IBM® Tivoli®

Deploying Cloud Components on POWER

- Introducing cloud on POWER and the component model overview
- Installing and configuring the cloud components
- Learning how to manage virtual machines

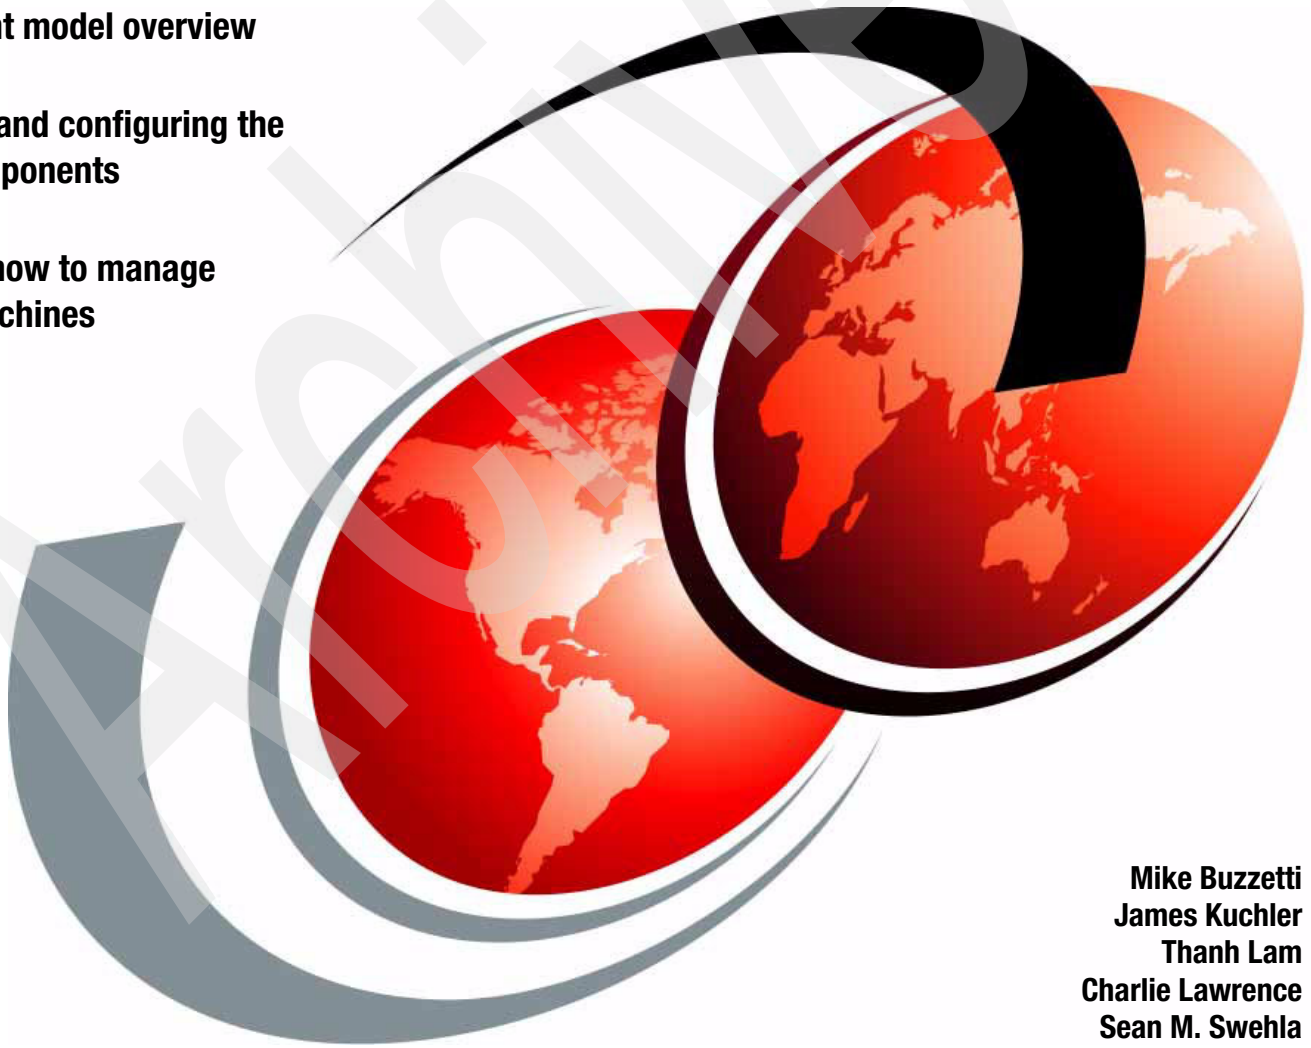

Mike Buzzetti
James Kuchler
Thanh Lam
Charlie Lawrence
Sean M. Swehla

Redbooks

ibm.com/redbooks

International Technical Support Organization

Deploying Cloud Components on POWER

November 2011

SG24-8009-00

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

First Edition (November 2011)

This edition applies to Version 6 Release 1 of AIX on POWER systems.

Contents

Notices

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

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

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not 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®
Blue Gene®
DB2®
GPFS™
IBM Systems Director Active Energy Manager™
IBM®
Jazz™
LoadLeveler®
POWER6®
PowerVM™
POWER®
Redbooks®
Redpaper™
Redbooks (logo) ®
System x®
System z®
Systems Director VMControl™
Tivoli®
z/Architecture®
z/OS®
z/VM®

The following terms are trademarks of other companies:

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

“The world is changing. A new reality is emerging for organizations of every size from every part of the planet. It’s called the cloud—a profound evolution of IT with revolutionary implications for business and society, creating new possibilities and enabling more efficient, flexible and collaborative computing models.”

_ http://www.ibm.com/cloud-computing/us/en/

This IBM® Redbooks™ publication applies to Version 6 Release 1 of AIX® on POWER® systems. This book is provided as an additional resource as you investigate or consider implementing and deploying a cloud in a POWER environment in the context of infrastructure as a service.

This book is intended for anyone who wants to learn more about cloud computing on POWER systems.

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. Three members of this team, Buzzetti, Kuchler, and Lawrence, also authored *Deploying a Cloud on IBM System z,* REDP-4711.

Mike Buzzetti is an IT Architect at the IBM Design Center with worldwide focus on client enterprise infrastructures. He began his IBM career in 2003 at the Test and Integration Center for Linux. In 2006, Mike joined the Design Center, where he helps clients design optimized IT Infrastructures. He designed a number of infrastructures that featured Linux on the mainframe and has had extensive experience helping clients use virtualization in complex environments. More recently, Mike has been a leader in implementing cloud computing.

Mike has authored a book on Java 2 Platform, Enterprise Edition (J2EE) on z/OS® security, as well as a number of whitepapers. He is a regular presenter at user conferences and a number of IBM-sponsored venues. Prior to joining IBM, Mike was a System Administrator and Programmer for the Chemung County (NY) Department of Social Services.

James Kuchler is a Systems Support Specialist currently working on projects related to cloud computing on System z® using z/VM® and Tivoli® products. Previously, James worked for six years as a Performance Tester for z/Linux and z/OS, where he supported the LSPR’s mission of assessing relative processor capacity for several benchmarks and System z models. In the role of performance tester, he also worked with another team to discover performance flaws in pre-release components of z/OS, while providing feedback to development and contributing to whitepapers. James holds a B.S. in Computer Science from the State University of New York at New Paltz.

Thanh Lam joined the IBM High Performance Supercomputer System Lab in 1988 and has been working in the areas of High Performance Computing (HPC) applications enablement and Cluster System Test. He possesses a broad range of skills in parallel applications and job scheduling and system resource management on large scale clusters, including AIX® on POWER systems and Linux on POWER systems and System x®.

In 2005, after co-authoring one of the IBM Redbooks® about problem determination on Blue Gene® L (BG/L), Thanh started working as a team lead for HPC Software System Testing on BG/L and BG/P. His accomplishments included bringing together the two test efforts of LoadLeveler® and General Parallel File System (GPFS™) into an integrated test process, which ensures better test coverage and shortens test duration.

Thanh enjoys exploring new technologies in virtualization and cloud computing. He also contributes to the technical communities by writing IBM Redbooks and participating in IBM Academy studies. Thanh is now an IT management consultant in Lab Services and Technical Training. He develops teaching materials for the cloud computing courses and will teach face-to-face or online instructor-led classes.

Charlie Lawrence joined IBM in 1966 and has been involved in various aspects of systems design, development, and testing, as well as in related curriculum and course development efforts that date back to the early days of System/360.

As an IBM instructor, much of Charlie's focus has been in training both new hires and experienced systems development programmers to prepare them for development and support positions in both the VM and z/OS lineage of systems. He served multiple assignments as a staff member of the IBM Kingston and IBM Mid-Hudson Valley education organizations, eventually based at IBM Poughkeepsie. He was also an adjunct instructor at Ulster County Community College (Stone Ridge, NY), where he taught various introductory computer courses, as well as Assembler Language Programming.

His more recent endeavors include contributing to z/Architecture® ESA Modal Extensions (ESAME) systems assurance and verification, participating in the design and testing of the Common Event Adapter, and also participating in the design and development of test cases to verify management by exception views of the Tivoli Enterprise Portal. Prior to joining the cloud computing project that is described in this paper, he was part of the Poughkeepsie-based IBM team that delivered the Predictive Failure Analysis (PFA) z/OS component.

Sean M. Swehla is an Architect at the IBM Design Center in Poughkeepsie. He began his IBM career in 2005 at the IBM Design Center. He left the Design Center for a period to contribute to Build Automation in High Performance Computing (HPC). Following that assignment, Sean participated in projects within the Linux Technology Center. He returned to the IBM Design Center in 2011 where he has been involved in various aspects of Cloud-related projects.

Sean holds dual degrees from the State University of New York at New Paltz: a B.A. in Computer Science and a B.S. in Music with a concentration in Jazz Studies.

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

Mike Ebbers
International Technical Support Organization, IBM Poughkeepsie

Jay Brennemen
Systems & Technology Group, IBM Poughkeepsie

Tyrone Dechalus
Systems & Technology Group, IBM Atlanta

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 or other IBM Redbooks publications in one of the following ways:

- Use the online **Contact us** review IBM Redbooks form found at:

  ibm.com/redbooks
- Send your comments in an email to:

  redbooks@us.ibm.com
- Mail your comments to:

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

Stay connected to IBM Redbooks

- Find us on Facebook:

  http://www.facebook.com/IBMRedbooks
- Follow us on Twitter:

  http://twitter.com/ibmredbooks
- Look for us on LinkedIn:

  http://www.linkedin.com/groups?home=&gid=2130806
- Explore new Redbooks publications, residencies, and workshops with the IBM Redbooks weekly newsletter:

  https://www.redbooks.ibm.com/Redbooks.nsf/subscribe?OpenForm
- Stay current on recent Redbooks publications with RSS Feeds:

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

1

Introduction and component model overview

Cloud has become a reality for IT shops.

Even if they have not implemented cloud, many IT shops, perhaps even yours, are competing with external cloud offerings, whether they realize it or not. This new reality has evolved as lines of business (LOBs) supplement their own IT shops by taking advantage of external providers of cloud offerings.

(LOBs) supplementing their own IT shops by taking advantage of external providers of cloud offerings not only affects IT, but also it can be bad for the business. Public cloud offerings can present exposures to the company that users generally do not see. Many users utilizing public cloud services have not considered the potential for security, compliance, and availability issues that often accompany the use of these services. Instead, IT needs to provide the kinds of services that LOBs seek in the public space in order to protect the business from the inherent dangers of those publicly offered services.

Cloud represents a new business model, in addition to being a deployment model. This new business and deployment model requires a process discipline, as well as the use of a corresponding set of technology. IT must operate differently to position itself to manage cloud resources effectively. The new model requires an understanding of the hardware configuration, software images, a virtualized storage infrastructure, and network management. The understanding and integration of these aspects of the computing environment are often handled by separate groups within the IT team.

This paper describes the base components that are needed to migrate the infrastructures that provide the current IT services to cloud services with a focus on POWER Systems and their related network, storage, and software.

Many of these components are not new technologies, and you are likely already familiar with them. The cloud aspect of this paper is presented in the context of using these components to “offer” infrastructure as a service (IaaS) on the POWER platform.

1.1 Component model approach to implementing cloud components

Before discussing the details of a POWER-based implementation of cloud components, this section first describes three component models: Managed From, Managed Through, and Managed To.

The Managed From component model provides the ability to utilize the Managed Through component model to instantiate virtual entities in the Managed To component model (sometimes referred to as guests, virtual machines, Linux instances, or AIX instances). The creation or instantiation is accomplished by the Managed Through component model at the request of and under the direction of the Managed From component model. The Managed From might be viewed as the point of contact for administrators to request and terminate instances of virtual entities or IaaS.

This section provides an introduction of the Managed From, Managed Through, and Managed To component models in the context of a POWER environment. This section includes the relationship between or interaction of the items that are depicted in Figure 1-1 Managed From - Through - To component models.

Table 1-1 on page 3 has a brief description of each item in Figure 1-1.

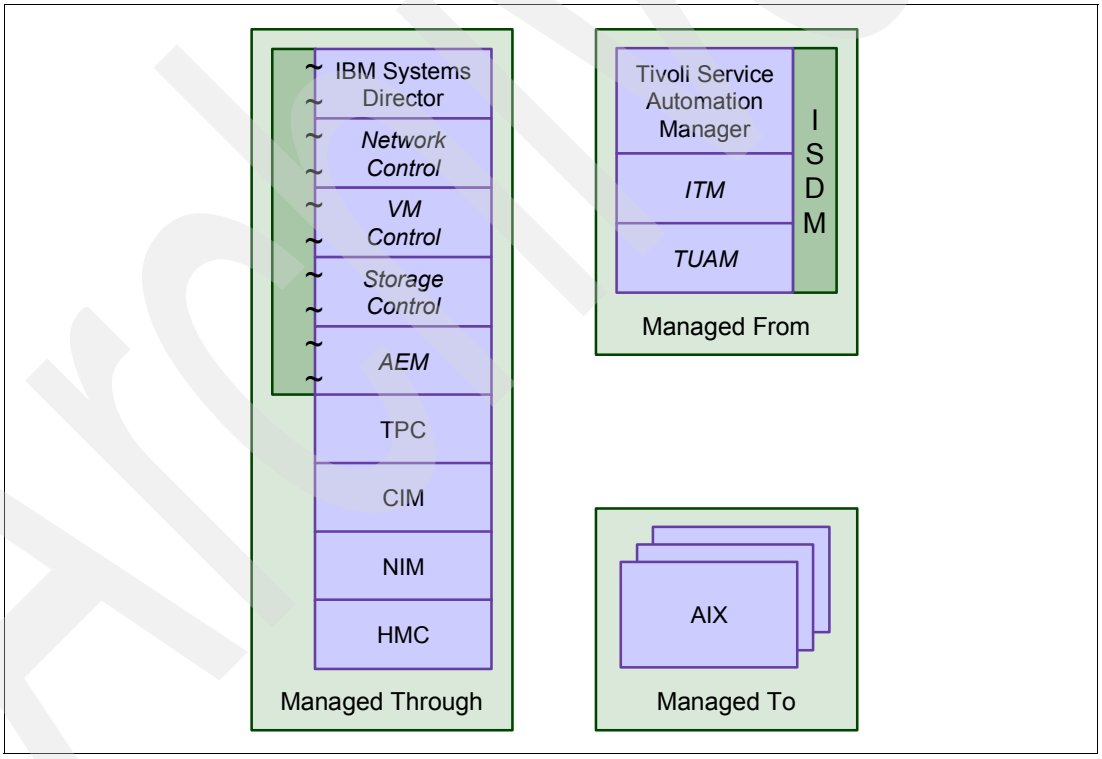

Figure 1-1 Managed From - Through - To component models

Table 1-1 Component model description and legend

Component	Description
Managed Through	
Systems Director	IBM Systems Director provides a single-point-of-control via a consistent user interface that facilitates common IT management tasks related to the total IT environment. The tasks relate to servers, storage, and network resources.
Network Control (Plug-in)	Provides a view across servers and storage facility management applications.
VM Control (Plug-in)	IBM Systems Director VMControl™ V2.3.0 facilitates the deployment of virtual appliances to create virtual servers. In our implementation, we configured the virtual appliances with the AIX and software applications needed.
Storage Control (Plug-in)	The Storage Control Plug-in (integrated with the VM Control Plug-in) provides support for storage provisioning for image creation, deployment, and cloning. You can obtain more information about Storage Control at http://www-03.ibm.com/systems/software/director/storage/.
AEM (Plug-in)	The Active Energy Manager (AEM) extends the IBM Systems Director capabilities in order to provide a unified view of, and single point of control for, the power usage of all the diverse components in the data center. You can obtain more information about AEM at http://www-03.ibm.com/systems/software/director/aem/.
TPC	IBM Total Storage Productivity Center (TPC) Suite reduces the complexity of managing storage environments. TPC works with storage networks, storage subsystems, and storage clients to provide an end-to-end management solution.
CIM	The Common Information Model (CIM) is a standardized way to describe objects and their relationships. CIM provides a common description of an environment to disjointed management products.
NIM	Network Installation Management (NIM) facilitates the installation and management of AIX file sets on machines over the network. You can obtain a detailed discussion of NIM in *NIM from A to Z*, SG24-7296.
HMC	Hardware Management Console (HMC).
Managed From	
Systems Director	In this context, you can use IBM Systems Director also as an administrative end point for controlling a cloud environment.
Tivoli Service Automation Manager	Delivered as part of the IBM Service Delivery Manager (ISDM). The Tivoli Service Automation Manager facilitates a catalog of available services and an associated portal that can be used to request and reserve services from the catalog.
ITM	IBM Tivoli Monitoring facilitates the monitoring of performance indicators.
TUAM	The Tivoli Usage and Accounting Manager facilitates tracking the usage of shared resources (down to the user level).
Managed To	
AIX	The base operating system built into the installation image.
Linux on POWER	A UNIX-like operating system that can be installed in a POWER logical partition (LPAR).

1.2 First steps toward implementing cloud components

The remainder of this document focuses on the first steps toward implementing cloud components.

Figure 1-2 provides a view of the actual environment implemented for this document. Note that this view varies somewhat from Figure 1-1 on page 2, in that the Managed From tasks were supported through direct contact or utilization of IBM Systems Director.

There are many ways to manage these cloud components. This document describes a simple set of tools that can be used to provide IaaS. It focuses on IBM Systems Director as the Managed From component, but there are a number of other options available. Other options include Tivoli Service Automation Manager and IBM Service Delivery Manager, which will be discussed in future publications.

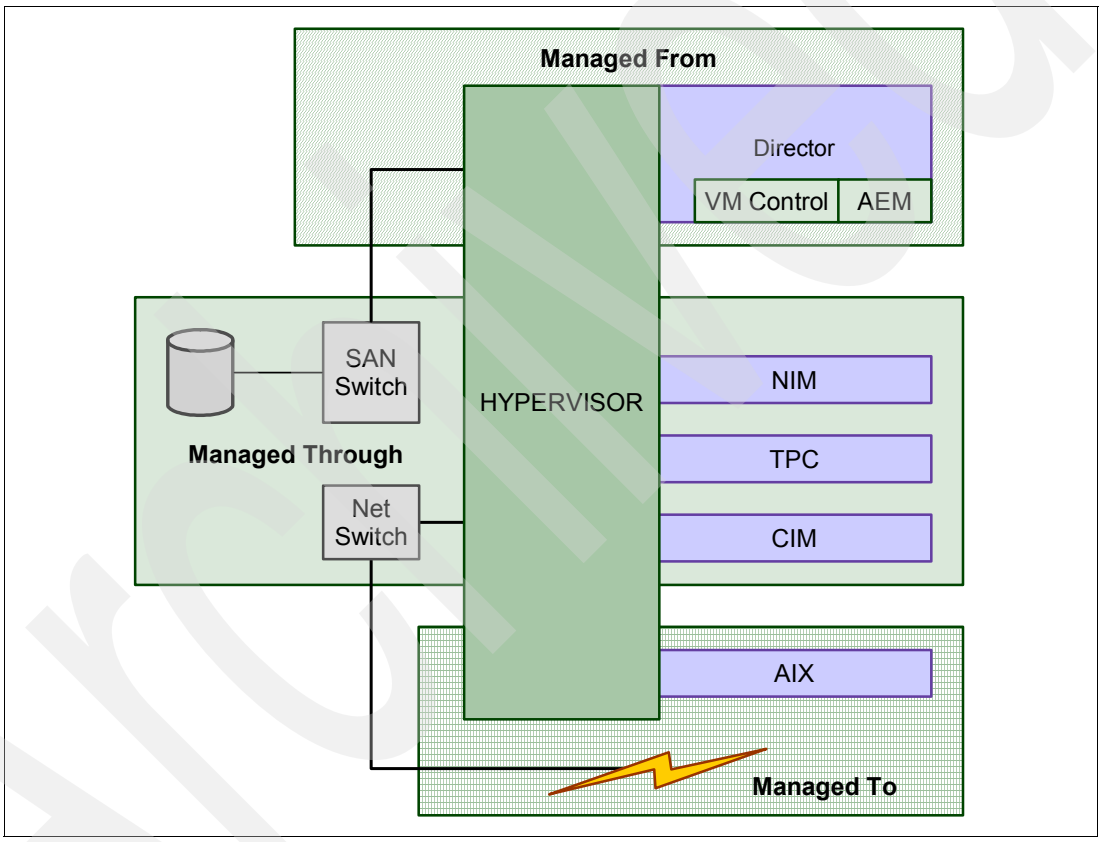

Figure 1-2 Managed From - Through - To component models implemented for this document

Note that in Figure 1-2, the boxes labeled Director, NIM, TPC, and CIM represent functionality that was implemented as individual LPARs. Figure 1-3 on page 5 provides another view of these LPARs in the context of the order they were created, from right to left as they appear in Figure 1-3 on page 5.

As each LPAR was created and the installation of functionality to be contained in each LPAR was completed, the environment was one step closer to a full implementation of the relationship needed to implement cloud. See Figure 1-3 on page 5.

The notations over each LPAR in Figure 1-3 on page 5 are intended to show how the LPARs and subsequent software installations were created. Ranging from labor-intensive and

administrator-intensive manual tasks to the point where it was possible to take advantage of previously implemented LPARs, again from right to left.

Figure 1-3 also shows certain relationships or connections between LPARs that facilitate communication and control.

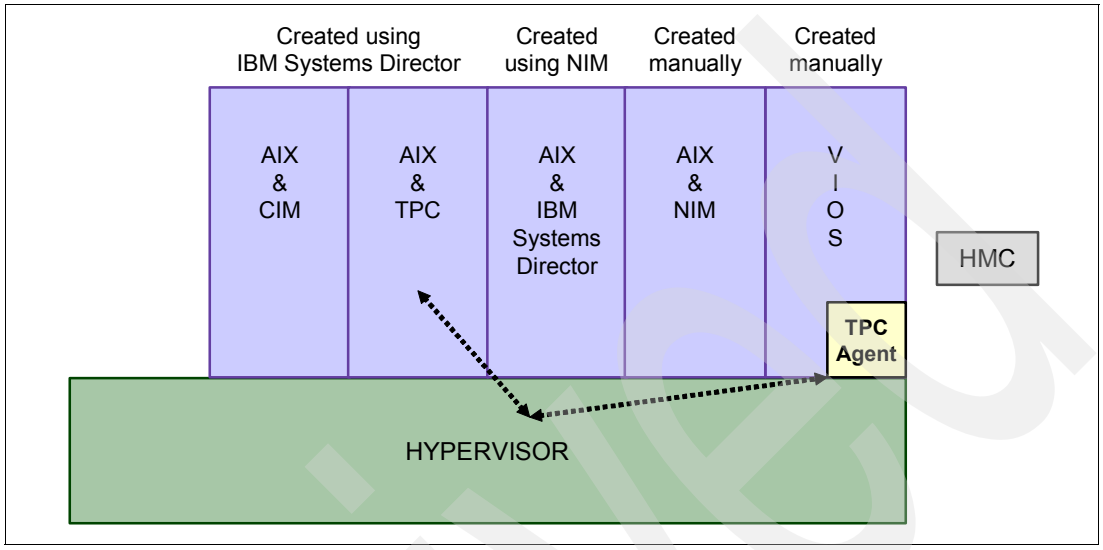

Figure 1-3 LPARs created and software installed (implemented from right to left)

You can obtain more details and discussion of the creation of these LPARs in 2.4, “Creating LPARs” on page 17. If you are unfamiliar with the concepts or tasks that are associated with LPAR creation, read 2.3, “Introduction to LPAR concepts” on page 14 before reading 2.4, “Creating LPARs” on page 17.

2

Initial steps

This chapter provides an overview of the base infrastructure that is available for the implementation of cloud components for the purposes of this paper. Next, we provide a brief discussion of LPAR concepts in the context of sharing resources. The contents of that section contribute to your understanding and appreciation of Chapter 3, “Installing and configuring components” on page 23, and it helps you to prepare to implement your own cloud components on POWER.

In addition, this chapter also contains a discussion of how we configured the network and storage using the Hardware Management Console (HMC) to provide resource pools for the cloud services.

The remainder of this chapter discusses the LPARs that were created and why. If you are already familiar with LPAR concepts and virtual I/O (VIO), you might want to skip to 2.3, “Introduction to LPAR concepts” on page 14.

2.1 Hardware inventory using the HMC

The two systems in the environment that is described in this paper differed in the number of processors and amount of installed physical memory. To distinguish between the two, a non-technical but easily understood naming convention of Little Box and Big Box was used. These names were never officially specified in any configuration or networking definitions. However, in conversations, it was quite clear which box was being referenced. We use this same naming convention in this paper.

Before any extensive project planning began, we completed an inventory of the hardware that was available for this effort. During that process, we identified two POWER6® systems. Using the HMC, we accessed windows that are similar to the window that is shown in Figure 2-1 through Figure 2-4 on page 10. These windows provided the ability to determine the resources attached to each box, as summarized in Table 2-1 on page 11. We completed many of the initial setup and configuration tasks using the HMC.

In Figure 2-1, both systems are POWER6 Model 570 (9117-MMA). This model supports logical partitioning (LPAR). We explain the advantages of using LPARs in 2.4, “Creating LPARs” on page 17.

Using the HMC, you can display the server hardware properties dialog by expanding the System Management navigation menus on the left side. Click the plus symbol (**+**) next to System Management. After you click +, a window that is similar to Figure 2-1 opens.

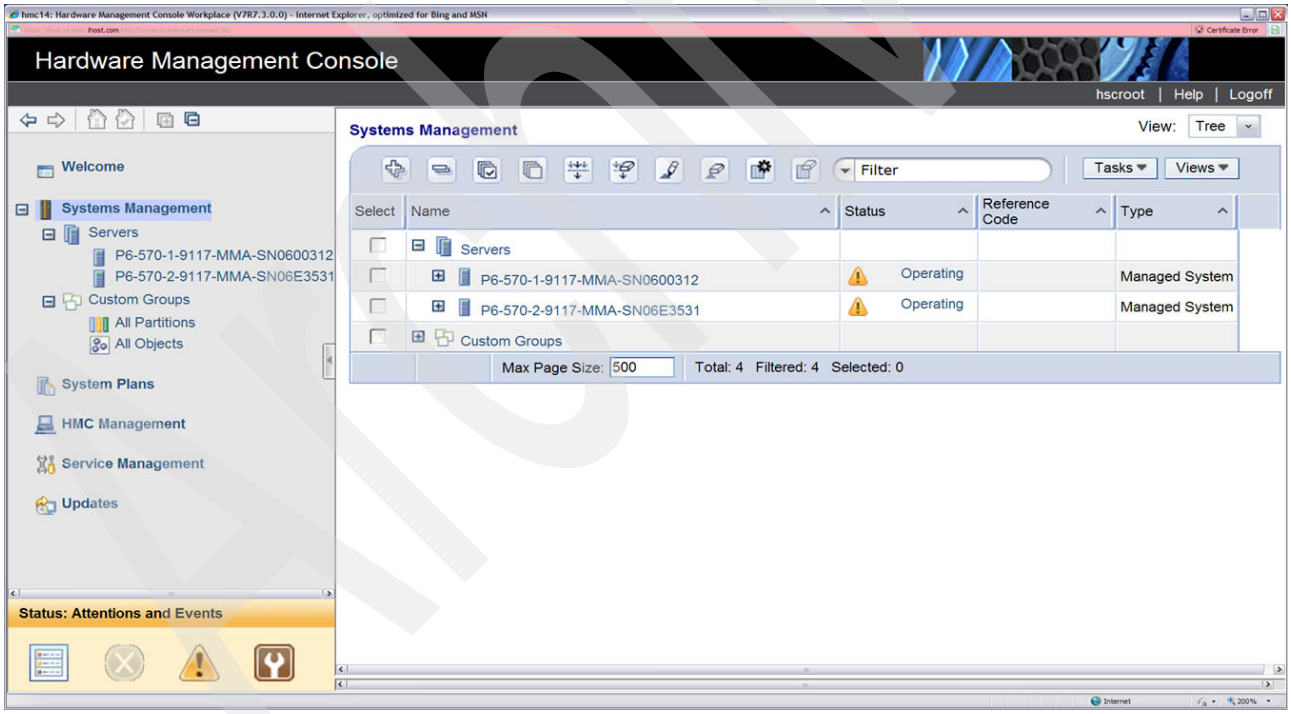

Figure 2-1 Server hardware properties dialog

By expanding the **Custom Groups** menu and selecting **All Objects,** the right side of the window shows the list of the hardware servers and LPARs, as shown in Figure 2-2 on page 9.

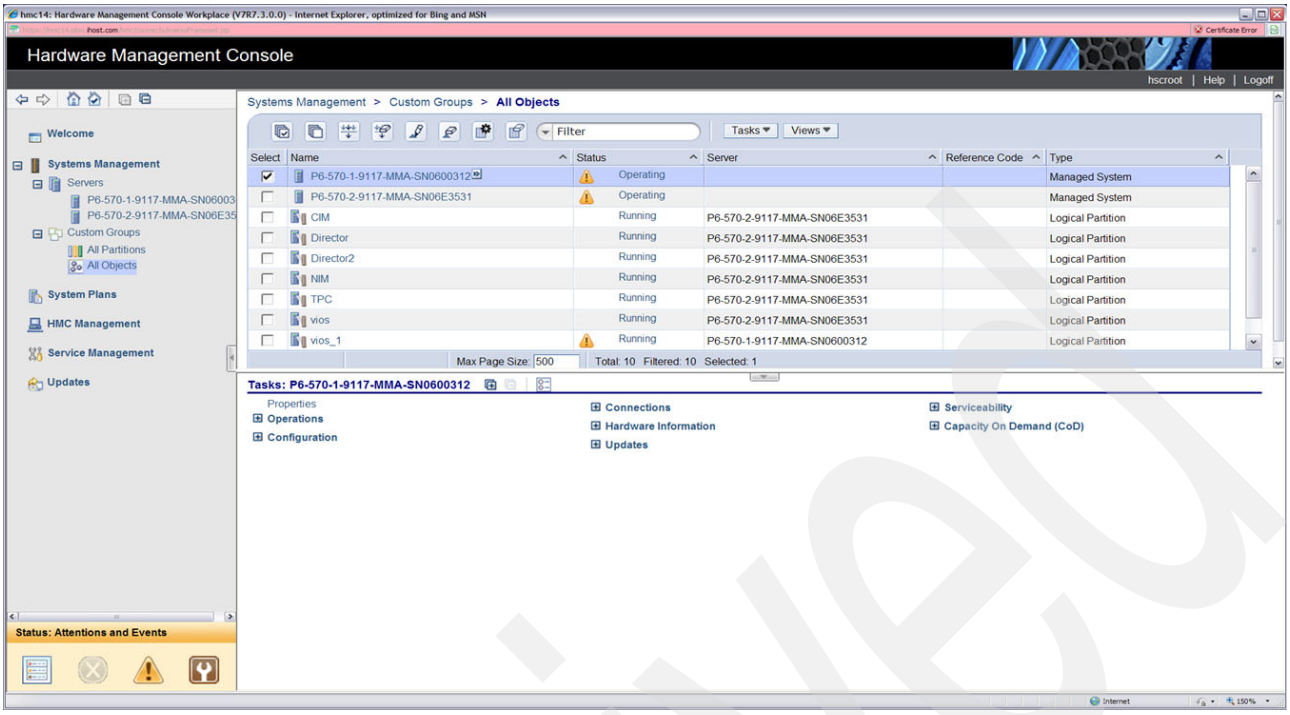

Figure 2-2 List of the hardware servers and LPARs viewed via HMC

At this point, click the **arrow** symbol at the end of the server name (see the highlighted row in Figure 2-2). A list of options appears, as shown in Figure 2-3.

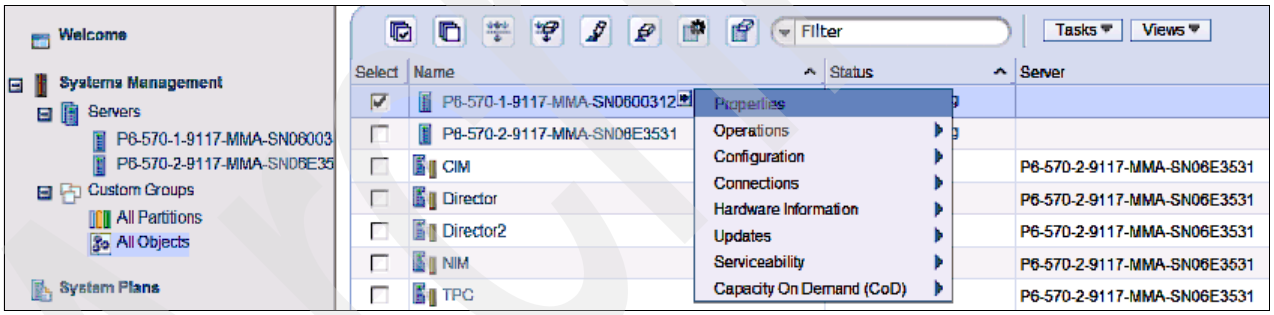

Figure 2-3 List of server-related functions

Subsequently selecting **Properties** (top item on list) produces the dialog that is provided in Figure 2-4 on page 10. Figure 2-4 on page 10 shows that the resulting dialog provides tabs that are labeled for access to processors, memory, and other attributes of the server described on the General tab. In this case, the Big Box (P6-570-1-9117-MMA-SN0600312) is in the Operating state.

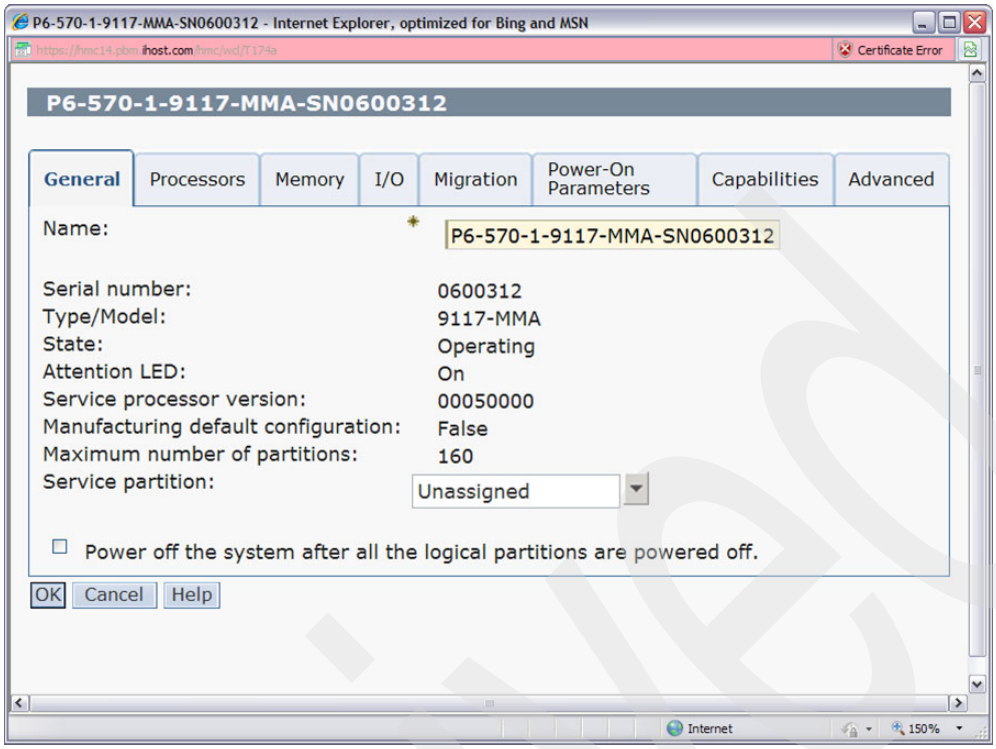

Figure 2-4 Tabs for viewing server properties

When looking further under the **Processors** tab (in Figure 2-4), 16 processors are shown as installed and configurable. In addition, pools were supported. We provide a summary of this investigation in 2.2, “Hardware planning” on page 10.

2.2 Hardware planning

Table 2-1 on page 11 contains a summary of the resources that are available to us on the little box and the big box. The little box was a single-node POWER6 570 and the big box was a quadruple-node POWER6 570. The actual machine identifiers are P6-570-2-9117-MMA-SN06E3521 and P6-570-1-9117-MMA-SN0600312.

In hindsight, it is better to reverse the usage of the Little Box and Big Box so that the “managed through” environment benefits from the redundancy and resources that are available in the Big Box. However, the Little Box became the basis for the “managed through” environment. We defined LPARs and associated them with virtual I/O (VIO), Network Installation Management (NIM), and other functions on the Little Box. We used the Big Box to support the “managed to” entities, with resources being allocated to either AIX or Linux instantiations.

Table 2-1 Little Box and Big Box summary

Little Box **P6-570-2-9117-MMA-SN06E3521**	**Big Box** **P6-570-1-9117-MMA-SN0600312**
Single Node: ► DQD08H9 Processors: 4 Memory: 32 GB Internal disks: disks 0 through 5: capacity = 73, 73, 300, 300, 300, 300 GB	**4 Nodes:** ► DQD08G0 ► DQD08G1 ► DQD08G6 ► DQD08G3 Processors: 16 Memory: 192 GB Internal disks: ► **DQD08G0** disks 0 through 5: capacity = 73, 300, 300, 300, 300, 300 GB ► **DQD08G1** disks 6 through 11: capacity = 300, 300, 300, 300, 300, 300 GB ► **DQD08G6** disks 12 through 17: capacity = 73, 300, 300, 300, 300, 300 GB ► **DQD08G3** disks 18 through 23: capacity = 300, 300, 300, 300, 300, 300 GB
Adapters: ► SCSI: 1 ► Ethernet: 1+ ► Fiber Channel: 1	**Adapters:** ► SCSI: 1 ► Ethernet: 1+ ► Fiber Channel: 1
Storage area network (SAN): ► Switch: Brocade 2005 B64 Version 10 ► DS4800 ► Cables: Connections between POWER6s, DS4800, and switch	

2.2.1 Inventory of resources

In preparation for building cloud components on POWER, we completed an inventory of the available resources. We suggest this inventory and documentation effort as a contributor to a successful implementation.

We physically inspected the POWER boxes in the front and rear to determine the component connections. As an example, Figure 2-5 on page 12 shows a summary of node DQD08G0. (Note that this is one of the nodes in the big box that was previously described in Table 2-1.) The information in this figure corresponds to the photograph of the back of the unit that is provided in Figure 2-6 on page 12.

In Figure 2-5 on page 12, there is a summary of the state of the cabling and connections on the back panel of one node, where the node identifier is DQD08G0. The list that follows Figure 2-5 on page 12 provides a correlation between the elements of Figure 2-5 on page 12 and Figure 2-6 on page 12. Although it was recorded, we have omitted the information about the cable connections on the other end from Figure 2-5 on page 12.

4 NODE SYSTEM:		
DQD08G0		
SLOT	IEEE ADDR	
P1-C1-T1	00-14-5E-74-89-AC	
P1-C1-T2	00-14-5E-74-89-AD	
P1-C2-T1	00-14-5E-74-89-B6	
P1-C2-T2	00-14-5E-74-89-B7	
P1-C3-T1	00-14-5E-74-89-B8	
P1-C3-T2	00-14-5E-74-89-B9	
P1-C4-T1	00-00-C9-74-E6-A8	SAN6 PORT 20
P1-C4-T2	00-00-C9-74-E6-A9	
P1-C5-T1	00-00-C9-76-40-D6	SAN6 PORT 01
P1-C5-T2	00-00-C9-76-40-D7	
P1-C6-T1	00-14-5E-74-86-DE	
P1-C6-T2	00-14-5E-74-86-DF	

Figure 2-5 Cabling and connections for node DQD08G0 (one of four nodes in Big Box)

The following descriptions relate to the arrows in Figure 2-6:

- Arrow A points to a label that contains the text P1-C1 to indicate the location of slot P1-C1 that contains two cable termination points or ports (T1 and T2). These ports correspond to the “slot” column entries, P1-C1-T1 and P1-C1-T2, in Figure 2-5.
- Arrow B points to the cable that is plugged into port T1 for the slot/port ID of P1-C1-T1. The other end of this cable connects to the SAN switch.
- Arrow C points to the cable that is plugged into port T1 for the slot/port ID of P1-C2-T1. The other end of this cable connects to the SAN switch.
- Arrow D points to the cable that is plugged into port T1 for the slot/port ID of P1-C4-T1.
- Arrow E points to the cable that is plugged into port T1 for the slot/port ID of P1-C5-T1.

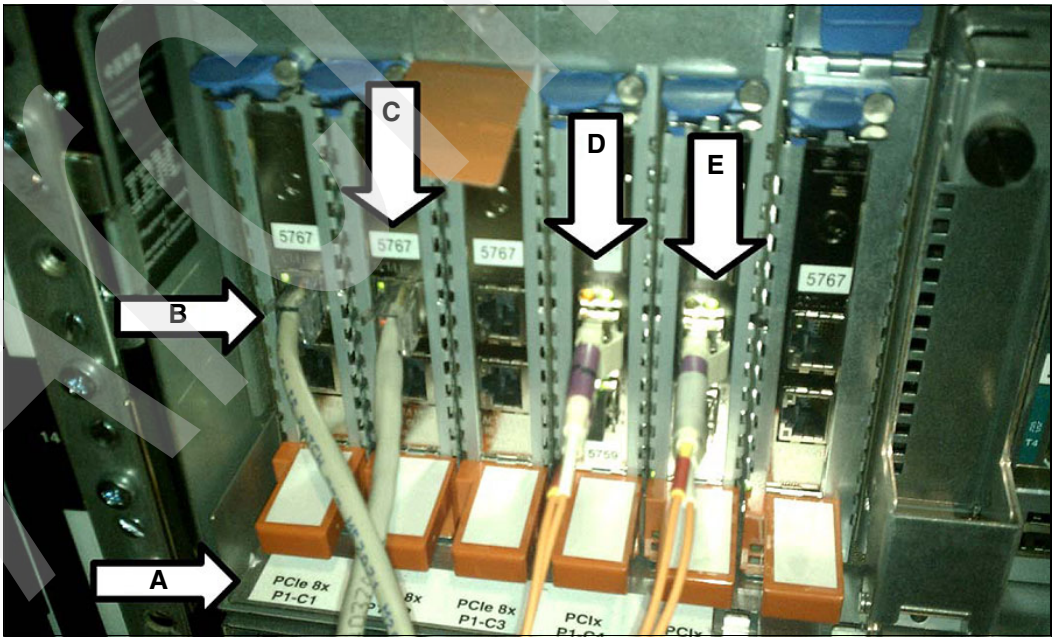

Figure 2-6 Annotated back panel

Completing this inventory of cables, nodes, and boxes (as shown in Figure 2-5 and Figure 2-7 on page 13 through Figure 2-10 on page 14) facilitated HMC and LPAR configuration tasks. You need to perform an inventory before you start any installation and configuration steps.

4 NODE SYSTEM:		
DQD08G1		
SLOT	IEEE ADDR	
P1-C1-T1	00-14-5E-74-87-A6	
P1-C1-T2	00-14-5E-74-87-A7	
P1-C2-T1	00-14-5E-74-88-4A	
P1-C2-T2	00-14-5E-74-88-4B	
P1-C3-T1	00-00-C9-73-C9-11	SAN6 PORT 07
P1-C4-T1	00-00-C9-76-45-FA	SAN6 PORT 13
P1-C4-T2	00-00-C9-76-45-FB	
P1-C5-T1	00-00-C9-76-19-4C	SAN6 PORT 29
P1-C5-T2	00-00-C9-76-19-4D	
P1-C6/8-T1	00-00-C9-73-C1-CA	SAN6 PORT 12

Figure 2-7 Big Box node DQD08G1

4 NODE SYSTEM:		
DQD08G6		
SLOT	IEEE ADDR	
P1-C1-T1	00-00-C9-74-E1-24	SAN6 PORT 06
P1-C1-T2	00-00-C9-74-E1-25	
P1-C2-T1	00-00-C9-74-E0-C6	SAN6 PORT 05
P1-C2-T2	00-00-C9-74-E0-C7	
P1-C3-T1	00-00-C9-74-E1-22	SAN6 PORT 28
P1-C3-T2	00-00-C9-74-E1-23	
P1-C4-T1	00-1A-64-91-8E-F4	
P1-C4-T2	00-1A-64-91-8E-F5	
P1-C5-T1	00-1A-64-91-91-D2	
P1-C5-T2	00-1A-64-91-91-D3	
P1-C6/8-T1	00-00-C9-74-DE-E4	SAN6 PORT 04
P1-C6/8-T2	00-00-C9-74-DE-E5	

Figure 2-8 Big Box node DQD08G6

4 NODE SYSTEM:		
DQD08G3		
SLOT	IEEE ADDR	
P1-C1-T1	00-14-5E-74-85-94	
P1-C1-T2	00-14-5E-74-85-95	
P1-C2-T1	00-14-5E-74-85-3A	
P1-C2-T2	00-14-5E-74-85-3B	
P1-C3-T1	00-14-5E-74-84-BC	
P1-C3-T2	00-14-5E-74-84-BD	
P1-C4-T1	00-00-C9-76-40-E0	SAN6 PORT 21
P1-C4-T2	00-00-C9-76-40-E1	
P1-C5-T1	00-00-C9-76-43-F2	SAN6 PORT 02
P1-C5-T2	00-00-C9-76-43-F3	
P1-C6-T1	00-14-5E-74-85-38	
P1-C6-T2	00-14-5E-74-85-39	

Figure 2-9 Big Box node DQD08G3

SINGLE NODE SYSTEM:		
DQD08H9		
SLOT	IEEE ADDR	
P1-C1-T1	00-14-5E-74-7E-54	
P1-C1-T2	00-14-5E-74-7E-55	
P1-C2-T1	00-00-C9-71-DD-4A	SAN6 PORT 19
P1-C2-T2	00-00-C9-71-DD-4B	SAN6 PORT 0

Figure 2-10 Little Box node DQD08H9

Figure 2-11 depicts the relationship between the boxes and the DS4800 storage that is accessed through DCSAN6.

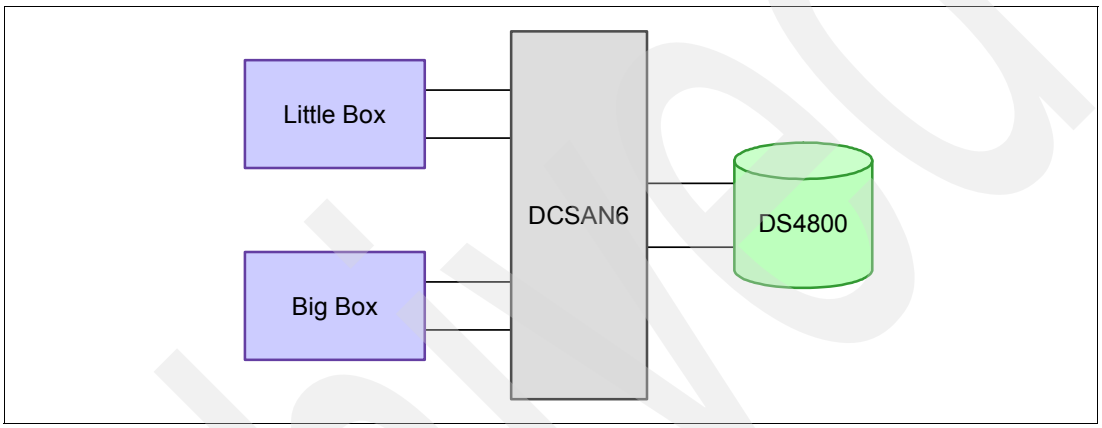

Figure 2-11 Relationship between the boxes and DS4800 storage accessed through DCSAN6

Figure 2-12 depicts the networking relationship between the boxes and the VLAN74 and VLAN43 network switches. VLAN74 provided access via the HMC, and VLAN43 provided networking access to the public local area network (LAN).

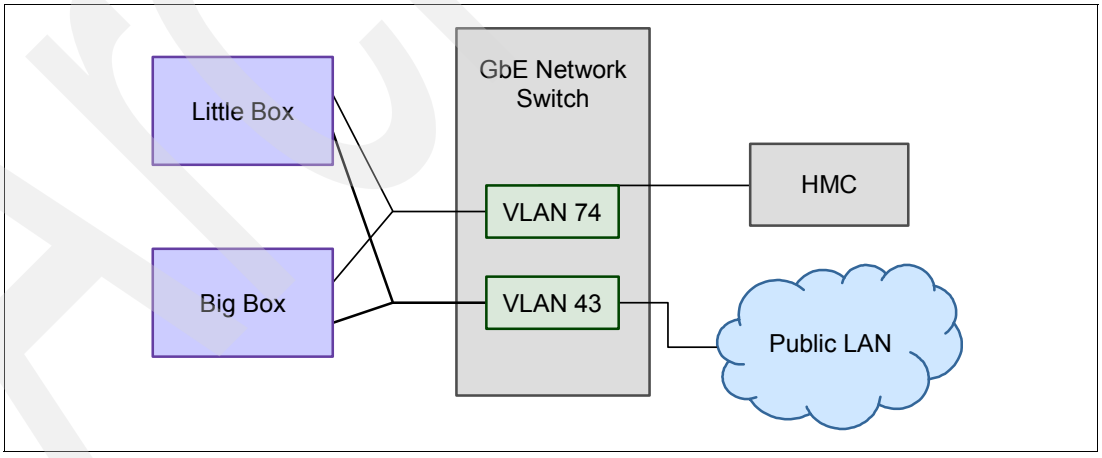

Figure 2-12 Relationship between the boxes and the network switches VLAN74 and VLAN43

2.3 Introduction to LPAR concepts

If you are already familiar with LPAR concepts and sharing I/O resources using VIO, skip this section and proceed to 2.4, “Creating LPARs” on page 17.

The initial tasks that we outlined in 2.1, “Hardware inventory using the HMC” on page 8 involved using the Hardware Management Console (HMC) to determine which resources were associated with our Little Box and Big Box.

After completing the initial HMC steps, we defined the logical partitions (LPARs) and installed the appropriate software for each LPAR, such as VIOs, the NIM, and more. These definitions established the access to the network, local storage (Small Computer System Interface 9SCSI)), and other devices using Fibre Channel (FC). In 2.4, “Creating LPARs” on page 17, we provide an overview of the LPARs that we created. We discuss each LPAR in more detail in Chapter 3, “Installing and configuring components” on page 23.

You need to understand the relationships between LPARs, PowerVM™ devices, and other concepts before implementing cloud components on POWER. Figure 2-13 through Figure 2-15 on page 17 provide a progression from simple to advanced topics.

Figure 2-13 depicts an LPAR in relation to the Hypervisor and the network, SCSI, and FC with which the LPAR might need to communicate to perform I/O operations. In Figure 2-13, the LPAR presents a request to the Hypervisor, which in turn facilitates or performs the I/O operation on behalf of the LPAR. Note that LPARs cannot perform I/O directly to the device. The Hypervisor needs to facilitate direct I/O access to the device.

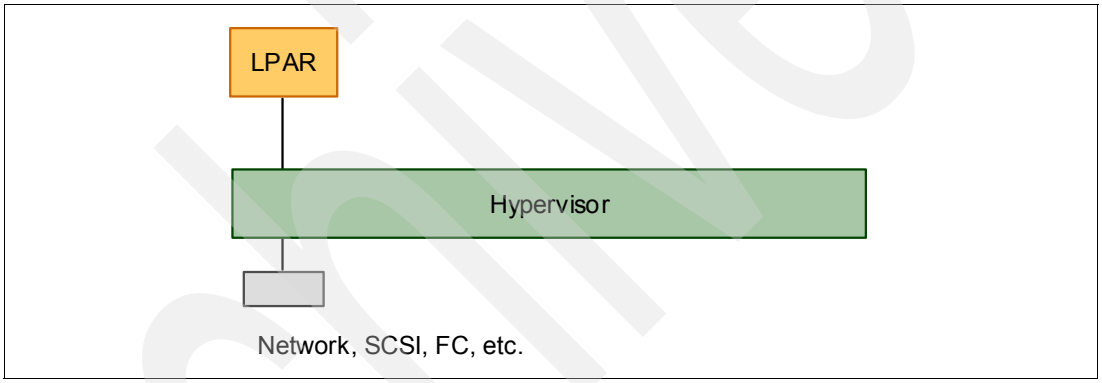

Figure 2-13 Single LPAR with access to network, SCSI, or FC

In contrast to Figure 2-13 on page 15, we add three more LPARs to our discussion, as shown in Figure 2-14. Note that LPAR 1 and 2 have the capacity to perform I/O operations to devices that are not shared. We show this capability by the solid lines from the LPARs to the Hypervisor and from the Hypervisor to the devices at the bottom of Figure 2-14. The dashed line represents LPAR 2 and LPAR 4 attempting to (directly) share a device. To facilitate this sharing, you must add an entity that supports the use of common or shared devices across LPARs. Figure 2-15 on page 17 shows this environment, in which LPAR 3 takes on the role of a Virtual I/O Server.

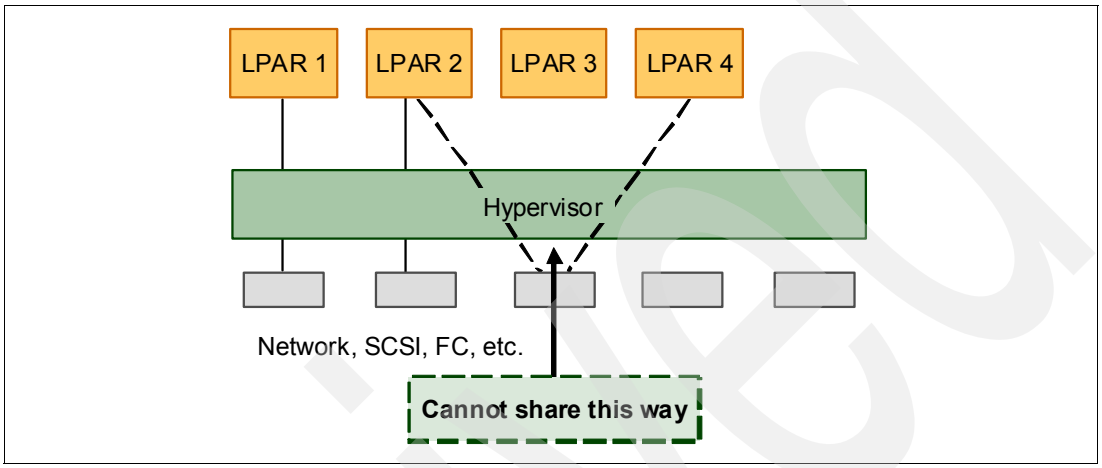

Figure 2-14 LPARs attempting to share access to Network, SCSI, or FC

In Figure 2-15 on page 17, you see that LPAR 2 and LPAR 4 share access to a device through LPAR 3 (the VIO LPAR). The donut-shaped objects in Figure 2-15 on page 17 represent virtual SCSI adapters. The VSCSI adapters that are shown in LPAR 2 and LPAR 4 are referred to as *VSCSI client adapters*. The adapters in the VIO LPAR are called *VSCSI server adapters*.

The role of the virtual I/O LPAR is to manage requests for operations and communications with devices that are shared across LPARs. Essentially (*simplified for the sake of brevity*) I/O requests by LPAR 2 and LPAR 4 (where the I/O is shared) are accomplished through a sequence of steps similar to the following list:

- The requesting LPAR (LPAR 2 or LPAR 4 in Figure 2-15 on page 17) presents a request for input or output through the VSCSI client adapters, channeled through the Hypervisor. Note that the relationship between LPAR 2 and the device and the relationship between LPAR 4 and the device need to be defined. We discuss these points further in Chapter 3, “Installing and configuring components” on page 23. The relationship between the VIO LPAR and the actual device needs to be defined. We provide the appropriate details in Chapter 3, “Installing and configuring components” on page 23.
- The Hypervisor facilitates the request that is reflected to the VIO LPAR through the corresponding VSCSI server adapter.
- The VIO server manages the requested I/O operation, scheduling or queuing the operation if the device is busy.
- When the device is not busy, the VIO server (LPAR 3 in Figure 2-15 on page 17) presents a request for input or output to the device via the Hypervisor.
- The actual I/O takes place and the requesting LPAR continues processing, as needed.

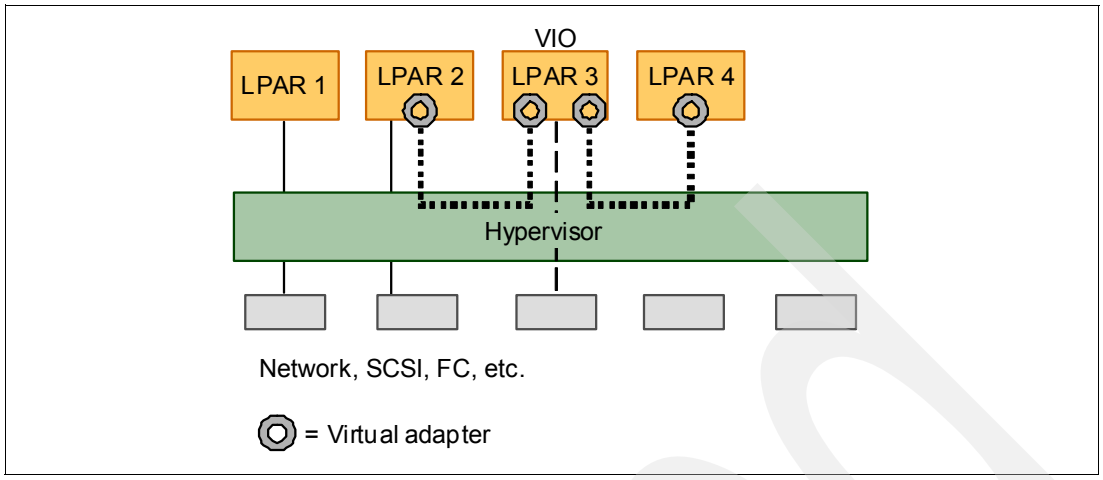

Figure 2-15 LPARs sharing access to Network, SCSI, or FC using virtual SCSI via VSCSI adapter

In 3.1.4, “Preparing storage for client LPARs” on page 40, we discuss the use of commands, such as **mklv** and **mkvdev**, on the VIO server. Those discussions demonstrate the tasks that need to be completed within the realm of the VIO (or LPAR 3) so that it can support virtual devices that appear real to the requesting LPARs (LPAR 2 and LPAR 4 in our example).

On the VIO server, you use **mklv** to make a logical volume and then use **mkvdev** to make a virtual device. Later, when you create the profile for LPAR 2, you use the associated dialog as discussed in Chapter 3, “Installing and configuring components” on page 23.

2.4 Creating LPARs

Figure 2-16 on page 18 through Figure 2-20 on page 20 depict the order of the creation of the LPARs in this implementation. This order is important, because as each LPAR was created and configured, it became part of a collective and growing scaffolding of function and capabilities. It eventually grew to a point where a functioning Managed Through component model existed.

We created the LPARs on the Little Box in this order:

- VIOS
- NIM
- IBM Systems Director
- Tivoli Productivity Center
- Common Information Module (CIM)

The first LPAR created was a VIO server LPAR. Creating the VIO server LPAR first creates an environment that permits sharing I/Os. We installed the VIOS LPAR in Figure 2-16 on page 18 without the aid of any automation or related services. You can read about the details associated with this process in 3.1, “Installing virtual I/O” on page 24. For now, remember that we accomplished these tasks using physical media and manually interacting with the HMC dialog. After creating the first two LPARs, subsequent tasks that are associated with the creation of LPARs can use the first two LPARs.

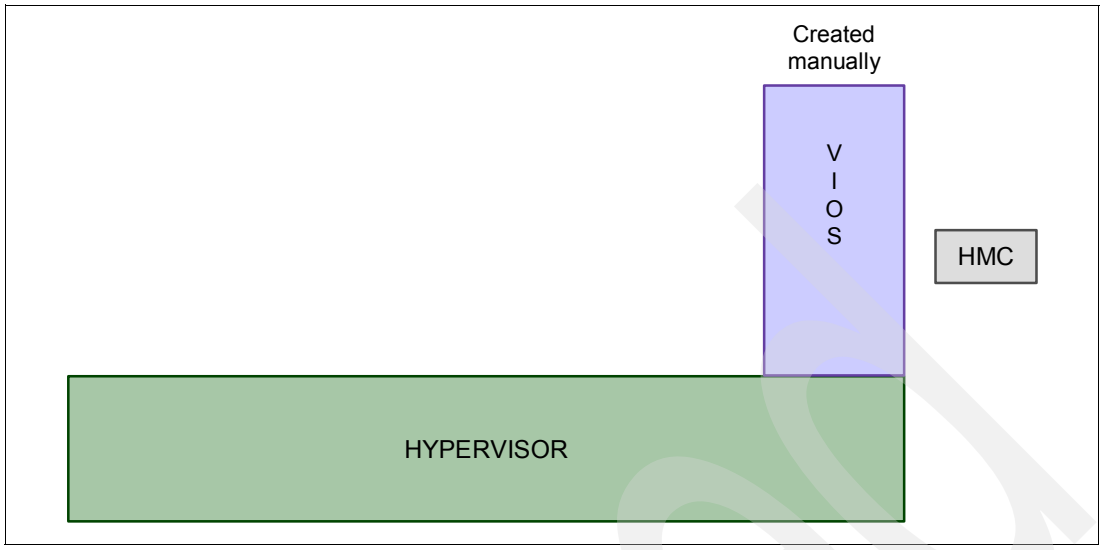

Figure 2-16 The VIOS LPAR was created and VIOS was installed manually

Figure 2-17 shows that the next LPAR (labeled NIM) was also installed without the aid of any automation or related services. With the addition of this NIM LPAR, other installations can take advantage of the functionality provided by the NIM. Therefore, manual installations are no longer necessary, as Figure 2-18 on page 19 demonstrates. The IBM Systems Director LPAR is provided in part through the power of NIM.

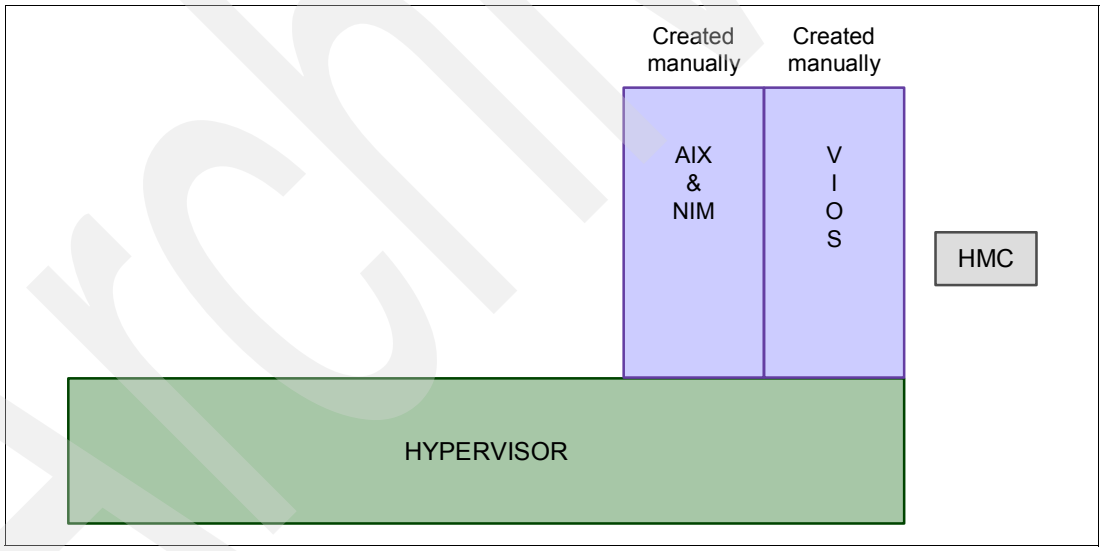

Figure 2-17 The NIM LPAR was created and NIM was installed without the aid of automated services

In Figure 2-18 on page 19, note that AIX and Systems Director can communicate with the NIM LPAR, as well as with the VIOS LPAR and HMC. From this point forward, IBM Systems Director can be used to perform provisioning of other LPARS. Essentially, a Managed Through environment has evolved.

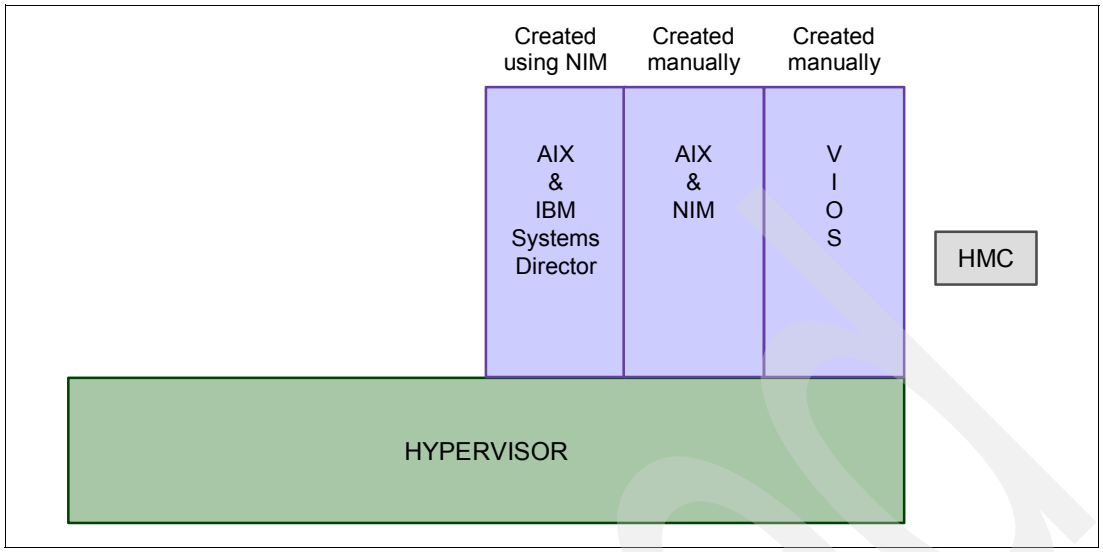

Figure 2-18 Installation of IBM Systems Director was accomplished using NIM

Next, as shown in Figure 2-19, the TPC LPAR is created. This LPAR has the capacity to interact with the TPC storage resource agent on the VIOS LPAR. The AIX & TPC LPAR was created using IBM Systems Director.

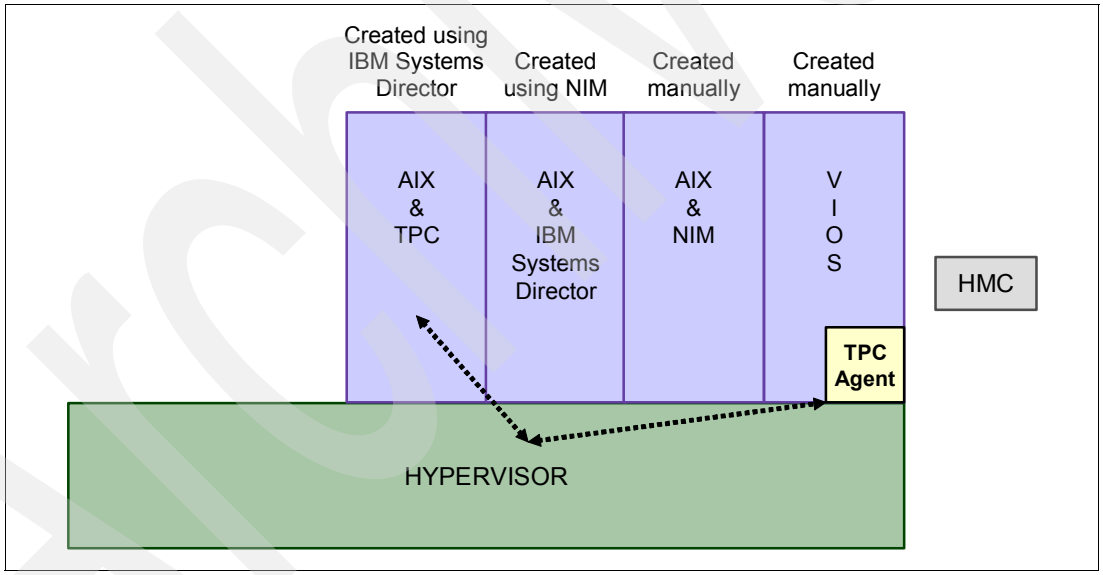

Figure 2-19 The TPC LPAR is created through the IBM Systems Director

Finally, in Figure 2-20, the AIX & CIM LPAR appears. This LPAR was created through the facilities of the IBM Systems Director.

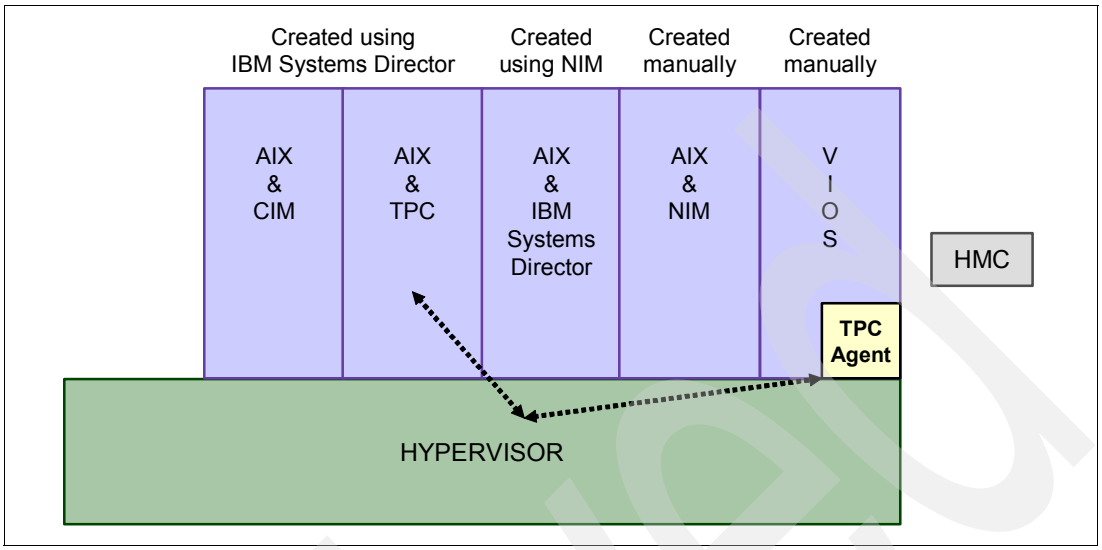

Figure 2-20 The CIM LPAR is created through the IBM Systems Director

At this point, further LPAR creation can be facilitated by the IBM Systems Director. Therefore, a Managed Through capability now exists, as shown in Figure 1-1 on page 2.

2.5 Network-attached storage

For this project, a Network File Server (NFS) already existed and was accessible using network-attached storage (NAS).

Having this universally accessible data store greatly simplified sharing files among nodes and LPARs. The NFS served as a repository for installation files for AIX and similar collateral required for VIOS, CIM Storage Management Initiative Specification (SMI-S) providers, and IBM Systems Director.

It is important to realize that an NFS is not an absolute requirement. But by having an NFS, many installation tasks can be simplified. For other implementation projects, you can substitute whatever file store or file transfer method appears be most advantageous.

2.6 Network configuration

Figure 2-21 on page 21 depicts a simplified network diagram of the environment that was built. There are three main components that need network connectivity. In order for the HMC to properly administer the Little Box and Big Box, a private network needed to be created between the components. To create this network, we created a virtual LAN (VLAN) numbered 74. This VLAN was configured on the switch along with VLAN 43, which was the general purpose network. VLAN 43 also connects to the public network connecting the Managed From to the Managed To.

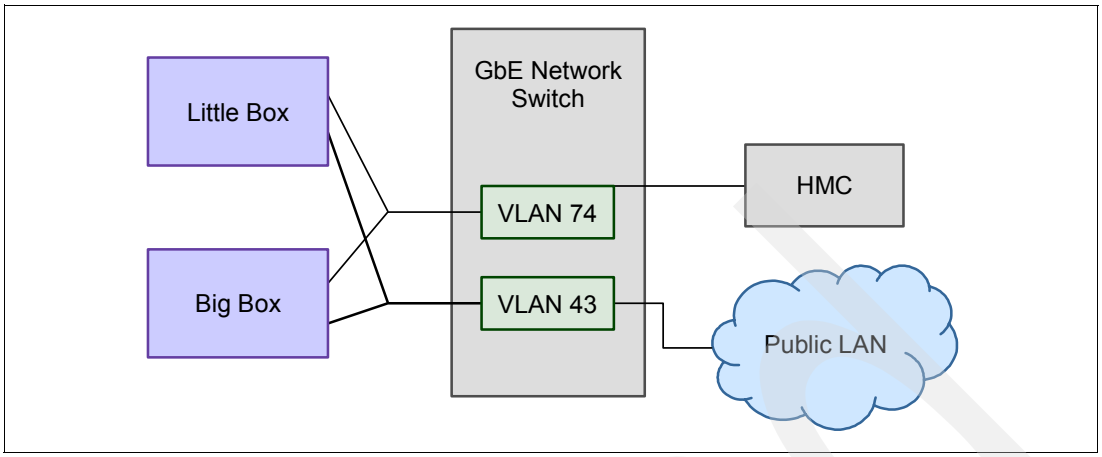

Figure 2-21 Relationship between the boxes and the network switches VLAN 74 and VLAN 43

2.7 Storage-attached storage

We used a fairly simple storage-attached network (SAN) layout for this project. As shown in Figure 2-22 on page 22, a DS4800 serves as the storage subsystem connected to a Brocade SAN switch. Both the Little Box and the Big Box plug directly into a Brocade SAN switch.

Each node (four nodes on the Big Box and one node on the Little Box) has multiple Fibre Channel (FC) adapters. Most nodes had multiple ports, as shown in the inventory described earlier, starting with Figure 2-5 on page 12. In cabling the Big Box, an effort was made to ensure that multiple adapters, which connected to each VIOS, were properly cabled into the switch.

Because there were fewer cables than FC ports, it seemed preferable to cable at least one port on more adapters rather than every port on a given adapter. As we describe in Chapter Chapter 3, “Installing and configuring components” on page 23, this design provided the necessary paths to the VIOS LPAR.

The storage subsystem was not dedicated to this POWER environment. A few logical unit numbers (LUNs) were previously created in support of a set of unrelated servers connected to the same SAN switch. In order to accommodate this configuration, we created zones manually on the SAN switch to isolate the other environment and its associated storage from our POWER environment.

A dedicated environment can save the additional zoning steps that we took, but those steps were a small one-time effort performed by our IT support staff.

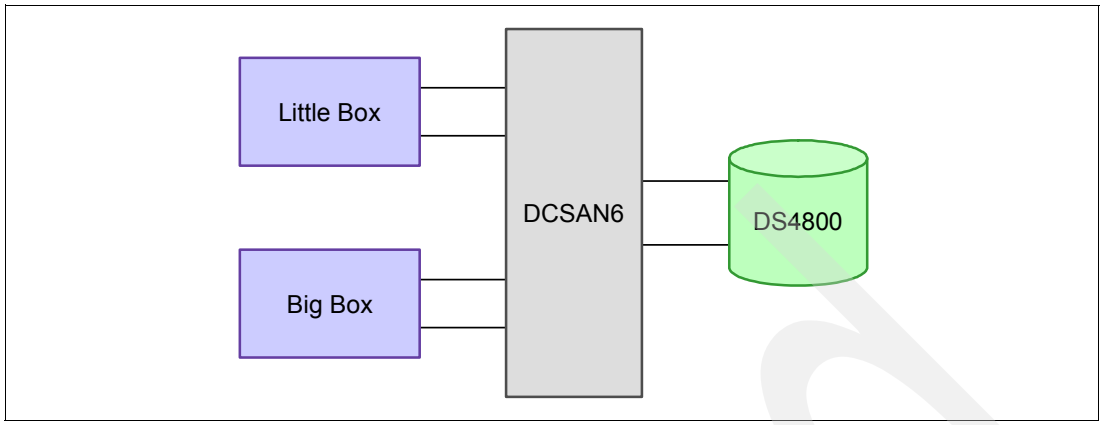

Figure 2-22 Relationship between the boxes and DS4800 storage accessed through DCSAN6

3

Installing and configuring components

We created and implemented the logical partitions (LPARs) on the Little Box in the following order:

- Virtual I/O Server (VIOS)
- Network Installation Management (NIM)
- IBM Systems Director
- CIM
- Tivoli Storage Productivity Center

This chapter provides a more detailed view of the tasks and steps that are involved in creating the LPARS and deploying cloud components in their respective LPARs. Each successive installation contributed to an evolving infrastructure that became part of a collective and growing scaffolding of function and capabilities, eventually growing to a point where a functioning *Managed Through* component model existed.

3.1 Installing virtual I/O

Virtual I/O (VIO) allows LPARs to exchange information without storing it on external devices, thus providing an increase in performance.

3.1.1 Creating a partition profile for the VIO LPAR

We used the Create LPAR Wizard, as shown in Figure 3-2 on page 25, to create a VIO server LPAR. As a result, LPARs were able to share access to physical I/O resources, as depicted in Figure 3-1.

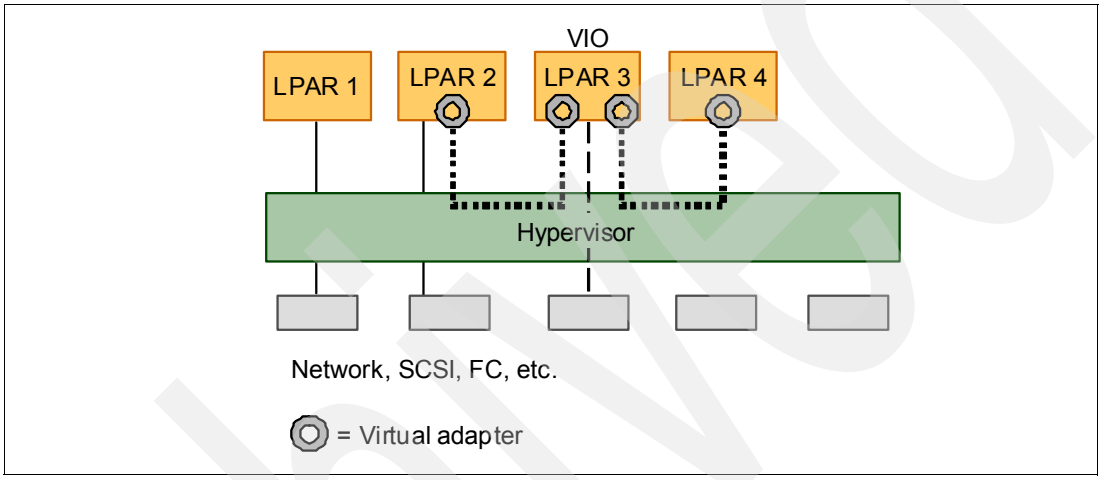

Figure 3-1 LPARs sharing access to network, Small Computer System Interface (SCSI), or Fibre Channel (FC)

We create a partition profile to describe the LPAR labeled LPAR3 and VIO in Figure 3-1 by specifying partition attributes using the Create LPAR Wizard. Before discussing the sequence of these specifications, it is appropriate to have a quick preview of the end result, as shown in Figure 3-2 on page 25.

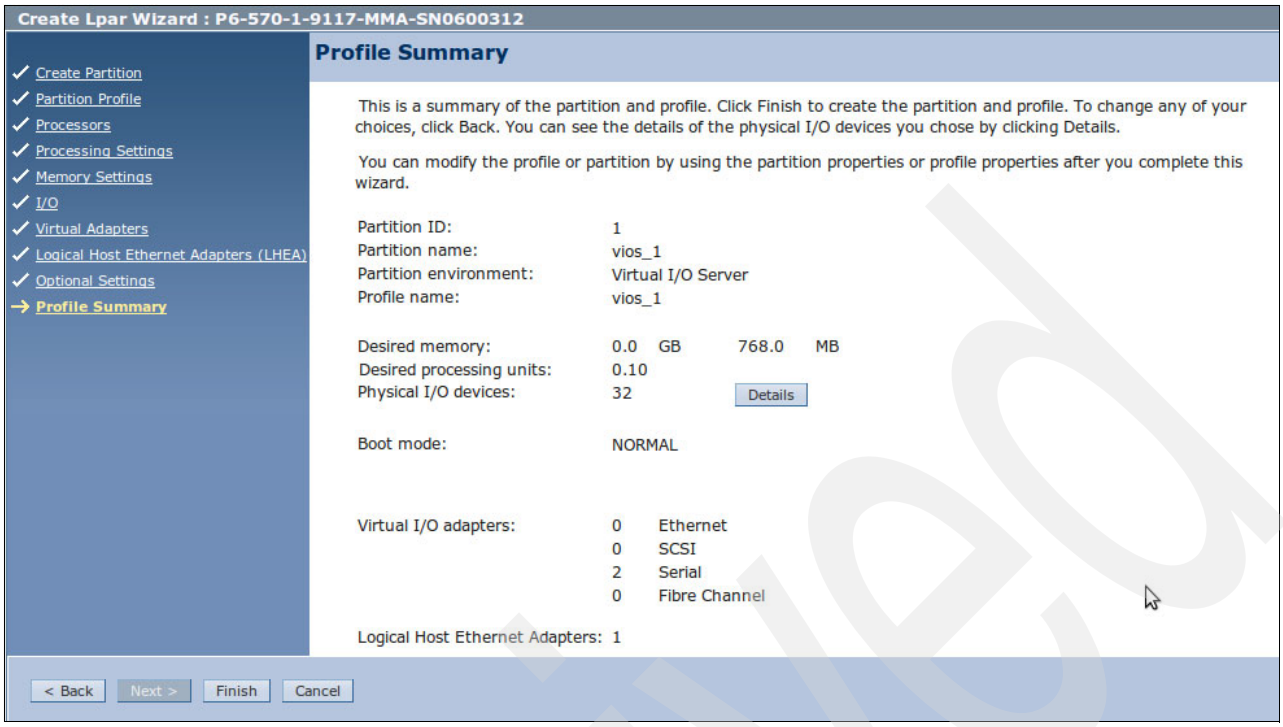

Figure 3-2 Profile Summary

The preceding Profile Summary is for the first profile that was created. Its Partition ID is 1. We specified the name vios_1 for the profile. Look at the memory and processing requirements. The boot mode and the virtual I/O adapters that are shown represent the information that was specified to the Create LPAR Wizard.

Figure 3-3 on page 26 and Figure 3-4 on page 26 show the Create LPAR Wizard windows for the specification of the system name, partition name, partition ID, and profile name. In Figure 3-4 on page 26, the “Use all the resources in the system” box is not selected. Selecting Finish in each of these windows continues the Create LPAR Wizard specification process.

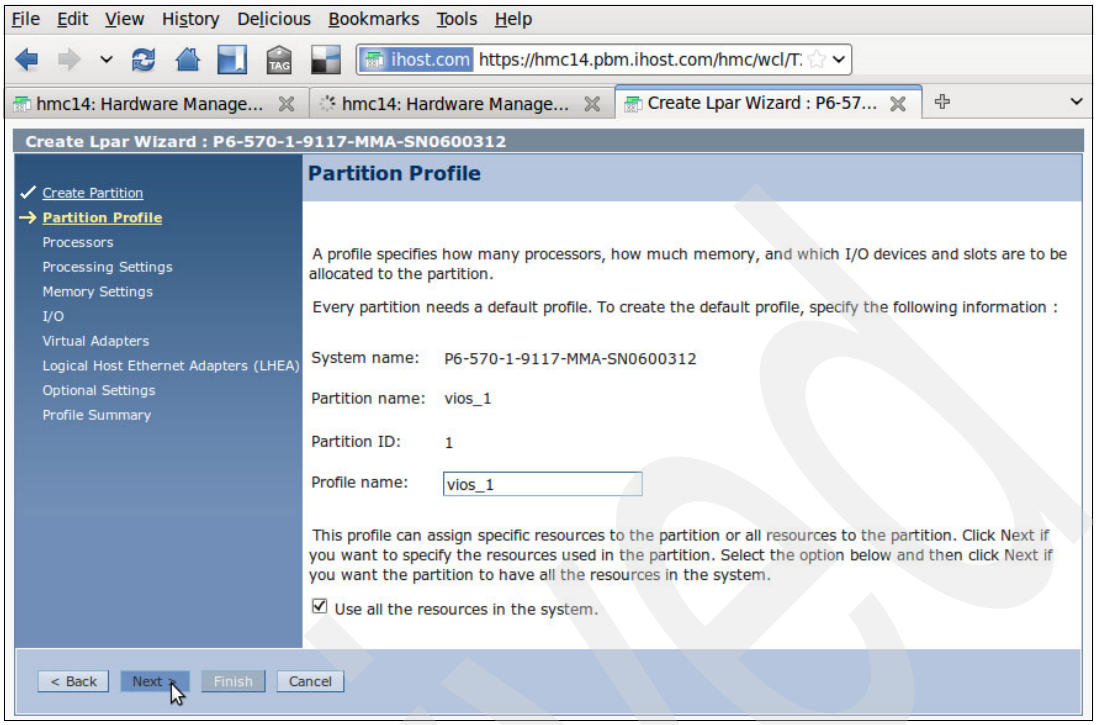

Figure 3-3 Specify Profile name

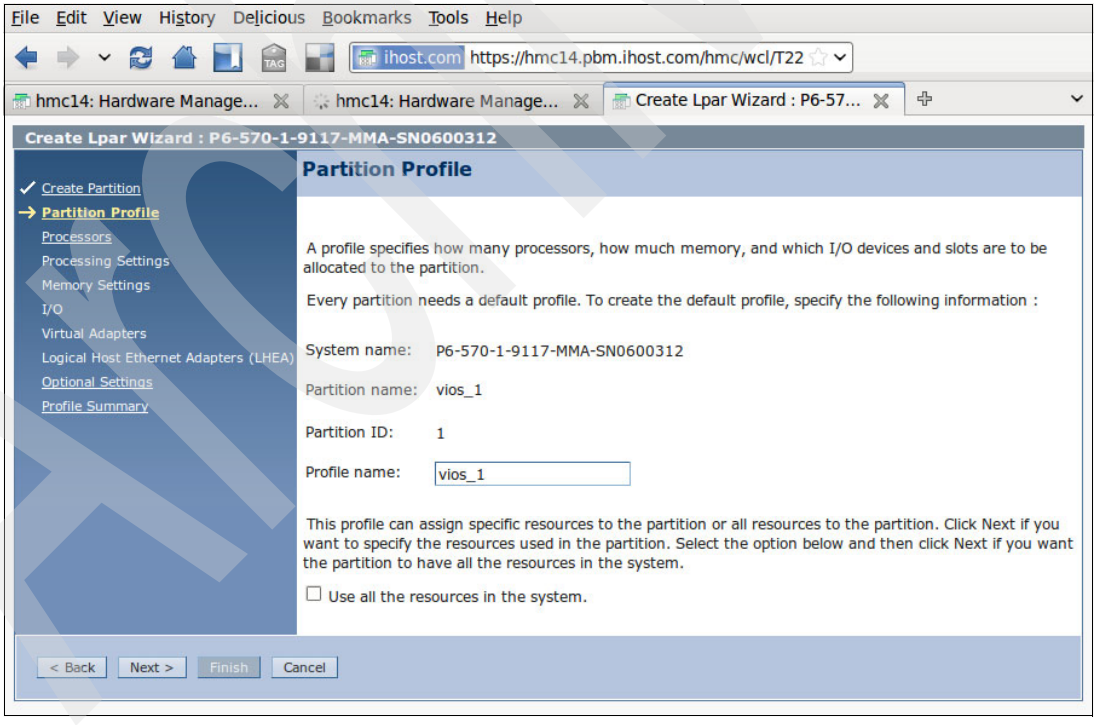

Figure 3-4 Creating a partition profile without using all system resources

We provide the specifications, as shown in Figure 3-5 on page 27 and Figure 3-6 on page 28, in the context of shared versus dedicated processors. If the processor is shared, the window shows the degree to which the processor is shared:

- Shared processor pool

- Minimum processing units
- Desired processing units
- Maximum processing units
- Virtual processor metrics:
  - Minimum virtual processors
  - Desired virtual processors
  - Maximum virtual processors
  - Uncapped weight

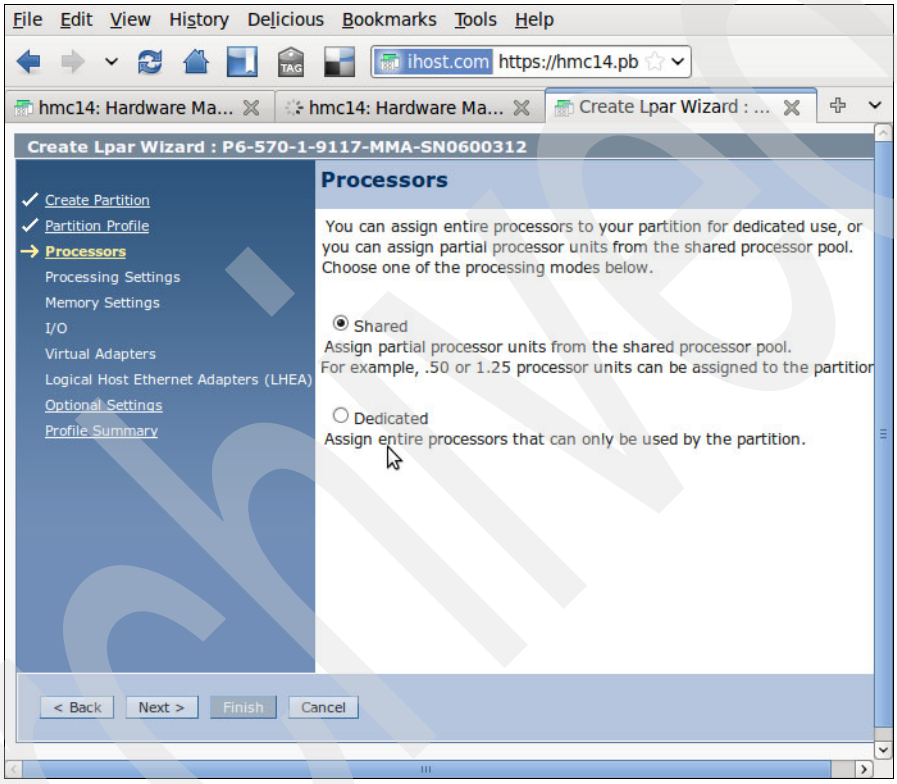

Figure 3-5 Create LPAR Wizard and processor specifications

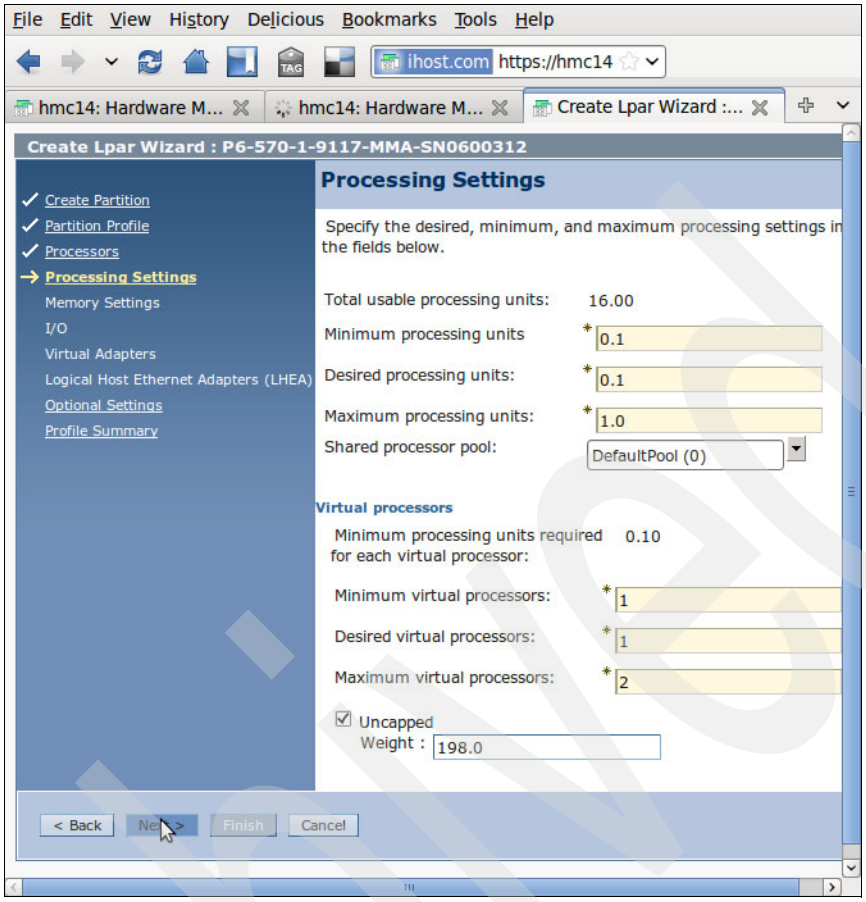

Figure 3-6 Create LPAR Wizard and processor settings

Figure 3-6 shows the details for the processor settings. Figure 3-6 indicates how many processors or how many processing units in the shared processor pool are to be assigned to this LPAR.

In Figure 3-7 on page 29, we specify the amount of memory to be allocated to this virtual machine:

- Minimum memory = 512 MB
- Desired memory = 768 MB
- Maximum memory = 1 GB

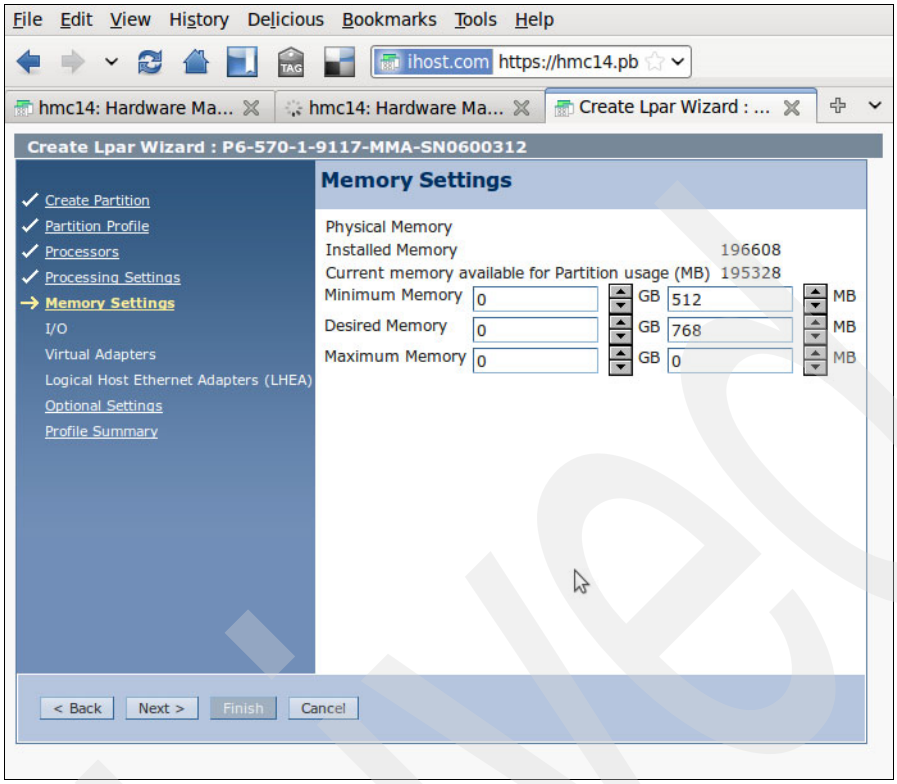

Figure 3-7 Create LPAR Wizard and Memory Settings

Before discussing how to specify information about physical and virtual adapters, you must understand the needs of the LPAR in the context of the physical adapters and how they are assigned. In addition, you must consider your virtual adapters needs so that they can be created (defined) and associated with an LPAR.

But first, consider what physically existed in this implementation, as described in the text that is associated with Figure 3-8 on page 30.

Figure 3-8 Selecting I/O to be associated with this LPAR

Figure 3-8 demonstrates the selection of I/O to be associated with this LPAR. The location code column identifies the location of the I/O line item relative to the annotated image in Figure 3-9.

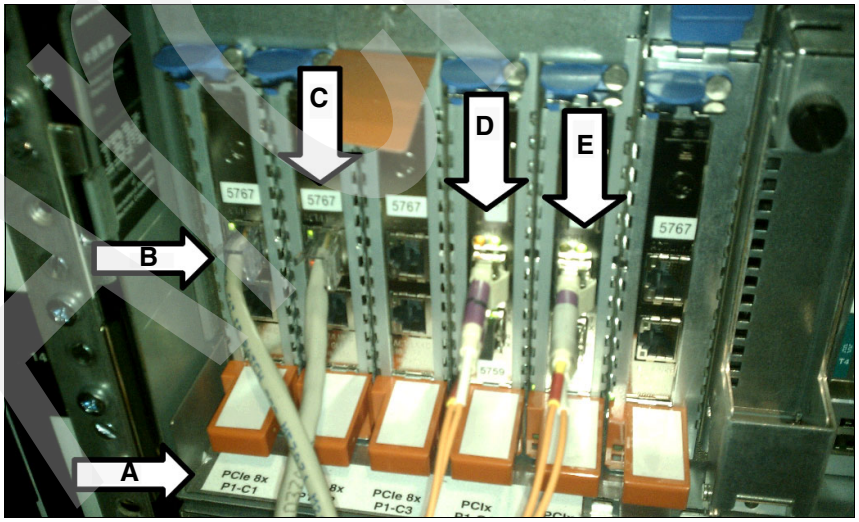

Figure 3-9 Annotated photograph of node DQD08G0

Note that the location code in Figure 3-8 can be related back to the inventory of cables and connectors and repeated, as shown in Figure 3-10 on page 31. In the first eight lines of the window that is shown in Figure 3-8, the DQD08G0 refers to a node in the Big Box that is described in Figure 3-10 on page 31.

The P1-T1 and P1-T3 do not appear in Figure 3-9 on page 30, because the USB and RAID controller connections are internal to the box. The P1-C4 and P1-C5 correspond to the P1-C4-T1 and P1-C5-T1 or arrows D and E respectively. These are Dual Port Fibre Channel dynamic device reconfiguration (DDR) adapters.

DQD08G0 contains four Ethernet controller cards. Two of them have Ethernet cables plugged in at locations B and C, as shown in the picture of the back panel (Figure 3-9 on page 30). They are listed in the inventory that is shown in Figure 3-10 as slots P1-C1-T1 and P1-C2-T1. Clearly, having an inventory of slot and port usage available and documented before getting to this point is helpful.

4 NODE SYSTEM:		
DQD08G0		
SLOT	IEEE ADDR	
P1-C1-T1	00-14-5E-74-89-AC	
P1-C1-T2	00-14-5E-74-89-AD	
P1-C2-T1	00-14-5E-74-89-B6	
P1-C2-T2	00-14-5E-74-89-B7	
P1-C3-T1	00-14-5E-74-89-B8	
P1-C3-T2	00-14-5E-74-89-B9	
P1-C4-T1	00-00-C9-74-E6-A8	SAN6 PORT 20
P1-C4-T2	00-00-C9-74-E6-A9	
P1-C5-T1	00-00-C9-76-40-D6	SAN6 PORT 01
P1-C5-T2	00-00-C9-76-40-D7	
P1-C6-T1	00-14-5E-74-86-DE	
P1-C6-T2	00-14-5E-74-86-DF	

Figure 3-10 Cabling and connections for node DQD08G0 (one of four nodes in the Big Box)

After completing this process, it is time to define the virtual adapters. The correlation between the client ID, virtual adapters, VIO server adapter IDs, and actual physical devices can be quite complex. In Figure 2-10 on page 14 in Chapter 2, the SLOT column shows physical addresses for the Little Box. These addresses have three components: type of address, card number, and port number. These devices logically attach to the VIO server. For instance, in this implementation, the Ethernet Adapter P1-C1-T1 was attached to the VIO server that is represented as `ent0`. To verify this correlation, use the **`lsdev`** command, as shown in Example 3-1. Note the bold **`ent0`** and **`P1-C1-T1`**.

Example 3-1 lsdev command to verify device correlation

```
$ lsdev  -dev ent0 -vpd
  ent0             U789D.001.DQD08H9-P1-C1-T1  2-Port 10/100/1000 Base-TX
PCI-Express Adapter (14104003)
2-P NIC-TX PCI-e:
        EC Level.....................D75975
        Part Number.................10N6845
        Manufacture ID..............YL1026
        FRU Number..................10N6845
        Network Address.............00145E747E54
        ROM Level.(alterable).......EP0140
        Hardware Location Code......U789D.001.DQD08H9-P1-C1-T1
PLATFORM SPECIFIC
Name:  ethernet
    Node:  ethernet@0
    Device Type:  network
    Physical Location: U789D.001.DQD08H9-P1-C1-T1
```

Example 3-1 on page 31 deals with a network device. The VIO can now present that device to the client LPARs, which we describe in detail in 3.1.3, “Preparing networks for the client LPAR” on page 40.

Virtualizing storage adapters with VIO increases the complexity of this correlation by adding multiple new layers of abstraction. For instance, logical volume management enables the partitioning of the storage to be used by the client LPARs without relying on the actual physical adapters or physical disks. The internal storage of the Little Box is used for part of the Managed Through components. Using **`lvm`**, we can partition the storage for use by the components.

All the components of the Managed Through have an ID that corresponds to the LPAR on which they reside. Each LPAR also has a set of IDs for its virtual adapters. Each of these client virtual adapters aligns with a virtual adapter of the VIO.

The VIO server on the Little Box has a client ID of 1. In 3.2.3, “Creating the partition profile for the NIM LPAR” on page 46, another LPAR is created for NIM. This LPAR has a client ID of 2. All of the internal disks and their associated physical adapters are assigned to the VIO. In order for NIM and any other LPAR to be able to use the internal storage, you must configure the logical volume management (*lvm*) on the VIO server, which introduces another set of IDs. Each LPAR (or client) is given a logical volume (*lv*) from the volume group (*vg*) that is created later. When creating the Partition Profile for these LPARs, a virtual adapter is generated for this LPAR and connected to the logical volume. Figure 3-11 shows the relationship between the VIO adapters and the client.

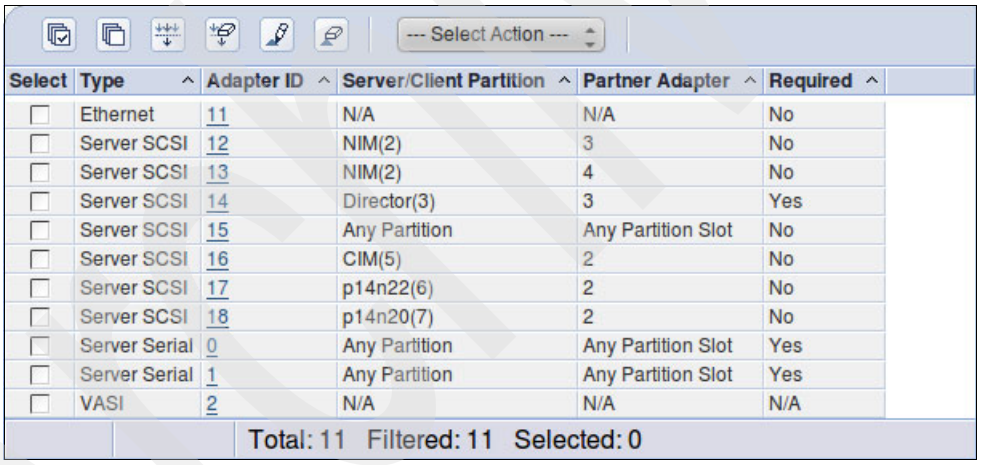

Figure 3-11 View of the relationship between client and server adapters

Figure 3-11 is the VIO server’s view of the relationship between the client and server adapters. To continue the example using NIM, the second row shows that NIM(2), where 2 in parentheses is the Client Partition ID, has an adapter with an ID of 12.

This server is partnered with the VIO adapter (Server SCSI) with the ID of 12. The ID 12 corresponds to the “physical” location code within the VIO server. To verify this relationship on the VIO server, use **`lsmap`**. The **`lsmap`** option of **`-plc`** requests the full physical location code, which is generated from the machine ID and the adapter ID. In this case, the physical location code is U9117.MMA.06E3531-V1-C12, where C12 is the adapter ID. Example 3-2 shows the lsmap for U9117.MMA.06E3531-V1-C12.

Example 3-2 lsmap for U9117.MMA.06E3531-V1-C12

```
$ lsmap -plc U9117.MMA.06E3531-V1-C12
SVSA            Physloc                                     Client Partition ID
```

```
-------------- ------------------------------------------ ------------------
vhost0         U9117.MMA.06E3531-V1-C12                   0x00000002

VTD                   vnimsrv_dvg
Status                Available
LUN                   0x8200000000000000
Backing device        nimsrv_dvg
Physloc
Mirrored              N/A

VTD                   vnimsrv_rvg
Status                Available
LUN                   0x8100000000000000
Backing device        nimsrv_rvg
Physloc
Mirrored              N/A
```

The output of the **lsmap command** in Example 3-2 on page 32 shows the physical location, which contains the VIO adapter ID and the client ID. It also shows the name of the virtual adapter that creates this relationship: vhost0. Vhost0 has two backing devices: nimsrv_dvg and nimsrv_rvg. Both backing devices are actually logical volumes.

When defining virtual adapters using the process that is shown in Figure 3-12 on page 34 and Figure 3-13 on page 35, the maximum number of virtual adapters is specified. On the window that is shown in Figure 3-12 on page 34, clicking Actions (top center of window) presents a dialog, which is shown in Figure 3-13 on page 35.

We type 13 as the Virtual SCSI Adapter ID. You might ask, “Where did the 13 come from?” After setting up the LPAR and installing VIO, which included the use of the **mklv** command, we issued the **lsmap** command. The output of the **lsmap** command provides the number that is needed (13 in this case), as shown in Example 3-3. Note that the line that begins with `vhost1` also contains the text `U9117.MMA.06E3531-V1-C13`. The 13 at the end is the number that is specified as the adapter number. Note that this line also contains `0x00000002`. Because this sample **lsmap** command sequence was issued after defining the client Virtual SCSI Adapter (see Figure 3-12 on page 34), this value represents the client adapter ID.

If the **lsmap** command were issued prior to using the GUI to define the adapter, the value on that line will be `0x00000000` instead. See Example 3-3.

Example 3-3 lsmap for virtual adapter vhost1

```
$ lsmap -vadapter vhost1
SVSA Physloc Client Partition ID
-------------- ------------------------------------------ ------------------
vhost1          U9117.MMA.06E3531-V1-C13                   0x00000002

VTD vcd
Status Available
LUN 0x8100000000000000
Backing device cd0
Physloc U789D.001.DQD08H9-P4-D1
Mirrored N/A

VTD vtopt0
Status Available
LUN 0x8200000000000000
```

```
Backing device /var//VMLibrary/6.1
Physloc
Mirrored N/A
```

In Figure 3-12, note that the maximum number of virtual adapters is specified. This restriction minimizes contention or any attempts to overutilize a physical adapter. By clicking the **Actions** pull-down list, you can choose the “Create Virtual SCSI adapter” option, resulting in the dialog that is shown in Figure 3-13 on page 35.

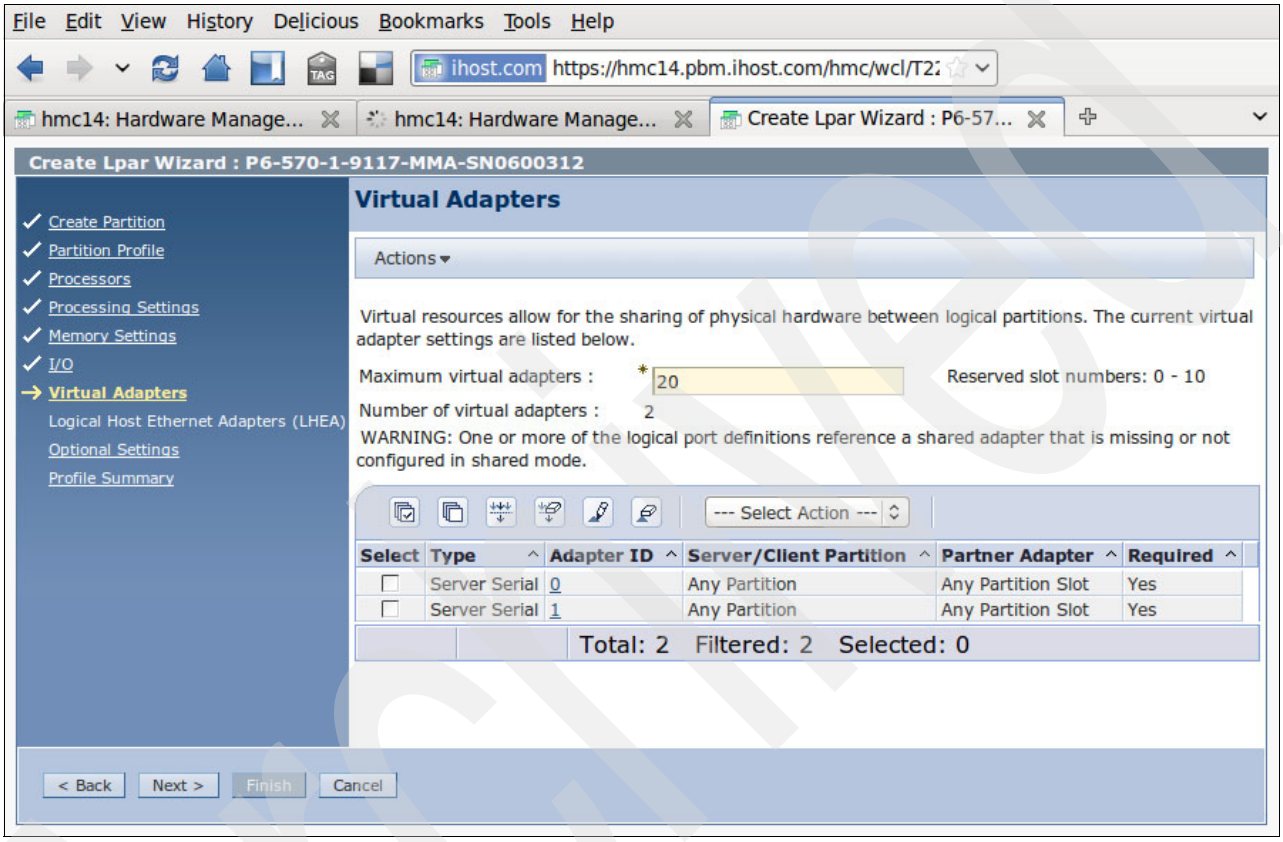

Figure 3-12 Virtual Adapters settings

In the window that requests the Create Virtual SCSI Adapter task (Figure 3-13), which, in this case, is for a Server Adapter, click **Any client partition can connect**. Later in this document, we use the same dialog to create a Virtual SCSI Adapter (Figure 3-14). At that point, the VSCSI Server Adapter can be associated with the VSCSI Client Adapter. Type the Client Adapter ID that you got from the output of an **`lsmap`** command that was issued on the VIO LPAR.

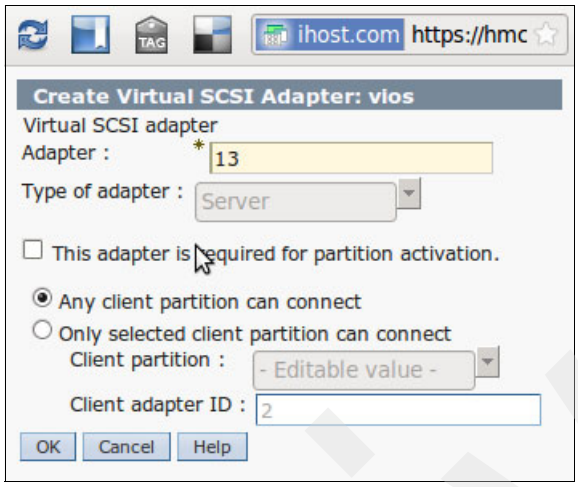

Figure 3-13 Create Virtual SCSI Adapter

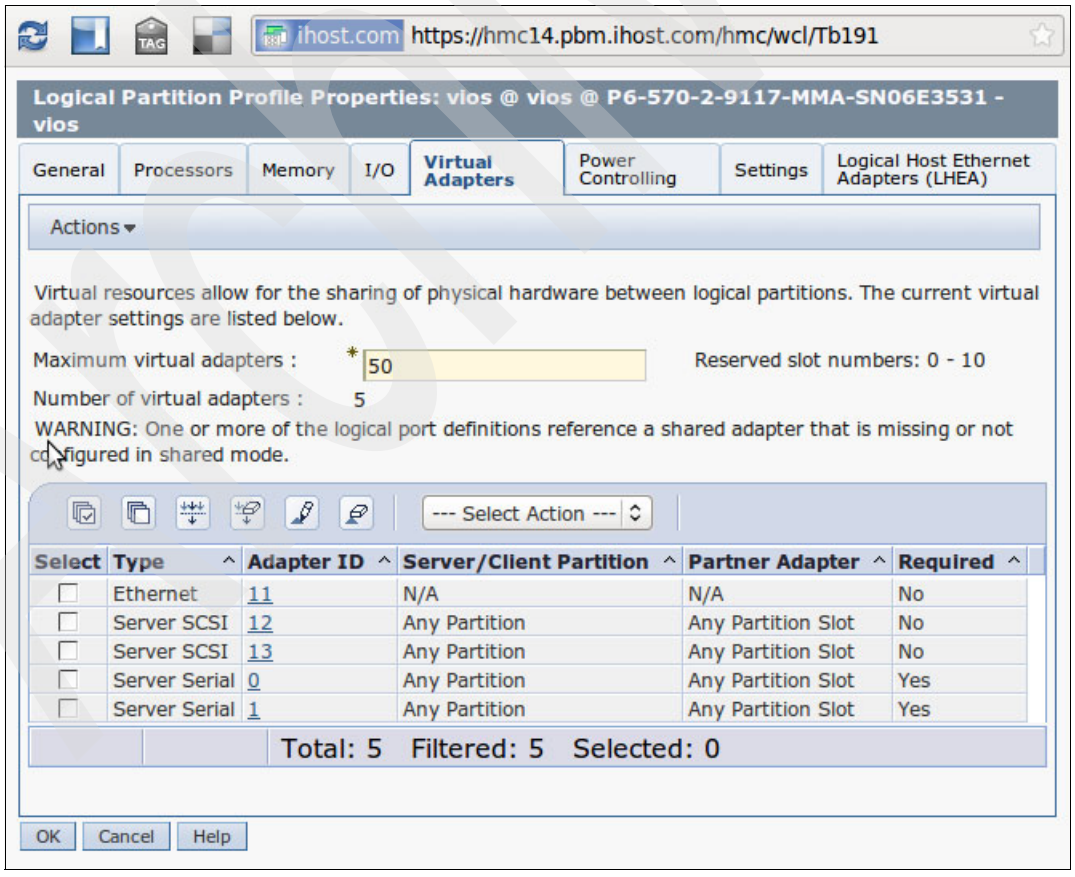

Figure 3-14 Logical Partition Profile Properties Virtual Adapters tab

We did not select Logical Host Ethernet adapters. Instead for this implementation, we selected **Next** at this point to continue to the Optional Settings window, as shown in Figure 3-16 on page 37.

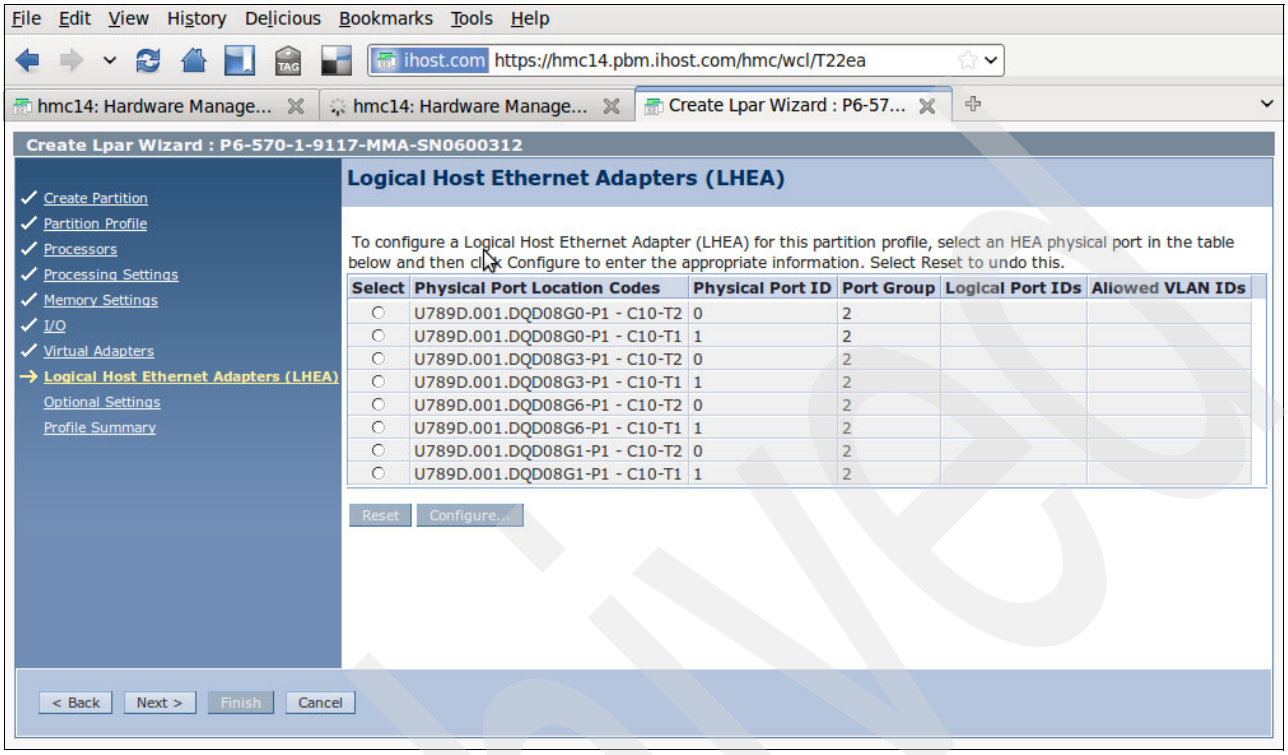

Figure 3-15 Logical Host Ethernet Adapters window was not used

In Figure 3-16, we specified **Automatically start with managed system**. We also selected **Normal** for Boot modes.

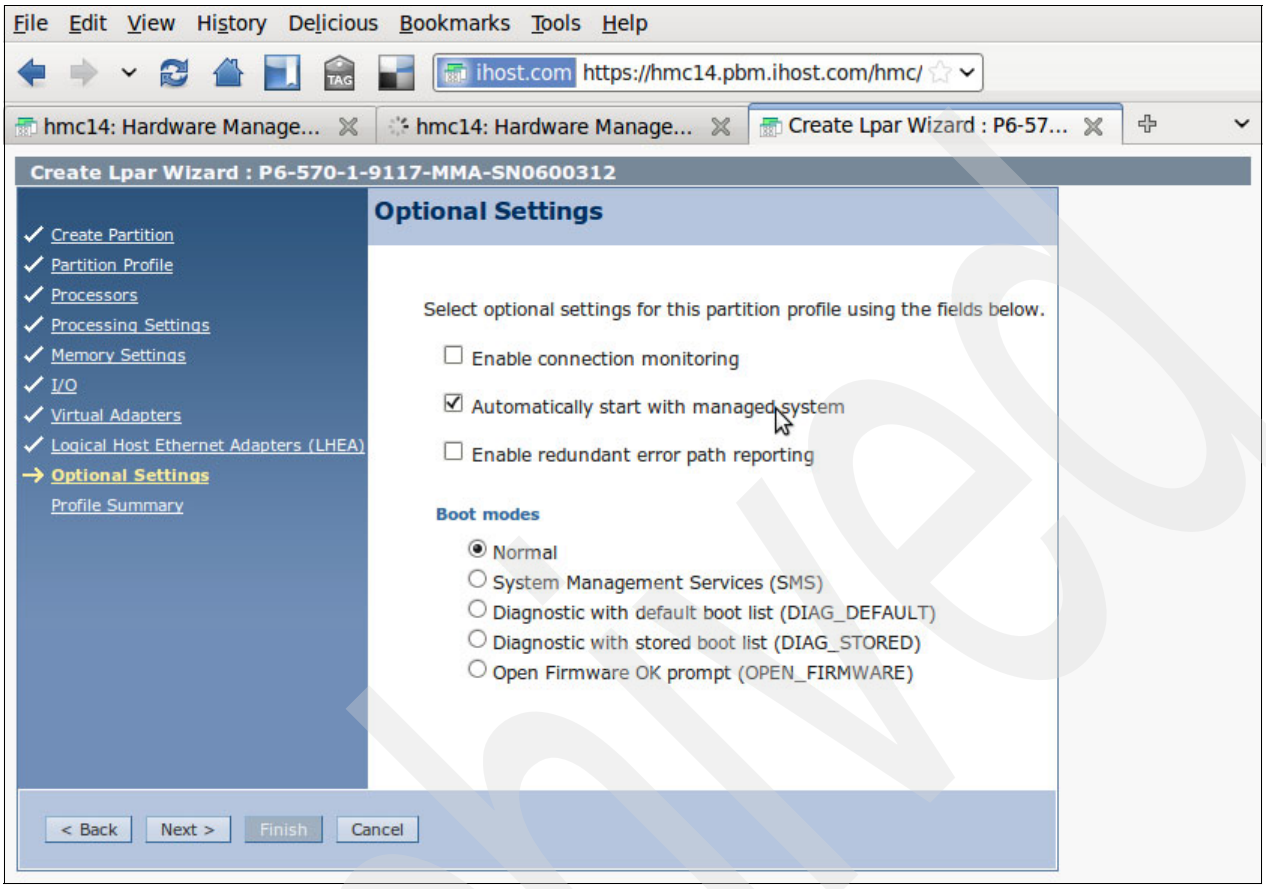

Figure 3-16 Choosing Automatically start with managed system

Figure 3-17 shows the Profile Summary window. Clicking **Finish** ends this phase.

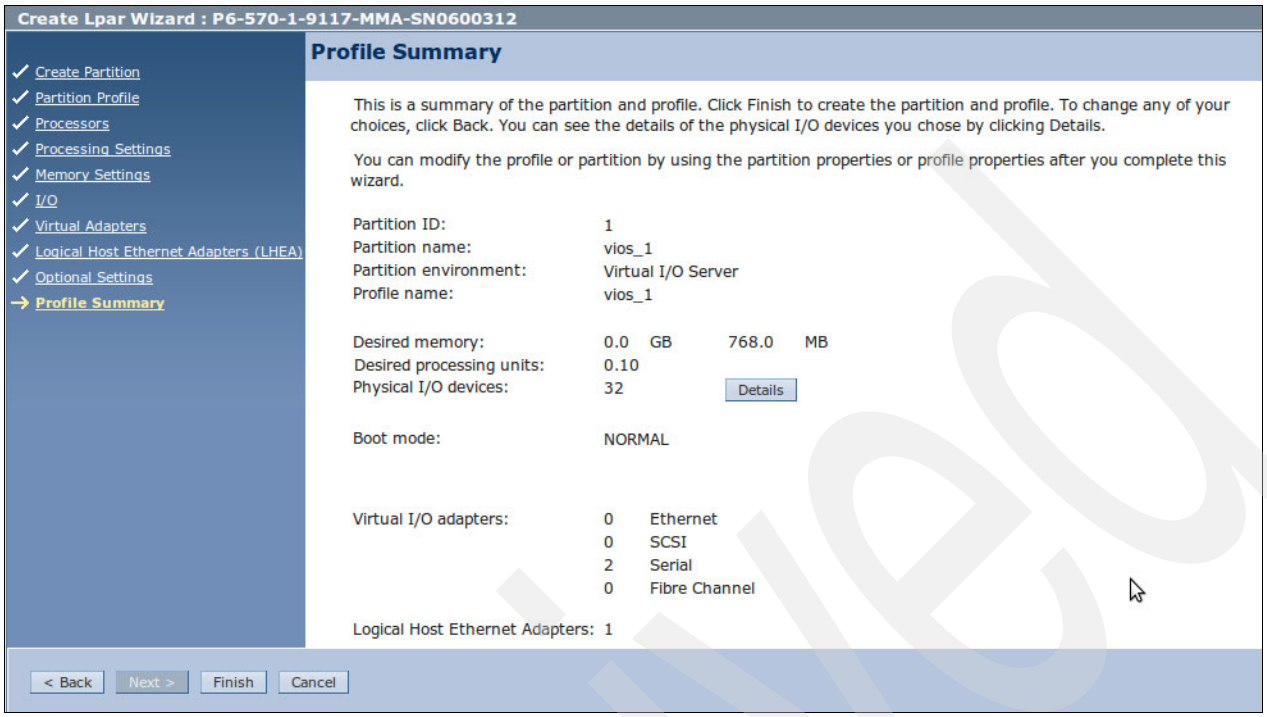

Figure 3-17 The Profile Summary window

3.1.2 Installing VIO server in the LPAR

After creating the partition profile for the VIO server LPAR, the next step is the installation of a VIO server. we used the System Management Services, which provided menus with numbered options, as shown in the sample Main Menu window in Figure 3-18 on page 39.

The following list contains the responses for this paper's implementation to the series of windows up to the point where there is a prompt to confirm exiting System Management Services to begin starting the VIO in the LPAR:

1. We chose "Select Boot Options" by typing menu item number 5 at the lower left of the window.
2. We chose "Select Install/Boot Device" by typing 1 at the bottom of the window.
3. We used "N" to scroll through multiple windows to locate the value, in this case, 27, that was associated with the device to be the "install from" device. Then, we typed that value (27) at the bottom of the window.
4. We chose "CD/DVD" by typing 3 at the bottom of the window.
5. We chose "SATA" by typing 5 at the bottom of the window.
6. We chose "a Media Adapter" by typing 1 at the bottom of the window.
7. We chose "SATA CD ROM" (the only device that was listed) by typing 1 at the bottom of the window.
8. We chose "Normal Mode Boot" by typing 2 at the bottom of the window.
9. We chose "Yes" to exit System Management Services by typing 1 at the bottom of the window.

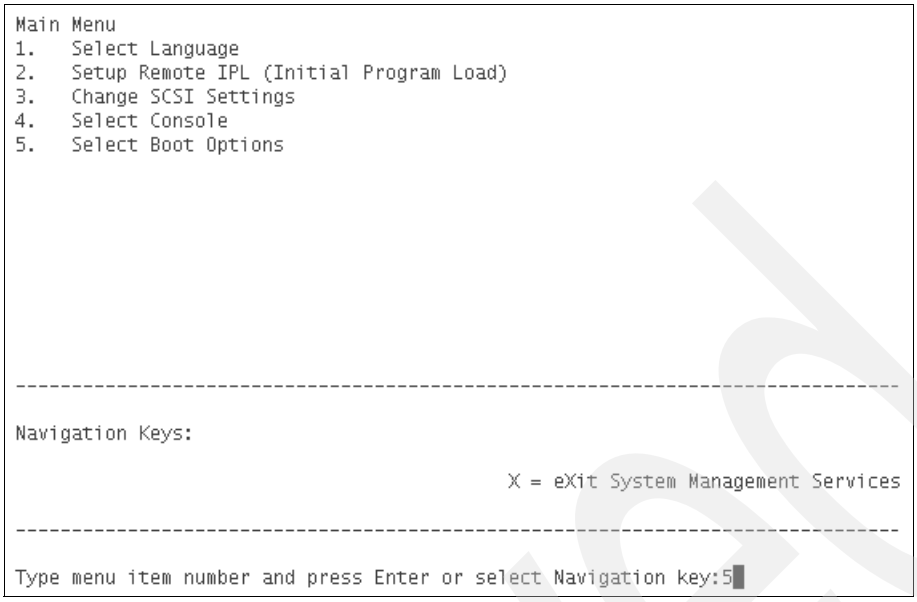

Figure 3-18 Sample Main Menu panel

After we completed the steps in the previous list, which began with the sample Main Menu window that is provided in Figure 3-18, we ended with a window filled with IBM. The software loading process had begun, as shown in Figure 3-19.

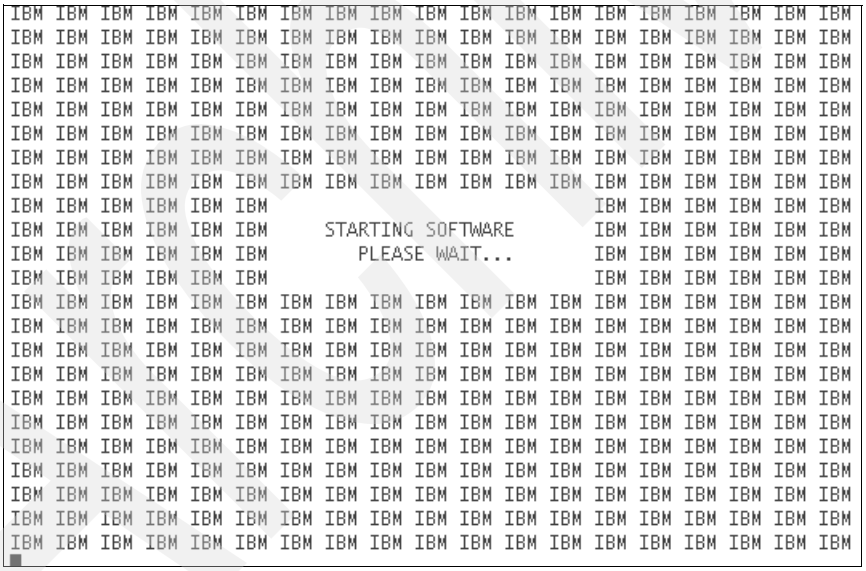

Figure 3-19 A panel filled with IBM shows that the loading process has begun

The previous window (Figure 3-19) is followed by a series of text menus that initiate the installation of VIO. The following list summarizes the responses for this paper's implementation:

- Reply 1 to use this terminal as the system console.
- Reply 1 to have English during install.
- Reply 1 to Start Install Now with Default Settings.
- Reply 1 to Continue with Install (there was also the option to specify 88 for Help, or 99 to go back to a previous window to change specifications).

Eventually, a window that states "Welcome to Virtual IO Server" opens. At this point, the boot process is completed, and it is time to log in using the padmin ID, which is a superuser ID. This initial login does not require a password, but a new password to be used for future logins is required.

Note that a reminder is also displayed at the time of the first login that states that you must accept the license agreement. Changing the password is considered acceptance of the license agreement.

3.1.3 Preparing networks for the client LPAR

Log in to the VIO LPAR as padmin to prepare the networks for the client LPAR.

Table 3-1 contains the commands that are used to prepare the networks for use by client LPARs. We took additional steps in preparing the networks for client LPARs, which we discuss later in sections that are specific to those LPARs.

Table 3-1 Commands that are used to prepare networks for use by client LPARs

Comments	Commands and responses
Specify network cards ent0 through ent11 (evenly distributed across the nodes of the Big Box) and aggregate them as a local adapter. The local adapter is known as ent12 (the response).	**mkvdev -lnagg ent0 ent2 ent4 ent6 ent8 ent9 ent10 ent11** ent12
The **mkvdev** command (make a virtual device), which, in this example, is a virtual adapter (a shared Ethernet adapter (SEA)), is called ent14. It is associated with the logical ent12. **ha_mode** is auto (high availability) ent13 is used for communication (heartbeat) between the two VIOs.	**mkvdev -sea ent14 -vadpater ent12 -default ent12 -defaultid 1 -attr ha_mode=auto ctl_chan=ent13**
Configure hostname p14n04.pbm.ihost.com to these IP addresses.	**mktcpip -hostname p14n04.pbm.ihost.com -interface ent15 -inetaddr 129.40.157.4 -netmask 255.255.255.0 -gateway 129.40.157.254 -nsrvaddr 129.40.106.1 -nsrvdomain pbm.ihost.com -start**

3.1.4 Preparing storage for client LPARs

This section provides an overview of the steps that need to be taken to permit VIO to manage SCSI devices in a way that allows sharing access to I/O by client LPARs. Log in to VIO LPARs as padmin. We used the VIO commands presented here to perform the following steps:

- Assign the serial-attached SCSI (SAS) adapter to the LPAR.
- Mirror the VIO.
- Create the disk array.
- Define the volume groups.
- Define the logical volumes within the volume groups.
- Create the virtual devices to expose logical volumes to client LPARs as hdisks.
- Create the virtual optical drive that is backed by the ISO.

Assigning the SAS adapter to the LPAR

During the creation of the partition profile for the VIO server, assign all of the I/O adapters to the VIO LPAR. Figure 3-20 shows the Hardware Management Console (HMC) panel for associating I/O adapters to the partition profile. All of the adapters shown in Figure 3-20 were assigned to the VIO server, including H9-P1-T3, the SAS adapter.

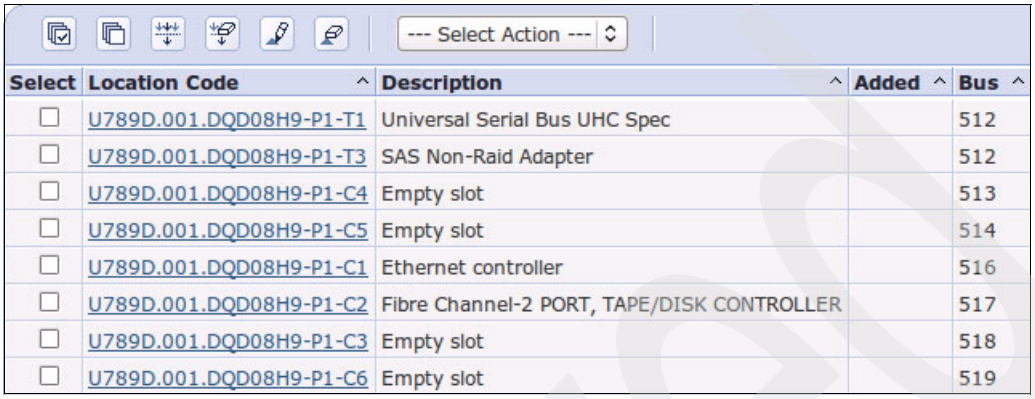

Figure 3-20 HMC panel for associating I/O adapters to the partition profile

Mirroring the VIO

When VIO was installed, a volume group was created on hdisk0. This volume group was called *rootvg*. VIO suggested practices state that this volume group must be mirrored to prevent data loss in the event of a physical disk failure.

Table 2-1 on page 11 in 2.2, “Hardware planning” on page 10 states that the Little Box has two 73 GB drives, which are identical. The first 73 GB drive is defined as hdisk0. The second 73 GB drive is defined as hdisk1.

For this implementation, it was necessary to change the volume group to add hdisk1 as an additional disk. Then, we extended the volume group by the size of the additional disk (hdisk1). Finally, we ran the **mirrorios** command to add hdisk1. Example 3-4 shows the commands that we used to accomplish this task.

Example 3-4 Mirroring hdisk0

```
$ chvg -factor 3 rootvg
$ extendvg -f rootvg hdisk1
$ mirrorios hdisk1
```

Creating the disk array

Before beginning the task of creating a disk array, we suggest that you research the content of the three online sources in the following list:

- This source provides information about accessing and using the disk array manager through the System Management Interface Tool (SMIT), or the AIX command line:

  http://publib.boulder.ibm.com/infocenter/powersys/v3r1m5/topic/arebj/sasusingthesasdiskarraymanager.htm

  This source describes the following steps to access the IBM SAS Disk Array Manager:

  a. At the command prompt, type `smit`, and press Enter.

  b. Select **Devices**.

  c. Select **Disk Array**.

d. Select **IBM SAS Disk Array**.

e. Select **IBM SAS Disk Array Manager** from the menu with options for configuring and managing the IBM SAS RAID Controller.

- This source provides information that you need relative to formatting sectors. This information includes the requirement that sectors must be 528 bytes in size:

  http://publib.boulder.ibm.com/infocenter/powersys/v3r1m5/topic/arebj/sasprepariingdisksforuseinsasdiskarrays.htm

  This online resource item describes the following steps to prepare disks for SAS disk arrays:

  a. Start the **IBM SAS Disk Array Manager**.

  b. Select **Create an Array Candidate** pdisk and format to 528 Byte sectors.

  c. Select the appropriate controller.

  d. Select the disks that you want to prepare for use in the SAS disk arrays. This step invokes a warning that this option will format the disks and that all data on the disks will be lost.

- This source discusses data redundancy (RAID 5, 6, and 10), zeroed state, and protection against failures:

  http://publib.boulder.ibm.com/infocenter/powersys/v3r1m5/topic/arebj/sascreatingasasdiskarray.htm

  This online resource item discusses disk arrays with data redundancy (RAID 5, 6, and 10), zeroed state vs. not zeroed state (considered to be in the rebuilding state), and when protection for disk failures becomes effective. RAID 0 is also discussed in the context of no data redundancy, and that it is not important to consider whether these pdisks are in the zeroed state.

Defining the volume groups

Table 3-2 shows the use of the make volume group (**mkvg**) command to create a volume group (rootvg_clients). We used this command to create a volume group for use by client LPARs.

Table 3-2 Using mkvg to create a volume group

Command	Response
`$ mkvg -f -vg rootvg_clients hdisk2`	rootvg_clients 0516-1254 mkvg: Changing the PVID in the Object Data Manager (ODM).

Defining the logical volumes within the volume groups

Table 3-3 shows the use of the make logical volume (**mklv**) command to create a logical volume called nimserv_rvg within the rootvg_clients volume group. The size of the logical volume nimsrv_rvg is specified as 10 G.

Table 3-3 Using mklv to create a logical volume

Command	Response
`$ mklv -lv nimsrv_rvg rootvg_clients 10G`	nimsrv_rvg

This logical volume becomes visible to the client LPAR on which NIM is installed. The client LPAR sees this logical volume as its root volume group (rootvg), which resides on a virtual hard disk: hdisk0.ing

Creating the virtual devices

Table 3-4 shows the use of the **mkvdev** command to create a virtual device called vnimsrv_rvg. This virtual device is associated with the logical volume to allow us to expose logical volumes to client LPARs. Note that LPAR-related and VLPAR-related operands are shown in generic form. For other implementations, replace `lpar_rvg` and `vlpar_rvg` with the appropriate names for each LPAR that is required.

Later, when you use the HMC and the Create LPAR Wizard to establish the NIM LPAR, associate its Virtual SCSI Adapter with these virtual devices. This association allows the client LPAR to access the logical volumes that have been created as disks.

Table 3-4 Using the mkvdev command to create a virtual device

Command	Response
$ mkvdev -vdev lpar_rvg -vadapter vhost0 -dev vlpar_rvg	vlpar_rvg Available
$ lsmap -vadapter vhost0	SVSA Physloc Client Partition ID --------------- -------------------- ------------- vhost0 U9117.MMA.06E3531-V1-C12 0x00000000 VTD vlpar_rvg Status Available LUN 0x8100000000000000 Backing device lpar_rvg Physloc Mirrored N/A

Creating the virtual optical drive that is backed by the ISO

Table 3-5 contains the steps, commands, and responses that are associated with the task of creating a virtual optical drive that contains the AIX image that is needed to support the installation on the NIM LPAR.

Table 3-5 Creating a virtual optical drive

Comments	Commands
Enter non-restricted shell in order to mount an Network File System (NFS) volume. According to the specifics of each implementation, replace the appropriate text with the name and location of the AIX install ISO.	**$ oem_setup_env** **$ mkdir /mnt/dcnas** **$ mount dcnas1.pbm.ihost.com:/vol/dcwork/ /mnt/dcnas** **$ exit**
On the root volume group, make a repository for the ISO files to back the virtual CD drive.	**$ mkrep -sp rootvg -size 10g** **Virtual Media Repository Created** **Repository created within "VMLibrary" logical volume**
Register the virtual ISO as AIX6.1.	**$ mkvopt -name aix6.1 -file //mnt/dcnas/aix6.1/base/CZZ81ML.iso**
Create the virtual optical drive.	**$ mkvdev -fbo -vadapter vhost1** **vtopt0 Available**
Virtually load the CD into the drive.	**$ loadopt -disk aix6.1 -vtd vtopt0**

3.2 Installing NIM

Now, it is time to install the Network Installation Manager (NIM). Recall from the discussion of the Managed Through component model that NIM facilitates the installation and management of AIX filesets over the network.

We accomplished this installation by performing the following steps:

1. Define storage to be used for rootvg for AIX (logged into VIO).
2. Define storage for NIM code (logged into VIO).
3. Create the partition profile for the NIM LPAR.
4. Install AIX on the NIM LPAR (taking advantage of the virtual optical drive that was defined and backed by the ISO, as described in Table 3-5 on page 43).
5. Use Systems Management Interface Tool (smitty) to install NIM.

Next, we describe the details of each of these steps.

3.2.1 Defining storage to be used for rootvg for NIM

The steps described in this section were taken while logged on to the VIO server as `padmin`.

Example 3-5 shows the commands that are used to define the storage to be used as rootvg for AIX. The storage for AIX was carved out using the **`mklv`** command to create a new 10 GB logical volume called nimsrv_rvg within rootvg_clients.

After successfully making the logical volume, which is called nimsrv_rvg, use the **`mkvdev`** command to create a virtual target device that maps the virtual server adapter to the target device. This action associates nimsrv_rvg with the virtual adapter vhost0. The name of the resulting virtual device is vnimsrv_rvg.

Example 3-5 The commands and responses used to define storage to be used as rootvg for AIX

```
$ mklv -lv nimsrv_rvg rootvg_clients 10G
nimsrv_rvg
$ mkvdev -vdev nimsrv_rvg -vadapter vhost0 -dev vnimsrv_rvg
vnimsrv_rvg Available
```

3.2.2 Defining storage to be used for datavg for NIM

The steps that are described in this subsection are performed while logged on to the VIO server as padmin. Example 3-6 shows the commands that are used to define the storage to be used as datatvg for NIM. The storage for NIM was carved out using the **`mklv`** command to create a new 200 GB logical volume that is called nimsrv_dvg within rootvg_clients.

After successfully making the logical volume (called nimsrv_dvg), use the **`mkvdev`** command to create a virtual target device, which maps the virtual server adapter to the target device. This action associates nimsrv_dvg with the virtual adapter vhost0. The name of the resulting virtual device is vnimsrv_dvg.

Example 3-6 The commands used to define storage to be used as datavg for NIM

```
$ mklv -lv nimsrv_rvg rootvg_clients 200G
nimsrv_dvg
```

After defining the rootvg and datavg storage, use the **lsmap** command to query the results of these previous commands. Use this command to confirm and gather two key pieces of information that are needed to associate a client server with a virtual server.

Example 3-7 shows the result of the **lsmap** command. Note that the fourth line of text (beginning with vhost0) contains this information:

```
U9117.MMA.06E3531-V1-C12                    0x00000002
```

If you use the **lsmap** command prior to defining the partition profile, the partition ID shows as all zeroes. If the same command is issued after creating the partition profile, the partition ID is the value of, in this case, the NIM partition or 2. Use the **lsmap** command to confirm that the proper relationship has been established.

Example 3-7 The result of the lsmap command

```
$ lsmap -vadapter vhost0

SVSA Physloc Client Partition ID

--------------- -------------------------------------------- ------------------

vhost0          U9117.MMA.06E3531-V1-C12                     0x00000002

VTD vnimsrv_dvg

Status Available

LUN 0x8200000000000000

Backing device nimsrv_dvg

Physloc

Mirrored N/A


VTD vnimsrv_rvg

Status Available

LUN 0x8100000000000000

Backing device nimsrv_rvg

Physloc

Mirrored N/A
```

Figure 3-21 shows how to associate the VSCSI adapter 12 with the server adapter 2.

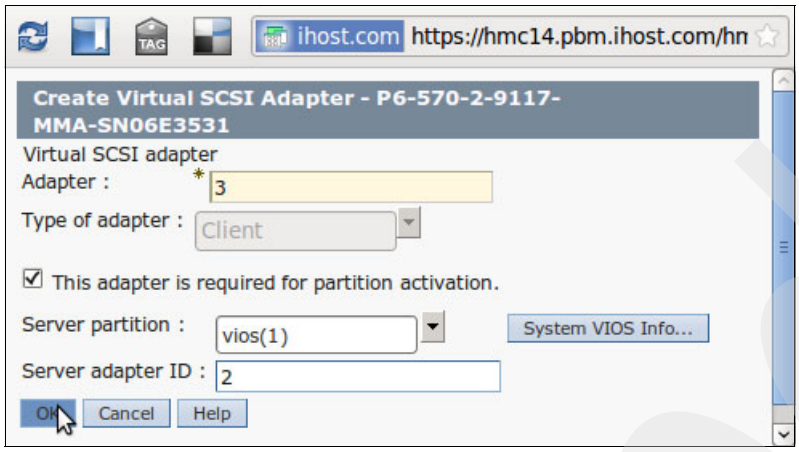

Figure 3-21 Associating VSCSI adapter with the server adapter

Where did the 12 come from? After setting up the LPAR and installing the VIO, which included the use of the **mklv** command, issue the **lsmap** command. The output of the lsmap command provides the number that is needed, which is 12 in this case, as shown in Example 3-8. Note the line that begins with `vhost0`, which also contains the text `U9117.MMA.06E3531-V1-C12`. The `12` at the end is the number to be specified as the adapter number. Note that this line also contains `0x00000002`. Because this example **lsmap** command sequence was issued after defining the client virtual SCSI adapter, this value represents the client adapter ID.

When issuing the **lsmap** command prior to defining the adapter, the value on that line is `0x00000000` instead.

Example 3-8 The output of the lsmap command provides the adapter number

```
$ lsmap -vadapter vhost0
SVSA Physloc Client Partition ID
--------------- -------------------------------------------- ------------------
vhost0          U9117.MMA.06E3531-V1-C12                     0x00000002
VTD vnimsrv_dvg

Status Available
LUN 0x8200000000000000
Backing device nimsrv_dvg
Physloc
Mirrored N/A

VTD vnimsrv_rvg
Status Available
LUN 0x8100000000000000
Backing device nimsrv_rvg
Physloc
Mirrored N/A
```

3.2.3 Creating the partition profile for the NIM LPAR

Access the Create LPAR Wizard from the HMC by selecting the hardware server (Little Box) and choosing **Create a new LPAR profile**.

Figure 3-22 provides a summary of the profile for the NIM LPAR. Note that the summary shows the Partition ID and Partition name, Partition environment (shown as AIX or Linux), Profile name, Desired memory, Desired processing units, and Boot mode. All of these values were specified through a series of Wizard windows.

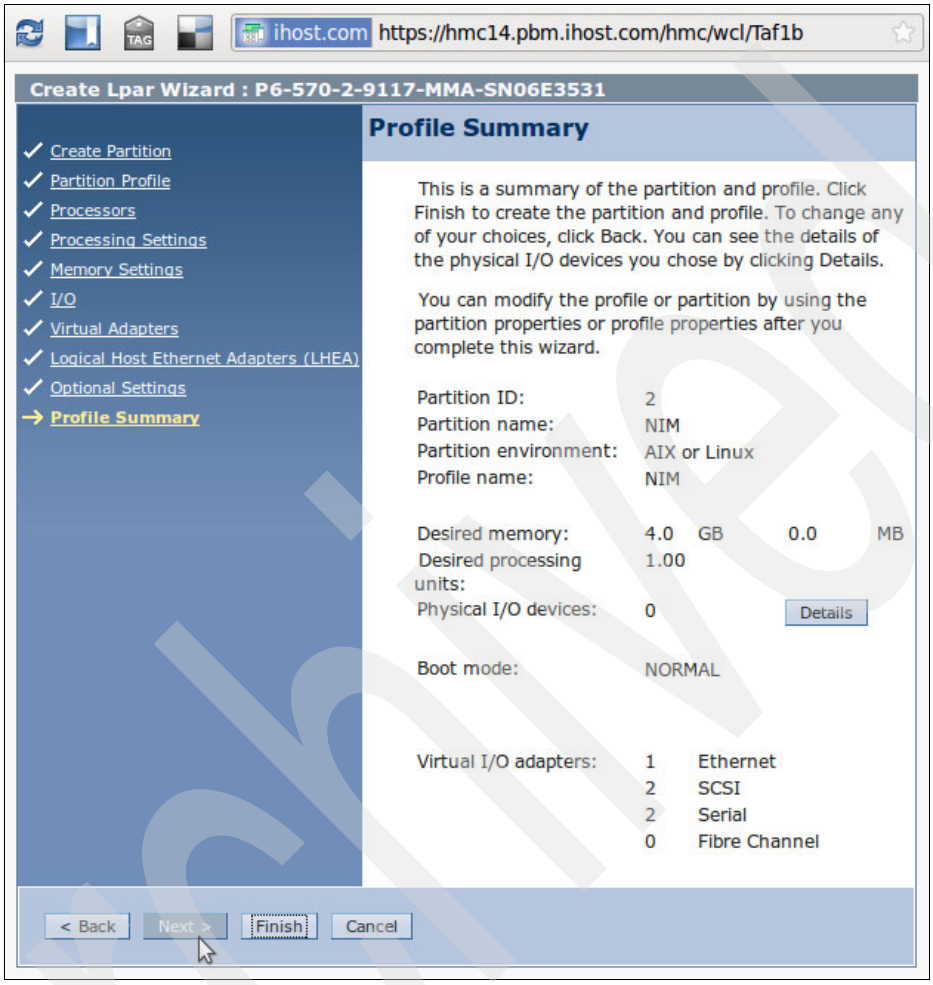

Figure 3-22 The summary of the profile for the NIM LPAR

The following series of Create LPAR Wizard windows show the steps that have been taken to define the NIM LPAR. Note that in Figure 3-23, no physical I/O was to be associated with the LPAR. Click **Next** to proceed to the Virtual Adapters window.

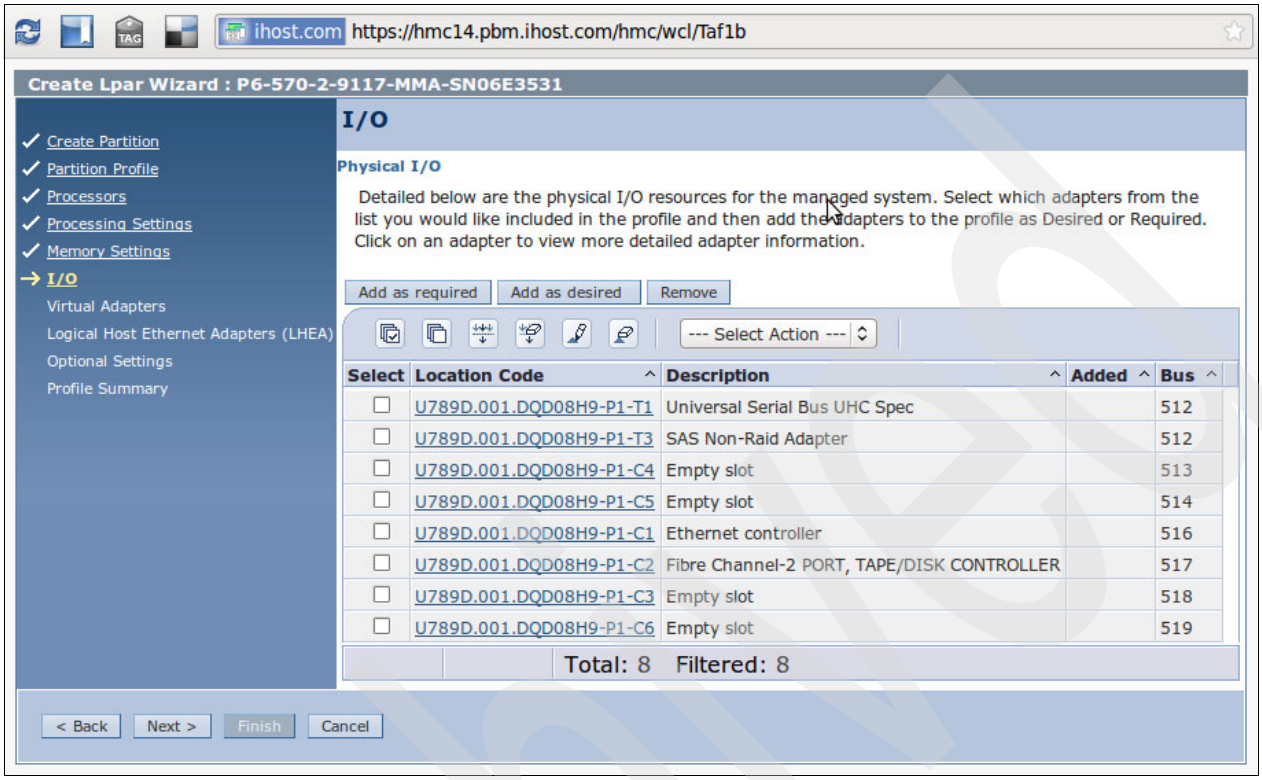

Figure 3-23 Do not select any physical I/O

After clicking Next, the window that is shown in Figure 3-24 on page 49 opens. We selected Create Virtual Adapter and then Ethernet Adapter. A pop-up window that is shown in Figure 3-25 on page 49 opens. The specifications for Ethernet0 (default) are shown.

Then (as shown in Figure 3-26 on page 50), we selected Create Virtual Adapter and selected SCSI Adapter. Recall that the output of **lsmap**, as shown in Example 3-9, can be used to verify the relationship between U9117.MMA.06E3531-V1-C12 and 0x00000002.

Example 3-9 Verifying the relationship between U9117.MMA.06E3531-V1-C12 and 0x00000002

```
$ lsmap -vadapter vhost0
SVSA Physloc Client Partition ID
--------------- -------------------------------------------- ------------------
vhost0          U9117.MMA.06E3531-V1-C12                     0x00000002
VTD vnimsrv_dvg

Status Available
LUN 0x8200000000000000
Backing device nimsrv_dvg
Physloc
Mirrored N/A

VTD vnimsrv_rvg
Status Available
LUN 0x8100000000000000
```

```
Backing device nimsrv_rvg
Physloc
Mirrored N/A
```

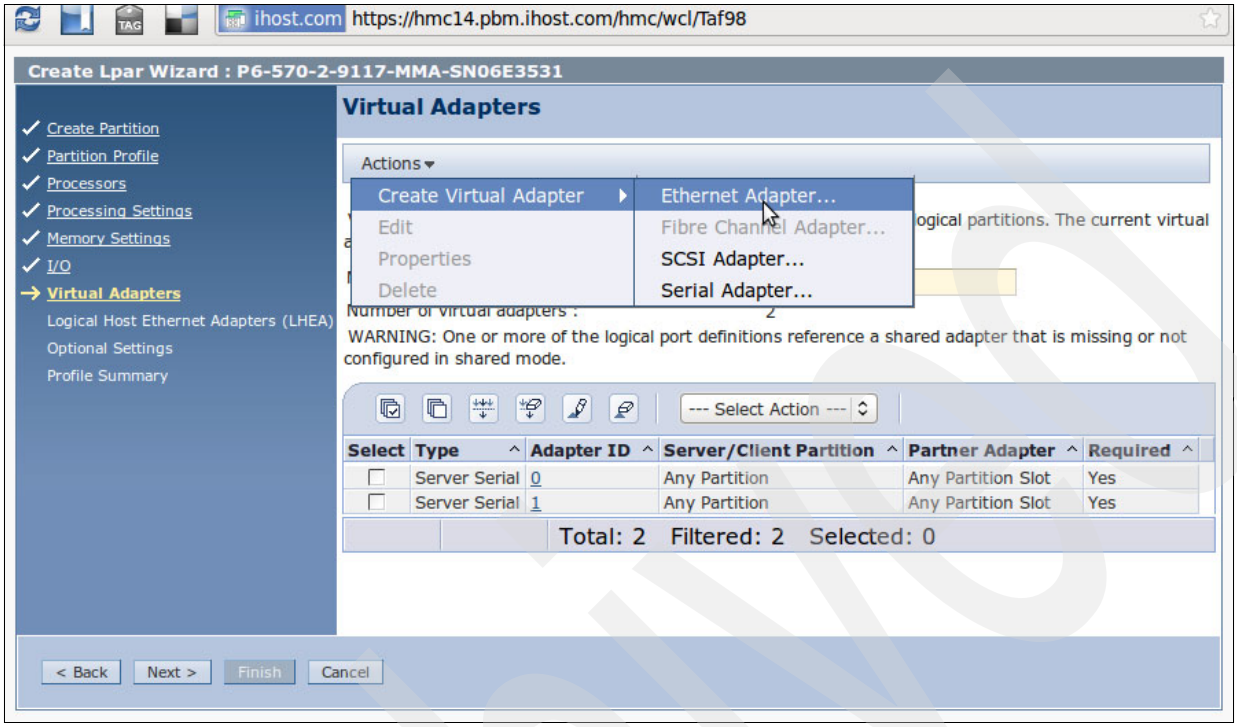

Figure 3-24 Choose Create Virtual Adapter then Ethernet Adapter

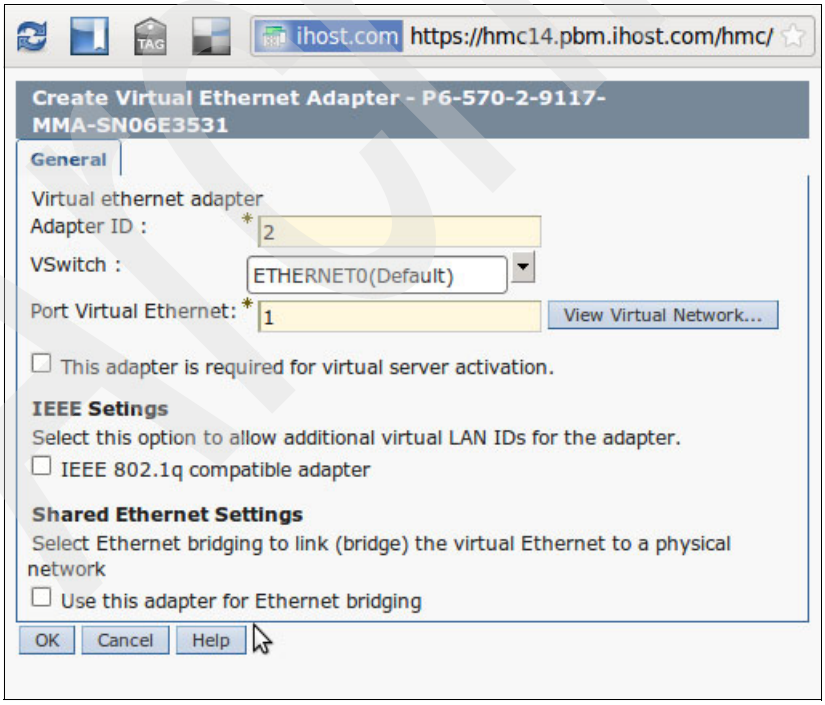

Figure 3-25 Create virtual Ethernet adapter

After providing the Ethernet Virtual Adapter information, we selected **Create Virtual Adapter** again. This time, we selected **SCSI Adapter**.

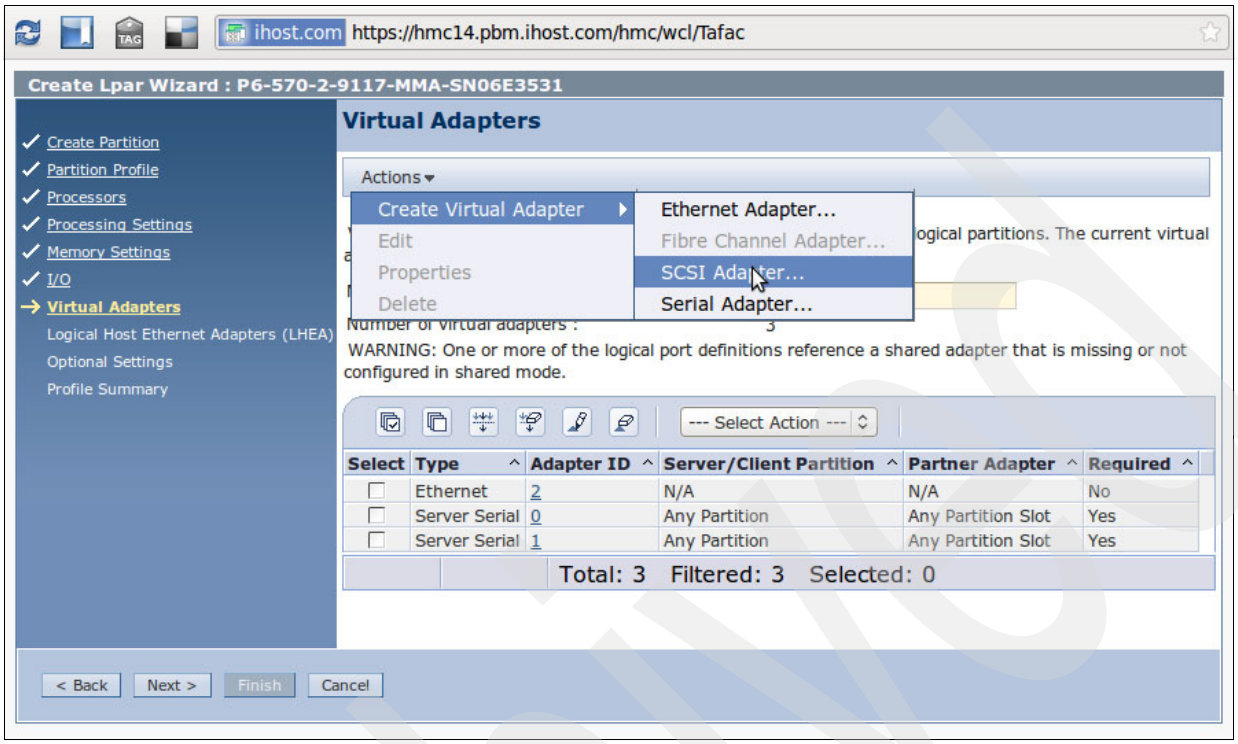

Figure 3-26 Choose SCSI adapter

After selecting Create Virtual Adapter and SCSI Adapter, the window that is shown in Figure 3-27 opened. At this point, we selected **vios(1)**. The value for the Virtual SCSI adapter adapter was prepopulated with the value 3 and the value of the Server adapter ID was prepopulated with the value 12.

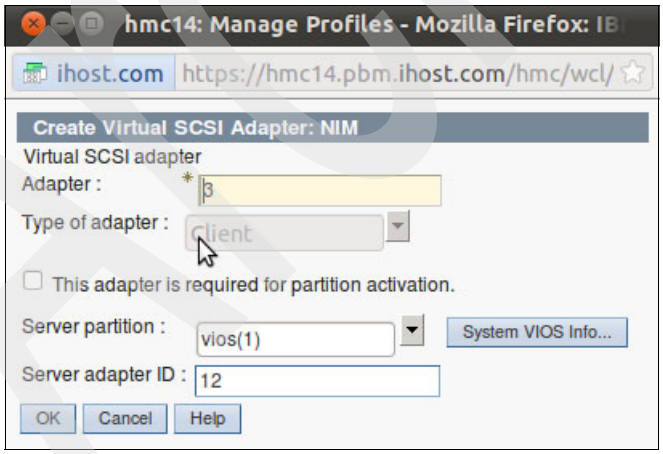

Figure 3-27 Create Virtual SCSI Adapter for NIM

As a result of clicking **OK** in the Create Virtual SCSI Adapter panel, a summary is provided, as shown in Figure 3-28.

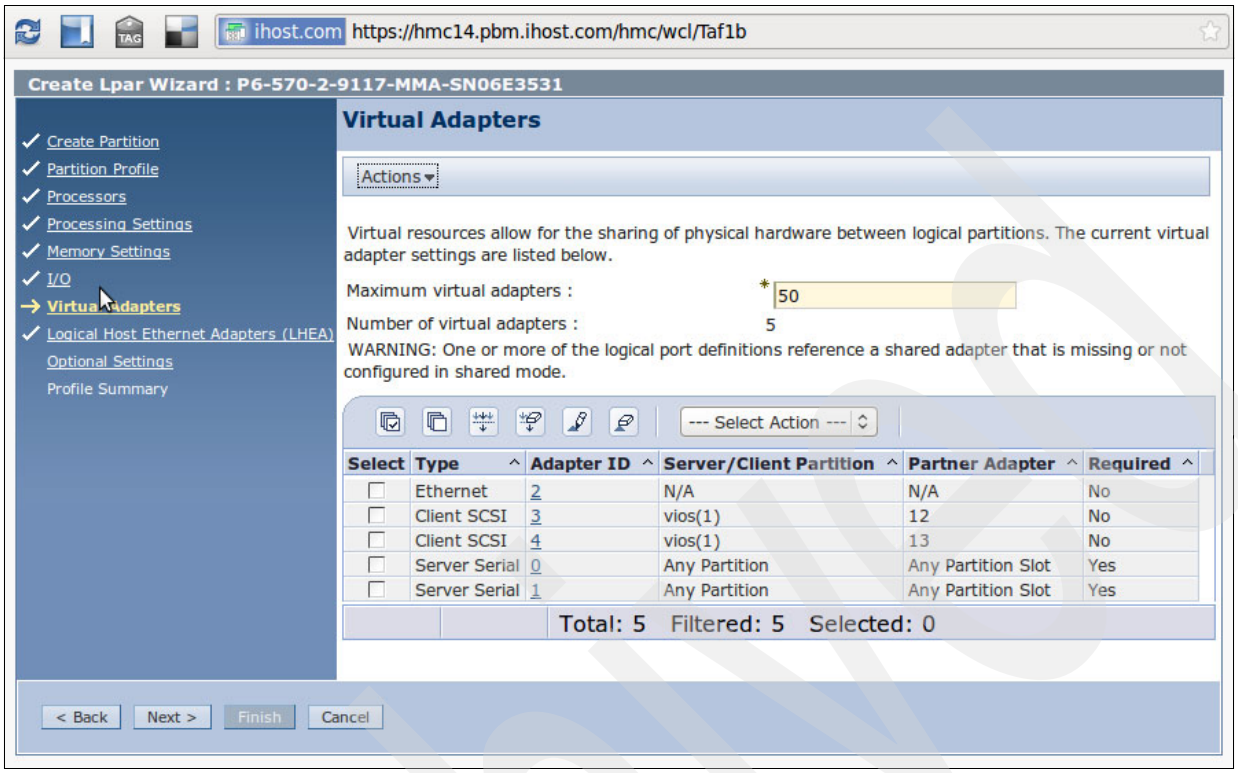

Figure 3-28 Virtual Adapter settings

In Figure 3-28, the second row (Client SCSI type) lists the attributes of virtual adapter ID 3 as VIOS(1) with a Partner Adapter of 12, which are the values specified earlier in Figure 3-27 on page 50.

In Figure 3-29, we selected no Logical Host Ethernet Adapter (LHEA) information. We clicked **Next**.

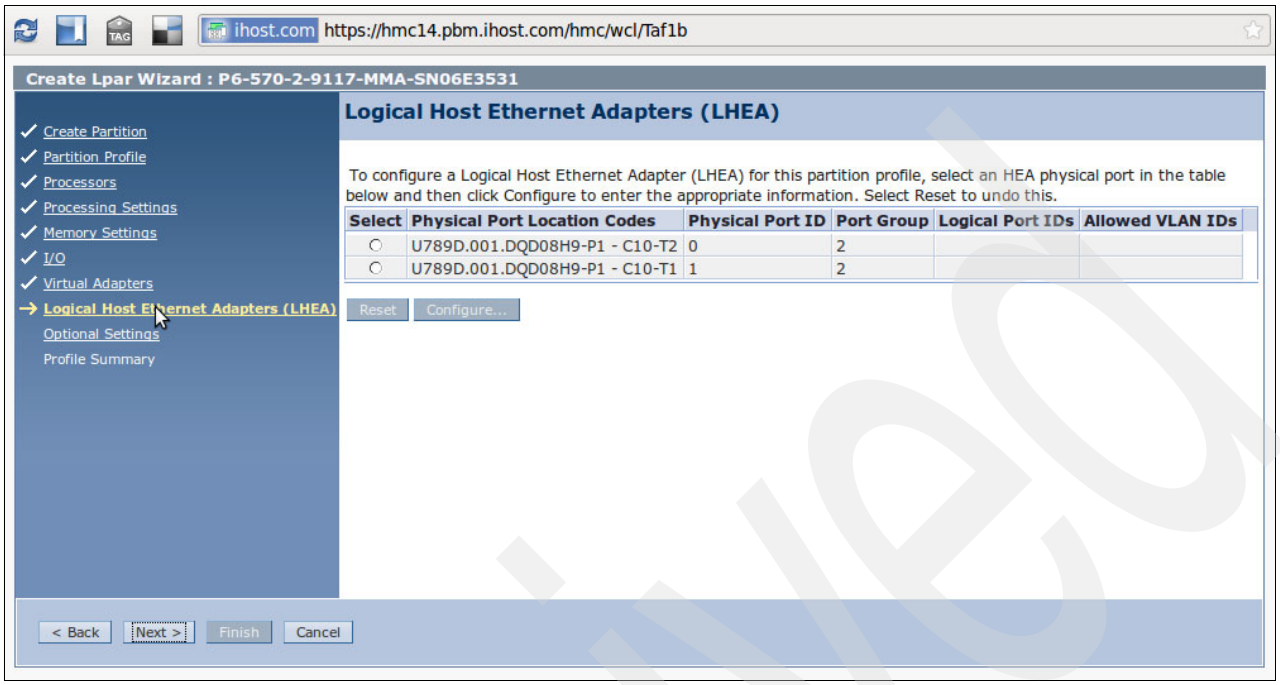

Figure 3-29 No Logical Host Ethernet Adapter (LHEA) information is specified

In Figure 3-30, we selected Boot mode **Normal**. We selected **Next**.

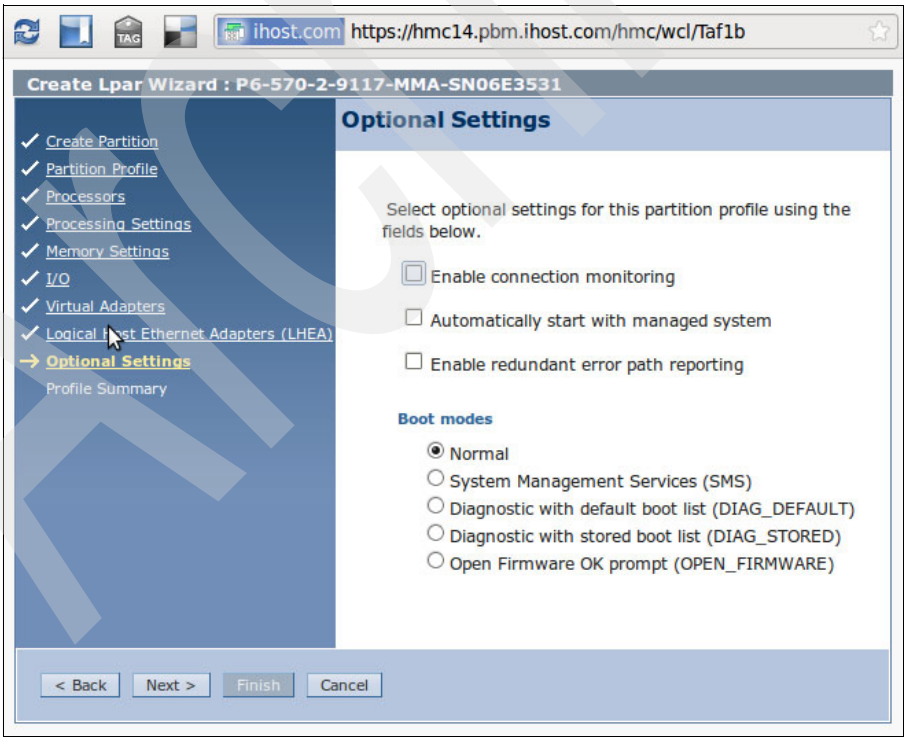

Figure 3-30 Optional Settings window

As a result, the Profile Summary window opens, as shown in Figure 3-31. We clicked **Finish**.

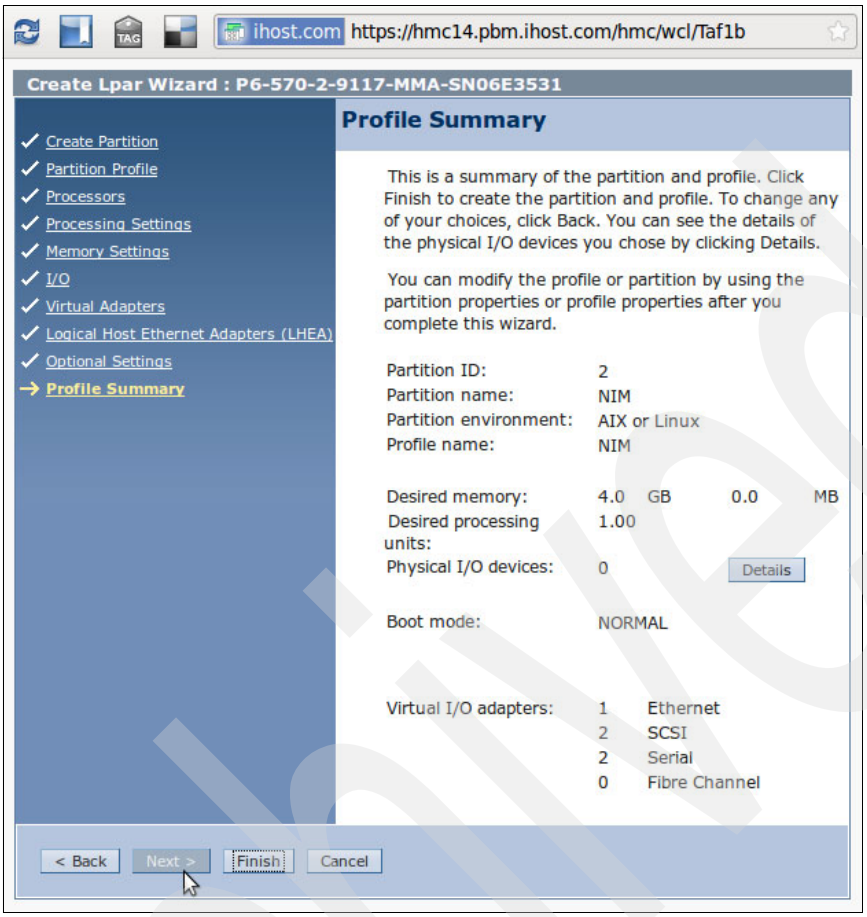

Figure 3-31 Profile Summary

3.2.4 Installing AIX in the NIM LPAR

To install AIX on the NIM LPAR, open a terminal window on the HMC by activating NIM LPAR and selecting **Open a terminal window or console session**. The console session that opens contains the Main Menu, as shown in Figure 3-32 on page 54. The “Select Boot Options” allowed us to take advantage of the Virtual Optical Drive backed by the ISO that was previously established. The following list describes the responses for this paper’s implementation to a series of windows beginning with Figure 3-32 on page 54. The end result was that we exited System Management Services, and AIX in the LPAR started.

1. We chose “Select Boot Option” by typing option 5 at the bottom of the window.
2. We chose “Select Install/Boot Device” by typing 1 at the bottom of the window.
3. We used “N” to scroll through multiple windows to locate the value (in this case “27”) that was associated with the device to be specified as the “install from” device. Then, we typed that value (27) at the bottom of the window.
4. We chose “CD/DVD” by typing 3 at the bottom of the window.
5. We chose “SATA” by typing 5 at the bottom of the window.
6. We chose “a Media Adapter” by typing 1 at the bottom of the window.
7. We chose “SATA CD ROM” (the only device listed) by typing 1 at the bottom of the window.

8. We chose “Normal Mode Boot” by typing 2 at the bottom of the window.
9. We chose “Yes” to indicate exit System Management Services by typing 1 at the bottom of the window.

Figure 3-32 Main Menu

If you follow the steps in the previous list beginning with the sample Main Menu window that is provided in Figure 3-32, you see a window filled with IBM as an indicator that the loading process has begun, as shown in Figure 3-33.

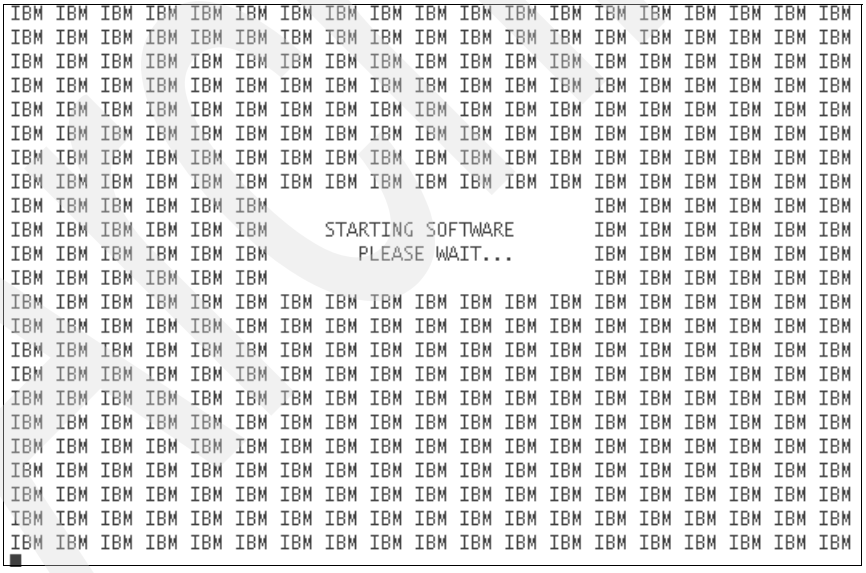

Figure 3-33 A panel filled with IBM indicates that the loading process has begun

Soon, the window full of IBMs is replaced by what is shown in Figure 3-34. At this point, the option to define the system console appears. Typing 1 indicates this terminal will be used as the system console.

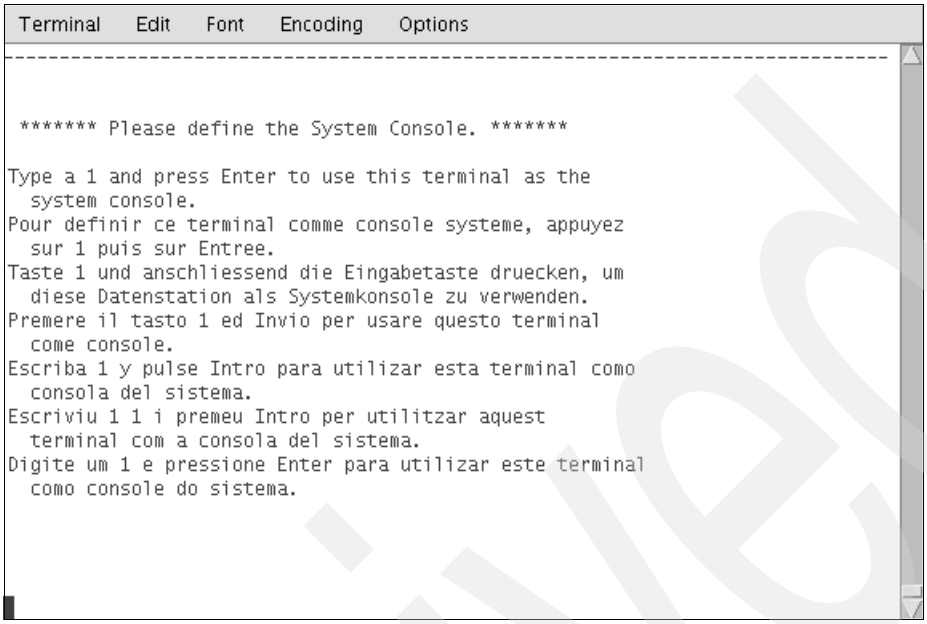

Figure 3-34 Specifying the system console

After specifying the system console, informational messages appear that include copyright information, such as the information that is shown in Figure 3-35.

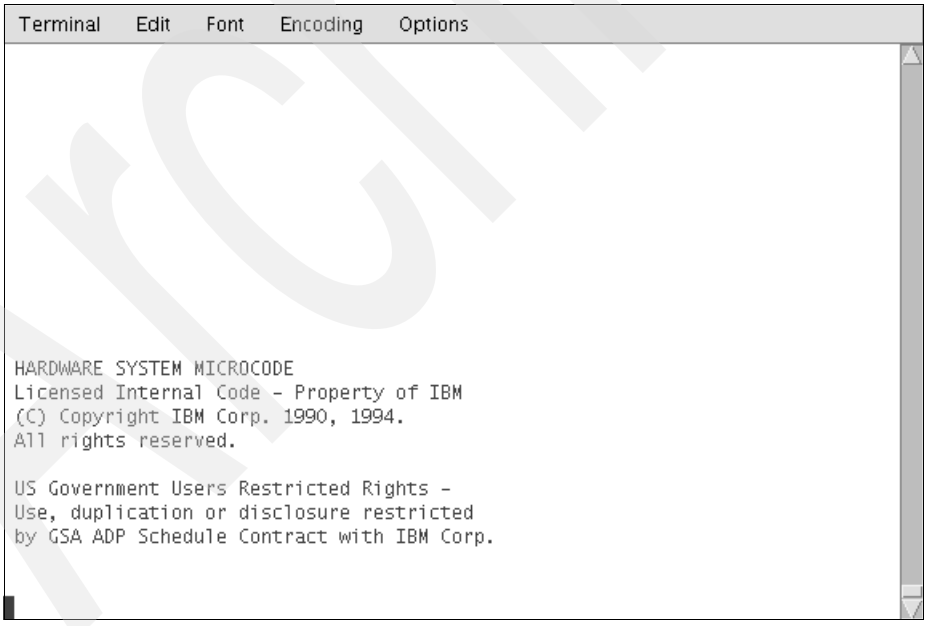

Figure 3-35 When boot process begins, messages appear that include copyright information

Shortly, the contents of Figure 3-36 on page 56 appear.

In Figure 3-36, we chose English for the language by typing 1.

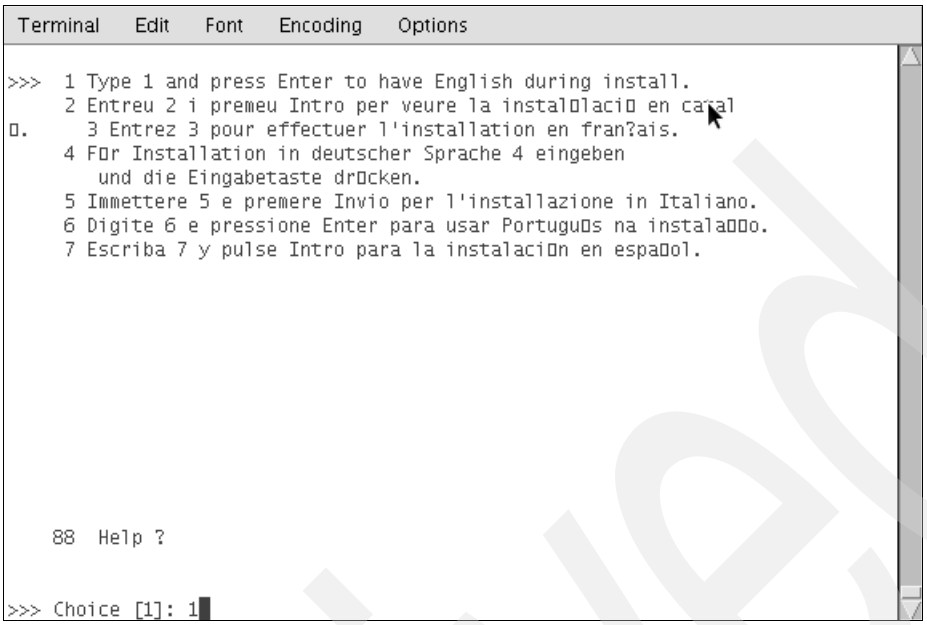

Figure 3-36 Specifying English as the language

After specifying the installation language, the Base Operating System Installation and Maintenance Menu appears, as shown in Figure 3-37. Again, type 1 to initiate an installation with default settings.

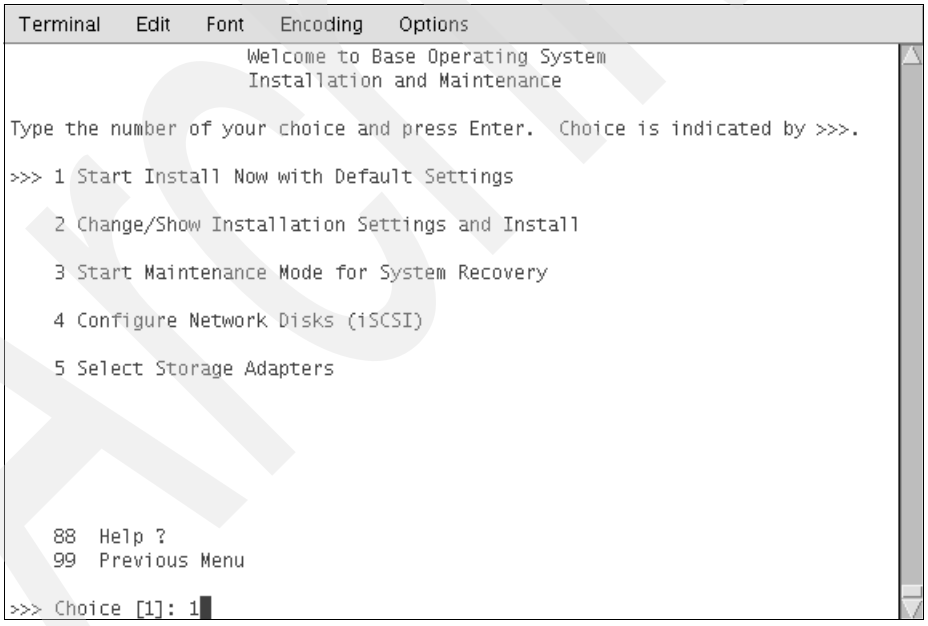

Figure 3-37 Base Operating System Welcome and Installation and Maintenance menu

After typing 1 to begin the installation with default settings, the System Backup Installation Summary is displayed, as shown in Figure 3-38 on page 57, with a warning in the lower right corner that continuing will destroy or impair recovery of ALL data on the destination disk hdisk0. By typing 1, the installation continues.

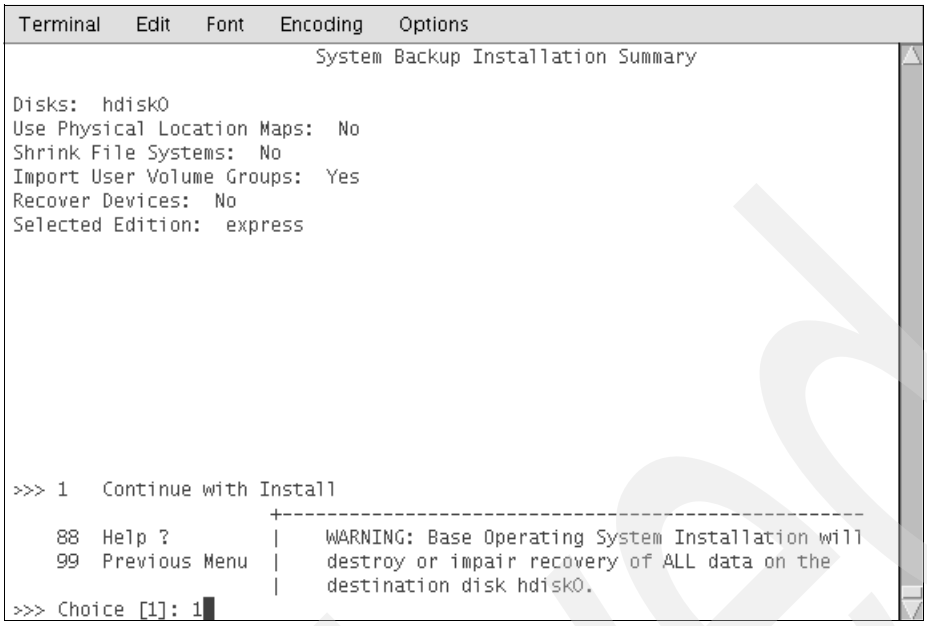

Figure 3-38 Warning BOS and data will be destroyed on hdisk0

At this point, as shown in Figure 3-39, the installation progress is displayed. The approximate percentage of tasks completed and the elapsed time (in minutes) are updated as the installation progresses. An appropriate set of messages follows this window to indicate that the installation is complete.

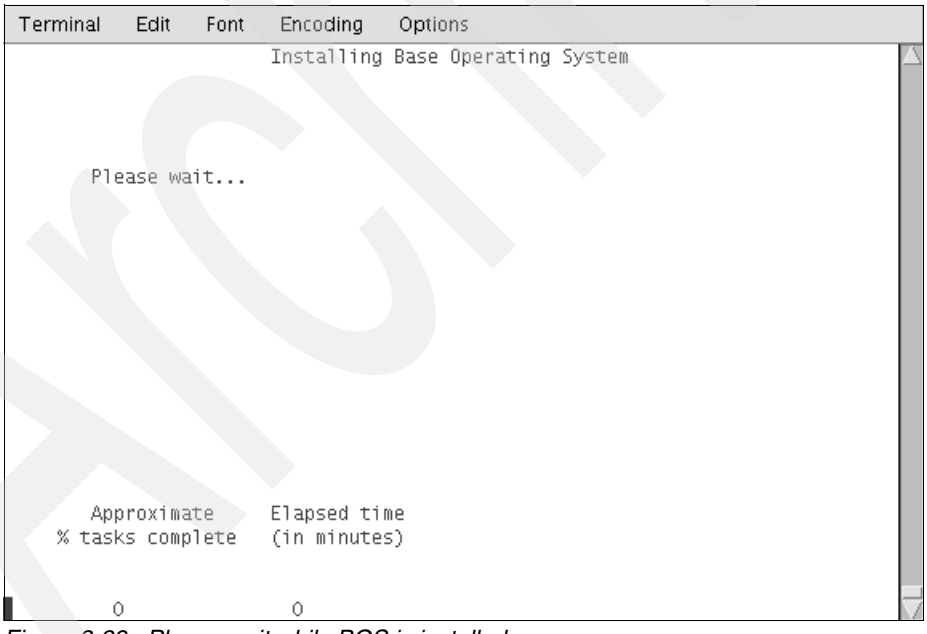

Figure 3-39 Please wait while BOS is installed

3.2.5 Installing and configuring NIM

Log on to AIX on the NIM LPAR as root. Use the Systems Management Interface Tool (smitty) text version. This tool presents a series of menu windows (starting with the sample window that is shown in Figure 3-40 on page 58. Responses are selected by using the cursor

movement keys to highlight the desired response and then pressing Enter. The responses that are used are shown in the following sample smitty window in Figure 3-40.

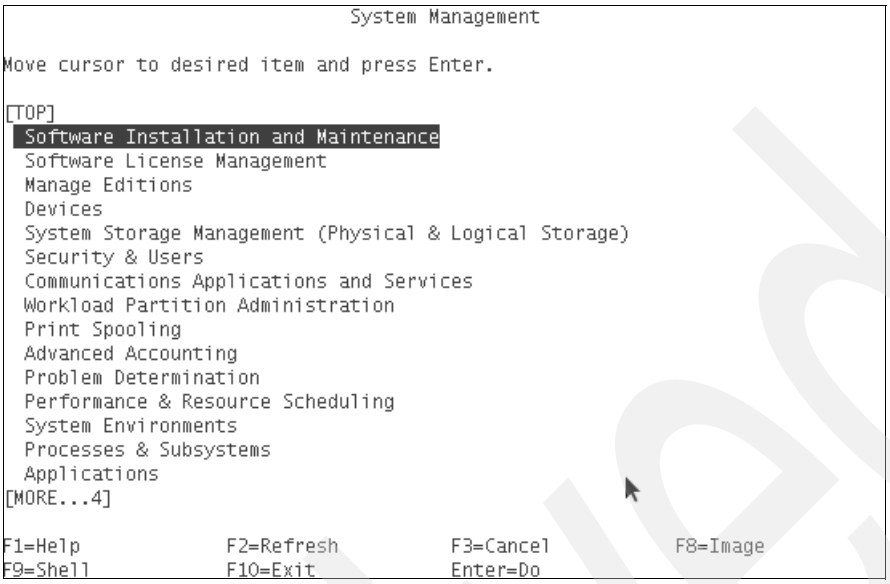

Figure 3-40 Smitty menu

This list shows the responses to smitty for this paper's implementation:

- Software installation and maintenance
- EZ NIM (EZ NIM tool)
- Configure as a NIM master

  This option provided a window for the specification of Software Source [`cd1`], Volume Group for resources [`datavg`], Filesystem for Resources [`/export/eznim`], Create System Backup Image [`yes`], Create New Filesystem [`yes`], and Display Verbose Output [`no`].

When the installation process completes, a window similar to Figure 3-41 opens. It contains the message, "`NIM Master Setup Complete - Enjoy`".

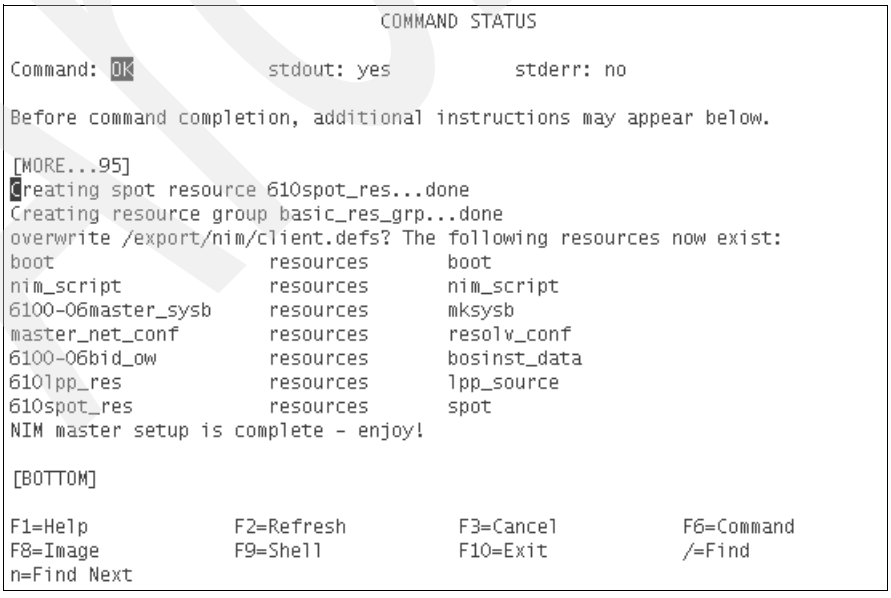

Figure 3-41 NIM master setup is complete

3.3 Installing IBM Systems Director

At this point, you can install IBM Systems Director. This installation can take advantage of using NIM to install AIX from the NFS.

Recall from the discussion of the Managed Through Component Model that was presented in Chapter 1, the IBM Systems Director provides a single point-of-control that "...facilitates IT management tasks related to the total IT environment". The significance is that after installing IBM Systems Director, its features and facilities can be used to automate the deployment of subsequent LPARs.

This installation was accomplished by taking the following steps:

- Define storage for rootvg
- Define storage for datavg
- Creation of the partition profile for the IBM Systems Director LPAR
- Take advantage of NIM to install AIX
- Use command line to install IBM Systems Director
- Install the following plug-ins:
  - Network Control
  - VM Control
  - Storage Control
  - AEM

3.3.1 Defining storage to use for rootvg for Systems Director

We performed the steps described in this subsection while logged on to the VIO server as `padmin`.

Example 3-10 shows the commands used to define storage to be used as rootvg for AIX in the Systems Director LPAR. The storage for AIX was carved out using the **`mklv`** command to create a new 20 GB logical volume called director_rootvg within rootvg_clients.

After successfully making the logical volume (called vdirector_rootvg), use the **`mkvdev`** command to create a virtual target device, which maps the virtual server adapter to the target device. This action associates vdirector_rvg with the virtual adapter vhost2. The name of the resulting virtual device is vdirector_rvg.

Example 3-10 Commands and responses used to define storage as rootvg for Systems Director LPAR

```
$ mklv -lv director_rootvg rootvg_clients 20G
director_rootvg

$ mkvdev -vdev director_rootvg -vadapter vhost2 -dev vdirector_rvg
vdirector_rvg Available
```

After defining the rootvg, use the **`lsmap`** command, as shown in Example 3-11, to query the results of these previous commands.

Example 3-11 Using lsmap to query vadapter vhost2

```
$ lsmap -vadapter vhost2
SVSA            Physloc                                     Client Partition ID
--------------- ------------------------------------------- ------------------
vhost2          U9117.MMA.06E3531-V1-C14                    0x00000003
```

```
VTD                  vdirector_rvg
Status               Available
LUN                  0x8100000000000000
Backing device       director_rootvg
Physloc
Mirrored             N/A
```

When using the **lsmap** command prior to defining the partition profile, the partition ID shows as all zeroes and Physloc is empty. After creating the partition profile, the **lsmap** command shows the Partition ID as the value of (in this case) the Director partition, which is 0x00000003 or 3.

3.3.2 Creating the partition profile for the IBM Systems Director LPAR

In the context of the Create Partition dialog, the following list provides a summary of the attributes and responses provided. Refer also to Figure 3-42 on page 61.

- Partition ID: 3
- Profile name: Director
- Partition environment: AIX or Linux
- Profile name: director
- Desired memory: 4.0 GB
- Desired processing units: 2
- Physical I/O devices: 0
- Boot mode: NORMAL
- Virtual I/O adapters: (use values here garnered from **$ lsmap -vadapter vhost0**):
  - 1 Ethernet
  - 2 SCSI
  - 2 Serial
  - 0 Fibre Channel

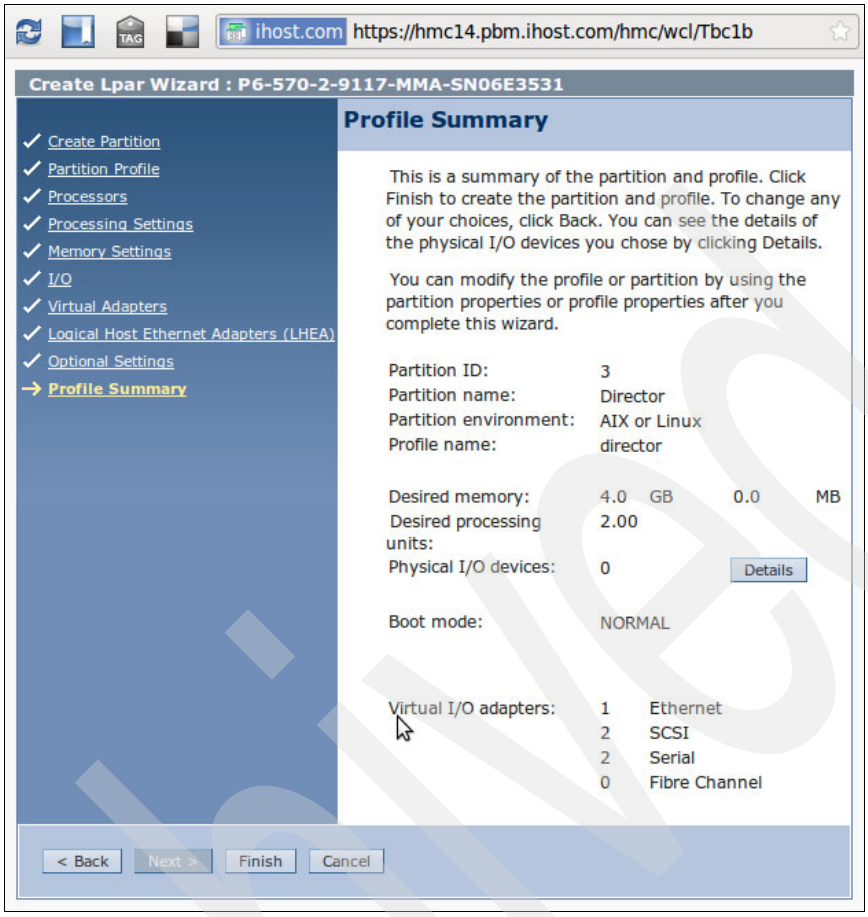

Figure 3-42 Profile Summary for the partition named Director

3.3.3 Using NIM to install AIX

Now that a NIM infrastructure has been built, you can automate the Base Operating System (BOS) deployment. In the previous section, we defined an LPAR profile to the system through the HMC. To install AIX on this LPAR, we defined the LPAR to NIM. By using smitty on the NIM Server that was created in *3.2, "Installing NIM" on page 44*, we initiated a BOS install.

Windows: The menus that are shown in the following figures have been cropped.

After entering smitty, we selected the **Software Installation and Maintenance** option, as shown in Figure 3-43.

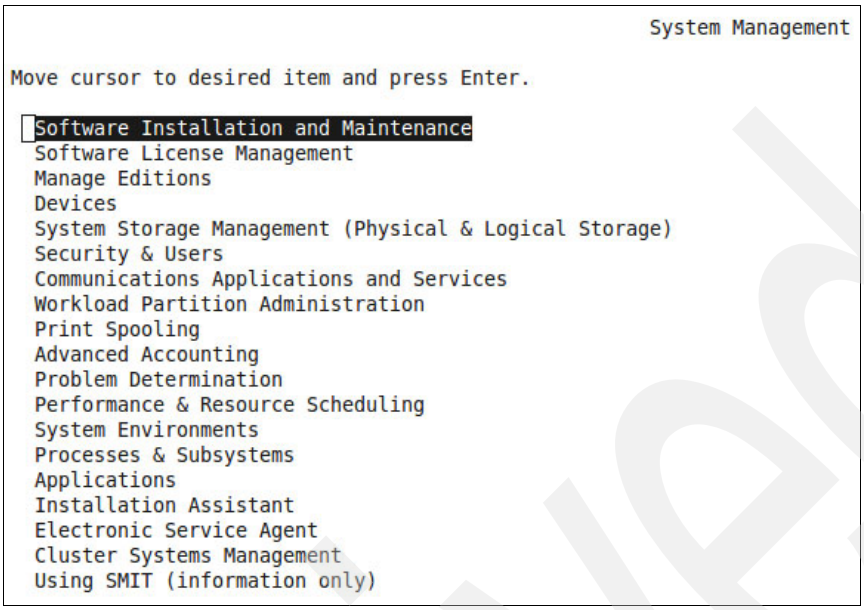

Figure 3-43 Choosing Software Installation and Maintenance

We selected **Install and Update Software**, as shown in Figure 3-44.

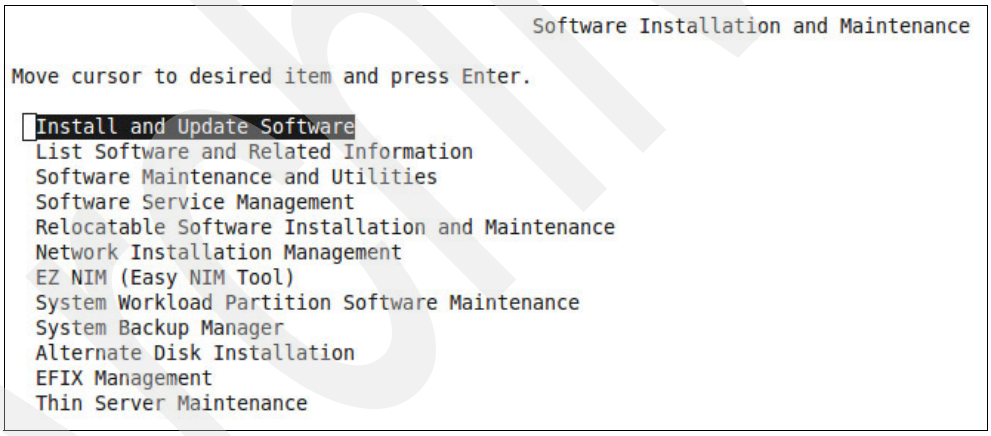

Figure 3-44 Choosing to Install and Update Software

Then, we chose **Perform NIM Software Installation and Maintenance Tasks**, as shown in Figure 3-45.

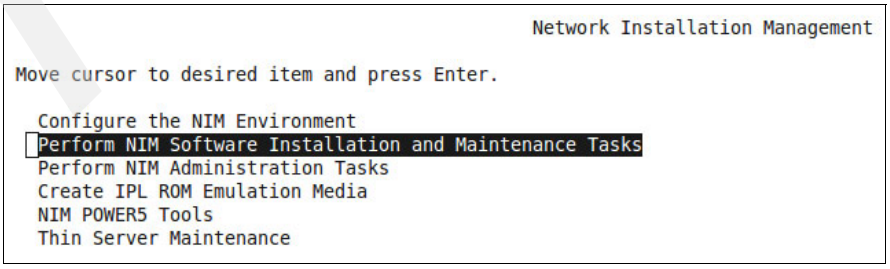

Figure 3-45 Choosing Perform NIM Software Installation and Maintenance Tasks

As shown in Figure 3-46, we selected **Manage Machines**.

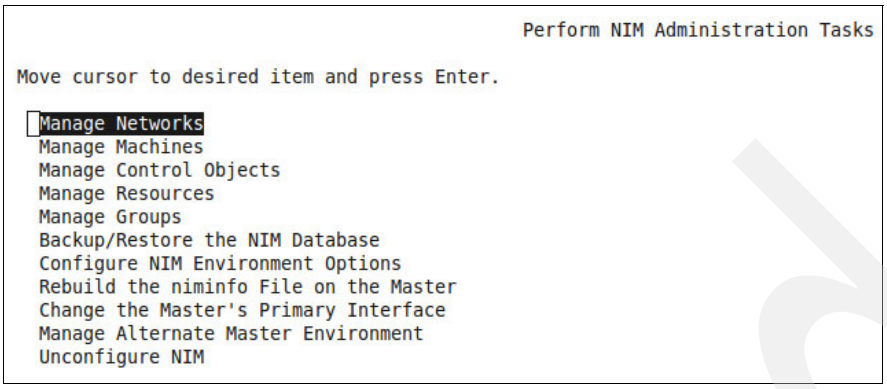

Figure 3-46 Choosing Manage Machines

Then, as shown in Figure 3-47, we chose **Define a Machine**.

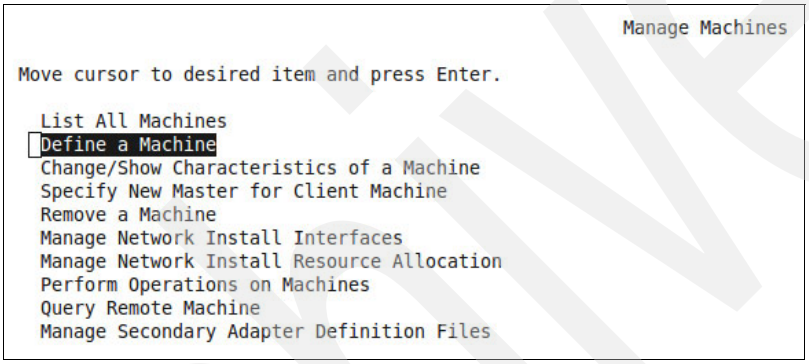

Figure 3-47 Choosing Define a Machine

The Define a Machine menu is shown in Figure 3-48. We entered the NIM Machine Name and Host Name. This data corresponds to the information that was provided to the HMC for the LPAR.

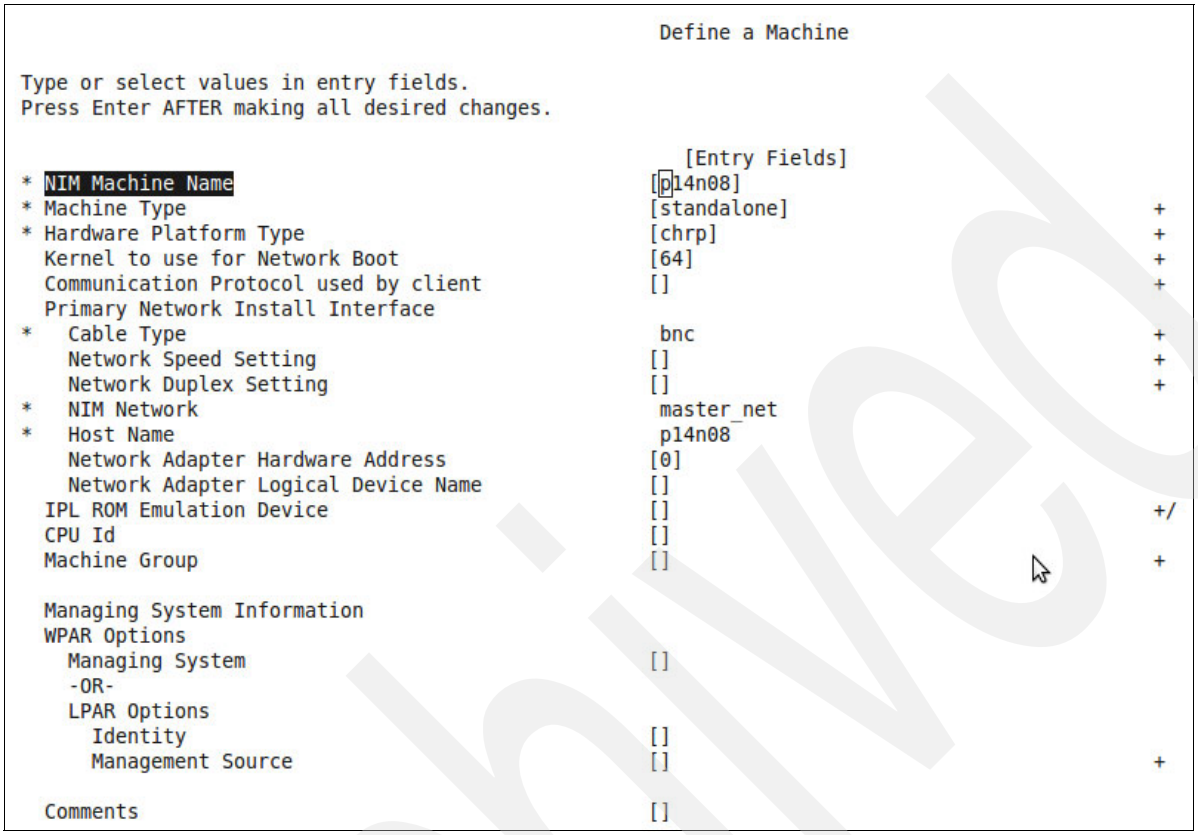

Figure 3-48 Define a Machine Menu

After the previous step completes, we went back to the main NIM menu and selected **Perform NIM Software Installation and Maintenance Tasks**.

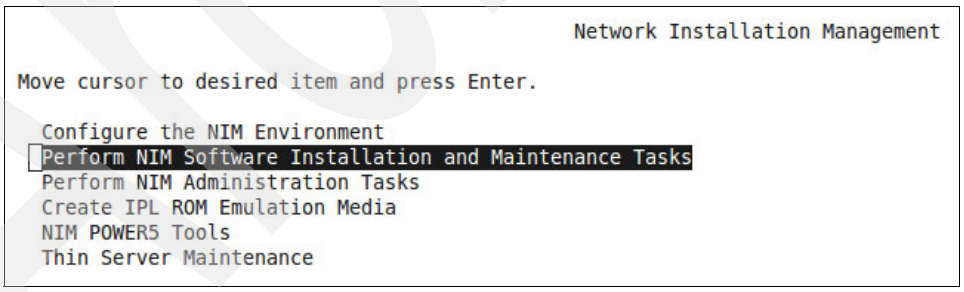

Figure 3-49 Selecting Perform NIM Software Installation and Maintenance Tasks

in Figure 3-50, we selected **Install and Update Software**.

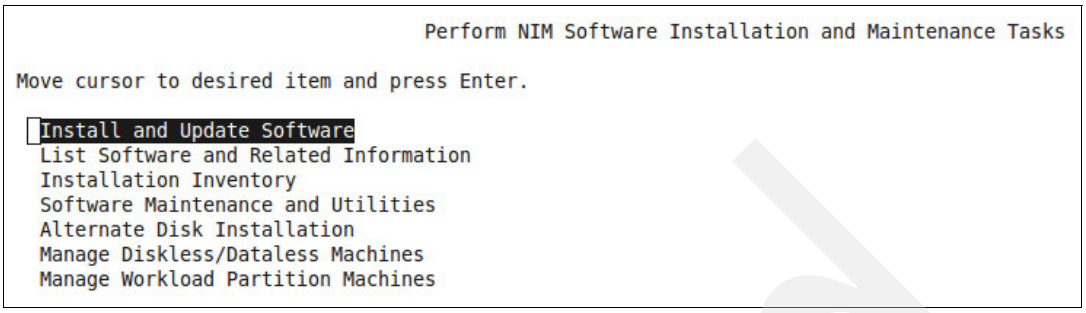

Figure 3-50 Selecting Install and Update Software

From the following menu, we chose **Install the Base Operating System on Standalone Clients**, as shown in Figure 3-51.

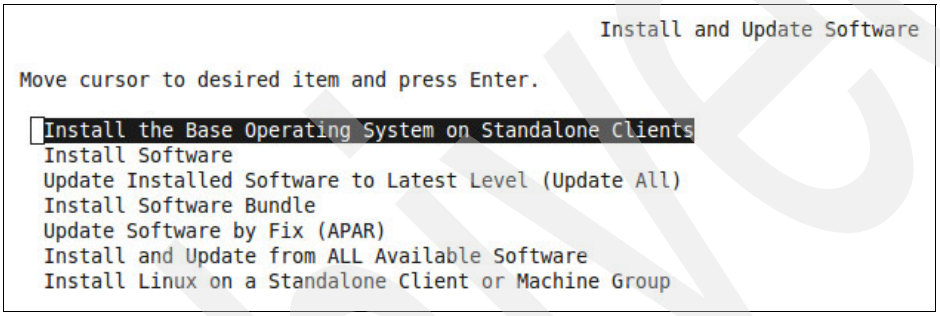

Figure 3-51 Install the Base Operating System on Standalone Clients

This action displays the Select a TARGET for the operation submenu, as shown in Figure 3-52. Only one standalone client has been defined.

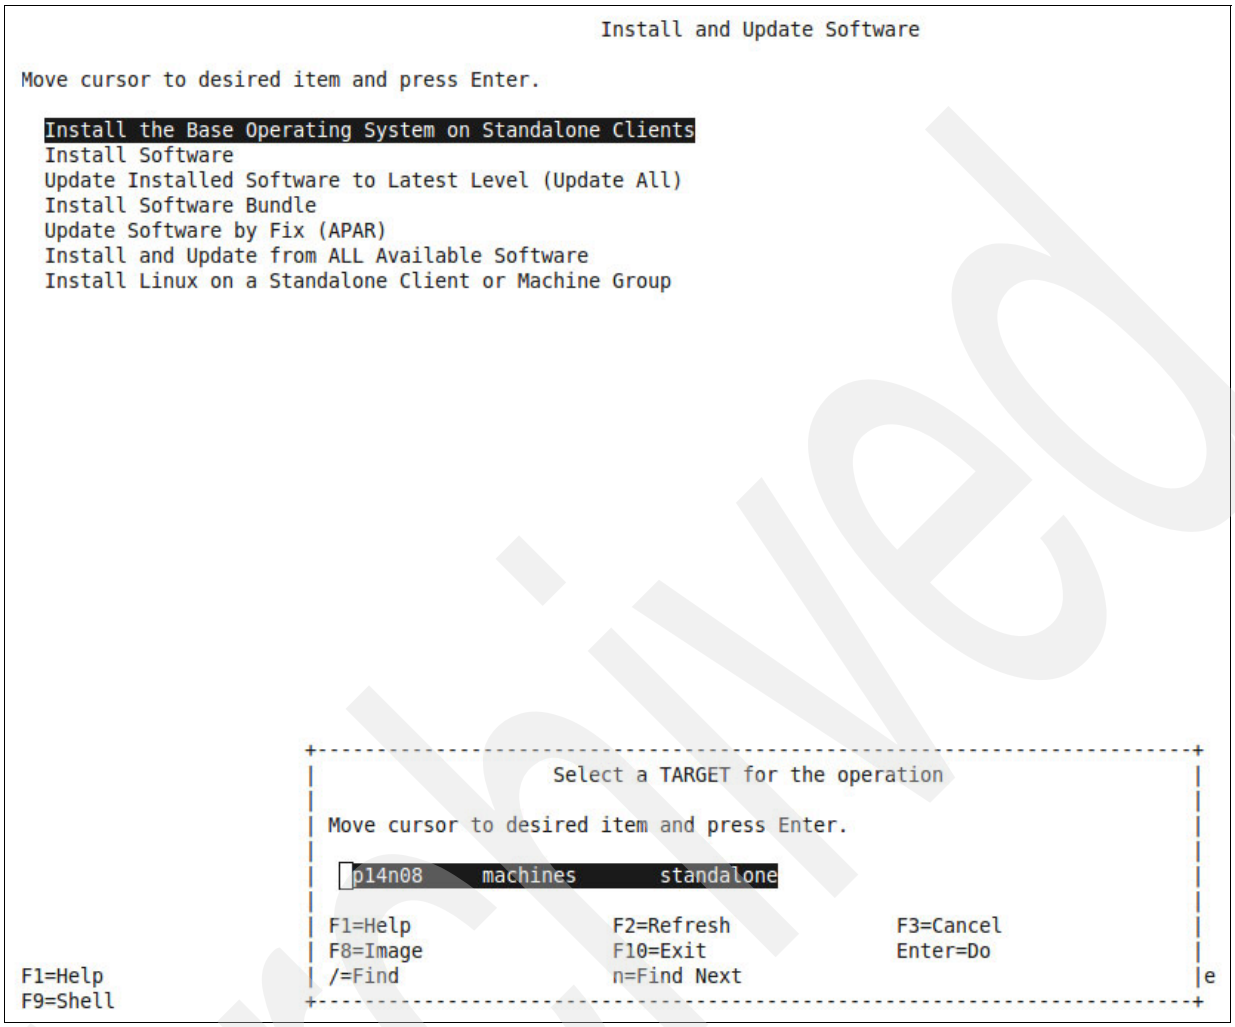

Figure 3-52 Select a TARGET for the operation submenu

Figure 3-53 shows that the Select the installation TYPE submenu is displayed and the mksysb option (Install from a mksysb) was selected.

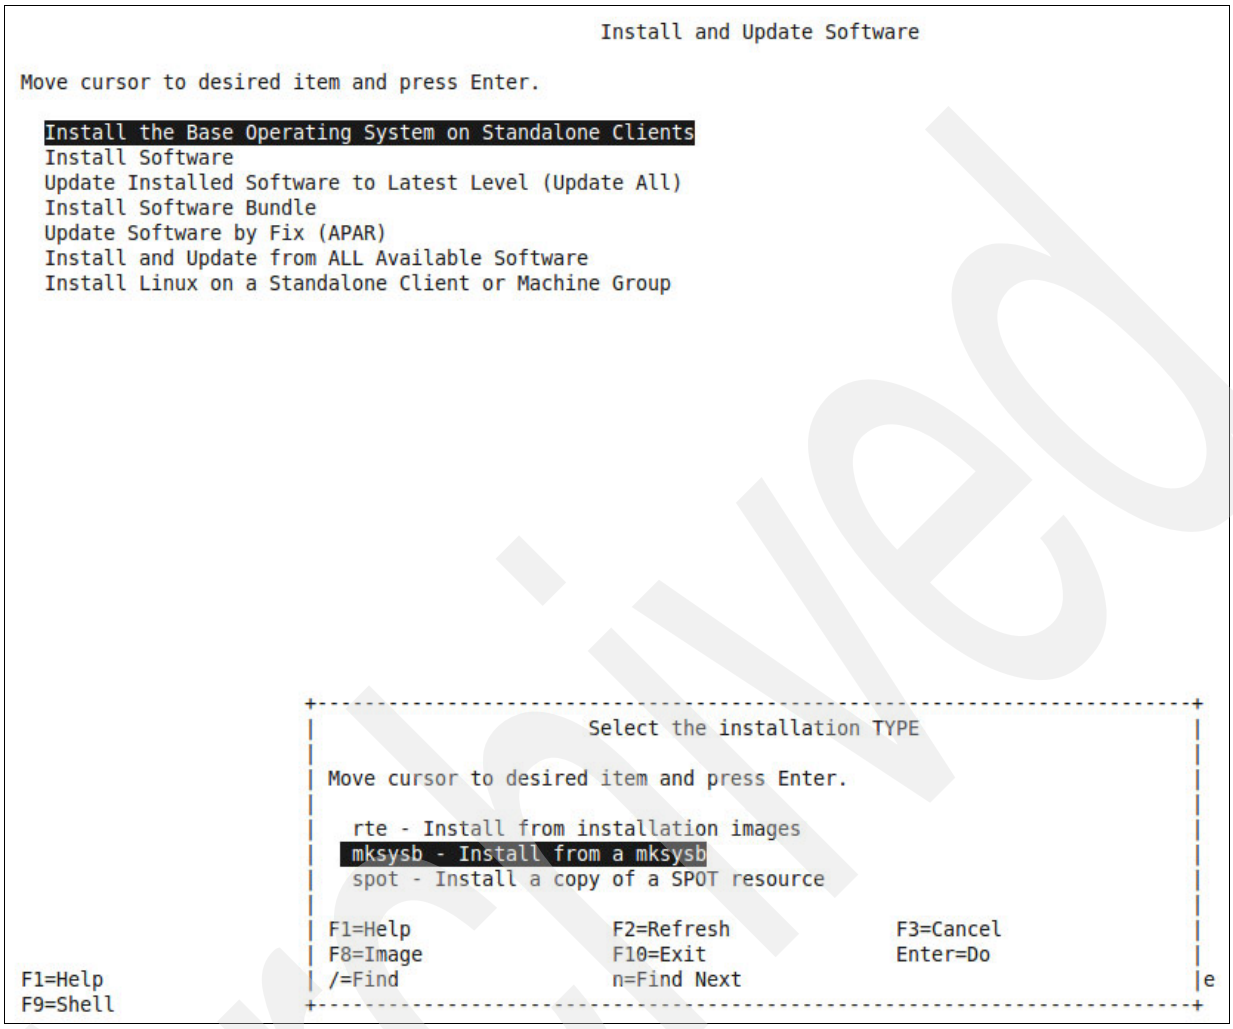

Figure 3-53 Select the installation TYPE submenu

This panel allows us to select from multiple mksysbs, but there was currently only one defined to NIM, as shown in submenu (Figure 3-54).

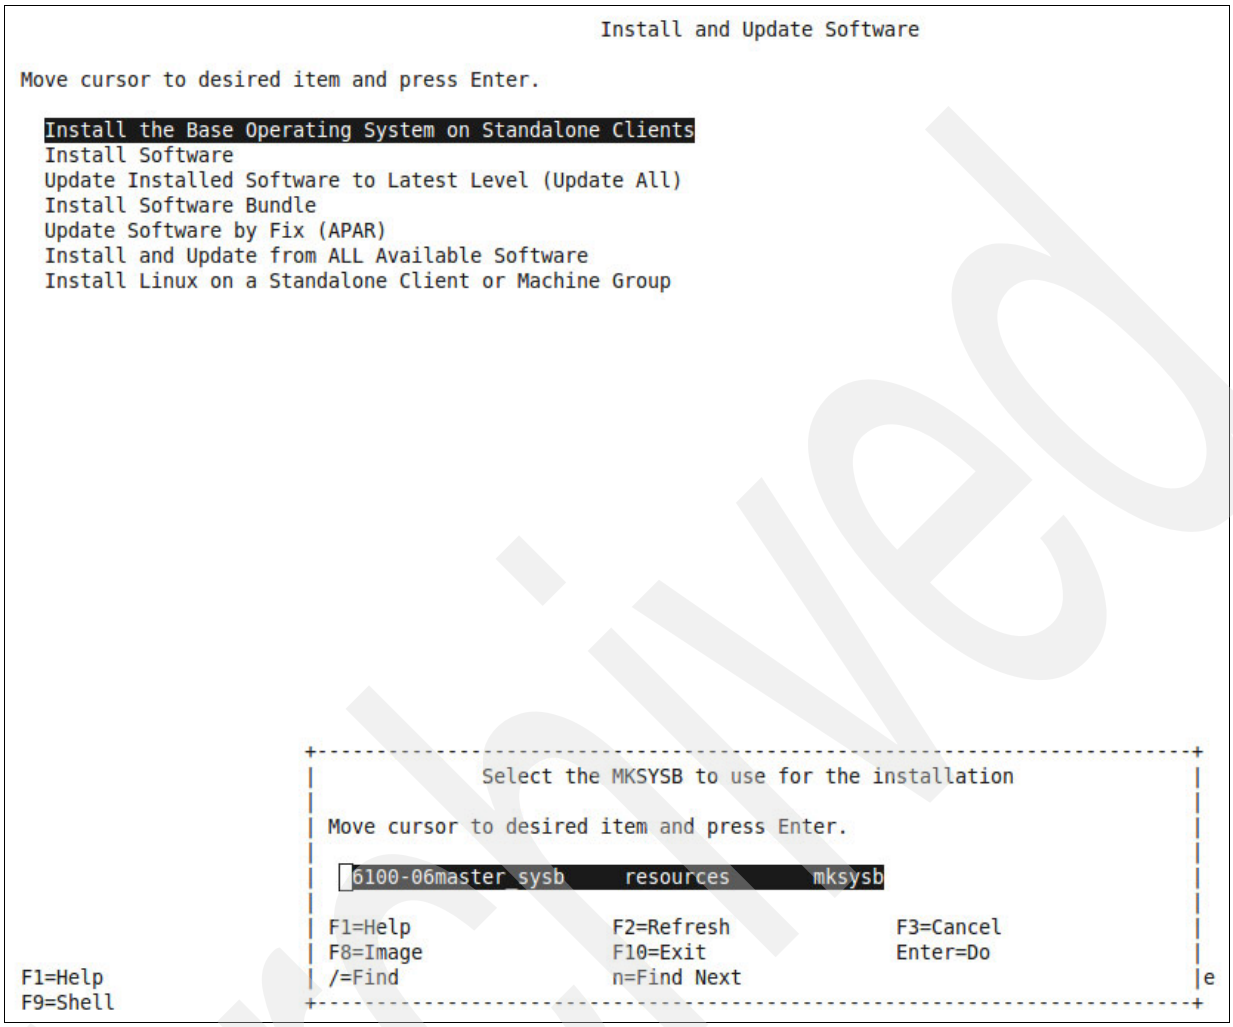

Figure 3-54 Select the MKSYSB to be used for the installation

Figure 3-55 shows the submenu that allows the selection from multiple SPOTs. Similar to the mksysb, there is only one defined; which was selected.

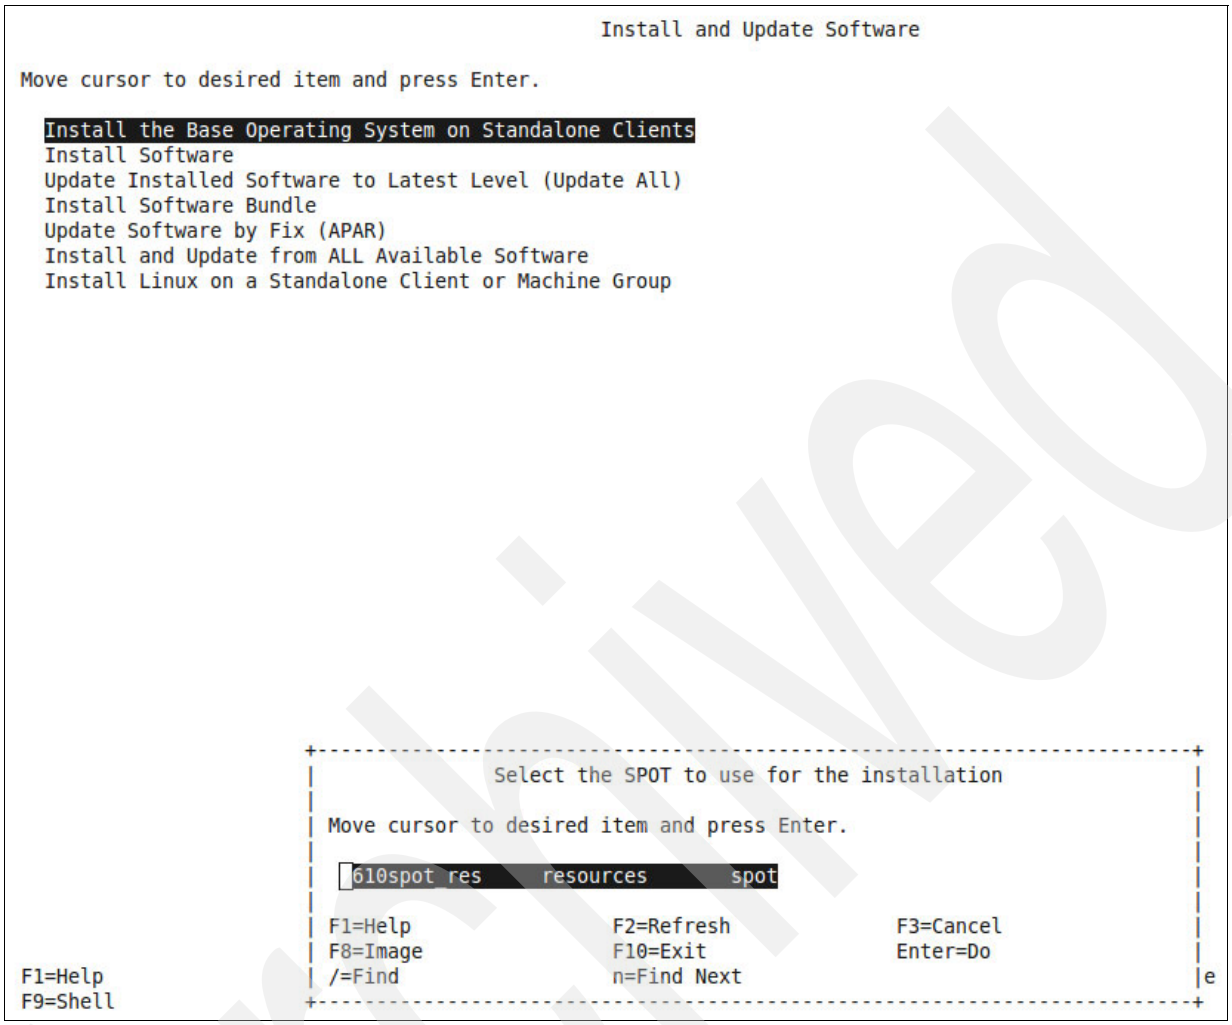

Figure 3-55 Select the SPOT to use for the installation

Figure 3-56 lists all of the required and optional data for the BOS install.

```
                        Install the Base Operating System on Standalone Clients

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

[TOP]                                                 [Entry Fields]
* Installation Target                                 p14n08
* Installation TYPE                                   mksysb
* SPOT                                               610spot_res
  LPP_SOURCE                                         []                                +
  MKSYSB                                             6100-06master_sysb

  BOSINST_DATA to use during installation            []                                +
  IMAGE_DATA to use during installation              []                                +
  RESOLV_CONF to use for network configuration       []                                +
  Customization SCRIPT to run after installation     []                                +
  Customization FB Script to run at first reboot     []                                +
    ACCEPT new license agreements?                   [no]                              +
  Remain NIM client after install?                   [yes]                             +
  PRESERVE NIM definitions for resources on          [yes]                             +
    this target?

  FORCE PUSH the installation?                       [no]                              +

  Initiate reboot and installation now?              [yes]                             +
    -OR-
  Set bootlist for installation at the               [no]                              +
     next reboot?

  Additional BUNDLES to install                      []                                +
    -OR-
  Additional FILESETS to install                     []                                +
    (bundles will be ignored)

  installp Flags
[MORE...23]

F1=Help                  F2=Refresh               F3=Cancel                F4=List
F5=Reset                 F6=Command               F7=Edit                  F8=Image
F9=Shell                 F10=Exit                 Enter=Do
```

Figure 3-56 Required and optional data for the BOS install

Figure 3-57 shows that the only option that was changed was the LPP_SOURCE to reflect what was defined in NIM.

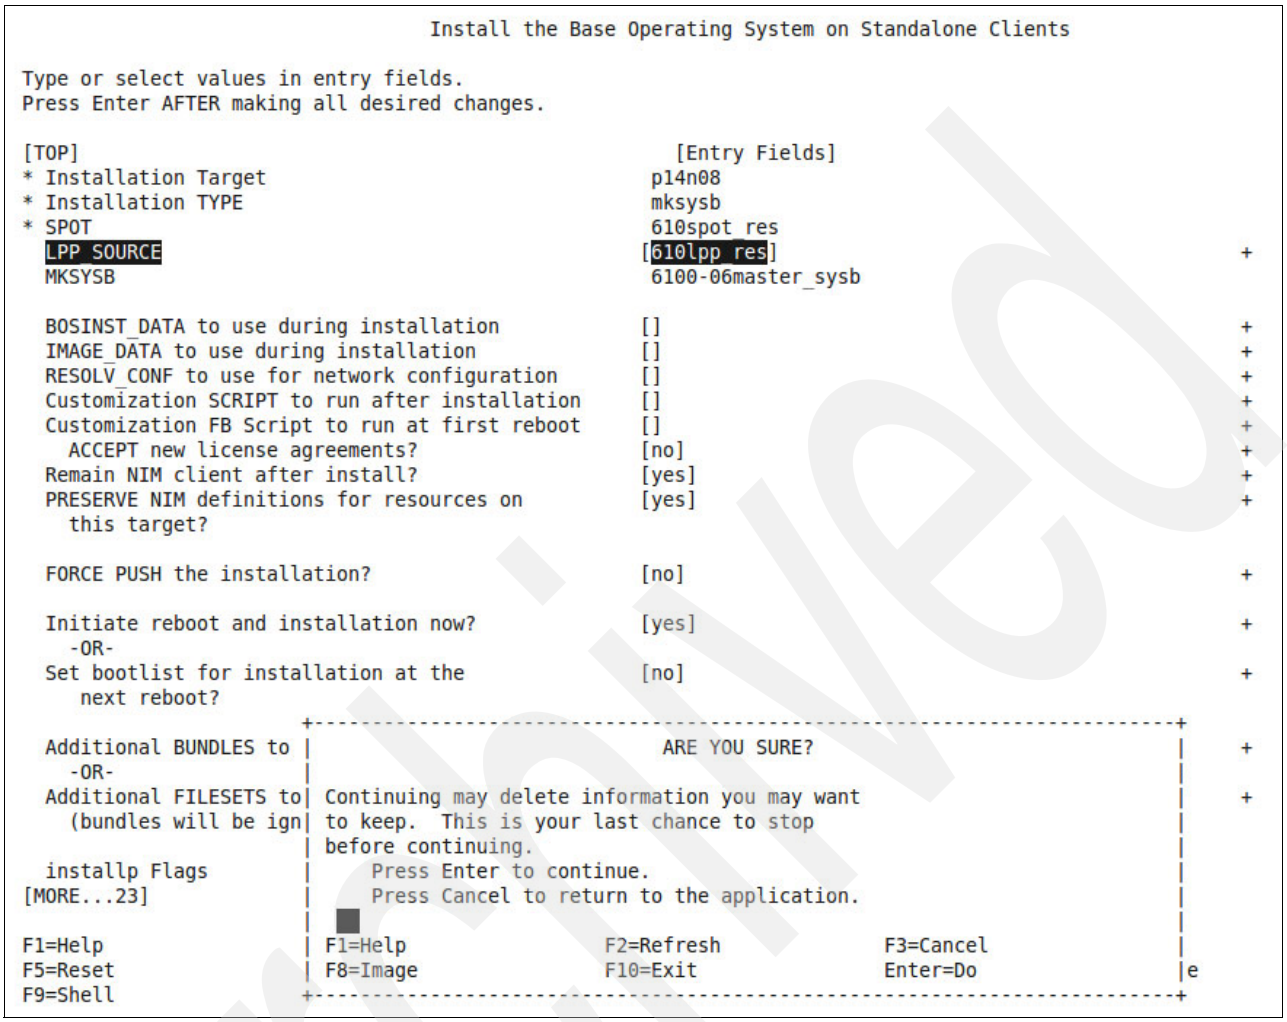

Figure 3-57 The only option that was changed was the LPP_SOURCE

Now that NIM is configured for the new LPAR, AIX can be installed.

The series of menus beginning with Figure 3-58 on page 72 show the input that is required to successfully boot the LPAR in installation mode. The responses or specifications are described in the text beneath each figure. Although not shown in these figures, each window had a response area provided at the bottom of the window. Actual responses vary depending on the implementation environment.

When the LPAR is first booted, the menu in Figure 3-58 is shown.

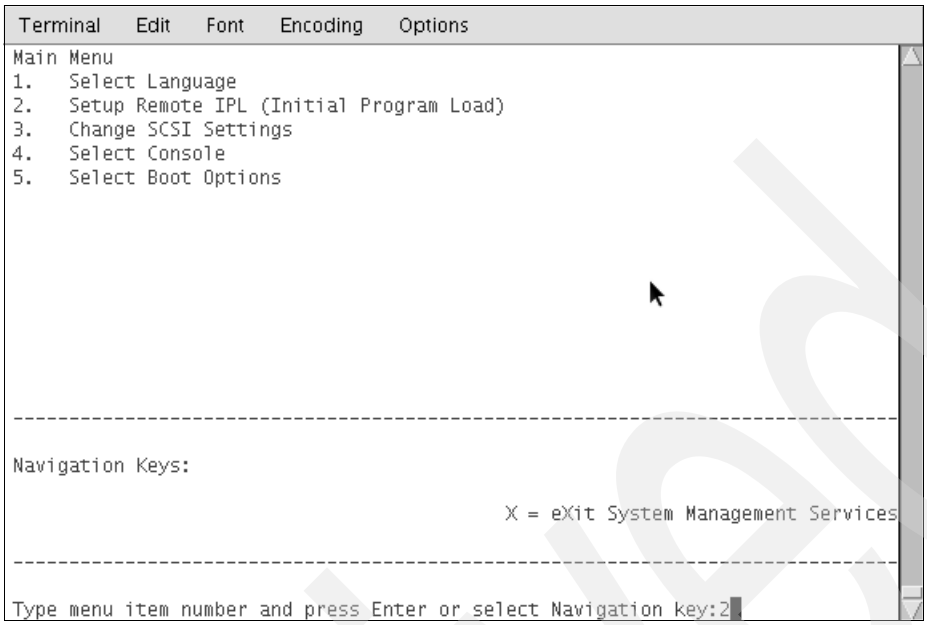

Figure 3-58 Setup Remote IPL (Initial Program Load) is selected

Because the BOS install is a remote install, we selected **Setup Remote IPL (Initial Program Load)**. Figure 3-59 shows selecting the Interpartition Logical LAN. For Figure 3-59, type `1`, which was the only Interpartition Logical LAN that was available.

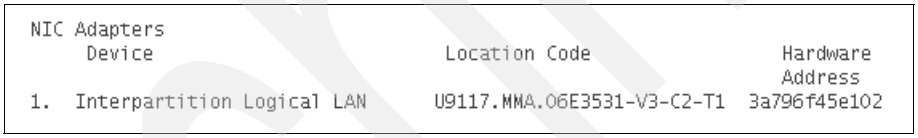

Figure 3-59 Select Interpartition Logical LAN

Figure 3-60 shows selecting the Internet Protocol Version. For Figure 3-60, we typed `1` to indicate IPv4.

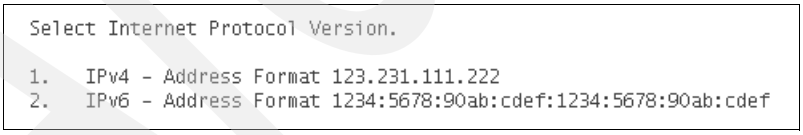

Figure 3-60 Specify IPv4

Figure 3-61 shows selecting the Network Service. For Figure 3-61, we chose BOOTP.

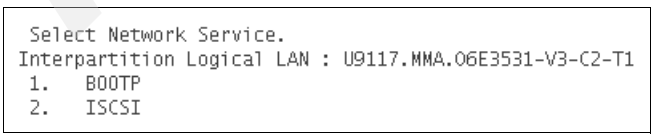

Figure 3-61 Choose BOOTP

Figure 3-62 on page 73 shows selecting the Network Parameters.

```
 Network Parameters
Interpartition Logical LAN : U9117.MMA.06E3531-V3-C2-T1
 1.   IP Parameters
 2.   Adapter Configuration
 3.   Ping Test
 4.   Advanced Setup: BOOTP
```

Figure 3-62 Choose IP Parameters

We selected the first option in Figure 3-62, which is 1. IP Parameters. This action results in the menu that is shown in Figure 3-63.

```
 IP Parameters
Interpartition Logical LAN : U9117.MMA.06E3531-V3-C2-T1
 1.   Client IP Address                    [129.40.157.8]
 2.   Server IP Address                    [129.40.157.7]
 3.   Gateway IP Address                   [129.40.157.254]
 4.   Subnet Mask                          [255.255.255.0]
```

Figure 3-63 Entering the IP Addresses and Subnet Mask of this LPAR

Figure 3-63 shows that we entered the IP Addresses and Subnet Mask of this LPAR. We also entered the Server IP Address that corresponds to the NIM Server. Figure 3-64 is shown next. We chose to run the Ping Test.

```
 Network Parameters
Interpartition Logical LAN : U9117.MMA.06E3531-V3-C2-T1
 1.   IP Parameters
 2.   Adapter Configuration
 3.   Ping Test
 4.   Advanced Setup: BOOTP
```

Figure 3-64 Choose Ping Test

Figure 3-65 (Option 3. Ping Test) verified that the LPAR was able to ping the IP addresses that were provided earlier. This action ensures communication between the LPAR and the NIM Server. Figure 3-65 shows that the ping test was successful.

```
.-----------------.
|  Ping  Success. |
`-----------------'
```

Figure 3-65 Ping test is successful

Next, we selected option 4. Advanced Setup: BootP, as shown in Figure 3-66.

```
 Network Parameters
Interpartition Logical LAN : U9117.MMA.06E3531-V3-C2-T1
 1.   IP Parameters
 2.   Adapter Configuration
 3.   Ping Test
 4.   Advanced Setup: BOOTP
```

Figure 3-66 Choosing Advanced Setup

After we selected the advanced setup, we accepted the default advanced setup specification values that are shown in Figure 3-67.

```
 Advanced Setup: BOOTP
Interpartition Logical LAN : U9117.MMA.06E3531-V3-C2-T1

 1.   Bootp Retries    5
 2.   Bootp Blocksize  512
 3.   TFTP  Retries    5
```

Figure 3-67 Default Advanced Setup specifications

After accepting the Default Advanced Setup Specifications that are shown in Figure 3-67, the window shown in Figure 3-68 opens.

```
FILE        : /tftpboot/p14n08.pbm.ihost.com
BLKSIZE     : 512
TFTP-RETRIES: 5
TFTP BOOT ----------------------------------------------------
Server IP.....................129.40.157.7
Client IP.....................129.40.157.8
Gateway IP....................129.40.157.254
Subnet Mask...................255.255.255.0
( 1  ) Filename................./tftpboot/p14n08.pbm.ihost.com
TFTP Retries..................5
Block Size....................512


FILE        : /tftpboot/p14n08.pbm.ihost.com
BLKSIZE     : 512
TFTP-RETRIES: 5
PACKET COUNT = 3600
```

Figure 3-68 Default Advanced Setup specifications

Figure 3-69 shows the beginning of the booting sequence.

```
 ******* Please define the System Console. *******

Type a 1 and press Enter to use this terminal as the
  system console.
```

Figure 3-69 Booting begins

Figure 3-70 shows a message to wait while the BOS loads.

```
                      Installing Base Operating System



        Please wait...


       Approximate     Elapsed time
     % tasks complete   (in minutes)


          12              0      7% of mksysb data restored.
```

Figure 3-70 Wait message and progress report

Figure 3-71 shows the software starting.

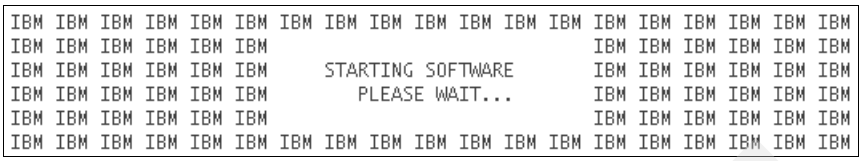

Figure 3-71 Software is starting

3.3.4 Installing the Systems Director

To install the Systems Director silently, first download and extract the Systems Director install packages. For silent installs, a response file needs to be provided to the install program. IBM Systems Director ships with an example response file. This response file contains all the potential options (as defaults) for installing Systems Director. For simple installation, you can use this response file without modification by using this command:

```
./dirinstall.server -r dirserv.rsp
```

The command runs until the successful installation message appears, as shown in Example 3-12. The sample output has been shortened for simplicity.

Example 3-12 Output of ./dirinstall.server -r dirserv.rsp

```
Starting server runtime...
The server runtime has been started.
Attempting to install features.

This installation log file can be found in /var/log/dirinst.log.
Installation of IBM Systems Director Server completed successfully.
```

After the IBM Systems Director Server install completes, the next step is to configure the Agent Manager. You execute the **/opt/ibm/director/bin/configAgtMgr.sh** script.

This script prompts for a series of responses, as shown in Example 3-13.

Example 3-13 Configuration script prompts and responses

```
Enter 1 to use the Agent Manager installed with this server (recommended)
Enter 0 to use an existing Agent Manager (advanced)
1
Enter the Resource Manager user ID that you would like to set for the Agent
Manager:
root
Enter the Resource Manager password to set for the Agent Manager:
<password entered but not displayed>
Verify the Resource Manager password to set for your Agent Manager:
<password entered but not displayed>
Enter the Agent Registration password to set for your Agent Manager:
<password entered but not displayed>
Verify the Agent Registration password to set for your Agent Manager:
<password entered but not displayed>
```

This configuration process generates extensive output. On completion, the following message is displayed:

```
Agent Manager configuration completed successfully.
```

Now, you can start the Systems Director, as shown in Example 3-14.

Example 3-14 Starting IBM Systems Director

```
/opt/ibm/director/bin/smstart
Starting IBM Director...
The starting process may take a while. Please use smstatus to check if the server
is active.
```

While Systems Director starts, you can use the **smstatus** command to verify the operating state. Example 3-15 shows the Inactive state.

Example 3-15 Using the smstatus command to query the status

```
# smstatus
Inactive
```

Example 3-16 shows the state after we issued the **smstart** command.

Example 3-16 State after smstart command was issued

```
bash-3.2# smstatus
Starting
```

Example 3-17 shows the status when Systems Director has been started and is operational.

Example 3-17 Systems Director status is Active

```
bash-3.2# smstatus
Active
```

To verify that the IBM Systems Director is completely functional, point your web browser at it. Figure 3-72 shows the login page, verifying that the Systems Director started successfully.

Figure 3-72 IBM Systems Director login page

3.3.5 Installing the VMControl plug-in

The VMControl plug-in for the IBM Systems Director provides many cloud-like operations for the POWER environment that is being built. The VMControl plug-in allows for the easy management of virtual servers, virtual appliances, and workloads that span multiple hardware platforms. For the purposes of this paper, our focus is on the Managed To AIX environment.

A response file must created to install the VMControl plug-in into IBM Systems Director using the silent install option. An example response file, which can be used as an example, is shipped with the installer. Example 3-18 shows an example of the required properties for the installation.

Example 3-18 Required properties for silent install of VMControl plug-in

```
INSTALLER_UI=Silent
LICENSE_ACCEPTED=true
START_SERVER=true
RT_TIMEOUT_MINUTES=default
```

After the response file is created, you can create the silent install, as shown in Example 3-19.

Example 3-19 VMControl Plug-in silent install

```
./Systems_Director_VMControl_Enterprise_Edition_AIX.sh -i silent
Preparing to install...
Extracting the JRE from the installer archive...
Unpacking the JRE...

Extracting the installation resources from the installer archive...
Configuring the installer for this system's environment...

Launching installer...

Preparing SILENT Mode Installation...

===============================================================================
IBM Systems Director VMControl                    (created with InstallAnywhere)
-------------------------------------------------------------------------------


 [==================|==================|==================|==================]
 [------------------|------------------|------------------|------------------]

Installation Complete.
```

You can verify the success of the install by logging in to IBM Systems Director and viewing the new panels, as shown Figure 3-73. The left navigation choices now contain VMControl, and the right side of Figure 3-73 shows that VMControl Enterprise Edition is now ready.

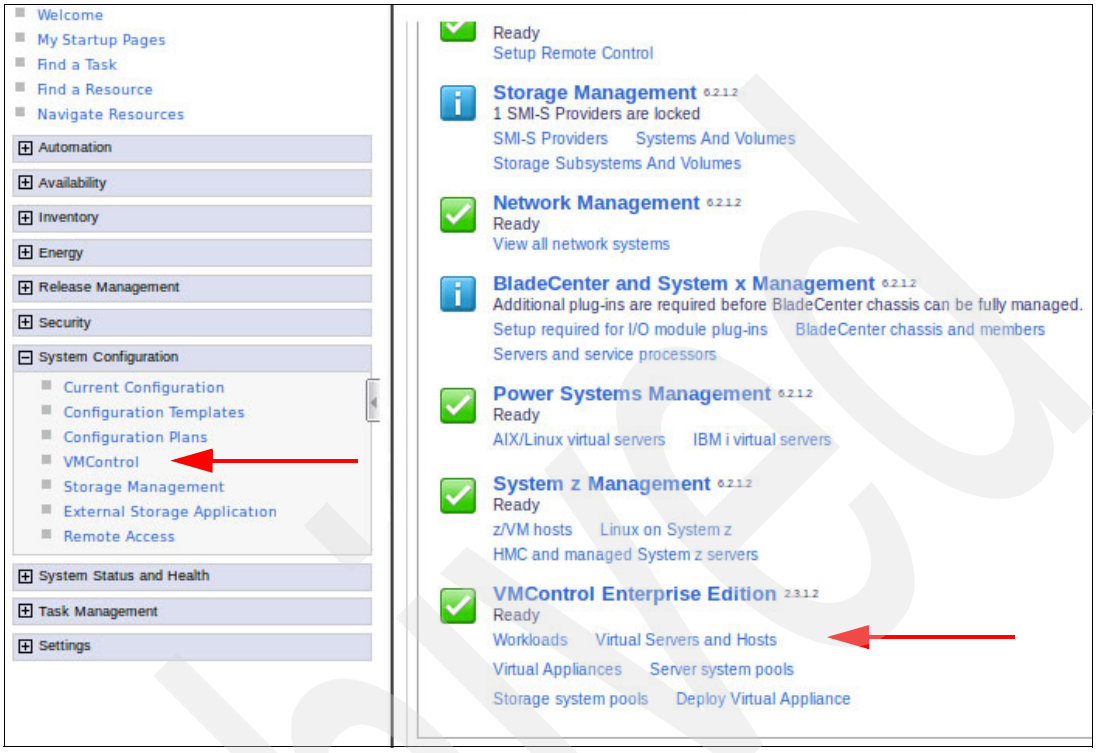

Figure 3-73 New VMControl functionality indicated by arrows

3.3.6 Installing and using the Active Energy Manager plug-in

The Active Energy Manager™ (AEM) plug-in works in concert with IBM Systems Director to provide a unified view of power consumption. The integration of AEM and the Systems Director is a new feature that is included with Systems Director 6.1 and AEM 4.1. Using the current versions of AEM and Systems Director simplifies the setup of the energy monitoring component of Cloud on the POWER environment by eliminating the need for a separate AEM console.

The particular Power hardware that is used (IBM Power 570 9117-MMA) only supports the Static Power Savings feature of AEM. The Static Power Savings feature reduces power usage by lowering the processor speed of a Systems Director-managed system (or resource group). The system consumes less energy, while still providing reasonable performance.

Installing the AEM plug-in

To install the AEM plug-in, log in to the Systems Director LPAR and launch the installable from the NFS-mounted directory. Then, respond to a series of GUI windows.

Figure 3-74 contains an image of the first window. The steps to perform are listed in the left column.

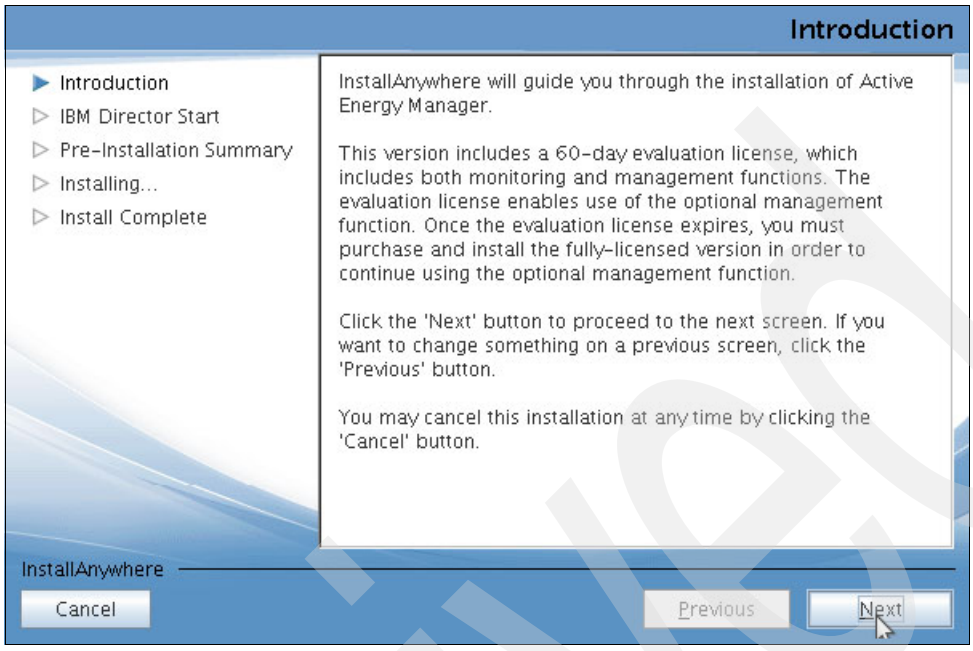

Figure 3-74 Introduction window

After clicking **Next**, the window that is provided in Figure 3-75 warns that AEM 4.3.1 requires IBM Systems Director 6.2.1 or later. Appendix B, “Using IBM Systems Director Update Manager” on page 173 describes the steps for using the IBM Systems Director Update Manager.

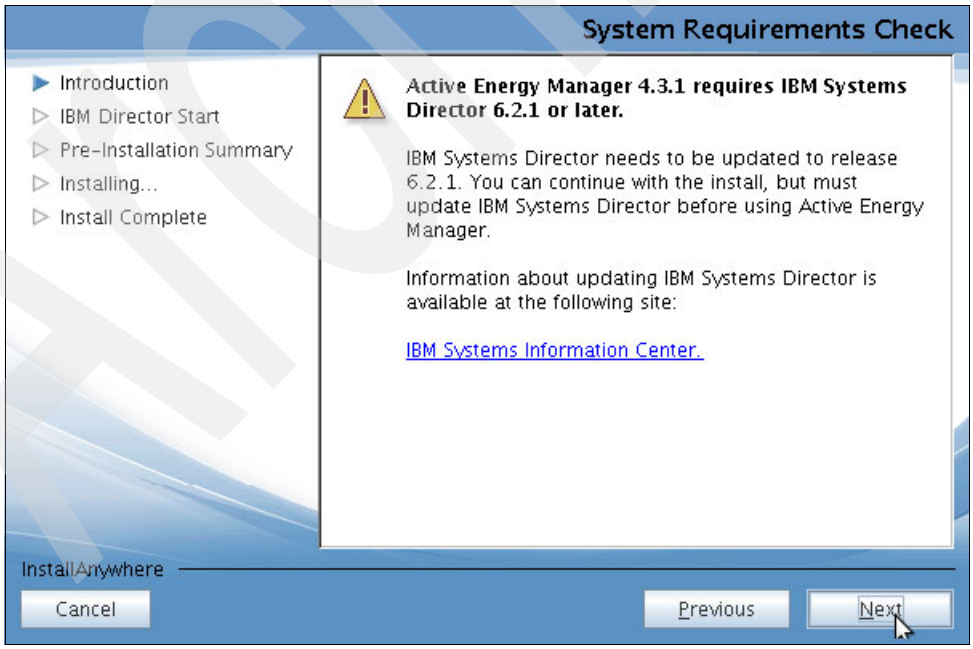

Figure 3-75 Active Energy Manager and IBM Systems Director release levels warning

After clicking **Next**, respond to the subsequent windows:

- Click Yes to accept the terms of the license agreement.
- Click Yes to enable default metering.
- Click Yes to restart IBM Systems Director automatically.

Then, as shown in Figure 3-76, the Pre-Installation Summary appears. At this point, click **Install**. The installation progress is indicated in Figures Figure 3-77 and Figure 3-78 on page 81.

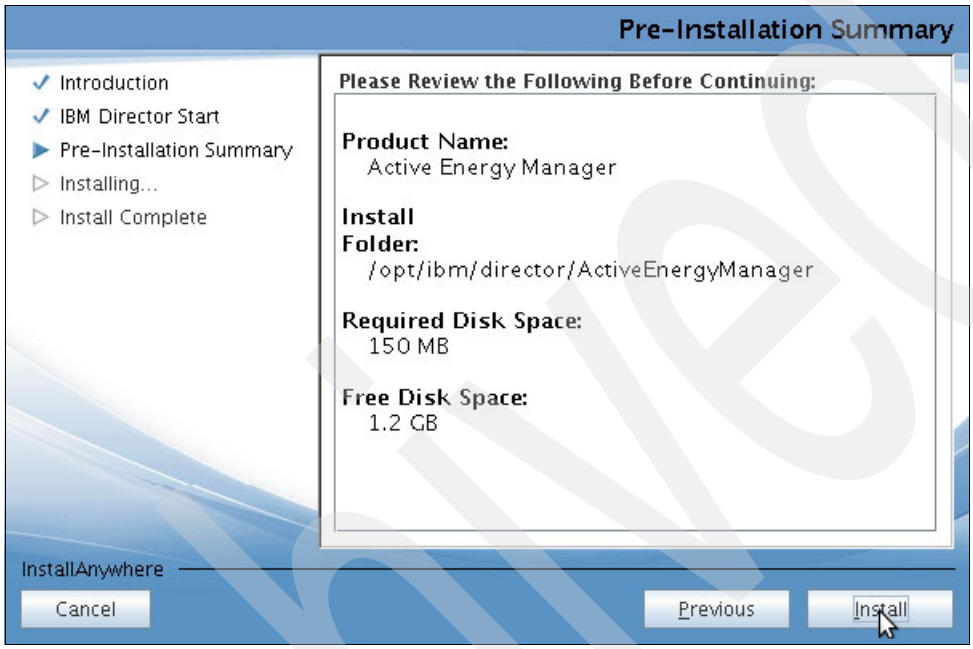

Figure 3-76 Pre-Installation Summary

A progress bar displays while AEM is installed (Figure 3-77).

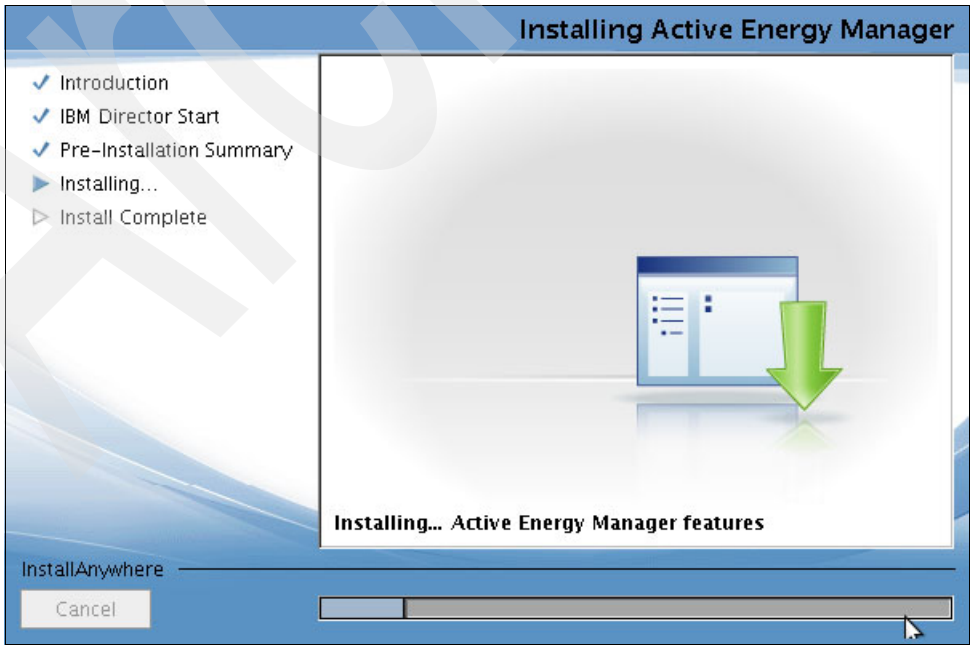

Figure 3-77 Installing Active Energy Manager features

When the installation is finished, Figure 3-78 opens.

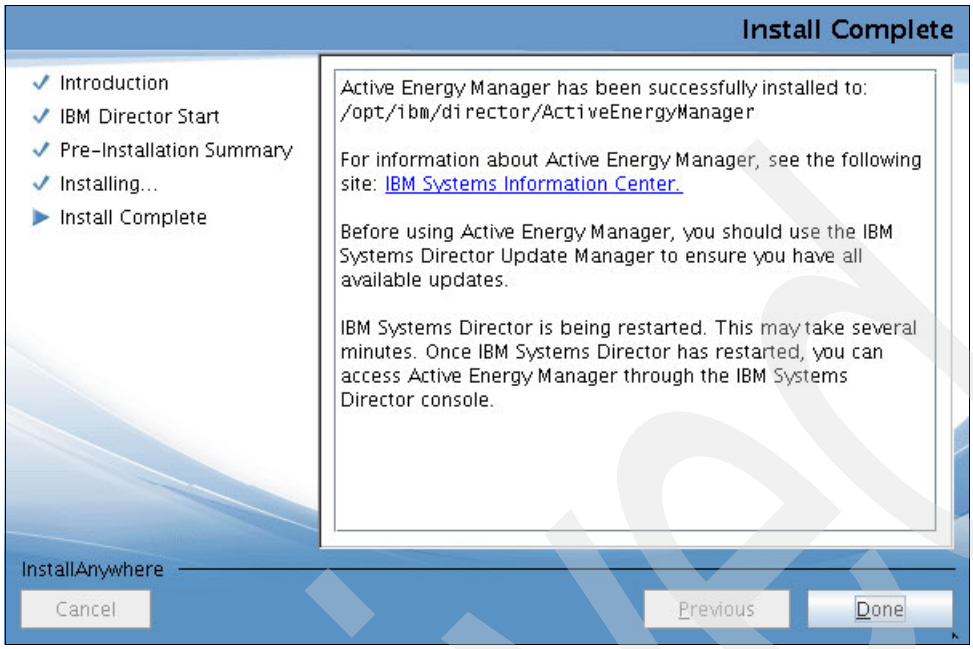

Figure 3-78 Install Complete

Using the AEM plug-in

When the IBM Systems Director Active Energy Manager starts, the following welcome window opens. For this implementation, we selected **English**.

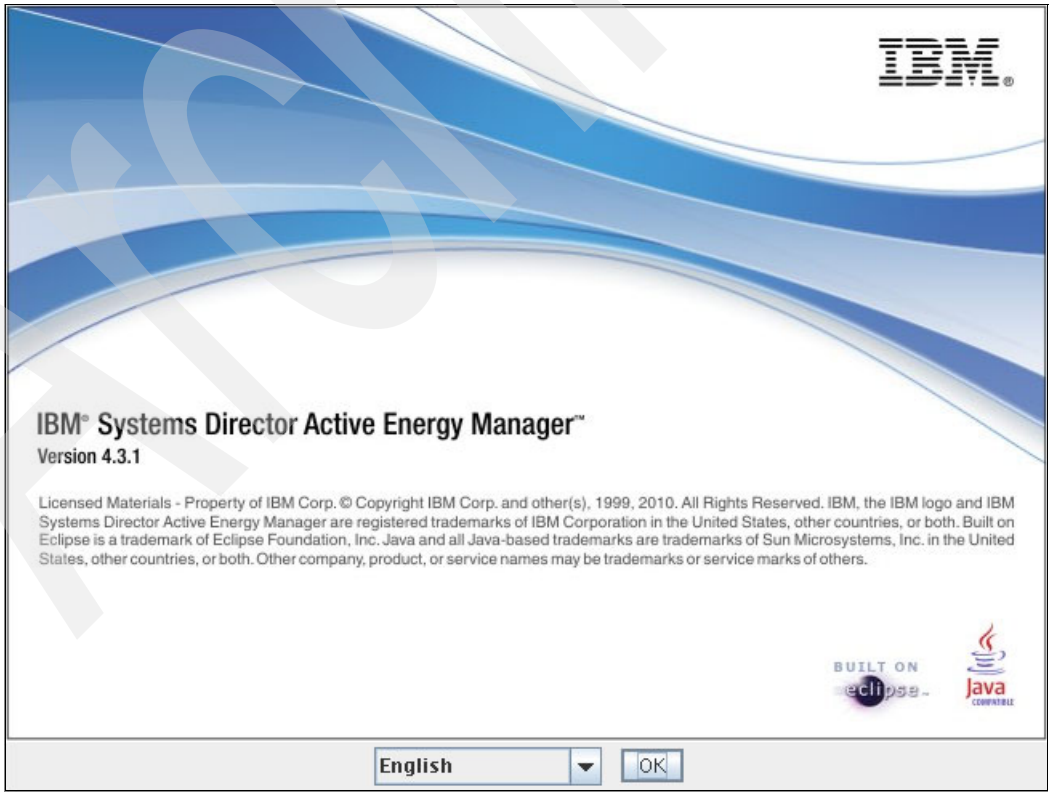

Figure 3-79 Active Energy welcome window

With the installation of the Active Energy Manager plug-in, a new navigation choice, Energy, appears in the IBM Systems Director navigation menu. Expanding this element presents a link to **Active Energy Manager**, as shown in Figure 3-80.

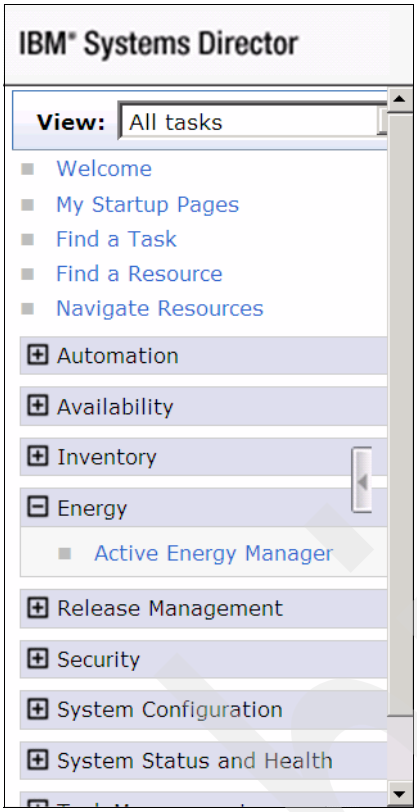

Figure 3-80 New navigation choice for Active Energy Manager

By clicking the **Active Energy Manager** link, the main Active Energy Manager interface is displayed, as shown in Figure 3-81. This interface displays information and options relative to energy monitoring and management. For this implementation, we selected the **Active Energy Managed Resources (2)** link, which displayed content that is shown in Figure 3-82 on page 83.

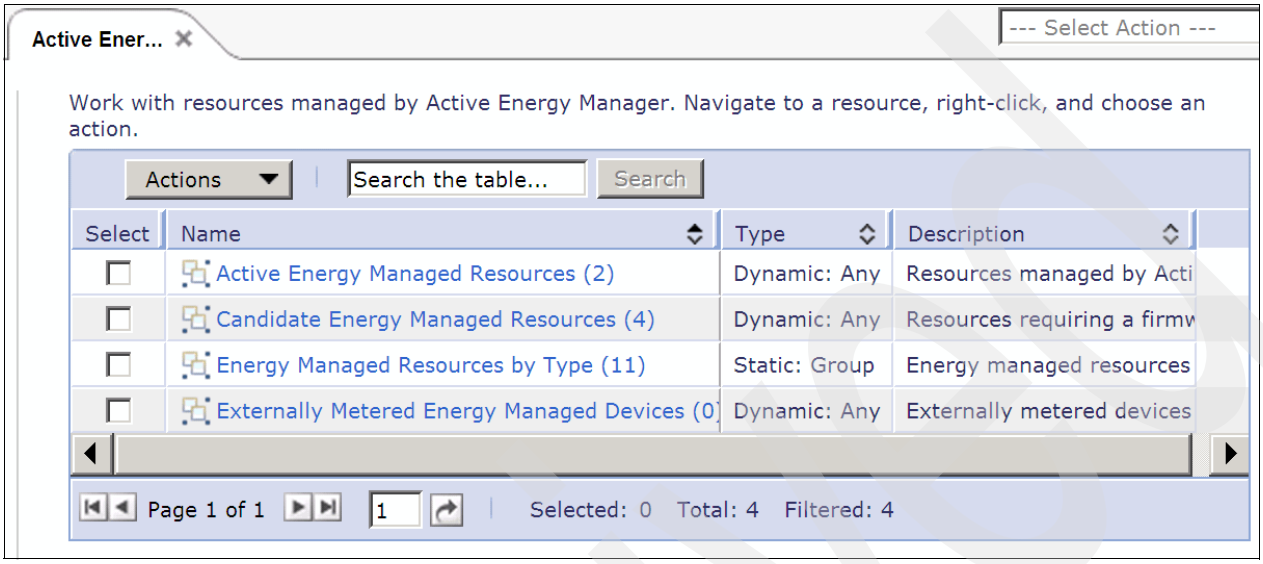

Figure 3-81 Active Energy Manager monitor options

Figure 3-82 shows the result of the selection that we made in Figure 3-81 on page 83. Note that the two POWER6 boxes are listed. We refer to these servers as Big Box and Little Box.

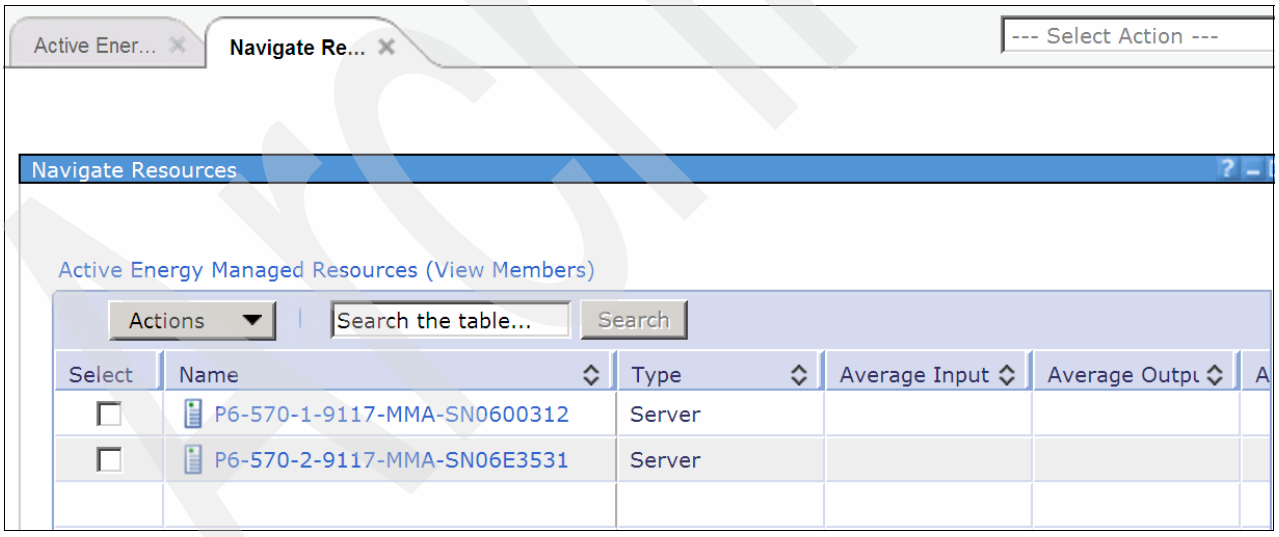

Figure 3-82 Navigate Resources via Active Energy Manager

After checking both boxes in the Select column in Figure 3-82, click and navigate through these choices **Active Energy Managed Resources** → **Energy** → **Manage Power** → **Power Savings**. Figure 3-83 shows these options.

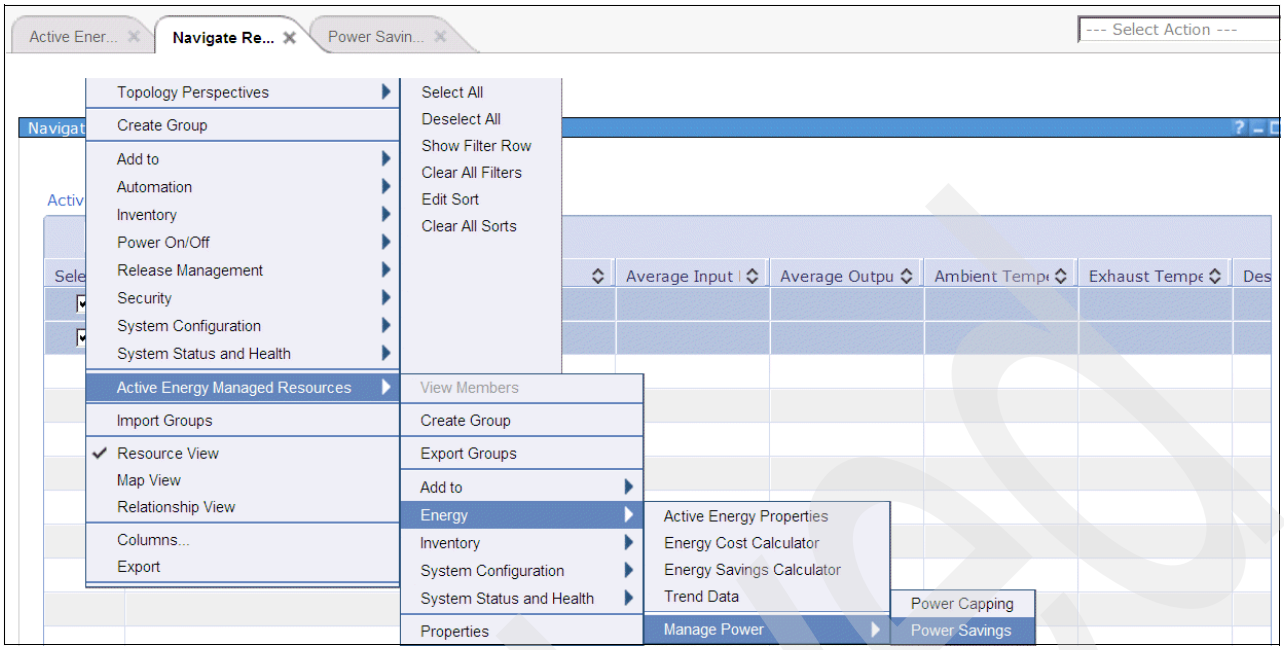

Figure 3-83 Navigating to Active Energy Manager Power Savings menu

Figure 3-84 shows the result of choosing Power Savings in Figure 3-83 on page 84. Active Energy Manager indicates that only one of the two boxes supports power savings.

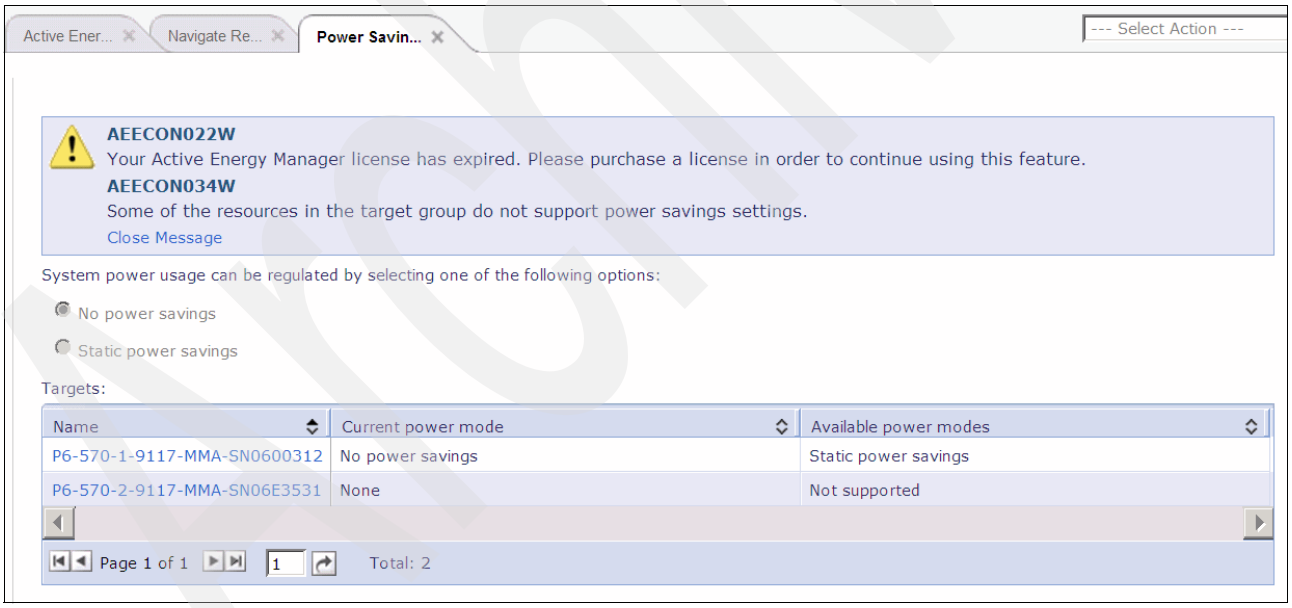

Figure 3-84 Active Energy Manager Power Savings tab

The two servers are not homogeneous with respect to power savings, which you also can confirm if you attempt to enable Static Power Savings, as shown in Figure 3-85.

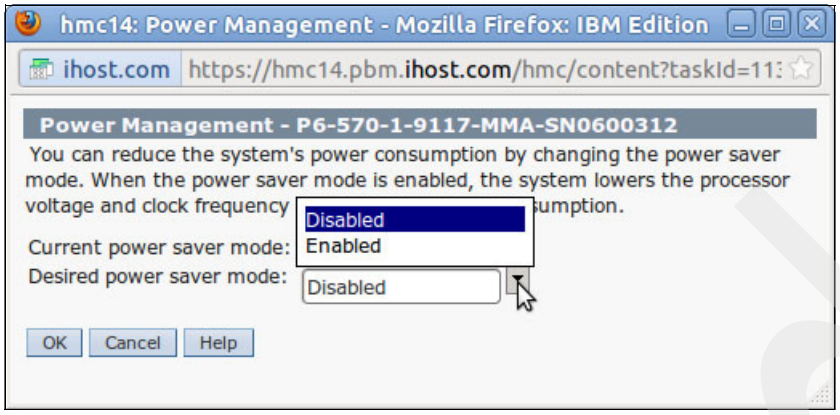

Figure 3-85 Active Energy Manager Power Management: Changing modes

When attempting to change the power savings mode of the two boxes by using the dialog that is shown in Figure 3-85, we received two separate responses. One attempt was successful and the other attempt was not, as shown in Figure 3-86 on page 85 and Figure 3-87 on page 85. One attempt failed, because the two boxes had separate levels of firmware.

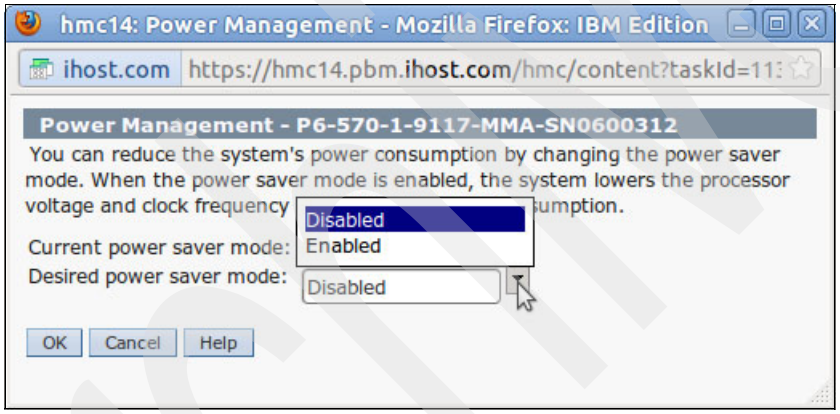

Figure 3-86 Active Energy Manager power settings successful

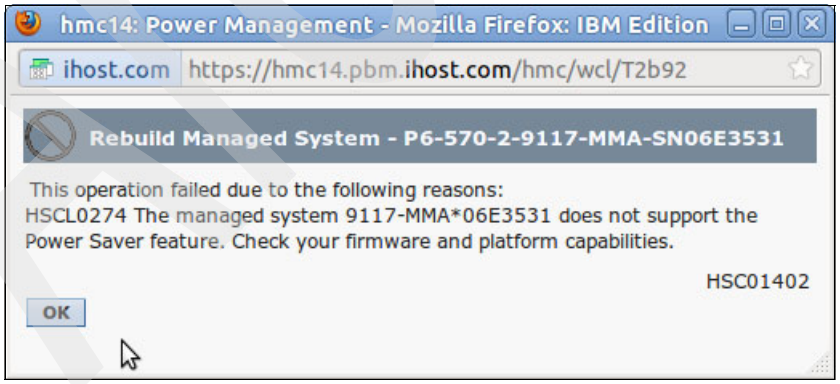

Figure 3-87 Active Energy Manager power settings unsuccessful

3.4 Installing CIM

The Common Information Model (CIM) is an open standard that defines how to represent the managed elements in an IT environment as a common set of objects and the relationships between them. CIM allows consistent management of these managed elements, regardless of their manufacturer or provider.

The storage environment that is implemented and described in this document uses CIM as the interface for both the IBM Tivoli Storage Productivity Center (TPC) and the IBM Systems Director. Before CIM can be used as an interface, the CIM LPAR must be prepared using the IBM Systems Director. Then, the CIM providers must be installed. Providers are designed so that they can be deployed in various environments, usually with no modifications.

In the context of this implementation, the two CIM providers were needed to support access to the DS4800 and Brocade. Each installed CIM server package was configured with its own port. One port was associated with the DS4800 provider, and the other port was associated with the Brocade provider.

3.4.1 Preparing the CIM LPAR using Systems Director

Use the following steps, as shown in Figure 3-88, to get to the Create Virtual Server Wizard to create the CIM virtual server. Using the Systems Director allows the use of automation to simplify the process of creating an LPAR and installing AIX. The advantage of this approach is that the steps that were described earlier in 3.3.3, “Using NIM to install AIX” on page 61 are no longer needed.

Follow these steps to get to the Create Virtual Server wizard to create the CIM virtual server:

1. Expand the **Inventory** menu (left navigation) in the Director window.
2. Expand the **VIEWs** window.
3. Click **Virtual Servers and Hosts**.
4. Check the box to the left of P6-570-2-9117-MMA-SN06E3531, which is the name of the Little Box.
5. Expand the **Actions** Menu, hover over **System Configuration**, and click **Create Virtual Server**.

After creating the virtual server, log in to AIX to complete the CIM installation.

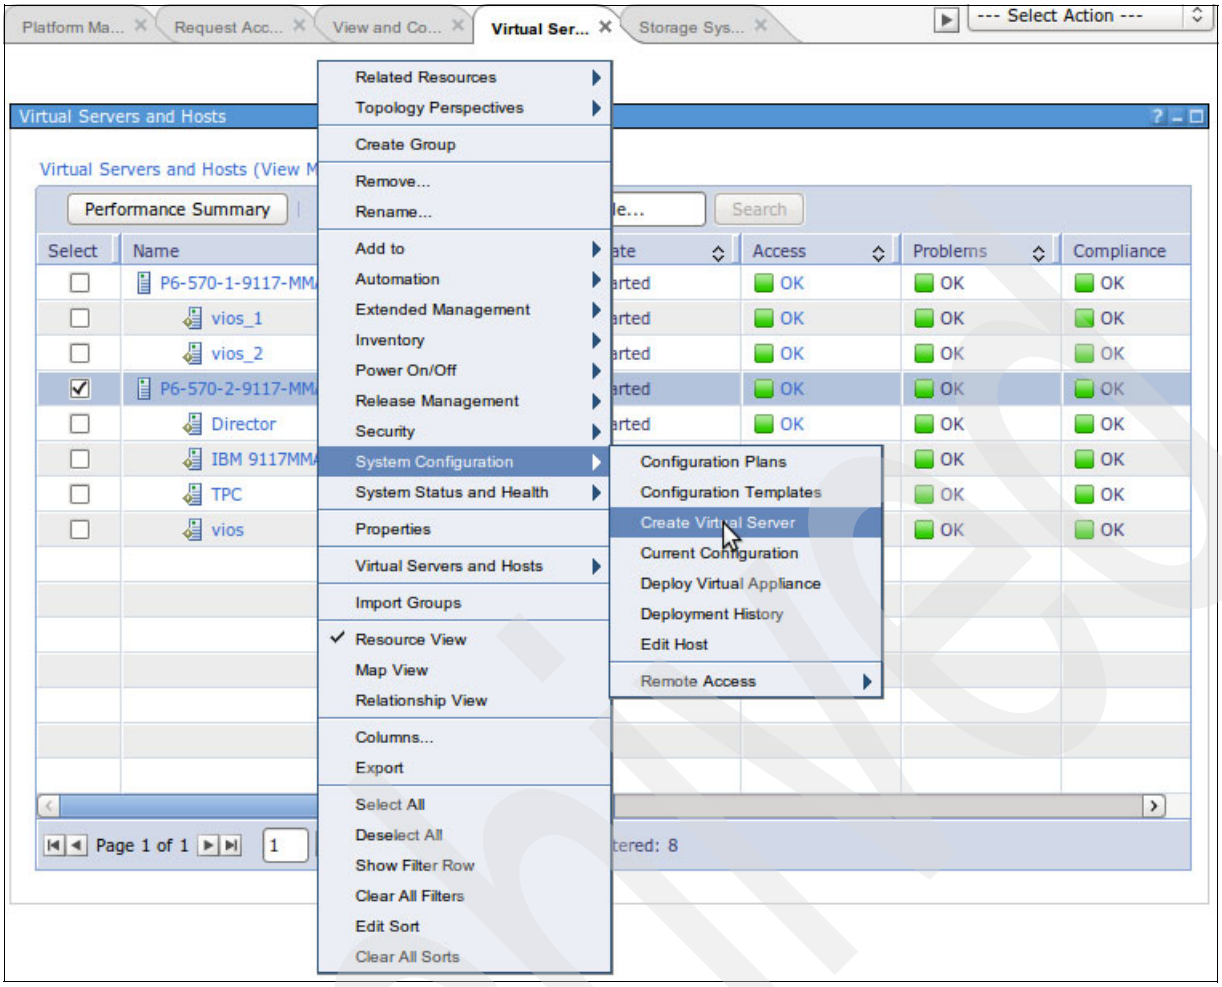

Figure 3-88 Navigating to the Create Virtual Server window on IBM Systems Director

Figure 3-89 opens.

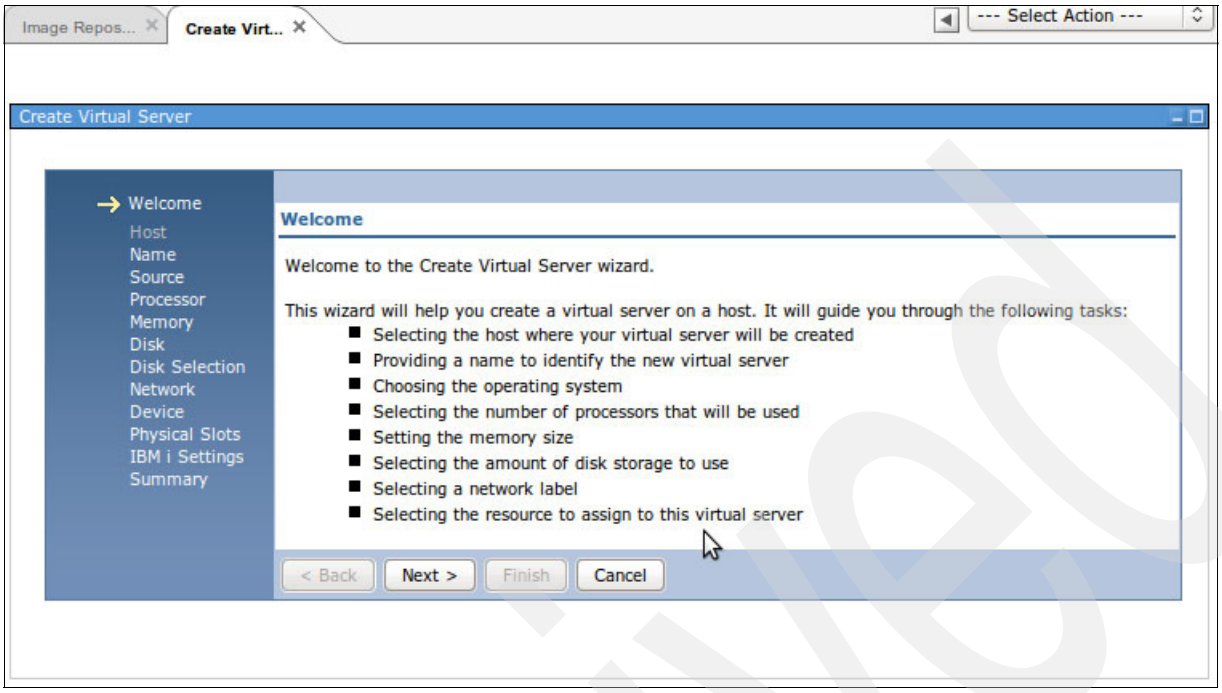

Figure 3-89 Create Virtual Server Welcome

In Figure 3-90, you type the name of the virtual server that you want to create.

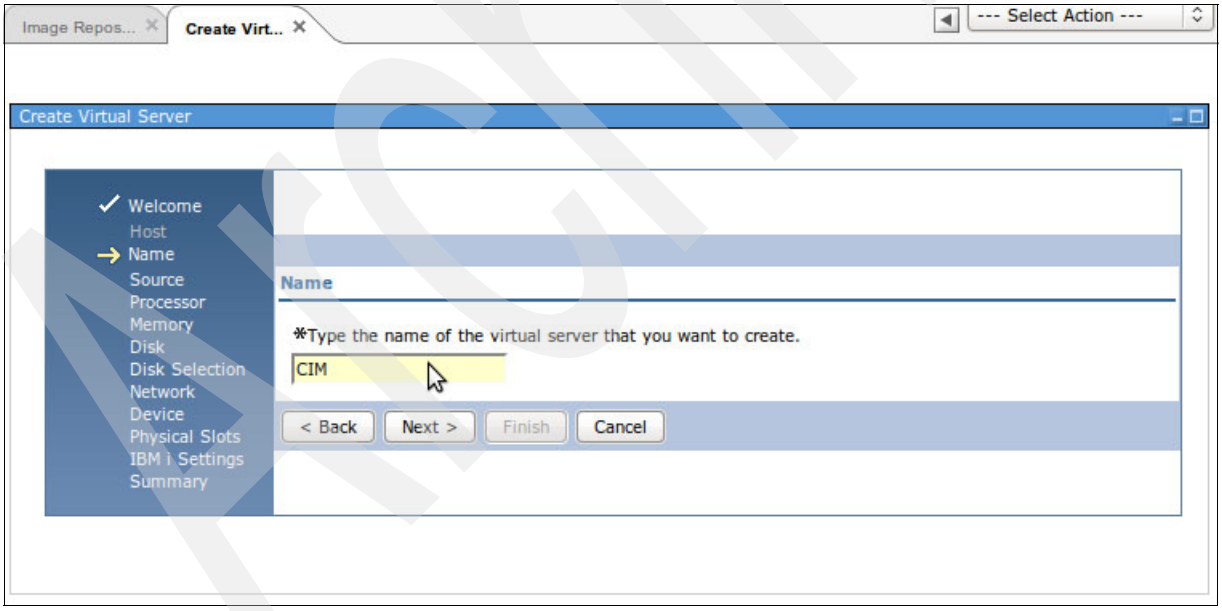

Figure 3-90 Type the name of the virtual server that you want to create

In Figure 3-91, you select the source that will be used as the basis for this new virtual server.

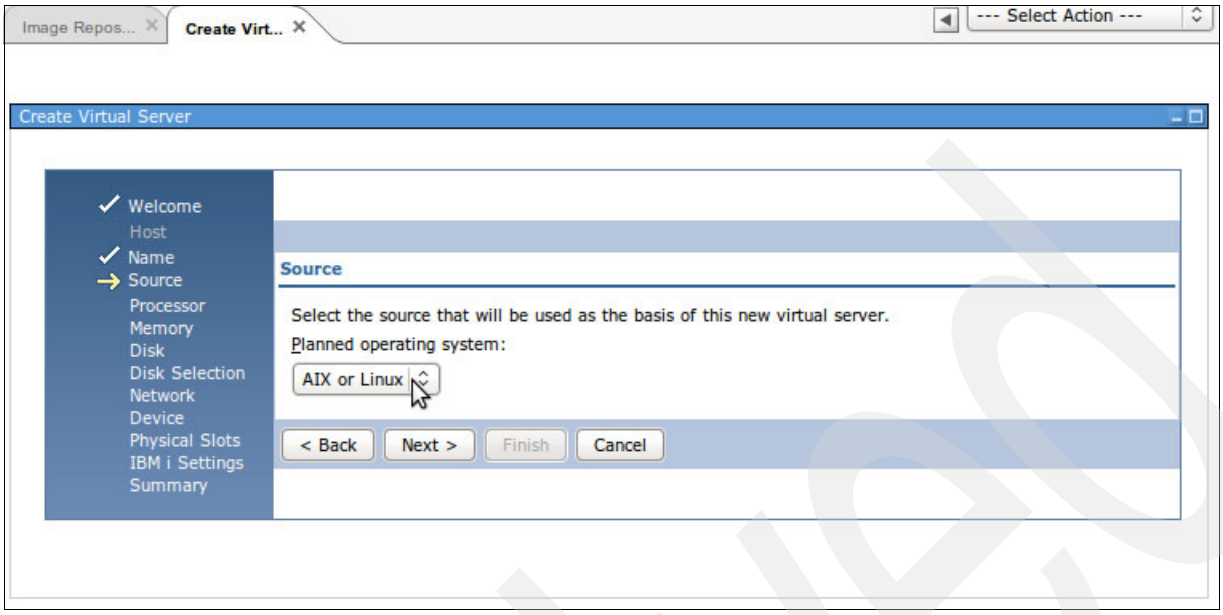

Figure 3-91 Select the source to use as the basis of this new virtual server

3.4.2 Installing Common Information Model providers

CIM: A complete discussion of CIM, CIM-XML, and Storage Management Initiative Specification (SMI-S) is beyond the scope of this document. You can obtain additional CIM-related online references in “Related publications” on page 177.

To contribute to the understanding of the terms and tasks that are addressed in this section, we briefly describe CIM, CIM-XML, and SMI-S. The Common Information Model (CIM) is a standardized way to describe objects and their relationships. CIM-XML defines a method for interacting with CIM through HTTP calls using XML. The Storage Management Initiative Specification (SMI-S) extends CIM to provide a standardized interface for the management, configuration, and monitoring of storage systems. The storage environment that was created for this paper makes use of these technologies to have an automated way to manage the storage infrastructure. The software that is discussed in this section provides a common interface for the IBM Tivoli Storage Productivity Center, a key component of the Managed Through environment. The Managed From environment, specifically IBM Systems Director, uses Tivoli Storage Productivity Center and CIM to provide the storage components of the Infrastructure as a Service (IaaS).

For the purposes of this document, note the following points:

- Tivoli Storage Productivity Center utilizes an SMI-S interface to communicate with the storage hardware, including the storage area network (SAN) switch.
- The SMI-S interface is presented by a set of providers and a CIM server.
- The CIM providers are supplied by the device manufacturer in an installable bundle, which includes a CIM server that is preconfigured with the appropriate providers. This CIM server provides CIM support for both the DS4800 and the SAN switch in the configuration.

The remainder of this section describes the tasks that are associated with the installation of Common Information Model (CIM) providers on the CIM LPAR. Each provider was needed to

provide the capacity to access storage components, a DS4800 storage array, and a SAN switch, as shown in Figure 3-92.

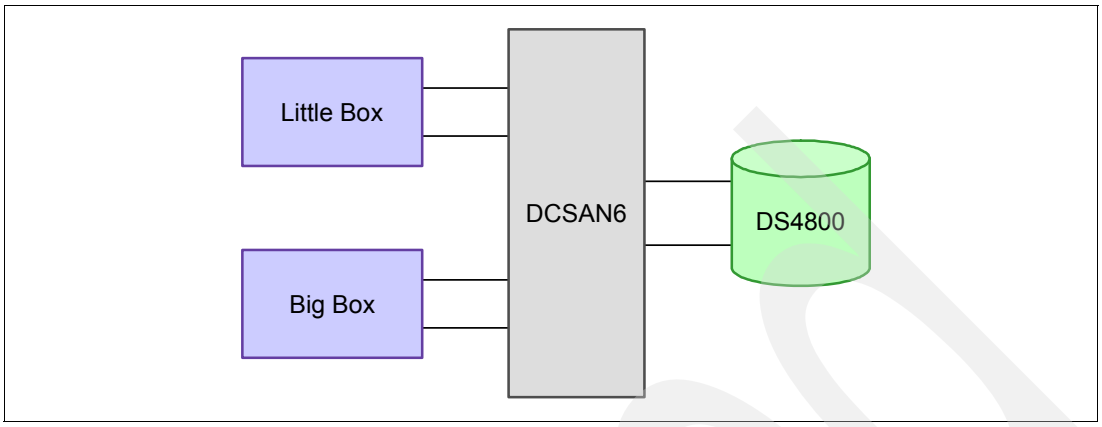

Figure 3-92 Accessing storage components

Because each set of providers is packaged with its own CIM server, both servers needed to be configured so that they did not interfere with each other, as well as with the CIM server that was installed with AIX.

Each CIM server package was installed into its own directory, and configured with its own port. Port 5988 was chosen for the DS4800 provider, and port 55988 was chosen for the Brocade provider. Other port specifications can be used as long as they do not conflict. The Tivoli Storage Productivity Center configuration must accurately reflect the ports that were chosen.

Steps for installing CIM providers for DS4800

To install the CIM providers for the DS4800, download the provider from LSI's website:

`http://www.lsi.de.com/storage_home/products_home/external_raid/management_software/smis_provider/index.html`

After downloading this package from the previous link, you must unpack it:

```
tar -xf AIX_Installer_10.50.G0.04.tar
```

After you unpack the installation package, run the installation script, as shown in Example 3-20.

Example 3-20 Running the CIM provider installation script

```
# ./install
+-----------------------------------------------------------------------------+
                    Pre-installation Verification...
+-----------------------------------------------------------------------------+
Verifying selections...done
Verifying requisites...done
Results...

--- output deleted for brevity ---

Restoring files, please wait.
919 files restored.
1623 files restored.
2978 files restored.
```

```
4079 files restored.
Finished processing all filesets.  (Total time:  4 mins 26 secs).

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

Installation Summary
--------------------
Name                          Level           Part        Event       Result
-------------------------------------------------------------------------------
EngenioProvider-10.50.6004. 10.50.6004.0    USR         APPLY       SUCCESS
EngenioProvider-10.50.6004. 10.50.6004.0    ROOT        APPLY       SUCCESS
```

Example 3-20 on page 90 shows the post installation tasks that are provided by **ProviderUtil**, where the CIMOM *username*, *password*, *input port*, *input operation*, *input IP,* and *Input Array password* are specified.

To configure the CIM provider for the DS4800, the **ProviderUtil** script is run. The responses to the script's prompts are shown in Example 3-21.

Example 3-21 Post-installation tasks: CIM providers for DS4800

```
# /opt/engenio/SMI_SProvider/bin/ProviderUtil
CIMOM username: root
CIMOM Password: ***password not displayed***
Input CIMOM Port [5988]:
Input operation:
'add' for AddDevice,
'rem' for RemoveDevice: add
Input IP or hostname for array: 129.40.150.139
Input Array password: ***password not displayed***
Attempting extrinsic method call.
The extrinsic call succeeded.
```

Steps for installing CIM providers for the SAN

To install the CIM providers for the SAN switch, download and unpack the providers from the Brocade website:

```
http://www.brocade.com/services-support/drivers-downloads/smi-agent/application_matrix.page
```

After downloading this package from the previous link, you must unpack it:

```
# tar -zxf smiagent120.9.0.tar.gz
```

Example 3-22 shows the use of the installation script.

Example 3-22 Installing CIM providers for the SAN

```
edit SilentInstallation.properties setting these parameters:
httpPort=55988, httpsPort=55989
run install as follows:
# ./install.bin -f ./SilentInstallation.properties
Preparing to install...
Extracting the JRE from the installer archive...
Unpacking the JRE...
```

```
Extracting the installation resources from the installer archive...
Configuring the installer for this system's environment...
Launching installer...
Preparing SILENT Mode Installation...
===============================================================================
Brocade SMI Agent                (created with InstallAnywhere by Macrovision)
-------------------------------------------------------------------------------
===============================================================================
Installing...
-------------
```

Example 3-23 configures the Brocade CIM server to start automatically at boot.

Example 3-23 Post-installation tasks: CIM providers for the SAN

```
install agent service
# cd /opt/SMIAgent/agent/server/jserver/bin
# ./install_agent_service
```

The configuration specifications for Example 3-24 were provided in Example 3-18.

Example 3-24 Configuration steps: CIM providers for the SAN

```
edit SilentInstallation.properties setting these parameters:
httpPort=55988, httpsPort=55989
```

3.5 Installing IBM Tivoli Storage Productivity Center

The IBM Tivoli Storage Productivity Center (TPC), in the context of this implementation, provides the functionality for the Discovery and Inventory. In the following subsections, we explain the process of discovering entities, such as storage and servers, with many screen captures in the order that they are encountered. In advance, you must prepare the Tivoli Storage Productivity Center LPAR by using the IBM Systems Director. Then, prior to being able to use Tivoli Storage Productivity Center for Discovery and Inventory, you must install and configure Tivoli Storage Productivity Center.

3.5.1 Preparing the Tivoli Storage Productivity Center LPAR using Systems Director

After replicating the steps that are provided in 3.4.1, “Preparing the CIM LPAR using Systems Director” on page 86, a new LPAR is available to be used for Tivoli Storage Productivity Center. Systems Director allows you to take advantage of automation, simplifying the process of creating an LPAR and installing AIX. The advantage is that it is no longer necessary to complete the steps that are described earlier in 3.3.3, “Using NIM to install AIX” on page 61.

3.5.2 Installing Tivoli Storage Productivity Center

After the Systems Director finalizes the installation of AIX on the LPAR, the installation of Tivoli Storage Productivity Center can begin. The first step is to install DB2®.

After starting the DB2 installation program, the Welcome window is presented (Figure 3-93).

Figure 3-93 DB2 installation program Welcome pane

To prevent potential installation issues, the installer presents links to the installation prerequisites, as shown in Figure 3-94.

Figure 3-94 DB2 installation prerequisites

After verification that the system meets all of the installation prerequisites, we clicked **Install a Product**, as shown in Figure 3-95.

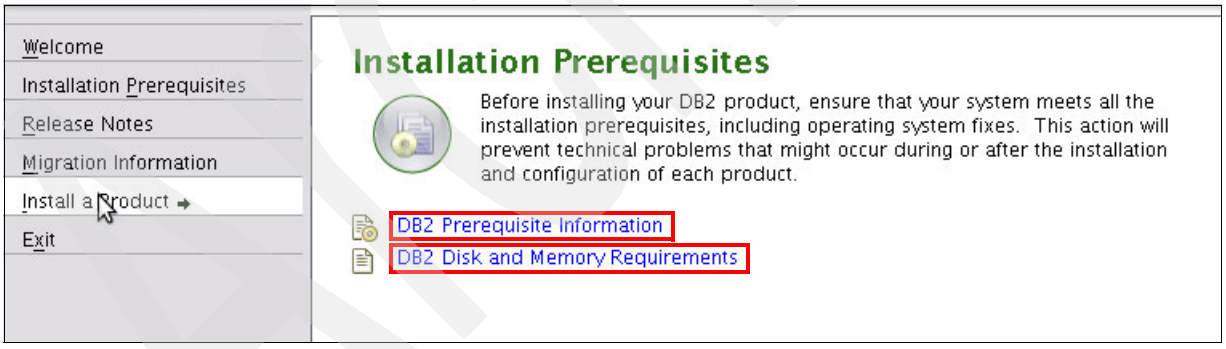

Figure 3-95 Choosing the Install a Product option on the left to install DB2

The DB2 installation program automatically detects if a previous version of the software is already installed. Because this LPAR was newly created, we selected **Install New**, as shown in Figure 3-96.

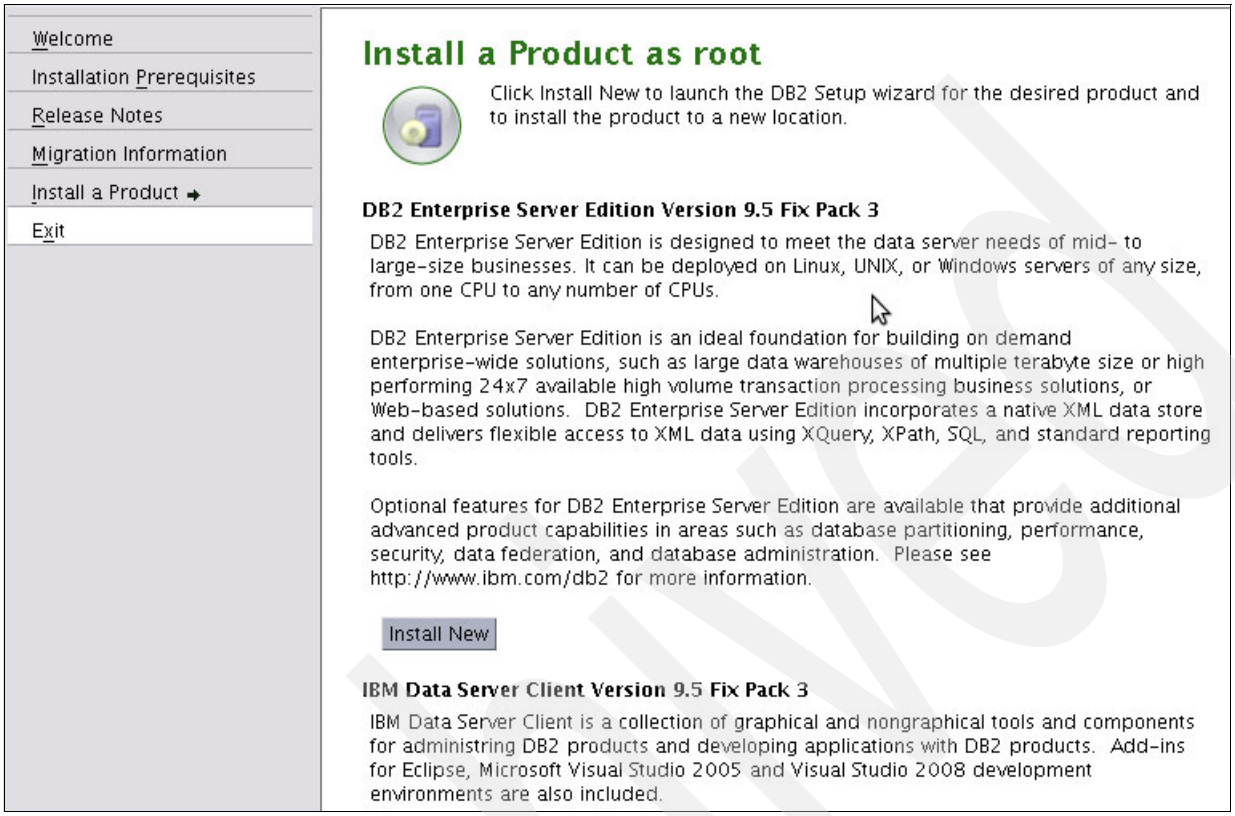

Figure 3-96 Install a Product as root window

The DB2 Setup wizard is displayed next, as shown in Figure 3-97. Click **Next**.

Figure 3-97 DB2 Setup wizard

Then, the Software License Agreement is presented. We accepted it.

Tivoli Storage Productivity Center uses a DB2 instance to store all the information about the storage environment. Tivoli Storage Productivity Center does not require a custom install; therefore, we chose **Typical**, as shown in Figure 3-98.

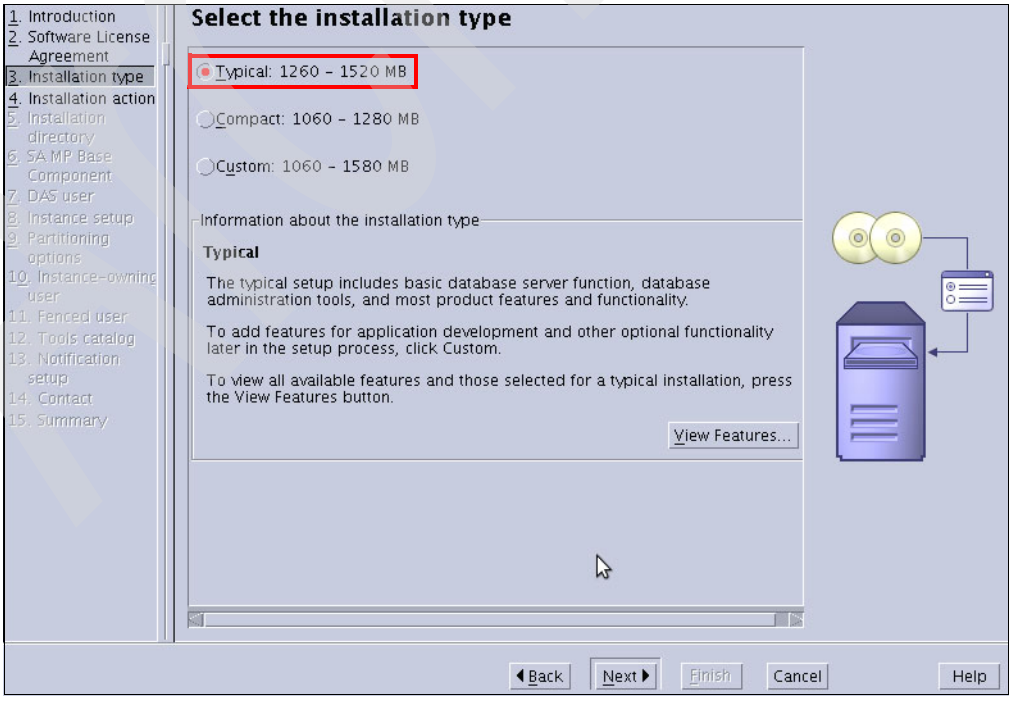

Figure 3-98 We chose a Typical installation

To facilitate potential future installations (or re-installations), all responses are saved in a response file, as shown in Figure 3-99.

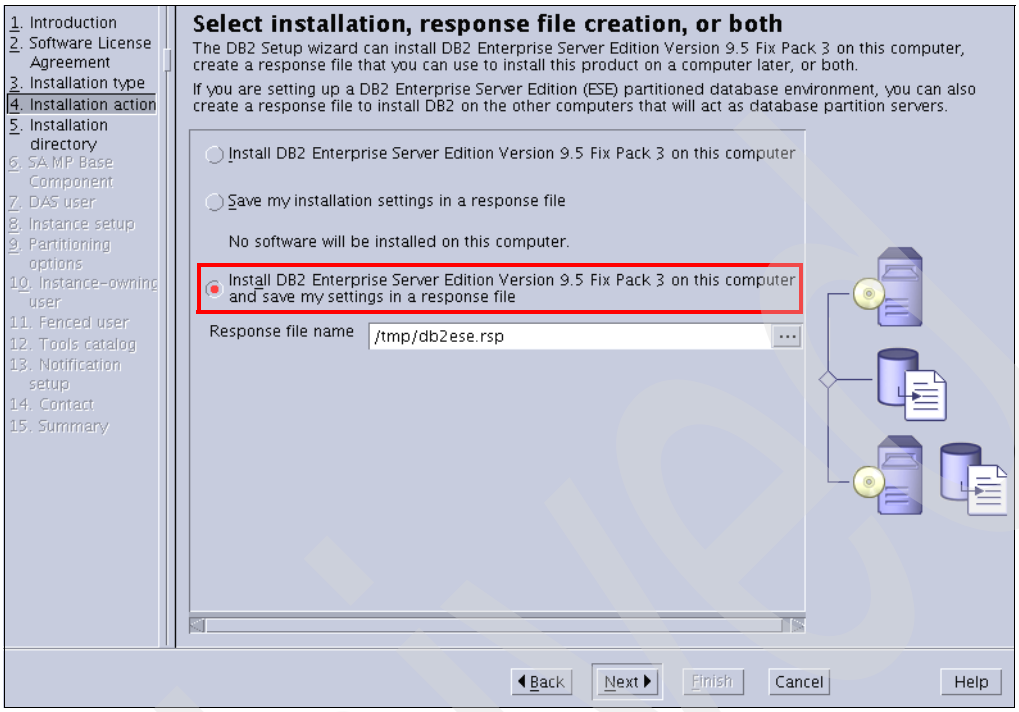

Figure 3-99 Choose to create an installation response file

The default installation directory has enough space to accommodate the installation so the default is selected, as shown in Figure 3-100.

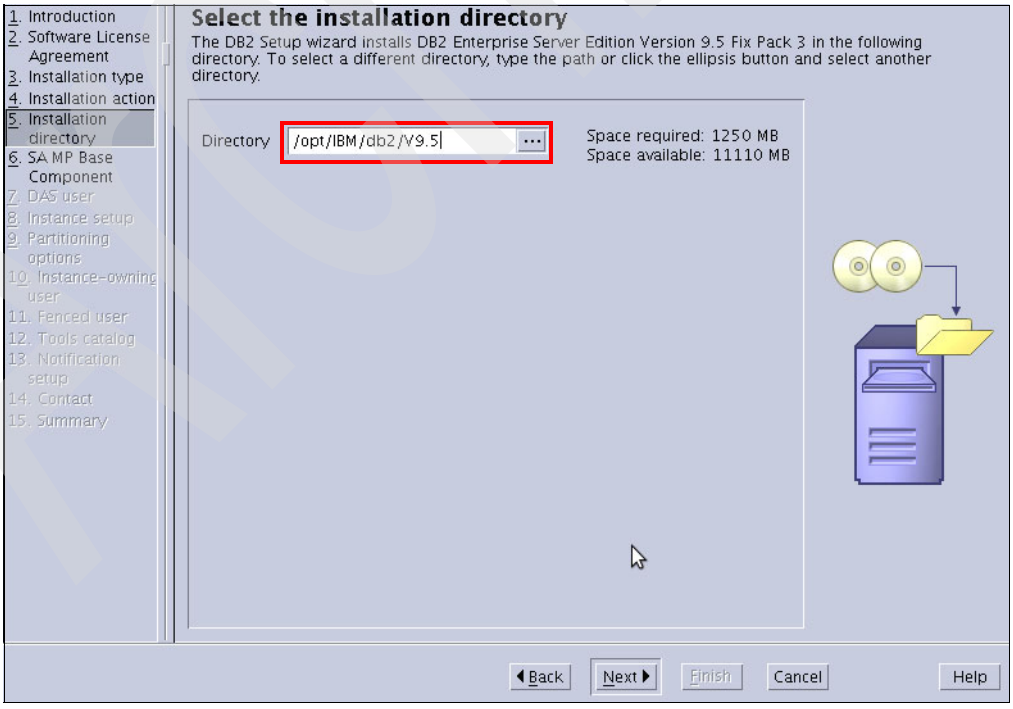

Figure 3-100 Choose the default installation directory

For the purposes of this paper, high availability was not a major concern. Therefore Tivoli System Automation for Multiplatforms Base Component (SAMP Base Component) is not installed, as shown in Figure 3-101.

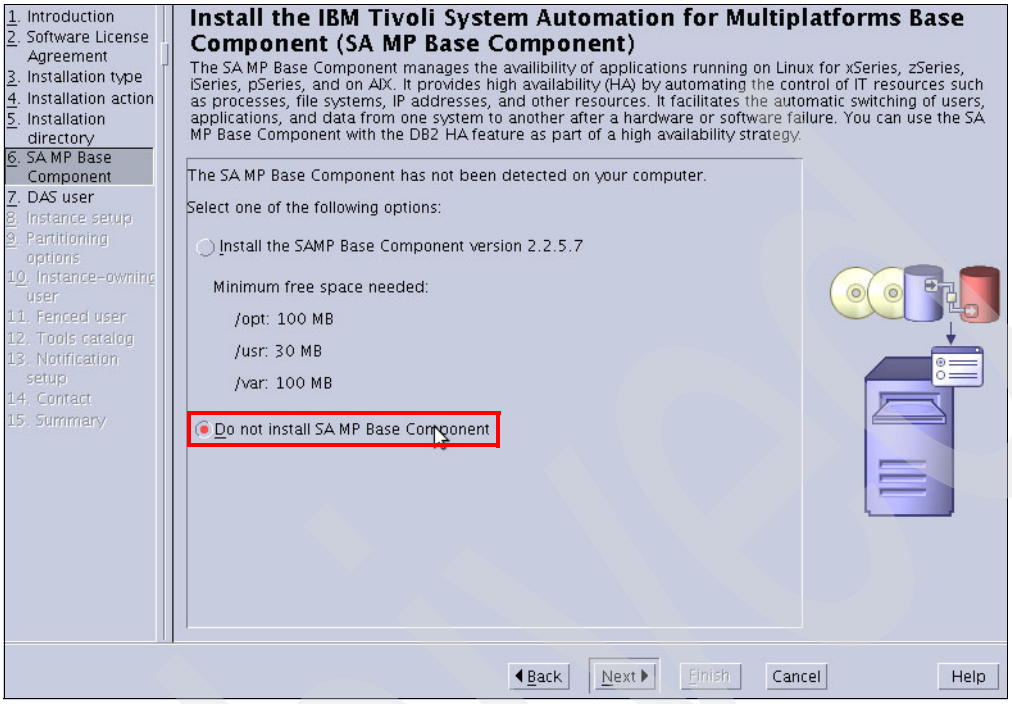

Figure 3-101 Tivoli System Automation for Multiplatforms Base Component is not installed

For DB2 to function properly, you must create a set of new users and groups. The DB2 installation program automates this process. The first user that is created is the DB2 Administration Server user, as shown in Figure 3-102.

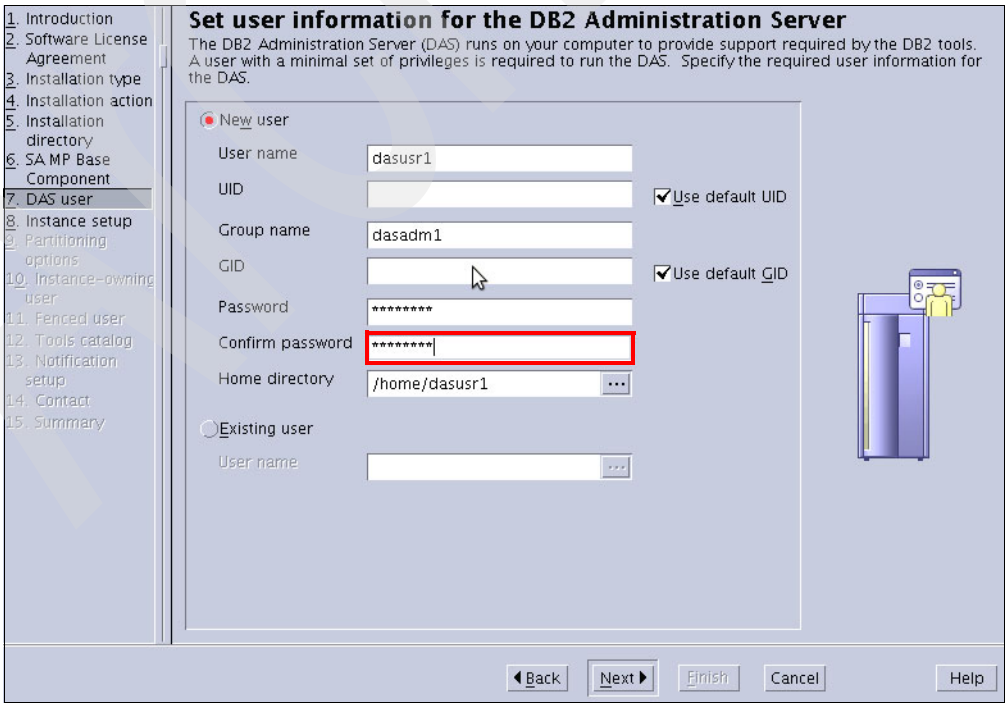

Figure 3-102 User information for the DB2 administrator

Tivoli Storage Productivity Center requires the creation of a DB2 instance. This option is chosen. See Figure 3-103.

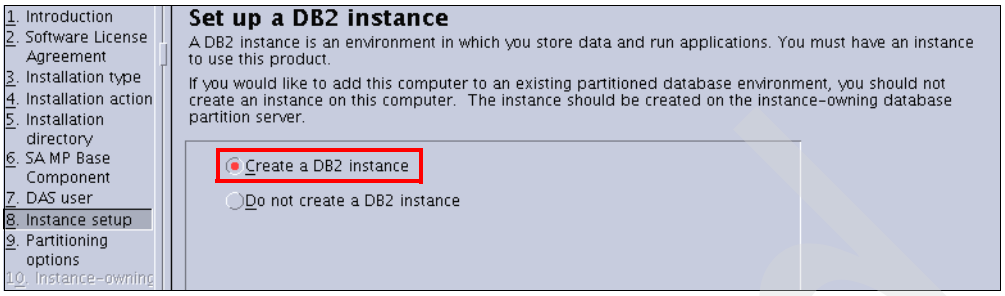

Figure 3-103 Set up a DB2 instance

Tivoli Storage Productivity Center does not require a partitioned instance, so a single partition instance is selected in Figure 3-104.

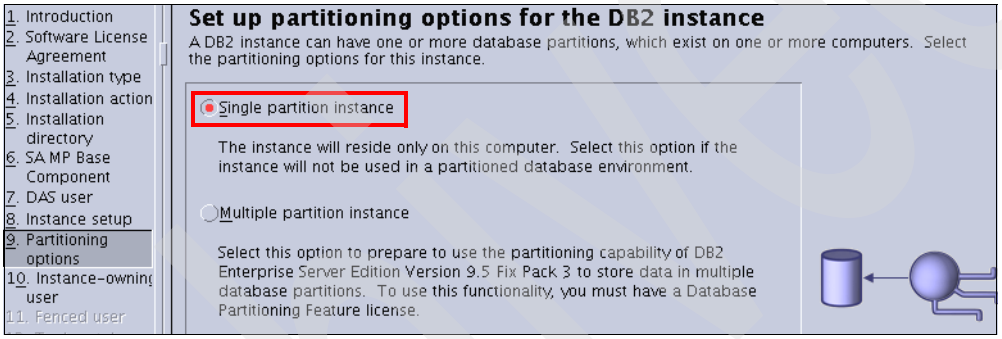

Figure 3-104 Set up a single partition DB2 instance

Figure 3-105 shows the prompt for information about the DB2 instance owner user name. This information is required during later steps of the Tivoli Storage Productivity Center installation.

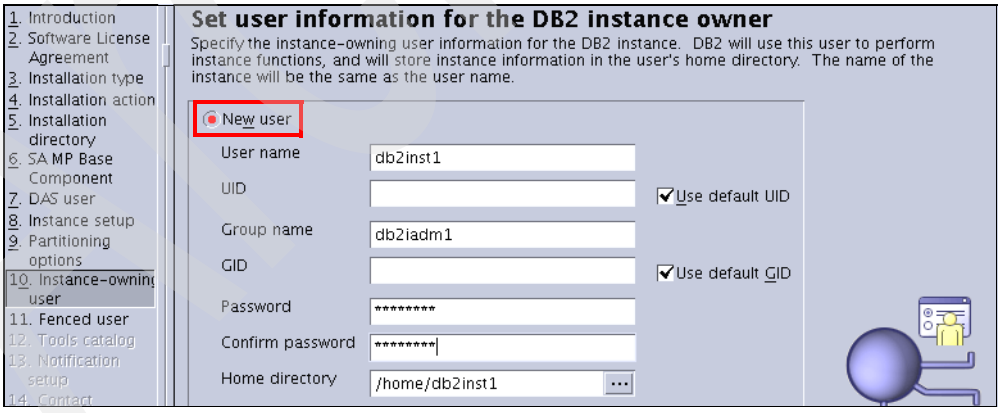

Figure 3-105 Set user information for the DB2 instance owner

The default fenced user information is entered, as seen in Figure 3-106.

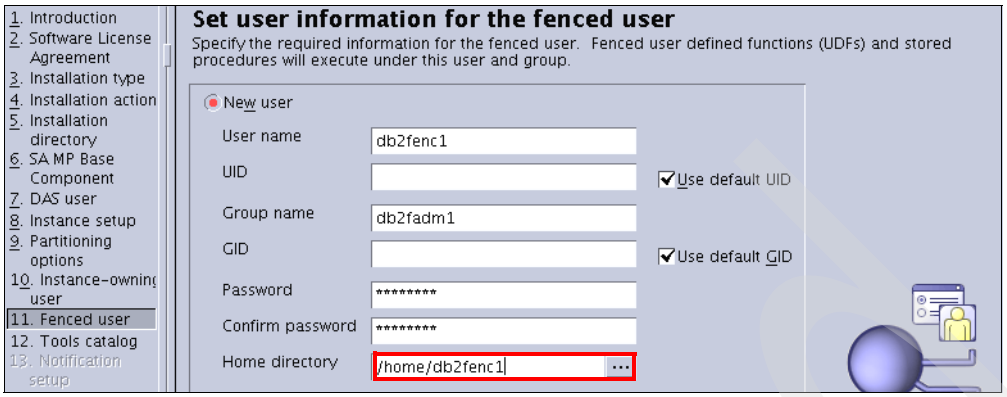

Figure 3-106 Set user information for the fenced user

There is currently no need to install the tools catalog so **Do not prepare the DB2 tools catalog** is selected, as shown in Figure 3-107.

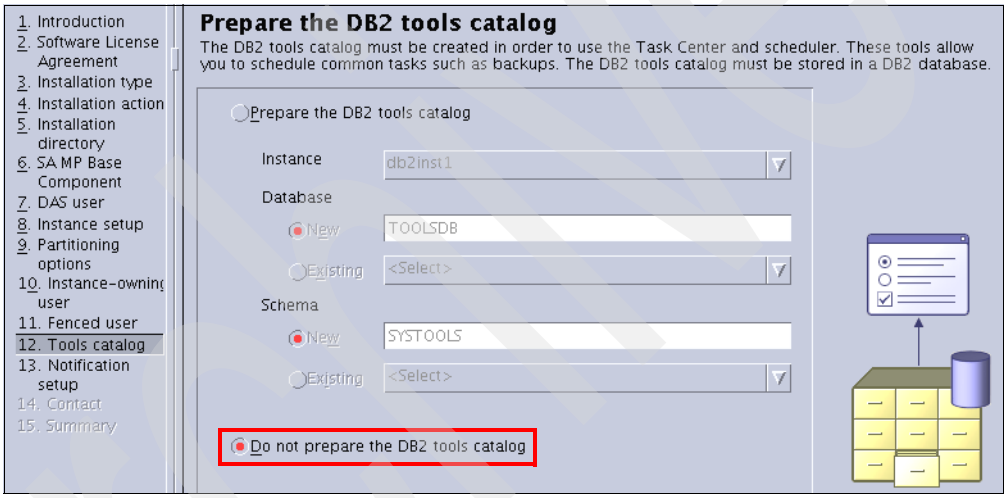

Figure 3-107 Do not prepare the DB2 tools catalog

The default notification settings are selected, as shown in Figure 3-108.

Figure 3-108 Set up default notification settings

The default administration contact for health monitor notification is selected, as shown in Figure 3-109.

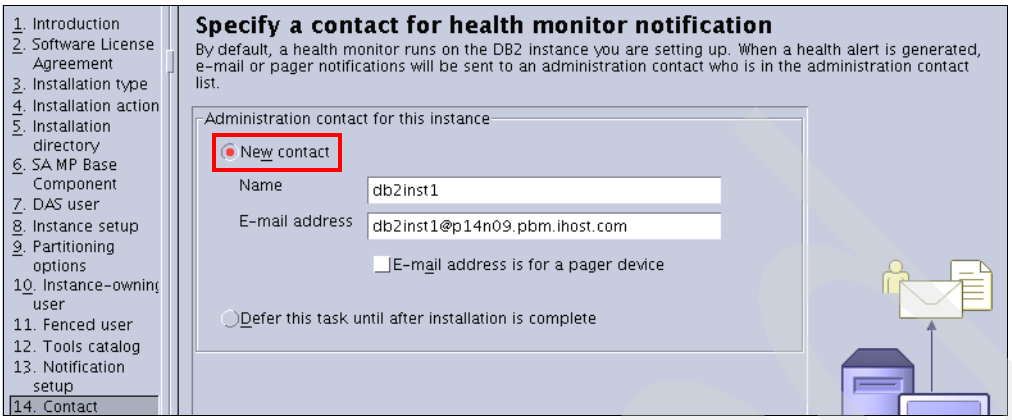

Figure 3-109 Specify a contact for health monitor notification

The summary window is shown before the installation begins, as depicted in Figure 3-110.

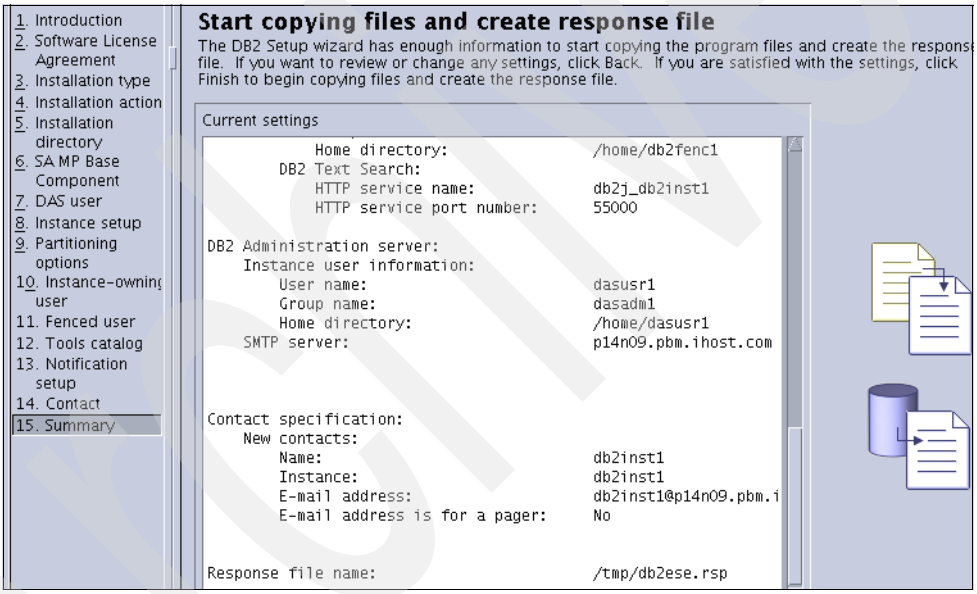

Figure 3-110 Installation summary

Figure 3-111 shows the progress bar that is displayed as DB2 installs.

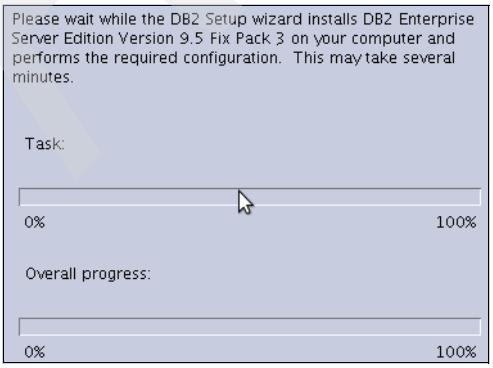

Figure 3-111 DB2 installation progress

Figure 3-112 shows that the setup has completed successfully.

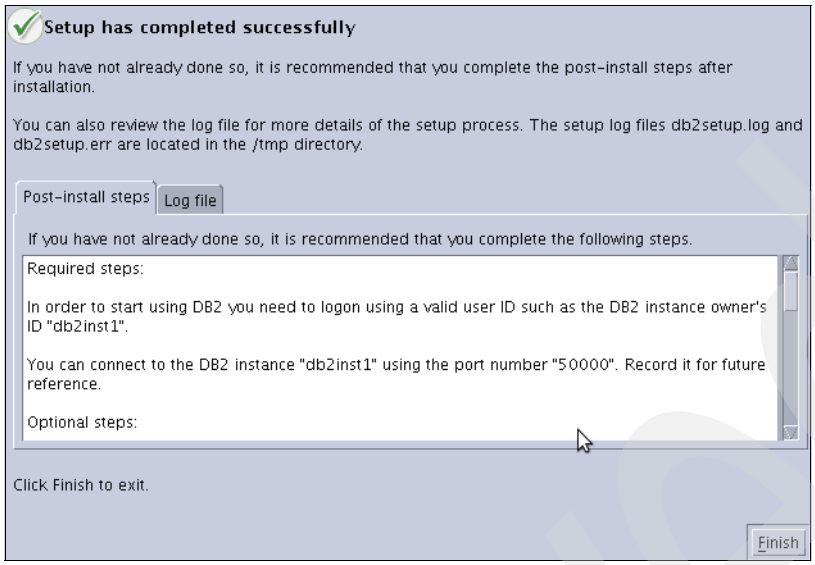

Figure 3-112 DB2 installed successfully

Now that DB2 is fully installed, the Tivoli Storage Productivity Center installer is started.

Setting variables: In order for the Tivoli Storage Productivity Center installation program to successfully communicate with the DB2 instance that was created in this section, the installation environment requires that you set certain variables. Set them by sourcing the DB2 profile, for example:

```
#. /home/db2inst1/sqllib/db2profile
#./setup.sh
```

Note the leading dot (.) and “space” in the first command.

When the Tivoli Storage Productivity Center installation starts, the first step is to accept the License Agreement, as shown in Figure 3-113.

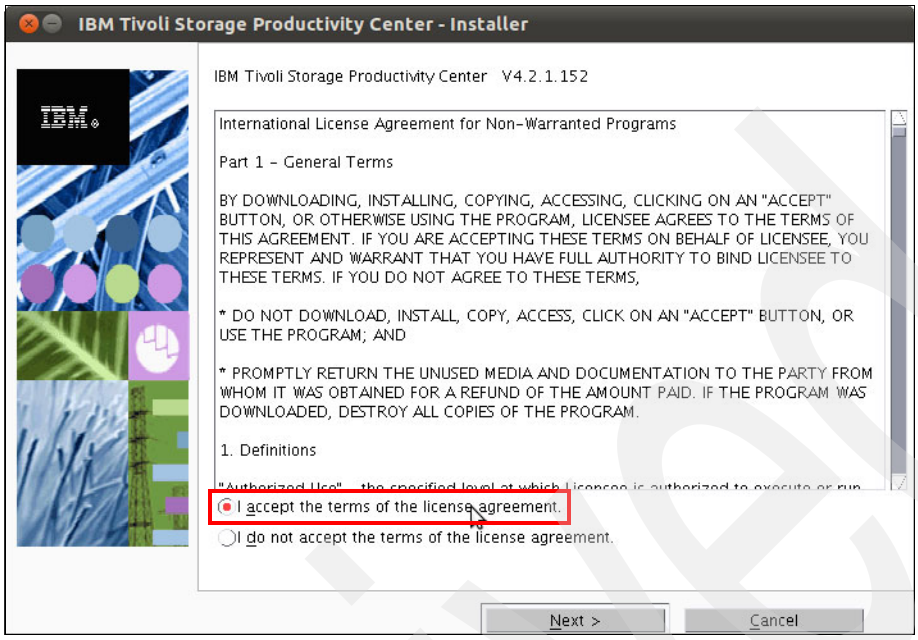

Figure 3-113 Accepting the IBM Tivoli Storage Productivity Center license agreement

Then, a typical installation is chosen, installing servers, clients, and storage resource agents, as shown in Figure 3-114.

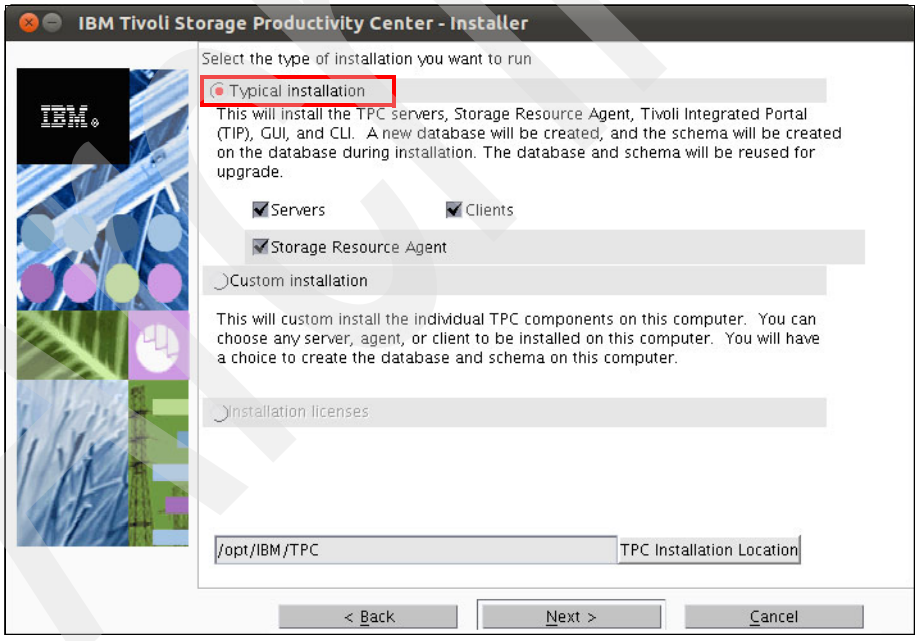

Figure 3-114 Choosing a Typical Tivoli Storage Productivity Center installation

The *user ID* and *password* combination of the DB2 instance owner is now required. You entered this same information when you installed DB2. The server name, server port, and agent port are left as the defaults (Figure 3-115).

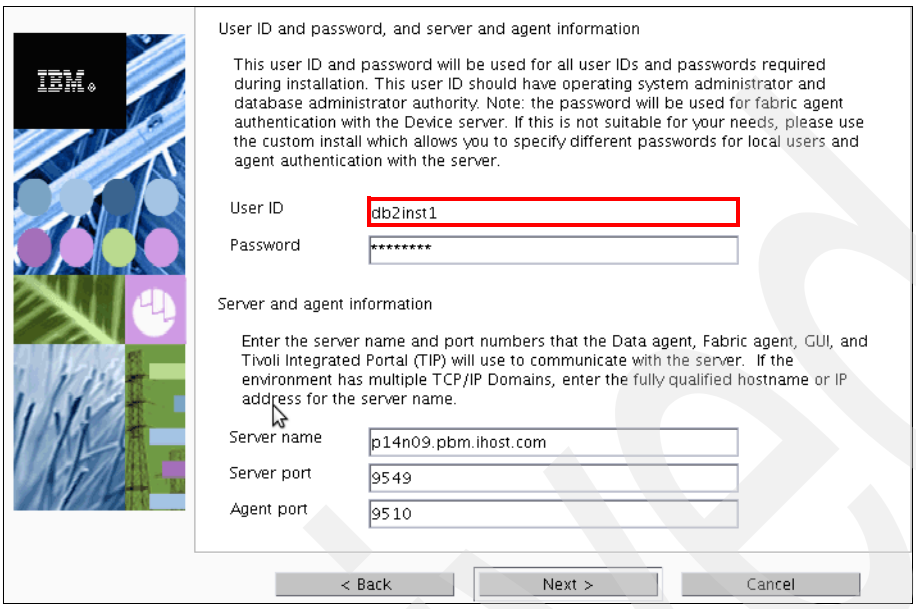

Figure 3-115 Specifying the DB2 instance owner to Tivoli Storage Productivity Center

The default Tivoli Integrated Portal (TIP) information is selected, and **Next** is chosen (Figure 3-116).

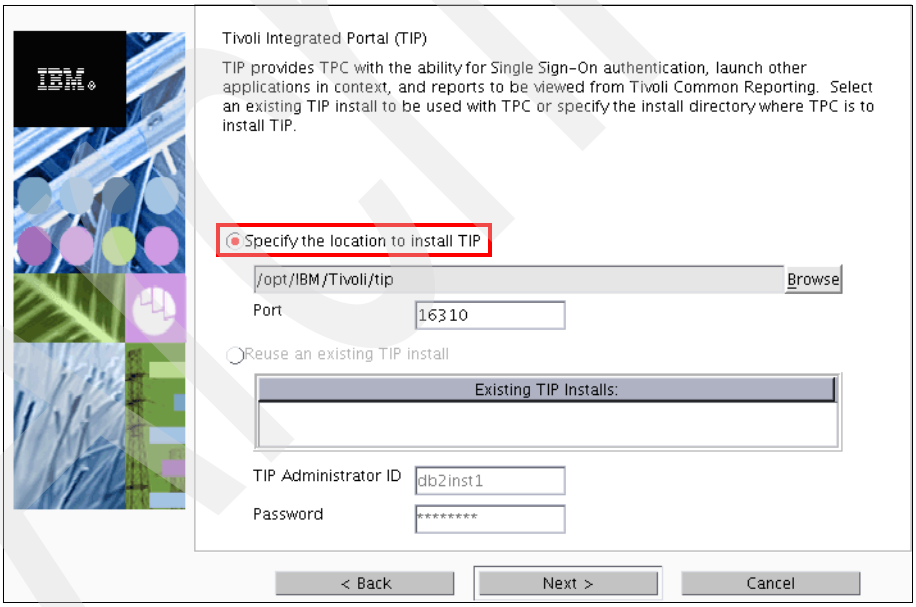

Figure 3-116 Specifying the default Tivoli Integrated Portal information

There is no Lightweight Directory Access Protocol (LDAP)/Active Directory in this environment, so the local operating system authentication is selected (Figure 3-117).

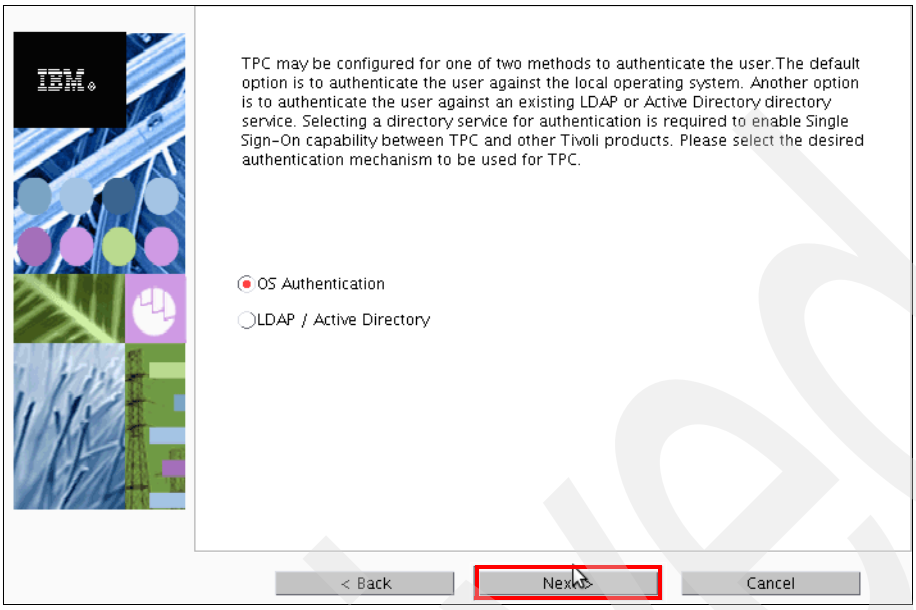

Figure 3-117 Specifying local operating system authentication vs. LDAP

The summary page is shown before installation begins (Figure 3-118).

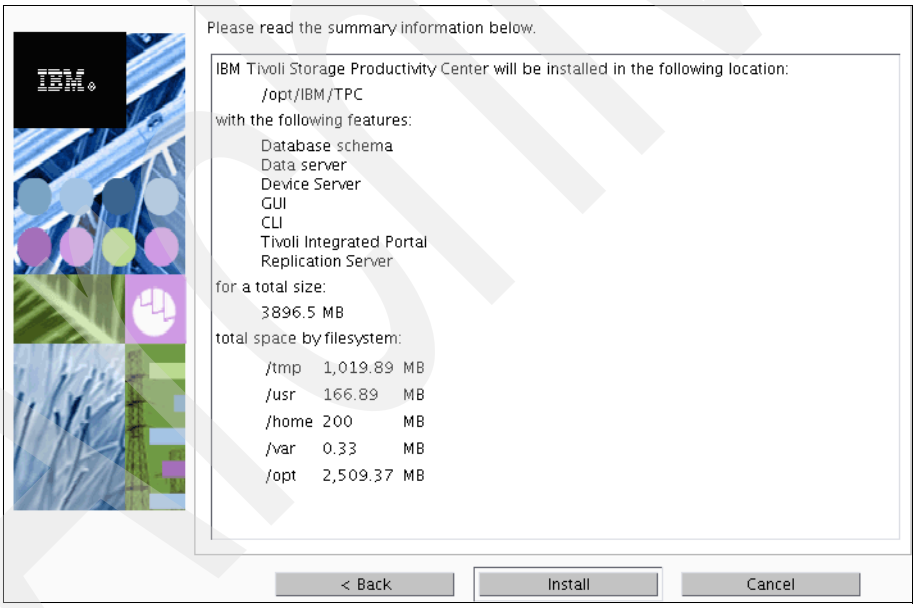

Figure 3-118 IBM Tivoli Storage Productivity Center installation summary

A progress bar is shown during installation (Figure 3-119).

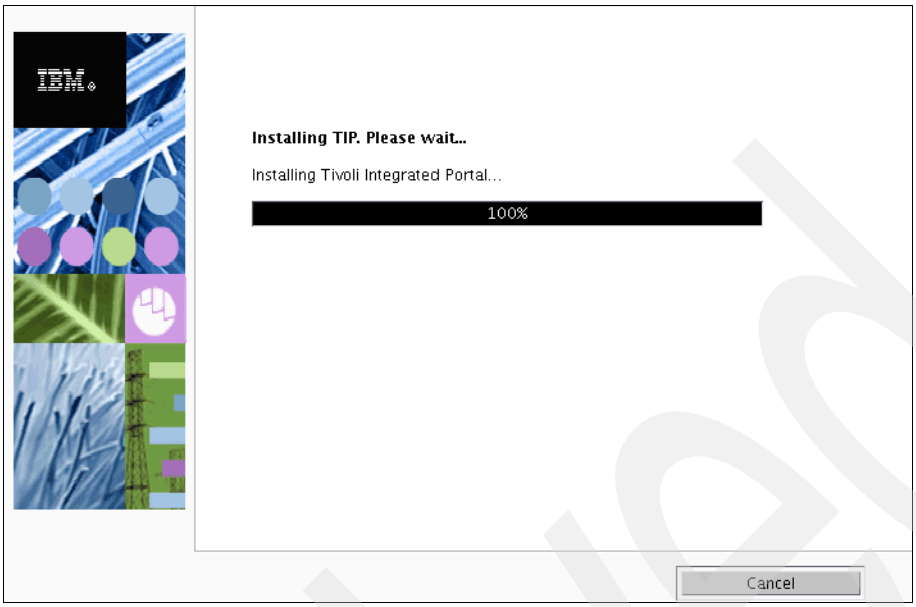

Figure 3-119 Tivoli Integrated Portal installation progress bar

Figure 3-120 shows the successful installation message from the InstallShield Wizard.

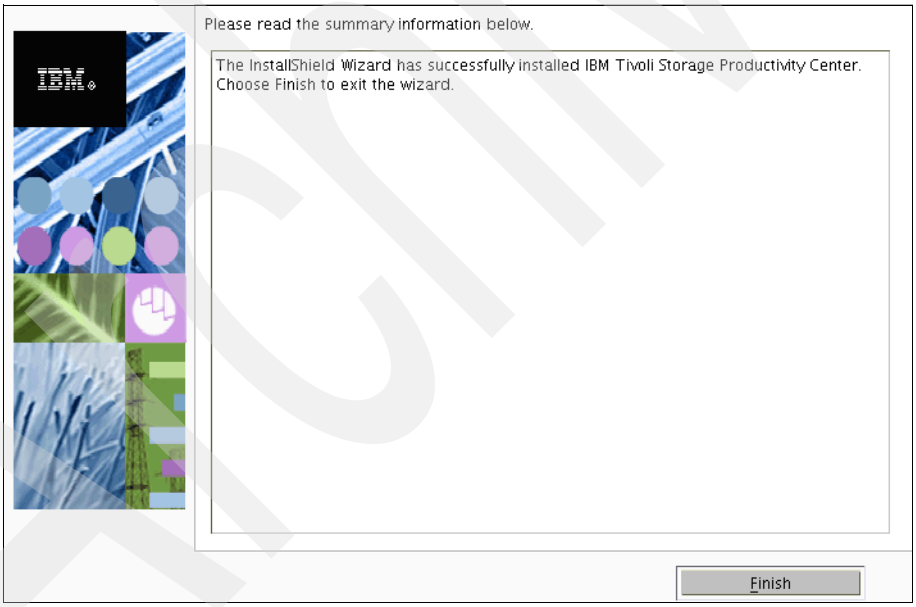

Figure 3-120 IBM Tivoli Storage Productivity Center successful installation message

3.5.3 Configuring Tivoli Storage Productivity Center

Now that Tivoli Storage Productivity Center is installed, you can use this command to launch the main Tivoli Storage Productivity Center interface window:

```
# /opt/IBM/TPC/gui/TPCD.sh
```

Enter the connection information in the window. During the installation, Local Operating System Authentication was chosen. Figure 3-121 shows the root user, password, and local host being specified.

Figure 3-121 Entering connection information

The Tivoli Storage Productivity Center dashboard appears, as shown in Figure 3-122.

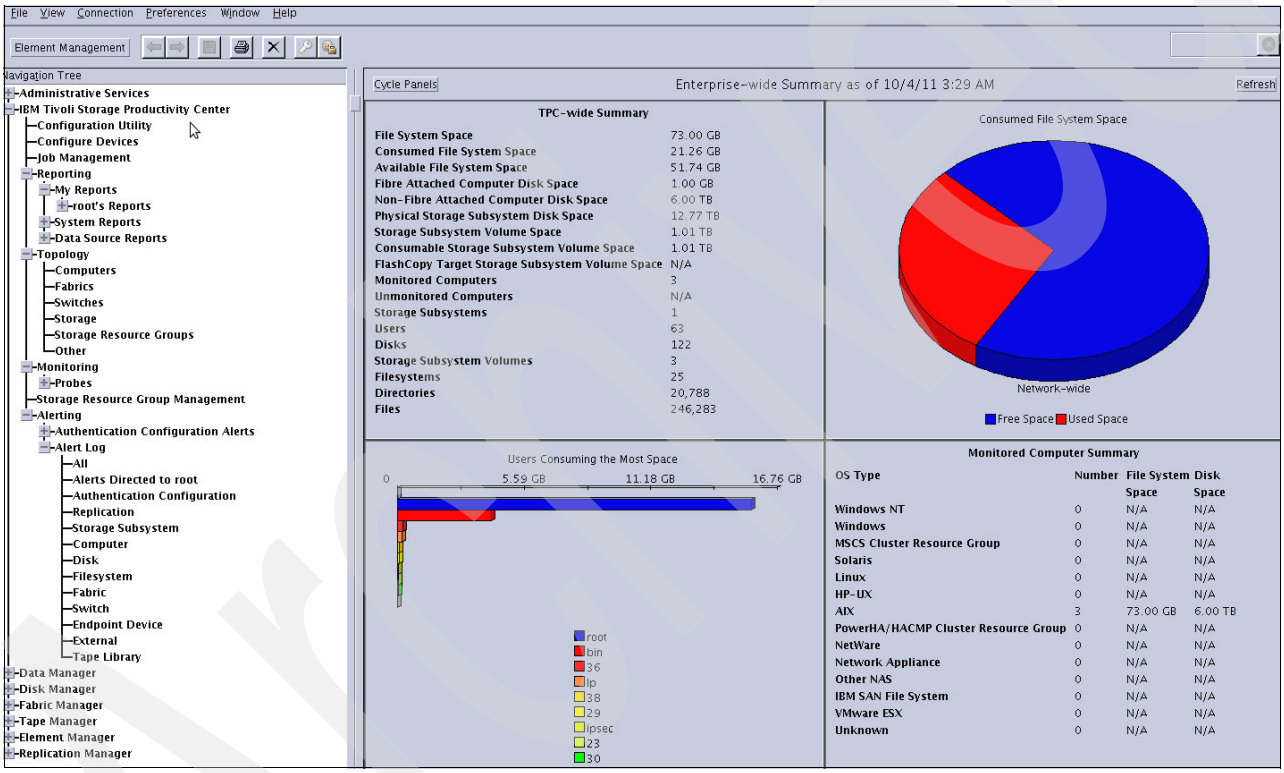

Figure 3-122 The Tivoli Storage Productivity Center dashboard

The first step in configuring Tivoli Storage Productivity Center to be able to correctly detect and correlate the various resources is to add information about the CIM providers (CIMOMs). Click **Add CIMOM** on the CIMOM Agents window.

Expand **Administrative Services** and then click **Data Sources**, as shown in Figure 3-123.

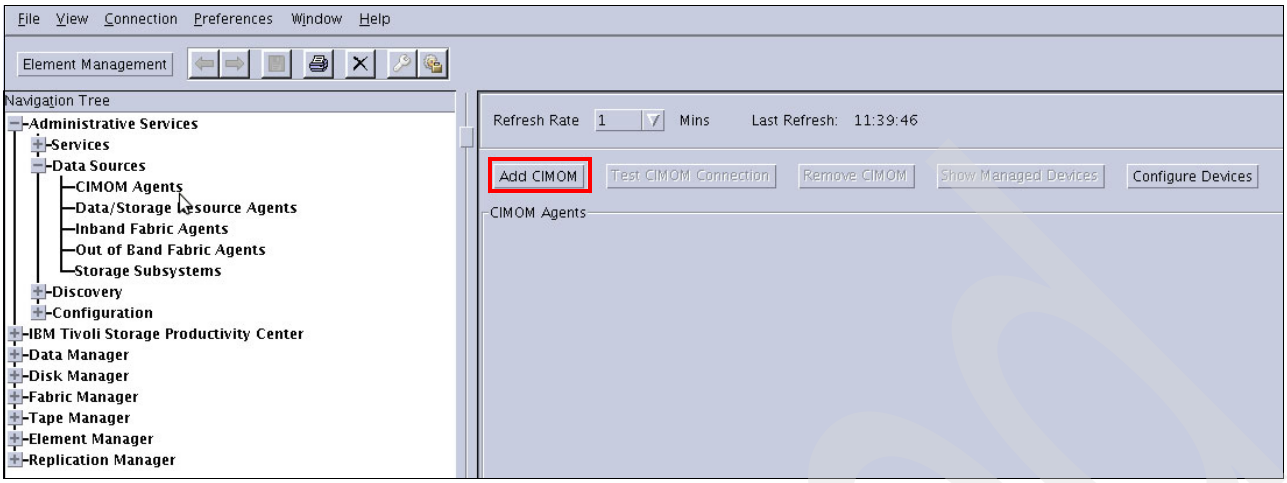

Figure 3-123 Data Sources submenu

In the window that opens (as shown in Figure 3-124), enter the connection information for the first CIMOM: the IP address of the CIM LPAR and the port for a particular provider. Enter 55988 for the Brocade FC switch CIM provider. This information was provided to the CIM Provider in 3.4, “Installing CIM” on page 86 (see Example 3-24 on page 92).

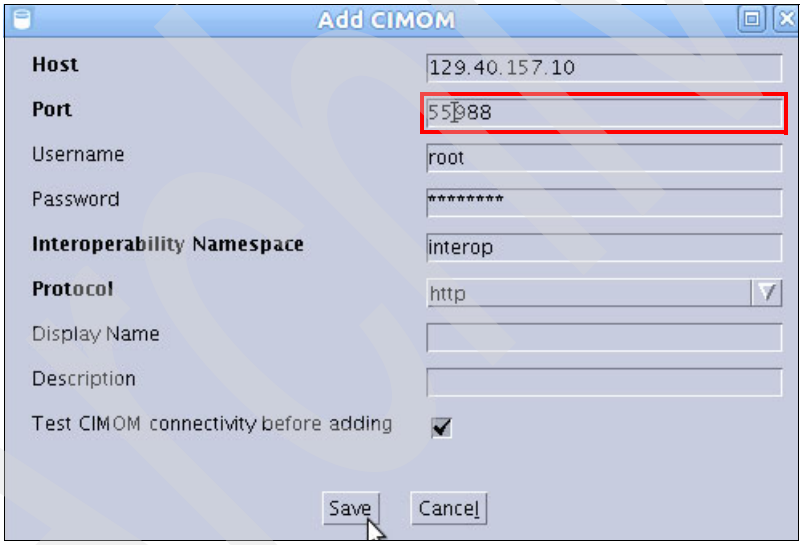

Figure 3-124 Connection information for the first CIMON

Click **Save** and the new CIMOM agent appears with the Connection Status of SUCCESS in the CIMOM Agents window, as shown in Figure 3-125.

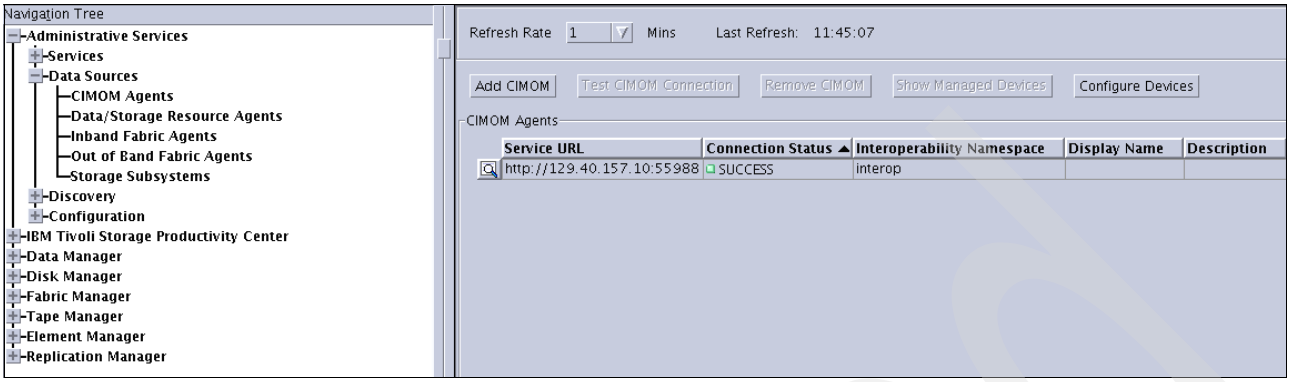

Figure 3-125 Data Sources submenu

Repeat the previous process for the other required CIMOM. Again, use the IP address of the CIM LPAR, but this time, use the port 5988 for the Engenio provider, which is used to communicate with the SAN subsystem, as shown in Figure 3-126.

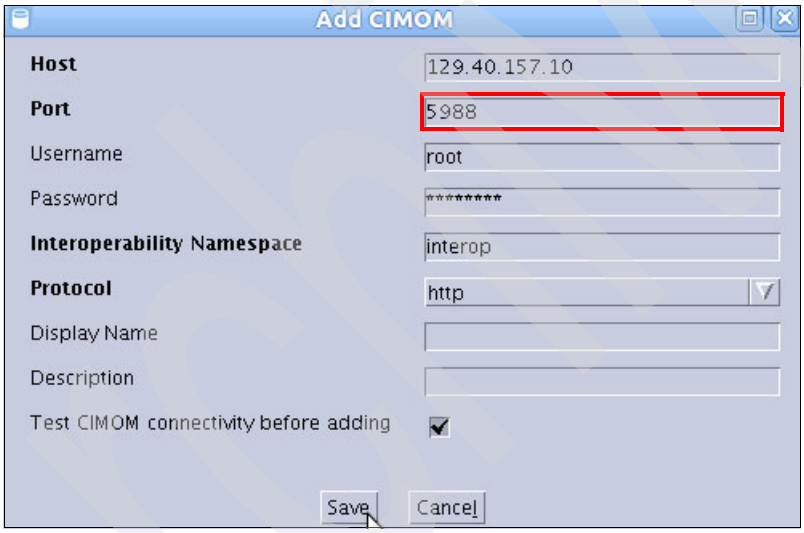

Figure 3-126 Providing port information for CIMOM

Click **Save**. Figure 3-127 shows that both CIMOMs are now displayed in the CIMOM Agents window.

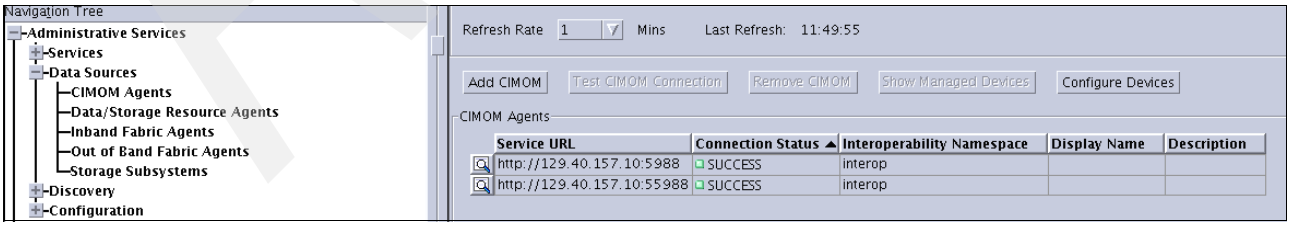

Figure 3-127 CIMOM Agents summary

The next step is to probe the SAN fabric and subsystem. In the Navigation Tree, expand the **Probes** menu, as shown in Figure 3-128.

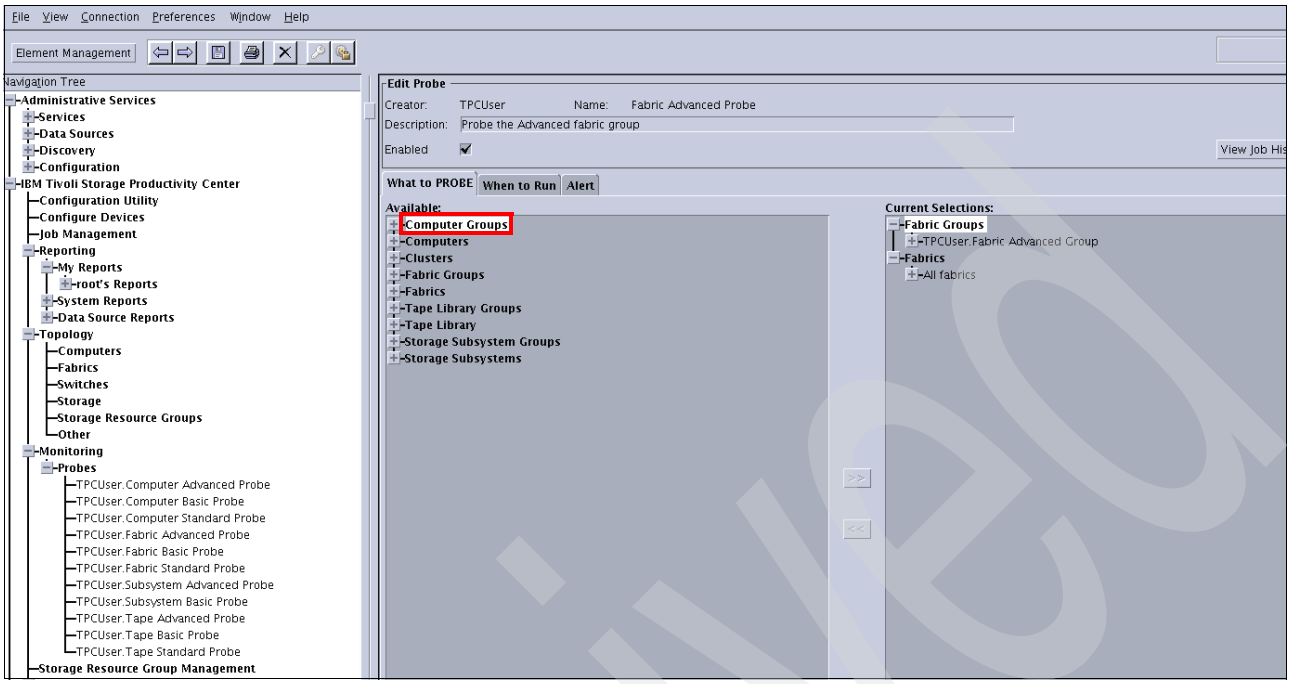

Figure 3-128 Edit the Fabric Probe

Make sure that TPCUser.Fabric Advanced Group appears in the Current Selections window. Then, click the **When to Run** tab. Ensure that the **Run Now** radio button is selected.

Submit the job by clicking the save/disk icon button, which is located above the Navigation Tree that is depicted in Figure 3-129.

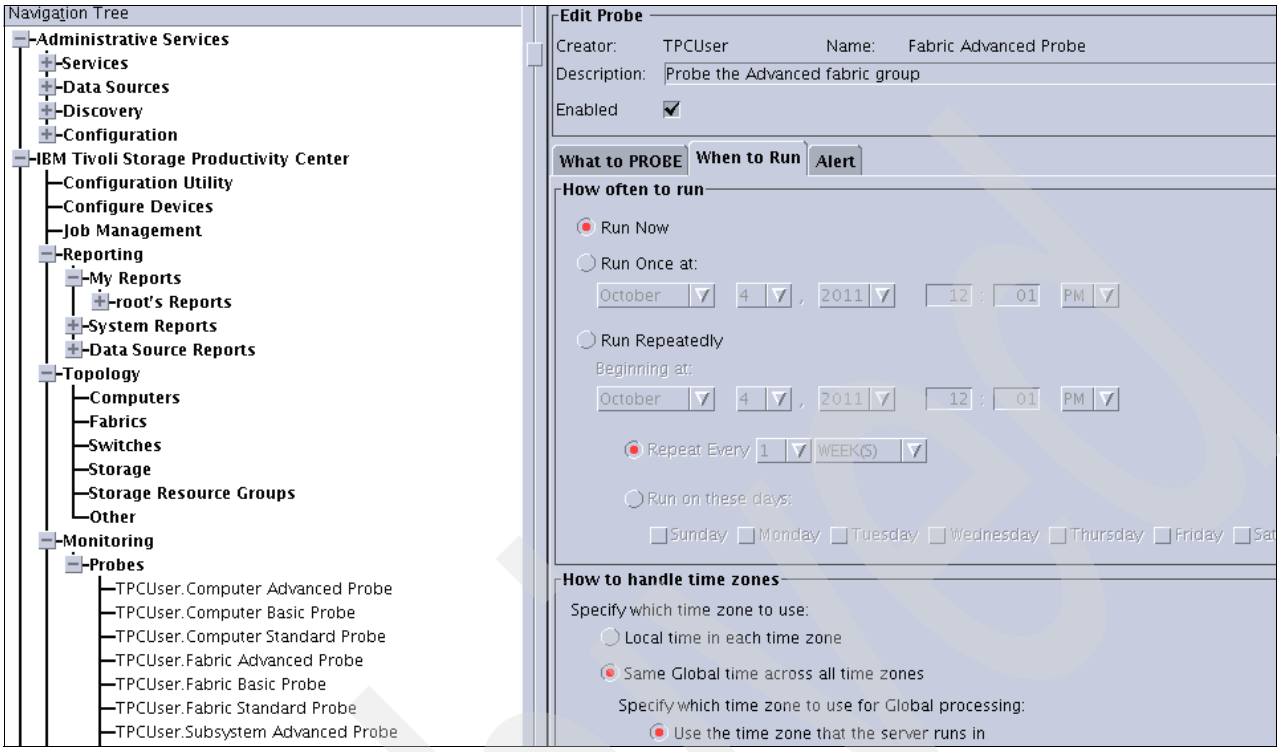

Figure 3-129 Schedule Fabric Probe

The confirmation window that is shown in Figure 3-130 opens. Click **Yes**.

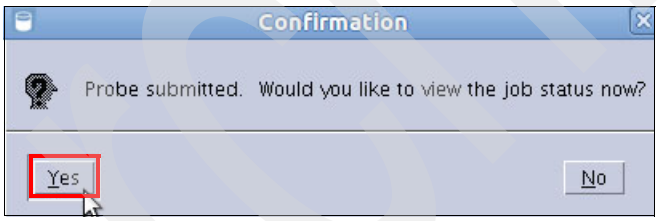

Figure 3-130 Probe submitted confirmation

Edit the Probe, as shown in Figure 3-131. Next, run the probe job for the Subsystem Advanced Probe, as shown in Figure 3-131 through Figure 3-133 on page 112.

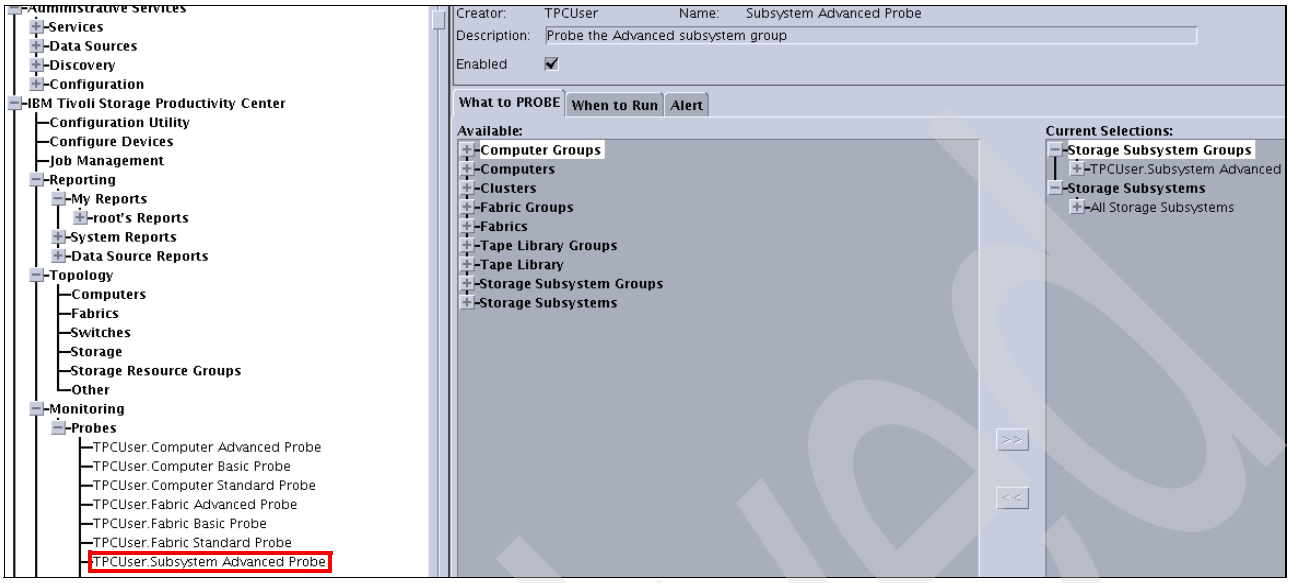

Figure 3-131 Edit the Subsystem Advance Probe

Figure 3-132 shows scheduling when to run the probe.

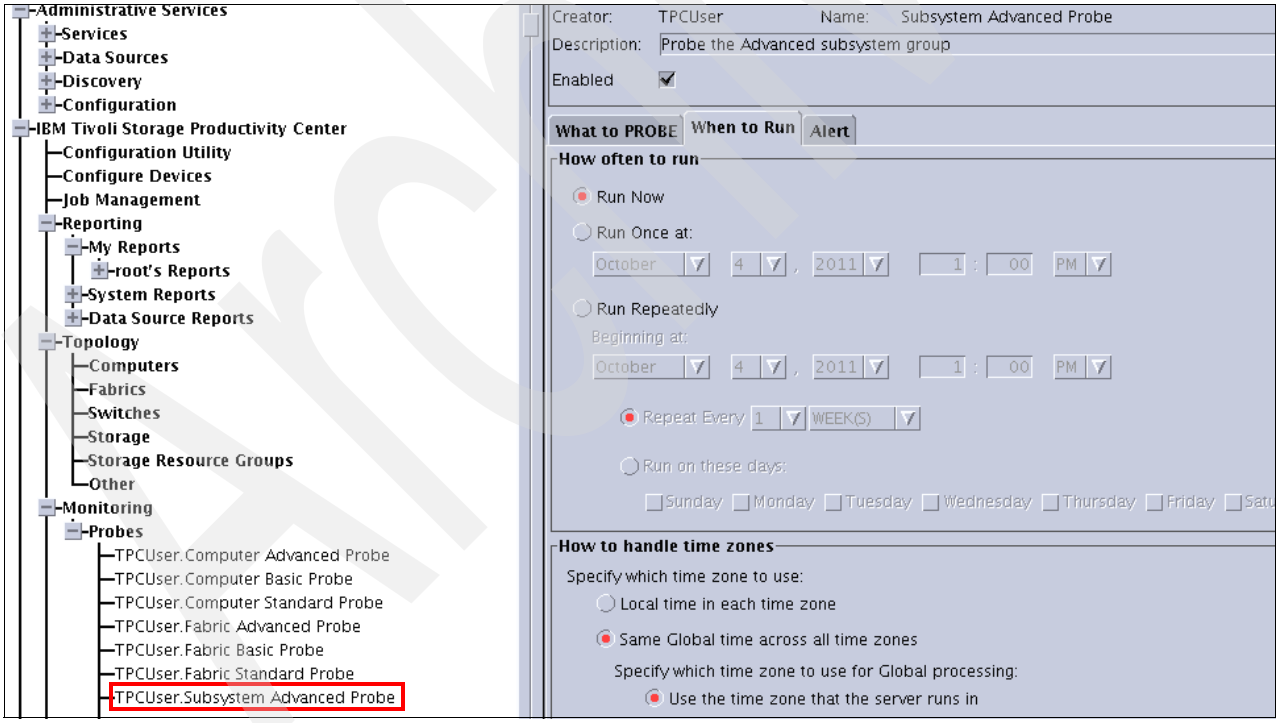

Figure 3-132 Schedule the Subsystem Probe

Figure 3-133 shows the confirmation window.

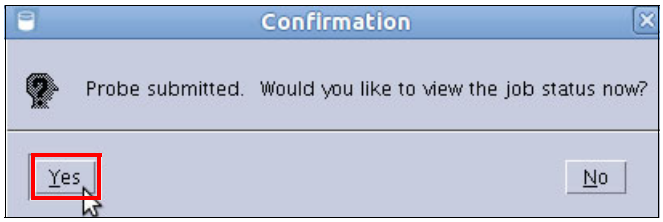

Figure 3-133 Probe submitted confirmation

Click **Yes** to display the job status. Note the status in lower-right portion of the panel. The status appears as completed successfully, as shown in Figure 3-134. Make sure that the jobs complete successfully.

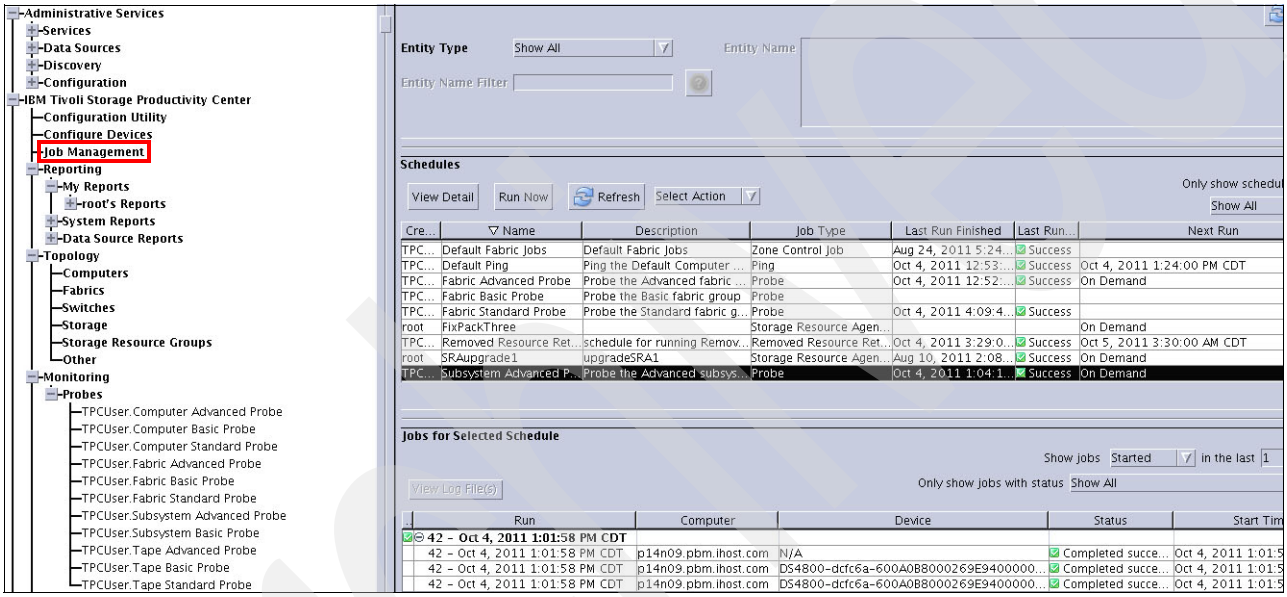

Figure 3-134 Job status

Now that the SAN fabric and storage subsystem are known to Tivoli Storage Productivity Center, the only remaining task for configuring Tivoli Storage Productivity Center is to connect it to the three VIO servers. The communication between Tivoli Storage Productivity Center and VIO is performed using storage resource agents (SRAs). The agents must be installed individually on each of the three VIOs. The easiest way to install the agents is to mount the extracted TPC install package to the VIO servers, for example, using NFS.

The following command extracts the TPC install package into the `/tmp` directory:

```
$ tar -xvf CZR3TML.tar  -C /tmp
```

While logged into one of the VIO servers as `padmin`, issue the following command:

```
$ oem_setup_env
```

Then, create a directory, such as `/opt/IBM/TPC/forSRA`, to contain the installed agent data:

```
# mkdir /opt/IBM/TPC/forSRA
```

Run the command in Example 3-25 to install the agent in the directory that we created previously.

Example 3-25 Command to install storage resource agent

```
# /tmp/data/sra/aix_power/bin/Agent -INSTALL -INSTALLLOC /opt/IBM/TPC/forSRA
-SERVERIP 129.40.157.9 -SERVERPORT 9549 -DEBUG MAX -COMMTYPE DAEMON -AGENTPORT
9519
```

Substitute the values for *-INSTALLLOC*, *-SERVERIP*, and *-AGENTPORT* with the directory that we created, the IP address of the Tivoli Storage Productivity Center server, and for any port that is available on the VIO server. This installation process needs to be completed on each VIO server.

As each storage resource agent is installed, it is automatically detected by Tivoli Storage Productivity Center. It then appears in the Data/Storage Resource Agents window, as shown in Figure 3-135.

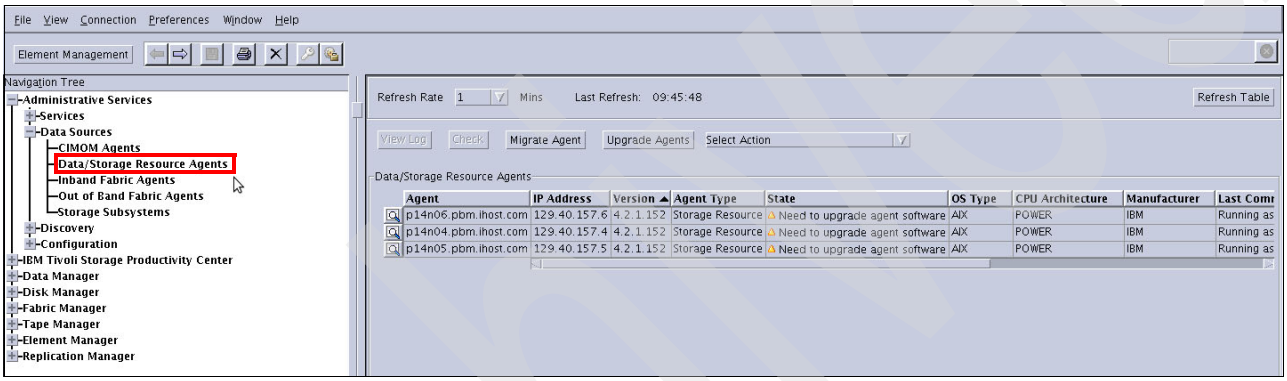

Figure 3-135 Determine if Data/Storage Resource Agent upgrade is available

The State column in Figure 3-135 shows that the agent software needs to be upgraded. If your agent software needs upgrading, you click **Upgrade Agents** near the top of the window in Figure 3-135. Clicking the Upgrade Agents button brings up the window that is shown in Figure 3-136. Ensure that each of the computers with an SRA to be upgraded is added to the Current Selections list, and then click the **Disk/save** icon to submit the SRA upgrade job.

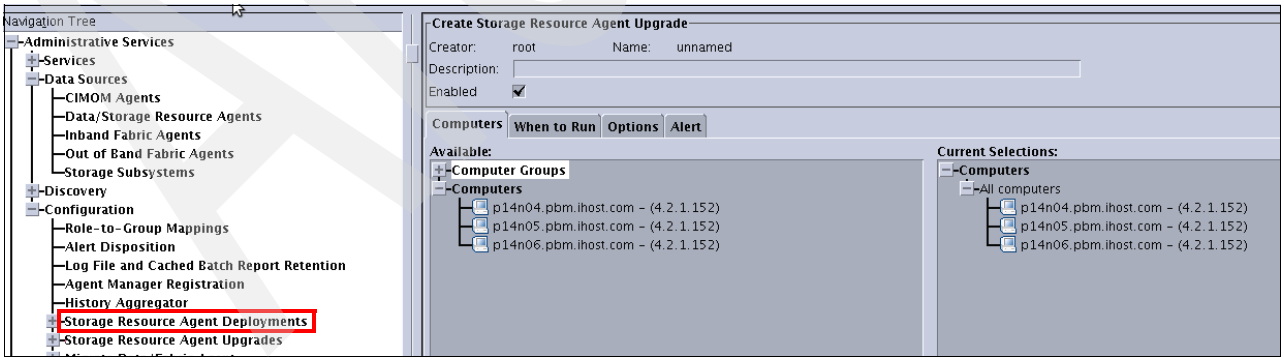

Figure 3-136 Submit Upgrade

Click **Yes** on the dialog that appears to show the job status (Figure 3-137).

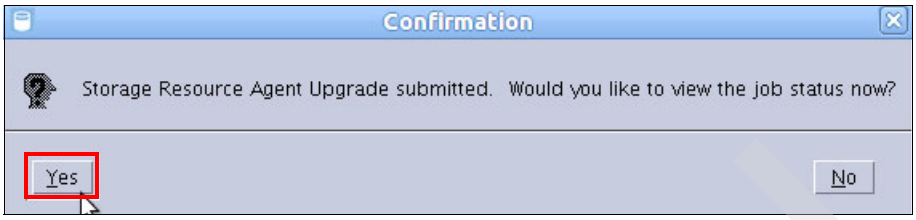

Figure 3-137 Upgrade submitted

The log that is shown in Figure 3-138 indicates that the upgrades were applied.

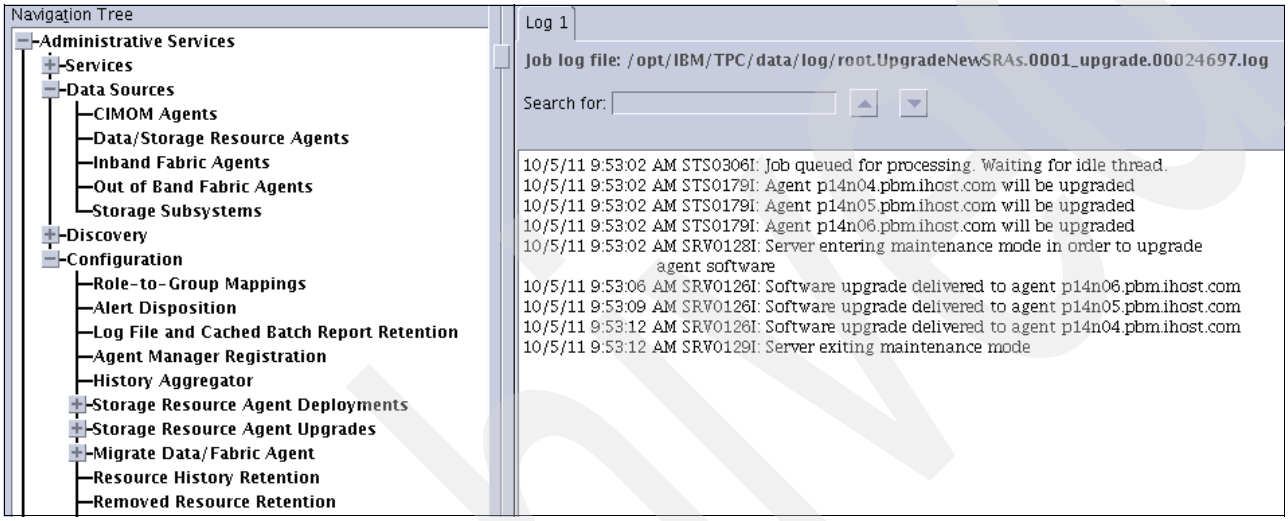

Figure 3-138 View Upgrade Job Status

The State column in Figure 3-139 now reads Up.

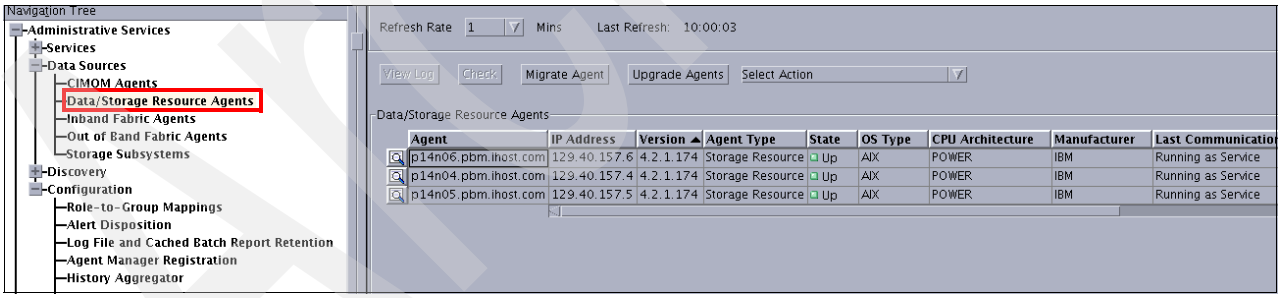

Figure 3-139 Upgrade complete and running as a service

The configuration is now complete. To verify that all the storage components have been correctly detected and correlated, navigate to the Topology window, as shown in Figure 3-140.

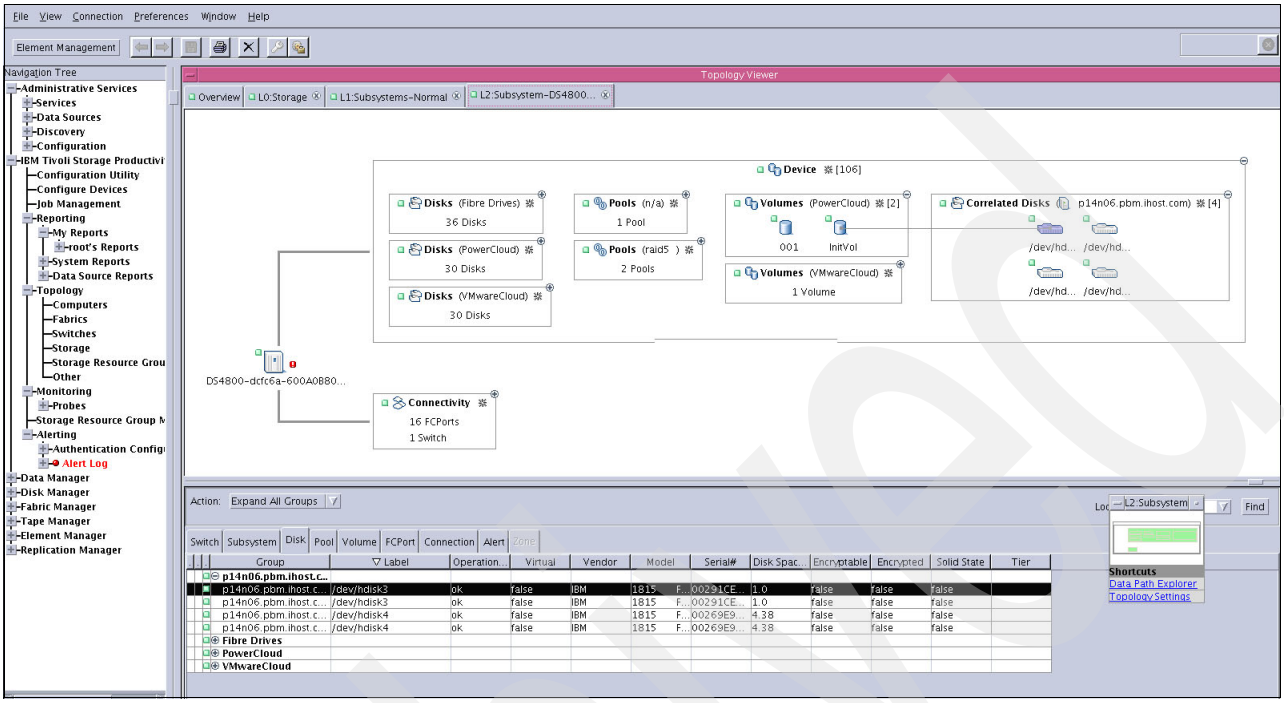

Figure 3-140 Topology window

3.5.4 Discovery

Change to the IPV4 addresses that correspond to the addresses of the LPARs that were created earlier, as shown in Figure 3-141. Then, click **Discover Now**.

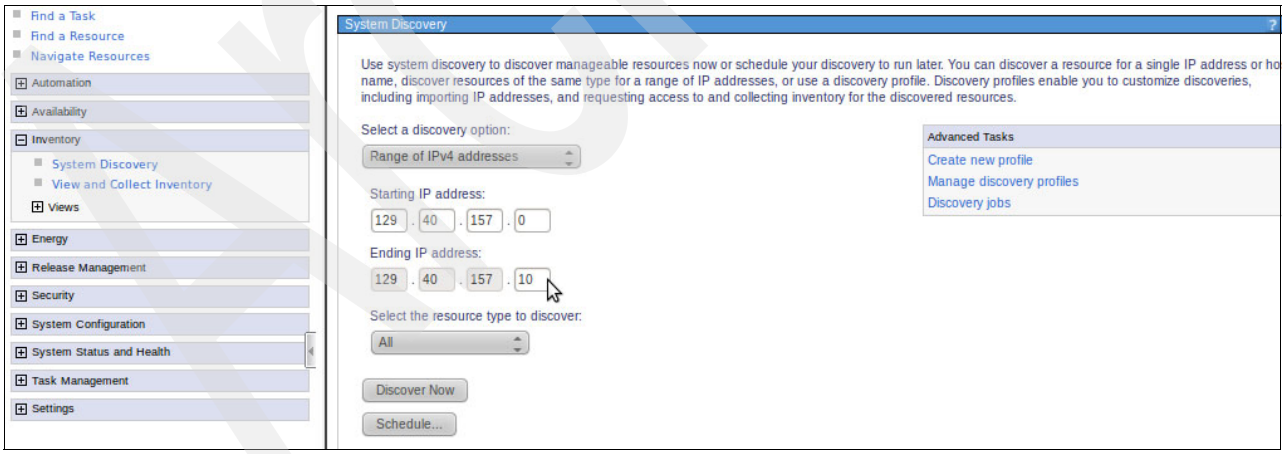

Figure 3-141 Discover a range of IPv4 addresses corresponding to the LPARs that were created earlier

After clicking Discover Now, the Active and Scheduled Jobs panel that is shown in Figure 3-142 provides the Discovered Systems tab. This tab provides information about the discovered systems. When the discover job has completed successfully, it is time to move on to the Access steps, as shown in the next section.

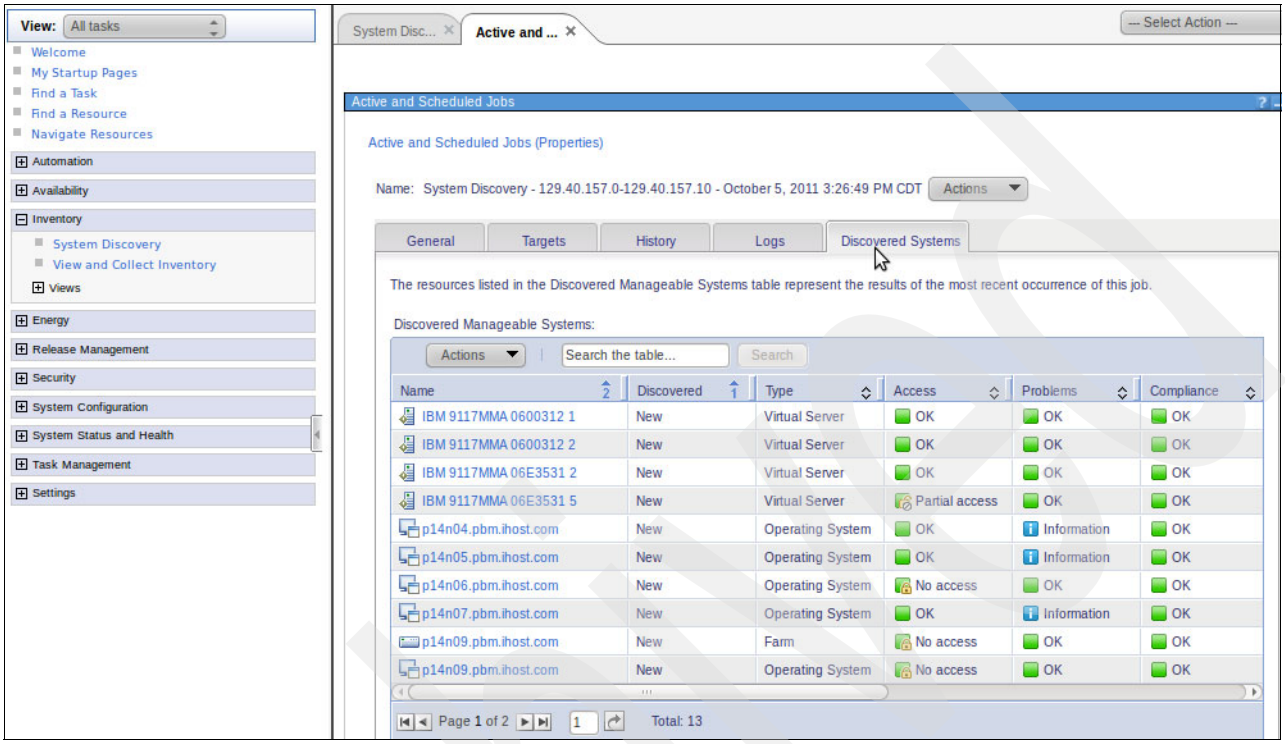

Figure 3-142 View discovered systems

3.5.5 Access

From the Active and Scheduled Jobs window (Figure 3-143), drill down into the operating systems that are listed by using the **Actions** drop-down menu.

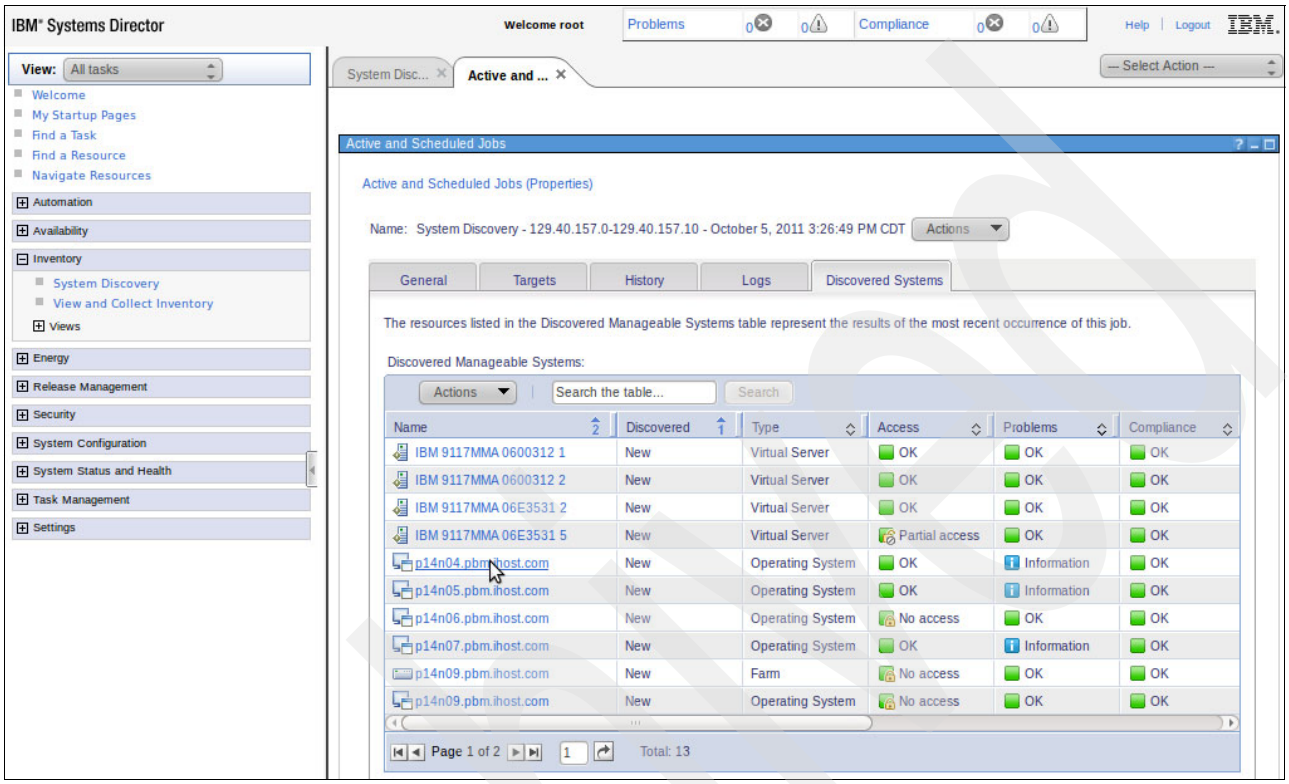

Figure 3-143 Display active and scheduled jobs

Select **Security** → **Request Access**, as shown in Figure 3-144.

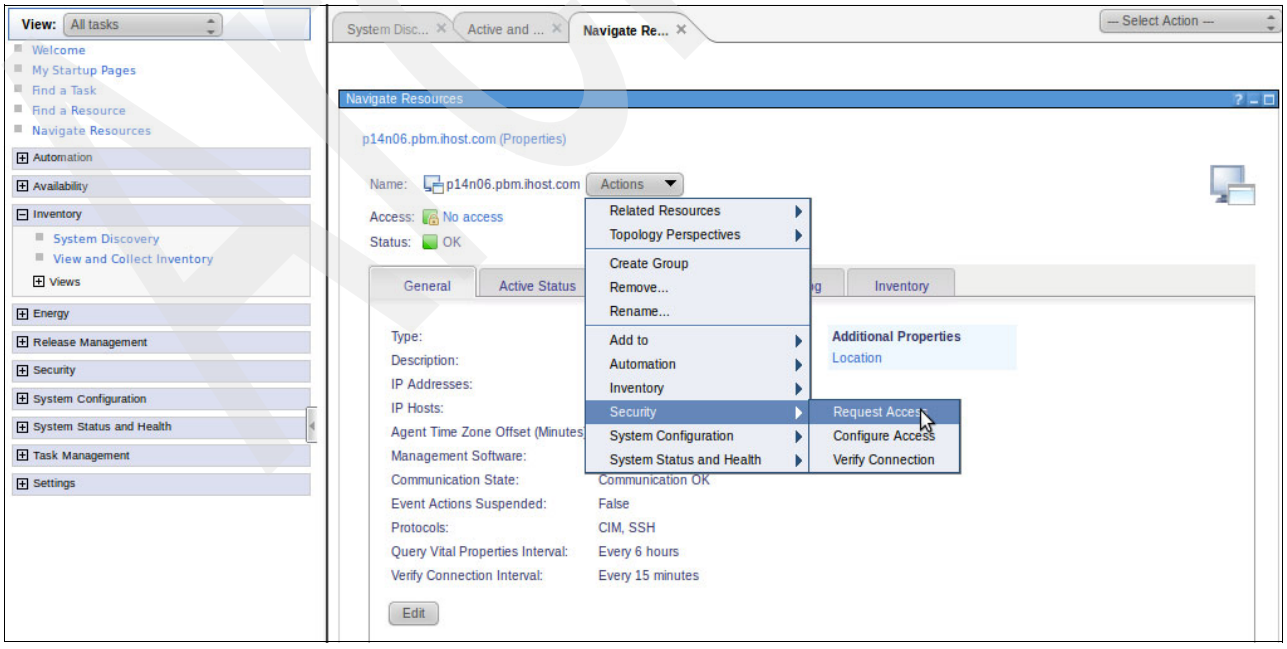

Figure 3-144 Navigating to Request Access

Using the **Request Access** tab, enter the *user ID* and *password* to enable access between the target system and IBM Systems Director, as shown in Figure 3-145.

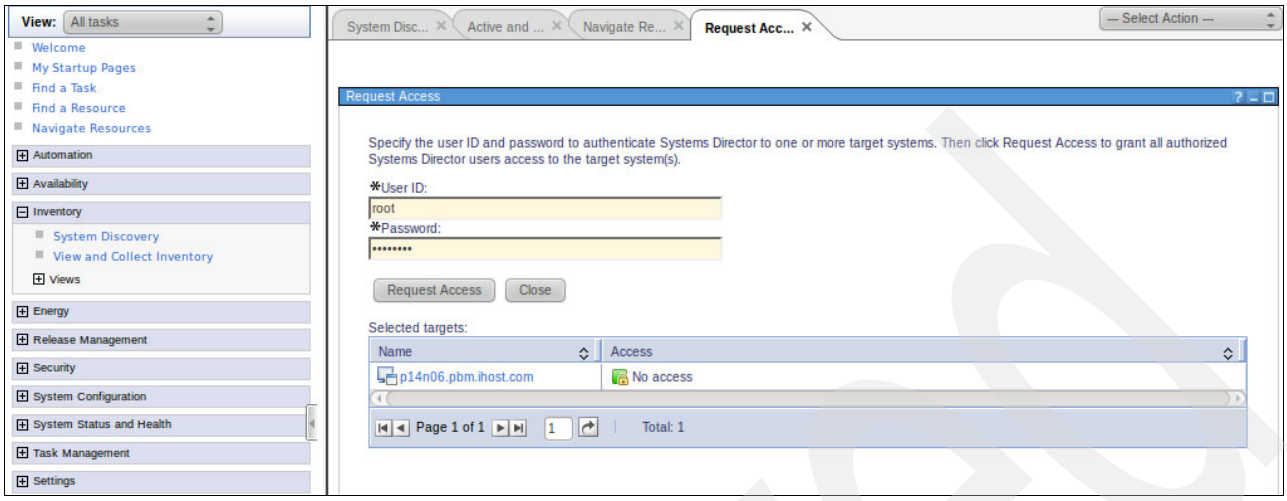

Figure 3-145 Specify the user ID and password for access between the target system and IBM Systems Director

The access column shows OK when access is enabled (Figure 3-146).

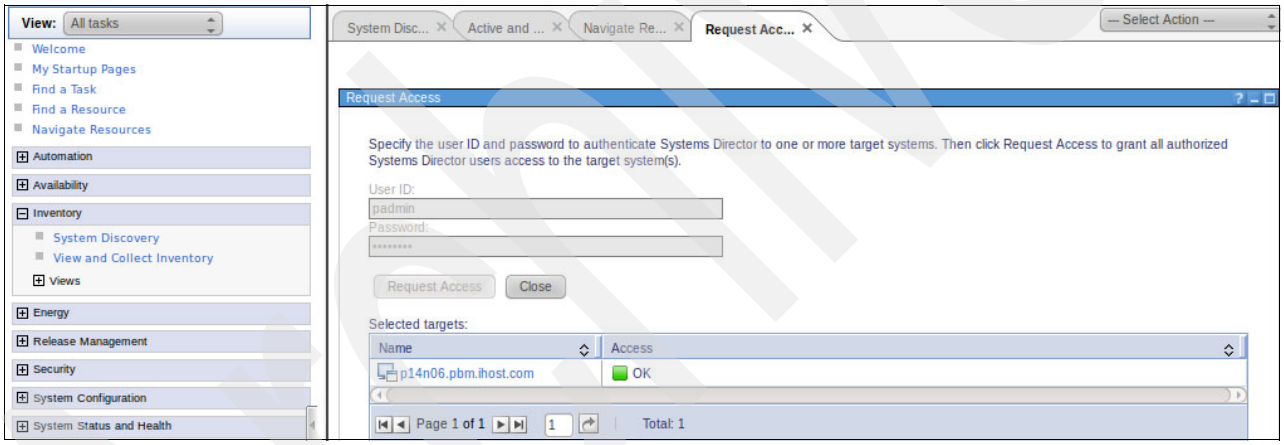

Figure 3-146 Access between target system and IBM Systems Director is now enabled

3.5.6 Inventory

When access has been provided for all discovered items, the inventory must be collected, as shown in Figure 3-147, by clicking the **View and Collect Inventory** tab.

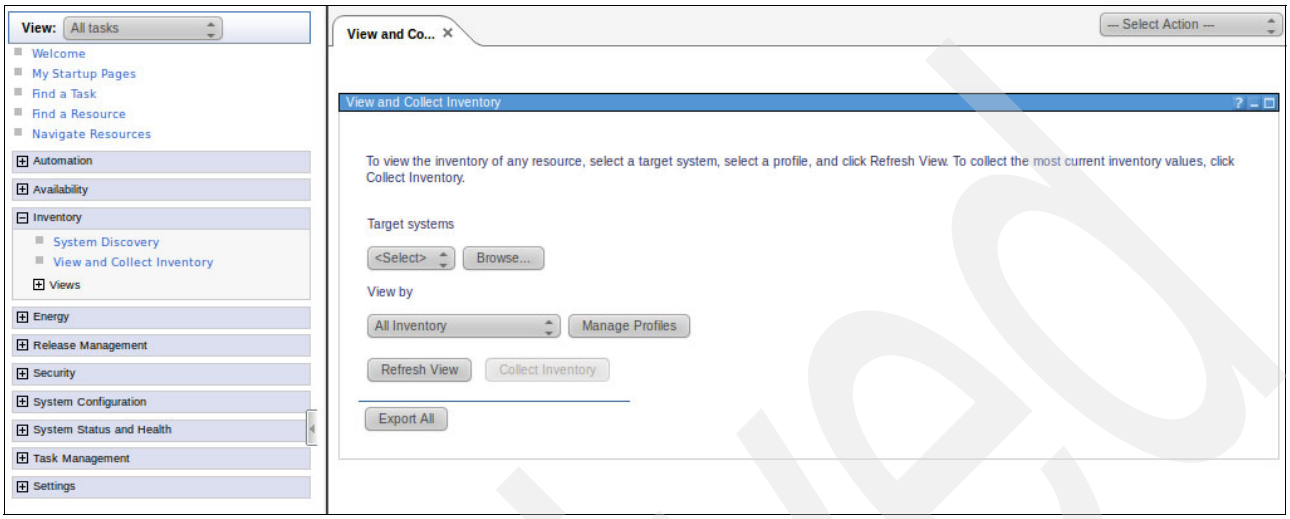

Figure 3-147 View and collect the inventory of discovered items

In Figure 3-147, by clicking **Browse** under Target systems, the following window opens (see Figure 3-148).

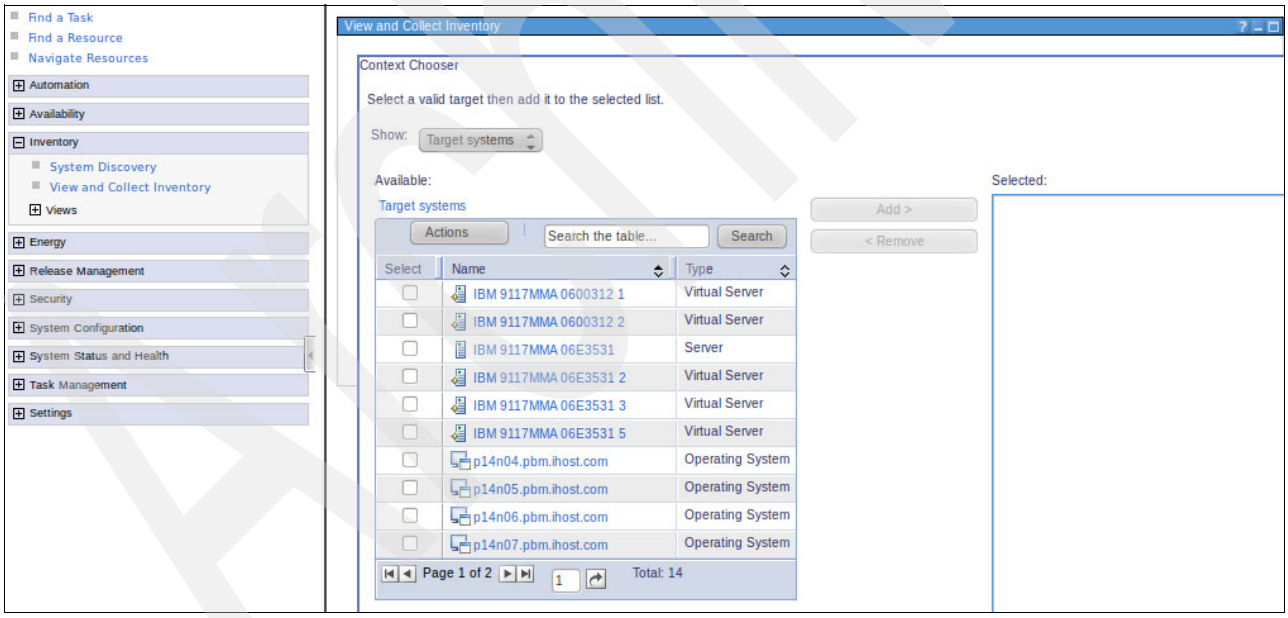

Figure 3-148 Using View and Collect Inventory to select target systems (page 1 of 2)

All of the available Target Systems are listed (as shown in Figure 3-148 page 1 of 2 and also in Figure 3-149 on page 120 page 2 of 2). Target systems are selected by clicking the box next to their name.

Clicking **Add**, as shown in Figure 3-149, causes the entire set of selected system names to appear on the right side in the Selected area of the window.

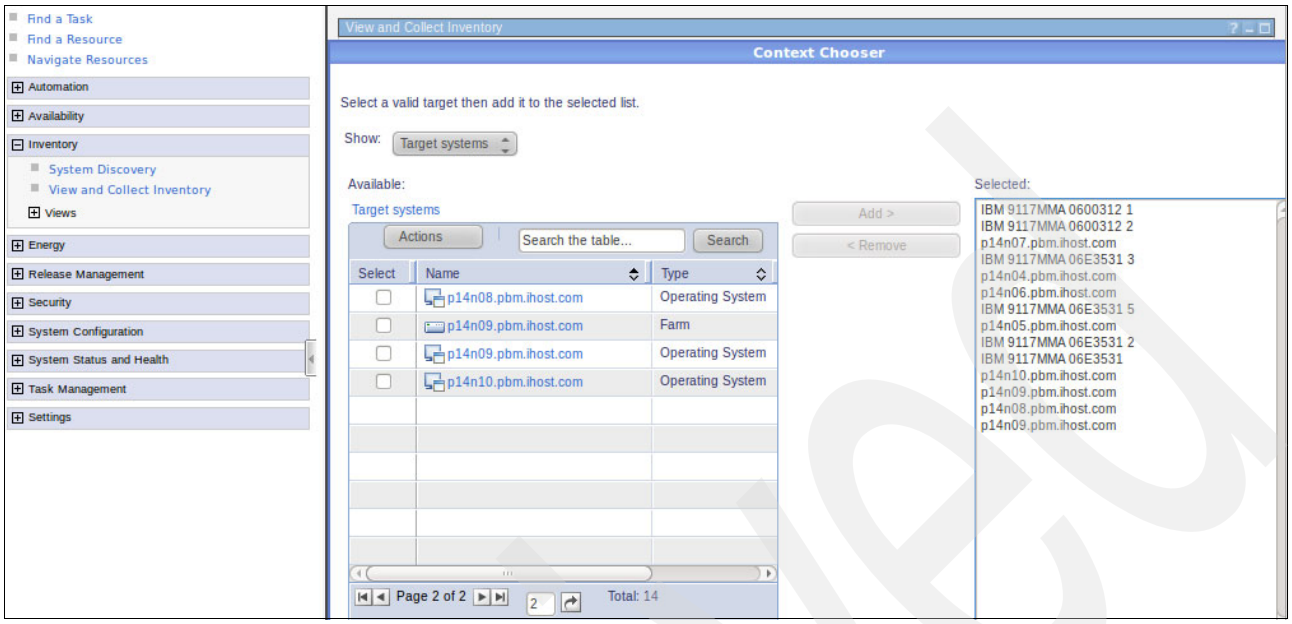

Figure 3-149 Target systems have been added to the selected column

By clicking **Collect Inventory** (see Figure 3-150), a new job scheduling window opens, as shown in Figure 3-151 on page 121.

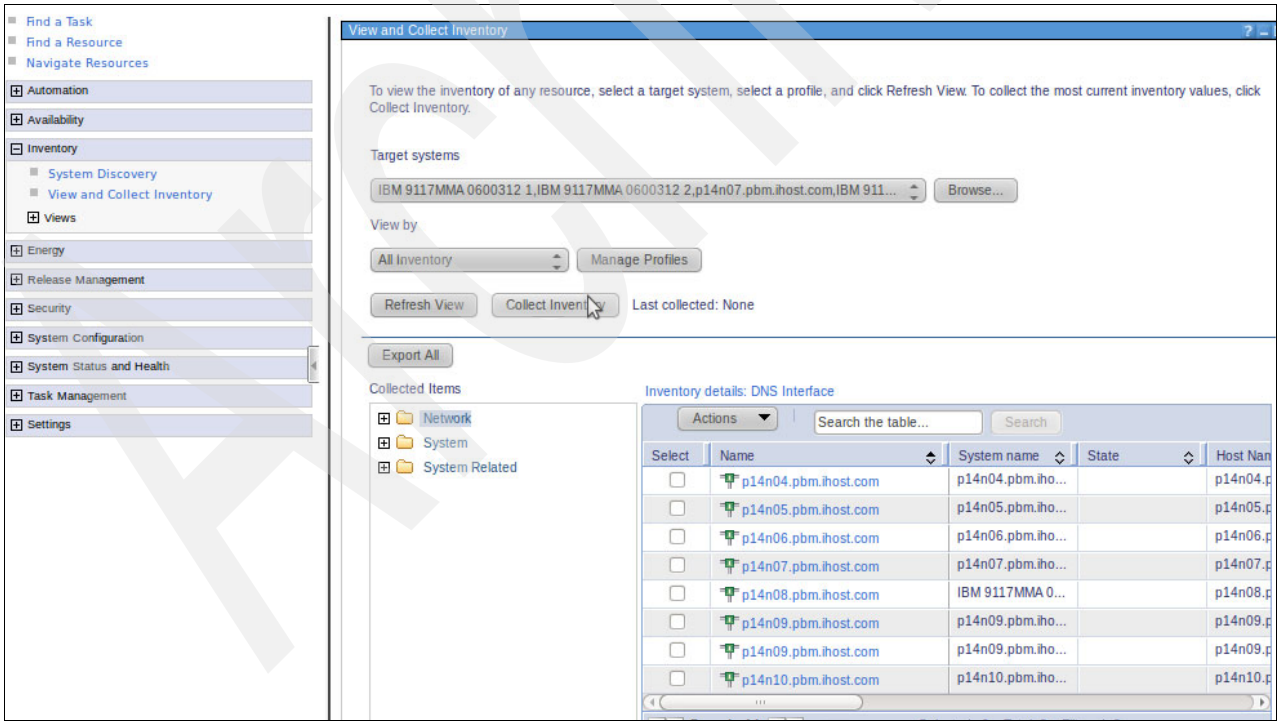

Figure 3-150 Using the Collect Inventory function

The Task Launch Dialog (Figure 3-151) shows that the Job Name `Collect Inventory - October 5, 2011 3:40:18 PM CDT` can either be Run Now or Scheduled. Choose **Run Now** and click **OK**. You can monitor the progress of the job from the Active and Scheduled Jobs section, which can be accessed by using the **Task Management** link that is located on the left portion of the window.

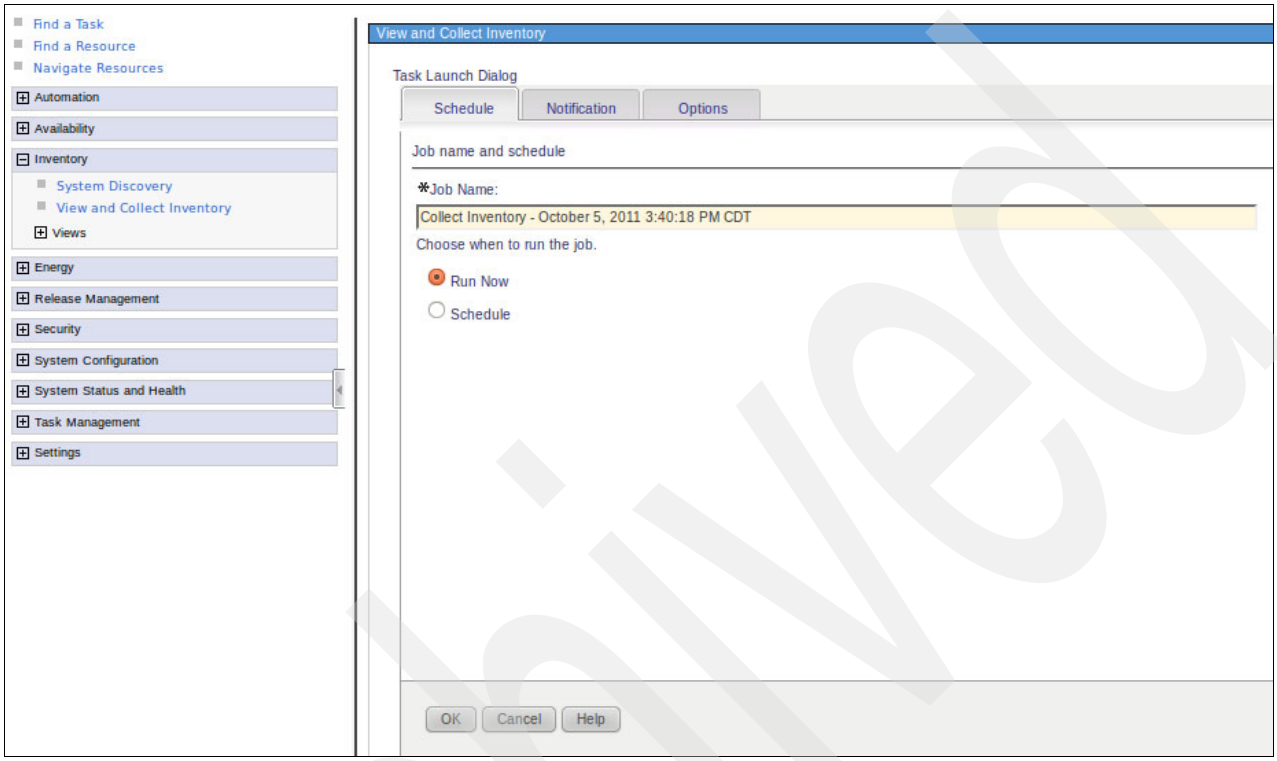

Figure 3-151 Task Launch Dialog of the View and Collect Inventory window

3.5.7 HMC discovery

In 3.5.6, “Inventory” on page 119, we discovered and inventoried the defined LPARs, operating systems, and physical machines. However, the process that was described in 3.5.6, “Inventory” on page 119 used a network range that did not include the range for the HMC. Therefore, we must discover and inventory the HMC separately.

Figure 3-152 shows a single IP address discovery. After typing the IP address, click **Discover Now**.

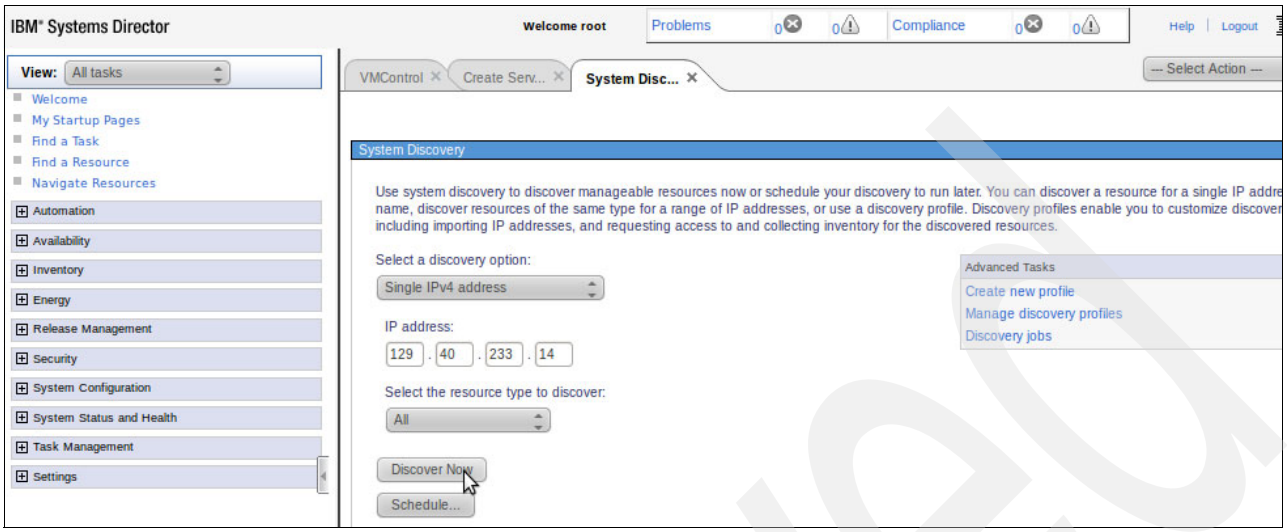

Figure 3-152 Discovering the HMC

After the discovery completes, both the physical machine and operating system have been discovered, as shown in Figure 3-153.

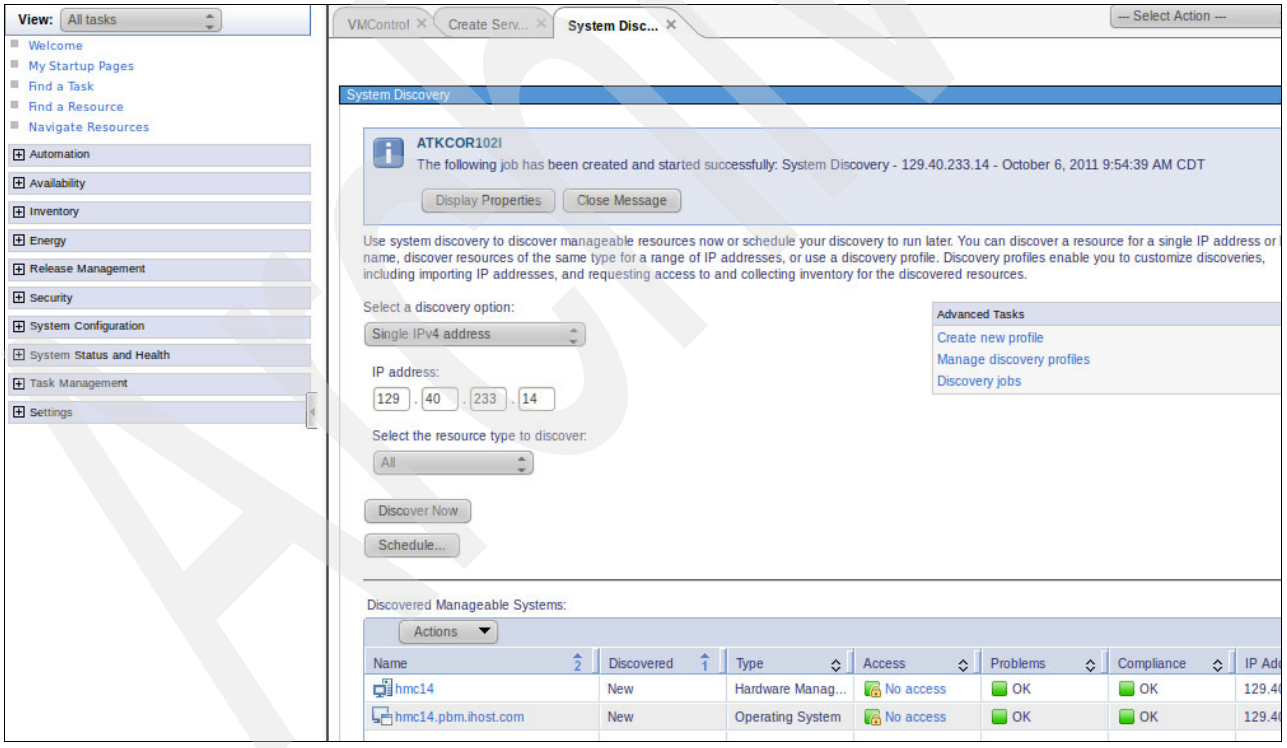

Figure 3-153 The HMC and operating system have been discovered

To allow access to the newly discovered HMC, you must provide a *user ID* and *password*. The *user ID* hscroot was provided, as shown in Figure 3-154.

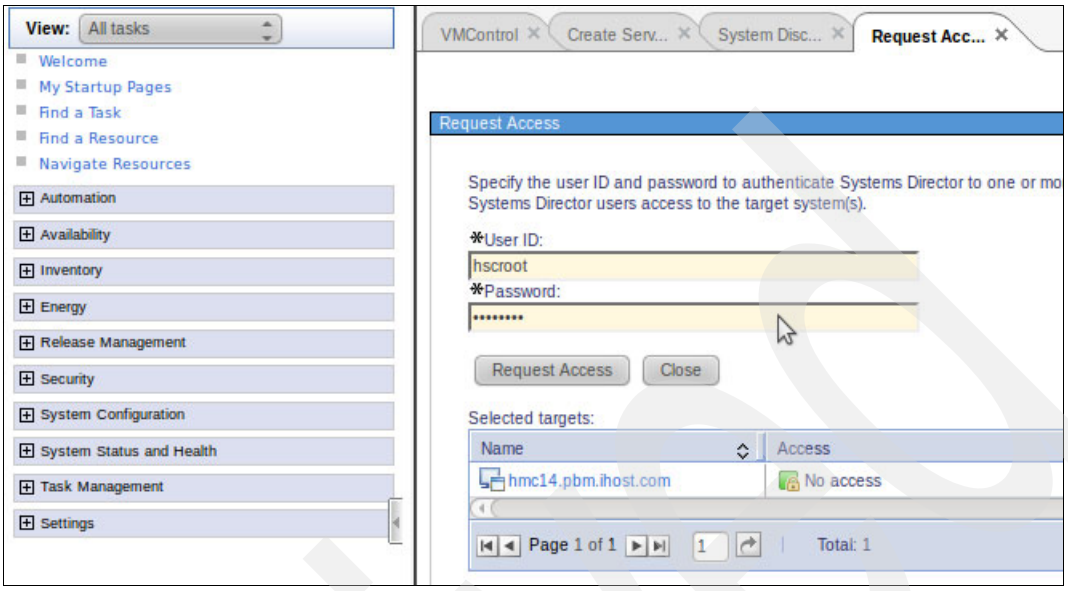

Figure 3-154 Specify the user ID and password for the newly discovered HMC

After you click **Request Access** in Figure 3-154, the access field is changed to OK, as shown in Figure 3-155.

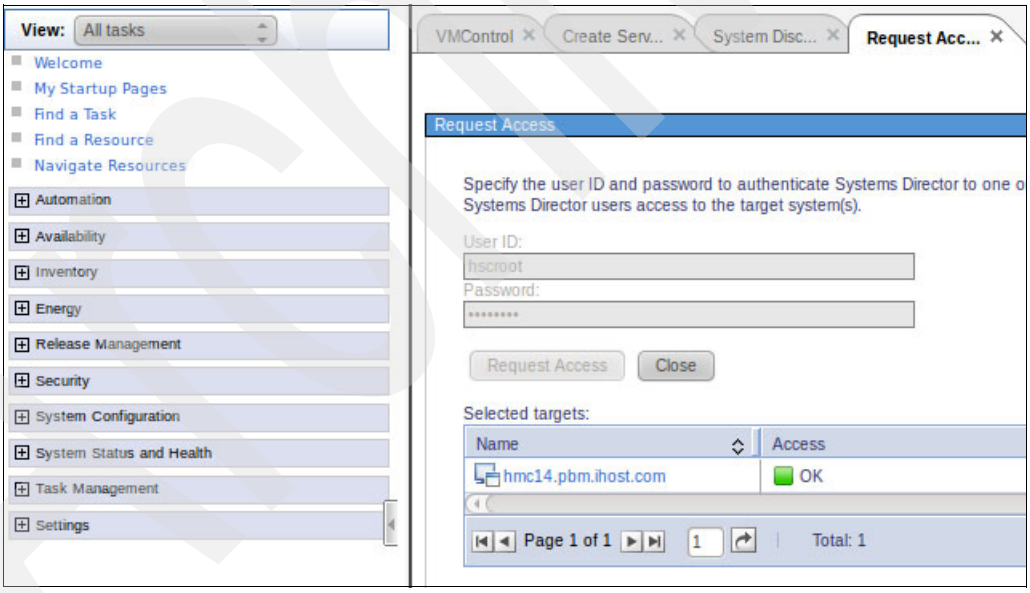

Figure 3-155 Requested access to HMC is now OK

3.5.8 HMC subagent install

VMControl needs a special agent that is installed on the NIM server. The first step of this installation process is to select **Install Agents** from the Common Tasks panel, which is located in the lower-right corner of the main VMControl Enterprise Edition window, as shown in Figure 3-156.

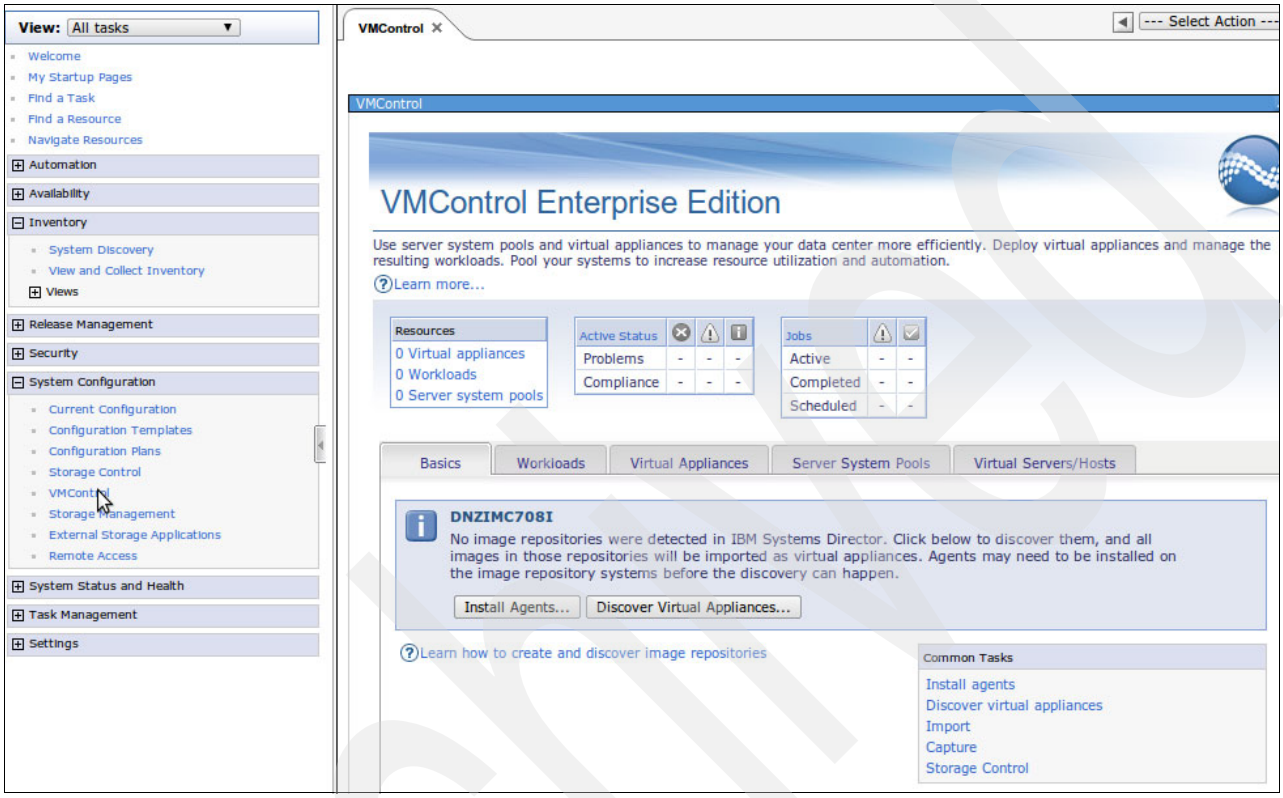

Figure 3-156 Select Install Agents from the Common Tasks panel

The **Agent Installation** wizard is started (Figure 3-157).

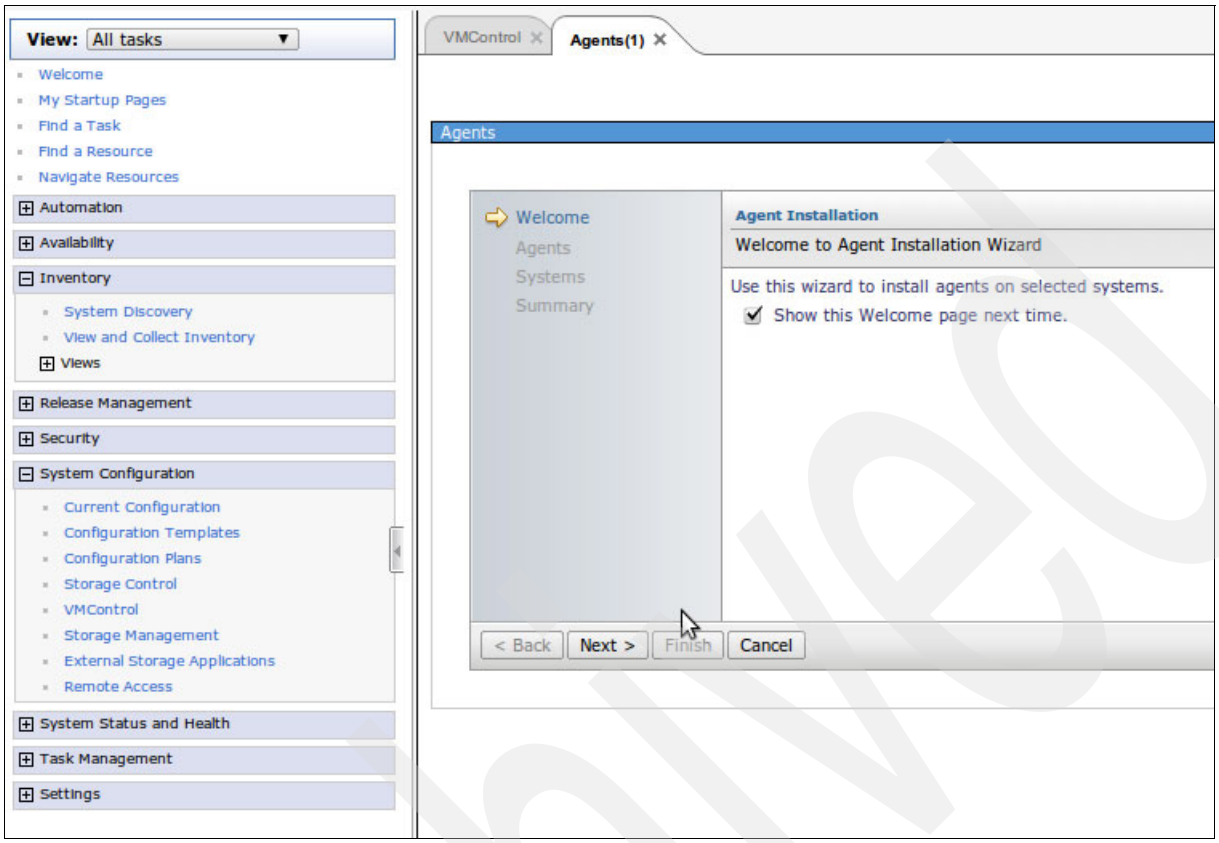

Figure 3-157 The Agent Installation wizard is started

There are many types of agents that can be installed from this wizard. The **Common Agent Subagent Packages** is selected and then added to the Selected area of the window, as shown in Figure 3-158.

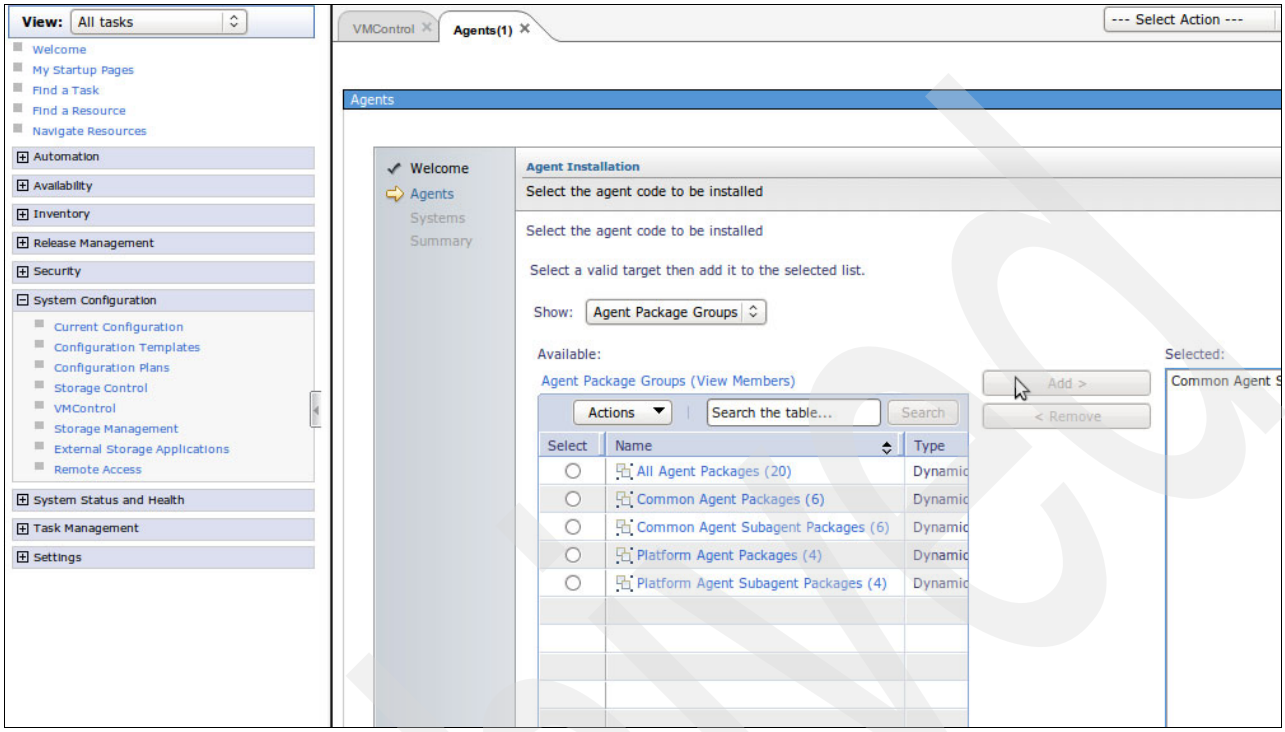

Figure 3-158 Selecting Common Agent Subagent Packages

In the same way that the Common Agent Subagent Packages were selected and added, **CommonAgentSubagent_VMControl_NIM** can now be added to the Selected area of the window, as shown in Figures Figure 3-159 and Figure 3-160 on page 128.

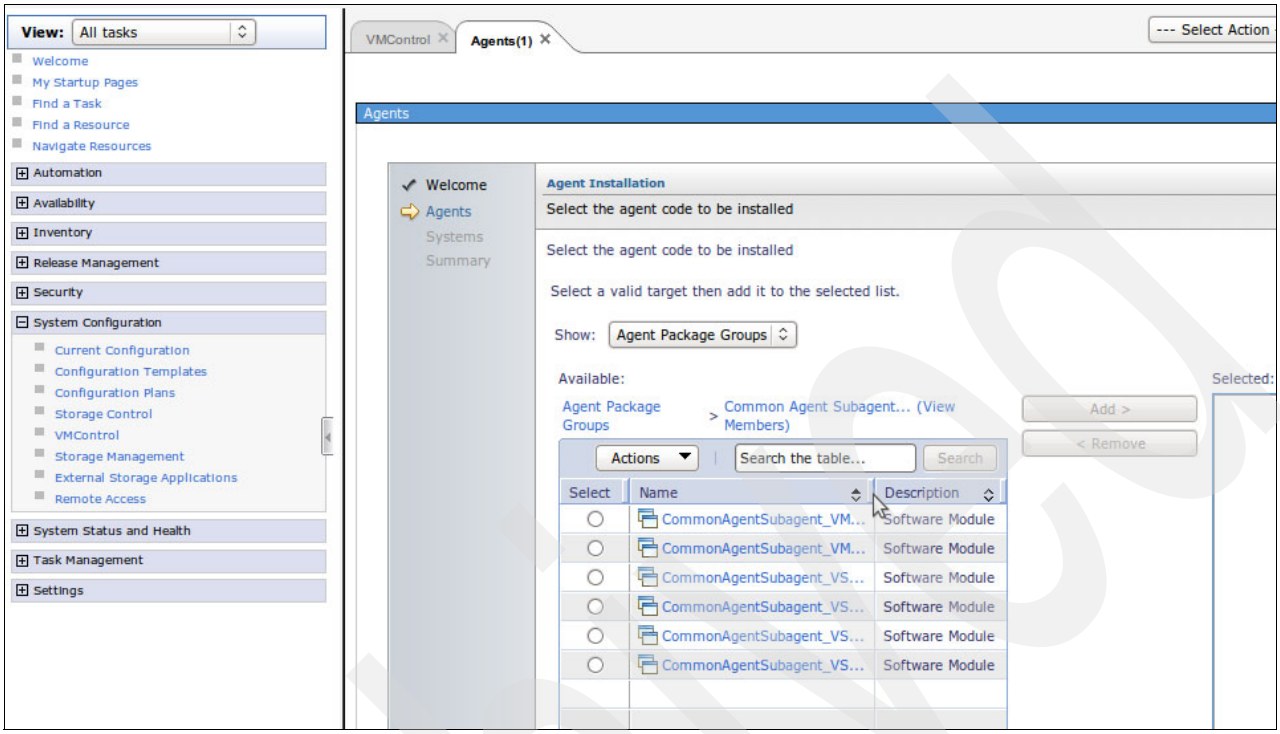

Figure 3-159 CommonAgentSubagent_VMControl_NIM can now be added to the Selected area of the window

The wizard presents a list of available operating systems on which to install the subagent. These operating systems were previously discovered, and proper access was granted to the Systems Director. In Figure 3-160, the NIM server for this environment is selected.

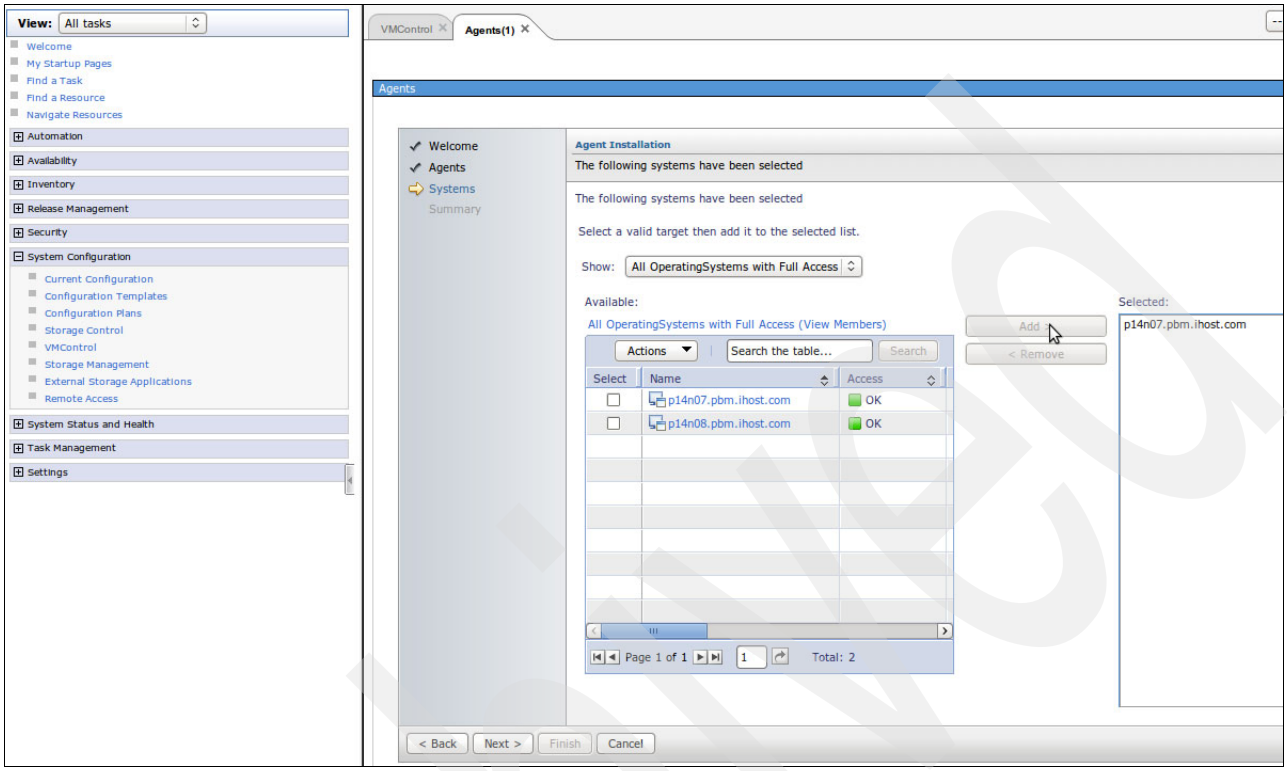

Figure 3-160 NIM server for this environment is selected

3.5.9 HMC inventory

Now that there is access to the discovered HMC, and the subagent is properly installed, the inventory must be collected. Expand the **Inventory** node and choose **View and Collect Inventory**, as shown in Figure 3-161.

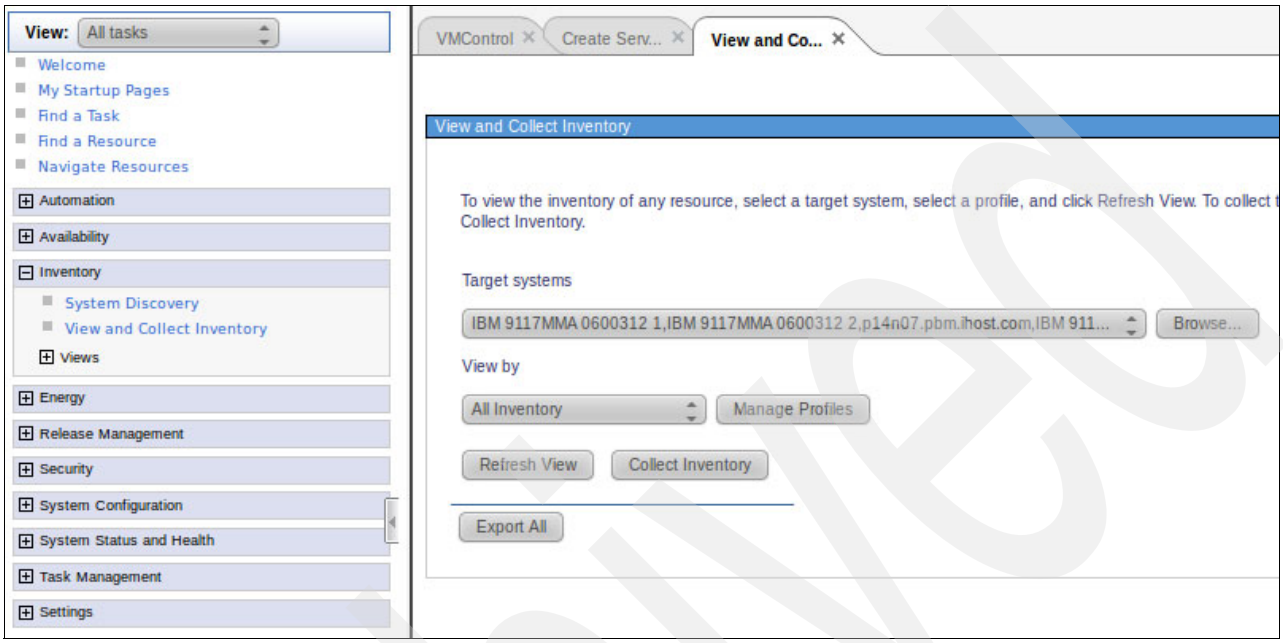

Figure 3-161 Collecting inventory for the discovered HMC

After the two items representing the HMC are added to the Selected area and **OK** is clicked, the following window (Figure 3-162) opens.

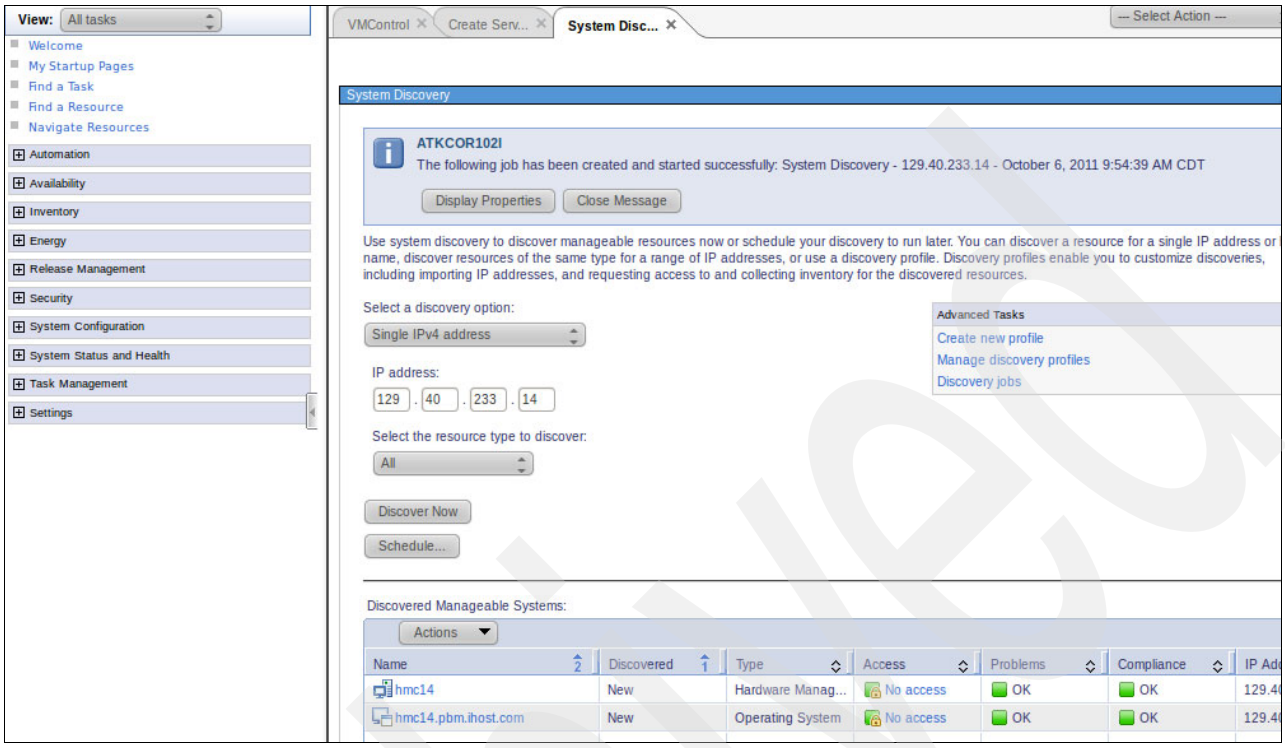

Figure 3-162 Selecting the HMC items

From the System Discovery window, **Create new profile** is selected from the Advanced Tasks box. This action starts the discovery profile wizard, as shown in Figure 3-163. We clicked **Collect Inventory** to run the inventory job.

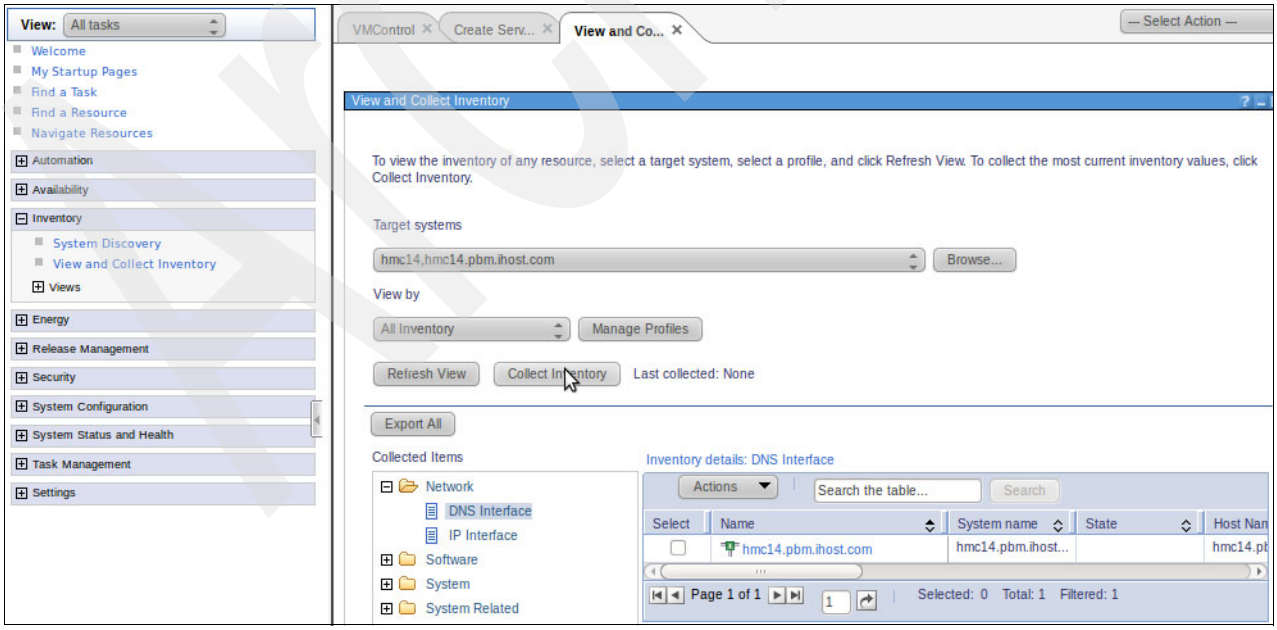

Figure 3-163 The discovery profile wizard View and Collect tab

From here, the profile is given a name. The resource type is set to **Farm** and a simple description is added.

3.5.10 Tivoli Storage Productivity Center discovery

At this point, a majority of the *Managed Through* has been defined through discovery and inventory collection.

You can add the storage infrastructure now. Because you are using Tivoli Storage Productivity Center to administer and manage both the storage subsystem and SAN fabric, you need to create a new discovery profile using the **System Discovery** tab, as depicted in Figure 3-164.

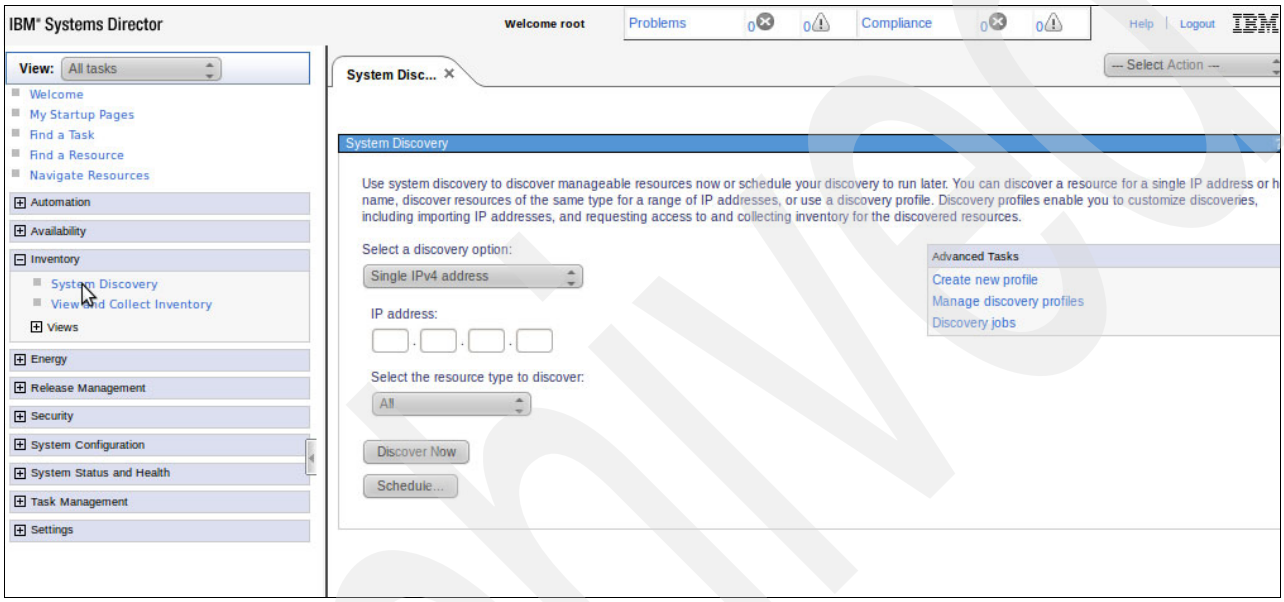

Figure 3-164 A discovery profile is needed for both the storage subsystem and SAN fabric

After choosing **Create new profile** from the Advanced Tasks menu in Figure 3-164, the following Create and edit discovery profiles panel (Figure 3-165) appears.

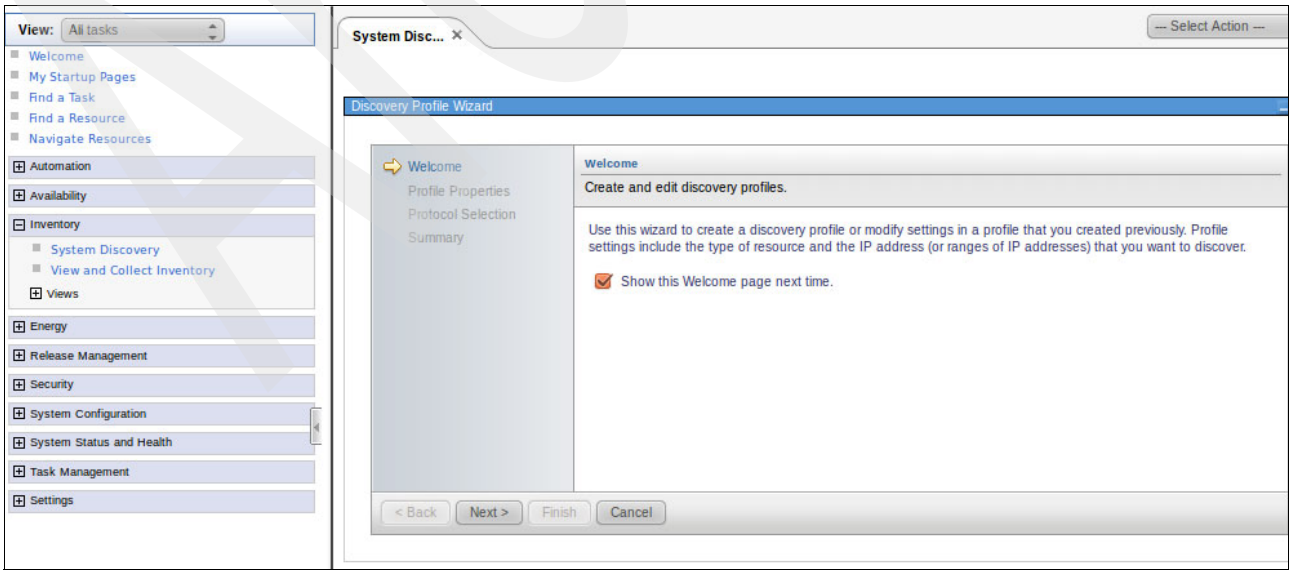

Figure 3-165 Create and edit discovery profiles

The next step is to specify the protocol. The default, **TPC Inventory Discovery**, was selected, as shown in Figure 3-166.

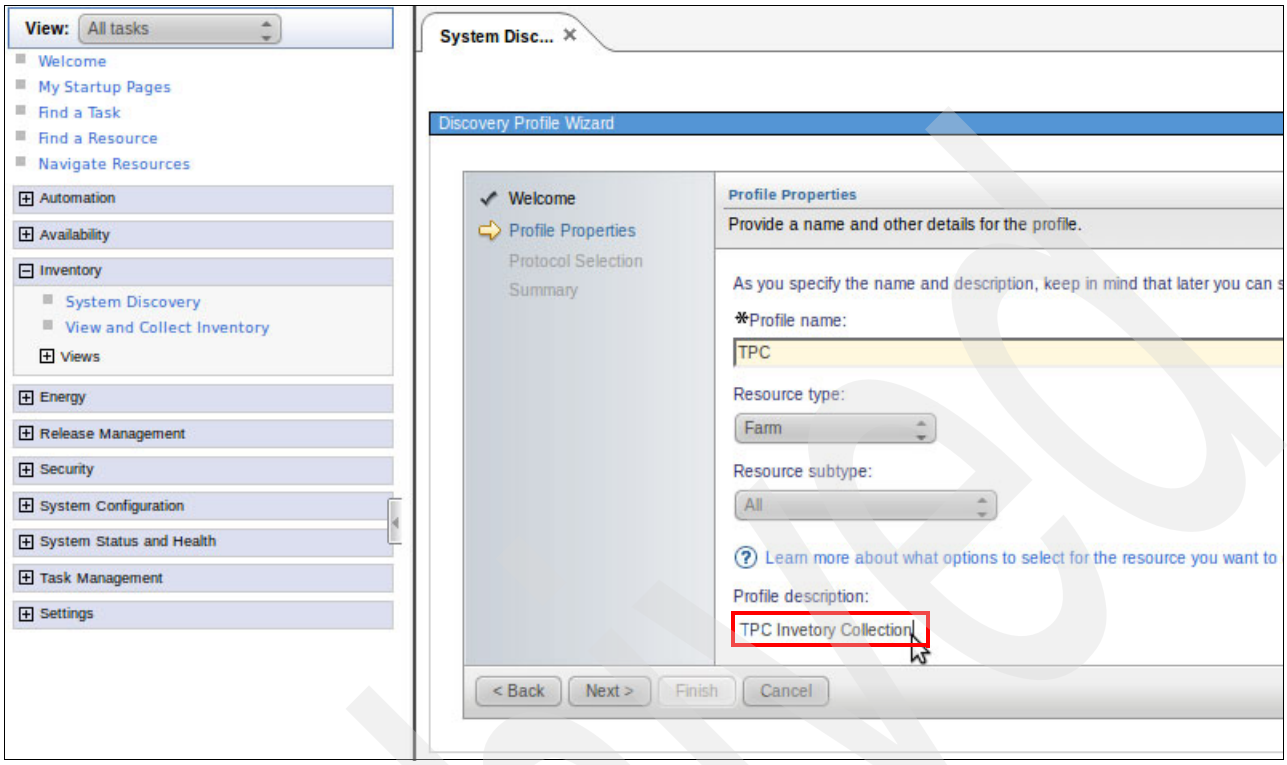

Figure 3-166 Use the TPC default profile

The options that are shown in Figure 3-166 represent the data that was entered when Tivoli Storage Productivity Center was installed, which includes the database name, user ID, password, and port. The Tivoli Storage Productivity Center user information is entered. This information was provided when installing Tivoli Storage Productivity Center.

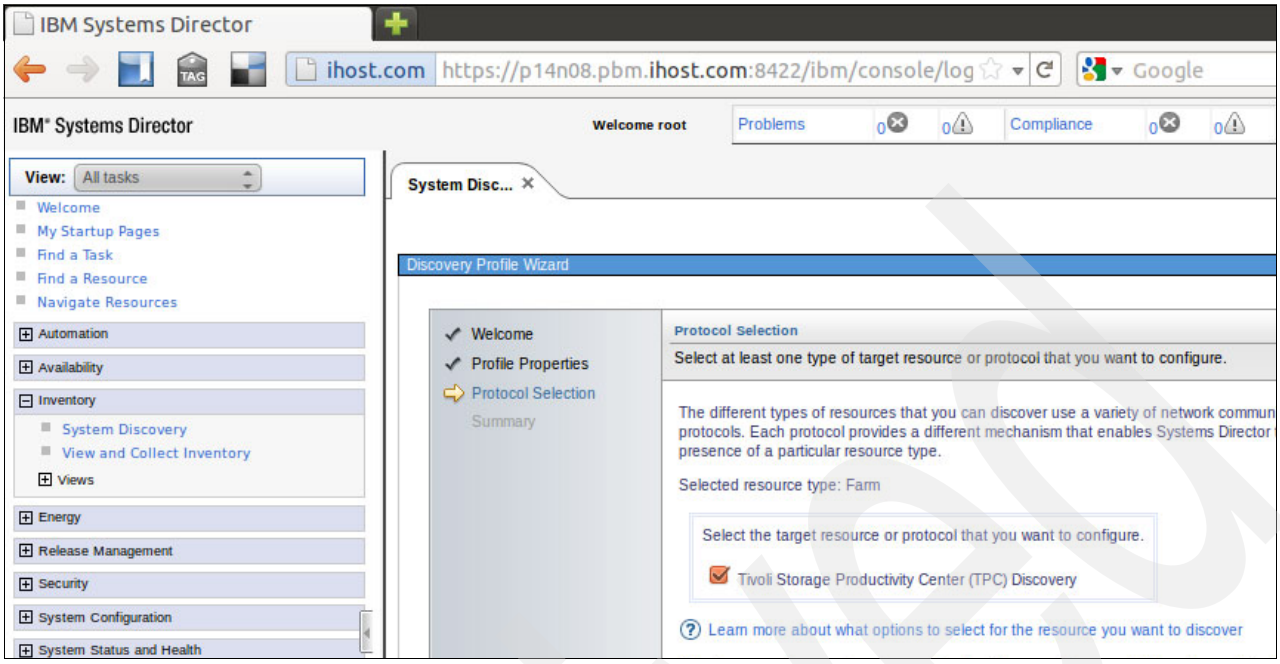

Figure 3-167 Select at least one type of target resource or protocol that you want to configure

The resource group is selected next, as shown in Figure 3-167. For the purposes of this document, all resources are selected. The window that is shown in Figure 3-168 opens.

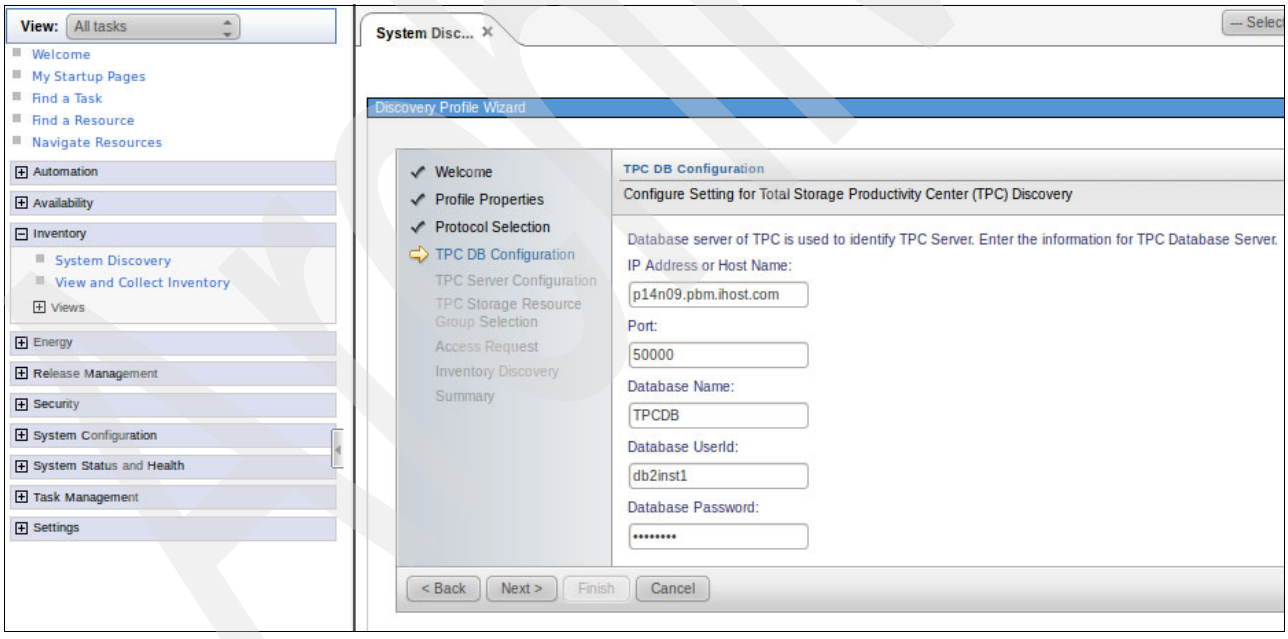

Figure 3-168 Configure Settings for Total Storage Productivity Center (TPC) Discovery

After Discovery, access must be obtained. The window that is shown in Figure 3-169 allows you to enter the information as part of the discovery process.

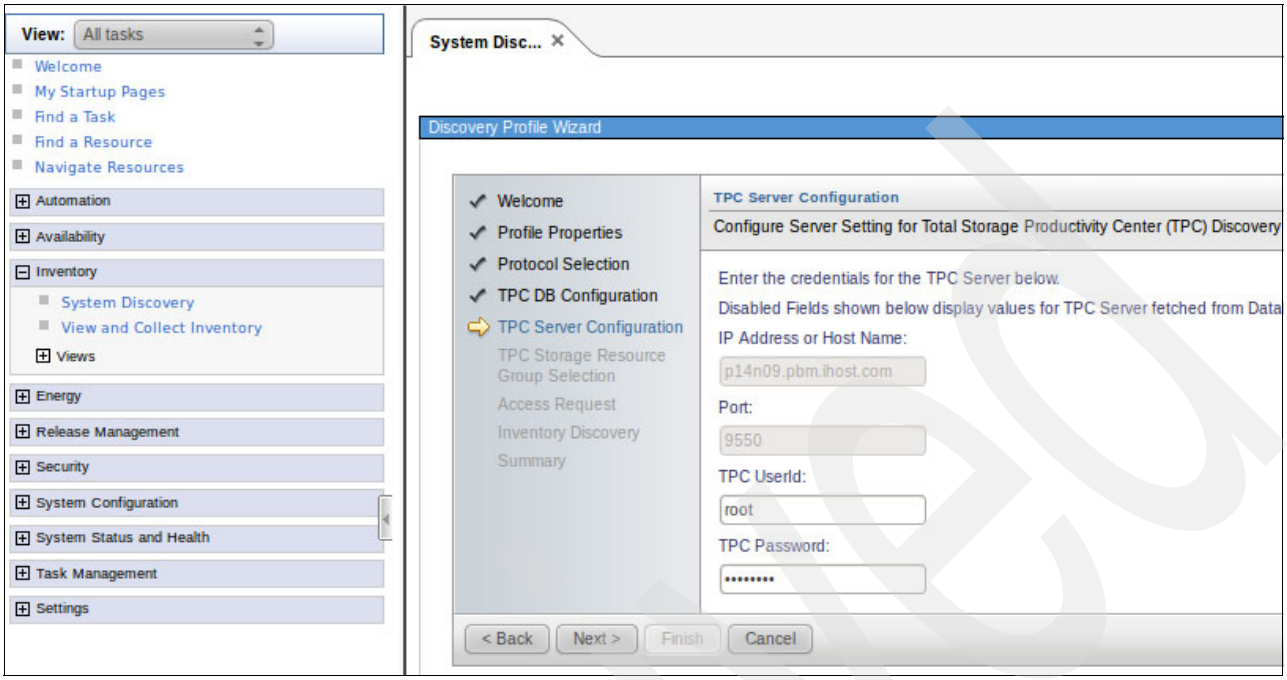

Figure 3-169 TPC Server Configuration

For simplicity, inventory is collected automatically. It probes Tivoli Storage Productivity Center and imports the storage area network configuration. Figure 3-170 shows the panel where you select the system resource group to be associated with this virtual system pool.

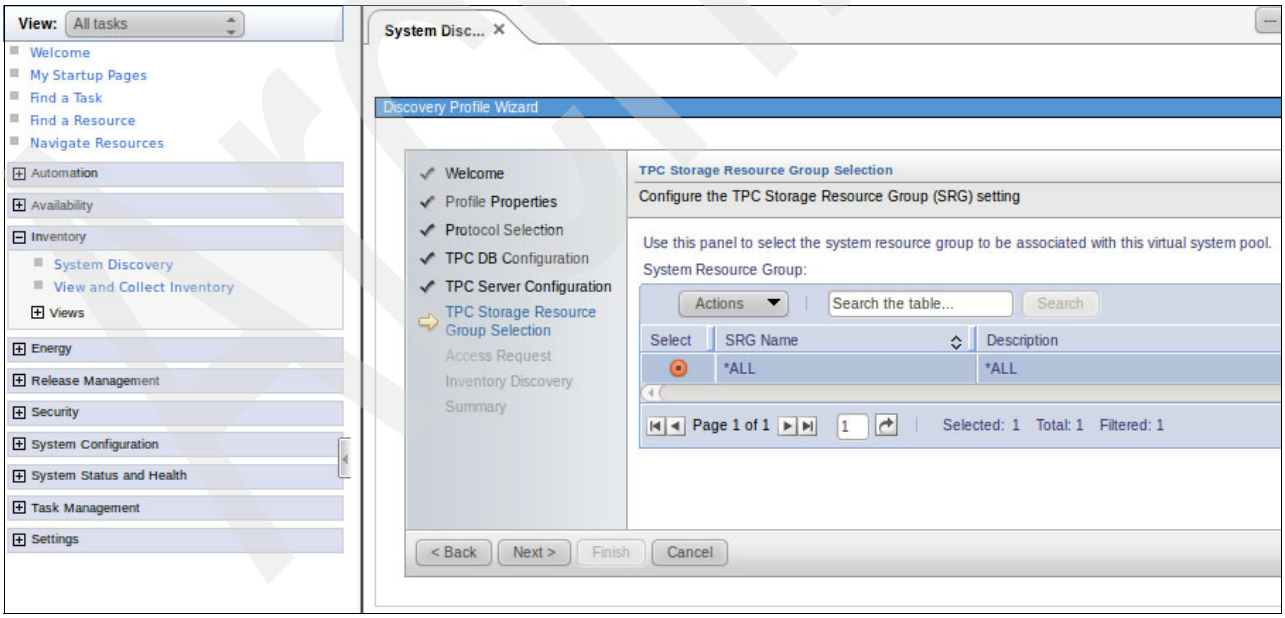

Figure 3-170 Configure the TPC Storage Resource Group (SRG) setting

Figure 3-171 allows you to configure settings to automatically request access to the discovered systems.

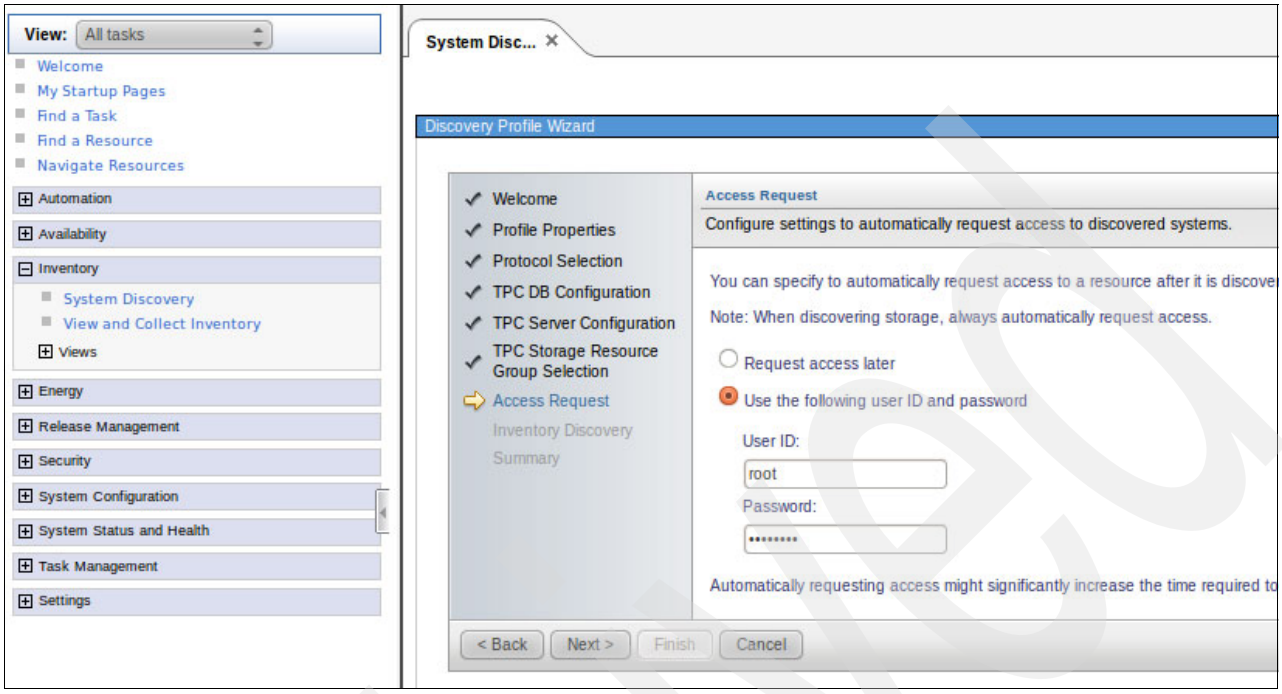

Figure 3-171 Configure settings to automatically request access to discovered systems

In Figure 3-171, the user ID `root` and corresponding password combination is provided to allow access to the discovered system.

The settings for automatic discovery are specified in Figure 3-172 by clicking **Automatically collect inventory, using the All Inventory default inventory profile**.

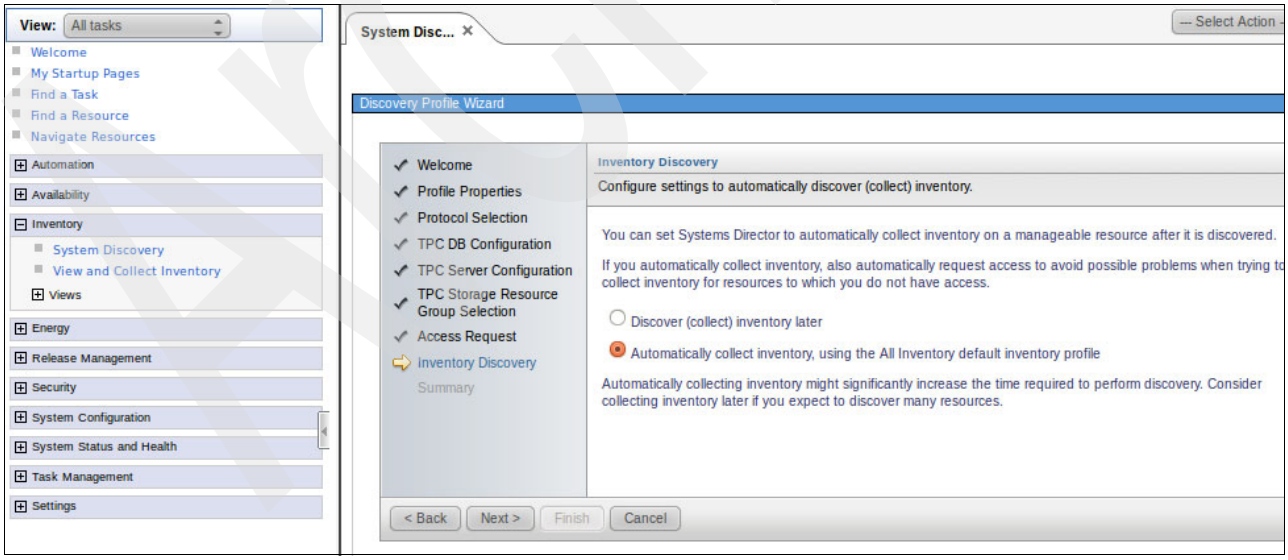

Figure 3-172 Automatically collect inventory using the All Inventory default inventory profile

After clicking **Next** in Figure 3-172 on page 135, a summary of your profile selections is displayed, as shown in Figure 3-173.

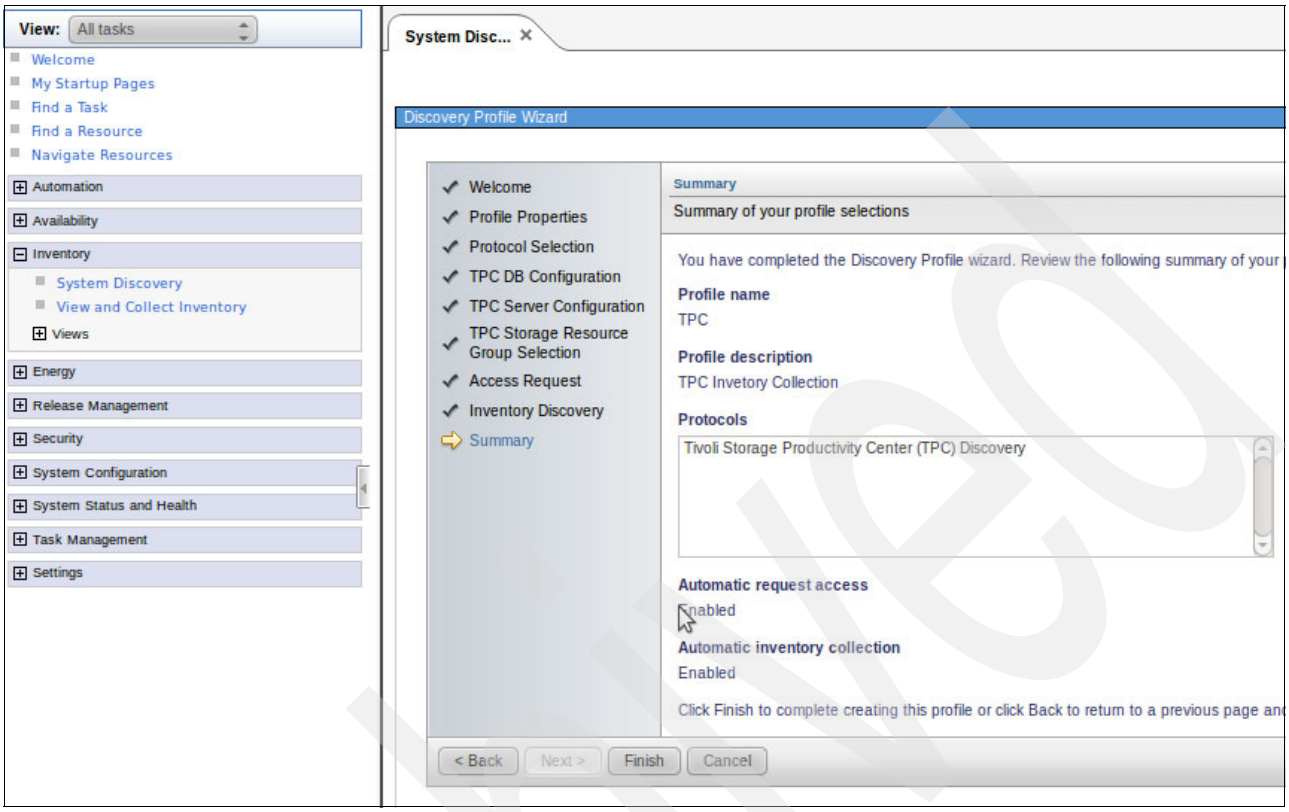

Figure 3-173 Summary of profile selections

Now that the discovery profile is successfully created, **Discover Now** can be run, as shown in Figure 3-174.

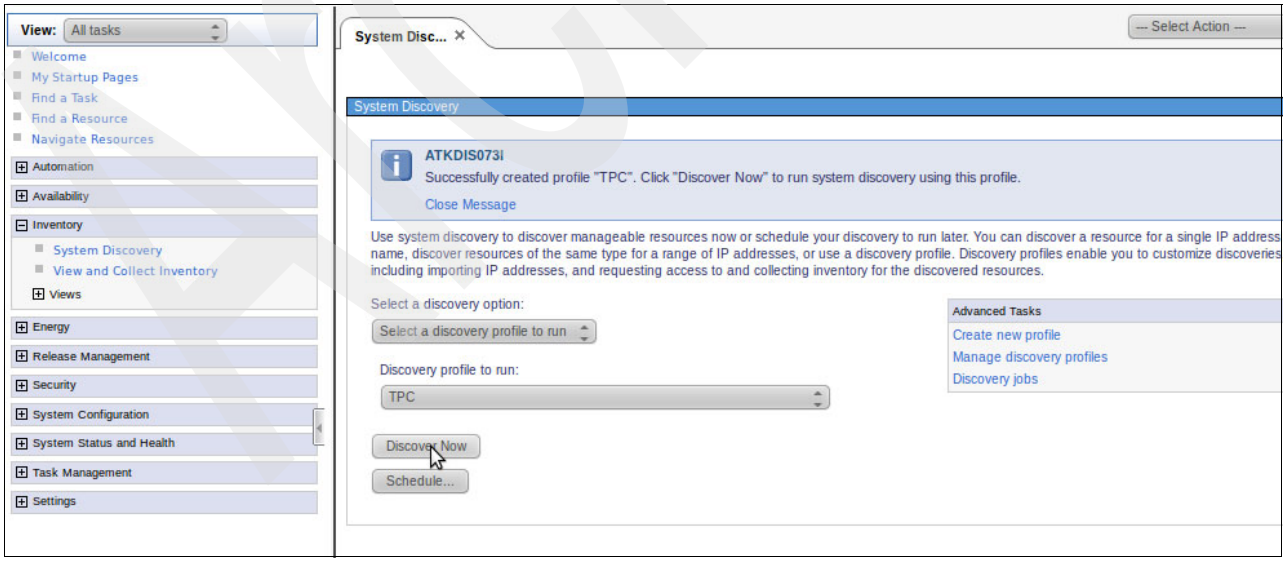

Figure 3-174 Running Discover Now using the TPC discovery profile

After the discovery completes, the inventory can be collected. In Figure 3-175, **p14n09.pbm.ihost.com** (of Type = Farm) is selected.

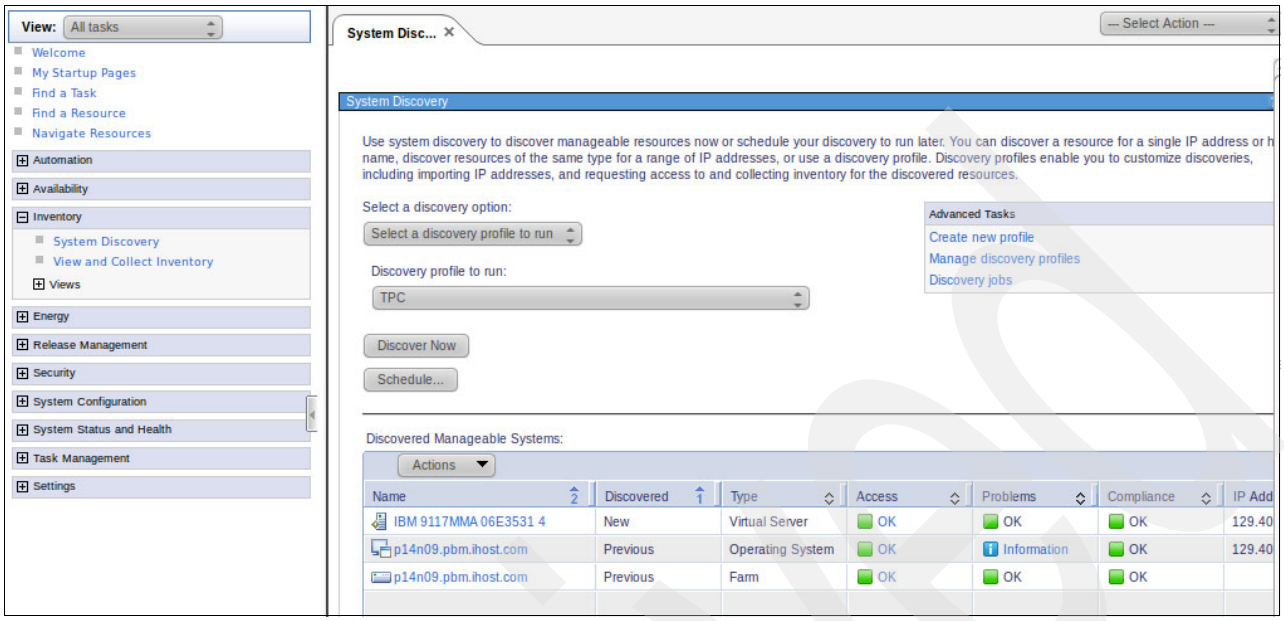

Figure 3-175 Select p14n09.pbm.ihost.com (of Type = Farm)

Figure 3-176 shows the discovered Tivoli Storage Productivity Center items that represent the SAN fabric and storage system.

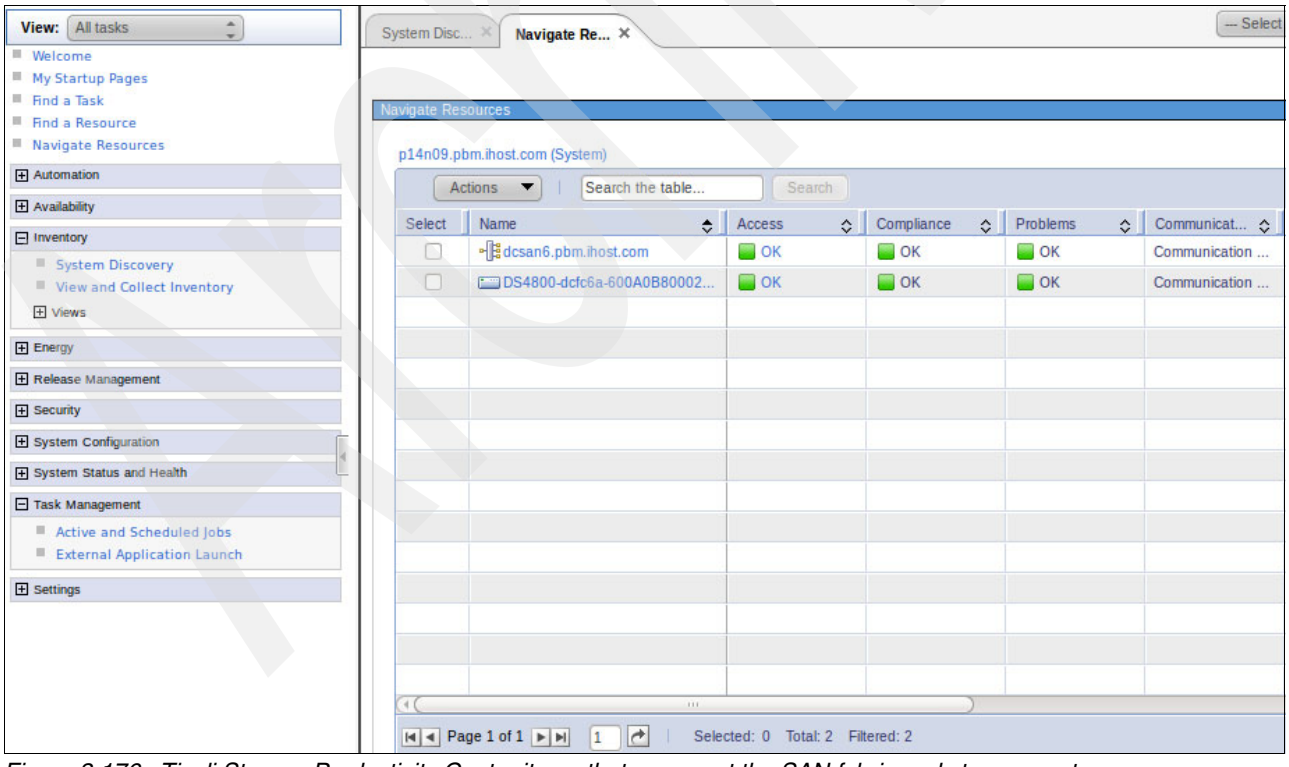

Figure 3-176 Tivoli Storage Productivity Center items that represent the SAN fabric and storage system

To verify that the information that the Director inventoried is equal to that of Tivoli Storage Productivity Center, the topology perspective can be viewed. From the Actions menu, select **Topology Perspectives** → **Storage**, **as** shown in Figure 3-177.

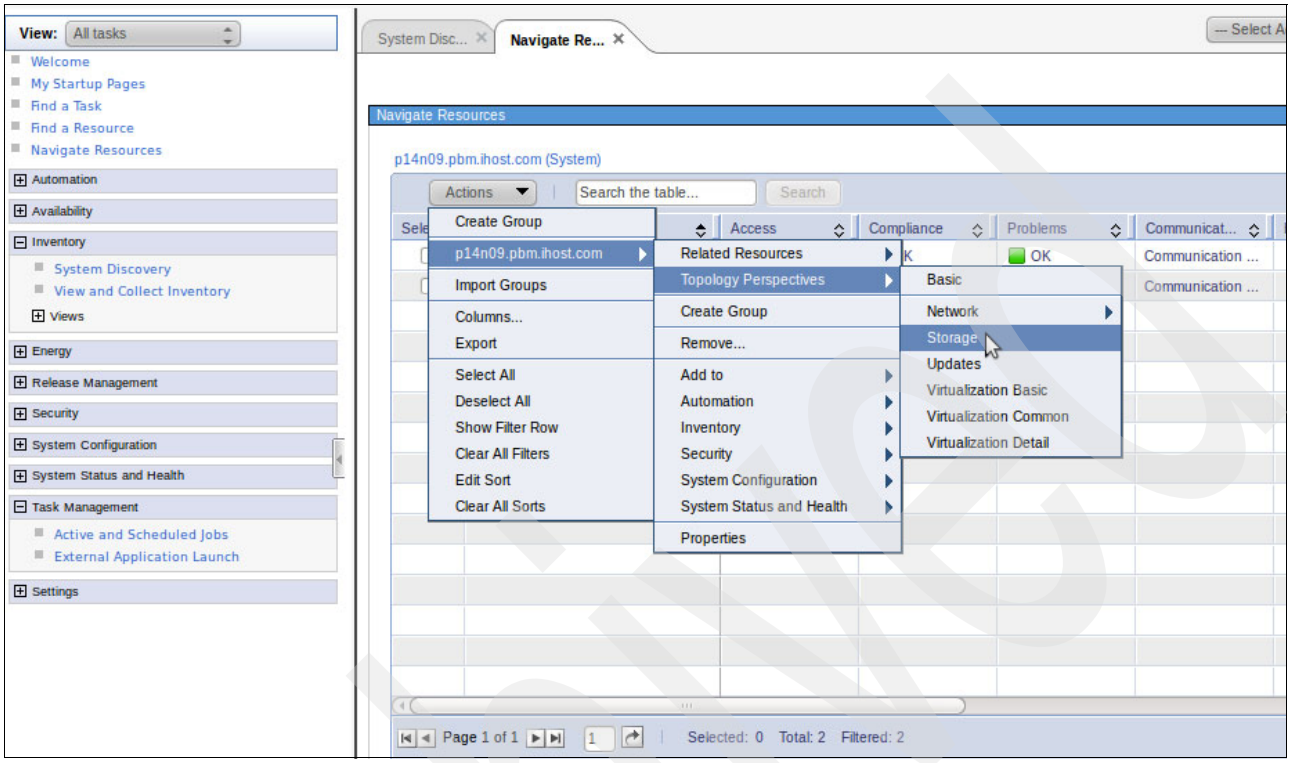

Figure 3-177 Verify that the information that the Director inventoried equals the Tivoli Storage Productivity Center information

Figure 3-178 shows the topology of the storage area network. This information is nearly identical to that shown in Figure 3-140 on page 115.

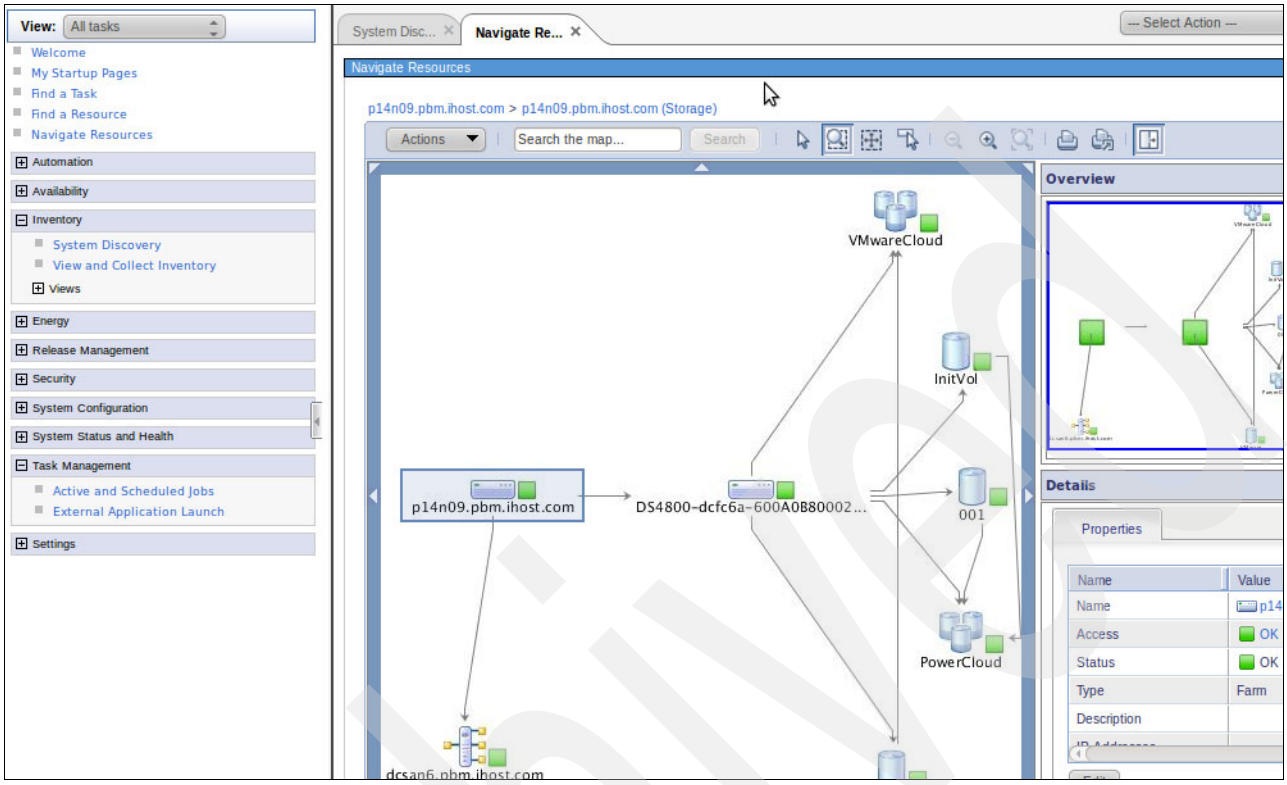

Figure 3-178 The topology of the storage area network

3.5.11 Storage pool

At this point, the entire infrastructure is defined to the IBM Systems Director. VMControl provides the capability to have single access to the storage environment. This concept is called *storage system pools*.

To create a system pool for the items that were inventoried in the previous section (3.5.10, “Tivoli Storage Productivity Center discovery” on page 131), open the main **VMControl** tab, and then, select **Storage system pools** from the Virtualization tasks menu in the lower-right corner, as shown in Figure 3-179.

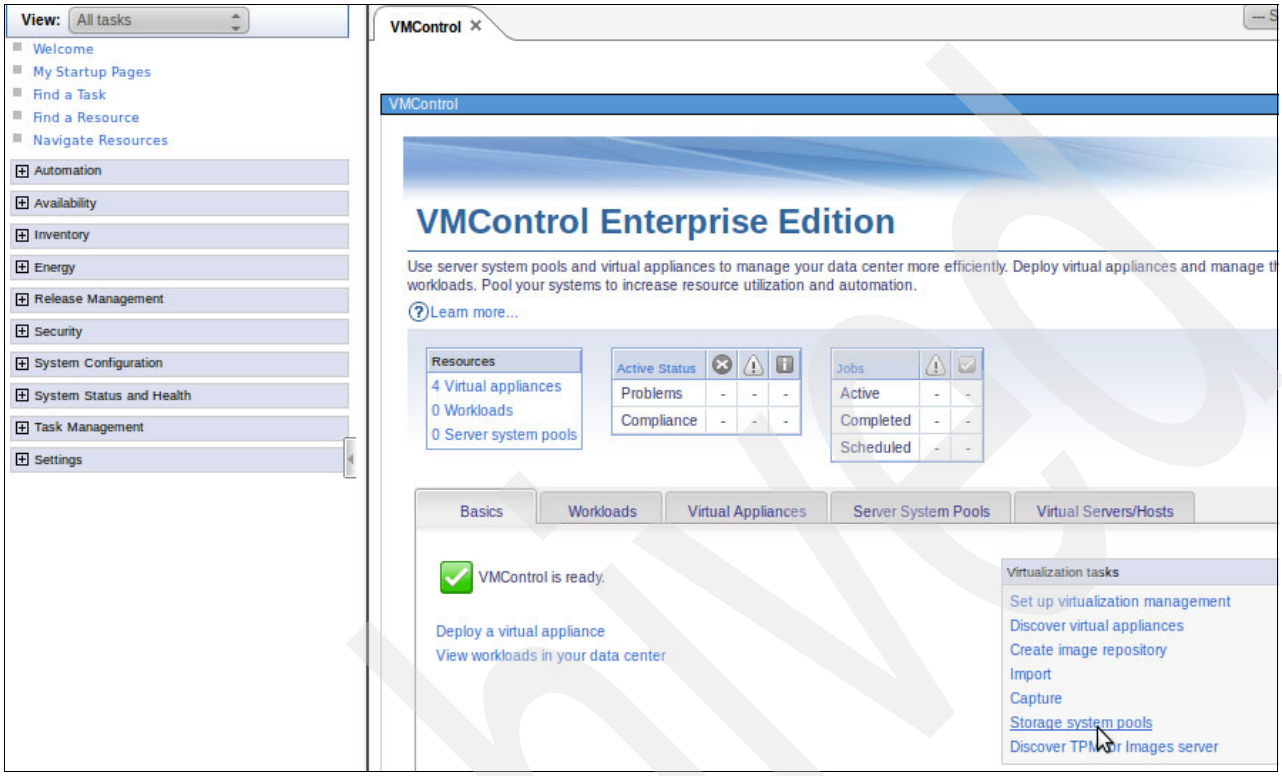

Figure 3-179 Selecting storage system pools from the VMControl tab Virtualization tasks menu

The Storage system pools tab opens, as shown in Figure 3-180.

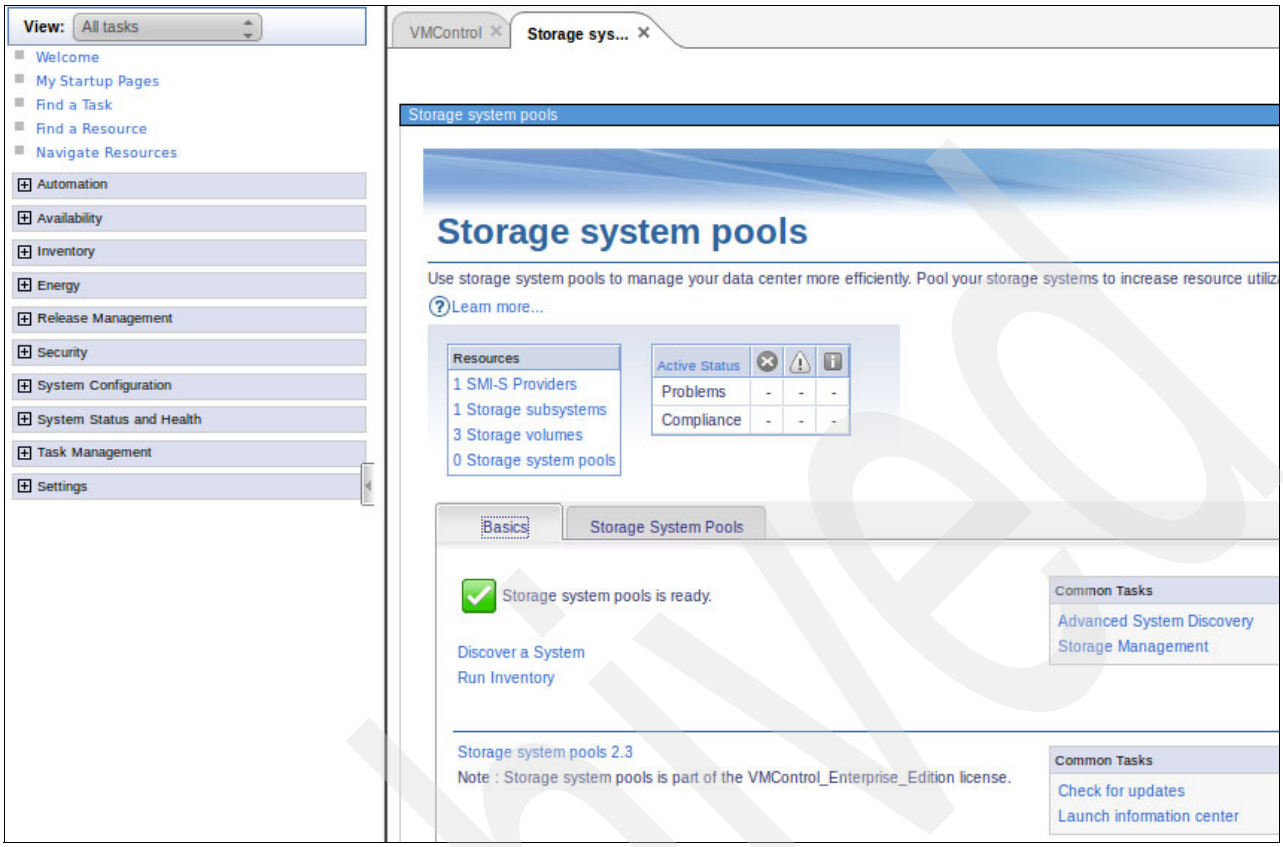

Figure 3-180 Select the Storage system pools tab

On the Storage system pools page, on the bottom task bar, select **Create**, as shown in Figure 3-181.

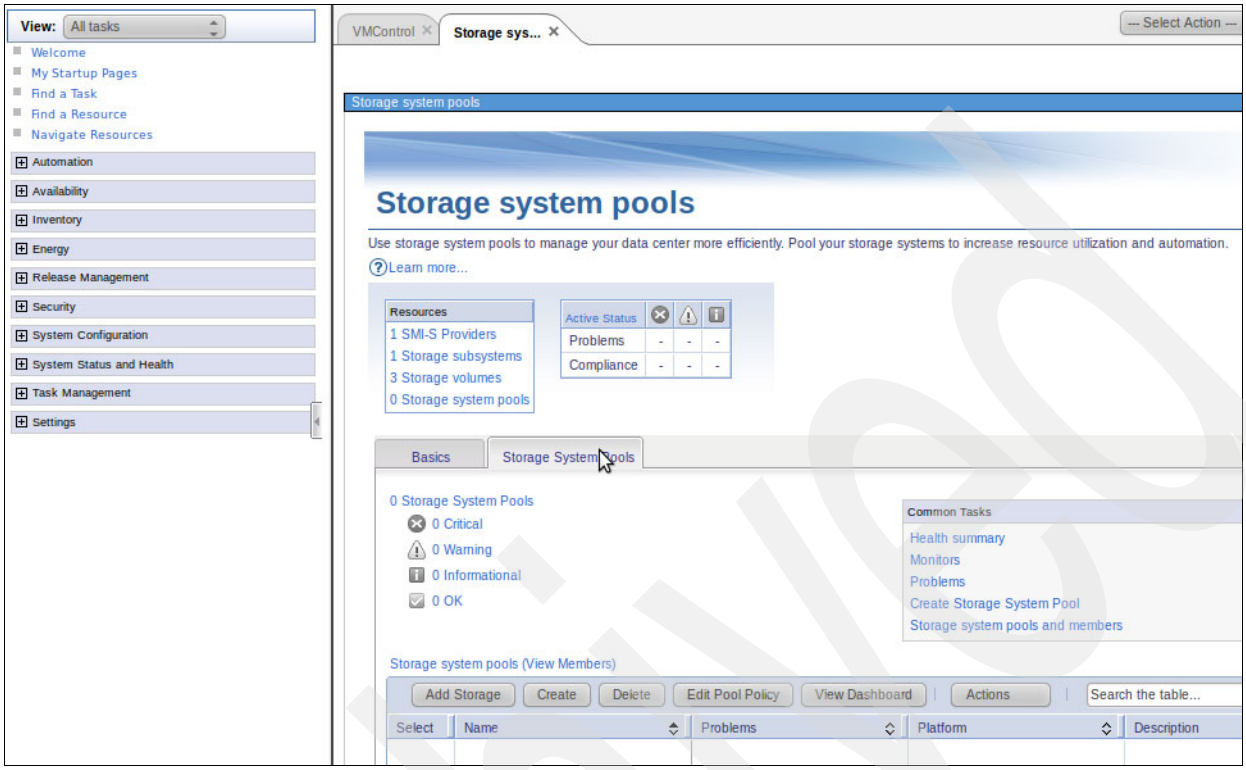

Figure 3-181 Click Create on the Storage system pools tab

The **Create** Storage System Pool wizard is started, as shown in Figure 3-182.

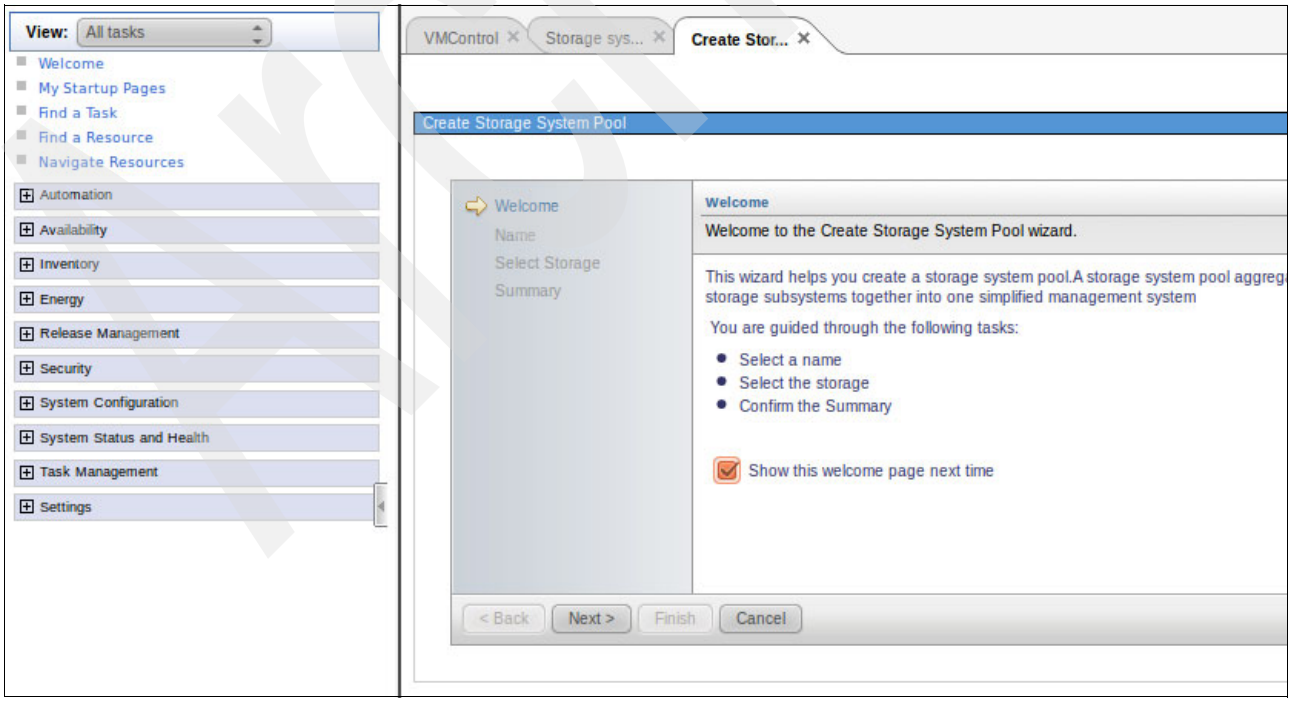

Figure 3-182 Storage System Pool wizard

An appropriate name for the pool to be created is provided, as shown in Figure 3-183. For this paper, we entered `PowerCloudStorage`. A description of this pool can also be entered.

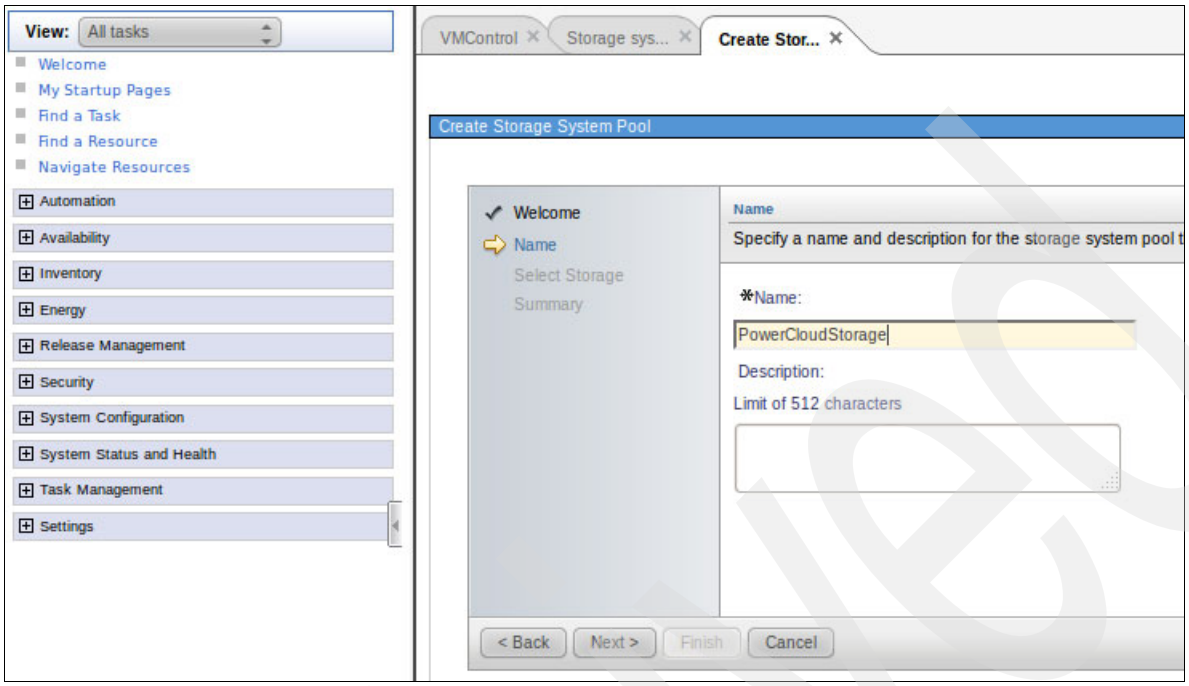

Figure 3-183 Specify an appropriate storage system pool name

In Figure 3-184, the information in the Available Storage Subsystems area comes from the Tivoli Storage Productivity Center environment that was discovered in 3.5.10, “Tivoli Storage Productivity Center discovery” on page 131. For this environment, the DS4800 is chosen from the **p14n06** subsystem, as shown in Figure 3-184.

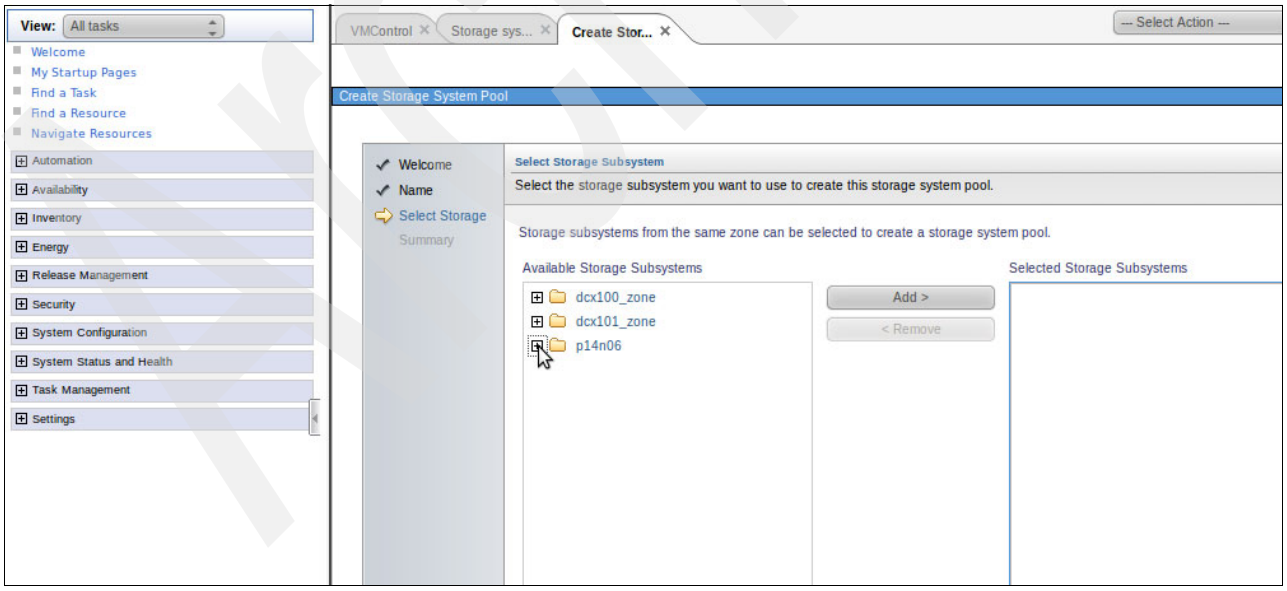

Figure 3-184 Information in the Available Storage Subsystems area comes from the Tivoli Storage Productivity Center discovered environment

In Figure 3-185, the storage subsystem is selected.

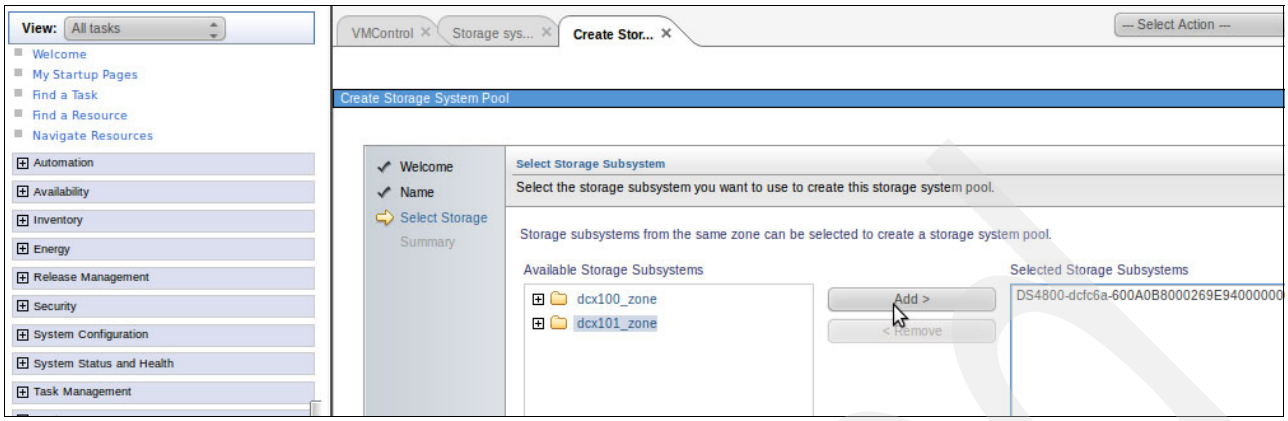

Figure 3-185 DS4800 is chosen from the p14n06 subsystem

A summary window is shown to verify the choices, as shown in Figure 3-186.

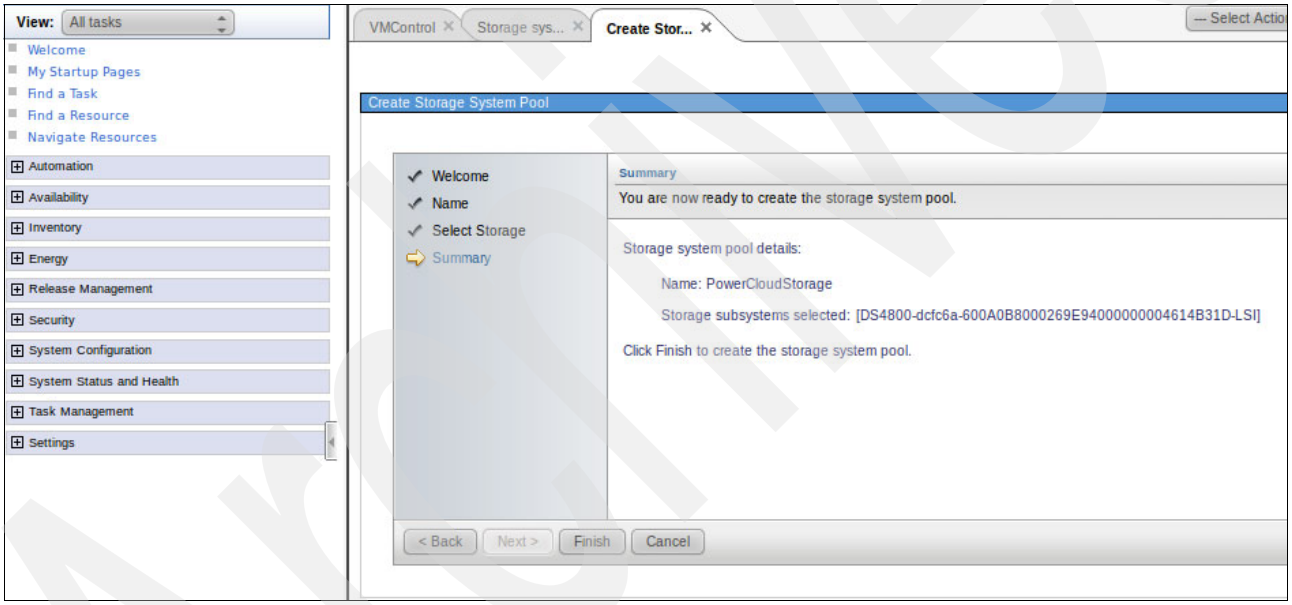

Figure 3-186 A summary window is shown to verify the choices

After the creation of the storage pool, the Storage system pools table is updated to reflect the change, as shown in Figure 3-187.

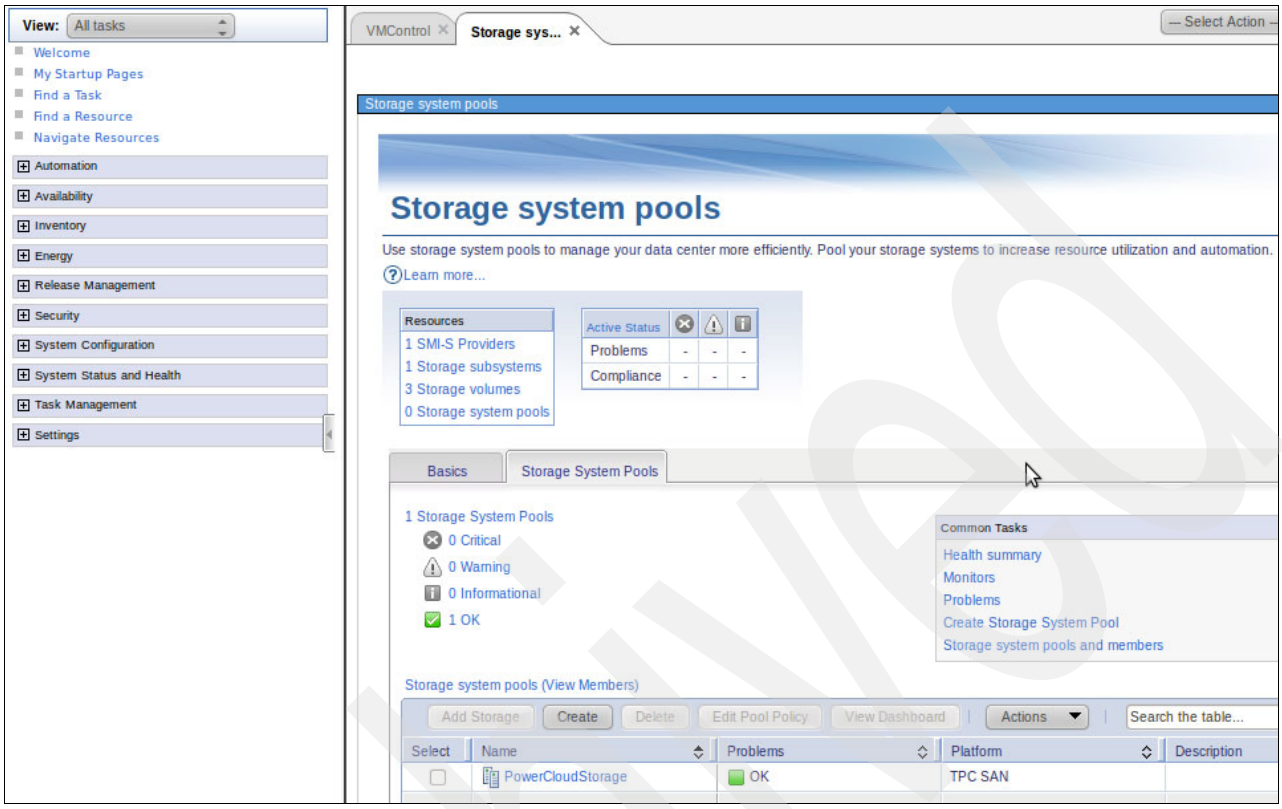

Figure 3-187 Storage system pools table

3.5.12 Server pool

Similar to the storage system pool, VMControl also provides the capability to pool virtual servers. When creating virtual servers, VMControl uses the information about the physical machines in the pool to determine placement. This pooling capability is an essential characteristic when considering cloud components.

From the main VMControl window, select the **Server System Pools** tab, as shown in Figure 3-188.

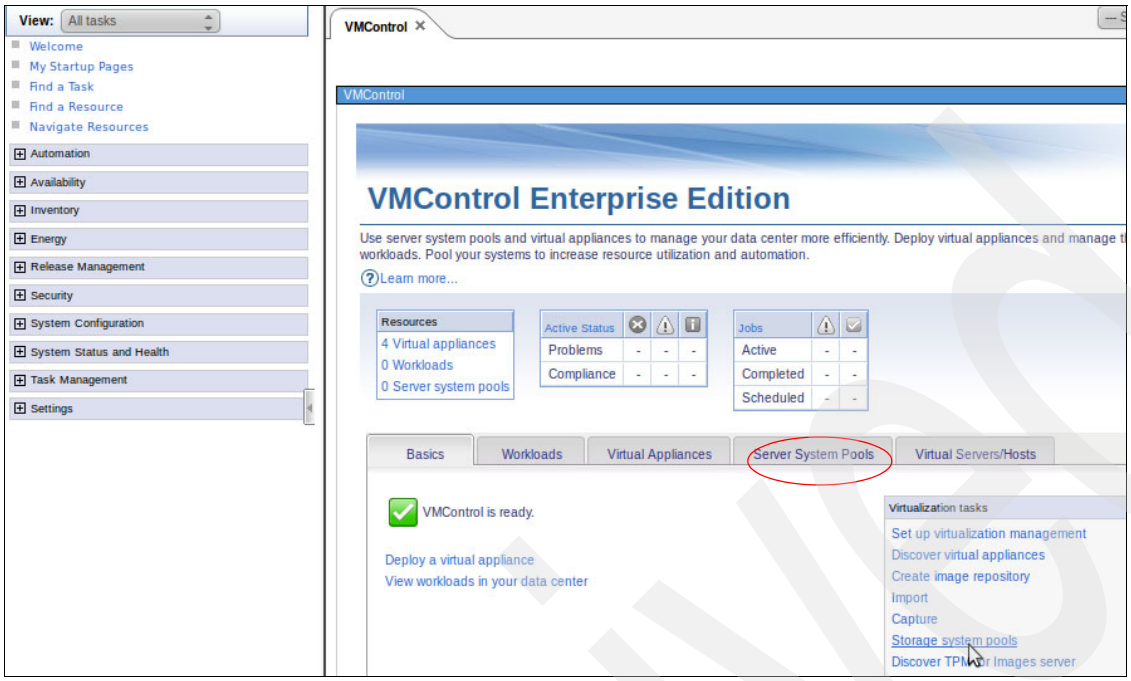

Figure 3-188 Server system pools tab

From this tab, select **Create** from the task bar on the Server Systems Pools table, as shown in Figure 3-189.

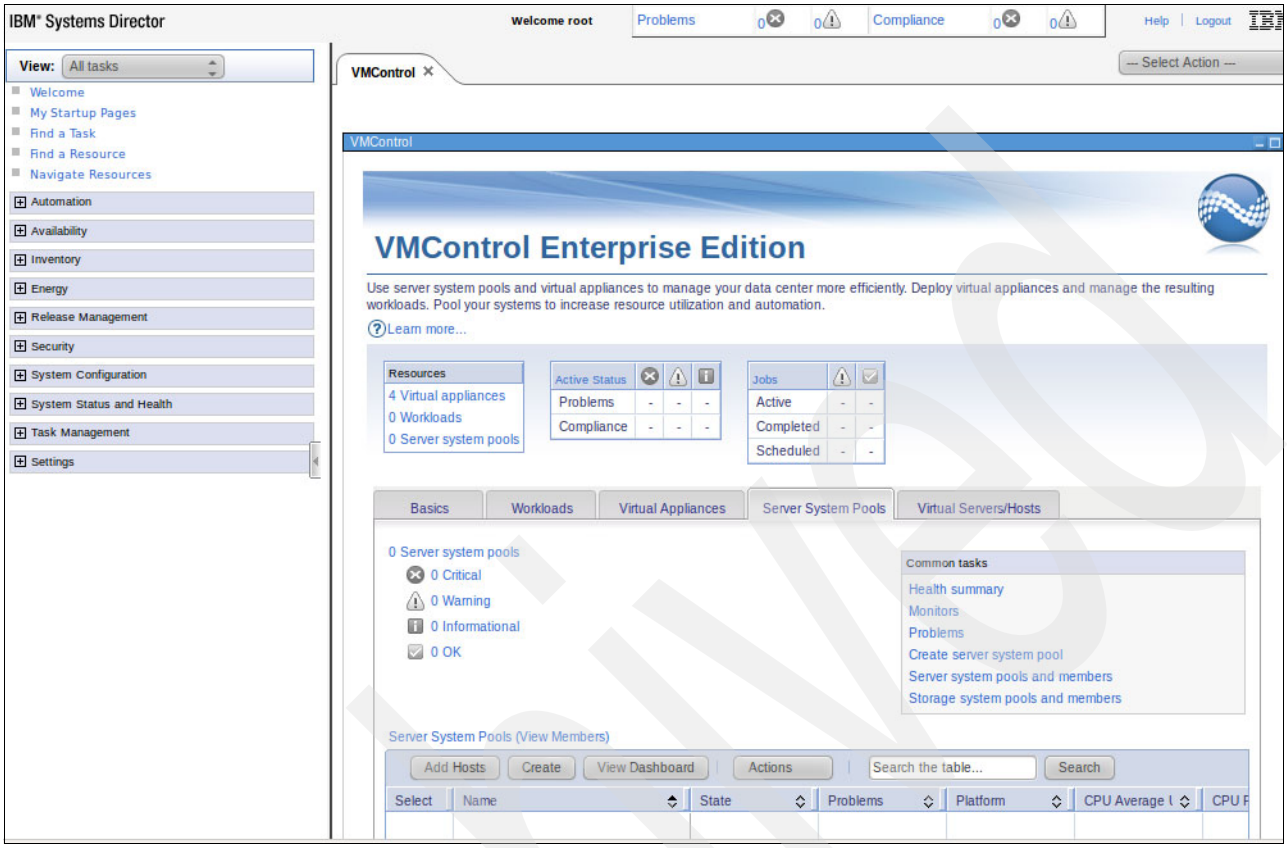

Figure 3-189 Click Create to start the Create Server System Pool wizard

The Create Server System Pool Wizard is started, as shown in Figure 3-190.

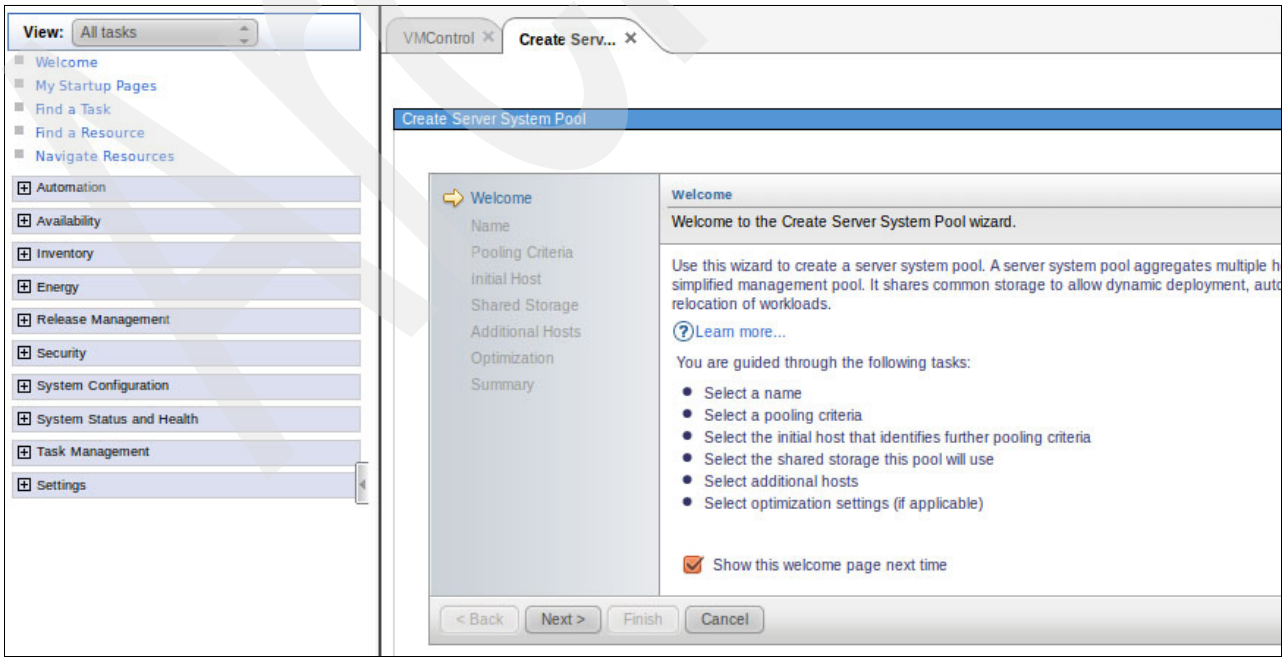

Figure 3-190 Create Server System Pool wizard welcome panel

A name and description can be entered for this server system pool. In Figure 3-191, we entered the name `Power Cloud Server Pool`.

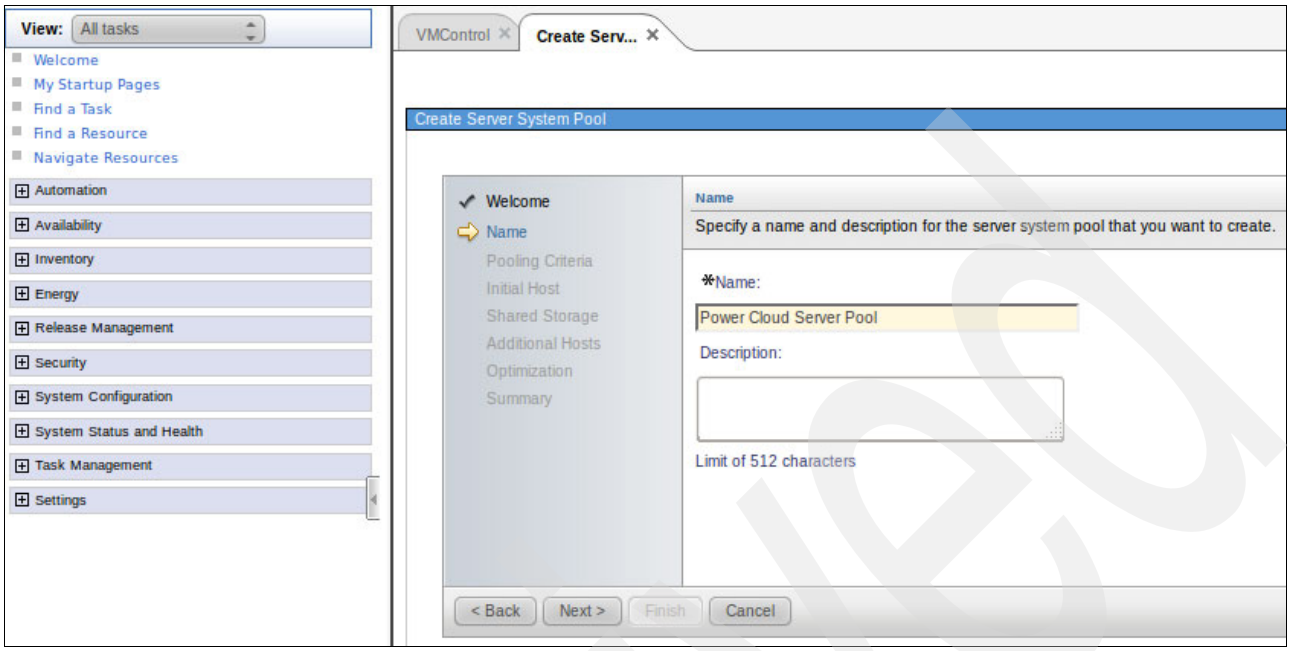

Figure 3-191 Specify name and description of the server system pool

High availability is not a main goal for this environment. Therefore, the check mark for the Resilience criteria is cleared (Figure 3-192).

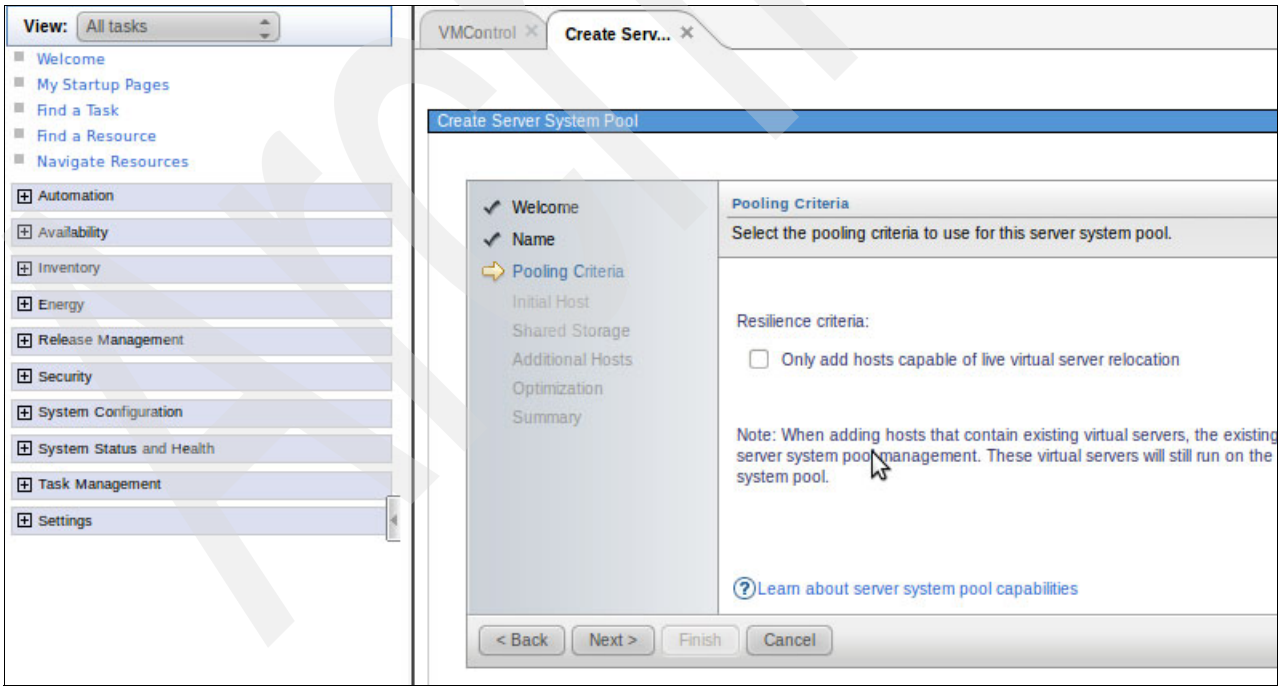

Figure 3-192 Resilience criteria check box is cleared

The wizard then asks for the initial physical host system to include in this pool. In this case, the Little Box is selected (Figure 3-193).

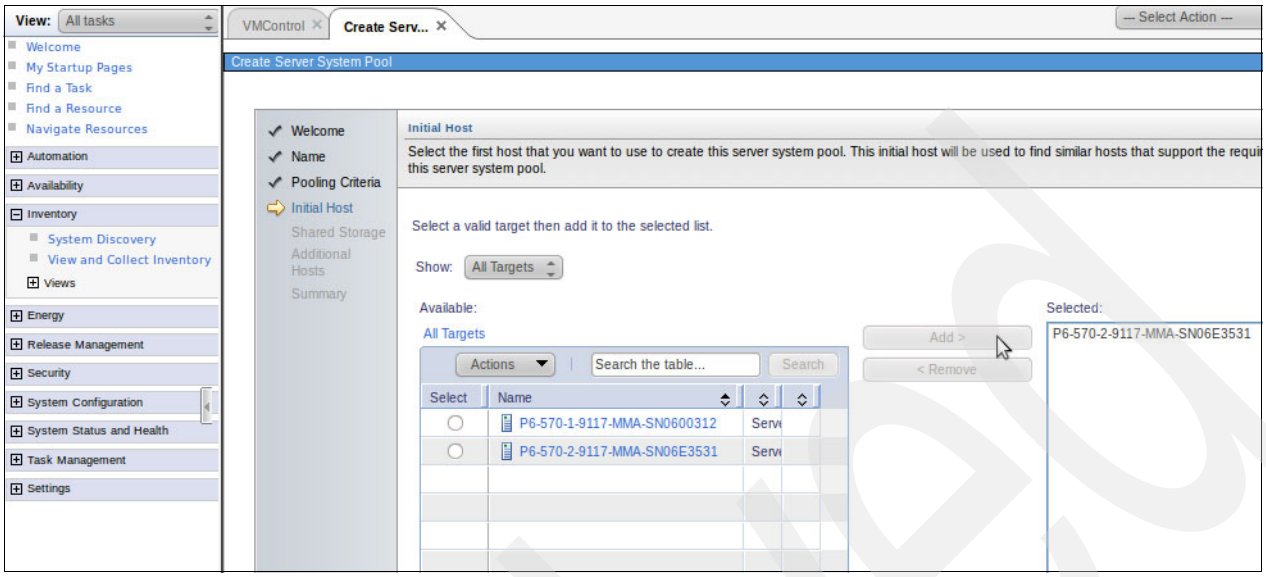

Figure 3-193 Choosing the Little Box for the initial physical system to be included in the pool

The shared storage that is used for this server system pool is the storage pool that was created in the 3.5.11, “Storage pool” on page 139, as shown in Figure 3-194.

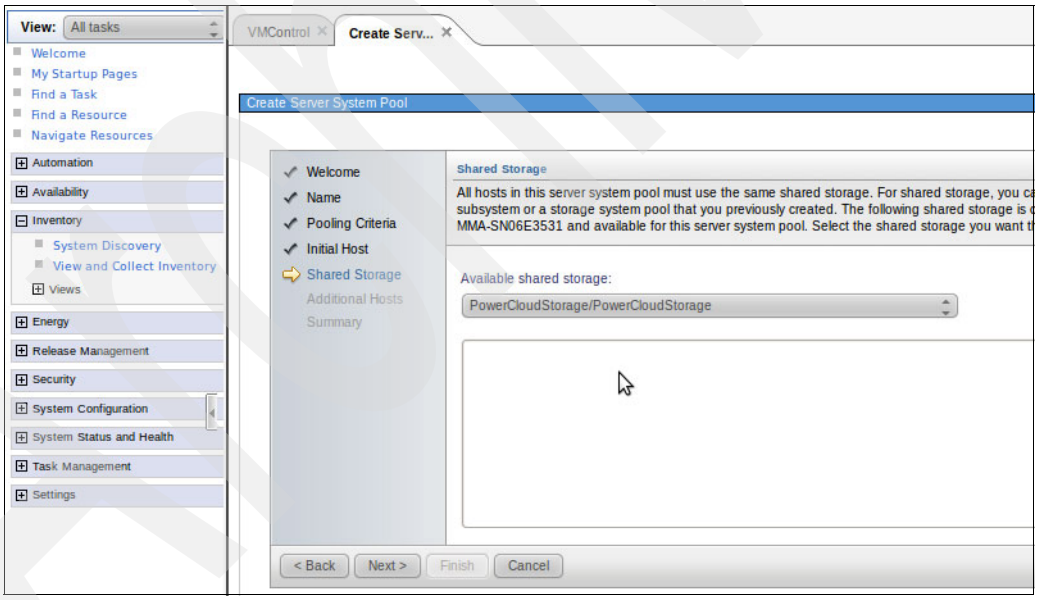

Figure 3-194 Create Server System Pool

Additional hosts can be added to the pool if required (Figure 3-195).

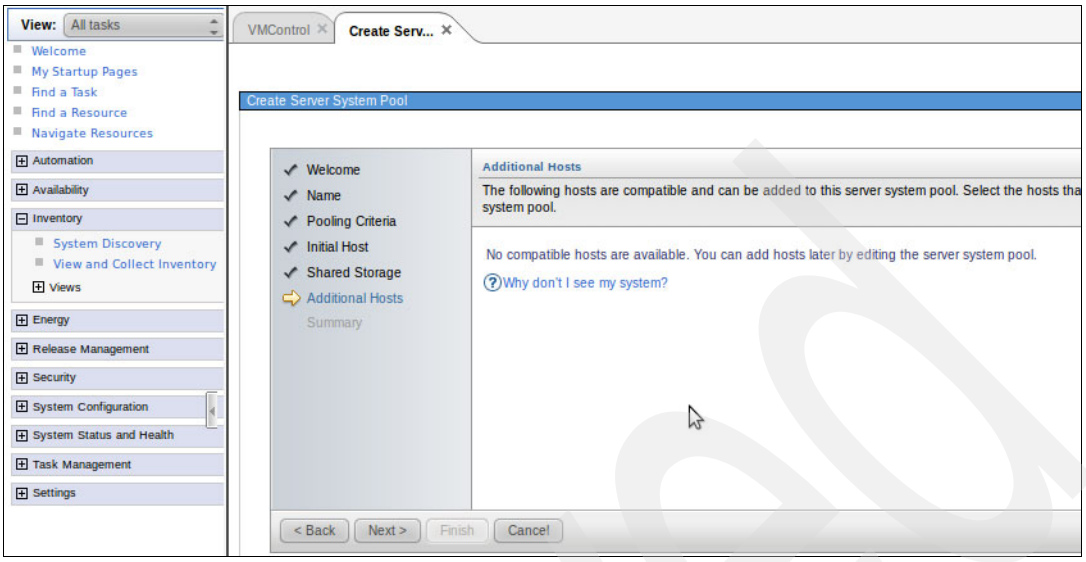

Figure 3-195 Additional hosts can be added to the server system pool if required

The summary window that is provided in Figure 3-196 shows which selections were made.

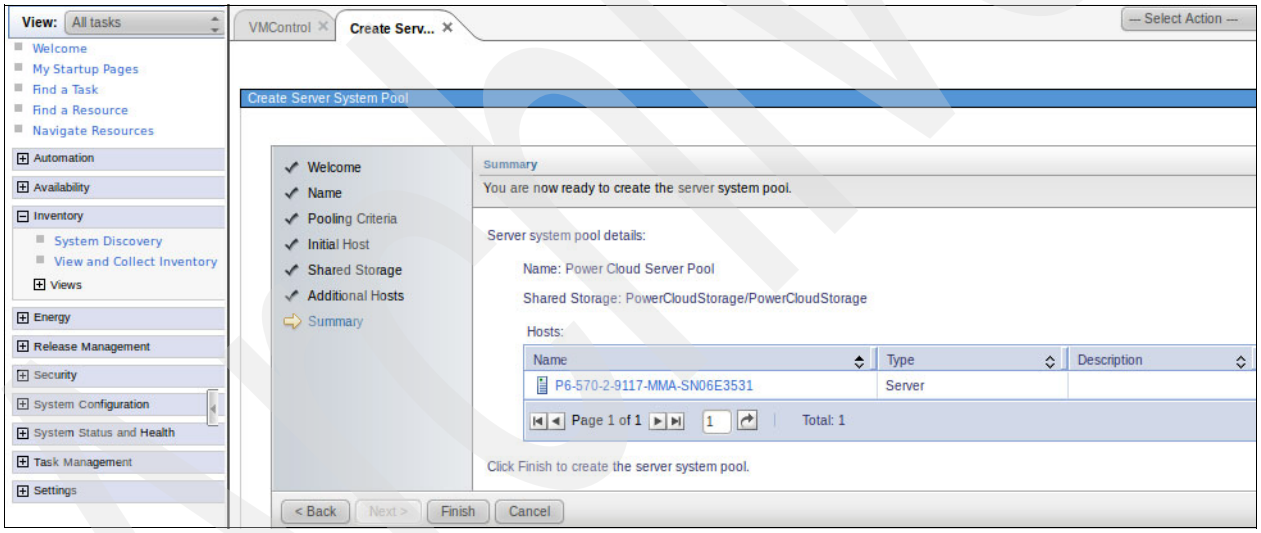

Figure 3-196 Server system pool details displayed as a summary

After the creation of the server system pool, the Server System Pools table is updated to reflect the change (Figure 3-197).

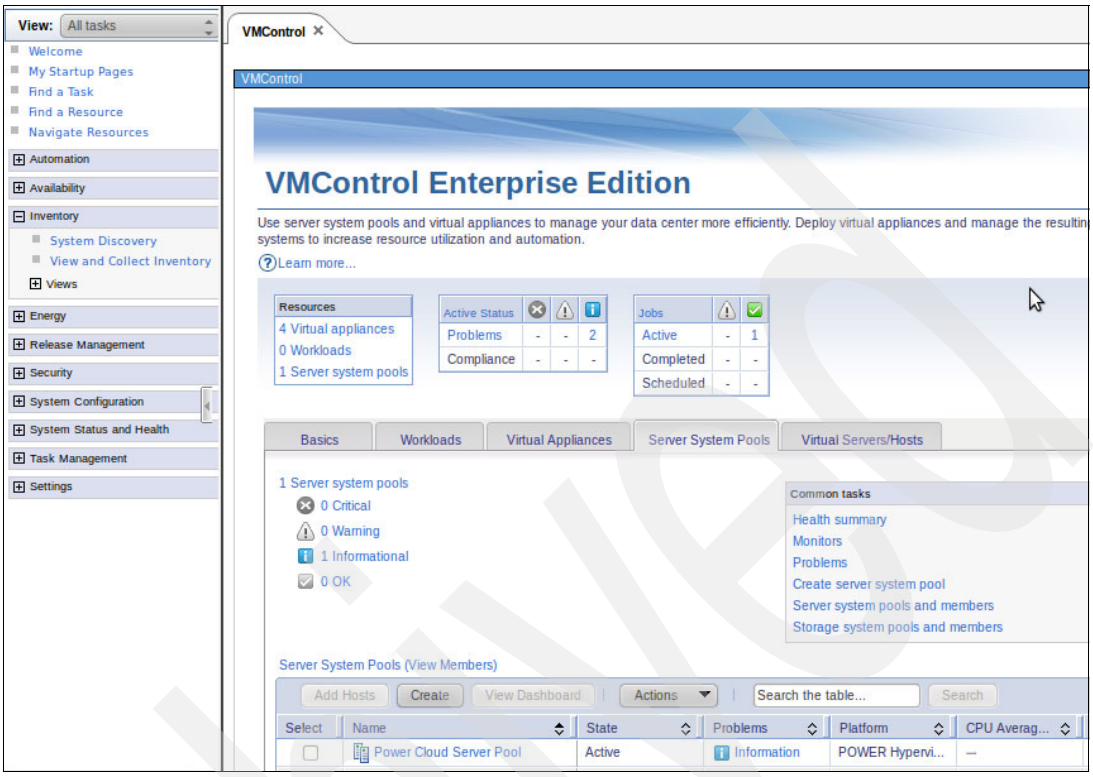

Figure 3-197 Server system pool has been created

4

Managing virtual machines

With IBM Systems Director installed, you have several of the features you typically look for in a cloud. You can capture running virtual machines and subsequently deploy them as virtual appliances. You can also export them for use in another environment or import images that you have captured in another environment. The remainder of this chapter provides an overview of this functionality. If you need a more thorough explanation of the capabilities, refer to the IBM Systems Director Information Center.

IBM Systems Director Information Center:

http://publib.boulder.ibm.com/infocenter/director/v6r2x/topic/com.ibm.director.editions.doc_6.2.1/editions_power_main.html

You can use these features to create reusable virtual appliances to simplify the support of your business processes.

4.1 Capturing a virtual machine

After you have deployed a virtual machine using IBM Systems Director and Network Installation Management (NIM), you can customize it with software stacks that are more useful in your infrastructure and properly support your business and processes. You can harness the power of IBM Systems Director to deploy the base operating system quickly. Other software packages can be deployed in a variety of ways: manual installation, deployment through NIM, and so on. Whatever method you use to configure your virtual machines, you can capture an installed system in order to quickly deploy other virtual machines with this configuration later.

To capture an installed machine, log in to your IBM Systems Director interface. Then, navigate to VMControl using the navigation links on the left side of the page by expanding **System Configuration** and clicking **VMControl**, as shown in Figure 4-1.

Figure 4-1 Expand System Configuration and click VMControl

Then, click **Capture** under the heading Virtualization tasks, as shown in Figure 4-2.

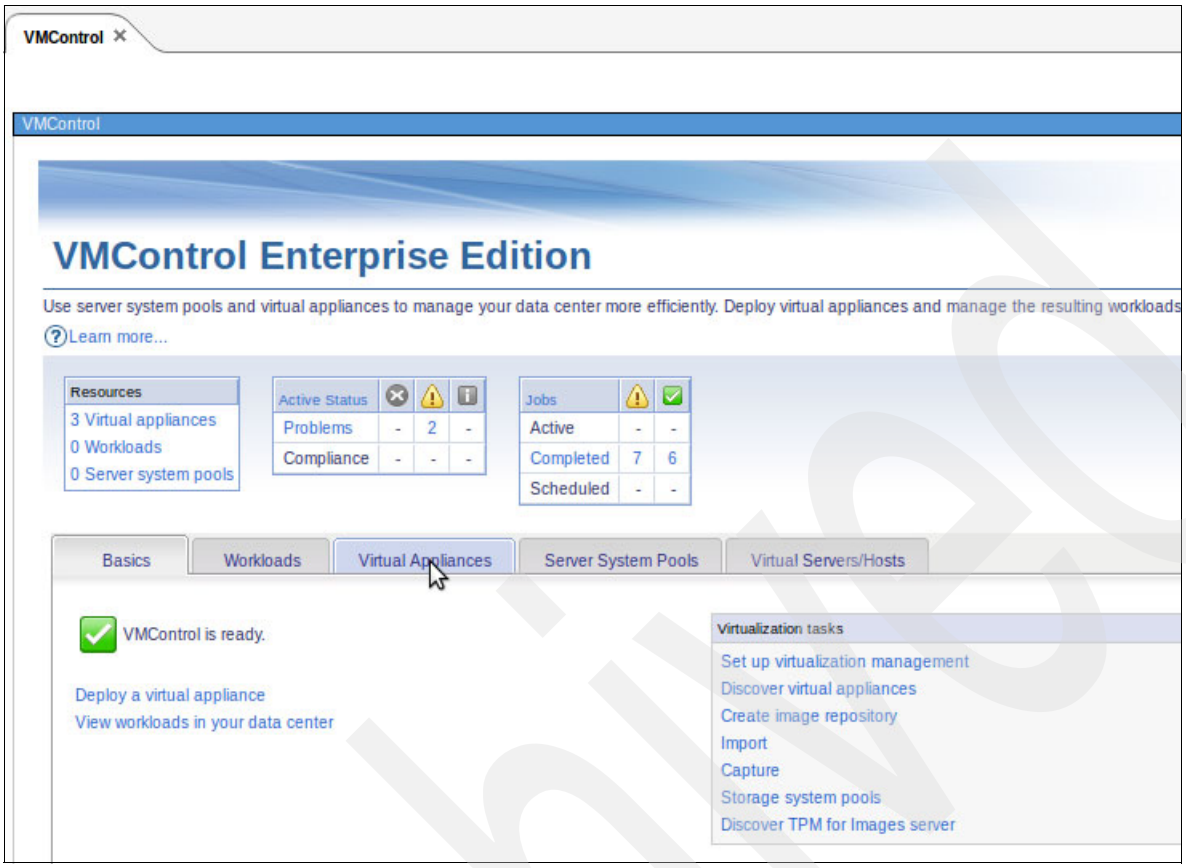

Figure 4-2 The Capture link listed under Virtualization tasks provides access to the Capture Wizard

Next, you see a window, as shown in Figure 4-3, with a message welcoming you to the Capture wizard. This wizard guides you through the process of capturing a virtual machine. Click **Next** to begin.

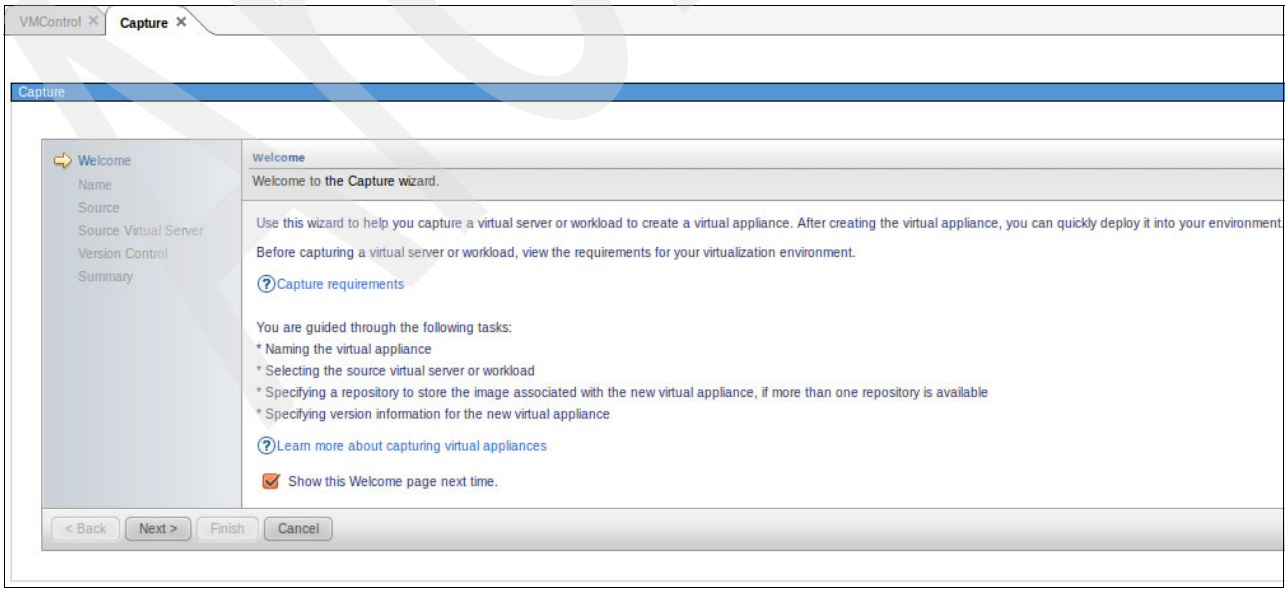

Figure 4-3 Welcome to the Capture wizard

As shown in Figure 4-4, you need to provide a machine name and brief description. Then, click **Next**.

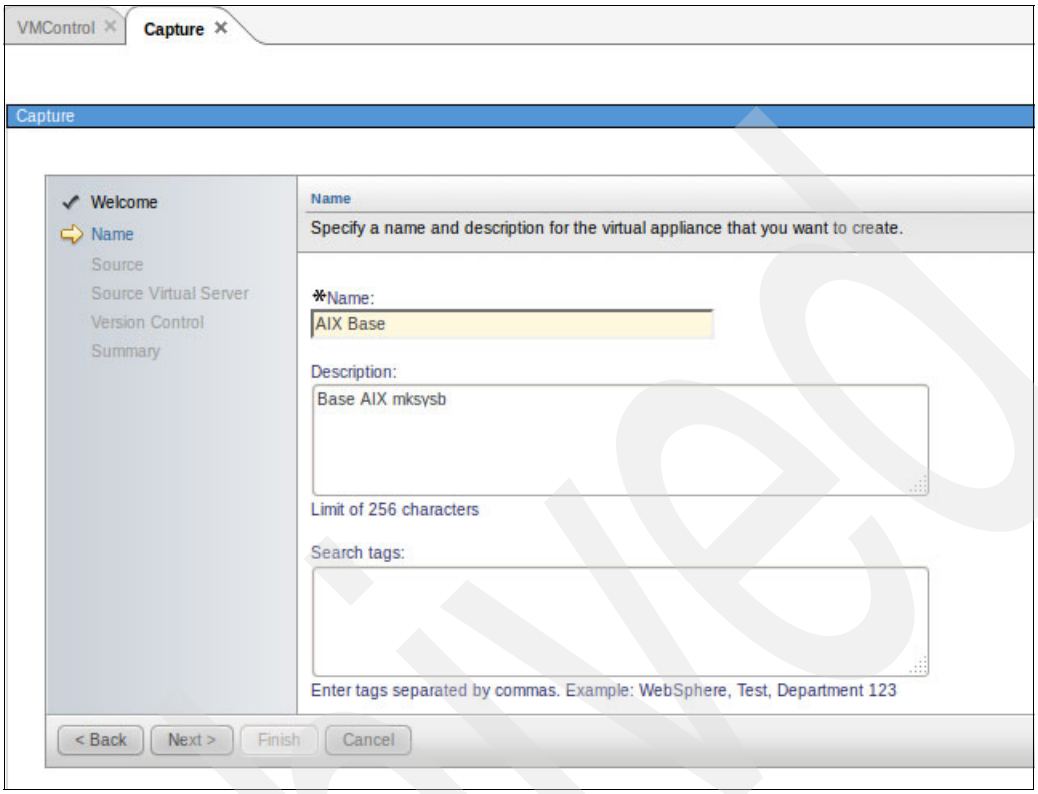

Figure 4-4 Specifying the name and description of the virtual device that you want to capture

Next, the type of entity being captured must be specified. In Figure 4-5, click **Virtual Server**. Click **Next** to continue.

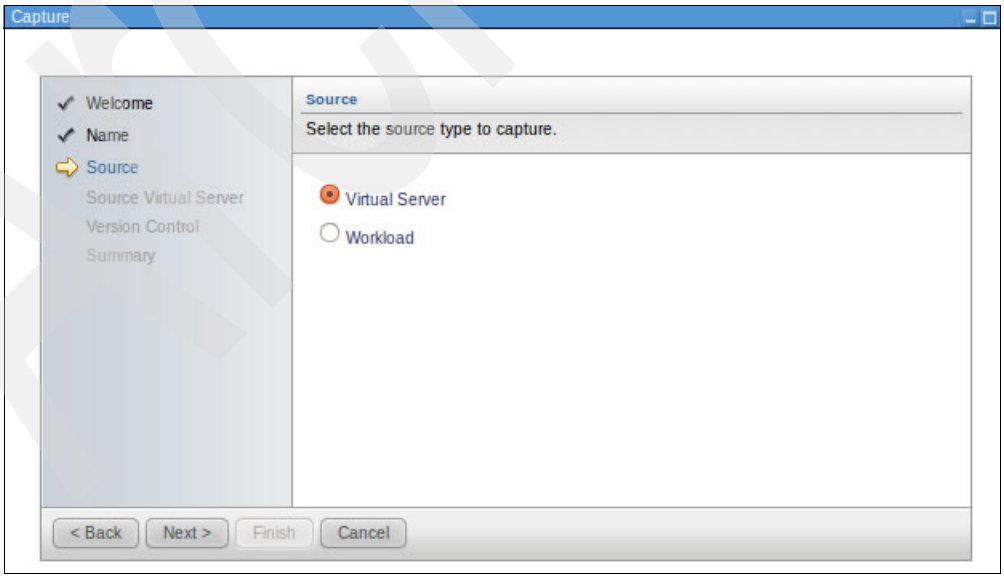

Figure 4-5 Tell the wizard what type of entity is being captured

At this point, as shown in Figure 4-6, a list of the virtual machines that have been already configured is presented. Choose a virtual machine and click **Add** to place it in the list.

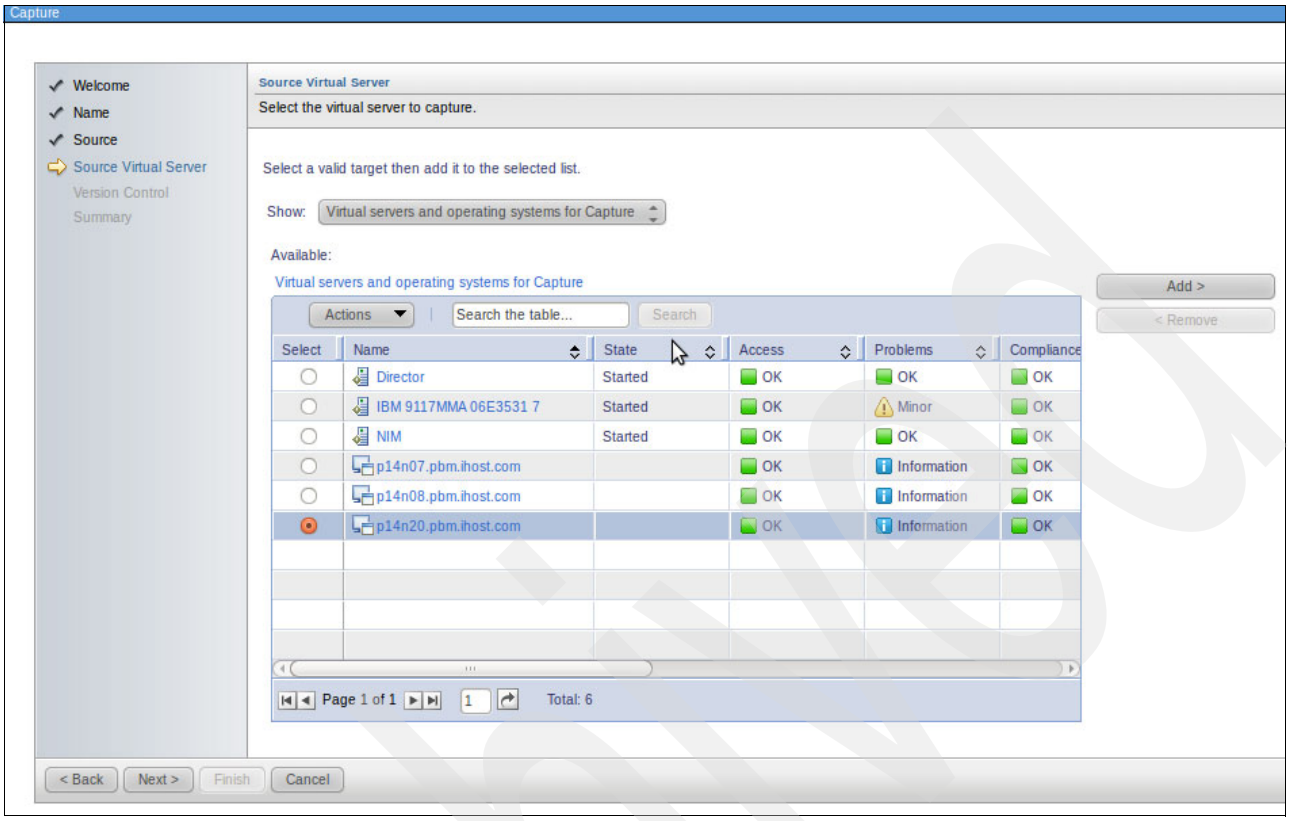

Figure 4-6 Tell the wizard which virtual servers to capture by clicking Add

Figure 4-7 shows the result of choosing **ADD**. The Selected portion of the panel will now contain the name of the server that was added.

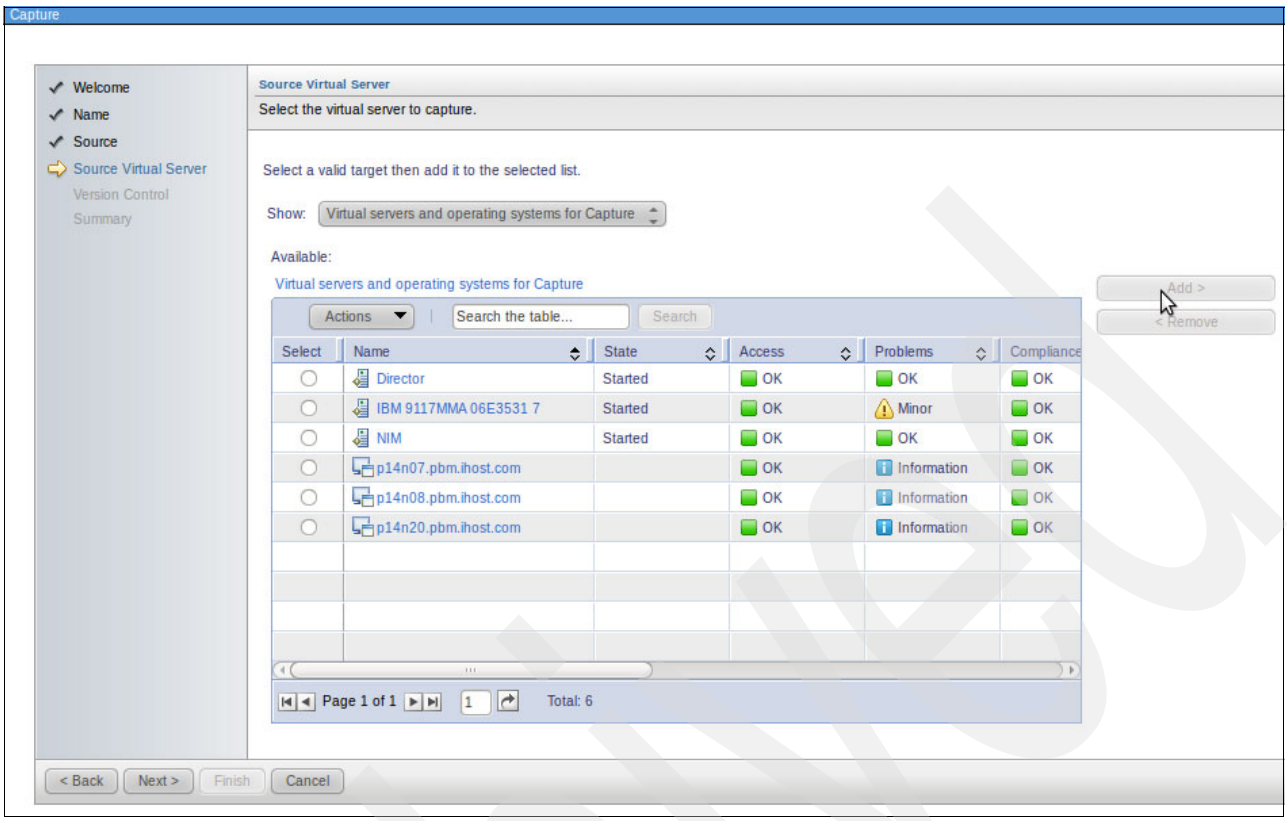

Figure 4-7 After clicking Add, the selected server is listed in the Selected column

After adding a name or names to the Selected column, Click **Next**.

Figure 4-8 shows the virtual networks to which the captured virtual machine is attached and gives you a chance to describe the virtual networks.

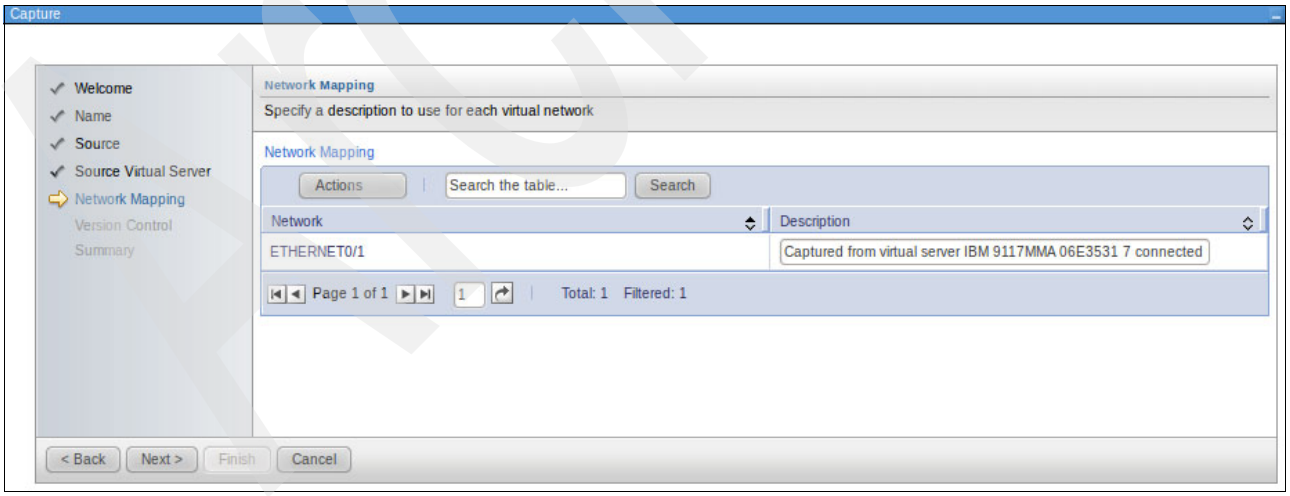

Figure 4-8 A description can be provided for each captured machine

Figure 4-9 shows how VMControl provides a mechanism for version control of the captured machines. For more information, see “Working with virtual appliance versions” in the information center.

Revisions to virtual appliances: Go to this website for a description of how IBM Systems Director VMControl manages revisions to virtual appliances:

`http://publib.boulder.ibm.com/infocenter/director/v6r2x/index.jsp?topic=%2Fcom.ibm.director.vim.helps.doc%2Ffsd0_vim_r_revisions.html`

In Figure 4-9, **Create a new version tree with the new virtual appliance as the root** was selected.

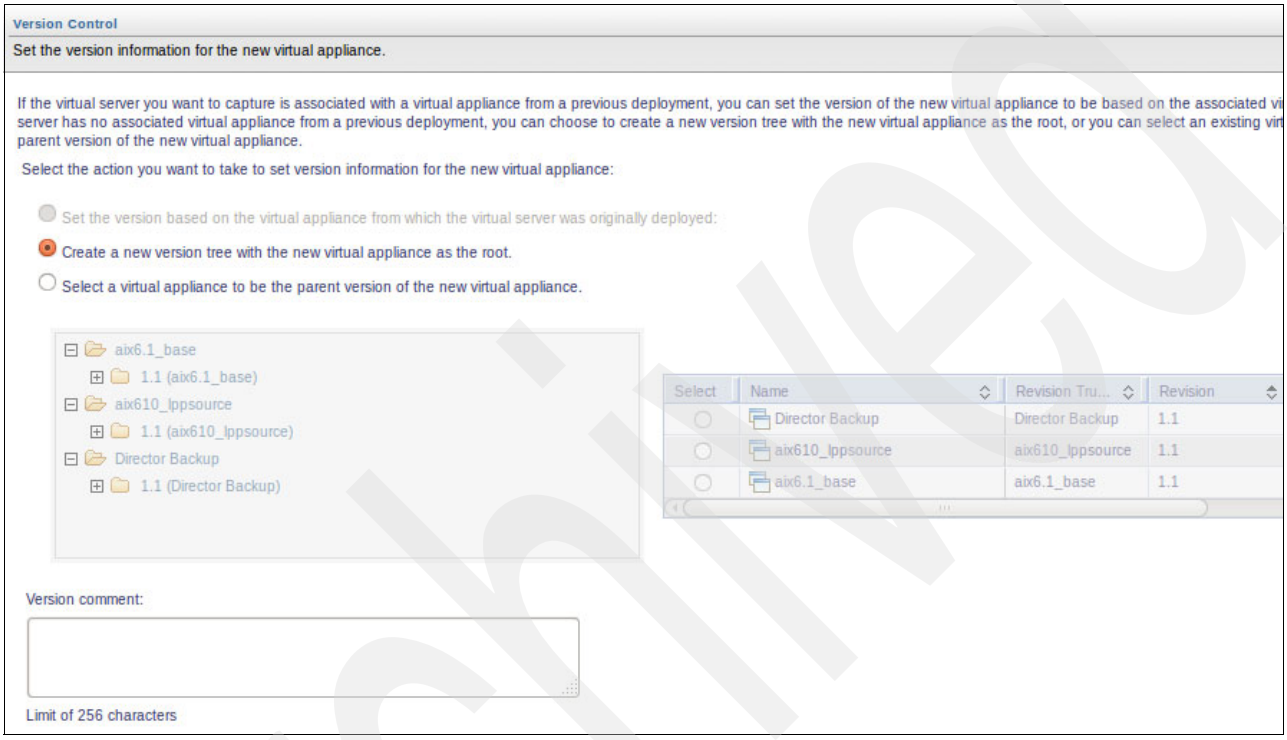

Figure 4-9 Create a new version tree with the new virtual appliance as the root

After reviewing the summary page that is shown in Figure 4-10, click **Finish**.

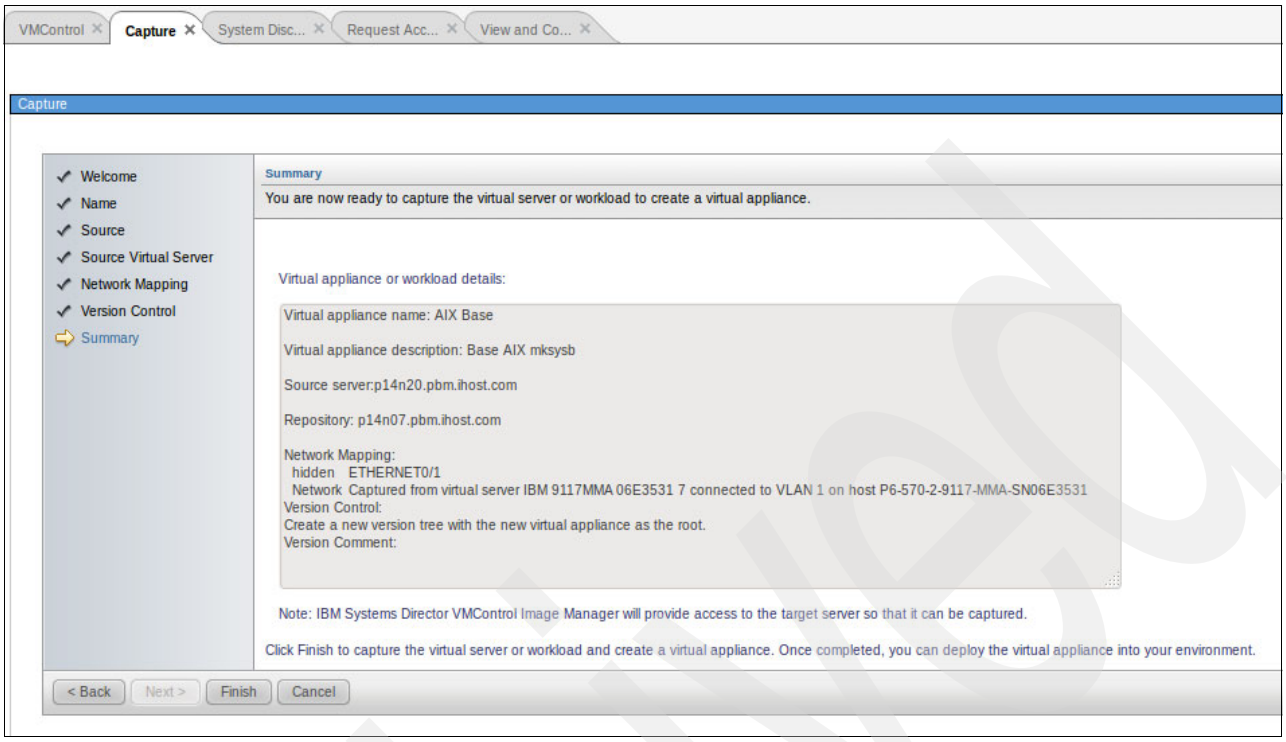

Figure 4-10 Summary of the specifications to capture a virtual server to create a virtual appliance

To submit the job immediately, click **Run Now**, as shown in Figure 4-11.

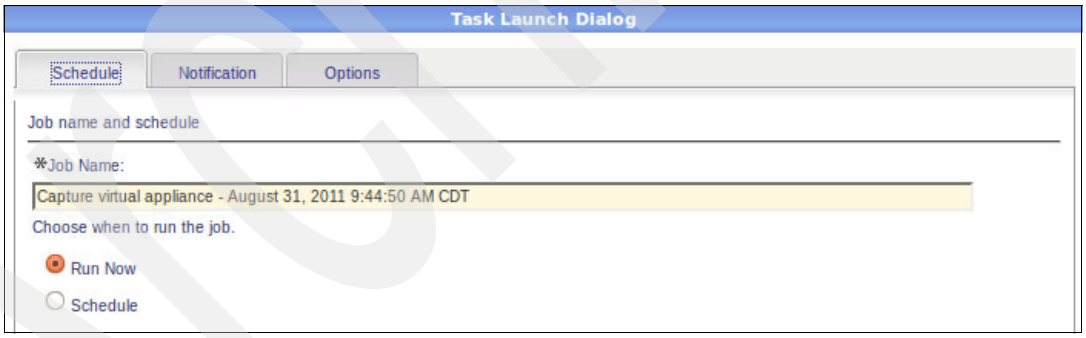

Figure 4-11 Run the Capture virtual appliance job now

After submitting the job, you can monitor its progress in the same way that you monitor any other job in the Systems Director, as shown in Figure 4-12. Both active and scheduled jobs, including the progress of the capture job, can be monitored, as shown in Figure 4-12.

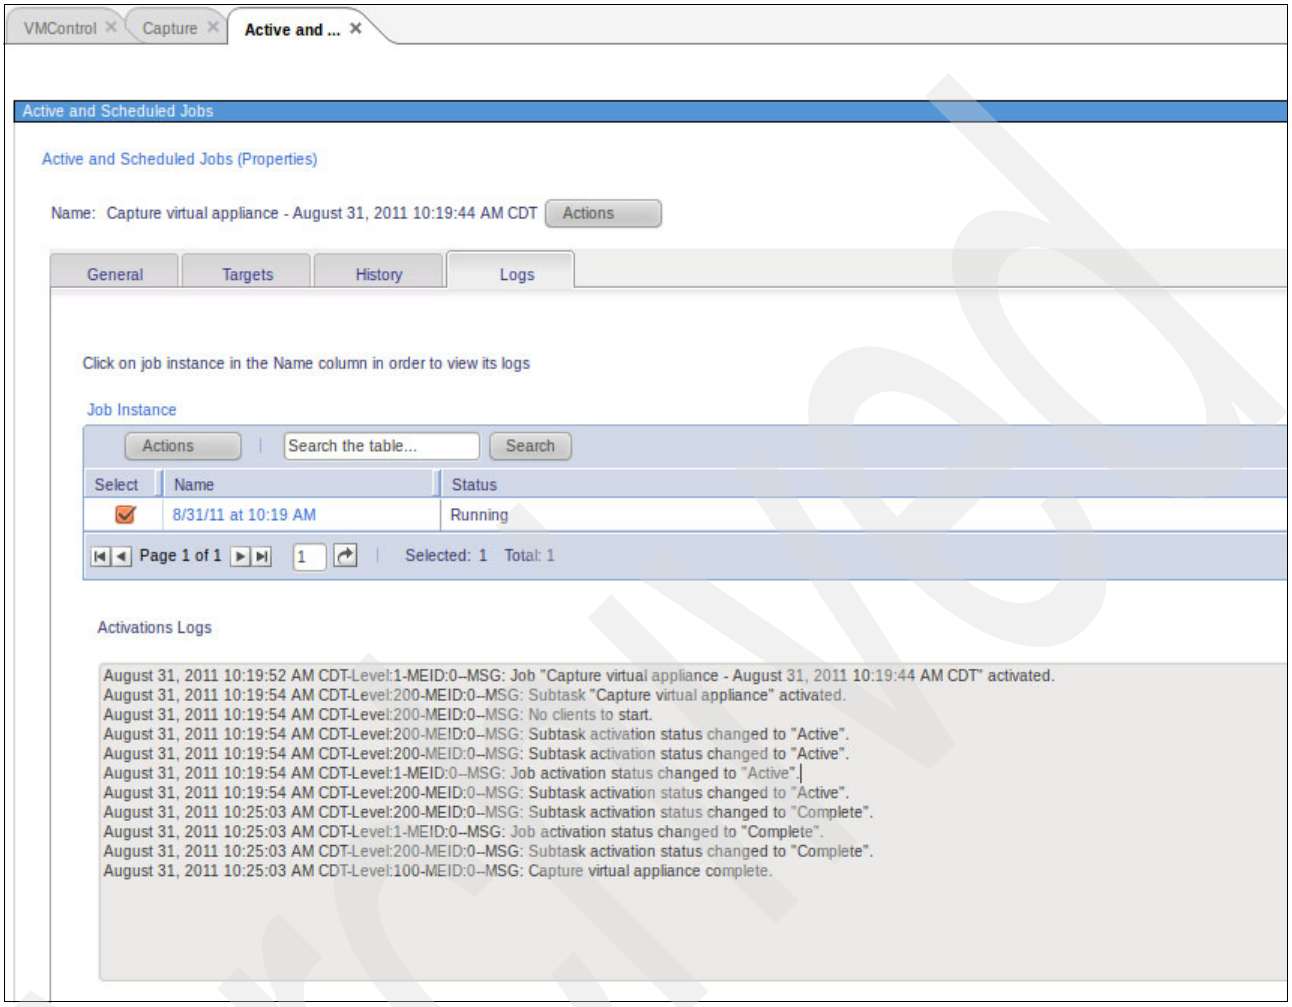

Figure 4-12 Monitor the progress of the job to see that the capture virtual appliance job is complete

4.2 Deploying a virtual machine

After completing the tasks of installing, configuring, and capturing a virtual appliance, you can deploy as many instances of it as you need.

To deploy instances of a captured virtual machine, return to the VMControl summary page and click **Deploy virtual appliance** (Figure 4-13 on page 162). In the wizard, select the storage pool that you want to use for the guest, as well as the network resource. As you can see in Figures 4-15 through 4-19, storage was assigned from our Power Cloud pool, and we attached the guest to the main virtual local area network (VLAN) that was configured earlier.

Specify a hostname and IP settings for your guest. Leave the NIM-specific settings blank. After verifying the information on the summary window, click **Finish** to submit the job.

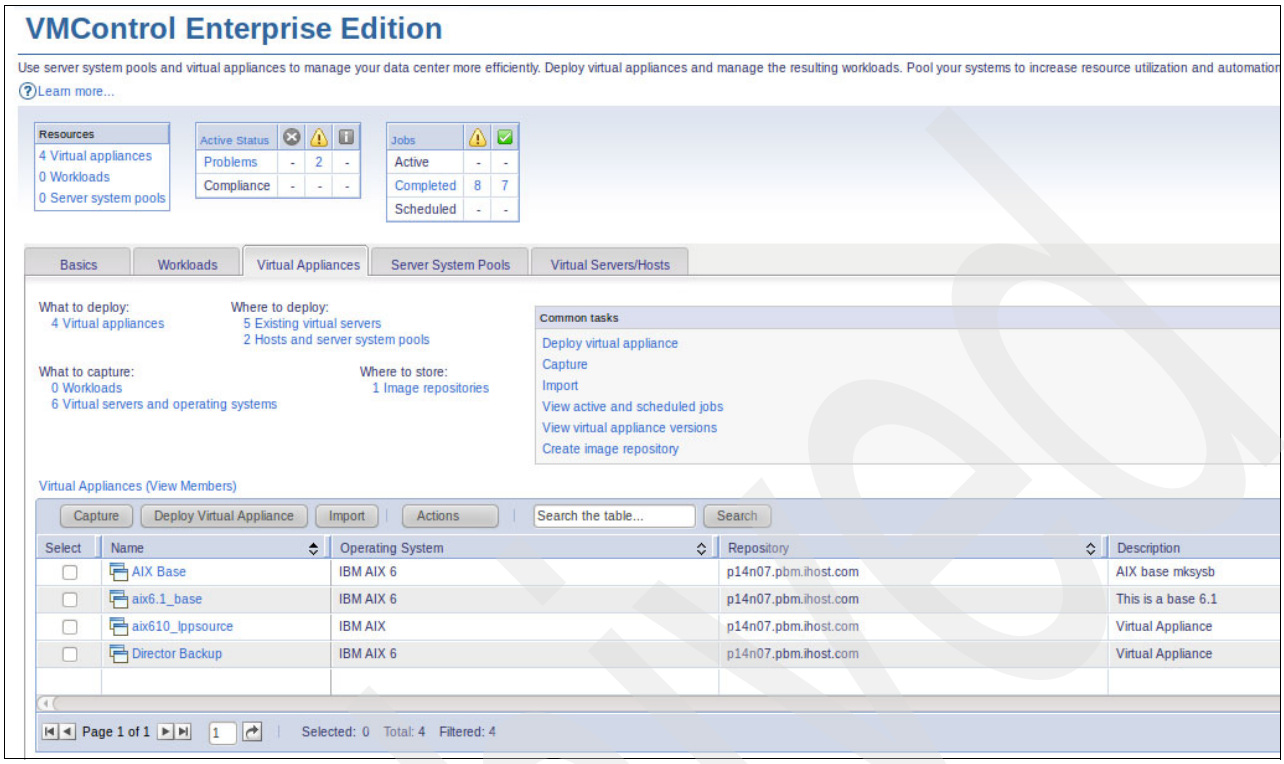

Figure 4-13 Common tasks include the Deploy virtual appliance job

Then, the Deploy Virtual Appliance wizard starts, as shown in Figure 4-14.

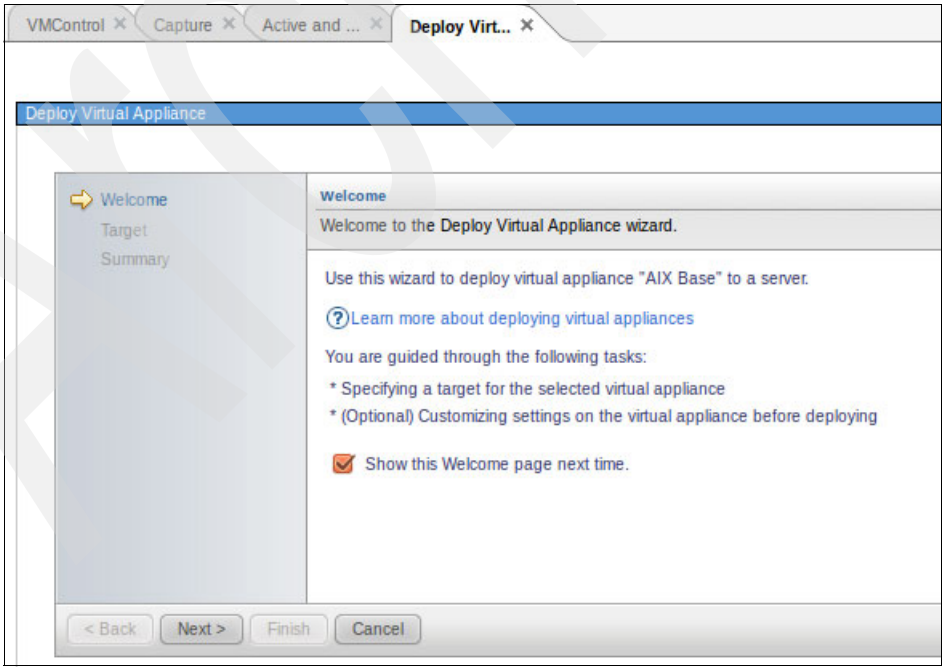

Figure 4-14 Deploy Virtual Appliance wizard

The first step in deploying a virtual appliance is to assign storage. VMControl creates a new disk to use for this appliance and assigns it to a storage pool or storage volume. See Figure 4-15.

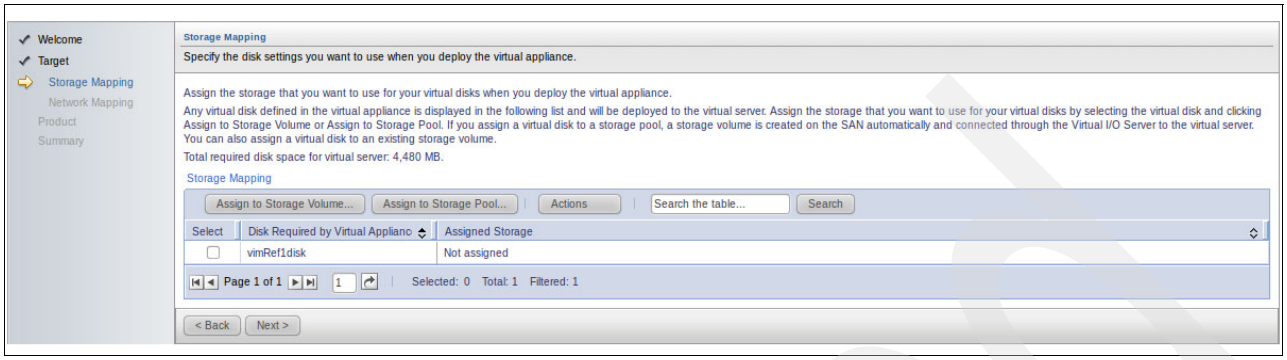

Figure 4-15 Specifying the disk settings for virtual appliance deployment

In 3.5.11, “Storage pool” on page 139, a storage pool was created for use by VMControl. This same pool is used to deploy this appliance (Figure 4-16).

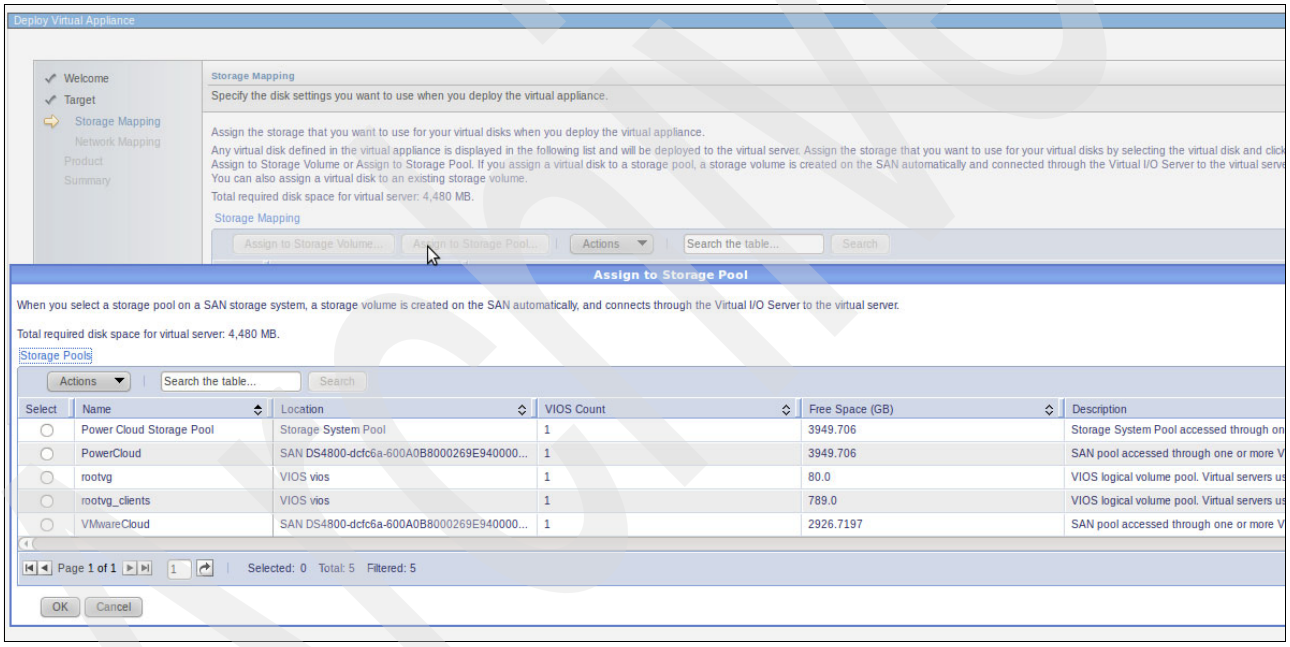

Figure 4-16 Assign to Storage Pool

In Figure 4-17, the panel is now updated with the assigned storage.

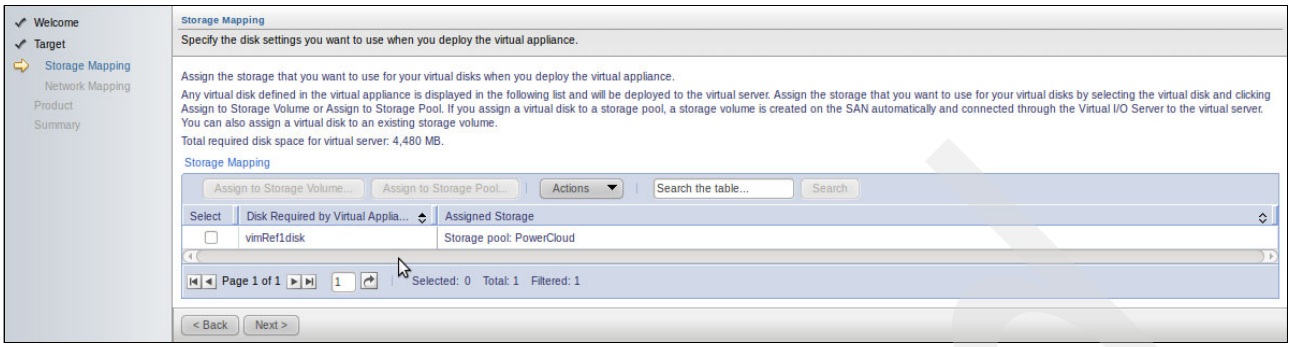

Figure 4-17 Assigned storage is updated

The next step is to assign a network to the virtual appliance. This environment has a simple network topology, so the default was chosen, as shown in Figure 4-18.

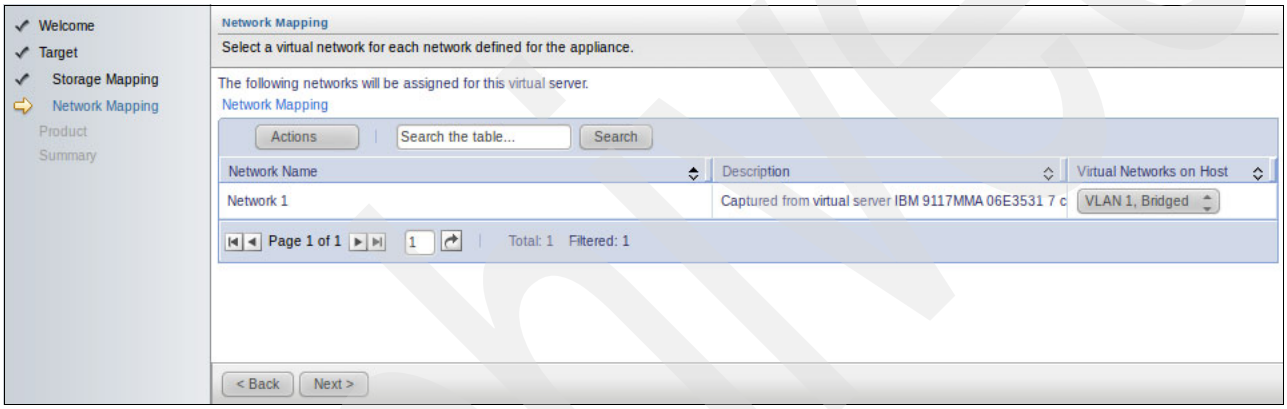

Figure 4-18 Assign a network to the virtual appliance

The last panel that requires answers allows you to specify the unique characteristics of this virtual appliance. You enter data, such as special NIM settings, host name, and IP address, as shown in Figure 4-19.

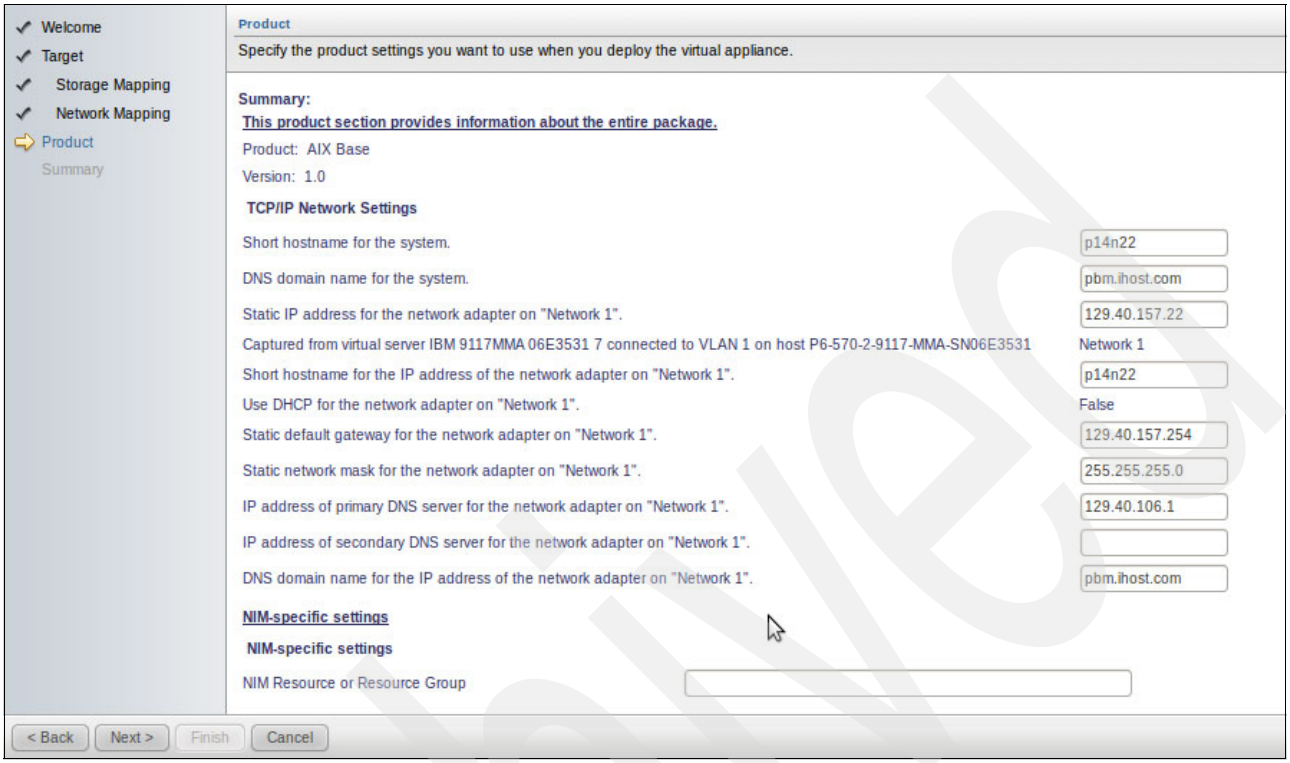

Figure 4-19 Specifying the unique characteristics of the virtual appliance

The summary in Figure 4-20 shows the selected options for this virtual appliance.

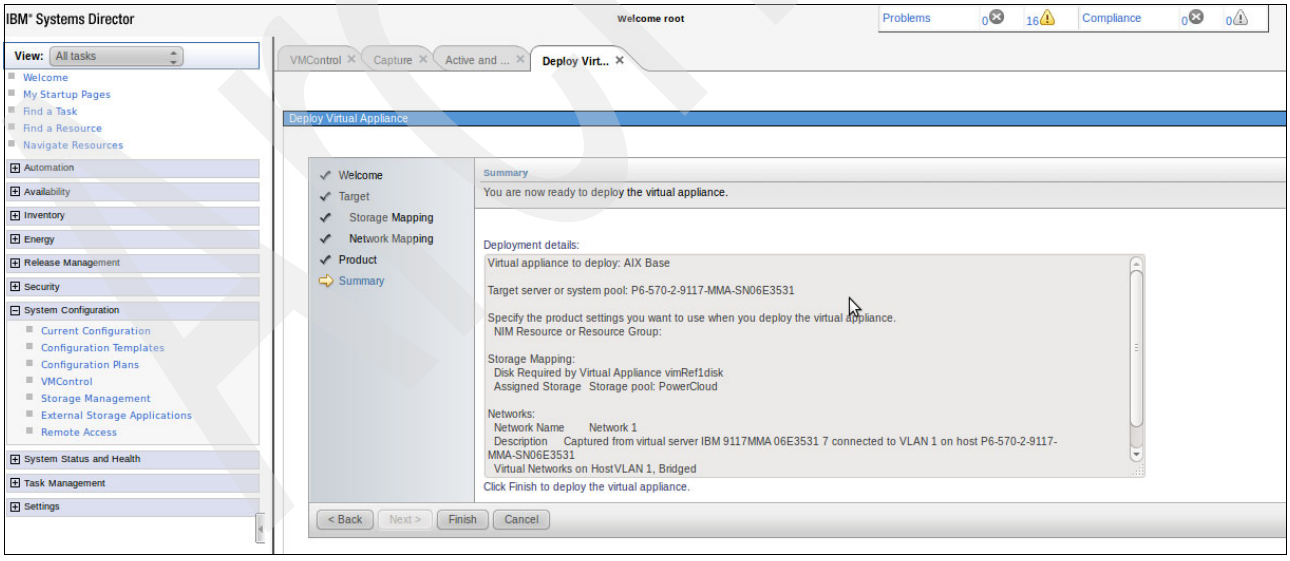

Figure 4-20 Summary of the selected options for this virtual appliance

A new job is created for this action. Because this job might run for a while, you can schedule it for a more convenient time, as shown in Figure 4-21.

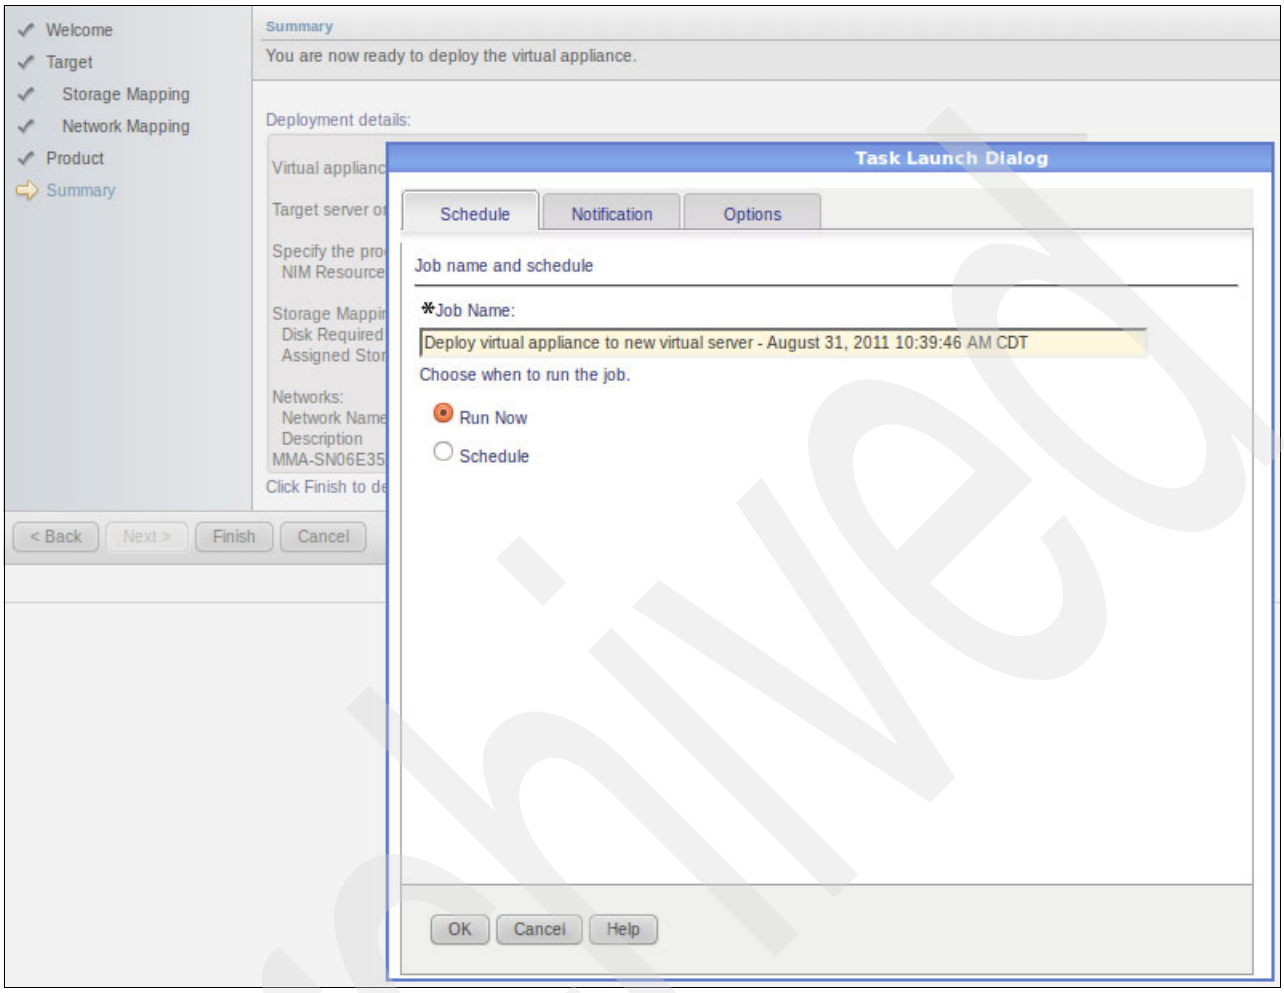

Figure 4-21 Use the Task Launch Dialog to run the deployment now or schedule the deployment for a later time

The progress window (Figure 4-22) shows how much of the deployment has occurred.

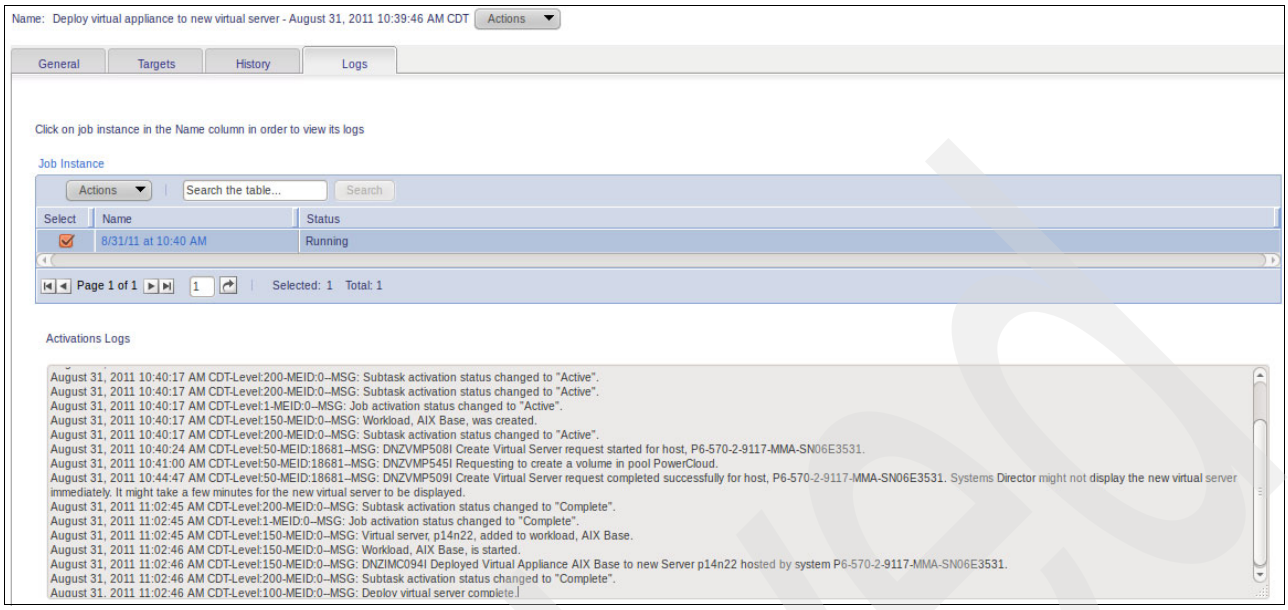

Figure 4-22 Monitoring the progress of deployment

Now the appliance is deployed successfully, which is confirmed by the message at the bottom of Figure 4-22.

4.3 Importing a virtual machine

With your environment configured and virtual machines deployed, you might want to share virtual machine images between multiple environments. You can also have a Hypervisor Edition of a software package.

If you have obtained an Open Virtualization Format (OVF) or Open Virtual Archive (OVA) for a virtual appliance, for example by exporting it from another environment or by purchasing a Hypervisor Edition of a software product, you can easily import it for use in your POWER setup. To import a virtual appliance, click **Import** on the VMControl summary page. Specify the location of the ovf or ova file corresponding to the appliance that you want to import and click **Next**. Give your appliance a name and specify the repository into which you want to import it. Select **Create a new version tree with the new virtual appliance as the root**. View the summary and click **Finish**.

4.4 Exporting a virtual machine

Exporting a virtual appliance is slightly more complicated than importing. You need to collect a set of resource definitions and identifiers into a form that facilitates the eventual task of exporting the virtual machine.

From the Virtual appliances tab in the VMControl summary window, select the appliance that you want to export. In the Actions drop-down list, select **Related Resources** → **Software Image**. The beginning part of the software image name is the identifier that you need to find the ovf file that is associated with this image.

Identifier: For example, if your software image is named `8ca52cd3-8aef-408d-ac27-668a3499d01c.vimRef1`, you will need the ID `8ca52cd3-8aef-408d-ac27-668a3499d01c`.

On the NIM server, this ID helps you locate the files that are related to the software image. If your repository is located in the `/export/nim` directory, `/export/nim/appliances/<ID>` contains an ovf and **`mksysb`** for your appliance. To export an appliance, you need to copy these files and any other XML files in the directory to the location where you want to store the exported appliance.

Modification: As of this writing, the ovf file requires a minor modification that is described in the information center:

`http://publib.boulder.ibm.com/infocenter/director/v6r2x/topic/com.ibm.director.vim.helps.doc/fsd0_vim_t_exporting_virtual_appliance_parent.html`

Refer to the information center for more information.

5

Epilog

This document describes the components that can be used to create an infrastructure as a service-based cloud. The components might already be familiar to you. Generally, they are not new, but are being used in a new way to provide ease of deployment. These components allow for the new service and deployment model that is associated with cloud computing. There are many ways to augment the components that are covered in this book.

IBM offers an advanced product called Tivoli Service Automation Manager:

http://www-01.ibm.com/software/tivoli/products/service-auto-mgr/

Tivoli Service Automation Manager provides another layer of abstraction in addition to the layer that IBM Systems Director's VMControl plug-in offers. These enhancements bring more cloud characteristics to the environment. Self-Service is already included with VMControl, but Tivoli Service Automation Manager provides a simple way to use a web-based interface that allows cloud service requests an easier way to request infrastructure.

There is also proper service management that can be integrated into an approval-based service flow. Tivoli Service Automation Manager provides ready-to-use integration with Tivoli Usage Account Manager. This integration provides extremely complex metering and charge-back capabilities. IBM Tivoli Monitoring and Tivoli Service Automation Manager can be configured so that a monitoring agent is installed along with the infrastructure that is provisioned.

A

Problems and solutions

Table A-1 and Table A-2 on page 172 describe a set of problems that we encountered while implementing and deploying Cloud Components on POWER. We also provide the corresponding solutions.

Table A-1 TPC Installation message references to the /etc/hosts file

Message or symptom	The TPC install returned a message ExitCode= 0; however, we were unable to start the agents. Further investigation determined that the error log contained this message: “IC65746: TPC AGENT INSTALLER FOR AIX VIO SERVERS MUST CHECK /ETC/HOST FOR VALID LOCALHOST ENTRY”
Cause	We paraphrase the associated error description that was online at this website: https://www-304.ibm.com/support/docview.wss?uid=swg1IC65746 “tpc agent installation on aix vio servers will fail if there is not a valid localhost entry in the /etc/hosts file.”
Solution	Inspect and edit the /etc/hosts file to confirm that there is a valid localhost entry or add a valid host entry prior to running the install. A valid localhost entry has the following format: *yourlocalhost_ipaddr* *yourfullyqualified_localhost_domainname* *short_hostname*

Table A-2 CIM providers are not working as expected

Message or symptom	CIM providers not working
Cause	Wrong or back-level provider
Solution	Refer to and follow the instructions that are provided in Appendix B, “Using IBM Systems Director Update Manager” on page 173.

B

Using IBM Systems Director Update Manager

B.1 Using IBM Systems Director Update Manager

IBM Systems Director is capable of automatically downloading and installing updates and fixes to the core Director Server components, as well as to the installed plug-ins. After logging into Systems Director, use the navigation panel to go to the Release Management Section and select **Updates**. The Update Manager is displayed. Click **Update IBM Systems Director**, as shown in Figure B-1.

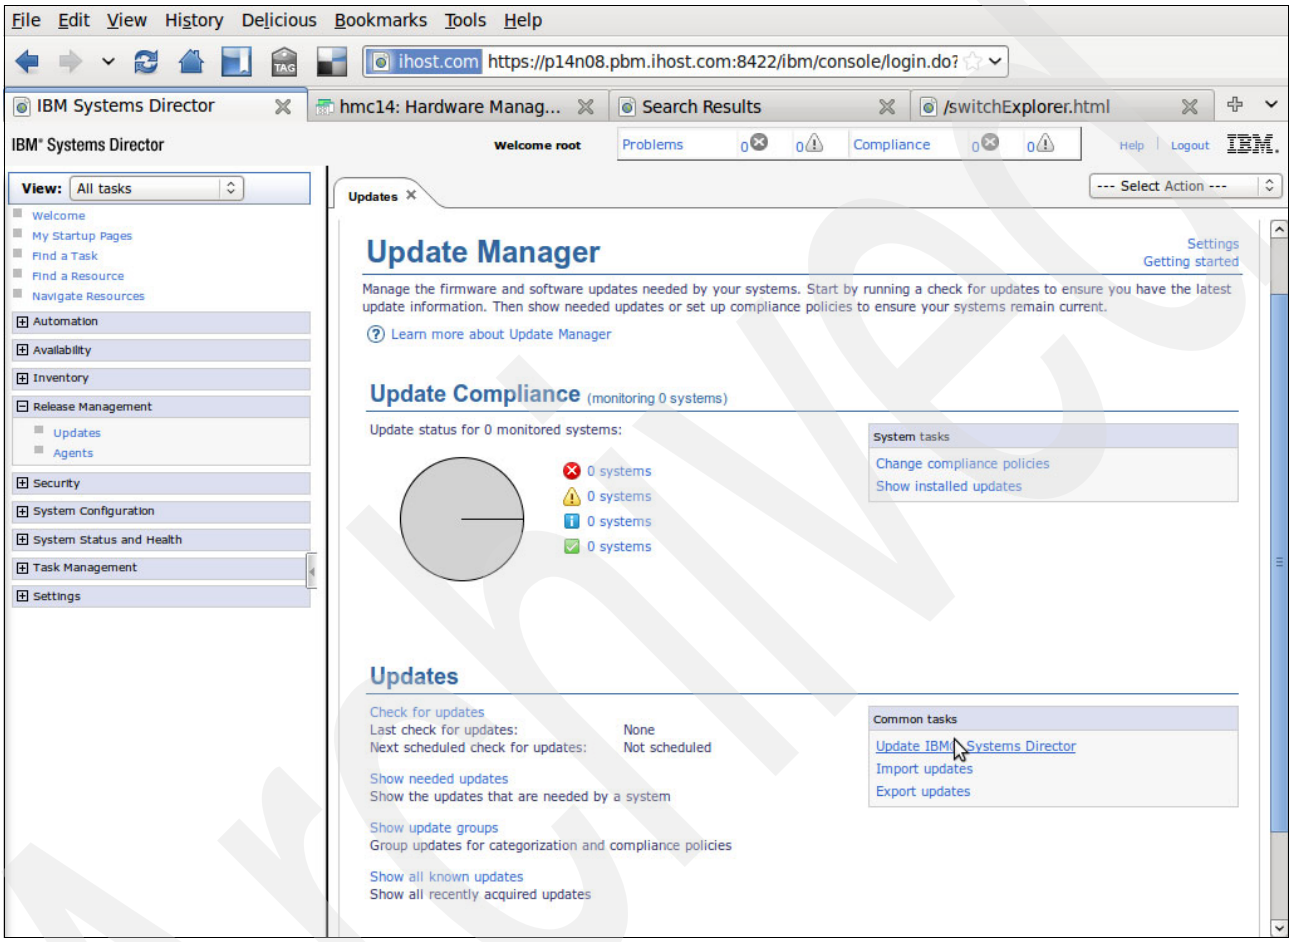

Figure B-1 Update Manager

The Update Manager searches for available updates for the installed components (Figure B-2).

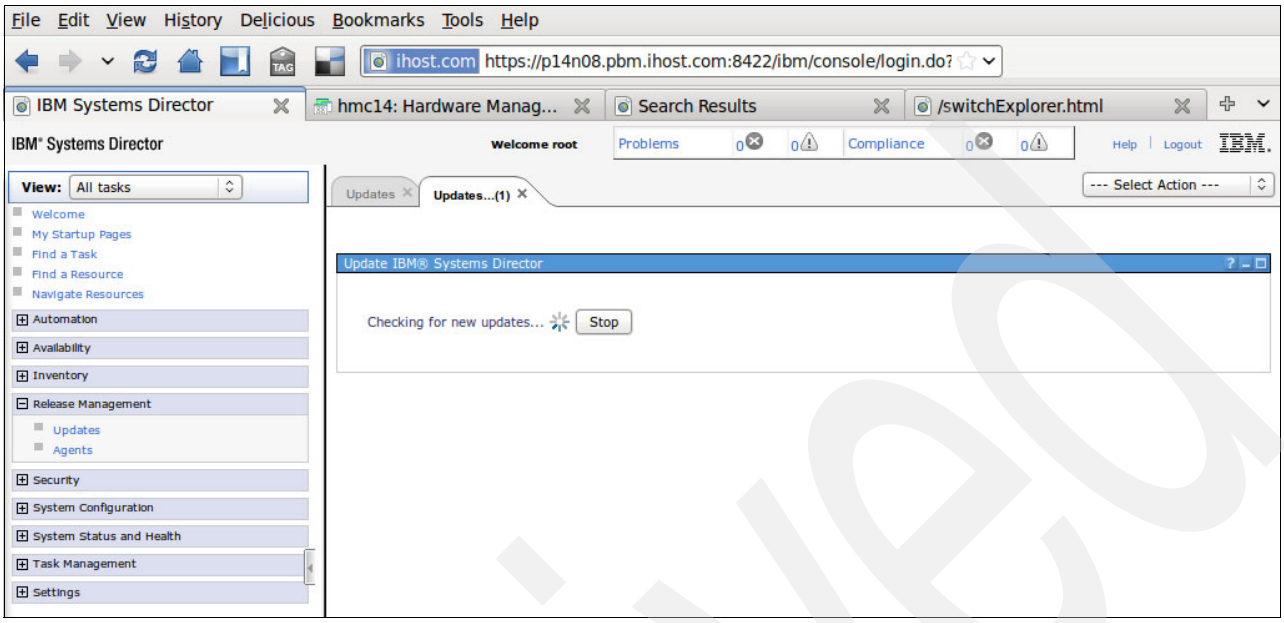

Figure B-2 Searching for available updates to the installed components

The Update Manager displays the currently installed and currently available versions, as shown in Figure B-3.

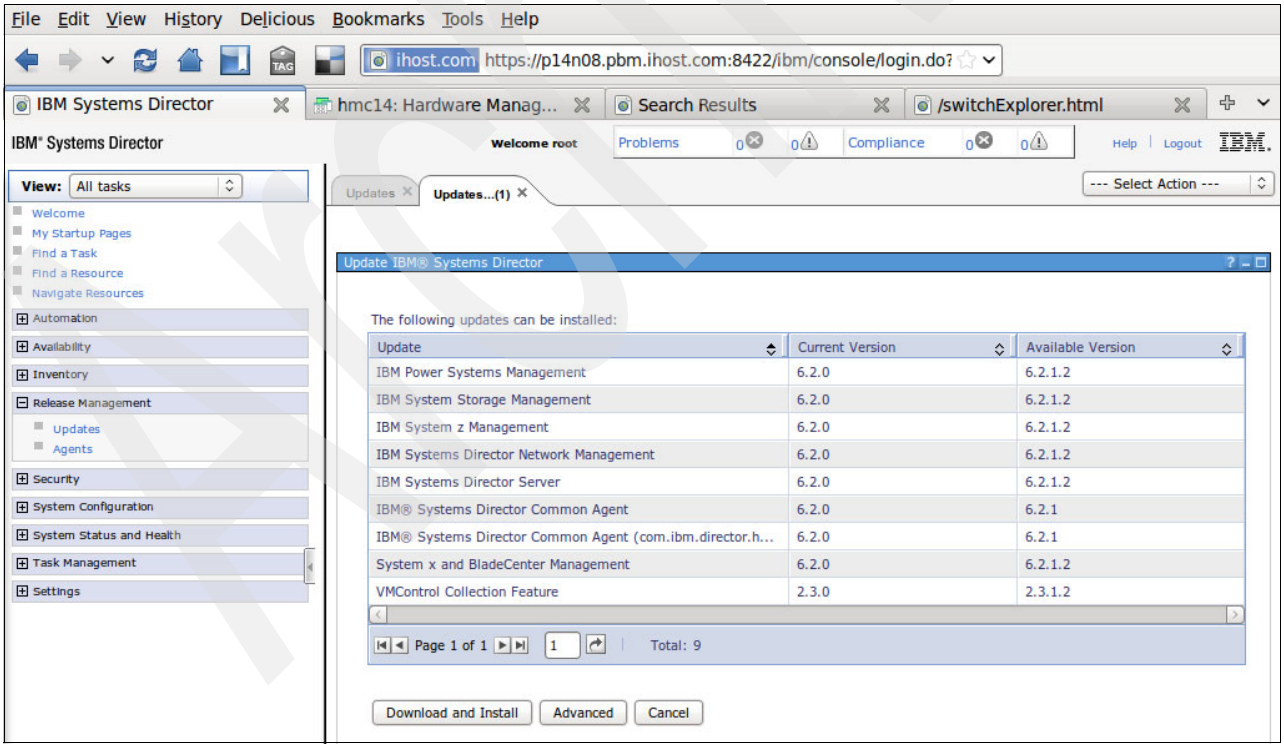

Figure B-3 List of available updates

To install all available updates, click **Download and Install**, which creates a new IBM Systems Director task. When completed, it will require that you restart IBM Systems Director.

To stop the IBM Systems Director, use the **smstop** command, as shown in Example B-1.

Example B-1 smstop command

```
/opt/ibm/director/bin/smstop
Shutting down IBM Director...
```

To start the IBM Systems Director, use the **smstart** command, as shown in Example B-2.

Example B-2 smstart command

```
/opt/ibm/director/bin/smstart
Starting IBM Director...
The starting process may take a while. Please use smstatus to check if the server
is active.
```

You can use the **smstatus** command to verify that the IBM Systems Director is active, as shown in Example B-2.

Example B-3 smstatus command

```
/opt/ibm/director/bin/smstatus
Active
```

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 certain publications referenced in this list might be available in softcopy only.

- *IBM Power 570 and IBM Power 595 (POWER6) System Builder*, REDP-4439
- *IBM Tivoli Storage Productivity Center V4.2 Release Guide*, SG24-7894
- *IBM Systems Director VMControl Implementation Guide on IBM Power Systems*, SG24-7829
- *Implementing IBM Systems Director 6.1*, SG24-7694
- *NIM from A to Z*, SG24-7296
- *IBM PowerVM Virtualization - Introduction and Configuration*, SG24-7940
- General IBM Redbooks portal access to POWER documents:

  `http://www.redbooks.ibm.com/portals/power`

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.

Online resources

These websites are also relevant as further information sources:

IBM Systems Director Information Center

- IBM Systems Director Information Center:

  `http://publib.boulder.ibm.com/infocenter/director/v6r2x/topic/com.ibm.director.editions.doc_6.2.1/editions_power_main.html`

- IBM Systems Director V6.1.x log file locations for troubleshooting:

  `http://publib.boulder.ibm.com/infocenter/director/v6r2x/topic/com.ibm.director.editions.doc_6.2.1/editions_power_main.html`

- How to modify Open Virtual Machine Format (ovf) files:

  http://publib.boulder.ibm.com/infocenter/director/v6r2x/topic/com.ibm.director.vim.helps.doc/fsd0_vim_t_exporting_virtual_appliance_parent.html

- How IBM Systems Director VMControl manages revisions to virtual appliances:

  http://publib.boulder.ibm.com/infocenter/director/v6r2x/index.jsp?topic=%2Fcom.ibm.director.vim.helps.doc%2Ffsd0_vim_r_revisions.html

- The Storage Control Plug-in:

  http://www.ibm.com/systems/software/director/storage/

- The Active Energy Manager:

  http://www.ibm.com/systems/software/director/aem/

CIM publications

- Steps for installing CIM providers for DS4800:

  http://www.lsi.de.com/storage_home/products_home/external_raid/management_software/smis_provider/index.html

- Steps for installing CIM providers for the SAN:

  http://www.brocade.com/services-support/drivers-downloads/smi-agent/application_matrix.page

- CIM agents:

  http://publib.boulder.ibm.com/infocenter/tivihelp/v4r1/topic/com.ibm.tpc_V41.doc/fqz0_t_managing_cim_agents.html

Virtualization and logical partitions (LPARs)

- IBM Virtual I/O Server (VIOS) installation over Network Installation Management (NIM):

  http://www.ibm.com/developerworks/wikis/display/LinuxP/IBM+VIOS+Installation+over+NIM

- How to set up the IBM Virtual I/O Server:

  http://www.ibm.com/developerworks/aix/library/au-aix-vioserver-v2/index.html#vioc1

- The VIO cheat sheet—Consolidate your AIX systems:

  http://www.ibm.com/developerworks/aix/library/au-viocheatsheet/index.html

- Virtual I/O Server and Integrated Virtualization Manager commands:

  http://publib.boulder.ibm.com/infocenter/powersys/v3r1m5/index.jsp?topic=/iphcg/iphcgkickoff.htm

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services

Index

S

T

U

V

Deploying Cloud Components on POWER

Introducing cloud on POWER and the component model overview

Installing and configuring the cloud components

Learning how to manage virtual machines

"*The world is changing. A new reality is emerging for organizations of every size from every part of the planet. Its called the cloud - a profound evolution of IT with revolutionary implications for business and society, creating new possibilities and enabling more efficient, flexible and collaborative computing models.*"

_ http://www.ibm.com/cloud-computing/us/en/

This IBM Redbooks publicationpublication applies to Version 6 Release 1 of AIX on POWER systems. This book is provided as an additional resource as you investigate or consider implementing and deploying a cloud in a POWER environment in the context of infrastructure as a service.

This book is intended for anyone who wants to learn more about cloud computing on POWER systems.

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

SG24-8009-00 ISBN 0738436240