future. It is designed to work as a standalone appliance or it can be integrated into the administrative framework of IBM Systems Director having the same look and feel and providing a common interface for systems administration across the data center.

SDMC enables administrators to work with a more simplified high-level view of systems, which makes it possible to organize tasks in a single panel instead of using different menus as the HMC or IVM do. The transition began in the second half of 2011, because new virtualization features will only be supported by the SDMC.

The SDMC is available as a software appliance that will replace the Integrated Virtualization Manager and a hardware appliance for management of midrange systems and high-end systems. The software appliance is a virtual machine that runs Linux as the base operating system and can run only on top of Red Hat KVM or VMware ESX/ESXi. The hardware appliance will come preinstalled on a System X and uses RHEV-H hypervisor.

For more details about SDMC, refer to *IBM Systems Director Management Console: Introduction and Overview*, SG24-7860.

The IBM Systems Director Update Manager plug-in is installed by default during the installation of IBM Systems Director and is used to acquire, install, manage updates, and monitor your systems to ensure that they are up-to-date. A check for updates has to be scheduled to run periodically and the administrator has to check systems for compliance. In case of noncompliance the systems administrator has to take action by installing required updates.

By integrating the IBM Systems Director Management Console with IBM Systems Director, the Update Manager gains more functionalities and you are able to perform firmware updates on Power Systems. Some differences in updating Power Systems firmware compared with traditional HMC are:

- ▶ In distinction of HMC, SDMC requires more update files for an image than just the .xml and .rpm files; it also requires the .dd.xml, .pd.sdd, and .readme.txt files.
- ▶ Updates and upgrades are done through the same installation flow pages.
- ▶ Management operations (Accept or Reject) are performed through different page flows than the installation.
- ▶ To update a Power System's firmware, it has to be discovered in SDMC to collect its inventory and to import the firmware updates into SDMC.
- ▶ The target system must be associated with the desired firmware package. If this relationship is not present, then the Install button will not be enabled when attempting to launch the install. By collecting inventory on the system the relationship is re-established.

As the scope of this book is only to help you to transform your CSM cluster into IBM Systems Director and to reveal its new functionalities further, we will not describe a step-by-step firmware update procedure but present the main steps.

For a step-by-step process, refer to Chapter 7, "Firmware Updates," in *IBM Systems Director Management Console: Introduction and Overview*, SG24-7860.

The Power Systems firmware update process consists of the following actions:

- ▶ **Readiness check** detects whether the system is in a good state before a code update. The system is checked for open serviceable events, operational state of the system, and possible connectivity issues.
- ▶ **Getting the updates** automatically with Update Manager by scheduling periodical checks against IBM FixCentral sites. There is also an option to manually download the updates and import them to IBM Systems Director from a local file system or from an FTP server. The command for manual import is:

```
smcli importupd -r /mnt/cdrom
```
- ▶ **Collecting the inventory** with the scope to collect the current system's firmware version installed on the system.
- ▶ **Install the updates** with the Update Manager which has a function that enables the user to compare the firmware version currently installed on the system with the latest updates downloaded earlier either manually or automatically. If it detects that the installed firmware is older than the downloaded one, then you have the option to go through a wizard and to update your firmware.
- ▶ **Checking the results** of the firmware updates.

4.4 Security

Securing your cluster environment is an important task to prevent unauthorized access to resources within the managed environment. For both performance and security reasons, it is

crucial to understand and control the data that is sent around in your network. If the data you transmit is sensitive and not encrypted, consider isolating that data on a dedicated network. Security policies and controls must be in place to ensure proper control of hardware commands, such as remote console or remote power, is authenticated and authorized, as well as user profiles and grouping, and access lists.

This purpose of this section is to highlight some of the security differences between CSM and Systems Director. However, it is not within the scope of this section to document all the differences between the two products because IBM Systems Director is so much more than just a management tool that CSM was considered to be. Since the security differences are mainly found in the authentication and authorization process, we primarily focus on these:

- ▶ Authentication and authorization
- ▶ Communication topology
- ▶ Roles and responsibilities
- ▶ Users and groups

4.4.1 Authentication and Authorization when using CSM

CSM uses the security functions of RSCT to ensure that the software components in your cluster can properly interact and authenticate the identity of clients, peers, or RSCT subcomponents. This determination is made in such a way that the cluster software component can be certain that the identity is genuine and not forged by some party trying to gain unwarranted access to the system.

Authentication

Authentication allows CSM to send and receive message keys (Message Authentication). Each key is signed by the sender and the signature is verified by the receiver. The information needed to verify the signature of the message is contained in a key and the authentication is done by key exchange, also known as credentials-based authentication.

Be aware that authentication differs from authorization.

Authorization

Authorization used in CSM uses an access control list-based authorization that provides access control to resource classes. The ACL is kept in a stanza format which is used to verify whenever a given host has the permissions to access the resource class or instance for the class named by the stanza. RMC is the only RSCT component that implements authorization.

CSM uses the security provided by the operating system. It does not come with any security hardening software nor any agent-based encryption protocols, which leaves room to ensure some of the more generic security concerns, especially since CSM utilizes the distributed shell (**dsch**) or remote shell command. It is therefore the system administrator's responsibility to fulfill any security obligations that this environment may require, because CSM does not provide any specific security configuration.

Kerberos:

Kerberos is considered one of the few leftovers from the PSSP transition paths because in the early days of PSSP neither the RSCT stack nor AIX provided any robust security authentication stack. Kerberos was therefore needed to enforce and strengthen the security within the management domain clustering environment. CSM did include its own security layer (CtSec) that could be used out of the box, but based on the transition path you might or might not have used it. In this book we do not cover Kerberos authentication and authorization to Systems Director.

Communication topology used by CSM

The communication topology used by CSM is centered around RSCT. As shown in Figure 4-11, there are several communication paths between these subsystems and the client commands that access the RMC infrastructure. We will not go into any great detail because this is not within the scope of this book. However, just to name a few communication interactions they could be RMC to RMC, RMC to RM, or RMC to the RSCT client.

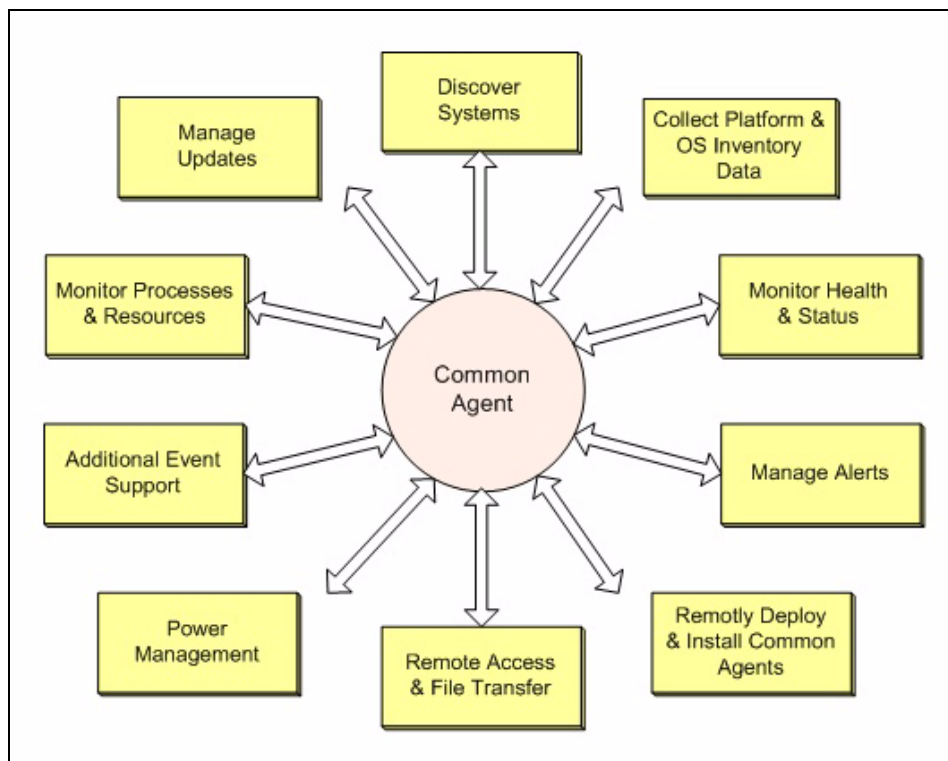


Figure 4-11 Basic CSM topology communication overview

RMC to RMC

All communication between RMC daemons (running on different machines) is done through UDP, using service `rmc=657/udp`. Such daemon-to-daemon communication takes place if the target of the RSCT command is a management server. The RMC daemon on the management server will connect to the RMC daemons on all the managed nodes through daemon-to-daemon communication.

RMC to RMs

The RMC daemon directly connects only to resource managers (RMs) residing on the same machine. For this connection, it always uses UNIX-Domain Sockets (UDS). If access to remote resources is required, the RMC daemon contacts the remote RMC through UDP, and the remote RMC daemon in turn connects to its local RM to access the resource in question.

RMC to RSCT client

The RMC daemon directly connects only to resource managers (RMs) residing on the same machine. For this connection, it always uses UNIX-Domain Sockets (UDS). If access to remote resources is required, the RMC daemon contacts the remote RMC through UDP, and the remote RMC daemon in turn connects to its local RM to access the resource in question.

Using the RMC subsystems

The RMC infrastructure consists of the following:

- ▶ Each AIX machine runs the RMC daemon `rmcd`, which is under SRC control.
- ▶ The CtSec library starts the Cluster Authentication Service (CAS) daemon `ctcasd` on demand when CtSec authentication is required. This daemon is also under SRC control. Both are in the SRC group `rsct`.

The data that CSM stores in the system registry is managed through CSM-specific resource managers and mainly resides on the management server. However, the RMC/SR infrastructure can be exploited in many different ways, such as on individual machines.

Furthermore, RMC itself only provides an abstract view of resource classes and resource instances. Different resource managers (RMs) provide access to the actual resources. Some RMs are shipped with RSCT, while other components such as CSM implement their own RMs. All RMs are under SRC control, in the SRC group `rsct_rm`. Running the command `lsrsrc -ls ctrmc` will produce a detailed status report of the RMC subsystem, including the active resource managers.

For more detailed information about RSCT and CSM, visit the Libraries Related to Cluster Products section at the IBM Cluster Products information center website or the following books.

SA23-1343-05 - *CSM for AIX and Linux V1.7.1: Administration Guide*

SA23-1344-05 - *CSM for AIX and Linux V1.7.1: Planning and Installation Guide*

SA22-7889-20 - *RSCT: Administration Guide*

SA22-7890-20 - *RSCT for AIX: Technical Reference*

4.4.2 Authentication and authorization when using IBM Systems Director

Systems Director offers a number of security features by using the authentication and user administration provided. System administrators can specify user privileges for specific tasks of resources. This user registry integration, integrity and secure data transmission are key elements of the basic security model. Systems Director is controlled by two basic independent processes, authentication and authorization.

Authentication

Systems Director uses authentication to determine the identity of the user. It verifies and validates that identity. When using default registry settings users are authenticated using their user IDs and passwords, which reside locally on the operation which is then verified against the information stored in the user registry that is configured by Systems Director. It also uses the group information. Besides default registry, options like Lightweight Directory Access Protocol (LDAP) are also supported,

Authorization

Systems Director uses authorization to validate the roles and privileges of the authenticated user. This occurs when already authenticated users use Systems Director to perform a specific task on a resource. If the task or role exists and contains the authorization necessary to complete the task or role that was specified, then it proceeds. If a user wants to run additional commands using `smcli` then additional roles need to be authorized.

In conjunction with this, Systems Director heavily relies on role-based access control (RBAC), which allows administrators to create and customize sets of permissions, also known as roles, and assign them to individual users or groups; refer to Figure 4-12.

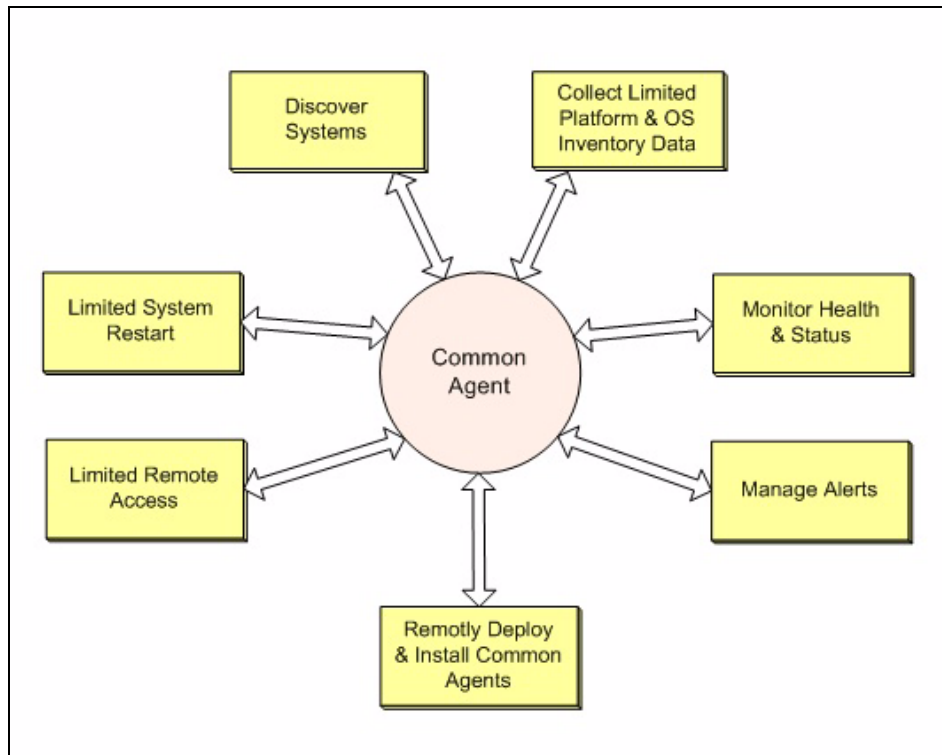


Figure 4-12 Basic IBM Systems Director communication topology

The Systems Director Server by default uses a Secure Sockets Layer (SSL) for the communication between the web management console and the server, and depending on its target such as agent, network devices, or storage devices, different communication types may be utilized. IBM Systems Director stores all of its data in a database repository, which also contains credentials such as userid and passwords that are used for accessing remote systems. All of the sensitive data, including credentials, is encrypted using the 3DES algorithm.

Notes:

- ▶ The default certificate used by Systems Director should be replaced by either a self-signed certificate or by using one signed by a certificate authority (CA). This is to ensure data privacy. The keystore password should also be changed.
- ▶ An SSH server is not provided by the Systems Director Server software, so use the SSH server included by the operating system or a third party.

Using the agent manager

The agent manager provides the authentication and authorization for installed common agent resources and maintains a registry of configuration information about the common agent managed systems. It also provides the core agent manager functionality services:

- ▶ Service catalog
- ▶ Credential manager
- ▶ Agent registry
- ▶ Querying service

Table 4-2 General security comparison

Security Topics	Cluster System Management	IBM Systems Director
User IDs and passwords	HMCs, console servers, and RSAs all require users to authenticate before executing any commands. This includes the CSM management server. User IDs and passwords for each console server, RSA, and HMC in the cluster are stored in the CSM database.	Similar to CSM but the authorization mechanism compares the user account, or the group to which the user belongs, to the role-based access control (RBAC). The agent manager then interacts with the user registry, where user- and group-related information is stored using SSL.
Resource Monitoring and Control access control lists	Commands, such as rpower (which can power on or off nodes and get their power status) and lshwinfo (which reports on the hardware in a node), use the security functions of RMC to determine who is allowed to run them. Access to the hardware control classes, and to actions on these classes, is controlled by stanzas in the <code>/var/ct/cfg/ctrmc.ac1s</code> file.	rpower in IBM Systems Director works similar to the CSM counterpart. However, due to the security differences it is necessary that the HMC managing the resource is discovered and the user requesting this command is properly authenticated. The agent manager is then responsible for authentication and authorization services between the management server, HMC, and common agents.
Console server security	The rconsole command opens a console window for a node. It uses the Conserver open source package to provide support for multiple read-only consoles on a single node.	The dconsole command works similar to the rconsole command, and the security authentication path is the same as for rpower .
Group Service and Topology Services	Group services and topology services, although being part of RSCT, are not used in the management domain structure of CSM. These two components are used in peer domain clusters for applications and are often referred to as hats, hags, high availability Group Services daemon (hagsd) and high availability Topology Service daemon (hatsd).	Not used by IBM Systems Director.

Security Topics	Cluster System Management	IBM Systems Director
Remote command execution	By default, dsh relies on the “classic” rsh command for remote execution. Unfortunately, rsh provides only a minimum security level. The authorization is based on the .rhosts file stored in the user’s home directory. The data exchanged between the management server and the nodes is not encrypted.	By default this relies on the secure shell (ssh). Should the remote system ssh server not respond to the request, then the remote command will try regular Telnet. Both TCP (default) and UDP are supported.

Tip: When accessing an agentless managed system from Systems Director, it is considered best practice to configure access using a user account other than root. This way you can limit the functions performed by the user account and enhance the information provided for audit purposes.

Table 4-3 lists the different agents used by IBM Systems Director when managing IBM Power-based systems. Be aware that other communication types might be used for different hardware architectures, which is not covered in this book.

Table 4-3 IBM Systems Director agent communication protocol for Power-managed systems

Managed System	Communication Type	Encryption used
Agentless	Simple Network Management Protocol (SNMP) v1 and v2.	Not encrypted.
	Simple Network Management Protocol (SNMP) v3.	Advanced Encryption Standard (AES) or Data Encryption Standard (DES).
	Secure Shell (SSH).	Encryption algorithm is negotiated.
Platform Agent	Agentless.	Supports the communication protocols and encryption listed for the agentless managed systems.
	Common Information Model (CIM).	If configured, using SSL on port 5989.
Common Agent	Tivoli Common Agent Service 6.x.	Encrypted Web Service (SSL).
Other	Service Location Protocol (SLP).	Not encrypted.

4.4.3 Security consideration when defining users and groups

System administration is an important aspect of daily operations, and security is an inherent part of most system administration functions. Also, in addition to securing the operating environment, it is necessary to closely monitor daily system activities.

Most environments require that different users manage different system administration duties. It is necessary to maintain separation of these duties so that no single system management user can accidentally or maliciously bypass system security.

When using CSM

CSM relies on the traditional AIX user management approach: using a single system administrator account named root that can perform all privileged system administration tasks on the system. Reliance on a single user for all system administration tasks is often seen as a problem in regard to the separation of duties. While a single administrative account is acceptable in certain environments, many environments require multiple administrators, with each administrator responsible for different system administration tasks.

In order to share the administration responsibilities with multiple users of the system, the historical practice was to either share the password of the root user or create another user with the same UID as the root user. This method of sharing system administration duties presents security issues, since each administrator has complete system control and there is no method to limit the operations that an administrator can perform. Since the root user is the most privileged user, root users can perform unauthorized operations and can also erase any audits of these activities, making it impossible to track these administrative actions.

As each user logs in to the system, the user supplies the user name of an account and a password if the account has one. If the password is correct, the user is logged in to that account and acquires the access rights and privileges of the account. The `/etc/passwd` and `/etc/security/passwd` files maintain user passwords.

CFM can be used to manage user accounts by using symbolic links to add the password and group files to the CFM repository. User accounts can be replicated to all nodes in the cluster. The steps required to set this up are shown. If you use indexed password files, you can distribute the `.idx` files using the same method of symbolic links, or you could write a `.post` script to have CFM run `/usr/sbin/mkpasswd` after the initial files have been distributed.

You can also keep user accounts on your cluster nodes separate from the accounts defined on your management server. To accomplish this, you would copy the actual password and group files into the CFM repository instead of creating symbolic links. This method would require additional work to define where the master files reside, and copy them to the management server before running CFM. Identification and authentication are used to establish a user's identity.

Most environments require that different users manage different system administration duties. It is necessary to maintain separation of these duties so that no single system management user can accidentally or maliciously bypass system security. With AIX 6.1 IBM added additional security improvements with features like Role Based Access Control (RBAC), Trusted AIX, and Trusted Execution, allowing system administrators to implement the much needed user separations.

When using Systems Director

In IBM Systems Director, users and user groups are based on users and groups that are defined in the configured registry, which is associated with either the IBM Systems Director Systems Management Guide operating system or Lightweight Directory Access Protocol (LDAP). IBM Systems Director uses the user and group information for the purpose of authentication and authorization.

IBM Systems Director does not provide the capability to create, update, or delete users or groups in a user registry regardless of where the registry resides. To manage users or groups in the user registry, use instead the appropriate tool associated with the registry in which the

users or groups are stored. IBM Systems Director does, however, give you the ability to enter and edit information for each user or group that describes each in the context of IBM Systems Director.

Access to particular resources or tasks is done by restrictions based on the user ID or user group membership and the roles that are defined for each user. For a user to access IBM Systems Director Server, one of the following conditions must exist:

- ▶ The user is a member of a user group that is authorized for IBM Systems Director Server.
- ▶ The user is a root user on the management server.

Authorization is the process that determines whether an authenticated user or group has the necessary privileges to access specific resources. With user authorization, IBM Systems Director users can perform tasks on specific resources by using the IBM Systems Director web interface. Use IBM Systems Director to configure the authorizations that provide access to IBM Systems Director tasks and resources. [The authentication flow is depicted in Figure 4-13.](#)

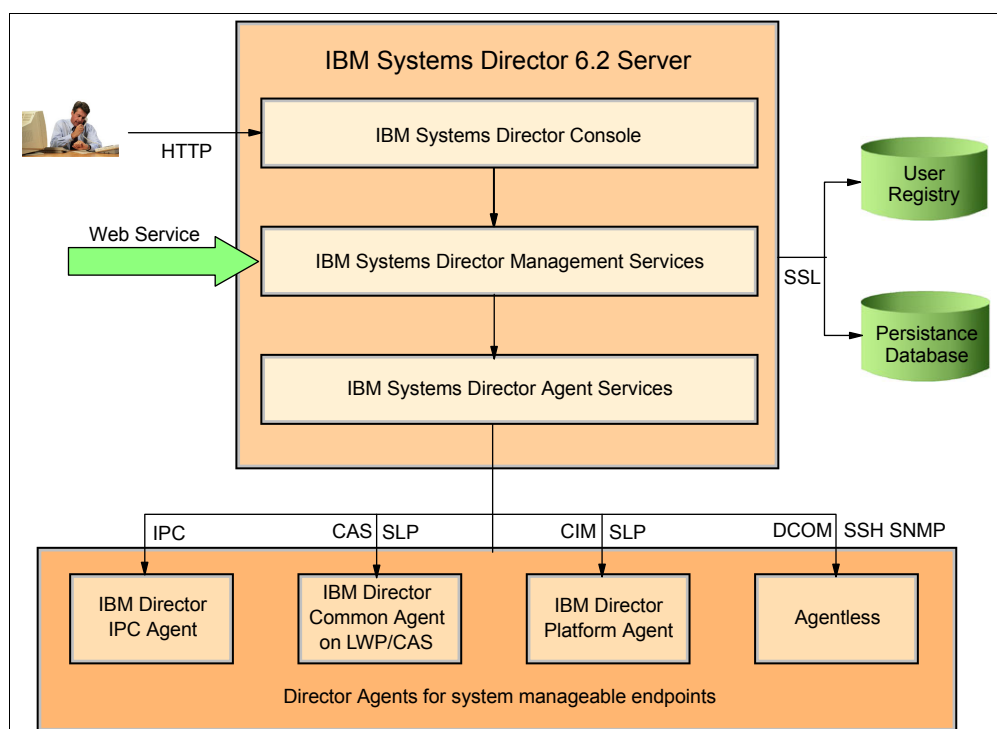


Figure 4-13 Authentication flow

To log in to the IBM Systems Director web interface and manage the resources that are discovered by IBM Systems Director, you must have a user account that is associated with a role that has the appropriate authority.

The following steps are required to authorize an IBM Systems Director user to manage resources (see [Figure 4-13 on page 103](#)):

1. If the user account that is needed does not already exist, create it on the operating system of the system that you want to manage or on the Lightweight Directory Access Protocol (LDAP) server.
2. Log in to IBM Systems Director as an SMAdministrator.
3. Assign an appropriate role to the user account or group to which the user account belongs and associate it with the resources that you want the account to manage. You can use any

of the existing initial role groups (smadmin smmgr, smmon, smuser) or you can create a new role that includes the privileges necessary to access the appropriate resource.

4. Our user wants to access a resource managed by the IBM System Director by using the web user interface.
5. The user login request is handled by the user authentication process, which verifies whether the user credentials allow access to the system. A user must authenticate by logging in with a user registry level account for the management server.
6. IBM Systems Director then interacts with the user registry where user- and group-related information is stored.
7. If the user wants to access a specific resource or perform a specific task on another managed system, then the appropriate credentials and passwords are also verified. If Single Sign-on (SSO) is used, then Credential Transformation Service (CTS) is utilized.
8. Depending on the agent communication scheme used, credentials and mappings are created for the agent access points. Should the resource be listed as offline, then access cannot be requested.
9. After the system is accessed, additional tasks and status are now available. Should this be the first time the system is accessed, then the padlock icon disappears.

Note: The only role that is automatically assigned is to the administrator user ID that installed IBM Systems Director. So, initially, no other user is associated with a role. The IBM Systems Director administrator must then associate the other users with roles.

Assigning a role other than smadmin, which is done with the operating system, occurs within the IBM Systems Director web interface.

If you want to use Lightweight Directory Access Protocol (LDAP) or another tool that the registry supports, you might need to manually create all of these user groups and assign users to them.

4.4.4 Roles

In Systems Director you assign roles that will control users and their access to resources. Based on the role defined you can control the tasks they can perform. The authorities that you configure for a role determine the level of access granted to each user assigned to this role. IBM Systems Director comes with a number of predefined roles:

► **SMAAdministrator (administrator role)**

The SMAAdministrator role has full authority to perform all tasks and functions, and has full control over permissions. A user assigned to this role can perform all operations (including security administration, product installation, and configuration) with any resource.

► **SMMManager (manager role)**

The SMMManager role can perform management operations, which are a subset of the functions that a member of the SMAAdministrator 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, but it has full access to all of the IBM Systems Director functions included within a functional manager or feature. The list of accessible functions includes those within the discovery manager, status manager, configuration manager, and update manager.

► **SMMonitor (monitor role)**

The SMMonitor role can access those 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 inventory, and viewing hardware status. This role cannot, for example, create virtual servers or reconfigure the IBM BladeCenter.

► **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, 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 to see the group contents.

The hierarchical order of the predefined roles is: the SAdministrator role has the highest authority and the SMUser role has the lowest authority. Roles with higher authority are permitted to run all operations that roles with lower authority are permitted to run. This means that if the execution of an operation is permitted for SMUser, then all of the other roles with higher privileges are also permitted to run the operation.

All users or groups of users that access IBM Systems Director must have a user role assignment. IBM Systems Director comes with predefined user roles that correspond directly with what IBM Systems Director installs at the operation system level. These cannot be deleted or modified—you can copy or create new ones to match your requirements.

4.5 OS deployment

The OS deployment can be divided into the following steps:

- Plan OS deployment
 - Go through all the steps in the planning phase because some of the chosen methods can determine other steps in a very early phase of the deployment.
- Prepare the deployment infrastructure
 - Prepare the installable image, which can be on DVD, predefined NIM resources in case of an already existing NIM environment for most production CSM clusters, or a backup from a running operating system.
 - Prepare the network on the deployment infrastructure.
 - Allocate the installable resources, such as LPP source, backups, OS images, and customization scripts.
- Prepare the target HW
 - Set up the install network.
 - Allocate storage space for the operating system.
 - Set up virtualization according to your requirements.
 - This can determine what deployment infrastructure you can use.
 - LPAR configuration.
- Install the operating system using the deployment infrastructure.

- Automated deployment requires HW control-related features available on most systems management infrastructures, such as remotely manage the power on the system or partitions and set the installable partition to boot from a specific device.
- Control the installation process so that it will use the correct target storage space and other preset parameters.
- Run post-install customization.
 - Configure network settings such as hostname, IP address, gateways, and name resolution.
 - Set up additional network adapters.

In some environments this can be a final step, because the new server cannot be alive on the public or production network until all the security hardening and other activation procedures are done.
 - Set up NTP.
 - Import and activate additional volume groups.
 - Install necessary management agents.
 - Customize security settings, OS hardening.
 - Run any other post-install scripts necessary for your environment.

In the following sections, we compare the main deployment steps that are available for CSM or Systems Director.

4.5.1 Comparison of installation services

In case of CSM, the main OS deployment infrastructure is the NIM feature of AIX.

CSM controls the creation of the NIM machine definition and the handling of NIM customization scripts, which is necessary to set the installed NIM client to be a CSM node. In case SSH is used for remote command and copy services, the post-install customization script transfers the CSM management server's SSH public key to the client.

In case you use Systems Director's deployment infrastructure, the install source can be a valid NIM resource captured by VMControl, or a system image that is used by VMControl's Storage Copy Services.

In the next sections, we explain more about how to use NIM as standalone installation service, or use VMControl features together with NIM to handle sources for OS deployment.

Standalone NIM environment used with Systems Director

It is possible to use NIM separately from Systems Director. That way you save all the efforts you have spent so far to define and handle NIM resources and automation.

When you have a stable environment and all base Systems Director features are used according to your needs, you can start using additional Systems Director extensions and plug-ins, such as VMControl with NIM integration.

In a standalone NIM configuration you keep the system backups, or install sources as they are already defined in NIM.

► NIM clients

All the former CSM nodes can remain defined as standalone NIM machines, when we remove them from the CSM cluster.

- Backup images

Our experience shows that system backups are usually created in a form of AIX mksysb and are stored on the NIM servers. For every mksysb image file a NIM resource is defined, which makes it possible to restore the backup to any of the NIM clients.

- Install and update sources

An LPP source NIM resource can be created from install DVDs or update packages and stored in NIM. You have to integrate this with Systems Director because it is using NIM for AIX updates even if VMControl is not used for OS deployment.

- Management objects for HW control

The new features in NIM enable HW control as well to automate the first phase of an installation. You can define management objects in NIM for HMCs and CECs where the LPAR is created, and NIM can start the LPAR in open firmware mode and configure it to boot from the network and start the OS installation.

- Post-install customization scripts

Because you want to integrate the installed NIM client to an existing Systems Director environment, you have to be sure that a correct version of the Systems Director Common Agent is installed at OS installation time (so that the filesets are part of the LPP source or it is installed on the NIM client on which the mksysb is created), and the post-install customization scripts will register to the appropriate Director server.

NIM integration - VMControl Standard Edition

NIM is necessary to install updates on an AIX-based client and it can be used with a Systems Director VMControl plug-in for base operating system deployment.

For Systems Director VMControl integration the Systems Director Common Agent and the VMControl NIM subagent have to be installed on the NIM server. The NIM server is discovered as image repository in VMControl.

Example 4-6 shows a NIM server as an image repository in VMControl.

Note: You have to use the value shown as OID for VMControl operations run on the command line (**smcli**).

Example 4-6 Listing image repository parameters

```
p55701p01(root)/export> smcli lsrepos -l -v
Wed Sep 28 16:18:04 EDT 2011 lsrepos Operation started.
Attempt to get repository criteria.
p55701p01 (NIM Repository)
  ClassName:com.ibm.usmi.datamodel.software.ImageRepository
  UniqueId:e56343ac-5c4b-4efe-85c5-d8c1857085f2
  ImageRepositoryType:1
  ChangedDate:2011-09-26T11:52:32-04:00
  SourceTokens:{ 'NO_IR_DELETE' }
  DisplayName:p55701p01
  CreatedDate:2011-09-26T11:52:32-04:00
  ImagingTool:DISCOVERY_NIM_REPOSITORY
  OID:8797
  Guid:F2C71133FD9439B6B5A155DBA442A600
  ObjectType:ImageRepository
```

If the NIM server is used for other purposes also, then create a separate file system or directory for those resources to allow VMControl full control of the resources handled in the VMControl-managed image repository. VMControl creates directories for appliances under the /export/nim/appliances directory.

Utilizing NIM from VMControl will use new features in NIM and also commands that are available in the Distributed Systems Management (DSM) filesset. These features make it possible to define so called management objects in NIM for HMCs and CECs and run commands to manage the boot process, which is necessary for an automated OS deployment. See the new NIM features in *IBM AIX Version 7.1 Differences Guide*, SG24-7910.

VMControl manages virtual appliances. A virtual appliance is created for installation sources such as LPP source and mkysyb image, and this virtual appliance is used later at OS deployment. For these virtual appliances a corresponding NIM resource is created:

► LPP source

If an already existing LPP source NIM resource is captured, then VMControl will not define a new one and the original LPP source directory is not duplicated.

The virtual appliance and the NIM resource are linked via a file in the appliances directory in a similar way as for mkysyb images.

► mkysyb

The captured mkysyb file itself is placed in the directory that is created for the appliance on the image repository (NIM server). A NIM mkysyb resource is also created, which points to the actual mkysyb file. If an already existing mkysyb NIM resource is captured, then VMControl will not define a new one and the original mkysyb file is not copied.

The virtual appliance and the NIM resource are linked via a file in the appliances directory, as shown in Example 4-7.

Example 4-7 Link between mkysyb-based virtual appliance and NIM resource

NIM resources:

```
p55701p01(root)/> lsrim|grep mkysyb
appliance-0_image-1      resources      mkysyb
p55701p03_mkysyb       resources      mkysyb
```

Captured from an mkysyb file:

```
p55701p01(root)/export/nim/appliances> ls -l 7*/*
-rw-r--r--    1 root    system      10040 Sep 28 17:37
76891ba1-db22-45d6-8bd1-311fcc51198c/76891ba1-db22-45d6-8bd1-311fcc51198c.ovf
-rw-r--r--    1 root    system    1336729600 Sep 28 17:37
76891ba1-db22-45d6-8bd1-311fcc51198c/p55701p02.mkysyb

76891ba1-db22-45d6-8bd1-311fcc51198c/vmcontrol/:
total 0
-rw-r--r--    1 root    system          0 Sep 28 17:37 mkysyb.appliance-0_image-1
```

Captured from a preexisting NIM mkysyb resource:

```
p55701p01(root)/export/nim/appliances> ls -l b*/*
-rw-r--r--    1 root    system      10037 Sep 28 18:17
bf7ad9de-fdc5-45fb-9bd2-ad6a726f0e01/bf7ad9de-fdc5-45fb-9bd2-ad6a726f0e01.ovf
```

```
bf7ad9de-fdc5-45fb-9bd2-ad6a726f0e01/vmcontrol/:
total 0
-rw-r--r-- 1 root system 0 Sep 28 18:17 mksysb.p55701p03_mksysb
-rw-r--r-- 1 root system 0 Sep 28 18:17 user_defined_resource.p55701p03_mksysb
```

► **Shared Product Object Tree (SPOT)**

The Shared Product Object Tree is used when the NIM client is network booted from the NIM server to provide base OS functionality during the install process.

A SPOT resource has to be created from the LPP source for a rte installation or from an mksysb if that is used to install the client. Depending on the LPP source handling and installation process, this can be done manually or via scripts when you perform CSM node installation.

SPOT creation is automatic in VMControl and the NIM resource remains defined after it is created and used. VMControl will reuse the SPOT at subsequent installations. VMControl removes a SPOT only when the virtual appliance for which the SPOT was created is deleted by the user.

VMControl operations

VMControl manages virtual appliances as part of its installation services. It provides the following operations regarding virtual appliances:

► **Import**

This operation is for importing an existing virtual appliance package to the actual Systems Director VMControl environment. The NIM resources are created according to the content of the appliance. This is useful for moving servers between Systems Director management environments.

► **Capture**

Capture an existing NIM resource such as mksysb or an LPP source. This is possible only on the command line using the **smcli captureva** command.

Example 4-8 shows how to capture an existing mksysb file. The command copies the file to the directory created for the new virtual appliance on the NIM server. The existing mksysb file is not removed.

Example 4-8 Capture an existing mksysb file

Capture command:

```
p55701p01(root)/> smcli captureva -n p55701p02_mksysb_file -r 8797 -F
repos://export/p55701p02/p55701p02.mksysb
```

Copy of mksysb file started in the background:

```
p55701p01(root)/> ps -ef|grep mksysb

root 10813628 11272230 0 17:35:53 pts/5 0:00 smcli captureva -n
p55701p02_mksysb_file -r 8797 -F repos://export/p55701p02/p55701p02.mksysb

root 10944652 12648512 61 17:35:56 - 0:03 cp -rf
/export/p55701p02/p55701p02.mksysb
/export/nim/appliances/76891ba1-db22-45d6-8bd1-311fcc51198c
```

NIM resource defined for mksysb:

```
p55701p01(root)/export/nim/appliances> lsnm|grep mksysb
```

```

alfreshsds resources mksysb
appliance-0_image-1 resources mksysb
p5570lp03_mksysb resources mksysb

p5570lp01(root)/export/p5570lp02> lsnim -l appliance-0_image-1

appliance-0_image-1:

    class = resources
    type = mksysb
    Rstate = ready for use
    prev_state = unavailable for use
    location =
    /export/nim/appliances/76891ba1-db22-45d6-8bd1-311fcc51198c/p5570lp02.mksysb
    version = 5
    release = 3
    mod = 12
    oslevel_r = 5300-12
    alloc_count = 0
    server = master
    creation_date = Wed Sep 28 17:37:22 2011

```

Virtual Appliance attributes as shown by lsva -l command:

```

p5570lp01(root)/export/nim/appliances> smcli lsva -l

p5570lp02_mksysb_file

    TrunkId:1
    ClassName:com.ibm.usmi.datamodel.virtual.VirtualAppliance
    RevisionVersion:1.1
    Description:
    ChangedDate:2011-09-28T17:37:25-04:00
    TrunkName:p5570lp02_mksysb_file
    DisplayName:p5570lp02_mksysb_file
    CreatedDate:2011-09-28T17:37:24-04:00
    SpecificationId:1
    SpecificationVersion:1.1
    OID:11496
    Guid:DB505DC488F83E86BD4E0D067F2B213C
    ApplianceId:76891ba1-db22-45d6-8bd1-311fcc51198c
    ObjectType:VirtualAppliance

```

Example 4-9 shows how to capture an existing NIM mksysb resource. Because the NIM resource and the backing mksysb file already exist, only a new virtual appliance definition is created and linked to the NIM resource.

Example 4-9 Capture an existing NIM mksysb resource

Capture command:

```

p5570lp01(root)/> smcli captureva -n p5570lp03_mksysb_NIM -r 8797 -F
repos:p5570lp03_mksysb

```

As NIM resource already exists, no new resource is defined:

```
p55701p01(root)/> lsrim -l p55701p03_mksysb
```

```
p55701p03_mksysb:
```

```
class      = resources
type       = mksysb
Rstate     = ready for use
prev_state = unavailable for use
location   = /export/p55701p03/p55701p03.mksysb
version    = 6
release    = 1
mod        = 6
oslevel_r  = 6100-06
alloc_count = 0
server     = master
creation_date = Wed Sep 28 18:16:45 2011
```

Virtual Appliance attributes as shown by lsva -l command:

```
p55701p01(root)/> smcli lsva -l -q name=p55701p03_mksysb_NIM
```

```
p55701p03_mksysb_NIM
```

```
TrunkId:2
ClassName:com.ibm.usmi.datamodel.virtual.VirtualAppliance
RevisionVersion:1.1
Description:
ChangedDate:2011-09-28T18:17:53-04:00
TrunkName:p55701p03_mksysb_NIM
DisplayName:p55701p03_mksysb_NIM
CreatedDate:2011-09-28T18:17:52-04:00
SpecificationId:1
SpecificationVersion:1.1
OID:11499
Guid:3B8E6D1FEA3233AA9C100D44F3E1EE5A
ApplianceId:bf7ad9de-fdc5-45fb-9bd2-ad6a726f0e01
ObjectType:VirtualAppliance
```

Link between the mksysb NIM resource and the virtual appliance is based on a file in VA-s vmcontrol directory:

```
p55701p01(root)/export/nim/appliances/bf7ad9de-fdc5-45fb-9bd2-ad6a726f0e01> ls
-l *
```

```
-rw-r--r-- 1 root system 10037 Sep 28 18:17
bf7ad9de-fdc5-45fb-9bd2-ad6a726f0e01.ovf
```

```
vmcontrol/:
total 0
```

```
-rw-r--r-- 1 root system 0 Sep 28 18:17
mksysb.p55701p03_mksysb
-rw-r--r-- 1 root system 0 Sep 28 18:17
user_defined_resource.p55701p03_mksysb
```

Check **smcli lscustomization** for attributes that can be changed and incorporated into the created virtual appliance. One of these attributes is the **disksize**, which tells the size of the virtual disk to be created in case of a new deployment. Example 4-10 shows the attributes that can be changed when capturing a NIM resource. The image repository object ID was 10826, which can be checked with the **smcli lsrepos -o** command.

Example 4-10 Customization attributes for the smcli captureva command

```
# smcli lscustomization -a capture -r 10826
cpushare
    Value: 1.0
    Min: 0.1
    Max: 0.0
    Increment: 0.1
    Description: Number of virtual processors

memsize
    Value: 1024
    Increment: 1
    IncrementType: LINEAR
    Description: Memory (MB)

disksize
    Value: 0
    Increment: 1
    IncrementType: LINEAR
    Description: Disk Size (bytes)
```

VMControl can also capture images as operating system backup for a running operating system that is discovered by Systems Director, as shown in Example 4-11. The Common Agent does not have to be installed on the server, but the HMC has to be managed (discovered and inventory collected) by Systems Director. This operation is managed by the NIM master, which calls remote NIM commands on the virtual server's operating system.

If the virtual server is already a NIM client, then that definition will be used to manage it. If not, then a new standalone machine is defined. In any case, after the capture the NIM machine definition is removed.

Example 4-11 Capture a running AIX system

Capture command:

```
p55701p01(root)/> smcli captureva -v -r 26950 -n p55701p03_mksysb_cap -s 0x549b
Tue Jan 10 09:06:29 EST 2012 captureva Operation started.
Get capture customization data
Call capture function
Call capture command executed. Return code= 27,211
Tue Jan 10 09:15:54 EST 2012 captureva Operation took 564 seconds.
```

NIM operations started to create NIM mkysyb resource:

```
p55701p01(root)/export/nim/appliances> ps -ef|grep p55701p03
```

```

root 3801240 7143634 0 09:06:47 - 0:00
/usr/lpp/bos.sysmgmt/nim/methods/m_mkbsi -aforce=yes -t mksysb -a server=master
-a
location=/export/nim/appliances/04a61ed1-1bd0-482b-b69c-3423636a5946/image1.mks
ysb -a mk_image=yes -a source=p5570lp03 -a mksysb_flags=X e appliance-4_image-1

```

```

root 7143634 6029548 0 09:06:47 - 0:00 nim -Fo define -t mksysb -a
server=master -a
location=/export/nim/appliances/04a61ed1-1bd0-482b-b69c-3423636a5946/image1.mks
ysb -a mk_image=yes -a source=p5570lp03 -a mksysb_flags=X e appliance-4_image-1

```

Process tree on client system to capture:

```

p5570lp03(root)/> proctree 11141300
3342474 /usr/sbin/srcmstr
11337850 /usr/sbin/nimsh -s
3145822 /usr/sbin/nimsh -s
7012430 /bin/ksh /usr/lpp/bos.sysmgmt/nim/methods/c_nimpush
/usr/lpp/bos.sysmgmt/nim/meth
11141300 /bin/ksh /usr/lpp/bos.sysmgmt/nim/methods/c_mkbsi
-aserver=p5570lp01 -alocation
5439522 /bin/ksh /usr/bin/savevg -i -f /tmp/11141300.mnt0/image1.mksysb -X
-e rootvg
4980752 backbyname -i -q -v -Z -p -U -f /tmp/11141300.mnt0/image1.mksysb
10223808 /usr/bin/cat /tmp/mksysb.5439522/.archive.list.5439522
11075626 /bin/ksh /usr/bin/savevg -i -f /tmp/11141300.mnt0/image1.mksysb -X
-e rootvg
10813538 /usr/bin/sleep 10

```

Virtual Appliance attributes as shown by lsva -l command:

```

p5570lp01(root)/> smcli lsva -l -q "name=p5570lp03_mksysb_cap"

```

```

p5570lp03_mksysb_cap
  TrunkId:6
  ClassName:com.ibm.usmi.datamodel.virtual.VirtualAppliance
  RevisionVersion:1.1
  Description:
  ChangedDate:2012-01-10T09:15:52-05:00
  TrunkName:p5570lp03_mksysb_cap
  DisplayName:p5570lp03_mksysb_cap
  CreatedDate:2012-01-10T09:15:52-05:00
  SpecificationId:1
  SpecificationVersion:1.1
  OID:27211
  Guid:3CD08D612A0B3CA8894497F8647BB034
  ApplianceId:04a61ed1-1bd0-482b-b69c-3423636a5946
  ObjectType:VirtualAppliance

```

NIM resource defined for mksysb:

```

p5570lp01(root)/> lsrim -l appliance-4_image-1

```

```

appliance-4_image-1:
  class = resources

```

```

type           = mksysb
arch           = power
Rstate         = ready for use
prev_state     = unavailable for use
location       =
/export/nim/appliances/04a61ed1-1bd0-482b-b69c-3423636a5946/image1.mksysb
version        = 6
release        = 1
mod            = 6
oslevel_r      = 6100-06
alloc_count    = 0
server         = master
creation_date  = Tue Jan 10 09:15:43 2012
source_image   = p55701p03

```

Link between the mksysb NIM resource and the virtual appliance is based on a file in VA-s vmcontrol directory:

```

p55701p01(root)/> ls -l
/export/nim/appliances/04a61ed1-1bd0-482b-b69c-3423636a5946
total 0
-rw----- 1 root system 13655 Jan 10 09:15
04a61ed1-1bd0-482b-b69c-3423636a5946.ovf
-rw-r--r-- 1 root system 3895910400 Jan 10 09:15 image1.mksysb
drwx--x--x 2 root system 256 Jan 10 09:15 vmcontrol/

p55701p01(root)/> ls -l
/export/nim/appliances/04a61ed1-1bd0-482b-b69c-3423636a5946/vmcontrol
total 0
-rw----- 1 root system 0 Jan 10 09:15
mksysb.appliance-4_image-1

```

New NIM resources (LPP source, mksysb) are created and a set of files and directories are also created as part of the appliance package in the `/export/nim/appliances` directory.

Removing the node from the CSM cluster does not remove the NIM machine defined for it, and even the CSM post-install customization script can remain allocated as a NIM resource. This will prevent a successful capture operation, because the machine definition cannot be removed from NIM (because VMControl removes it automatically) until there are resources allocated to it.

Unexport all parent directories and remove them from `/etc/exports` file which is under the default virtual appliance directory: `/export/nim/appliances` as the automatic NFS export which is started by the NIM master will fail and the **mksysb** command on the node will not be able to write into the new virtual appliance directory which has `rwxr-xr-x` (755) access rights.

Note: Capturing a running virtual server (LPAR) will create an mksysb that contains only the rootvg and even if the rootvg spans multiple disks, the virtual appliance definition will contain only one disk, albeit big enough to hold the whole rootvg.

► Deploy

Install the operating system using the appliance and the corresponding NIM resources. VMControl will call all necessary NIM operations and create additional NIM resource definitions to install the operating system.

Other DSM-provided operations are also started from the NIM server to control the installation process—for example, key exchanges between the NIM server and HMC, and MAC address collection.

Note: We had to remove the network interface definition from Systems Director, which belongs to the HMC.

In a CSM cluster beside others, the **systemid**, **getadapters**, **csm2nimnodes**, **csmsetupnim** and **netboot** commands are doing what is automated in VMControl.

The customizable attributes used by VMControl to set up the deployed operating system depend on the type of deployment. To check what attributes are available, use the **smcli lscustomization** command. Many of the attributes are explained in 4.5.2, “Post-install customization” on page 118.

There are two types of deployment operations:

- The **smcli -a deploy_new** command creates a new virtual server or LPAR definition, so this will need to set the processor, memory, and virtual disk information also. In this deployment mode the host or server system pool can be specified as a target with the **-s** flag. The creation of the new LPAR is controlled by the following customization methods:
 - Some of the attributes can be specified at capture time, which will determine the attributes used at deployment as shown in Example 4-12. These will be incorporated into the virtual appliance definition. You can find them in the OVF file created for the virtual appliance.
 - Additional customization attributes can be specified with the **-A** flag of the **smcli deploy -a deploy_new** command.
 - A customization file can be specified with the **-F** flag.

For further explanation and examples, see *IBM Systems Director VMControl Implementation Guide on Power Systems*, SG24-7829.

- The **smcli -a deploy_existing** command deploys an operating system onto an existing virtual server or LPAR. In this case the LPAR basic virtualization attributes can not be changed by the deployment operation. In this mode one or more virtual servers can be specified with the **-s** flag.

When you run a deploy operation, VMControl provides the IP addresses the NIM server should use for key exchanges. It is possible that the HMC's IP address used for communication with the managed systems is provided. The NIM server will not be able to communicate with HMC using that address and gives back the error that is shown in Example 4-12. In this case, HMC has the IP address 10.0.0.1 set for communication with the service processors of the managed systems.

Example 4-12 Deployment error because of failing communication between NIM server and HMC

```
p55701p01(root)/export/nim/cust> smcli deployva -a deploy_existing -V 11703 -s 10765 -F /export/nim/cust/cust_for_p55701p02_>
```

DNZIMC032E Error occurred during deploy operation.

DNZIMN105E An error occurred starting the virtual server to prepare it for deployment.

DNZIMN867E Could not exchange SSH key with 10.0.0.1 due to the following error:
2760-277 [dkeyexch] Timed out waiting for response from target: 10.0.0.1

We found two solutions for this problem:

- Remove the IP interface that belongs to the given IP address from the NIM server's inventory in Systems Director. Be careful because this will be added back at the next inventory collection.
- Exclude the IP address or IP range from the discovery as documented by the "Excluding IP addresses" chapter in *IBM Systems Director for AIX Planning, Installation and Configuration Guide Version 6.2.1*, GI11-8709-06.

Example 4-13 shows the NIM resources that are auto created by VMControl at deployment time.

Example 4-13 NIM resources defined by VMControl at deployment time

```
p5570lp01(root)/> ls nim
...
vmc_config                      resources      script
...
nimrf-000000000000000001-script  resources      script
H172_16_20_107                  management    hmc
cec-570-107CD9E                 management    cec
p5570lp02                        machines      standalone
nimrf-000000000000000002-spot    resources      spot
nimrf-000000000000000003-res_group  groups        res_group
nimrf-000000000000000003-bosinst_data  resources      bosinst_data
nimrf-000000000000000003-resolv_conf  resources      resolv_conf
nimrf-000000000000000003-image_data    resources      image_data
nimrf-000000000000000003-adapter_def    resources      adapter_def
```

► Delete

To delete the appliance use the VMControl delete operation. If the capture was done for a preexisting NIM resource or preexisting mkysyb file or directory that contains the LPP source, then these are not removed.

For a more detailed description and the consequence for NIM resources, see *IBM Systems Director VMControl Implementation Guide on Power Systems*, SG24-7829.

Note: Mkysyb images available as files (even on the NIM server) but do not have NIM resources defined for them have to be captured also, which will copy the mkysyb file under the VMControl-managed appliance directory in the image repository. This causes duplicate storage requirements, so remove the original mkysyb files after the capture operation.

A customer may use tape systems to store the mkysyb files outside of the disk-based storage devices. The backup procedure could need some changes in case of a VMControl-controlled NIM environment.

Because VMControl owns the NIM install procedure from end to end, the already set up post-install customization has to be integrated into the VMControl deployment process. For more details, see 4.5.2, "Post-install customization" on page 118.

VMControl Storage Copy Services

Starting with VMControl V2.3.0, a new concept was introduced entitled Storage Copy Services for virtual appliance management as part of a Server System Pool. The Storage Copy Services (SCS) capability uses a radically different method for virtual server capture and deployment from the NIM methodology discussed here. Because the existing CSM for AIX customer is most typically a NIM user, this document does not go into depth about this new capability. Further information about SCS can be found in *IBM Systems Director VMControl Implementation Guide on Power Systems*, SG24-7829.

With the Storage Copy Services model, the Virtual I/O Servers (VIOS) become the Image Repositories. IBM Systems Director must discover the Common Agent with the VMControl Common Repository subagent installed running on the VIOS. This model requires the management of the storage infrastructure (storage subsystems and SAN fabric) from the IBM Systems Director Server. See Figure 4-14 on page 117.

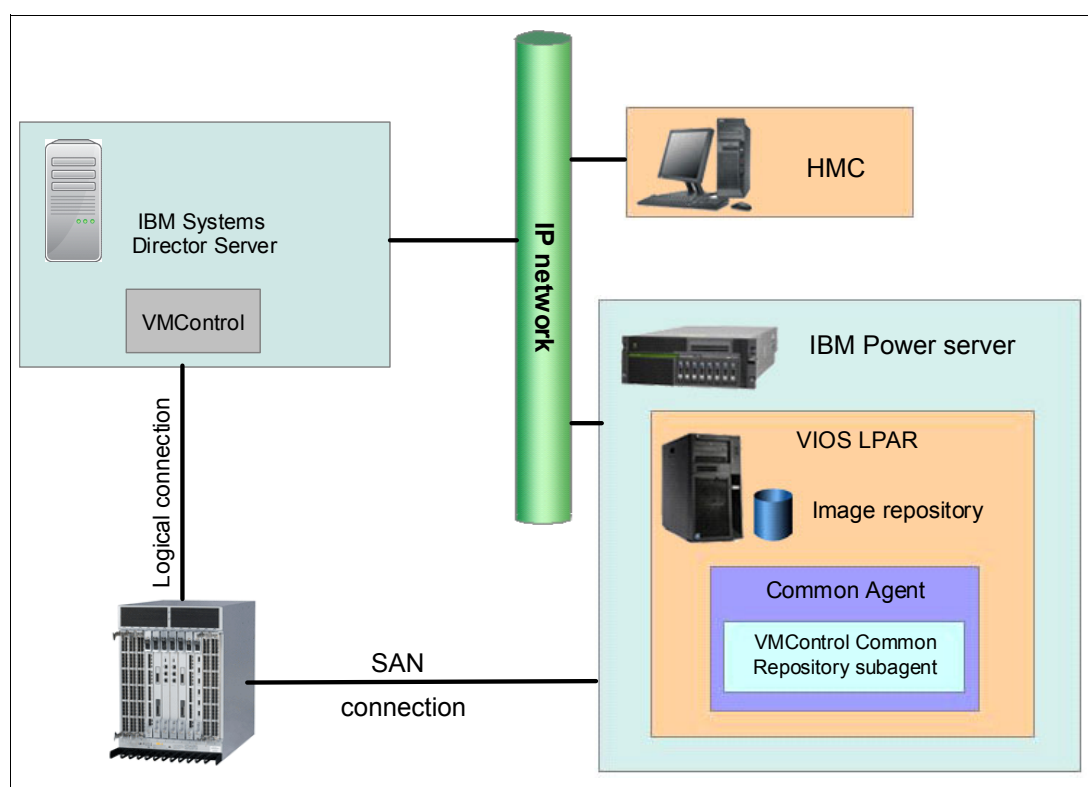


Figure 4-14 IBM Systems Director VMControl architecture

The basic VMControl Standard Edition virtual appliance tasks are different in an SCS environment:

- ▶ Capture operations use the storage management to provision a new LUN for the virtual appliance. A byte-level copy is made from the source operating system LPAR disk. This means only one disk and the OS must be shut down (LPAR powered off) during the capture.
- ▶ Deploy operations again will provision a new LUN for the new virtual server and create a byte-level copy from the virtual appliance “stored” on the Image Repository.
- ▶ Import operations use the raw disk image plus the metadata (OVF package) provisions a new LUN to store the virtual appliance.
- ▶ Both AIX and Linux on Power operating system images are supported.

NIM servers used by both Update Manager and VMControl

The base IBM Systems Director contains a built-in function known as the Update Manager. The Update Manager is used for documentation, compliance policy management, and installation of various types of updates, including AIX Technology Level (TL) and Service Pack (SP) updates. The updating of AIX requires that a single, designated NIM Server running the Common Agent be configured to provide support for this capability. This NIM server can also be an Image Repository for virtual appliances.

The Update Manager stores the AIX update filesets in the /export/um_lpp_source directory in a named subdirectory corresponding to the update, as shown in Example 4-14 on page 118.

Example 4-14 Listing of the Update Manager directories on a NIM server

```
ls -l /export/um_lpp_source
total 40
drwx--x--x  2 root    system      4096 Apr 21 10:54 7100.00.02.1041
drwx--x--x  2 root    system     16384 Aug 16 09:07 7100.00.03.1115
drwxr-xr-x  2 root    system       256 Jul 30 2010  lost+found
```

As explained, the directory for virtual appliances is different (/export/nim/appliances) and therefore the administrator must be cognizant of the total space requirements if using a NIM server.

The NIM resources created by the Update Manager will also differ from the resources used for capture and deployment, as shown in Example 4-15 (corresponding resources in bold).

Example 4-15 Listing of the Update Manager NIM resources

```
lsnim -t lpp_source
lpps61          resources      lpp_source
7100-00-03-1115_lppsrc  resources      lpp_source
7100-00-02-1041_lppsrc  resources      lpp_source
lpps71          resources      lpp_source
```

Both functions, Update Manager and VMControl Standard Edition, can happily coexist on the same AIX NIM server.

4.5.2 Post-install customization

This section provides post-installation and customization details.

Comparing post-install customization capabilities

Both CSM and Systems Director databases store information internally about the resources and virtual appliances that can be deployed to an LPAR, but they can use information stored in external customization files also.

As both products can utilize NIM for installation, they are able to transfer the necessary attributes to NIM, such as hostname, adapter data, script locations, and more. This is done when the node is defined into NIM and at the time of MAC address discovery.

VMControl defines the machine in NIM only for the time when a NIM operation is running. After that the node definition is removed from NIM.

Besides the install adapter information, additional installable filesets or fileset bundles can be allocated via NIM and pushed to the installed operating system. We suggest to create one or more NIM resource groups because this can be specified when we deploy a virtual appliance using VMControl GUI and the `smcli` command line as well.

Note: Because VMControl can handle only one entry for the NIM resource attribute, use resource grouping for all different kinds of installable objects and create more groups if different resources should be allocated to different insatiable operating system instances.

Adding install adapter-related TCPIP information

One of the final steps of restoring a system backup to a different OS instance is to reconfigure the hostname and IP addresses to avoid duplicate IP addresses on the network. We now show the different methods that could be used by CSM and VMControl.

► NIM

NIM machine definition contains the primary IP address and hostname, which is used to install and set up the operating system. When using only NIM to install AIX, you need also to enter the MAC address or hardware address information manually.

► CSM

Example 4-16 shows the node attributes for primary TCPIP information.

Example 4-16 Install adapter attributes in CSM

```
p55701p01(root)/> lsnode -l -a Hostname,Name,InstallAdapter* -n p55701p03
Hostname = p55701p03
InstallAdapterDuplex =
InstallAdapterGateway = 172.16.20.1
InstallAdapterHostname = p55701p03
InstallAdapterMacaddr = 523300003002
InstallAdapterName = en0
InstallAdapterNetmask = 255.255.252.0
InstallAdapterSpeed =
InstallAdapterType = ent
Name = p55701p03
```

The source of information depends on the way you define the node. Part of it can be an LPAR mapping file that is created with the `lshwinfo` command, or definition stanza files, which can define all node attributes.

The MAC address for the install adapter is added with the `getadapters` CSM command.

Before starting network installation, this information is transferred to NIM with the `csm2nimnodes` command.

► VMControl

When you deploy a virtual appliance, you can specify the IP addresses in the GUI manually or via customization files together with all other attributes, which is shown here.

Because only virtualized LPARs can be installed by VMControl at the time of writing this book, the adapter definition can be created and requested from the machines' virtualization infrastructure. In case of a preexisting LPAR the VIO server and the Power hypervisor already know about the adapter definitions and the virtual Ethernet mapping is already created. In case of a new LPAR, the mapping is defined based on the attributes you set in the GUI or in the deployment customization file.

Example 4-17 shows the adapter mapping attributes for deployment onto a new LPAR created automatically by VMControl.

Example 4-17 Virtual adapter mapping

```
p5570lp01(root)/> smcli lscustomization -a deploy_new -V 11703 -s 0x2d0a|grep
-p virtualnetworks
virtualnetworks
    Changeable Columns:
        Column Name*    CLI Attribute
        Virtual Networks on Host    hostVnet

p5570lp01(root)/export/nim/cust> grep virtualnetworks
cust_for_new_using_lp02_backup
virtualnetworks[Network 1]=hostVnet:ETHERNET0/1,
virtualnetworks[Network 2]=hostVnet:ETHERNET0/2,
```

IP addresses are specified at node deployment in the GUI or by a customization file, which is specified with the **smcli** deployment command. Example 4-18 shows the base TCPIP attributes for a node in one of our scenarios.

Example 4-18 TCPIP customization attributes

```
product.AIX1.com.ibm.ovf.vim.2.system.hostname=p5570lp02,
product.AIX1.com.ibm.ovf.vim.2.networkport.6.ip=172.16.20.202,
product.AIX1.com.ibm.ovf.vim.2.networkport.6.hostname=p5570lp02,
product.AIX1.com.ibm.ovf.vim.2.networkport.6.gateway=172.16.20.1,
product.AIX1.com.ibm.ovf.vim.2.networkport.6.netmask=255.255.252.0,
```

Name resolution

In case of CSM, you can set the name resolution by using file collections via adding a link to the `/etc/resolv.conf` file in the CFM's `/cfmroot` directory on the CSM management server. This is transferred to the nodes after initial installation and can be transferred with the **cfmupdatenode** command.

The other two methods are available on CSM and in VMControl as well:

- Post-install customization script

Create a customization script and allocate it as script NIM resource to the node, which can either copy or create a new `/etc/resolv.conf` file right after node installation.

- Allocate a `resolv_conf` NIM resource at the node installation.

This is automatic in VMControl because a `resolv_conf` NIM resource is created on the fly and used at deployment, and if there is DNS and domain information specified in the GUI, customization files or attributes of the **smcli** command, it will be used in the file referenced by the NIM resource.

Post-install scripts and other NIM resources and resource groups

One of the post installation scripts we use frequently is to customize the `/etc/netsvc.conf` file. You can create a script in the `/export/nim/cust` directory, for example, and define a script NIM resource for it, as shown in Example 4-19.

Example 4-19 NIM script definition

```
p5570lp01(root)/export/nim/cust> lsrim -l SDS_postinstall_script_1
SDS_postinstall_script_1:
    class      = resources
    type       = script
```

```
Rstate      = ready for use
prev_state  = unavailable for use
location    = /export/nim/cust/SDS_postinstall_script_1
alloc_count = 0
server      = master
```

NIM scripts and other predefined resources can be grouped into NIM resource groups. These resources and groups can be allocated to the node, when you run a `bos_install` operation in case of CSM. In VMControl you have two options to allocate NIM resources at node deployment:

- ▶ GUI

The last attribute on the panel that can be set at deployment time is for the NIM resources.

- ▶ Command line

Use the `product.AIX1.com.ibm.ovf.vim.2.nim.6.nim.Resource.1` deployment attribute either on the `smcli` command line with the `-A` flag or in the customization file specified with the `-F` flag. You can use the `smcli lscustomization` command to list the attributes that you can customize.

Adapter information

A secondary adapter configuration is handled as a post-install customization in case of a CSM cluster or in a standalone NIM-based management infrastructure.

Both CSM and NIM can handle adapter information and can use the stored information to set up additional noninstall adapters.

The information is stored in stanza files. CSM provides a node attribute, `AdapterStanzaFile`, to point to the file, while NIM provides an *adapters* class in which a resource can be defined specifying the locations of the files for the NIM clients.

After storing the information either in CSM or NIM and allocating to the node, it is used automatically at the next installation. CSM and NIM can force the nodes to reconfigure their adapters with the commands `updatenode -c -n "nodename"` and `nim -o cust -a adapter_def="adapter_def NIM resource" "nodename"`, respectively.

When the former CSM nodes are managed by Systems Director and inventory is collected for them, the adapter definitions are stored in the Systems Director database.

Note: It is advisable that all secondary adapter definitions that are stored in CSM are moved to NIM before the CSM cluster is deleted, if the NIM server is kept after removal of CSM. Use the `nimadapters` NIM command to process the CSM adapter stanza file and transfer definitions node by node to NIM.

For new installations or to reinstall the operating system the following methods are available to set up secondary adapters:

- ▶ NIM method

Assign the prepared adapter definition NIM resource to the node at deployment time using the method that was explained in 4.5.2, "Post-install customization" on page 118.

- ▶ VMControl method

List the additional adapter parameters such as IP address, IP label, and others in the customization file, on the command line with the `-A` flag or in the GUI. In this case,

VMControl creates a new adapter_def NIM resource and uses it while the deployment is running. After that the resource is removed from NIM.

4.5.3 Deployment in work

We used the following command to deploy a virtual appliance to an existing virtual server.

```
smcli deployva -a deploy_existing -V 11703 -s 10765 -F
/export/nim/cust/cust_for_p5570lp02_restore
```

The virtual appliance was created by capturing a running operating system managed by Systems Director.

Example 4-20 shows the customization file we used for deployment.

Example 4-20 VMControl customization file

```
p5570lp01(root)/export/nim/cust> cat cust_for_p5570lp02_restore
virtualnetworks[Network 1]=hostVnet:ETHERNET0/1,
virtualnetworks[Network 2]=hostVnet:ETHERNET0/2,
product.AIX1.com.ibm.ovf.vim.2.nim.6.nim.Resource.1=SDS_postinstall_script_1,
product.AIX1.com.ibm.ovf.vim.2.system.hostname=p5570lp01,
product.AIX1.com.ibm.ovf.vim.2.networkport.6.ip=172.16.20.202,
product.AIX1.com.ibm.ovf.vim.2.networkport.6.hostname=p5570lp02,
product.AIX1.com.ibm.ovf.vim.2.networkport.6.gateway=172.16.20.1,
product.AIX1.com.ibm.ovf.vim.2.networkport.6.netmask=255.255.252.0,
product.AIX1.com.ibm.ovf.vim.2.networkport.7.ip=10.10.6.2,
product.AIX1.com.ibm.ovf.vim.2.networkport.7.hostname=clprivlp02,
product.AIX1.com.ibm.ovf.vim.2.networkport.7.gateway=10.10.6.2,
product.AIX1.com.ibm.ovf.vim.2.networkport.7.netmask=255.255.255.0
```

Example 4-21 shows the log of the **netboot** command, which is started by the NIM server based on the order from VMControl.

Example 4-21 Netboot process

```
p5570lp01(root)/var/ibm/sysmgt/dsm/log> cat dnetboot.p5570lp02.log.272
Output log for dnetboot is being written to
/var/ibm/sysmgt/dsm/log//dnetboot.p5570lp02.log.272.

-----
dnetboot: Logging started Thu Sep 29 18:16:54 EDT 2011.
-----

dnetboot Status: Invoking /opt/ibm/sysmgt/dsm/dsmbin/lpar_netboot p5570lp02
18:16:54 dnetboot Status: Invoking /opt/ibm/sysmgt/dsm/dsmbin/lpar_netboot
p5570lp02
dnetboot Status: Invoking /opt/ibm/sysmgt/dsm/dsmbin/lpar_netboot -i -t ent -D -S
172.16.20.201 -G 172.16.20.201 -C 172.16.20.202 -m 523300002002 -s auto -d auto -F
/etc/ibm/sysmgt/dsm/config/passwd_hscroot_172_16_20_107 -j hmc -J 172.16.20.107 2
107CD9E 9117-570

-----
dnetboot: Logging stopped Thu Sep 29 18:16:54 EDT 2011.
-----

# Connected
# Checking for OF prompt.
# Timeout waiting for OF prompt; rebooting.
# Checking for power off.
```



```
# Client IP address is 172.16.20.202.
# Server IP address is 172.16.20.201.
# Gateway IP address is 172.16.20.201.
# Getting adapter location codes.
# /vdevice/l-lan@30000002 ping successful.
# Network booting install adapter.
# bootp sent over network.
# Network boot proceeding, lpar_netboot is exiting.
# Finished.
18:18:15 dnetboot Status: Was successful network booting node p55701p02.
```

VMControl creates an adapter definition NIM resource automatically based on the customization attributes. Example 4-22 shows the details of this resource.

Example 4-22 NIM secondary adapter definition

```
p55701p01(root)/> lsrim -l nimrf-0000000000000003-adapter_def
nimrf-0000000000000003-adapter_def:
  class      = resources
  type       = adapter_def
  comments   = <?xml version="1.0" encoding="UTF-8"?><nimrf><created>Thu Sep 29
17:09:36 EDT 2011</created><machine>p55701p02</machine></nimrf>
  Rstate     = ready for use
  prev_state = unavailable for use
  location   = /export/nim/adapter_def/nimrf-0000000000000003-adapter_def
  alloc_count = 1
  server     = master
p55701p01(root)/> cat
/export/nim/adapter_def/nimrf-0000000000000003-adapter_def/p55701p02.adapters
p55701p02:
  hostname=p55701p02
  machine_type=secondary
  network_type=en
  hostaddr=172.16.20.202
  location=U9117.570.107CD9E-V2-C4-T1
  secondary_hostname=c1privlp02
  netaddr=10.10.6.2
  subnet_mask=255.255.255.0
  cable_type=N/A
  media_speed=Auto_Auto_Duplex
  route="10.10.6.2::255.255.255.0::10.10.6.2"
```

Example 4-23 shows which NIM resources are allocated to a NIM machine at installation time.

Example 4-23 NIM machine definition at installation time

```
p55701p01(root)/> lsrim -l p55701p02
p55701p02:
  class      = machines
  type       = standalone
  comments   = <?xml version="1.0" encoding="UTF-8"?><nimrf><created>Thu Sep
29 17:04:41 EDT 2011</created></nimrf>
  adapter_def = nimrf-0000000000000003-adapter_def
  connect     = nimsh
  platform    = chrp
```

```

netboot_kernel = mp
ifl            = itsonet p5570lp02 523300002002
net_settings1  = auto auto
cable_type1    = N/A
mgmt_profile1  = H172_16_20_107 2 cec-570-107CD9E
Cstate         = Base Operating System installation is being performed
prev_state     = BOS installation has been enabled
Mstate        = in the process of booting
info           = BOS install 38% complete : 41% of mksysb data restored.
boot           = boot
bosinst_data   = nimrf-0000000000000003-bosinst_data
image_data     = nimrf-0000000000000003-image_data
mksysb         = appliance-1_image-1
nim_script     = nim_script
resolv_conf    = nimrf-0000000000000003-resolv_conf
script         = SDS_postinstall_script_1
script         = nimrf-0000000000000001-script
script         = vmc_config
spot          = nimrf-0000000000000002-spot
cpuid          = 00C7CD9E4C00
control        = master
Cstate_result  = success

```

All directories are NFS exported by VMControl, which contains NIM resources necessary for network installation.

After the successful deployment, only the NIM resources shown in Example 4-24 remain defined.

Example 4-24 Remaining NIM resources after deployment

SDS_postinstall_script_1	resources	script
appliance-1_image-1	resources	mksysb
vmc_config	resources	script
H172_16_20_107	management	hmc
cec-570-107CD9E	management	cec
nimrf-0000000000000002-spot	resources	spot

4.6 Agents

As stated in previous sections, IBM Systems Director communicates with and controls the various systems in its group via agents. These agents are responsible for communicating with IBM Systems Director. With these agents clients can keep their management systems informed of their current situation. Two sets of agents can be installed: the Common Agent and the Platform Agent.

The Common Agent provides the following functions (see Figure 4-15 on page 125):

- ▶ System discovery
- ▶ Comprehensive platform and operating system inventory collection
- ▶ Monitoring health and status
- ▶ Managing alerts

- ▶ Remotely deploy and install Common Agent
- ▶ Performing remote access, including transferring files
- ▶ Performing power management functions
- ▶ Event support
- ▶ Monitor processes and resources, and set critical thresholds for send notifications when triggered
- ▶ Managing operating system resources and processes
- ▶ Managing updates

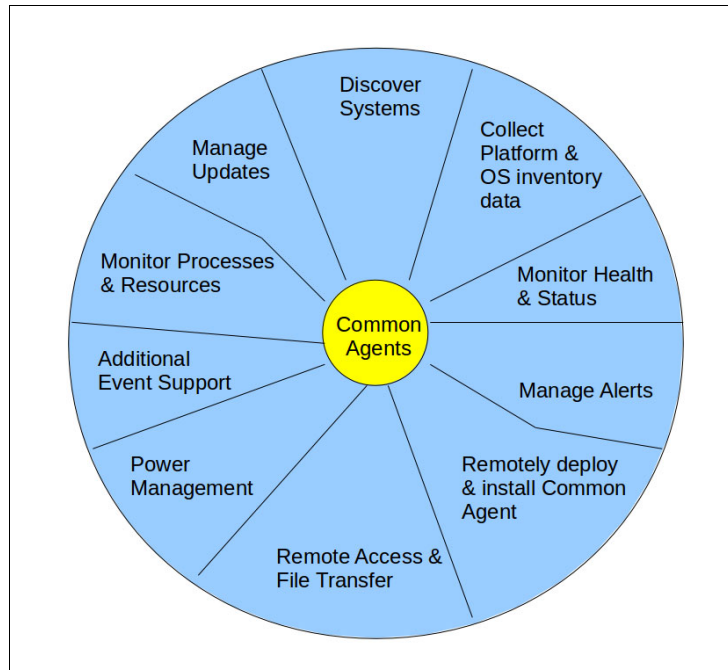


Figure 4-15 Common agents

Note: Throughout the IBM Systems Director documentation, the term Common Agent (with both words capitalized) refers to the IBM Systems Director Common Agent, which includes subagents that provide specific management capabilities for IBM Systems Director. IBM Systems Director can also discover and perform limited management on other common agents that use the common agent services (CAS) architecture. When referring to these common agents generically, lowercase text is used.

The Platform Agent provides a subset of Common Agent functions used to facilitate communication between the managed system and IBM Systems Director. It is an option for environments that require smaller footprints without having to sacrifice much of systems managability.

The Platform Agent provides the following functions (see Figure 4-16 on page 126):

- ▶ Discover systems
- ▶ Collect limited platform inventory data
- ▶ Monitor health and status
- ▶ Manage alerts

- ▶ Remotely deploy and install Common Agent
- ▶ Perform limited remote access
- ▶ Perform limited restart capabilities on the system

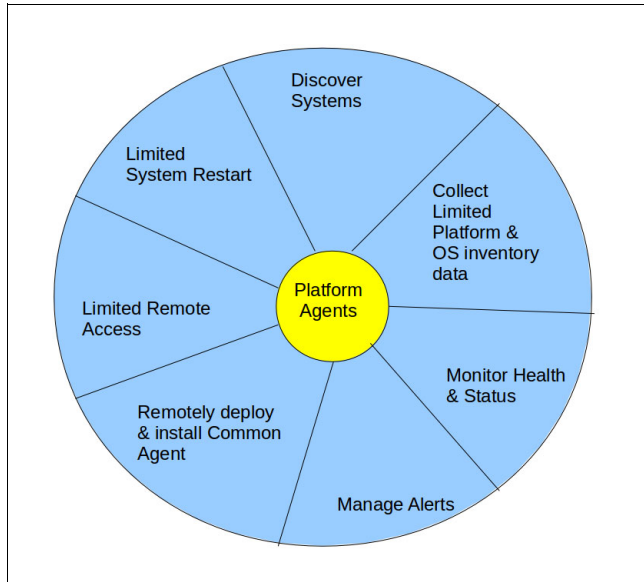


Figure 4-16 Platform agents

You cannot use just the Platform Agents in AIX. Both the Common Agents and Platform Agents will be part of your AIX installation (see Table 4-4). A default installation of the AIX operating system will install the Common Agents. By default, the Common Agent daemons are started at startup.

Table 4-4 AIX versions supported by IBM Systems Director

Operating System
<p>IBM AIX Version 5.3:</p> <ul style="list-style-type: none"> ▶ AIX 53 TL06 SP09 or later SP levels ▶ AIX 53 TL07 SP06 or later SP levels ▶ AIX 53 TL08 SP04 or later SP levels ▶ AIX 53 TL09 or later TL levels <p>Note:</p> <ol style="list-style-type: none"> 1. Agentless support includes only discovery, limited access, a limited subset of the Agent Installation Wizard task, and the ability to launch the AIX Management Console for IBM AIX Version 5.3 agents. 2. For IBM Systems Director Serve or Common Agent support provided by IBM AIX Version 5.3, newer TLs and Server Packs will be supported on the day of their general availability (GA). 3. Common Agent is installed with AIX; it is started by default.

Operating System
<p>IBM AIX Version 6.1:</p> <ul style="list-style-type: none"> ▶ AIX 61 TL00 SP07 or later SP levels ▶ AIX 61 TL01 SP03 or later SP levels ▶ AIX 61 TL02 or later TL levels <p>Note:</p> <ol style="list-style-type: none"> 1. Agentless support includes only discovery, limited remote access, a limited subset of the Agent Installation Wizard task, and the ability to launch the AIX Management Console for IBM AIX Version 6.1 agents. 2. For IBM Systems Director Server or Common Agent support provided by IBM AIX Version 6.1, newer TLs and Service Packs will be supported on the day of their general availability (GA). 3. Common Agent is installed with AIX; it is started by default.
<p>IBM AIX Version 7.1</p> <p>Note:</p> <ol style="list-style-type: none"> 1. Agentless support includes only discovery, limited remote access, a limited subset of the Agent Installation Wizard task, and the ability to launch the AIX Management Console for AIX Version 7.1 agents. 2. For IBM Systems Director Server or Common Agent support provided by AIX Version 7.1, newer TLs and Service Packs will be supported on the day of their general availability (GA). 3. Common agent is installed with AIX; it is started by default.

4.6.1 Daemons, filesets, and protocols

IBM Common Agents and Platform Agents running on AIX have modest installation prerequisites for minimum processor speed and disk space requirements, as shown in Table 4-5.

Table 4-5 Servers running AIX minimum hardware requirements

Requirements	Common Agent
Processor speed	Power5, Power6 or Power7
Memory (RAM)	512 MB (minimum)

Table 4-6 provides information concerning disk storage requirements for installing Common Agent on your system. Table 4-7 details the additional space necessary on the default AIX file systems, and also shows the additional space requirements depending on your chosen installation method.

Table 4-6 Disk storage requirements for installing Common Agent (including Platform Agent)

File System	Disk space required
/	<5 MB
/usr	250 MB
/var	300 MB
/tmp	200 MB
/opt	140 MB

Note: The /tmp space is only required for installation. Disk space is freed after installation completes.

Table 4-7 Additional disk space necessary for installation based on installation method

File System	Disk space required
Manual installation using SysDir6_2_1_Common_Agent_Installp_AIX_VIOS.tar.gz	
Web download file: SysDir6_2_1_Common_Agent_Installp_AIX_VIOS.tar.gz (Location of your choice)	Size specified on downloads website
Space needed to extract web downloaded file: SysDir6_2_1_Common_Agent_Installp_AIX_VIOS.tar.gz (Location of your choice)	290 MB
Manual installation using SysDir6_2_1_Common_Agent_AIX.tar.gz	
Web download file SysDir6_2_1_Common_Agent_AIX.tar.gz	Size specified on downloads website
Space needed to extract web downloaded file SysDir6_2_1_Common_Agent_AIX.tar.gz	290 MB
Space needed to extract dir6.2.1_commonagent_aix into /tmp/DirectorAgentselfextract.<identifier>	290 MB
Agent deployment with Agent Installation Wizard using SysDir6_2_1_Common_Agent_AIX.tar.gz	
Server system: Web downloaded file (SysDir6_2_1_Common_Agent_AIX.tar.gz)	Size specified on downloads website
Server system: Space needed to extract web downloaded file (SysDir6_2_1_Common_Agent_AIX.tar.gz)	290 MB
Target system: Depoly dir6.2.1_commonagent_aix file	290 MB
Target system: Space needed to extract dir6.2.1_commonagent_aix into /tmp/DirectorAgentselfextract.<identifier>	290 MB

4.6.2 Port conflicts - ITM?

IBM Systems Director communicates with the Common Agents installed on the client systems using the TCP protocol. These ports are labeled port, jport and nport and by default, are set to ports 9510, 9514, and 9515, respectively, on the client system. Before installation of the Common Agents make sure that these ports are not currently in use by other applications. You can use the **netstat** command shown in Example 4-25 to check whether these ports are in use by your AIX system.

Example 4-25 Search your AIX system for ports in use

p5570lp01(root)/> netstat -an grep LISTEN egrep "951(0 4 5)"					
tcp	0	0	*.9510	*.*	LISTEN
tcp4	0	0	127.0.0.1.9514	*.*	LISTEN
tcp4	0	0	127.0.0.1.9515	*.*	LISTEN

In Example 4-25, the ports are currently in use and the Common Agent on this system must be reconfigured to use other ports. This can be accomplished with the **configure.sh** command provided by the Common Agent installation:

```
configure.sh -unmanaged -port <newport> -jport <newport> -nport <newport>
-force
```

Where <newport> is the unused port you wish to change. Example 4-26 illustrates a change from the default 9510, 9514, and 9515 ports to the unused 9520, 9524, and 9525 ports.

Example 4-26 Change default ports

```
p5570lp01(root)/var/opt/tivoli/ep/runtime/agent/toolkit/bin>./configure.sh
-unmanaged -port 9520 -jport 9524 -nport 95520 -jport 9524 -nport 9525 -force
<
```

```
BTC7861I Checking if the common agent is already registered...done
BTC7862E The common agent is already configured (it contains certificates).
```

```
BTC7872I Checking the common agent state...the common agent is running
BTC7875I Stopping the common agent...done
```

```
BTC7876I Reporting 'uninstallation' status...done
```

Resolved configuration:

The common agent is going to be reconfigured - force option specified

The common agent connection information:

```
name = localhost
port = 9520
jport = 9524
nport = 9525
wport = disabled
wsport = disabled
```

The agent manager information:

Common agent not managed by the agent manager

The common agent service information:

```
noinstall = false
nostart = false
```

```
BTC7878I Validating the resolved configuration of the common agent...done
```

```
BTC7880I Configuring the common agent...done
```

```
BTC7881I Installing the common agent service...done
```

```
BTC7882I Starting the common agent...done
```

Note: In Example 4-26 on page 129, we recognized that our Common Agent cannot connect to the Director server because the Common Agent is already installed on the client system. In this case, we should use the shortest path to the script `/var/opt/tivoli/ep/runtime/agent/toolkit/bin/configure.sh` directory instead of this path:
`/usr/lpp/cas.agent/inst_root/var/opt/tivoli/ep/runtime/agent/toolkit/bin/configure.sh`.

The **configure.sh -unmanaged** function can be useful when we want to move a client to a different Director server (AgentManager).



Special topics

This chapter discusses and provides guidance about:

- ▶ “SDMC”
- ▶ “High Performance Computing”
- ▶ “High availability IBM Systems Director management server”
- ▶ “Backup and restore of the management server”

5.1 SDMC

This section provides a brief description to IBM Systems Director Management Console. It touches upon the advantages of using it over HMC and IVM. It also throws some light on high-availability options, and it covers the command line interface of the IBM Systems Director Management Console.

5.1.1 Overview of SDMC

The IBM Systems Director Management Console (SDMC) is the successor to the Hardware Management Console (HMC) and the Integrated Virtualization Manager (IVM).

The HMC is designed for management of Power Systems from entry level servers to high-end systems. Management of entry systems can also be done with the IVM, a function of the Virtual I/O server.

An HMC can administer up to 256 LPARs, whereas IVM is confined to the system the Virtual I/O Server is installed on. Moreover, an HMC does not allow the management of Power processor-based blades.

Figure 5-1 on page 133 shows one HMC (as a machine) managing some Power servers and an IVM (as a box) managing the whole blade system or only one blade. IVM is installed normally on one or more blade servers, but it can be installed on entry level Power servers also. IVM is packaged together with VIO Server as part of PowerVM.

Also see Figure 1-1 in the following publication:

<http://www.redbooks.ibm.com/abstracts/sg247860.html>

Figure 5-1 on page 133 shows the Management Consoles HMC and IVM managing Power servers and the Power blades, respectively.

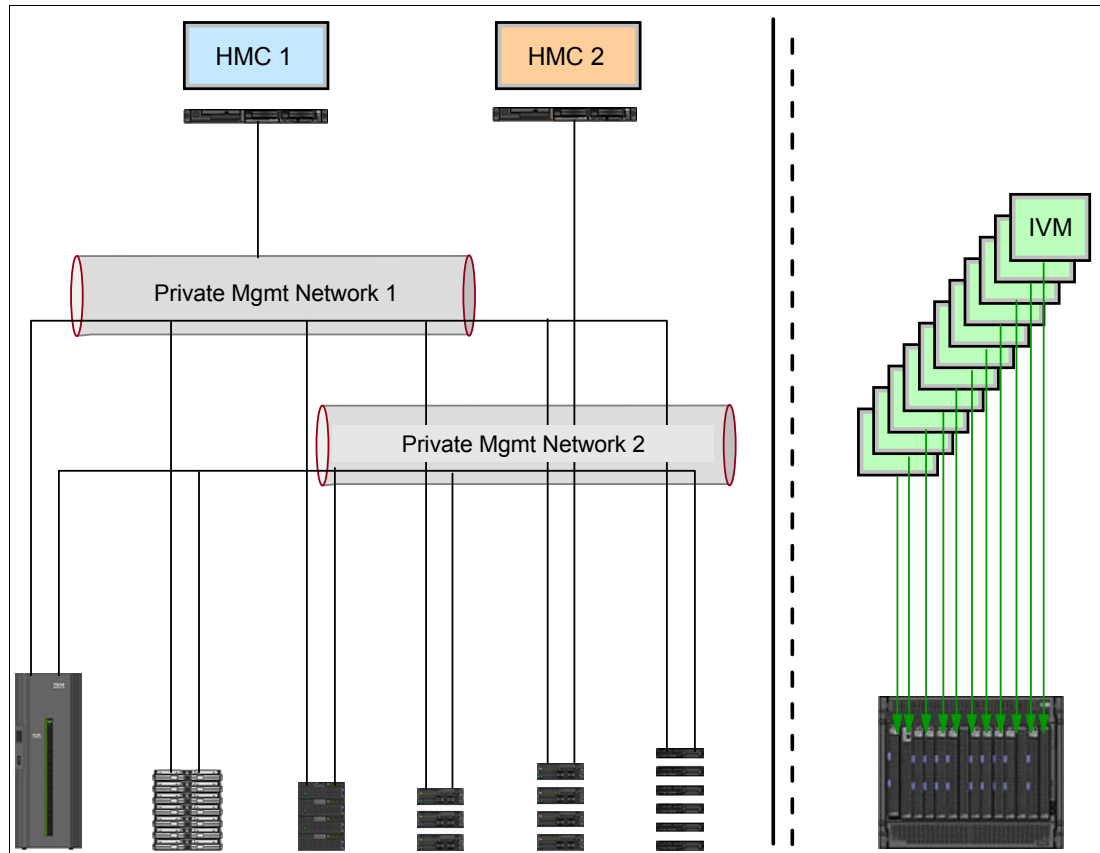


Figure 5-1 HMC and IVM management for the Power servers and the Power blades

With the introduction of SDMC, the need for two different management servers is withdrawn. It is capable of managing both the Power servers and the Power processor-based blades, as shown in Figure 5-2 on page 134.




Figure 5-2 SDMC managing both Power servers and Power processor-based blades

Note: SDMC is designed to Manage only Power Systems including Power blades. However, it supports only systems POWER6 and above.

Some of the key enhancements of SDMC in the area of virtualization management include:

- ▶ It provides a simplistic IVM-like user interface for the virtualization features, such as creating an LPAR, modifying its properties, deleting an LPAR, and so on.
- ▶ The views of Virtual Server properties and dynamic logical partitioning are combined to present a single view from where all the Virtual Server operations can be performed.
- ▶ Even in the stopped state, SDMC provides the ability to modify the resource assignment.
- ▶ SDMC can now manage virtual slots automatically, leading to enhanced VIOS management.

5.1.2 Parallel management with the HMC and the SDMC

Figure 5-3 on page 135 describes a parallel management of HMC or SDMC on a single POWER6 or POWER7 frame. It also shows a management configuration where the IBM Systems Director, with the advanced management plug-ins installed, leverages the HMC's management interface to the frame to facilitate use of the plug-ins.

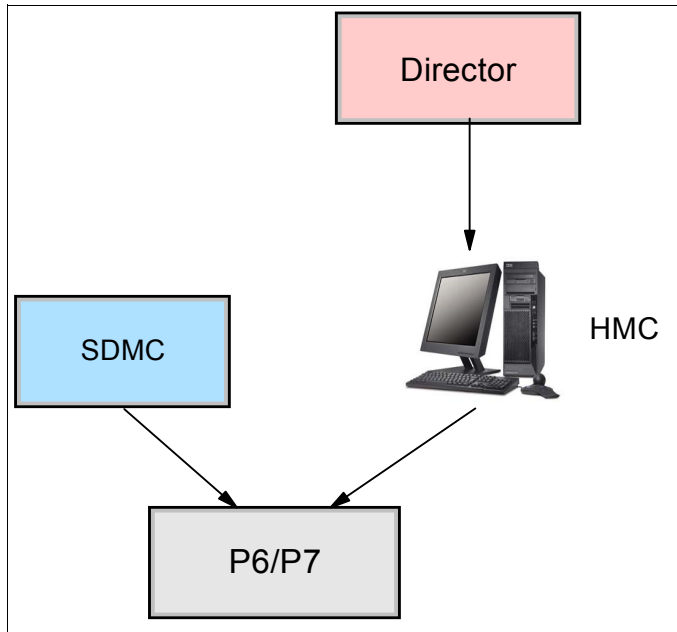


Figure 5-3 IBM Systems Director along with HMC and SDMC

As the HMC transitions out, the IBM Systems Director will be able to manage POWER6 and POWER7 either directly or hierarchically through SDMC.

Note: Here are the supported levels for the above configuration:

- ▶ The level of HMC is 7.3.5 or higher
- ▶ The level of IBM Systems Director is 6.2.1.2 or higher

5.1.3 High availability of SDMC

Users familiar with the redundant HMC setup in their Power Systems environment can use the same setup in SDMC, because multiple SDMCs can connect to and actively manage a single managed server.

Users can also implement SDMC High Availability, which provides active/passive failover capability, with one active SDMC and one passive SDMC on standby to take over in case of failure.

Both setups are shown in Figure 5-4 on page 136.

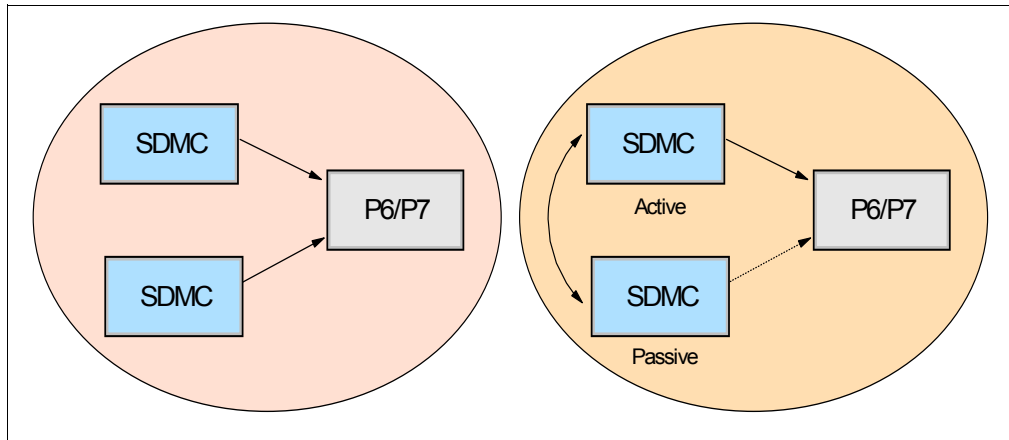


Figure 5-4 High availability of the SDMC

5.1.4 Command Line for SDMC as compared to HMC

The command-line interface has been kept almost the same in SDMC. Hence, it might not require changes to existing scripts that use the Hardware Management Console.

Note: For more information about SDMC installation and other features, refer to the IBM Redbooks publication provided in:

<http://www.redbooks.ibm.com/abstracts/sg247860.html>

5.2 High Performance Computing

IBM Cluster Systems Management provides High Performance Computing solution packs to assist with software installation, administration, and customization of the cluster nodes. These solution packs, delivered via additional AIX filesets, provide support for the following HPC stack applications:

- ▶ IBM General Parallel File System (GPFS)
- ▶ IBM Engineering and Scientific Subroutine Library (ESSL)
- ▶ IBM Parallel ESSL
- ▶ IBM Parallel Environment
- ▶ IBM Tivoli Workload Scheduler LoadLeveler®

The integration of the HPC stack applications in CSM was not implemented in IBM Systems Director. For the technical computing and HPC AIX customer, the recommended software product for systems management is Extreme Cloud Administration Toolkit (xCAT 2). xCAT includes single-point-of-control management functions including RMC for resource monitoring, the IBM Energy Management plug-in for Power and the HPC Hardware Server for FSP communication.

Note: For more information about xCAT, see the product website:

<http://xcat.sourceforge.net>

and the IBM Redbooks publication *Configuring and Managing AIX Clusters Using xCAT 2*, SG24-7766.

5.3 High availability IBM Systems Director management server

Building an IT infrastructure that addresses today's challenges and tomorrow's opportunities requires high availability and quality of existing services. Management of highly available servers requires that your management server also be highly available and fault-tolerant to ensure alerts are not missed because the IBM Systems Director server is unavailable.

A highly available IBM Systems Director environment is to deploy redundant IBM Systems Director servers, and have the resources managed by both of them. On the surface, this sounds like a simple solution. However, based on the types of resources being managed, and the tolerance of duplicate notification, this may become significantly more complicated.

To manage devices, simply discover, request access to, and collect inventory from all of the managed resources from each IBM Systems Director server, and then create an Event Automation Plan on each IBM Systems Director server to process the events. Events will flow from the managed resources to each IBM Systems Director server, where they will be processed based on the Event Automation Plan.

Figure 5-5 illustrates a simple redundant server configuration.




Figure 5-5 Simple redundant server configuration

The managed resources will be configured to send alerts to both IBM Systems Director servers once they have been discovered and authenticated using Request Access.

Managing common agents by multiple IBM Systems Director servers is much more complicated than managing other resource types. This is because IBM Systems Director uses the Common Agent Services (CAS) architecture, which provides a shared infrastructure for managing systems. In this architecture, one or more resource managers (for example, IBM Systems Director server) use an agent manager to communicate with the common agents that are installed on managed resources.

IBM Systems Director uses only one active agent manager at a time to communicate with common agents. Each common agent can use exactly one agent manager. However, the agent manager that is embedded with IBM Systems Director server can be used by any number of instances of IBM Systems Director server. This allows multiple IBM Systems Director servers to share a single agent manager for security and credential management, yet access the common agents directly for other types of management.

5.3.1 Additional considerations

When considering duplicate automation, there are some actions that tend to work better than others. For example, notification via email, or forwarding events to an Enterprise Systems Management application such as Tivoli would be types of actions where duplication would be successful. Actions such as start a program on the system that generated the event or start a task on a system that generated the event do not work well with duplication automation, since the same program or task would potentially be executed multiple times on the same system. If using these actions, the program or task being executed will need to perform a check to handle this redundancy.

5.3.2 Highly available IBM Systems Director management servers

The following are a few reasons for the requirement to have highly available management servers:

- ▶ When managing a HA/FT environment, your management server must be as highly available or more so than the servers being managed.
- ▶ Most MEP types use “fire and forget” alerting. Common Agent (CA) has limited capability to “store and forward”.
- ▶ Highly available Director management servers ensure critical alerts are not missed.
- ▶ Install the Director management server in a virtual server, and use virtualization tools to ensure availability.
- ▶ Use cluster technologies to ensure the Director management server is running.
- ▶ Active/Passive servers ensure that the Director management servers are always operational. Can be configured using native product capabilities.

5.3.3 Redundant IBM System Director management server options

The following are a few options for redundant management servers.

Redundant management servers with alert duplication (CAS)

This list provides reasons why this option can be of consideration:

- ▶ Multiple IBM Systems Director servers manage the same set of endpoints.

- ▶ Both Director servers process alerts.
- ▶ Useful when sending alerts to ESM that has the ability to detect and delete duplicates.

Redundant management servers without alert duplication (no CAS)

The following list provides reasons why this option can be useful:

- ▶ Multiple IBM Systems Director management servers manage the same set of endpoints.
- ▶ Only one IBM Systems Director management server processes alerts.
 - Secondary IBM Systems Director management server monitors the primary management server.
 - Secondary IBM Systems Director management server only processes alerts if the “primary” management server is unavailable.
- ▶ Useful for customers with an aversion to duplicate notification.

5.3.4 Conclusion

IBM Systems Director can be deployed in a highly available configuration ensuring hardware critical alerts are always processed. Although IBM Systems Director does not support typical clustering technologies, it can be deployed in a redundant server configuration, using customized scripts and Event Automation Plans to ensure critical alerts are always processed.

5.4 Backup and restore of the management server

This section provides instruction for the backup and restore of CSM and IBM Systems Director management server.

To protect the CSM and IBM Systems Director configuration data, both products are provided with commands to back up and restore in the event of corruption and disaster recovery.

Backup is also required to replace a defective management server, to move a management server to a different machine, to upgrade a management server's operating system, and finally to migrate management software.

Note: For both CSM and IBM Systems Director, there is no requirement to take client backups.

5.4.1 IBM Cluster Systems Management (CSM) backup and restore

Before starting the transformation of the CSM cluster into Systems Director, a good precaution method is to take a backup of the CSM management server. In case you are facing issues during the transformation and you need to roll back, then this backup is a must.

csmbackup

The **csmbackup** command copies persistent CSM data from the management server and stores the data in the directory specified. This command almost replicates the CSM state of the management server. The command backs up the nodes, node groups, conditions and responses, and condition and response associations. Customization scripts and DCEM scripts are also copied.

By default, the backup files are placed in the `/var/opt/csm/csmdata` directory of the management server, but to ensure that backup files are not lost in case of hard disk failure, copy the backup files to external disk, tape, or an nfs mounted file system.

Note: The command does not back up all data on the management server because some information needs to be copied to a backup storage device manually. However, additional files can be stored by the **csmbackup** command using the **-f** flag.

The **csmbackup -r** command can be run with the following parameters:

- ▶ DeviceHwCtrl
- ▶ DeviceGroup
- ▶ Sensor
- ▶ EventResponse
- ▶ Condition
- ▶ Association
- ▶ ManagedNode
- ▶ NodeGroup
- ▶ DmsCtrl
- ▶ User file name

These parameters back up data associated with its resource class. You can create a condition that monitors when one of the resource classes changes by monitoring 'ConfigChanged >= 0' on a specified resource class, then set up a response that runs the following script:

```
csmbackup -r Resource_Class> -d directory
```

To back up information to the default backup directory `/var/opt/csm/csmdata`, enter:

```
csmbackup
```

To back up all node groups to a mounted file system, enter (see Example 5-1):

```
csmbackup -d /work/csmbackup
```

Example 5-1 The csmbackup command

```
p55701p01(root)/> csmbackup -d /work/csmbackup
Backing up CSM data, this could take a few minutes.
CSM backup data completed.
Please read /work/csmbackup/tasks.csmbak for CSM backup/restore information
p55701p01(root)/>
```

After the backup is done you should archive this directory and keep it in a safe place in case you need to roll back. Another backup method would be to take a full system backup using the **smit mksysb** command. Of course, this method is much more time consuming.

csmrestore

In an unpleasant event when you have to roll back the configuration of your CSM cluster, you can restore from the backup using the command **csmrestore**.

The **csmrestore** command restores the files that were copied by the **csmbackup** command into the specified directory or the default directory `/var/opt/csm/csmdata`. The CSM management

server must be completely installed and configured by the administrator in advance; the **csmrestore** command does not install any CSM code on the management server.

The **csmrestore** command only restores the nodes, node groups, devices, device groups, conditions and responses, and condition and response associations on the management server. Customization scripts and DCEM scripts are also restored.

To restore data using the **csmrestore** command, CSM must be installed on the replacement or upgraded management server in the same manner that CSM was installed on the management server that ran the **csmbackup** command.

To restore information from the default backup directory `/var/opt/csm/csmdata`, enter:

```
csmrestore
```

To restore data from any other directory, use the **-d** flag (see Example 5-2):

```
csmrestore -d /work/csmbak
```

Example 5-2 The csmrestore command

```
[p55701p01]#csmrestore -d /work/csmbak
```

```
Restoring node definitions backed up on AIX 6.1.6 CSM version 1.7.1.6...
```

```
Remove all existing nodes to restore saved nodes? (Y/N): y
```

```
Defining CSM Nodes:
```

```
p61p02: Changing Mode from: "Managed" to: "PreManaged".
```

```
p61p03: Changing Mode from: "Managed" to: "PreManaged".
```

```
2 nodes have been defined successfully.
```

```
Restoring nodegroups backed up on AIX 6.1.6 CSM version 1.7.1.6...
```

```
Remove all existing nodegroups to restore saved nodegroups? (Y/N): y
```

```
Restoring hardware device definitions backed up on AIX 6.1.6 CSM
```

```
Remove all existing hardware devices to restore saved hardware devices? (Y/N): y
```

```
Restoring hardware device groups backed up on AIX 6.1.6 CSM version 1.7.1.6...
```

```
Remove all existing hardware device groups to restore saved groups? (Y/N): y
```

```
Restoring condition information backed up on AIX 6.1.6 CSM version 1.7.1.6...
```

```
Restoring Event Response information backed up on AIX 6.1.6 CSM
```

```
Restoring sensor information backed up on AIX 6.1.6 CSM version 1.7.1.6...
```

```
Restoring condition/response association backed up on AIX 6.1.6 CSM
```

```
Restoring least-privilege (LP) information backed up on AIX 6.1.6 CSM
```

```
Restoring CSM user customization scripts...
```

```
Restoring csmconfig information backed up on AIX 6.1.6 CSM version 1.7.1.6...
```

```
Restore of CSM data completed.
```

```
There are several files in the /work/csmbak directory that contain  
information on the state of the system when it was backed up.
```

```
You may find it useful to look at these files. They include:
```

```
    /work/csmbak/config.csmbak
```

```
    /work/csmbak/tasks.csmbak
```

```
[p55701p01]#
```

Note: When using the **csmbackup** and **csmrestore** commands across different CSM levels, use the **csmrestore** command with caution; overwriting resources can cause unexpected problems. The **csmrestore** command prompts you before removing resources; when prompted, press the n key to not remove those resources.

For further details about backup and restore steps, refer to the Administration Guide of the IBM Cluster Systems Management for AIX and Linux at:

<http://publib.boulder.ibm.com/epubs/pdf/a2313435.pdf>

5.4.2 IBM Systems Director backup and restore

Backup of IBM Systems director's file system and database is required to shield the data from a disaster. To simplify the backup and restore process, IBM Systems Director offers the **smsave** and **smrestore** commands.

Consider the following criteria during backup and restoration:

- ▶ System consideration
 - Ensure that you have enough free storage space to save your data with the **smsave** command.
 - Restoration of data is possible only on the same type of hardware as the previous installation (data backup from IBM Power systems cannot be restored on System x server).
 - Restore the data in the same version of IBM Systems Director.
 - Restore the data in the same version of the operating system.
- ▶ Database consideration
 - If the database server is remote, the disk space required for backup requirements is split across the management server and the database server. Database storage uses the majority of the **smsave** data. You do not need much space on the management server but on the database server.
 - Restore the database on the same version of the database application as the previous installation (backups cannot be moved from one database type to another database type or version).
 - If you have a remote database, the backup command produces two data sets: One at the location of the remote database server, and one on the IBM Systems Director management server system. The data sets are mated sets. You must maintain and restore these data sets together.
 - If your database is on a remote server, make a directory on that database server with permissions such that the IBM Systems Director has write access.
 - Enable password file authentication for the database system administrator user ID that is used for backup/restore.
 - Ensure that the database instance owner is in the same primary group as the IBM Systems Director user who performs the backup and restore operations. This enables the database server to read and write the database backup image to the backup directory.
 - If you chose to create the IBM DB2 database yourself, verify how you configured the rollforward option. If this option is set to require user approval, running **smrestore** will cause errors. Correct this issue now to avoid potential problems later with restoring data.

5.4.3 smsave

The **smsave** command saves a backup image of Systems Director persistent data, including file system data and databases. You must stop the IBM Systems Director Server before running the command.

Here are the options used with the **smsave** command:

- ▶ **targetDir**
Optional location to save backup files.
If not specified, the location will be `/Director/backup/`. Refer to Example 5-3.
- ▶ **dbTargetDir**
Optional location to save database backup files. Required if database is remote. If not specified, the location will be defaulted based on the database type.
- ▶ **dbUserName**
Optional user name for database access. Must be used in conjunction with **dbUserPwd**.
- ▶ **dbUserPwd**
Optional password for database access. Must be used in conjunction with **dbUserName**.

Note: For an AIX system, if **-dbUserName** and **-dbUserPwd** are not specified, and you have an IBM DB2 or Oracle database, you are prompted to provide a user name and password. For an Apache Derby database, the **-dbUserName** and **-dbUserPwd** options are optional.

Example 5-3 IBM Systems Director backup command

```
p55701p01(root)/> smsave -targetDir /Director/backup/
Command is running. Monitor live status and results in Director/log/smsave.log
ALR1325I: The lightweight runtime has started.
Command completed successfully
p55701p01(root)/>
```

5.4.4 smrestore

The **smrestore** command restores the IBM Systems Director persistent data, including file system data and databases. You must stop the IBM Systems Director server before running the command.

Syntax: **smrestore** [-h | --help | -?] -sourceDir [targetPath] -dbSourceDir [dbTargetPath] -dbUserName [User] -dbUserPwd [Password] -noPrompt [true | false]

- **sourceDir**
Location to backup files.
If not specified, the default location will be `/Director/backup/`. Refer to Example 5-4.
If the default location is used, then there can only be one set of save files in `/Director/backup/`.
More than one set of save files in `/Director/backup/` makes this parameter mandatory.

- dbSourceDir
Optional location to locate database backup files.
If not specified, the location will be defaulted based on the database type.
- dbUserName
Optional user name for database access. Must be used in conjunction with dbUserPwd. If not specified, current Director credentials will be used.
- dbUserPwd
Optional password for database access. Must be used in conjunction with dbUserName. If not specified, current Director credentials will be used.
- noPrompt
Optional. Run without a confirmation prompt.

Example 5-4 IBM Systems Director restore command

```
p5570lp01(root)/> smrestore -sourceDir /Director/backup/  
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 Director/log/smrestore.log  
ALR1325I: The lightweight runtime has started.  
Command completed successfully  
p5570lp01(root)/>
```

Note: If you set -noPrompt to true, you will receive a warning after issuing the **smrestore** command saying that the command execution is destructive and will completely replace the current IBM Systems Director data with the restored image. You will be prompted to confirm that you still want to run the operation.



isdstat script

This appendix provides the `isdstat` script to get the status of the cluster nodes in IBM Systems Director. This script correlates to the `csmdstat` command which provides the cluster node status in an IBM Cluster Systems Management environment.

isdstat script sample

Example A-1 provides a sample script to get node status in the IBM Systems Director. For explanation on what this script provides, refer to 2.7.3, “Managing node status information with csmstat” on page 22.

Example A-1 isdstat script

```
#!/bin/ksh

# Server headers
printf "%-12s%-12s%-12s%-12s%-12s%-12.11s%-12s\n" "Hostname" "HealthState"
"Communicate" "AccessState" "LPAR" "PowerState" "HostSystem" "PowerState"

# HMC header
printf "%36s%-12s%-12s%-12s\n" " " "HMC" "State" "Communicate" "AccessState"

printf "%-96.96s\n"
"-----"
"-----"

# List server status

for lPAR in `smcli lssys -A
InstalledOSDisplayName,SerialNumber,PowerState,OperatingState,CommunicationState,A
ccessState -t Server -w "HypervisorPlatformRole=VirtualContainer" -d:"|grep -v
Unsupported|sed "s/ //"` ; do

LPAR=`echo $lPAR|sed "s:/ /g"`
set -A LPARarr $LPAR

oS=`smcli lssys -A HealthState,CommunicationState,AccessState -t OperatingSystem
-w "DisplayName=${LPARarr[1]}" -d:"|sed "s/ //"`

OS=`echo $oS|sed "s:/ /g"`
set -A OSarr $OS

mACHINE=`smcli lssys -A
TWGMParentSet,PowerState,OperatingState,CommunicationState,AccessState -t Server
-w "HypervisorPlatformRole=HostPlatform AND SerialNumber=${LPARarr[2]}" -d:"|sed
"s/ //"`

MACHINE=`echo $mACHINE|sed "s:/ /g"`
set -A MACHarr $MACHINE

hMCoID=`echo ${MACHarr[1]}|sed "s/{///"|sed "s/}///"`

hMC1oID="NO HMC FOUND!!!"
hMC2oID="NO HMC FOUND!!!"
hMC1oID=`echo $hMCoID|cut -f1 -d","`
hMC2oID=`echo $hMCoID|cut -f2 -d","`

if [ ! -z "$hMC1oID" ] ; then
    hMC1=`smcli lssys -A OperatingState,CommunicationState,AccessState -t
HardwareManagementConsole -w "OID=$hMC1oID" -d:"|sed "s/ //"`

```



```

        HMC1=`echo $hMC1|sed "s:/ /g"`
        set -A HMC1arr $HMC1
else
        HMC1="      "
        set -A HMC1arr $HMC1
fi

# LPARarr array elements:
# 0: lPARnAME
# 1: hOSTnAME
# 2: sERIAL
# 3: lPARpOWERSTATE
# 4: lPARoPsTATE
# 5: lPARcOMsTATE
# 6: lPARaCCsTATE

# OSarr array elements
# 0: oSnAME
# 1: oShEALTH
# 2: oScOMsTATE
# 3: oSaCCsTATE

# MACHINE array elements
# 0: mACHnAME
# 1: mACHhMC
# 2: mACHpOWERSTATE
# 3: mACHoPsTATE
# 4: mACHcOMsTATE
# 5: mACHaCCsTATE

# HMC1arr array elements
# 0: hMC1nAME
# 1: hMC1oPsTATE
# 2: hMC1cOMsTATE
# 3: hMC1aCCsTATE

# HMC2arr array elements
# 0: hMC2nAME
# 1: hMC2oPsTATE
# 2: hMC2cOMsTATE
# 3: hMC2aCCsTATE

# Full status
# echo
$hOSTnAME\t$oShEALTH\t$oScOMsTATE\t$oSaCCsTATE\t$lPARnAME\t$lPARpOWERSTA
TE\t$lPARoPsTATE\t$lPARcOMsTATE\t$lPARaCCsTATE\t$mACHnAME\t$mACHpOWERSTA
TE\t$mACHoPsTATE\t$mACHcOMsTATE\t$mACHaCCsTATE\t$hMC1nAME\t$hMC1oPsTATE"
\t$hMC1cOMsTATE\t$hMC1aCCsTATE

# OS status with HW info
printf "%-12s%-12s%-12s%-12s%-12s%-12s%-12.11s%-12s\n" ${LPARarr[1]} ${OSarr[1]}
${OSarr[2]} ${OSarr[3]} ${LPARarr[0]} ${LPARarr[3]} ${MACHarr[0]} ${MACHarr[2]}

# HMC info

```

```

printf "%36s%-12s%-12s%-12s%-12s\n" "First HMC:  " ${HMC1arr[0]} ${HMC1arr[1]}
${HMC1arr[2]} ${HMC1arr[3]}

if [ "$hMC1oID" != "$hMC2oID" -a "$hMC1oID"<>"NO HMC FOUND!!!" ] ; then
    hMC2=`smcli lssys -A OperatingState,CommunicationState,AccessState -t
HardwareManagementConsole -w "OID=$hMC2oID" -d:"|sed "s/ //g"`
    HMC2=`echo $hMC2|sed "s/:/ /g"`
    set -A HMC2arr $HMC2
    printf "%36s%-12s%-12s%-12s%-12s\n" "Second HMC:  " ${HMC2arr[0]}
${HMC2arr[1]} ${HMC2arr[2]} ${HMC2arr[3]}
fi

done

```

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 VMControl Implementation Guide on Power Systems*, SG24-7829
- ▶ *Configuring and Managing AIX Clusters Using xCAT 2*, SG24-7766
- ▶ *IBM AIX Version 7.1 Differences Guide*, SG24-7910

You can search for, view, download or order these documents and other Redbooks, Redpapers, Web Docs, draft and additional materials, at the following website:

ibm.com/redbooks

Other publications

These publications are also relevant as further information sources:

- ▶ *IBM Systems Director VMControl Installation and User's Guide Version 2 Release 3*:
http://publib.boulder.ibm.com/infocenter/director/v6r2x/topic/com.ibm.director.vim.helps.doc/fsd0_vim_pdf.pdf
- ▶ *IBM Cluster Systems Management for AIX and Linux Planning and Installation Guide Version 1, Release 7.1*, SA23-1344-05
- ▶ *IBM Systems Director for AIX Planning, Installation and Configuration Guide Version 6.2.1*, GI11-8709-06
- ▶ *RSCT RMC Programming Guide*, SA23-1346

Online resources

These websites are also relevant as further information sources:

- ▶ CSM Planning and Installation Guide
<http://publib.boulder.ibm.com/epubs/pdf/a2313445.pdf>
- ▶ CSM Administration Guide
<http://publib.boulder.ibm.com/epubs/pdf/a2313445.pdf>
- ▶ IBM Systems Director publications
http://publib.boulder.ibm.com/infocenter/director/v6r2x/topic/com.ibm.director.main.helps.doc/fqp0_bk_install_gde_aix.pdf

- ▶ IBM Systems Director Redbooks publication
<http://www.redbooks.ibm.com/redbooks/pdfs/sg247694.pdf>
- ▶ IBM Systems Director and VMControl Infocenter
http://publib.boulder.ibm.com/infocenter/director/v6r2x/topic/com.ibm.director.main.helps.doc/fqm0_main.html
- ▶ AIX and NIM (V7.1) Infocenter
<http://publib.boulder.ibm.com/infocenter/aix/v7r1/index.jsp>

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services

Index

A

Agent Manager 8
Agentless managed systems 9

B

Baseboard Management Controller (BMC) 90

C

Command

bc 27
cfgdbcmd 20
cfmupdatenode 46, 89
chcondition 80
chevtautopl 85
chresmonthresh 80
chresmonthresh -A 80
chresmonthresh -D 80
cimssubscribe 20
configure.sh 129
csm2nimnodes 115, 119
csmbbackup 70, 139–140
csmbbackup -r 140
csmbbackup -r Resource_Class> -d directory 140
csmrestore 140
csmsetupnim 115
csmstat 22, 45
dconsole 100
dcp 85
dsh 85
dshbak 85
getadapters 115, 119
lscondition 85
lsconditions 80
lscondres 80
lscondresp 46, 85
lshwinfo 119
lspp 62
lsresmon 80
lsresmonthresh 80
lsresponse 80
lssrc -ls ctrmc 98
lssys 69, 81
mcli lscustomization 115
mkconditions 79
mkcondresp 80
mkevtautopln 84–85
mkflashfiles 92
mkresmonthresh 79–80
mkresponse 80
netboot 115
netstat 128
nim -o cust -a adapter_def="adapter_def NIM re-
source" "nodename" 121

nimadapters 121
rconsole 100
rdist 21
rflash 92
rflashscan 93
rmcondition 80
rmnode 69
rmresmonthresh 80
rpower 100
runresmon 79
shwinfo -p hmc -c HMC_IP_or_hostname 93
smcli 20, 119
smcli -a deploy_new 115
smcli accesssys 64
smcli accesssys -u hscdir -p 67
smcli accesssys -u root -p 64, 69
smcli captureva 109
smcli cleanupd -mFv -P Platform=Director 62
smcli collectinv -n -p "All Inventory" 67
smcli collectinv -n -p "All Inventory" 68
smcli collectinv -n -p "All Inventory" 69
smcli discover -i -t OperatingSystem 63
smcli discover -i -t OperatingSystem 67
smcli dsh --help 88
smcli importupd -r /mnt/cdrom 95
smcli installneeded -v -F 62
smcli lsevtautopln 84
smcli lssys 63
smcli lssys -l 23
smcli lssys -l -w "AccessState=Locked" 63
smcli lsver 62
smit mksysb 140
smreset 20
smrestore 20, 142–143
smsave 20, 142–143
smstart 20, 62
smstatus 20
smstatus -r 62
smstop 20, 62
startcondresp 80
startsvc director_agent 28
stopcondresp 80
stopsrc -s rsct_rm 71
systemid 115
uninstallms 71
updatehwdev -k 91, 93
updatenode -c -n "nodename" 121
viosvrcmd 27

Common Agent 8
Common Agent Services (CAS) 74, 138
Common Information Model (CIM) 101
Credential Transformation Service (CTS) 104

D

Daemon

sshd 81

Data Encryption Standard (DES) 101

Directory

/cfmroot 89

/csminstall/csm/fw 92

/export/nim/appliances 108

/opt/csm/bin 44

/var/opt/csm/csmdata 140

Distributed Systems Management (DSM) 108

domain management server 76

E

Encrypted Web Service (SSL) 101

Encryption Standard (AES) 101

Event Automation Plans (EAP) 82

F

File

/etc/group 46

/etc/hosts 46

/etc/passwd 46, 102

/etc/resolv.conf 120

/etc/security/limits 66

/etc/security/passwd 102

Filesystem

/nimrepo 66

/opt 62

H

Hardware Management Console (HMC) 91

I

IBM Cluster Systems Management (CSM) 1

IBM Engineering and Scientific Subroutine Library (ESSL)
136

IBM General Parallel File System (GPFS) 136

IBM Management Process Command-Line Interface (MP-
CLI) 90

IBM Systems Director 2

Integrated Virtualization Manager (IVM) 94

L

Licensed Internal Code (LIC) 91

Lightweight Directory Access Protocol (LDAP) 98

M

Managed System 9

Management Server 8

N

Network Installation Manager (NIM) 4

Network Management Protocol (SNMP) 101

P

Platform Agent 9

R

Redbooks website 149

Contact us xi

Reliable Scalable Cluster Technology (RSCT) 74

Remote Supervisor Adapter (RSA) 90

Role-based access control (RBAC) 98

S

Secure Sockets Layer (SSL) 99

Service Location Protocol (SLP) 101

Shared Product Object (SPOT) 109

smcli -a deploy_existing 115

T

Tivoli Provisioning Manager (TPM) 74

U

UNIX-Domain Sockets (UDS) 97

V

Virtual Network Computing (VNC) 91

VMControl Standard Edition 4

IBM CSM to IBM Systems Director Transformation Guide

(0.2"spine)
0.17"<->0.473"
90<->249 pages



IBM CSM to IBM Systems Director Transformation Guide



Provides planning guidelines

Includes features comparison

Discusses transformation scenarios

For many years, IBM Cluster Systems Management (CSM) provided a single point of management for IBM Power Systems servers running the AIX operating system. Now you can transform your environment to IBM Systems Director, which provides CSM clients with the next generation of Cluster Systems Management for their Power Systems servers.

The target audience for this IBM Redbooks publication includes technical professionals (IT consultants, technical support staff, IT Architects, and IT Specialists) responsible for planning and implementing the Cluster Systems Management software transformation from CSM to 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

SG24-8002-00

ISBN 0738436887