IBM

Implementing IBM SmartCloud Entry on IBM PureFlex System

- Introducing IBM PureFlex System and cloud
- Installing IBM SmartCloud Entry hardware and software
- Configuring and using IBM SmartCloud Entry

Mike Buzzetti
James Kuchler
Charlie Lawrence
Rick Lemelin
Jim T. Suto
Sean M. Swehla

Redbooks

ibm.com/redbooks

International Technical Support Organization

Implementing IBM SmartCloud Entry on IBM PureFlex System

June 2013

SG24-8102-00

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

First Edition (June 2013)

This edition applies to Version 2, Release 4, Modification 1 of IBM SmartCloud Entry and IBM PureFlex System.

Contents

Notices

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

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

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:
IBM Director of Licensing, IBM Corporation, North Castle Drive, Armonk, NY 10504-1785 U.S.A.

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

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

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

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

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

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

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

COPYRIGHT LICENSE:

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

Trademarks

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

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

AIX®
FlashCopy®
Global Technology Services®
IBM®
IBM Flex System™
IBM Flex System Manager™
IBM SmartCloud®
Micromuse®
POWER®
Power Systems™
POWER7®
POWER7+™
PowerVM®
PureApplication™
PureFlex™
PureSystems™
Redbooks®
Redpaper™
Redbooks (logo) ®
Storwize®
System z®
Tivoli®
VMready®
z/Architecture®
z/OS®
z/VM®

The following terms are trademarks of other companies:

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

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

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

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

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

Preface

Distributed computing has been transformed with the introduction of virtualization technology. This has driven a re-architecture of traditional data center workload placement. In 2012, IBM® announced IBM PureSystems™, an offering based on preconfigured software, servers, and storage that form an *expert integrated system*.

Expert integrated systems now combine traditional IT resources into a single optimized solution, with prepackaged components including servers, storage devices, networking equipment, and software. With this evolution of technology, we move from discrete, siloed, and underutilized IT resources to shared resource pools.

This IBM Redbooks® publication can help you install, tailor, and configure IBM SmartCloud® Entry on IBM PureFlex™ System. This book is intended for anyone who wants to learn more about cloud computing with IBM SmartCloud Entry and offerings based on IBM Flex System™ elements.

Authors

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

Several members of this team previously collaborated as authors for the following Redbooks:

- *Deploying a Cloud on IBM System z,* REDP-4711
  (Buzzetti, Kuchler, and Lawrence)
- *Deploying Cloud Components on POWER*, SG24-8009
  (Buzzetti, Kuchler, Lawrence, and Swehla)

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 several 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 Platform Enterprise Edition (Java EE) on IBM z/OS® security, and also several white papers. He is a regular presenter at user conferences and several 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 IBM System z® using IBM z/VM® and IBM Tivoli® products. Previously, James worked for six years as a Performance Tester for z/Linux and z/OS, where he supported the Large System Performance Reference's (LSPR) 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 prerelease components of z/OS, while providing feedback to development and contributing to white papers. James holds a B.S. in Computer Science from the State University of New York at New Paltz.

Charlie Lawrence joined IBM in 1966 and has been involved in various aspects of systems design, development, and testing, and also 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, and also Assembler Language Programming. His technical contributions and experience spans systems assurance and verification activities relating to IBM z/Architecture® ESA Modal Extensions (ESAME), the Common Event Adapter of z/OS, management by exception views of the Tivoli Enterprise Portal and also serving as a member of the IBM team that delivered the Predictive Failure Analysis (PFA) component of z/OS. As an Information Developer, Charlie has co-authored two previous IBM Redpaper™ publications.

Rick Lemelin is an Executive IT Architect and Chief Architect for Converged Infrastructure in the IBM Global Client Center with focus on developing next-generation enterprise architectures. He began his IBM career in 2001 in IBM Global Technology Services®, where he was an Architect and Complex Solution Leader, specializing in IT Service Management, IT Operations consulting, IT Optimization, Infrastructure Virtualization, and IT Technology architecture. Rick led projects for clients focused on determining optimal usage for their information technology, based on their unique business process requirements, to realize IT flexibility, IT to business alignment and sustainable cost savings. He has worked with dozens of companies in the Insurance, Banking, Financial Services, Industrial, Public, Health Care, Pharmaceutical, Retail, Energy, Consumer Products, Technology, and Communications industries. Prior to joining IBM, Rick worked 17 years at a U.S. based regional grocery chain, Hannaford Bros. Co., serving 10 years in IT Operations.

Jim T. Suto is a Cloud Technology Specialist at the IBM Power and Smart Analytics Benchmark Center where he is working on designing cloud solutions for internal and external customers using best-of-class technologies. Jim joined IBM in 2006 as a Network Systems Engineer for the IBM Tivoli Software Group development data center in New York City. In his role there, Jim worked on architecting the design and implementation of web-based technologies, enabling the monitoring of connectivity, resource availability, and environmental conditions of Tivoli labs worldwide. As part of an initiative to enhance the scalability of the infrastructure, Jim also worked with a global team deploying cloud computing capabilities at multiple anchor sites enabling software developers and testers to quickly get systems provisioned when projects required additional servers. Prior to 2006, Jim worked as Senior Systems Administrator for IBM Micromuse®.

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:

Robert J. Brenneman
Frank De Gilio
Dave Dubetsky
Robert Maher
Sharon Selzo
IBM Poughkeepsie

Mike Ebbers
Ella Buslovich
IBM ITSO, Poughkeepsie center

Now you can become a published author, too!

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

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

ibm.com/redbooks/residencies.html

Comments welcome

Your comments are important to us!

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

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

  ibm.com/redbooks

- Send your comments in an email to:

  redbooks@us.ibm.com

- Mail your comments to:

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

Stay connected to IBM Redbooks

- Find us on Facebook:

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

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

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

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

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

1

Introduction to IBM PureFlex System and cloud

Distributed computing has undergone a transformation over the past decade, with the introduction and mass adoption of hypervisor-based virtualization technology. As this technology has evolved, it has driven a re-architecture of traditional data center infrastructure layouts, enabling new approaches to workload placement.

Originally, deploying virtualization simply enabled multiple discrete workloads to be placed on the same physical server, leading to higher asset utilization and server consolidation. However, as this technology has matured, coupled with new industry standards driving new features, and higher virtual machine (VM) density per physical server was achieved, it has also begun to impact traditional interactions between server, storage, and network entities, including entity management and administration.

The industry was looking for a new approach. Hence, the concept of converged infrastructure, or expert integrated systems was born. Expert integrated systems package traditional IT resources into a single optimized solution, with prepacked components including servers, storage devices, networking equipment, and software for IT infrastructure management, automation, and orchestration. With this evolution of technology, we moved from discrete, siloed, and underutilized IT resources to shared resource pools. These resource pools, enabled by expert integrated systems, exploit inherent features of advanced hypervisor-based virtualization and cloud computing. To meet these new industry requirements, IBM developed a completely new integrated data center solution.

In 2012, IBM announced IBM PureSystems, an offering based on preconfigured software, servers, and storage that form an *expert integrated system*.

1.1 Introduction to IBM Flex System hardware

IBM Flex System hardware provides the building blocks for IBM PureFlex System and IBM PureApplication™ System.

IBM Flex System is at the heart of the new type of data center building blocks. The IBM Flex System can meet the growing needs of any organization that requires the most advanced blade system to date. The design of the Flex System hardware is focused on long term investments, performance, and flexibility. This focus allows Flex System to support client needs now and with future generations of technology.

IBM Flex System is based on four main integrated components:

- Compute nodes
- Storage
- Networking
- Management

IBM Flex System compute nodes support multiple compute architectures and up to four separate operating systems. The converged infrastructure and advanced management aim to improve utilization and productivity to ultimately lower the cost of delivering IT services. The storage capabilities are built for performance and scalability.

IBM PureFlex System offering is designed for private clouds. It is available in three configurations:

- Express
- Standard
- Enterprise

By combining the flexibility of the general purpose PureFlex System with SmartCloud Entry, clients can add cloud functionality to their converged environment. SmartCloud Entry provides enterprises with the capability of managing their multiple architecture cloud solution on a single system.

1.1.1 Three configurations

PureFlex combines the components of IBM Flex System hardware into expert integrated systems. They are factory-preconfigured in Express, Standard, and Enterprise configurations to meet a wide variety of needs. All of the PureFlex configurations include a PureFlex Enterprise Chassis incorporated into a PureFlex 42U Rack, and also a standard IBM Flex System Manager™ (FSM) as described in section 1.5. Each configuration brings a different level of power and cooling, appropriate to the size of the intended deployment.

IBM PureFlex System Express

The IBM PureFlex System Express configuration is the entry point in the PureFlex family. It includes, as standard, two power supplies (out of a maximum of six), and two fans (maximum of eight).

IBM PureFlex System Standard

The PureFlex Standard configuration adds two more power supplies, for a total of four.

IBM PureFlex System Enterprise

The PureFlex Enterprise configuration is preconfigured with the maximum of six power supplies allowed in an Enterprise chassis, and also all eight configurable fans. In this configuration, the maximum number of PureFlex compute nodes can be supported with N+N power supply redundancy, providing the highest levels of compute power and redundancy.

1.2 IBM Flex System x86 Nodes

IBM Flex System x86 nodes are built on Intel Xeon processors and support a wide range of virtualization technologies and operating systems. There are several members of the IBM Flex System x86 family. Only the IBM Flex System x240 systems were used in the sample environment for this writing.

1.2.1 IBM Flex System x240

The IBM Flex System x240 is a two-socket, single-bay Flex node based on Intel Xeon processors. The system supports up to 768 GB of RAM and up to 2 TB of internal storage. Integrated 10GbE networking and system management make the x240 an excellent choice for a wide variety of workloads, including mainstream virtualization.

1.3 IBM Flex Power Systems nodes

IBM Flex Power Systems™ nodes bring the Power architecture to the Flex family. The Flex Power Systems nodes are built on IBM POWER7® and IBM POWER7+™ technology, which bring faster frequencies and larger L3 cache sizes and improve performance on most workloads. POWER7 family processors are also extremely efficient and provide more performance per watt of energy consumed. Both Flex Power Systems nodes, described next, support IBM AIX®, IBM i, and Linux.

1.3.1 IBM Flex System p260 compute node

The Flex System p260 compute node is a single-wide node and is available with 8 or 16 cores in speeds in the range of 3.2 - 4.1 GHz. The p260 supports up to 512 GB of RAM.

1.3.2 IBM Flex System p460 compute node

The Flex System p460 compute node provides four sockets in a double-wide form-factor. This allows a p460 to have either 16 or 32 POWER7 cores, with 4 MB of L3 cache per processor core. With a maximum of 1 TB of RAM per node, the p460 is suitable for memory and compute-intensive workloads, and also for meeting the demands of heavy virtualization.

1.4 IBM Flex System V7000

The IBM Flex System V7000 is the IBM Storwize® V7000 adapted to the IBM Flex System ecosystem. It can house up to 240 drives for each control enclosure. Currently supported drives include solid-state drives (SSDs), nearline SAS, 10,000 RPM (10k) SAS, and 15k SAS drives. Built-in support for tiering and advanced replication features make the most efficient use of the installed drives and allow a variety of configurations to suit many environments.

1.5 IBM Flex System Manager

The IBM Flex System Manager (FSM) is the center of IBM Flex System management. The FSM provides a real-time interactive chassis map to simplify monitoring and management of resources in the Enterprise chassis. The FSM provides access to power and health monitoring of nodes and also configuration of installed nodes, and simplifies deployment of workloads across compute nodes. It is a central point of management for every element in the Enterprise chassis.

A single FSM can manage up to four Enterprise chassis.

1.6 IBM Flex System Enterprise Chassis

The IBM Flex System consists of one or more rack-mounted chassis. Each chassis can contain two Chassis Management Modules (CMM), up to 14 compute nodes, six power supplies, and 10 fans.

1.6.1 Chassis Management Module (CMM)

The CMM is a single-chassis manager that communicates with the individual compute note management controllers. CMM supports system monitoring, event recording, and alerts. It also provides an interface for the management of the chassis, its devices, and the compute nodes. Each chassis supports two CMMs for higher availability.

1.6.2 Midplane

The midplane in an IBM Flex System Enterprise chassis serves several functions:

- It is responsible for distributing power. It does this in a single power domain, distributing power to the compute nodes, the IO modules (scalable switch elements), and ancillary components.
- It provides the physical connectivity from network and storage adapters in each node to corresponding ports in the IO modules.
- It contains apertures which connect the cooling channels in the front of the enterprise chassis to the appropriate channels in the rear of the chassis to provide better airflow and more efficient cooling.

The IBM Flex System Enterprise chassis midplane design differs from the backplane design used in many blade chassis in that it eliminates active components. The midplane is designed for high-frequencies and will support the next generation of networking and storage signalling protocols.

1.6.3 Power

Power in the Flex System Enterprise chassis is provided by up to six power supplies that can supply 2500W each. This flexibility allows you to grow your power capacity according to your needs. The power supplies in an Enterprise chassis can be configured for N+1 or N+N redundancy to provide the combination of power and availability most appropriate for the target environment.

1.6.4 Cooling

The Enterprise chassis is designed to provide excellent cooling capacity to the installed nodes, I/O modules and management modules. Cooling is divided into two zones, each of which allows for up to four 80 mm fans.

1.7 IBM Flex System scalable switch elements

Scalable switch elements provide the Flex System nodes and the enterprise chassis with access to networking and storage resources. The Enterprise chassis provides space for up to four scalable switch elements. These switch elements can be any combination of network, storage, or converged IO devices. Several options are available. Following is a description of the modules employed for the purposes of this publication.

1.7.1 IBM Flex System Fabric EN4093 10 Gb Scalable Switch

The EN4093 is a 10 Gb Layer 2/3-managed Ethernet switch that provides up to 42 internal and 22 external ports. The basic configuration of the EN4093 provides 14 internal 10GbE ports and 10 external 10GbE ports. Two Features on Demand upgrades are available. With upgrade 1 applied, the switch provides 28 internal 10GbE ports, 10 external 10GbE ports, and two external 40GbE ports. With upgrade 2 applied, the switch provides 42 internal 10GbE ports, 14 external 10GbE ports, and four external 40GbE ports. This, as with several other features of the IBM Flex System products, allows the switch to grow to meet the demands of the data center without wasting capacity that is not needed. Support for advanced virtualization features such as Virtual Fabric and IBM VMready® make the switch well suited to highly virtualized environments.

1.7.2 IBM Flex System FC3171 8Gb SAN Switch

The FC3171 is an 8 Gb SAN switch that provides 14 full-duplex, internal ports and six external ports. The FC3171 is available as a pass-through switch, to expose only the host ports to the SAN, and as a regular SAN switch.

2

Overview of example IBM Flex System

The system used for this book is based on the IBM PureFlex offering. However, the IBM PureFlex offering can be created from base IBM Flex System components. This chapter details the contents of the chassis and the placement of the nodes and scalable switch elements (ScSE).

2.1 How the system is set up

Our Flex System Enterprise Chassis (7893-92X) is populated as described in this section.

Figure 2-1 shows the front view of chassis; Table 2-1 on page 9 lists the contents of the slots.

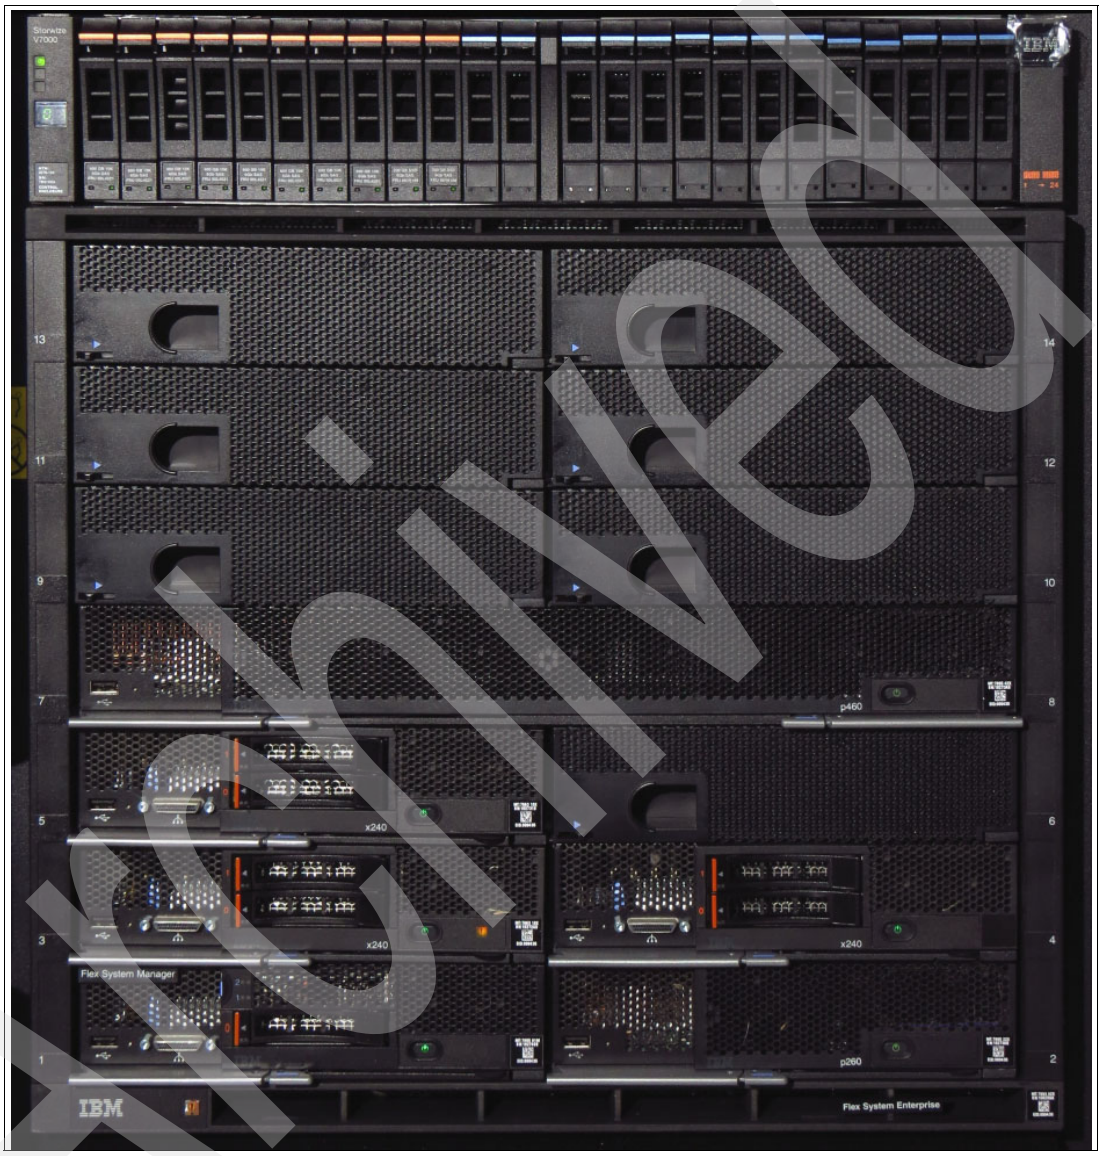

Figure 2-1 Front view of Flex Chassis and V7000

Table 2-1 Contents of the slots

Slot contents	Location
Flex System Manager (7955-01M)	Slot 1
Flex System p260 (7895-22X)	Slot 2
Flex System x240 (7863-10X)	Slot 3
Flex System x240 (7863-10X)	Slot 4
Flex System x240 (7863-10X)	Slot 5
Empty	Slot 6
Flex System p460 (7895-42X)	Slots 7 and 8
Empty	Slots 9 through 14

Figure 2-2 shows the rear view of the chassis; Table 2-2 on page 10 lists contents of the slots.

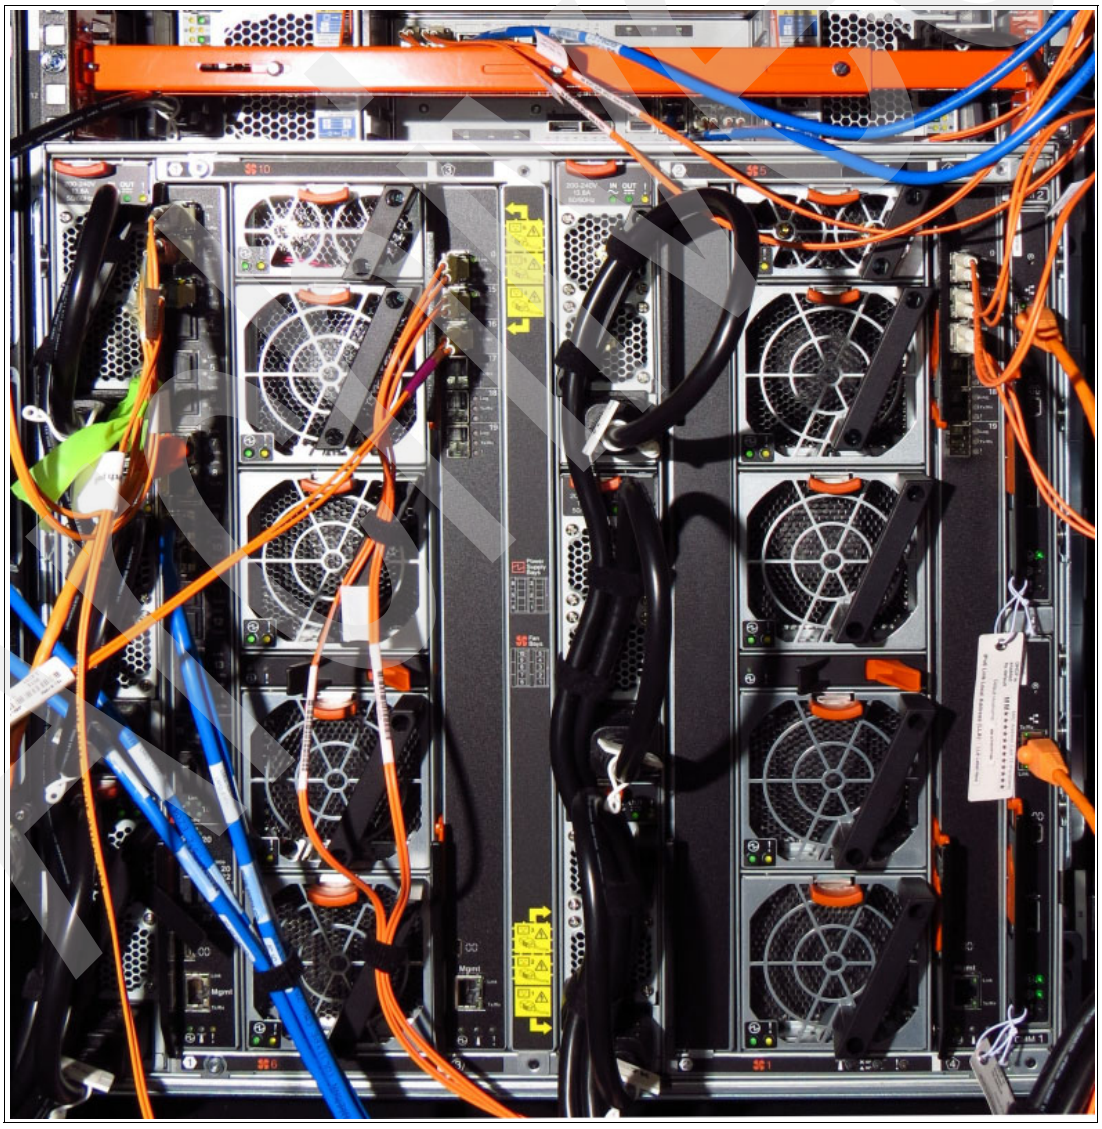

Figure 2-2 Rear view of Flex Chassis

Table 2-2 Contents of the slots

Slot content	Location
IBM Flex System EN4093 10 Gb Virtual Fabric Scalable Switch	Slot 1
Empty	Slot 2
IBM Flex System FC3171 8 Gb SAN Switch	Slot 3
IBM Flex System FC3171 8 Gb SAN Switch	Slot 4
Chassis Management Module	Mounted in CMM slot 1
Chassis Management Module	Mounted in CMM slot 2

2.2 Network setup

The Flex system has a complex networking environment. (In later sections of this book, such as 4.3, “Networking setup” on page 27 and 5.3, “Configuring cloud network pools” on page 151, this environment is described in detail.) There are two separate subnetworks configured for this environment: one for management of the nodes and one for the node usage.

NOTE: The network external to the chassis in this environment currently does not support IP version 6 (IPv6.) For the remainder of this document, IPv4 addresses are used. However, in some instances IPv6 addresses are shown because of the default network configurations; these interfaces are not used in this environment.

Management network

Table 2-3 shows the IP addresses of our management network. The gateway was set to 129.40.180.94 and netmask to 255.255.255.224 on all the components listed in the table.

Table 2-3 Management network slots and addresses

IP address	Location
129.40.180.66	Slot 1
129.40.180.82	Slot 2
129.40.180.83	Slot 3
129.40.180.84	Slot 4
129.40.180.85	Slot 5
Empty	Slot 6
129.40.180.87	Slots 7 and 8
Empty	Slots 9 through 14
129.40.180.65	CMM Slot1
129.40.180.70	I/O Bay 1
129.40.180.73	I/O Bay 3

IP address	Location
129.40.180.74	I/O Bay 4
129.40.180.68	V7000 Canister 1
129.40.180.69	V7000 Canister 2

Data network

Table 2-4 shows the IP addresses of our user data network. The gateway was set to 129.40.21.222 and netmask to 255.255.255.224 on all the components listed in the table.

Table 2-4 User network addresses and slots

IP Addresses	Location
129.40.21.201	Slot 1
129.40.21.202	Slot 2
129.40.21.203	Slot 3
129.40.21.204	Slot 4
129.40.21.205	Slot 5
Empty	Slot 6
129.40.21.207 129.40.21.208	Slots 7 and 8
Empty	Slots 9 through 14

3

IBM SmartCloud Entry introduction and concepts

IBM SmartCloud Entry for IBM Flex System is focused on providing private cloud services. SmartCloud Entry is designed to be simple to install and easy to use. It has a self-service web portal that eases the deployment of virtual machines. It supports a variety of deployment options and works with many cloud infrastructures. For the purposes of this document, the focus is Flex System compute nodes. SmartCloud Entry can manage both Power compute nodes and x86 compute notes. At the time of this writing, VMware is the only virtualization platform that is supported by IBM SmartCloud Entry on IBM Flex Systems Intel compute nodes. IBM has stated that it intends to support Hyper-V and KVM virtualization platforms in the future.

From the web portal, IBM SmartCloud Entry supplies administrators with a point of control to manage their virtualized environments. Administrators can deploy new virtual machines, modify the machine characteristics, and delete them when they are no longer needed. One can also enable reporting of usage and metrics to provide appropriate billing.

To explain IBM SmartCloud Entry cross-platform manageability, this chapter introduces several terms and concepts.

3.1 Appliance

An appliance is the basis for the cloud service that is to be provided. Appliances contain the operating system and software stack that are required for the service to be created. Appliances are created by either capturing a running workload or using one of the underlying cloud tools. For x86 compute nodes, a VMware Template is an appliance; for Power compute nodes, a VMControl captured virtual server is an appliance. Each appliance can be associated with a project. Appliances store the configuration information about the type of virtual machine to provision, its networking configuration, storage configuration, and a few other settings. Appliances also store information about the amount of processing and memory that the service requires.This information differs for each of the cloud pools that it represents. For instance, an appliance that is captured from a working IBM POWER® compute node has settings for processor pooling and the weight of the processor; VMware has different settings. IBM SmartCloud Entry user interface is shown in Figure 3-1.

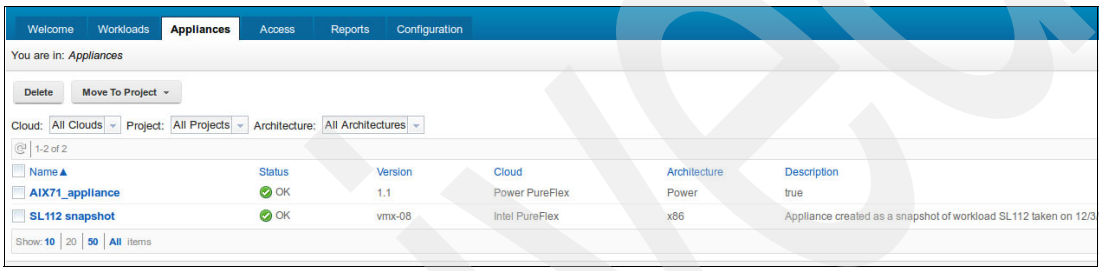

Figure 3-1 Appliances

When a user selects and deploys an appliance, that instance becomes a workload.

3.2 Workload

IBM SmartCloud Entry uses the term *workload* for both x86 and POWER based compute nodes. A workload is the combination of the metadata that describes a virtual machine and also the actual virtual machine installation data. IBM SmartCloud Entry stores workload information for POWER based compute nodes similarly to the workload information that is stored in the Flex System Manager (FSM), including the AIX virtual machines. Workloads on an x86 compute node are equivalent to VMware virtual machines. SmartCloud Entry stores additional information to facilitate high level cloud functions such as copying a workload and having approval information associated with a workload.

Figure 3-2 shows workloads.

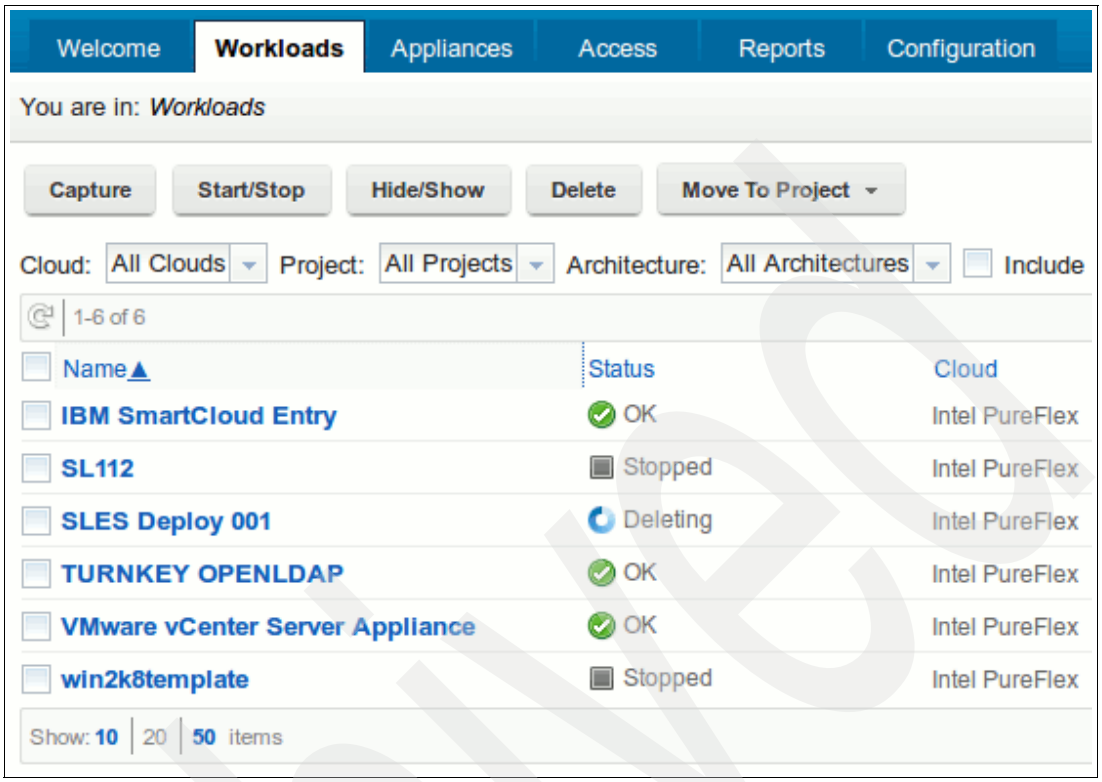

Figure 3-2 Workloads

3.3 Approvals

Users can request to deploy an appliance in the projects to which they have access. System administrators can enable approvals to ensure that the requests go through the proper channels. Approvals can be associated with projects or with clouds. Approvals can be set for the following actions:

- Workload initiation
- Workload expiration extend
- Workload resize
- Workload capture
- Virtual server attach storage request
- Virtual server detach storage request
- Virtual server save image
- Virtual server restore

When one of these actions is requested, a notification is sent to either the cloud administrator or the project owner. They can then approve the request and allow it to continue, or deny the request. The user is notified with the result and the action is automatically taken.

3.4 Metering

IBM SmartCloud Entry provides the capability to meter the services that are deployed as workloads. For each workload, there is an automated collection of resource usage metrics. This information consists of the following information:

- Elapsed time the workload has been available for in hours.
- Amount of processors (CPUs) consumed
- Amount of memory (in GB) consumed
- Amount of storage (in GB) consumed

Figure 3-3 shows the SmartCloud Entry User Interface usage metering report.

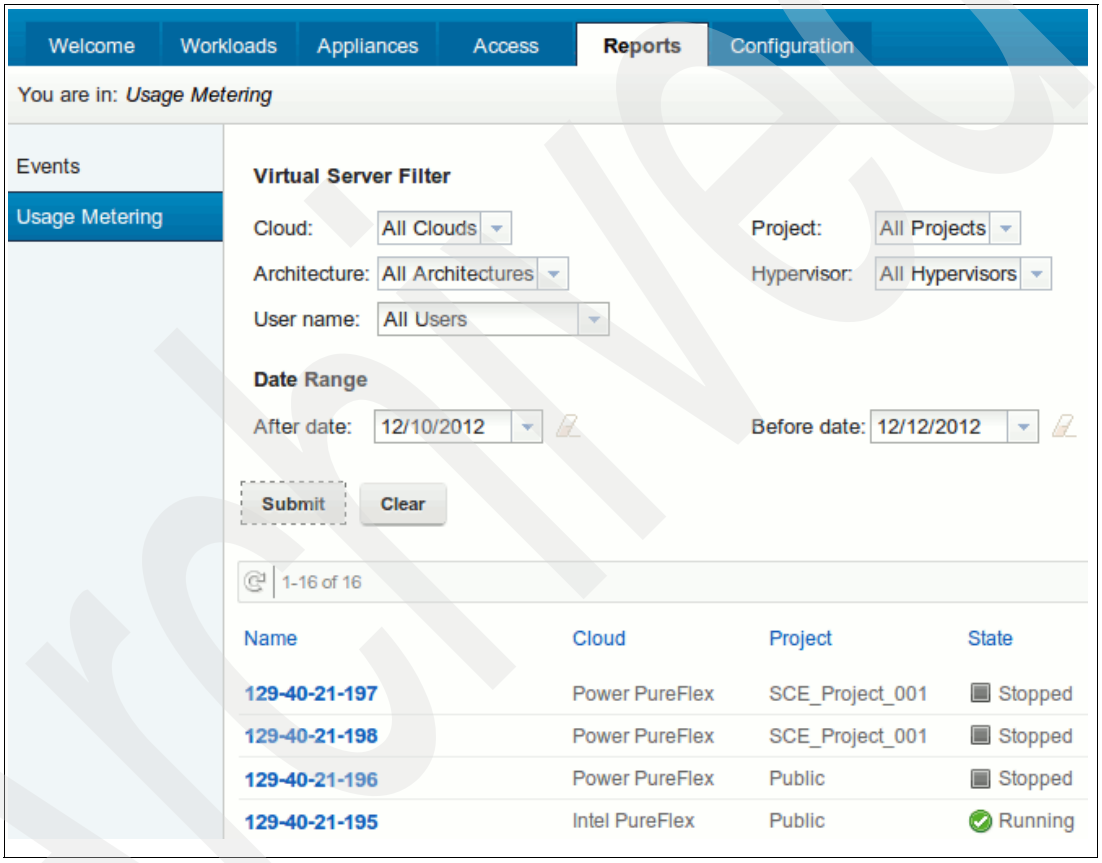

Figure 3-3 Usage Metering report

Administrators can create a bill for the services used with this information. IBM SmartCloud Entry can combine this information for all workloads in its own internal billing system.

3.5 Configuring

For the steps to configure these IBM SmartCloud Entry features, see Chapter 5, “SmartCloud Entry installation and configuration” on page 141.

4

Hardware setup and configuration

Each PureFlex System configuration includes built-in configuration patterns to enable simpler and faster converged infrastructure configuration, virtual machine (VM) deployment, management, and simplifying key tasks across all IT resources. Moreover, they can be highly customized for each client-specific target workload and environment.

4.1 Overview

Unique to IBM, the PureFlex solution delivers a hybrid compute architecture to provide extreme flexibility for the types of workloads it can handle within a single frame, including the following items:

- Choice of architectures: IBM POWER7 or x86 processors within the same systems.
- Choice of OS: AIX, IBM i, Microsoft Windows, and Linux from Red Hat or SUSE
- Choice of hypervisors: IBM PowerVM®, KVM, VMware, or Microsoft HyperV
- Designed for cloud: SmartCloud Entry included on Standard and Enterprise configurations
- Designed for simplicity: Integrated, single system management across physical and virtual resources

This architecture clearly differentiates IBM from the competition. Moreover, the forward-looking node architecture enables PureFlex to achieve higher VM density levels too.

From a memory, midplane, I/O, and processor or core standpoint, the nodes are more powerful than traditional blades although they are delivered in a similar footprint. Currently, clients are achieving VM compression ratios ranging from 4:1 to 30:1 on traditional blades. The node architecture should be able to attain higher VM compression ratios. On average, this can be two to three times the compression ratio of traditional blades.

The IBM Flex System is a rack-mounted chassis based system. Each chassis will contain compute nodes (POWER Systems or x86) and rear bays for power supplies, fan modules and scalable switch elements. A midplane resides in the chassis to facilitate airflow, power connections, and other connectivity between the compute nodes and the power and I/O connectors.

Compute nodes within the chassis are managed using four interfaces.

- The Integrated Management Module (IMM) is used to manage an x86 compute node. The IMM consolidates the service processor functionality, Super I/O, video controller, and remote presence capabilities in a single chip on the server system board.

  Several IMM standard features are as follows:

  - Access to critical server settings
  - Access to server vital product data (VPD)
  - Advanced Predictive Failure Analysis (PFA) support
  - Automatic notification and alerts
  - Continuous health monitoring and control
  - Simple Network Management Protocol (SNMP) support
  - User authentication using a secure connection to a Lightweight Directory Access Protocol (LDAP) server
- The Advanced System Management (ASM) is used to manage a Power Systems compute node.
- The Chassis Management Module (CMM) manages the devices in the Chassis in which it resides. The CMM provides single-chassis management. The CMM is used to communicate with the management controller in each compute node (IMMv2 in x86 processor-based compute nodes and FSP in POWER7 processor-based compute nodes) to provide system monitoring, event recording and alerts, and to manage the chassis, its devices, and the compute nodes. The chassis supports up to two CMMs. If one CMM fails, the second CMM can detect its inactivity and activate itself to take control of the system without any disruption. The CMM is central to the management of the chassis and is required in the Enterprise Chassis.

Several functions are as follows:

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

- The IBM Flex System Manager (FSM) is used to manage one or more chassis. It is a high performance scalable system management appliance. The FSM hardware is preloaded with systems management software, which enables you to configure, monitor, and manage IBM Flex System resources in up to four chassis, currently.

  Several high level features are as follows:

  - Supports a comprehensive, pre-integrated system that is configured to optimize performance and efficiency.
  - Automated processes triggered by events simplify management and reduce manual administrative tasks.
  - Centralized management reduces the skills and the number of steps it takes to manage and deploy a system.
  - Enables comprehensive management and control of energy utilization and costs
  - Automates responses for a reduced need for manual tasks: Custom actions / filters, configure, edit, relocate, automation plans.
  - Full integration with server views, including virtual server views enables efficient management of resources.

4.2 Chassis Management Module (CMM) setup

After the PureFlex System is installed in the rack and the cabling is complete, the next step is setting up the CMM. The following network information is required in order to finalize the CMM setup:

- Host name for the CMM
- IP address and netmask
- Gateway
- DNS server address
- NTP server address (optional)

The CMM can be accessed by connecting a notebook to the CMM using the serial or Ethernet port located on the back of the CMM. This requires either an Ethernet patch cable or a serial cable adapter (IBM Flex System Management Serial Access Cable Part 90Y9338). The default factory defined IP address of the CMM is 192.168.70.100.

To be able to connect to the CMM initially, the network interface on the notebook must be configured with an IP address on the same subnet as the CMM. For example, configure eth0

with the IP address: 192.168.70.10 and netmask 255.255.255.0. After the connection is established, open a web browser to the following URL:

`https://192.168.70.100`

Then, complete the following steps:

1. When the IBM Chassis Management Module login panel opens (Figure 4-1), enter the default user ID and password (USERID and PASSW0RD) to log in to the CMM for the first time.

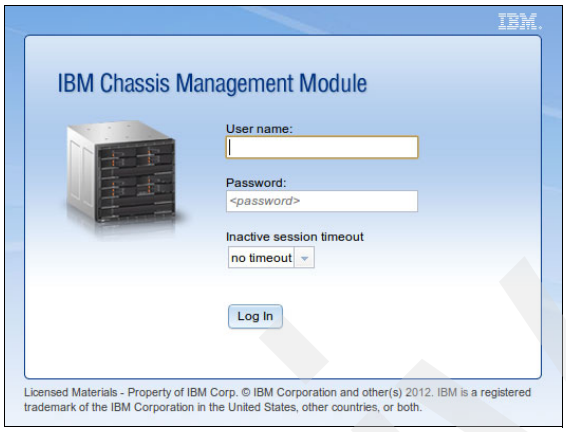

Figure 4-1 Chassis Management Module login page

2. The Initial Setup Wizard, Welcome page opens (Figure 4-2), which provides information to help you get started. Review this and click **Next**.

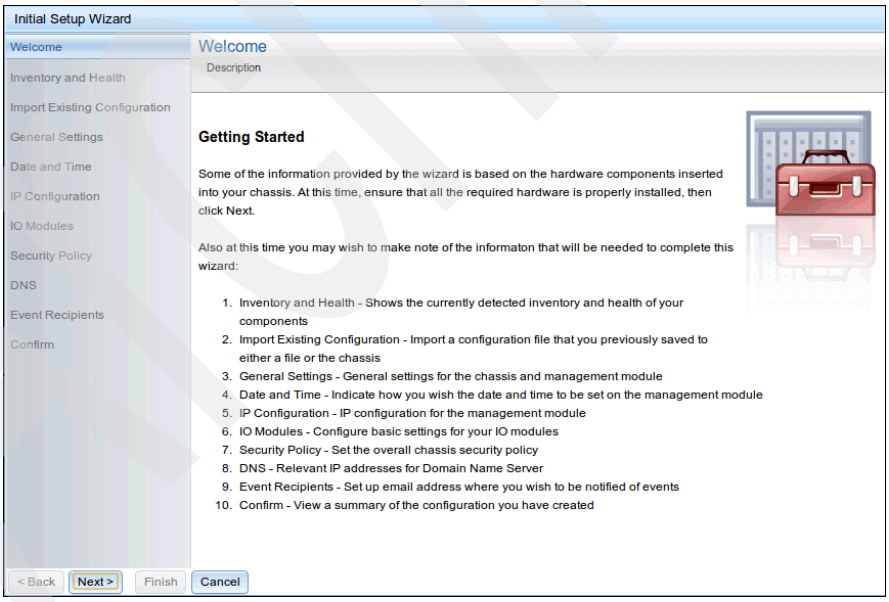

Figure 4-2 CMM Initial Setup Wizard: Welcome

3. The wizard's Inventory and Health page (Figure 4-3) opens. Warnings might be displayed. At this point, you can ignore the warnings. Click **Next**.

Figure 4-3 CMM Initial Setup Wizard: Inventory and Health

4. The Import Existing Configuration page opens (Figure 4-4). The following message is indicated near the top of the page:

   `If this is your first time setting up a chassis, you will not have a configuration to import.`

   Because there is no configuration that can be imported, manual steps are necessary to complete the configuration. Click **Next**.

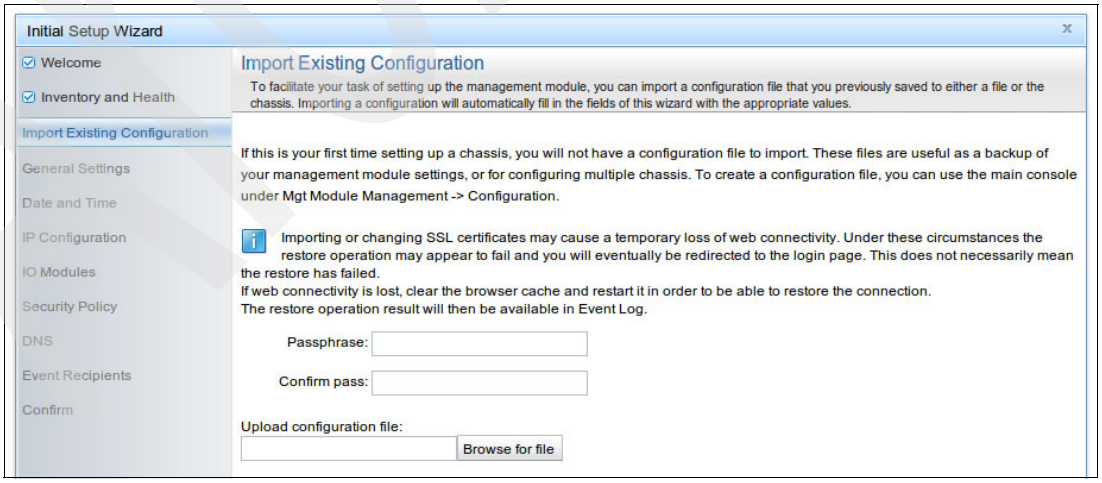

Figure 4-4 CMM Initial Setup Wizard: Import Existing Configuration

5. The General Settings page opens (Figure 4-5). Enter a management module name. The management module name is a unique identifier. This unique identifier can be the host name or, as shown in the figure, the default entry can also be used, which is the serial number of the CMM.

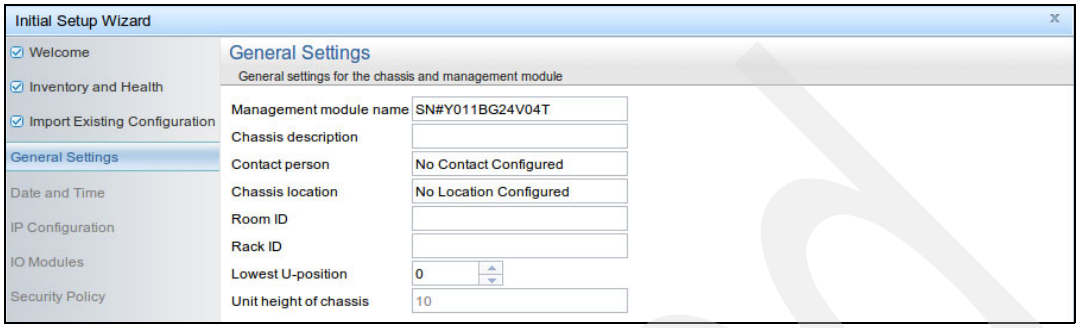

Figure 4-5 CMM Initial Setup Wizard: General Settings

The other General Settings fields are optional and specific to the CMM environment and installation. After providing the requested information, click **Next.**

6. The Date and Time page opens (Figure 4-6). On this page, enter the current date and time, select the appropriate time zone, and indicate whether or not the system should automatically adjust the clock for Daylight Savings Time. Click **Apply** to make activate the settings. Then, click **Next**.

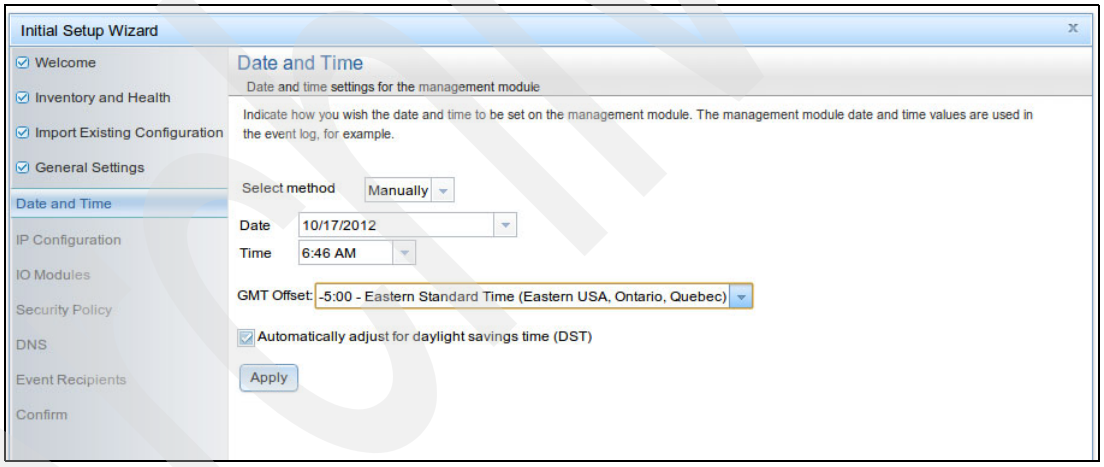

Figure 4-6 CMM Initial Setup Wizard: Date and Time

7. The IP Configuration page opens (Figure 4-7). Enter the network information for the CMM, including the host name, domain name, gateway and DNS server. If the host name specified does not resolve to an IP address, select the **Register this interface with DNS** check box. Click **Next**.

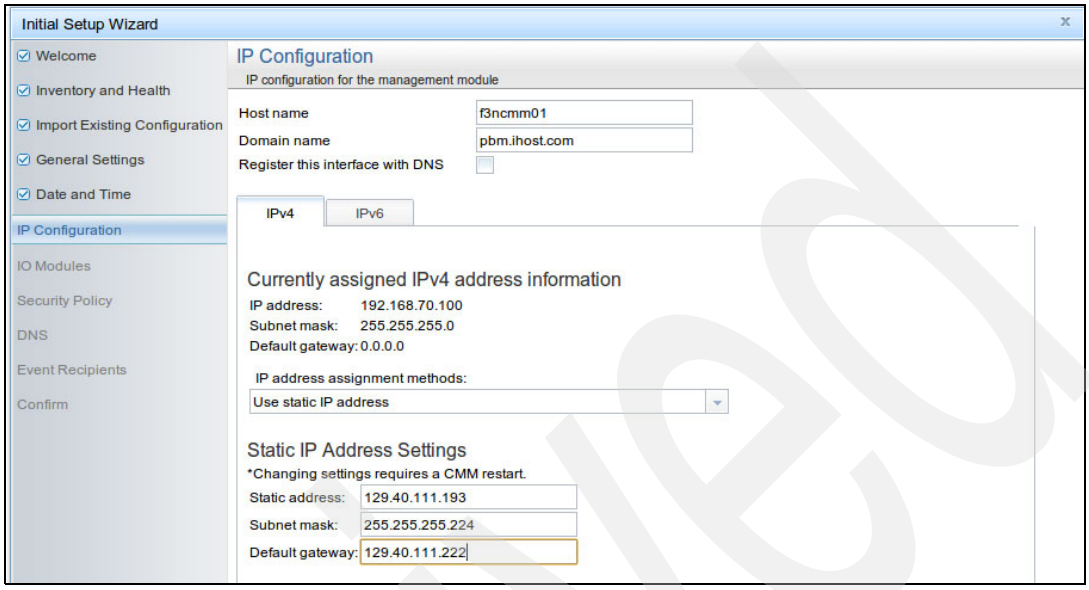

Figure 4-7 CMM Initial Setup Wizard: IP Configuration

8. The IO Modules page opens (Figure 4-8). The defaults on this page are acceptable. Click **Next**.

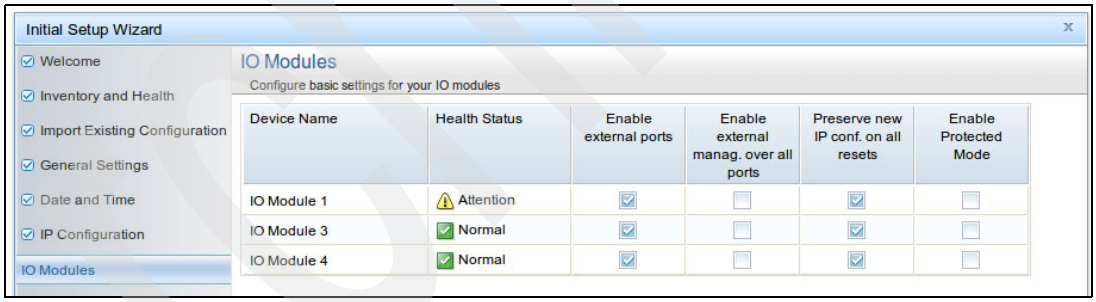

Figure 4-8 CMM Initial Setup Wizard: IO Modules

9. The Security Policy page opens (Figure 4-9). “Secure” is the recommended level for Security Policy as explained to the right of the slider bar. Select a security policy and click **Next**.

Figure 4-9 CMM Initial Setup Wizard: Security Policy

10. The DNS page opens (Figure 4-10). Configure the DNS. The DNS configuration tasks are referenced in Figure 4-10 and Figure 4-11 on page 25.

    The DNS page indicates an IPv6 network with DNS not enabled. Select the appropriate DNS address type (IPv6 or IPv4) from the drop-down list and, as needed, select the **Enable DNS** check box. To specify DNS information for IPv4, select IPv4 using the drop-down list.

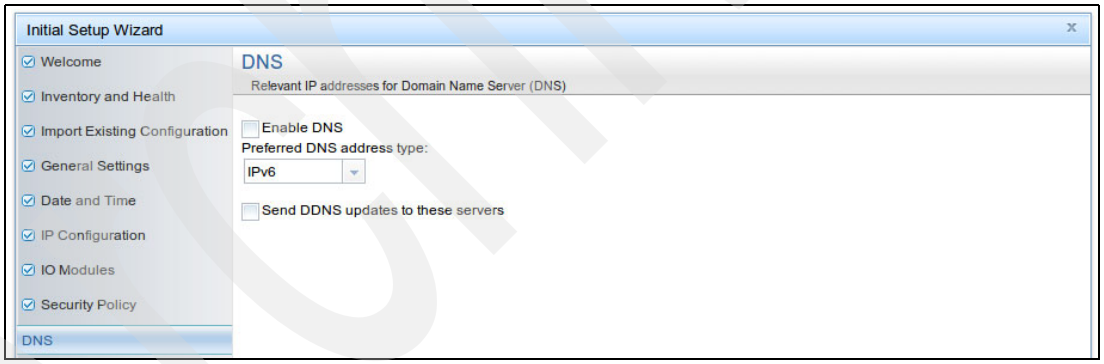

Figure 4-10 CMM Initial Setup Wizard: DNS Specifications

DNS information for IPv4 is specified in Figure 4-11. This figure also shows Primary and Secondary IPv4 addresses for the DNS servers. Click **Next**.

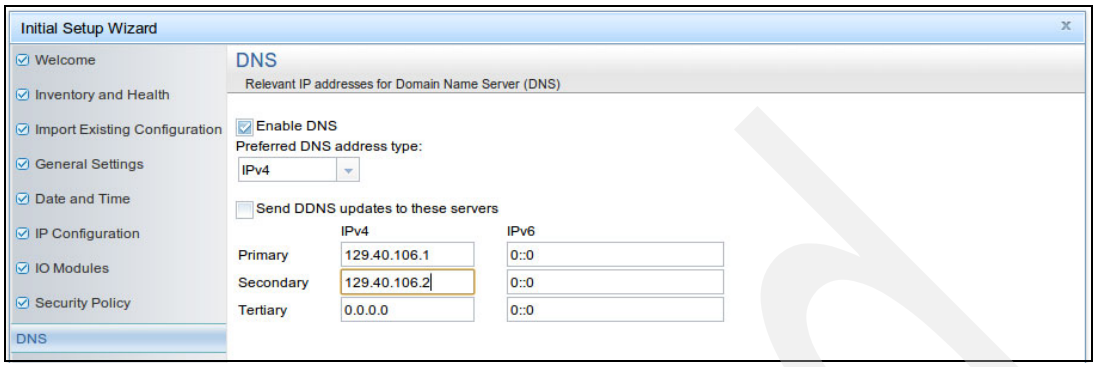

Figure 4-11 CMM Initial Setup Wizard: DNS IPv4

11. The Events Recipients page opens (Figure 4-12). This page is used to specify the email and SMTP addresses so that messages can be sent to the contacts. Click **Next**.

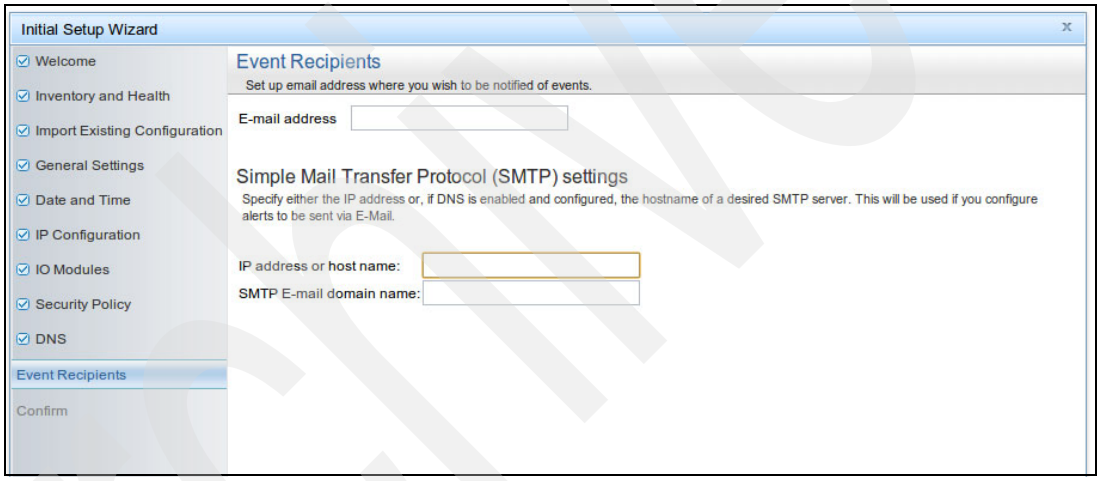

Figure 4-12 CMM Initial Setup Wizard: Specify Event Recipients

12. The Confirm page opens (Figure 4-13). Review the configuration. If the summary information is correct, click **Finish**. If necessary, click **Back** to return to a previous page or pages to make changes. Then, click **Next**.

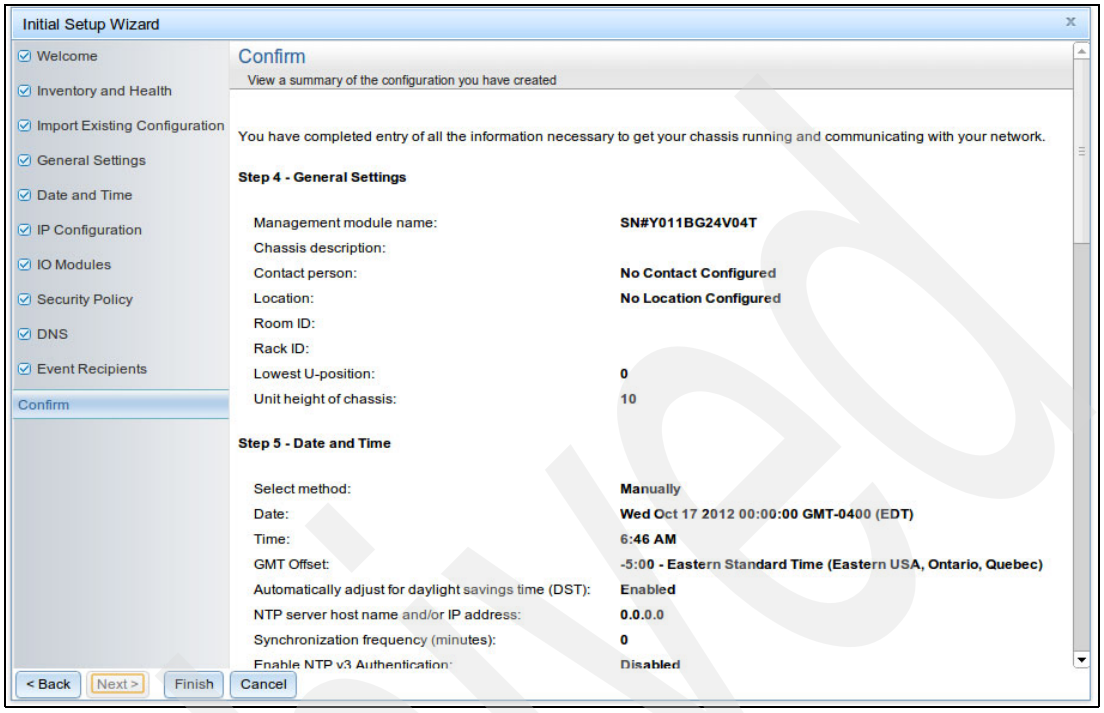

Figure 4-13 CMM Initial Setup Wizard: Confirm

13. In the pop-up window (Figure 4-14) click **OK** to continue with restart and to acknowledge that changes will take effect after the next restart of the CMM. You must wait for the CMM to restart, which can take up to 30 minutes.

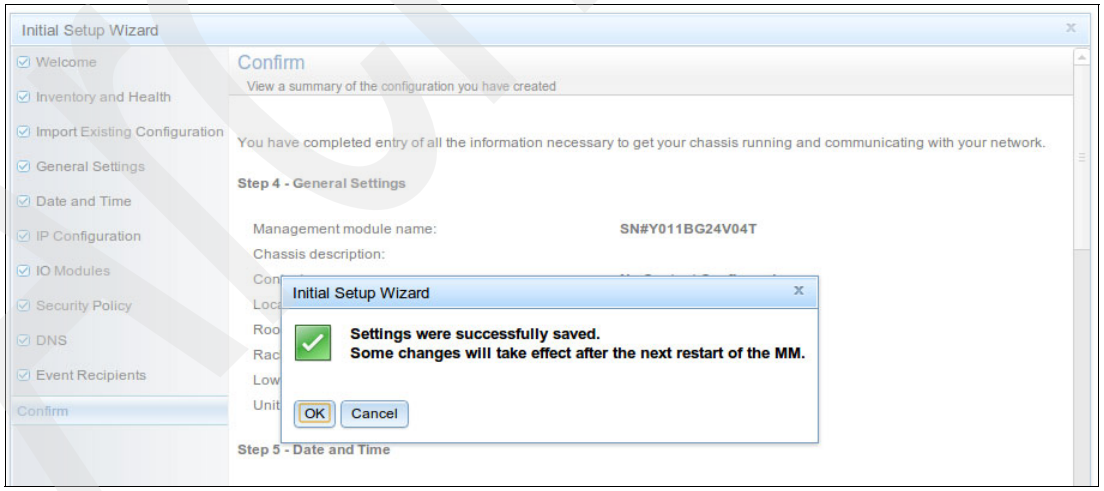

Figure 4-14 CMM Initial Setup Wizard: Settings are saved

After the CMM is configured and restarted, the next step in the PureFlex configuration process is the configuration of the FSM. Before doing that next step, read about setting up the network (4.3, “Networking setup” on page 27).

4.3 Networking setup

It is important to understand the complex networking setup that can be achieved with the Flex System hardware. The steps to achieve the setup can be divided into two logical parts:

- The first part describes the internal network between the chassis and the compute nodes. In addition, it describes the proper configuration of the Flex System Manager, which is vital for IBM SmartCloud Entry to provide robust cloud services on these compute nodes.
- The second part covers the configuration of the scalable switch elements.

4.3.1 Chassis and compute node setup

The Chassis Management Module (CMM) has its own physical connection to the external network. Other compute nodes within the chassis require a scalable switch element to connect to the external network. The configuration for these scalable switch elements is discussed in 4.3.2, “Configure scalable switch element” on page 28.

The Flex System Manager (FSM) adds more complexity. The FSM is delivered as a physical appliance that is similar to a standard compute node. However, the advanced management capabilities of the FSM require it to have slightly different hardware than a standard compute node. This is especially true in the network connections within a Flex System chassis.

Figure 4-15 is a greatly simplified logical diagram of the networking connections within the chassis.

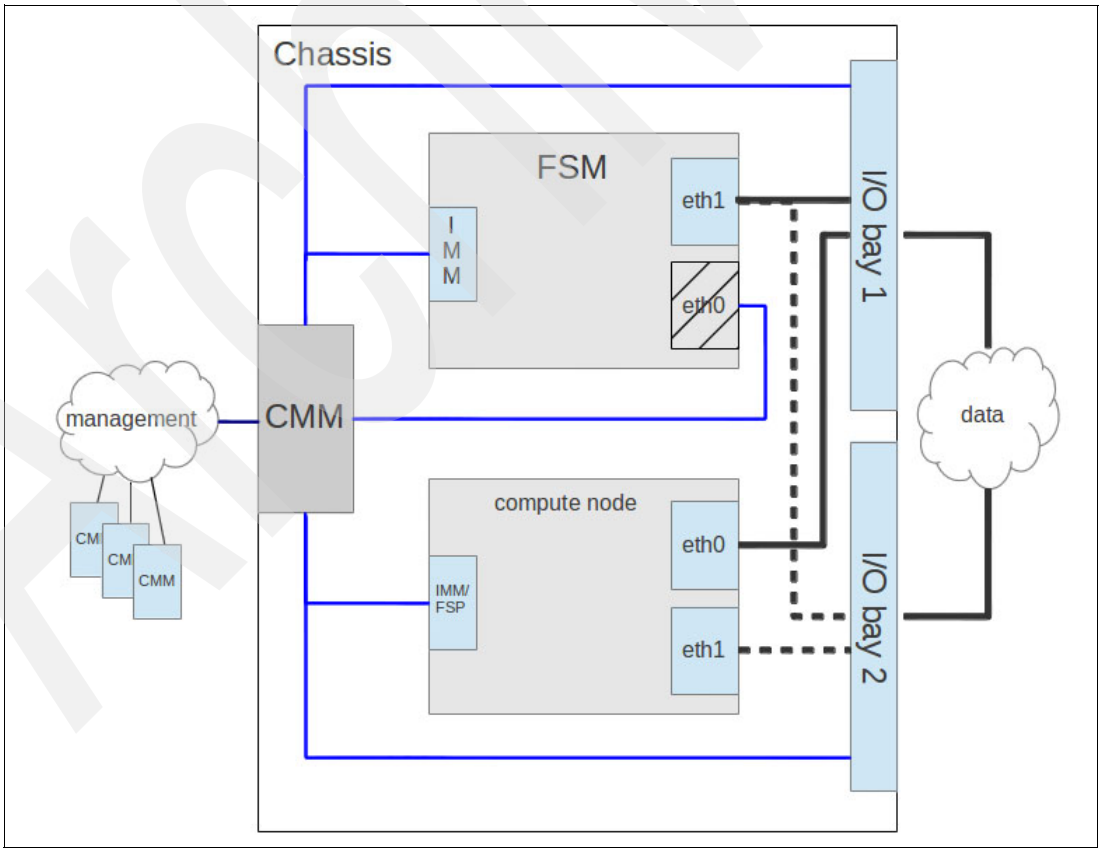

Figure 4-15 Network connections within the chassis

The figure shows the FSM's three main network connections. One of the network connections is through the FSM's IMM. This is identical to the IMMs of any compute node (Advanced System Management instead of IMM for Power nodes). The other two network connections (eth0, eth1) are differ slightly from other compute node connections. The connection represented at eth0 in Figure 4-15 on page 27 is part of the "management" network. In this context, the management network is an FSM construct that allows for the separation of management tasks (IMM configuration, ScSE configuration, storage subsystem configuration, and so on) from user tasks. This configuration can be selected during the initial setup of the FSM (discussed in 4.4, "FSM setup" on page 32). Regardless of the network topology chosen, select eth0 as the management network adapter and define a default gateway.

If separate networks are chosen for management and data traffic, eth1 must also be configured. The eth1 on the FSM is connected to an internal port on both I/O bays 1 and 2. Configure it with an IP address from the "data network." In the environment used for this book, there is only one GbE I/O module installed (in I/O bay 1) with one internal port licensed for each node. For completeness, the dashed lines indicate that there are internal connections to I/O bay 2, although there is no scalable switch element installed there.

If one network is selected for both management and data traffic, only eth0 is used. As stated previously, eth0 should be configured with the default gateway.

4.3.2 Configure scalable switch element

For each of the scalable switch elements installed in the chassis, a management IP needs to be set. This can be done by logging into the CMM and selecting **Chassis Management** → **Component IP Configuration**. This opens the Component IP Configuration for both I/O Modules and Compute Nodes.

Click the device name (see the row for Bay 1 in Figure 4-16) for each of the scalable switch elements to be configured. After clicking the device name, the window that opens is used to specify the Static IP Configuration for the switch. Click **Apply**.

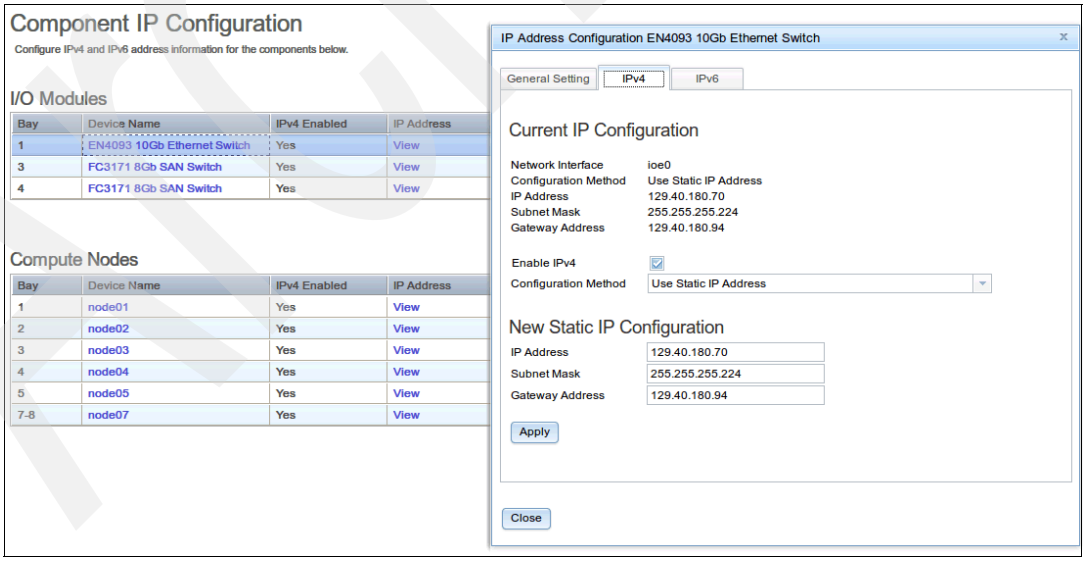

Figure 4-16 Networking: Specify Static IP information for scalable switch

Launch the I/O Module Console, as shown in Figure 4-17.

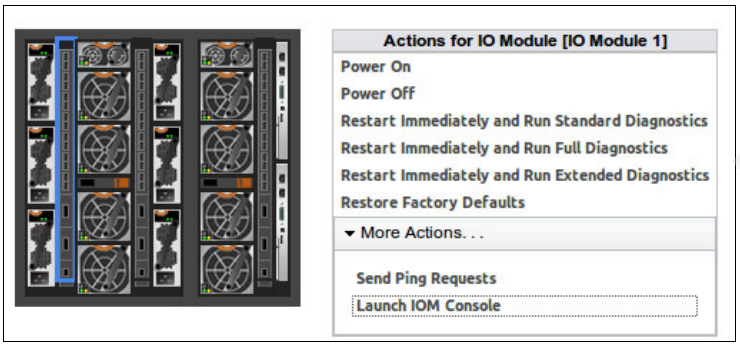

Figure 4-17 Networking: Launch IOM Console

A login panel opens (Figure 4-18). Enter the Username and Password and click **Submit**.

Figure 4-18 Networking: Log in to Scalable Switch

The Switch Dashboard opens (Figure 4-19). Click the **Configure** link near the IBM logo.

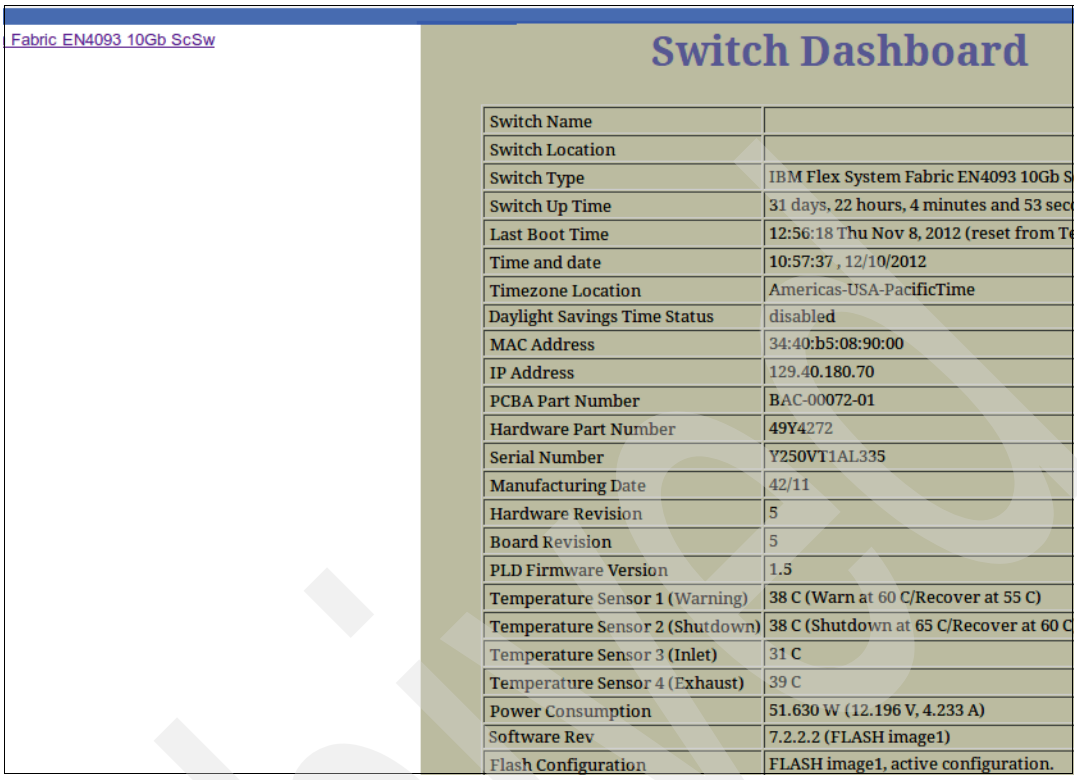

Figure 4-19 Networking: Scalable Switch Dashboard

The left side of the window populates with folder icons (Figure 4-20). Click the **Switch Ports** folder.

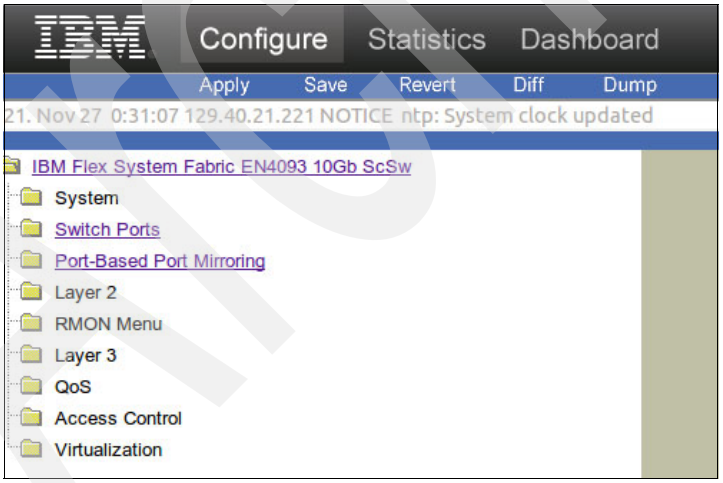

Figure 4-20 Networking: Scalable Switch Configure

A list of Switch Ports on the right side of the window is listed (Figure 4-21).

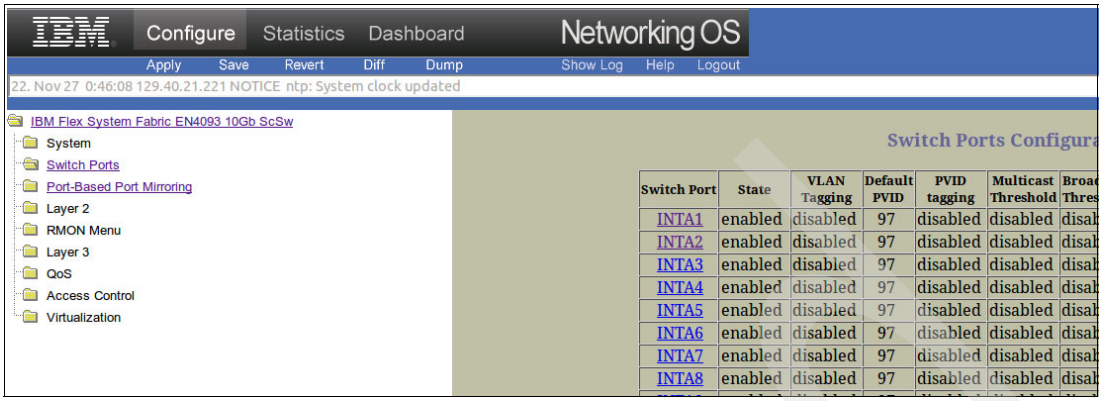

Figure 4-21 Networking: Scalable Switch Configure with Switch Ports folder opened

Click **INTA1** to view the Switch Port INTA1 Configuration. Then specify 97 for the **Default Port VLAN** as shown in Figure 4-22**.** This VLAN ID was provided by the networking staff for the environment used in this book.

Figure 4-22 Networking: Switch Port INTA1 Configuration

4.4 FSM setup

The CMM provided GUI Chassis Map can be used to access and configure the Flex System Manager's Integrated Management Module. Figure 4-23 depicts the physical hardware controlled by the CMM.

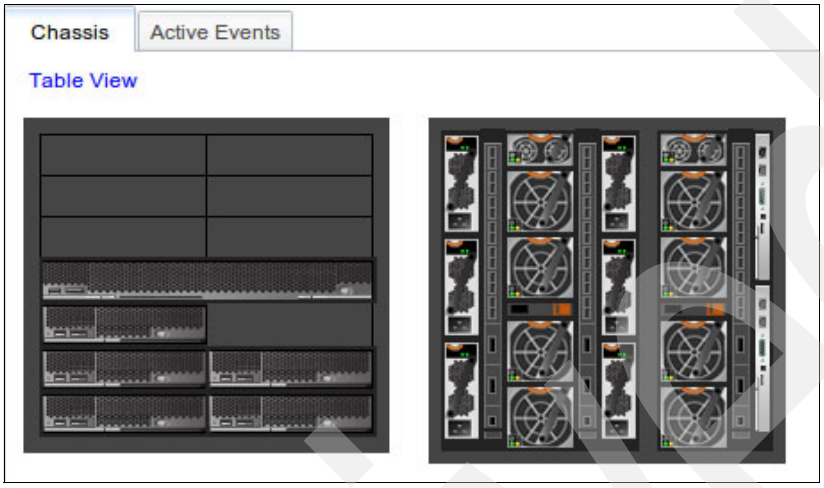

Figure 4-23 GUI Chassis Map

To configure the FSM, the IMM must first be configured for remote access. Tot do this, navigate to Chassis Management and select **Component IP Configuration** from the drop-down list (Figure 4-24).

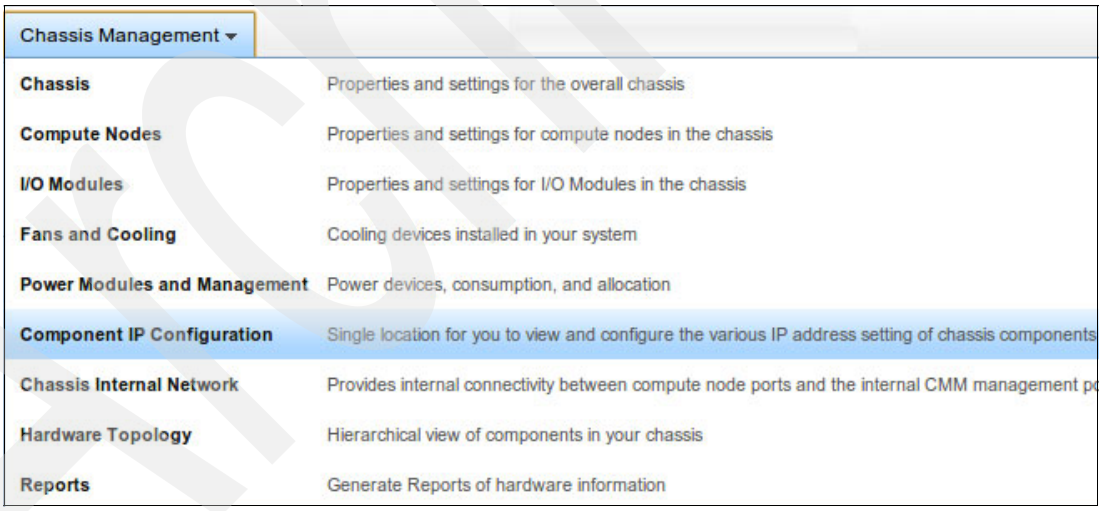

Figure 4-24 Chassis Management: Navigation and choices

I/O Modules and Compute Nodes are listed (Figure 4-25). The management network IPv4 information for the compute nodes and scalable switch elements can be set by clicking device names.

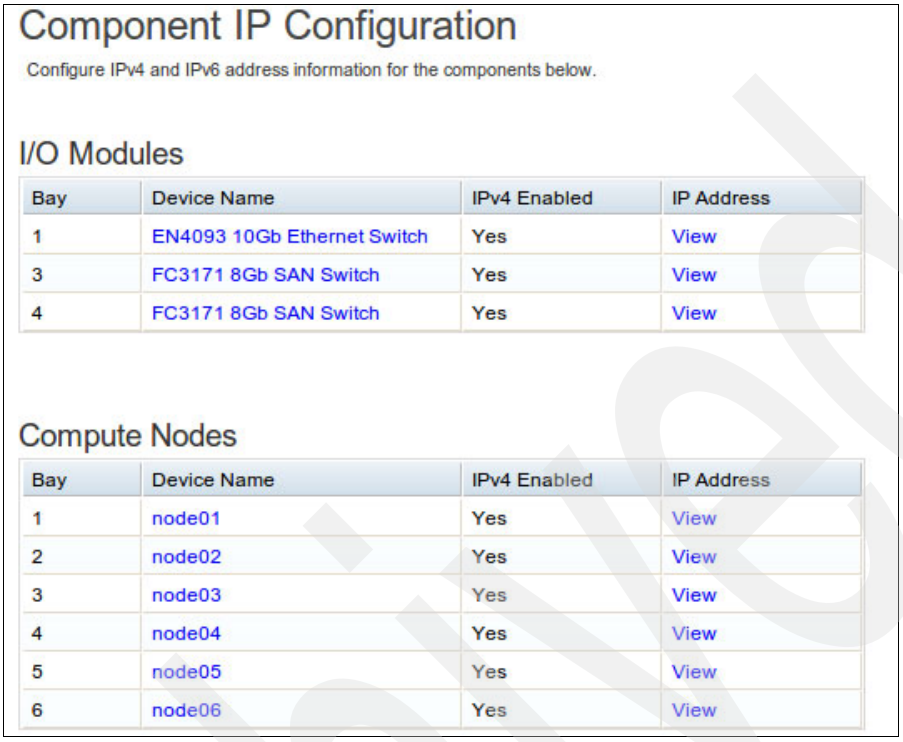

Component IP Configuration

Configure IPv4 and IPv6 address information for the components below.

I/O Modules

Bay	Device Name	IPv4 Enabled	IP Address
1	EN4093 10Gb Ethernet Switch	Yes	View
3	FC3171 8Gb SAN Switch	Yes	View
4	FC3171 8Gb SAN Switch	Yes	View

Compute Nodes

Bay	Device Name	IPv4 Enabled	IP Address
1	node01	Yes	View
2	node02	Yes	View
3	node03	Yes	View
4	node04	Yes	View
5	node05	Yes	View
6	node06	Yes	View

Figure 4-25 Component IP configuration

Click **node01** (in the Compute Nodes section), which accesses the Current IP Configuration for that node (Figure 4-26 on page 34).

New static IP information can be specified for node01, as shown in Figure 4-26. Provide any changes and click **Apply**.

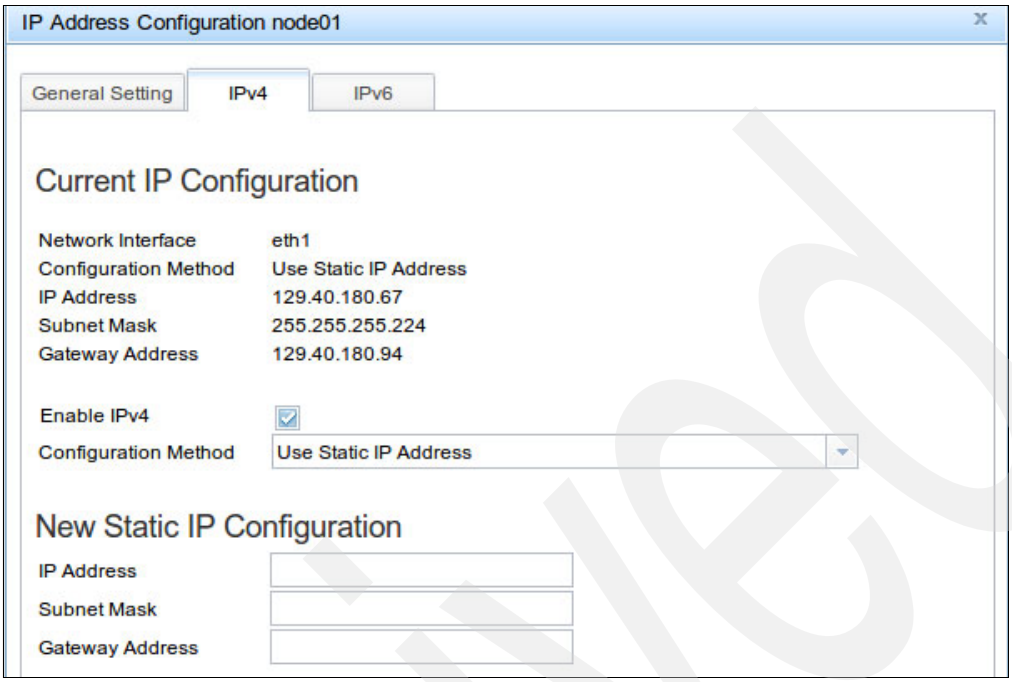

Figure 4-26 IPv4 tab: IP Address configuration

After the FSM's IMM has IPv4 network connections, a remote presence console can be opened. The FSM can be selected by clicking the image at the FSM's position (*the node in slot 1*) at the lower left of the graphical representation of the back of the chassis. The Actions for Compute Node [node01] menu is displayed to the right of the chassis (Figure 4-27). If the FSM is not already on, click **Power On.**

Expand the "More Actions" part of this menu, then click **Launch Compute Node Console**.

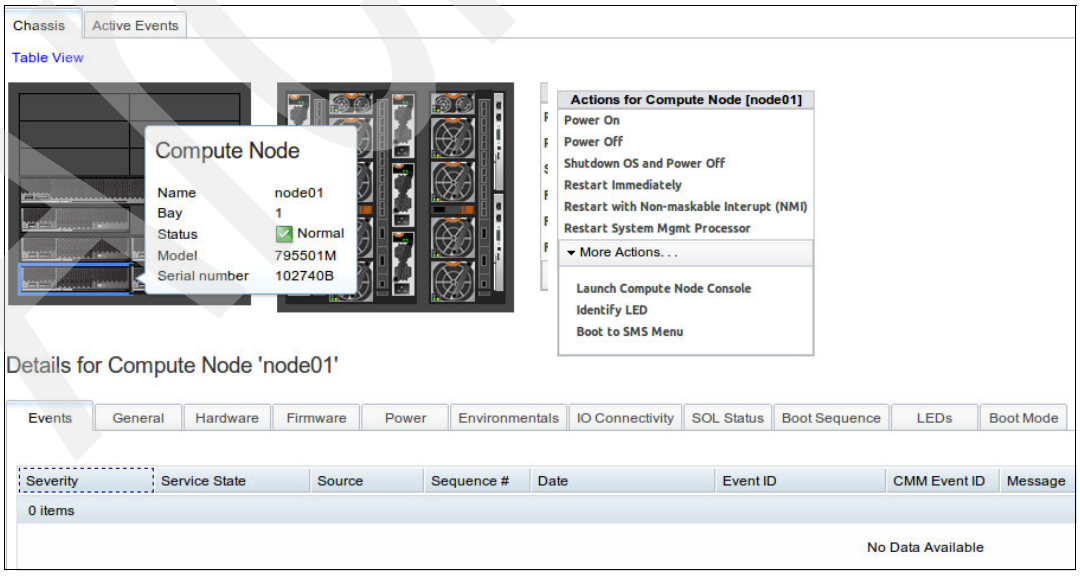

Figure 4-27 Launch Compute Node Console

In the Launch Node Console window that opens, enter the IP address of the FSM's IMM if it is not already there. Select **Remote Presence** for Context, and **HTTPS** for Protocol (Figure 4-28), and then click **Launch**.

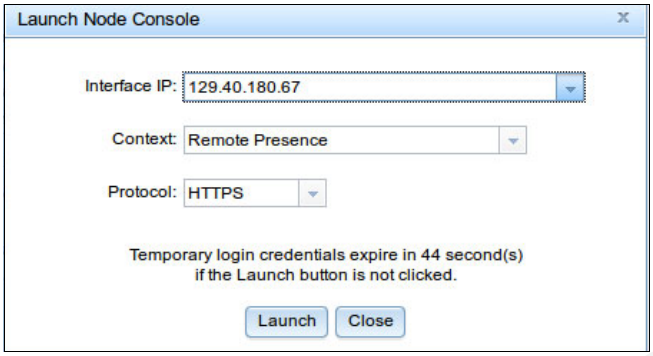

Figure 4-28 The Launch Node Console

The next window offers a choice between an ActiveX client or a Java client. Choose the preferred method by selecting the appropriate radio button, then click either **Start remote control**, depending on whether or not more than one person will be remote controlling the FSM.

When the remote control window opens, the IBM Flex Systems Manager License Agreement is displayed. If the FSM was recently powered on, the boot window opens prior to the License Agreement.

Review each of the licensing agreements by clicking the links in the left pane, then click **I agree** to advance to the next window. See Figure 4-29.

Figure 4-29 Flex System Manager: License Agreement

Wait for the FSM setup wizard to start. See Figure 4-30. This wizard is similar to the one used for the CMM Initial Setup. Specify the date, time, time zone, and optionally the host name or IP address of an NTP server.

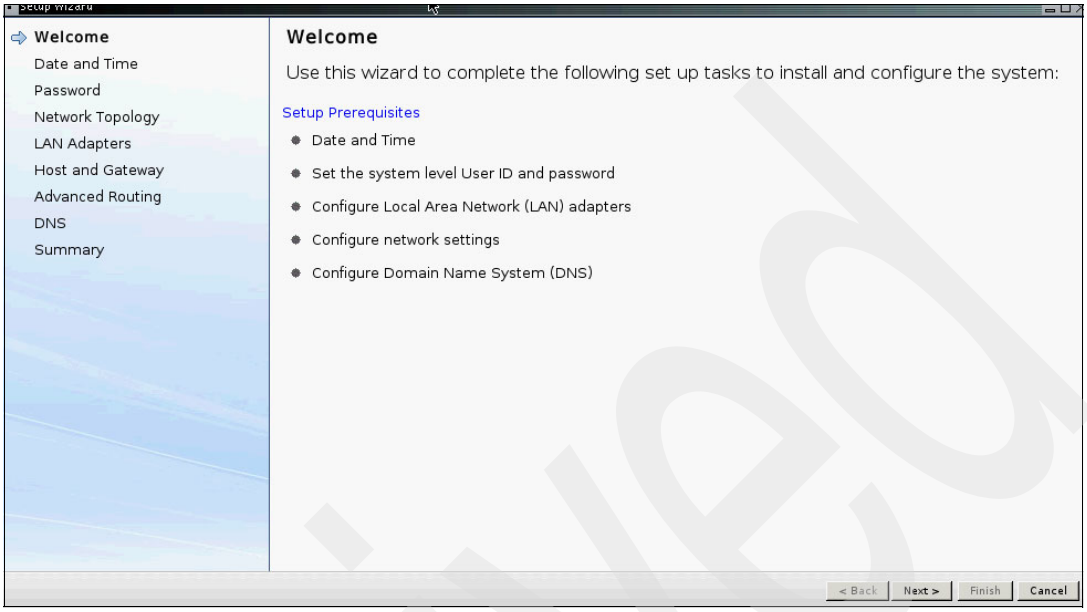

Figure 4-30 Flex System Manager: Welcome

On the next panel (Figure 4-31), enter the system-level user information.

Figure 4-31 Flex System Manager: System-Level User ID and Password

As discussed in Chapter 2, “Overview of example IBM Flex System” on page 7, the environment used to build this book had two networks (management and data). The figures that follow (Figure 4-32 on page 37 through Figure 4-37 on page 39) correspond to that environment.

The Network Topology window (Figure 4-32) has diagrams of the two possible network configurations. The FSM has two configurable network interfaces, eth0 and eth1. The eth0 interface is connected internally to the CMM and is accessible using an external jack. With this external jack, making a direct connection using a notebook becomes possible. This interface is intended to be used for the Management Network. This network is both internal to the chassis and external, because it connects to the top-of-rack switch. Click **Next**.

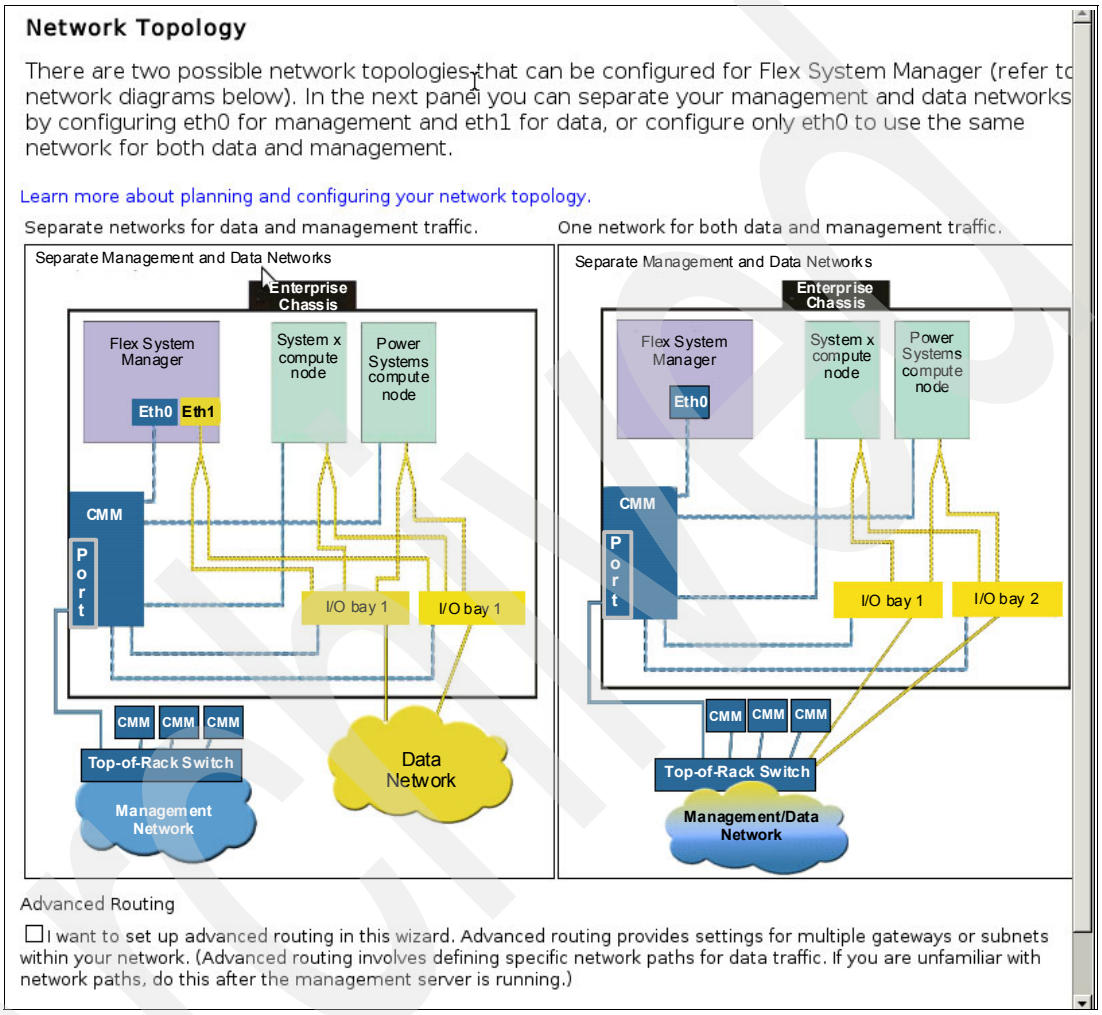

Figure 4-32 Flex System Manager: Network Topology

Select adapter eth0 (Figure 4-33). Then click **Next** to configure the Management Network.

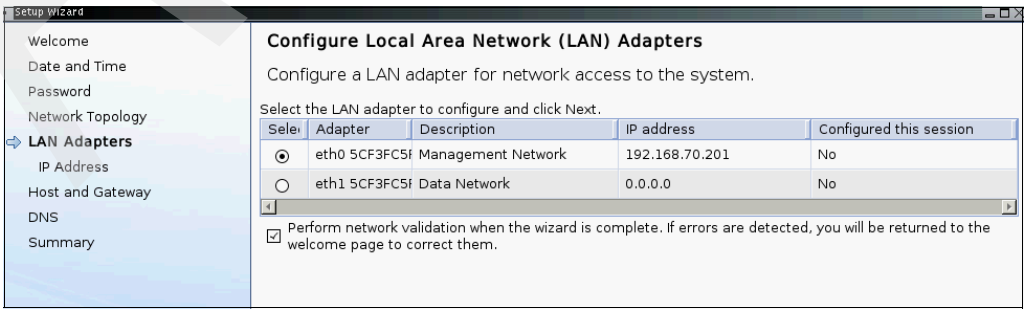

Figure 4-33 Flex System Manager: Configure Local Area Network (LAN) Adapters

On the panel shown in Figure 4-34, provide the IP address of the FSM Management Network adapter. This is the address that will be used to reach the FSM's graphical interface after the setup is complete. This address can optionally be DHCP if the environment supports that. Click **Next.**

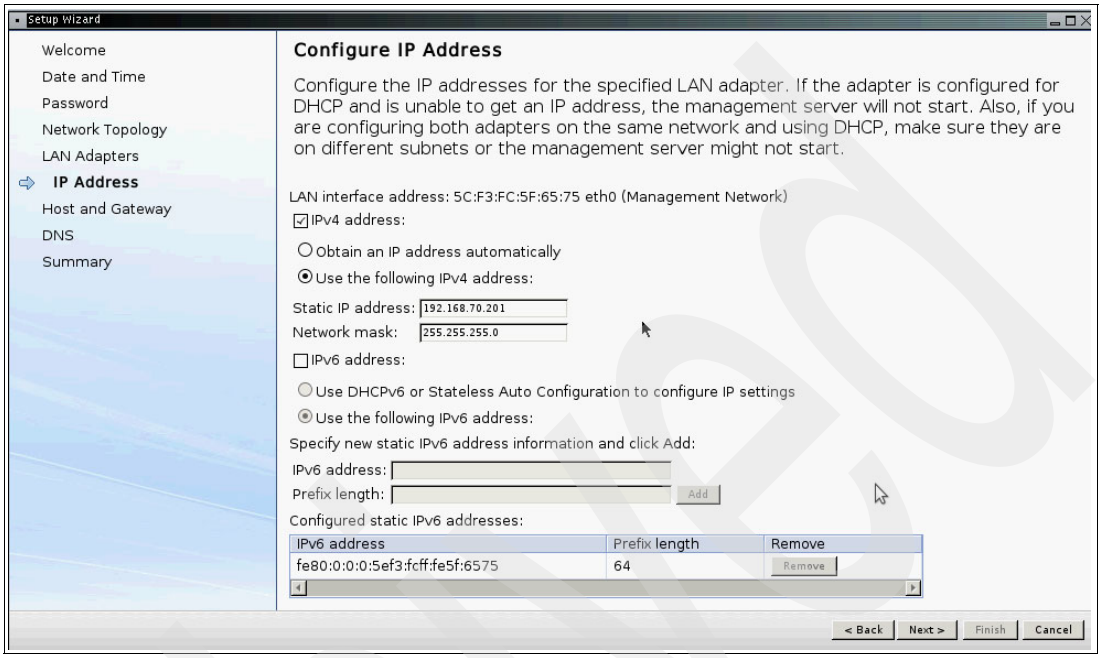

Figure 4-34 Flex System Manager: Configure IP Address

If the Data Network will be used, select **Yes, I want to configure another LAN adapter** (Figure 4-35). In the table that opens below the radio button, select the **eth1** adapter and click **Next.**

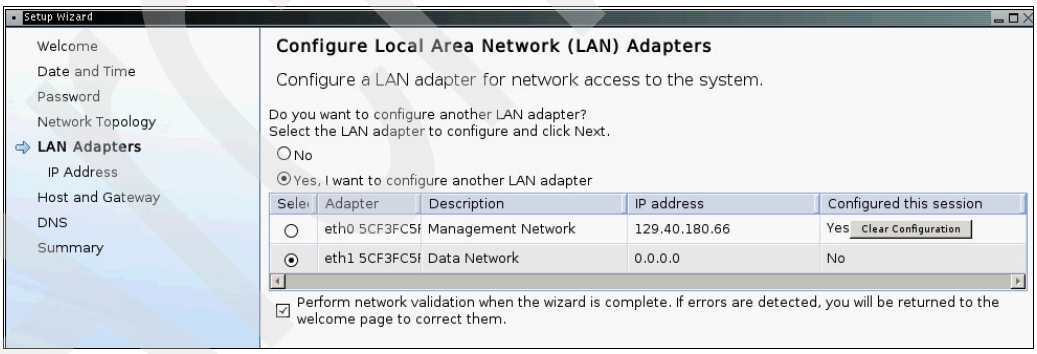

Figure 4-35 Flex System Manager: Configure another LAN adapter

The panel shown in Figure 4-36 shows how to configure the IP address of the FSM Data Network adapter. This is the address that will be used to reach the FSM. If enabled, the address may optionally be set by DHCP. Then, click **Next**.

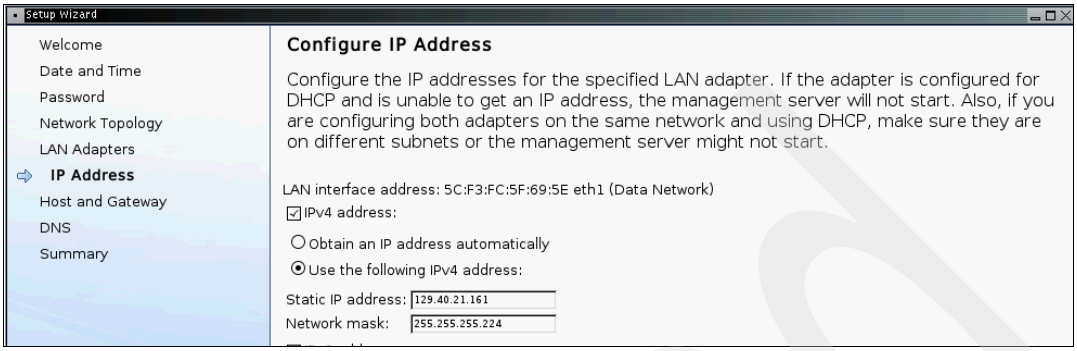

Figure 4-36 Flex System Manager: Configure IP Address for Data Network

The Configure LAN Adapters window opens again (Figure 4-37). Because there are no more LAN adapters, click the **No** radio button and then click **Next**.

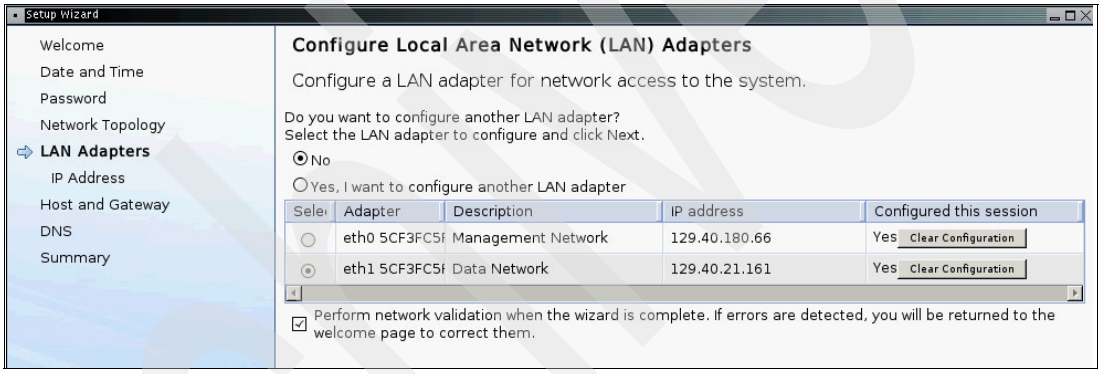

Figure 4-37 Flex System Manager: Completed LAN specifications

On the next panel (Figure 4-38), enter the short host name, domain name, and the default gateway address for the FSM. This is required. The data network is optional. Click **Next.**

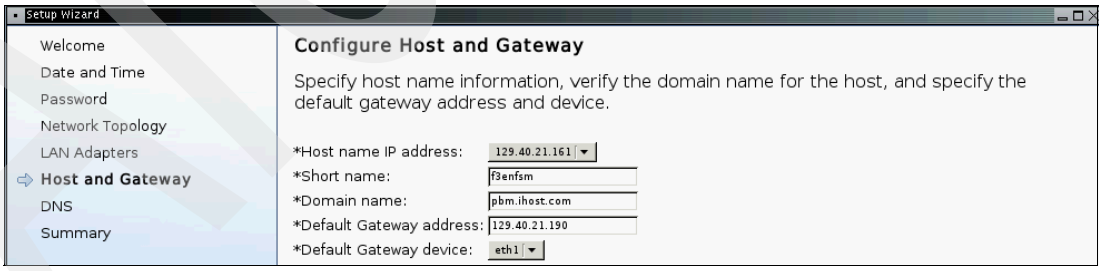

Figure 4-38 Setup Wizard: Configure Host and Gateway

Enable DNS services by selecting the appropriate check box (Figure 4-39). The FSM requires proper name resolution services. Ensure that the DNS services and suffix search order is correct. Click **Next** to continue.

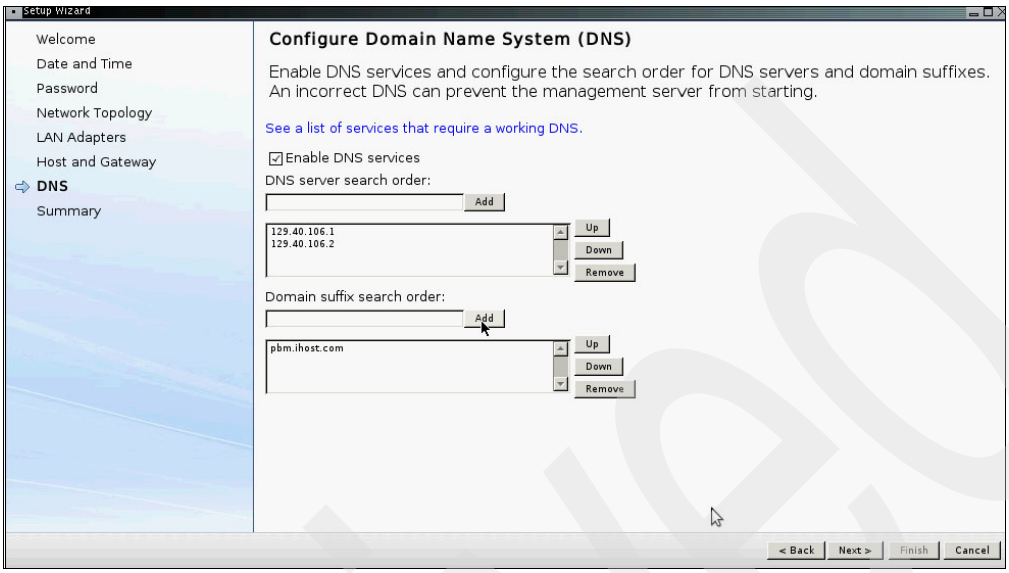

Figure 4-39 Setup Wizard: Configure Domain Name System

After the DNS specifications are complete, a Summary window opens (Figure 4-40). Verify that the settings are correct, and then click **Finish**.

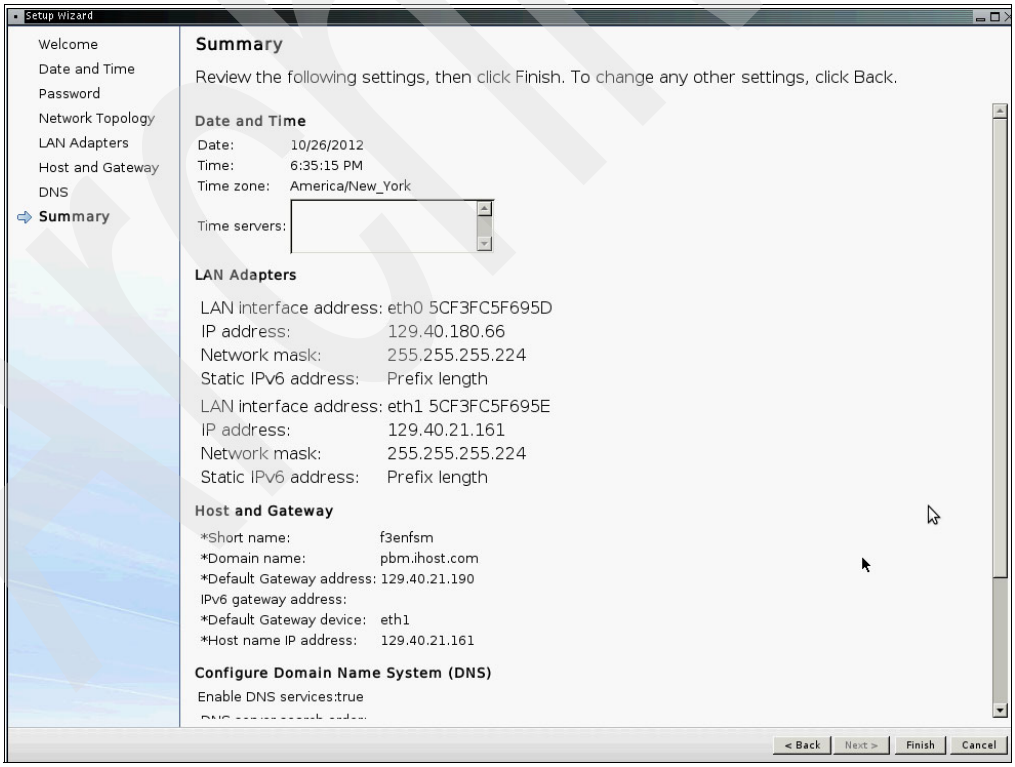

Figure 4-40 Setup Wizard: Summary

Figure 4-41 shows processing information, which indicates that the system setup is in progress.

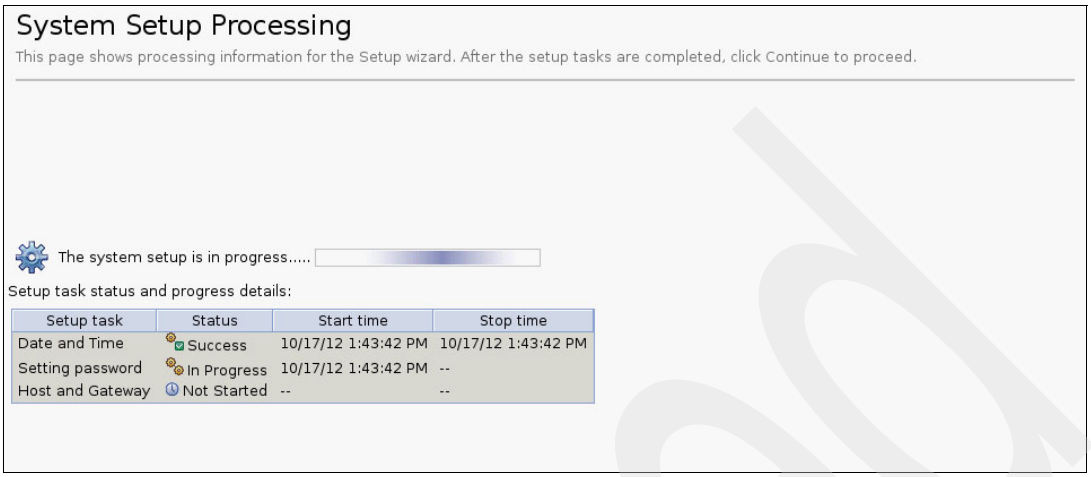

Figure 4-41 Setup Wizard: System setup processing

When the system is finished applying the settings, a pop-up window indicates that the setup tasks are completed (Figure 4-42). Click **Continue**.

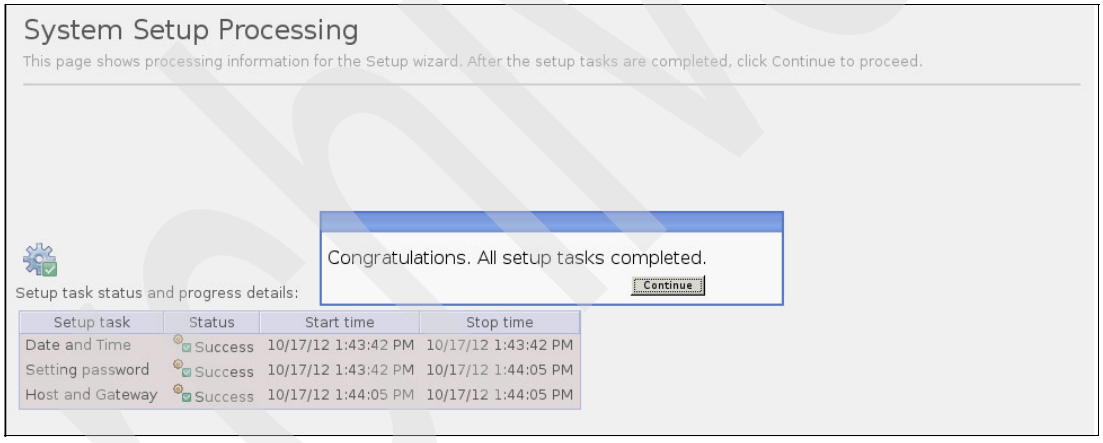

Figure 4-42 Setup Wizard: System setup processing complete

The network settings are applied. This involves a system restart which might last more than 30 minutes. A message about the server being restarted is displayed (Figure 4-43).

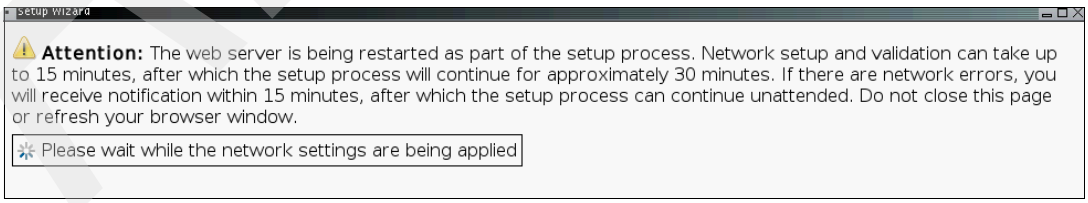

Figure 4-43 Setup Wizard: Network settings being applied

After the network settings are applied, a warning message states: `Do not reboot or power off the system` (Figure 4-44). The warning indicates that doing so before processing is complete might corrupt the installation, thus requiring a reinstallation.

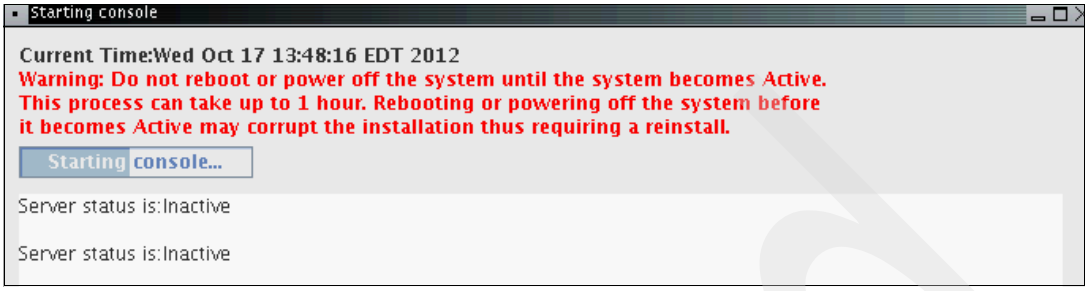

Figure 4-44 Starting console

At this point there is no longer a need to be physically attached to the CMM. Because the networking devices are now configured, the remainder of the setup can be completed from a browser on any machine that can access the management network.

To continue, point a browser to the following address; be sure to replace the IP address with the address entered previously for eth0 (as shown in Figure 4-36 on page 39):

`https://129.40.180.66:8422/ibm/console`

The login panel opens (Figure 4-45).

Figure 4-45 Log in to the IBM Flex System Manager

During the initial login to FSM, the Getting Started window opens (Figure 4-46). Click **Close** to continue to the main panel.

Figure 4-46 Flex System Manager Welcome: Getting Started

The main panel (Figure 4-47) provides links that are used to access and perform the following actions:

- Flex System Manager: Check and Update
- Flex System Manager Domain: Select Chassis to be Managed
- CMMs: Check and Update Firmware
- Compute Nodes: Check and Update Firmware
- I/O Modules: Check and Update Firmware

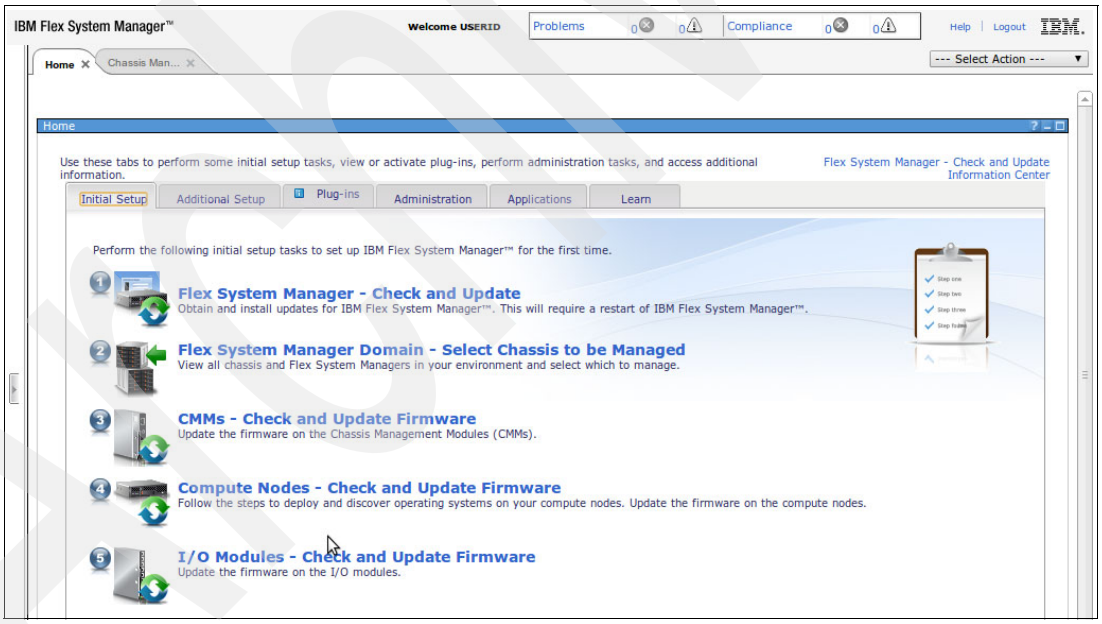

Figure 4-47 IBM Flex System Manager: Main panel

Click item 1, **Flex System Manager - Check and Update**.

Wait for the check for new updates to complete, as shown in Figure 4-48. This will work only if the FSM is able to access the Internet. After a short wait, any updates that are found are listed as indicated in the figure.

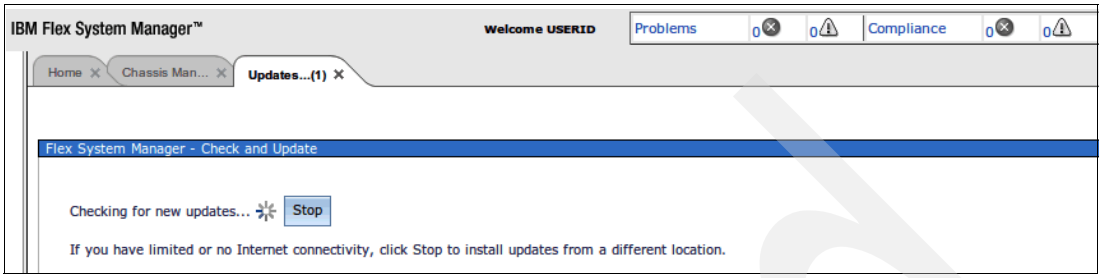

Figure 4-48 IBM Flex System Manager: Updates tab

If any available updates are listed (Figure 4-49) click **Download and Install**.

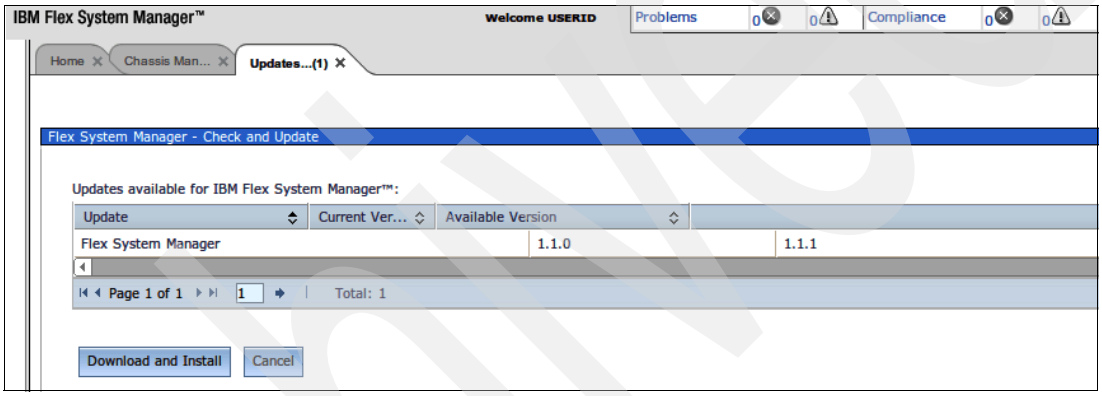

Figure 4-49 Choosing Download and Install: IBM FSM Updates Panel

A warning message recommends that you back up critical data (Figure 4-50). At this point, you can disregard this warning, because only minimal configuration has been done. Click **OK** (without doing a backup in this case).

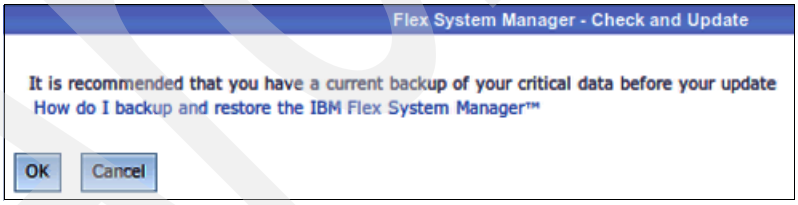

Figure 4-50 Warning to backup critical data

the **Job name and Schedule** panel opens (Figure 4-51). Select the **Run Now** radio button i and then click **OK** to submit the job.

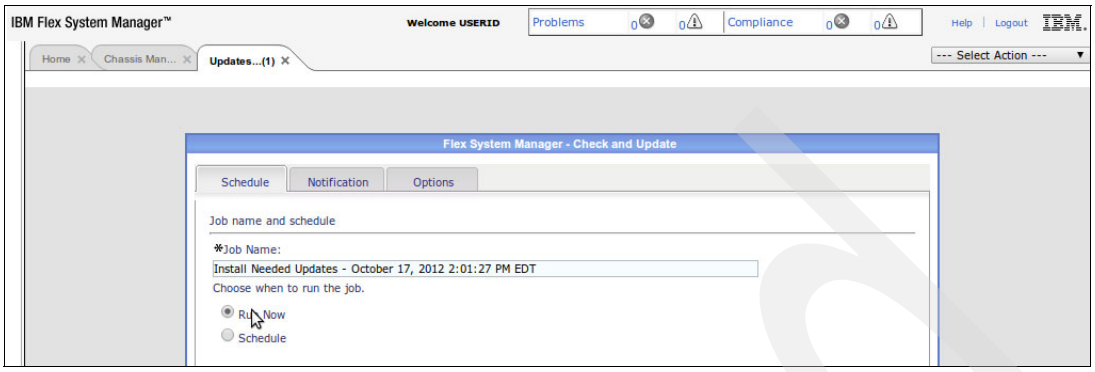

Figure 4-51 FSM Check and Update: Run Now

A message indicates that a job has been created and started successfully (Figure 4-52). Click **Display Properties** to observe the job progress.

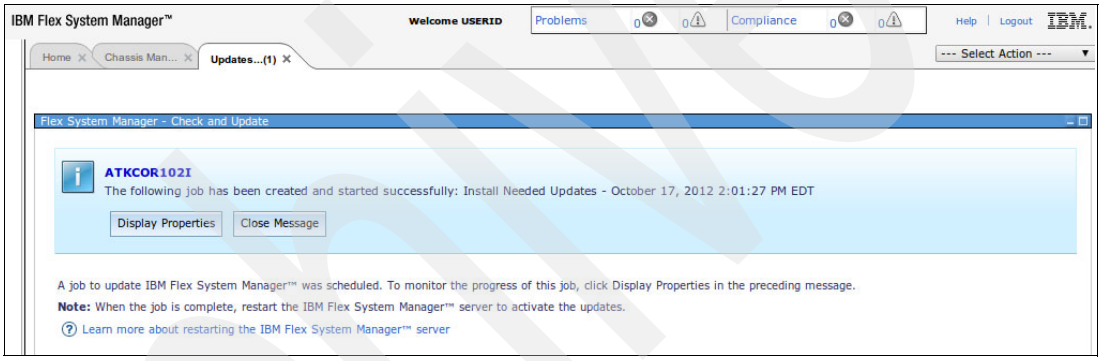

Figure 4-52 Job has been created and started for updates

Click the **Logs** tab (Figure 4-53). As the job progresses, this window continues to update.

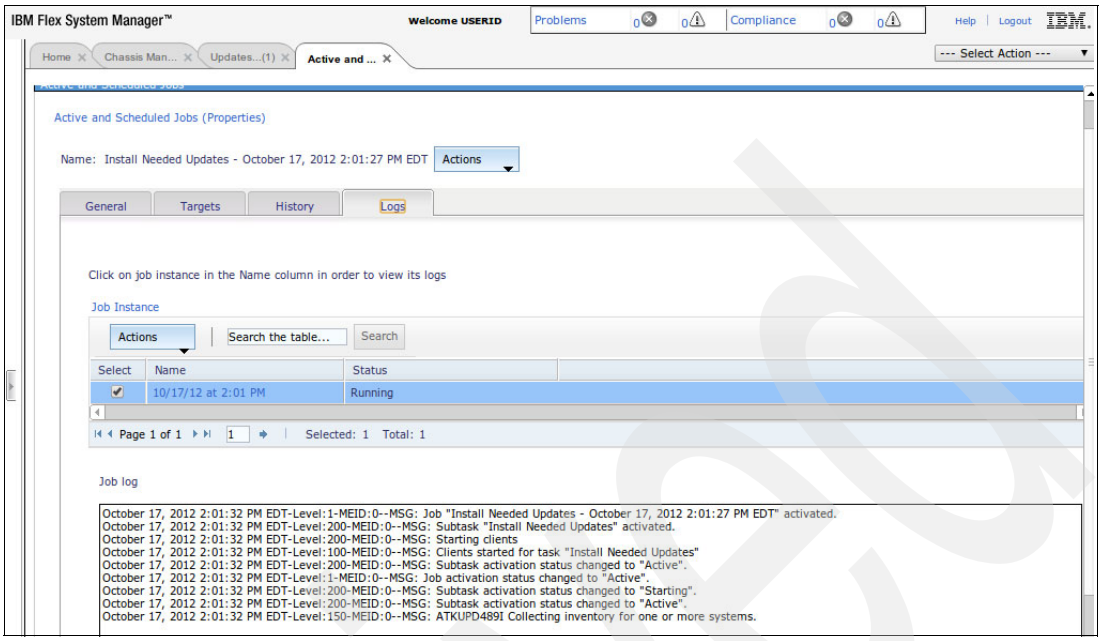

Figure 4-53 Update job logs

As indicated by the message in Figure 4-54, after the updates are installed Flex System Manager must be restarted. Depending on the load on the IBM Update Server, the updates might or might not complete with errors. If errors occur, wait a few moments and try the updates again.

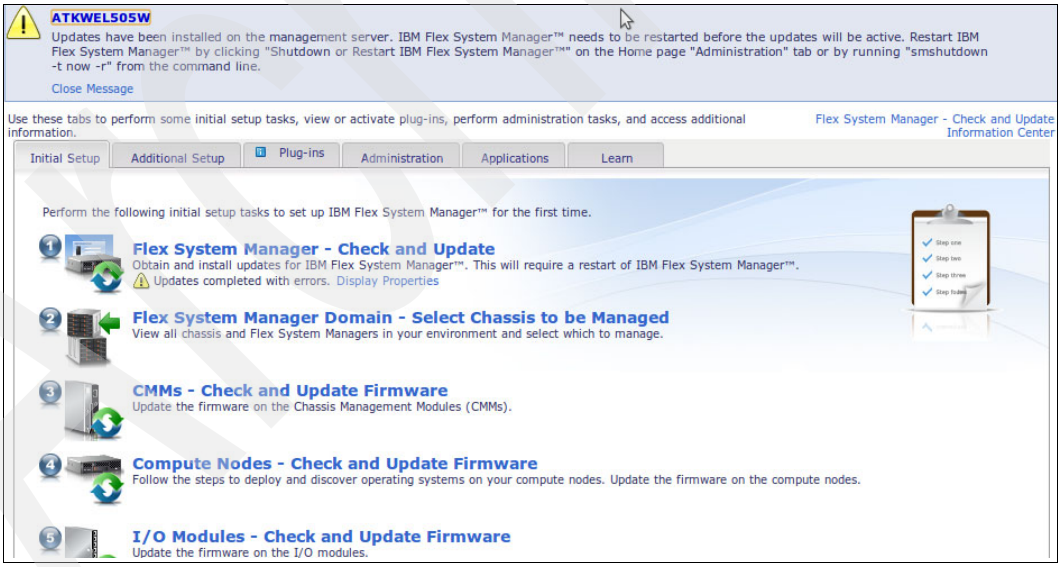

Figure 4-54 Updates have been installed

As Figure 4-55 shows, a check mark (in green box) appears next to the step 1 icon when Flex System Manager has been successfully upgraded.

The next step is to add the Flex System Enterprise chassis to the FSM's inventory. To initiate this task, click step 2, **Flex System Manager Domain - Select Chassis to be Managed**.

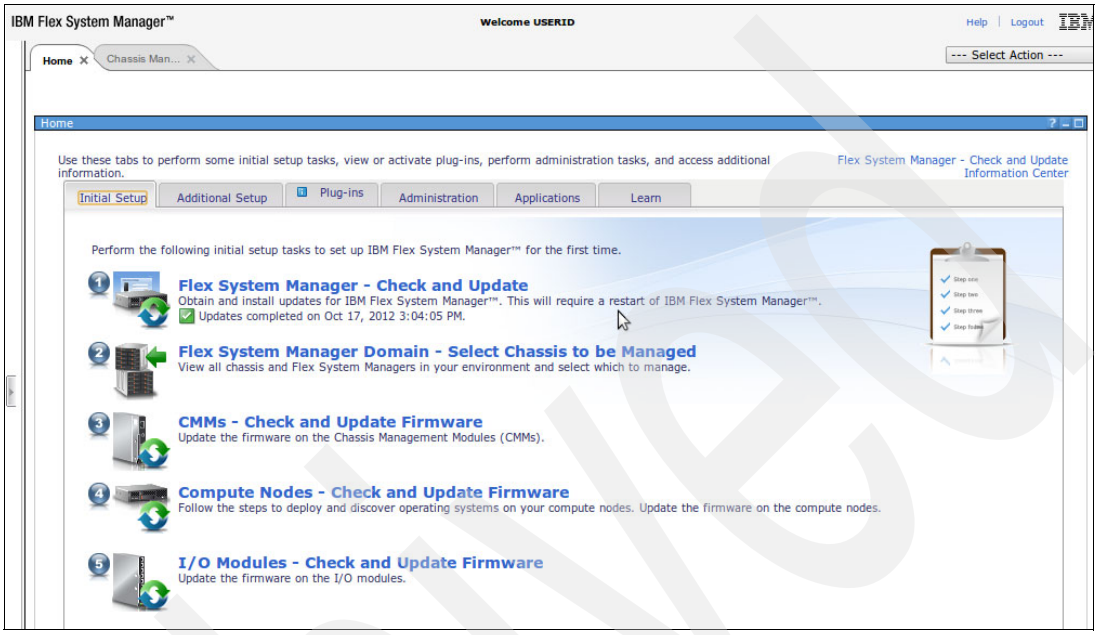

Figure 4-55 Green Check Icon: Upgrade was successful

A list of all discovered chassis is displayed (Figure 4-56). This figure shows the selection (note the check mark in the Selection column) of 120.40.180.65. Click **Manage** (to manage selected chassis).

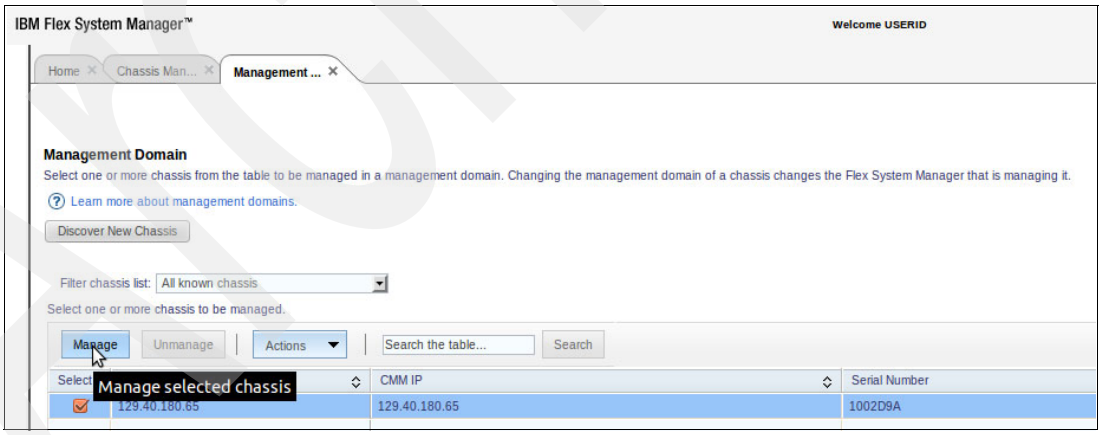

Figure 4-56 Discovered Chassis

A Manage Chassis confirmation panel opens (Figure 4-57). This panel indicates that the local Flex System Manager (in this case f3enfsm) will manage the chassis selected and that the management environment will not be effective for another 20 minutes. Click **Manage** to confirm this management request.

s

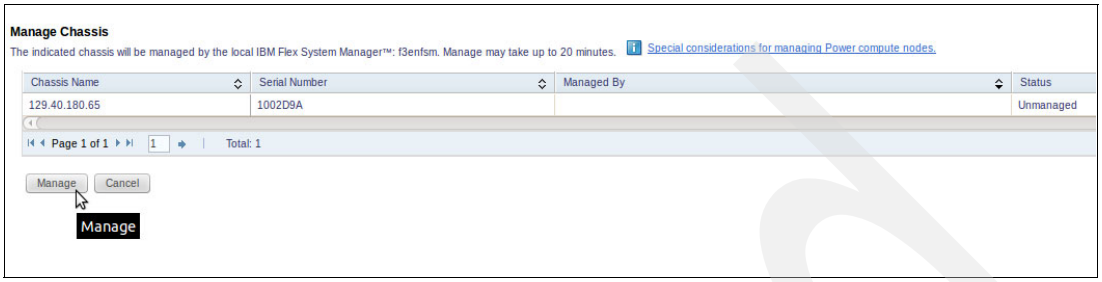

Figure 4-57 Request Manage Chassis

The Request access panel opens (Figure 4-58), which prompts for the user ID and password to authenticate access to, in this case, the selected chassis at 129.40.180.65. Provide the appropriate User ID and Password, and click **OK**.

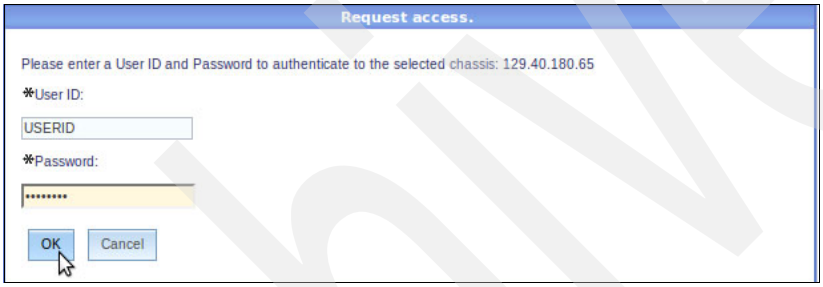

Figure 4-58 Request Access Panel

Wait for the chassis discovery operation to complete. The icon in the Status column (Figure 4-59) changes from Processing to Success. It might take as long as 20 minutes for this status change as noted on the panel. Click **Done** when the status indicates Success.

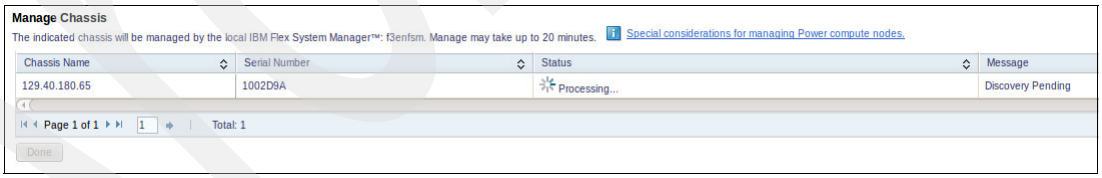

Figure 4-59 Chassis Discovery in process

Now that the FSM setup is complete, proceed to 4.5, “Storage setup” on page 48.

4.5 Storage setup

In our environment, we had an IBM Storwize V7000 storage subsystem, and also internal hard drives on each of the installed compute nodes. Most storage configuration covered in this book involves the V7000. This section involves configuring external storage only.

For the compute nodes to communicate with the V7000, a SAN fabric must be created. A SAN fabric is a collection of ports on one or more compute nodes, Fibre Channel switches

(scalable switch elements in I/O bays 3 and 4) and one or more storage subsytems (V7000). After the SAN fabric is created, SAN zones can be defined on the SAN fabric.

SAN zoning is a way to assign parts of a storage subsystem's available capacity to one or more hosts, while restricting access to others. Zoning can also provide redundant paths between hosts and subsystems to insure against failures of individual components. By using ports on both scalable switch elements (ScSE) and ports on both V7000 nodes, there will still be full access between hosts and storage in the event that one of them fails.

All compute nodes in a SAN fabric are represented by worldwide node names (WWNN). The ports on each mezzanine card in each compute node are represented by worldwide port names (WWPN), as are each port in the ScSE and the ports on the storage subsystem.

WWNNs and WWPNs are unique 16-hexadecimal digit identifiers. The WWPNs are grouped together in various combinations to form the zones that are defined on the ScSE. Aliases for these WWPNs can be defined on the ScSE to more easily configure the zoning.

This chapter covers the setup and configuration of all the components required to build a fully working SAN fabric:

- V7000 setup
- Scalable switch element setup for storage
- FSM setup for storage management

It is important to understand some key concepts before continuing to the listed steps.

The following examples use the half-wide POWER node, the scalable switch elements in bays 3 and 4, and the V7000 to showcase the complexity of a SAN fabric on IBM Flex Systems hardware.

Two Fibre Channel (FC) ports appear on the Virtual I/O Server (VIOS, or VIO Server) that is installed on the internal hard disk of the half-wide POWER compute node. Example 4-1 shows a VIO command that returns the device names for each of these ports: fcs0 and fcs1.

Example 4-1 Device names for ports

```
$ lsdev -dev fc*
name             status      description
fcs0             Available   Dual Port 8Gb FC Mezzanine Card (7710322577107501)
fcs1             Available   Dual Port 8Gb FC Mezzanine Card (7710322577107501)
```

The WWPNs for ports fcs0 and fcs1are listed in the Virtual Server's inventory on the FSM, as shown in Figure 4-60.

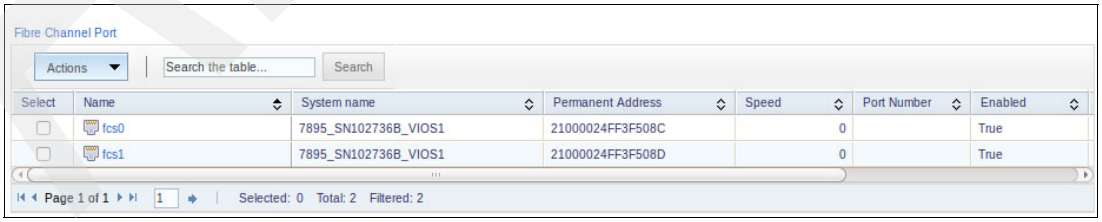

Figure 4-60 Fibre Channel Ports shown on Virtual Server's Inventory

Each of these port's WWPNs is accessible to the ScSE to which that port is connected. Unlike traditional infrastructures, these links are not using physical cables, but they are using the chassis internal midplane. The ports on this card are connected to separate I/O bays. Port 0, which has the name fcs0, is connected to an internal port on the scalable switch element that

is installed in I/O bay 3. Port 1, which has the name fcs1, is connected to an internal port on the scalable switch element that is installed in I/O bay 4. Only one of the two WWPNs appears on the first scalable switch element. In this case fcs0, which has a WWPN of 21000024FF3F508C, will show in the topology for the ScSE in bay 3.

The **show topology** command (Example 4-2) can be run on this ScSE (in I/O bay 3) command-line interface (CLI) to show the list of WWPNs detected.

Log in to the CLI by using SSH, as follows, and then enter the appropriate password:

```
ssh USERID@129.40.180.73
```

Example 4-2 Show topology

```
IBM8Gb: USERID> show topology
Unique ID Key
  -------------
  A = ALPA,  D = Domain ID,  P = Port ID
         Loc   Local                   Rem  Remote                  Unique
  Port    Type  PortWWN                 Type NodeWWN                 ID
  ----    ----  -------                 ---- -------                 ------
  Ext1:0  F     20:00:00:c0:dd:24:16:2d N    50:05:07:68:02:00:d3:72 020000   P
  Ext2:15 F     20:0f:00:c0:dd:24:16:2d N    50:05:07:68:02:00:d3:73 020f00   P
  Bay2    F     20:02:00:c0:dd:24:16:2d N    20:00:00:24:ff:3f:50:8c 020200   P
  Bay3    F     20:03:00:c0:dd:24:16:2d N    20:00:00:24:ff:3f:4f:5c 020300   P
  Bay4    F     20:04:00:c0:dd:24:16:2d N    20:00:00:24:ff:48:91:a6 020400   P
  Bay5    F     20:05:00:c0:dd:24:16:2d N    20:00:00:24:ff:20:49:d8 020500   P
  Bay7    F     20:07:00:c0:dd:24:16:2d N    20:00:00:24:ff:3f:4e:c2 020700   P
  Bay8    F     20:08:00:c0:dd:24:16:2d N    20:00:00:24:ff:3f:4f:02 020800   P
```

Example 4-2 shows two external ports Ext1:0 and Ext2:15. The Remote NodeWWN listed (50:05:07:68:02:00:d3:72 and 50:05:07:68:02:00:d3:73) belong to the two Fibre Channel devices on the V7000: one on the control enclosure and the other on the expansion enclosure.

The WWPN of port 0 (fcs0) on the mezzanine card shows up as the Remote NodeWWN for bay 2.

Both of the devices connected to the external ports with the listed WWNs have two additional WWPNs associated with them, one for each port on the device. These are the WWPNs that are grouped with the WWPNs of the compute node's Fibre Channel (FC) ports to define the zones.

This zoning configuration has a "single initiator" (host port) per zone, meaning that there will be as many zones as there are FC ports on all of the compute nodes. Example 4-3 shows the two zones containing the FC ports on the single-wide Power node in bay 2. The one zone for the compute node in bay 2 on the scalable switch element (ScSE) is installed in I/O bay 3 (Example 4-3).

Example 4-3 Zone information

```
SN102736B_VIOS1
          SN102736B_HBA1_PORT1 (Alias)
                    21:00:00:24:ff:3f:50:8c
          SN78N1XKA_NODE1_PORT1 (Alias)
                    50:05:07:68:02:10:d3:72
          SN78N1XKA_NODE2_PORT1 (Alias)
```

```
                    50:05:07:68:02:10:d3:73
```

The complementary zone on the scalable switch element is in I/O bay 4 (Example 4-4).

Example 4-4 Zone information

```
SN1O2736B_VIOS1
          SN78N1XKA_NODE2_PORT2 (Alias)
                    50:05:07:68:02:20:d3:73
          SN78N1XKA_NODE1_PORT2 (Alias)
                    50:05:07:68:02:20:d3:72
          SN1O2736B_HBA1_PORT2 (Alias)
                    21:00:00:24:ff:3f:50:8d
```

Notice that the WWPNs for the V7000 (known here as SN78N1XKA) are similar to the WWNs for the external ports listed by the **`show topology`** command. For the WWPNs, the third-to-last pair of digits are incremented by 10 as in the following examples:

- The WWN for the FC device on the control enclosure (known as NODE1):

  `50:05:07:68:02:00:d3:72`

- The WWPN for PORT1 of this device:

  `50:05:07:68:02:10:d3:72`

- The WWPN for PORT2 of this device:

  `50:05:07:68:02:20:d3:72`

- Similarly, the WWPNs for the remaining two ports (NODE2) on the V7000:

  `50:05:07:68:02:10:d3:73`
  `50:05:07:68:02:20:d3:73.`

The zoning for the rest of the compute nodes is constructed in the same way. One zone is created for each host FC port by grouping it with two of the four WWPNs on the V7000. The complete zoning configuration can be shown using the **`zoning list`** command on either FC switch's CLI. This is discussed further in 4.5.2, “Scalable switch element setup for storage” on page 52.

4.5.1 V7000 setup

One of the core characteristics of cloud computing is the pooling of resources. To this effect, the V7000 pools the internal physical hard drives. This is accomplished using RAID sets of the physical disk drives known as managed disks (MDisks). MDisks are assigned to storage pools. Storage pools can then be further subdivided into storage volumes and presented to the hosts. Because of this virtualization, storage volumes and their capacity can be dynamically altered, transparent to the host. In addition, this virtualization enables the V7000 to support advanced features such as thin provisioning, cloning, and snapshot capabilities.

Note: From the perspective of the V7000, a host is a named entity that is associated with one or more World Wide Port Name.

The V7000 might need to be configured so that the disks it contains can be used by IBM SmartCloud Entry, depending on how it was delivered. The first step in this configuration is setting up the IP networking information. One way to accomplish this is by inserting a blank, formatted USB flash drive into the USB slot on the back of the control enclosure. The controller that controls the V7000 is the one that reads 00 on the LED display on the front.

The firmware will detect the USB device and then write a text file containing information about the system. This file may then be edited by adding the desired IP networking information, and then re-inserting the drive into the control enclosure. This causes the V7000 to reconfigure itself with the new networking information. For more details about this process, see the following sources of information:

- IBM information center for the IBM Storwize V7000:

  http://pic.dhe.ibm.com/infocenter/storwize/ic/index.jsp

- Search for V7000 and PureFlex on the IBM Redbooks website:

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

After the IP address information is set on the V7000, the browser-based administration GUI can be accessed. Using the GUI, the pools, volumes and MDisks can be defined and also host definitions and subsequent mapping of volumes to those hosts. As with other sections in this book, depending on how the hardware was delivered, the amount of setup needed can vary. Therefore, this book does not cover the initial setup of the V7000. See Chapter 2, “Overview of example IBM Flex System” on page 7 for an overview of the setup on which this book is based.

The V7000 used in the discussion and examples initially has five volumes, defined as shown in Figure 4-61. These volumes are used for the Virtual I/O Servers and x86 environment.

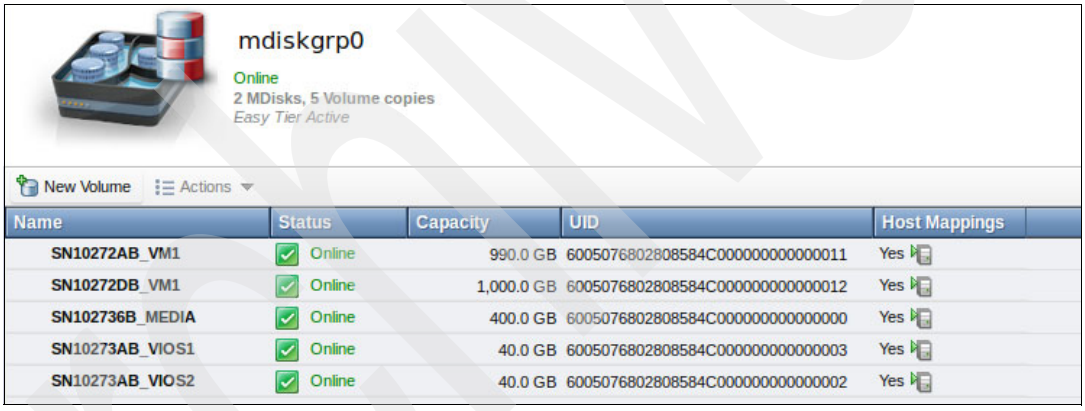

Figure 4-61 Initial volumes

An example of creating volumes and mapping volumes to hosts is in 4.6.2, “Create volume on V7000” on page 68.

4.5.2 Scalable switch element setup for storage

Configuration for the scalable switch element that connects the elements in the storage fabric is similar to the setup for the Ethernet ScSE. Section 4.3.2, “Configure scalable switch element” on page 28 shows how to set up an IP for the management side of the ScSE. It is the same process for the Fibre Channel ScSE. After the IP information is configured, the zoning can be created.

Figure 4-62 on page 53 shows the user interface when accessed over HTTP. Enter the Login Name and Password for this ScSE, and then click **Add Fabric**.

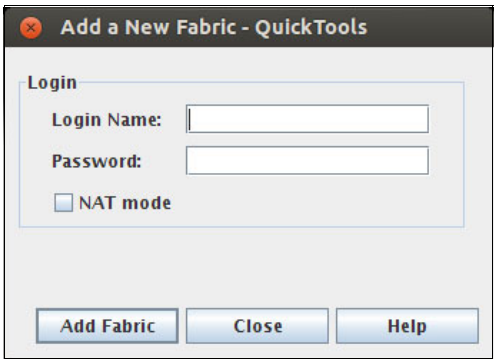

Figure 4-62 Add a New Fabric: QuickTools

After logged in to the ScSE, the scalable switch element's fabric is displayed, (Figure 4-63).

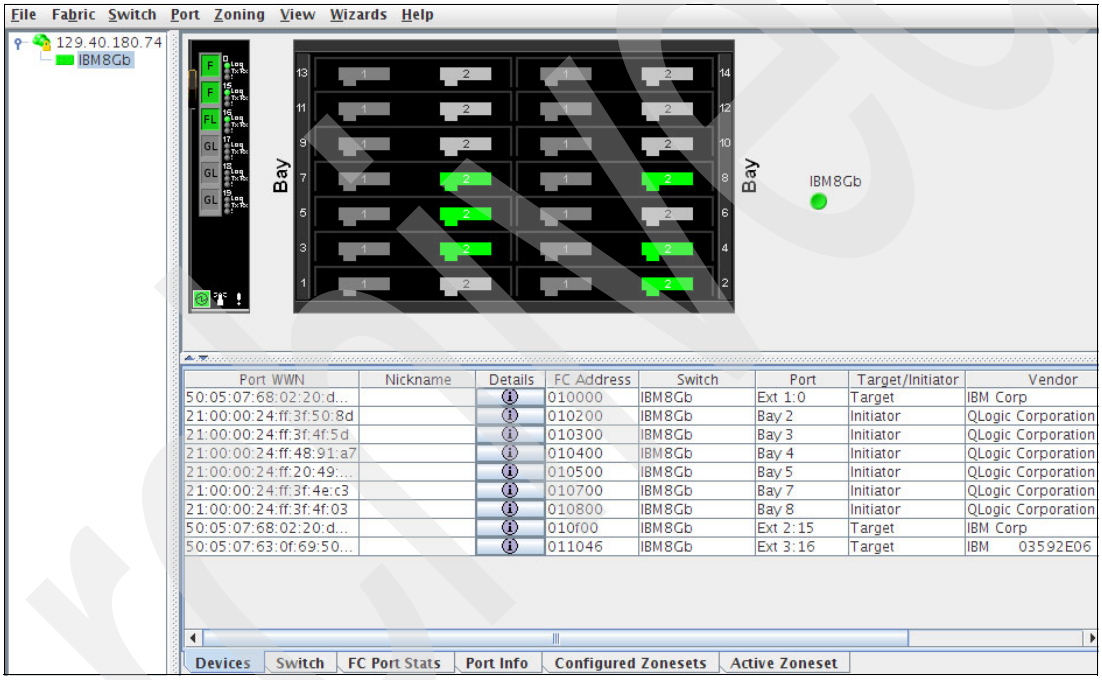

Figure 4-63 ScSE fabric displayed

Notice that the bays correspond to the compute nodes that are installed in the chassis. To edit the zone, click **Zoning** and then click **Zoning Edit**. The edit window opens (Figure 4-64).

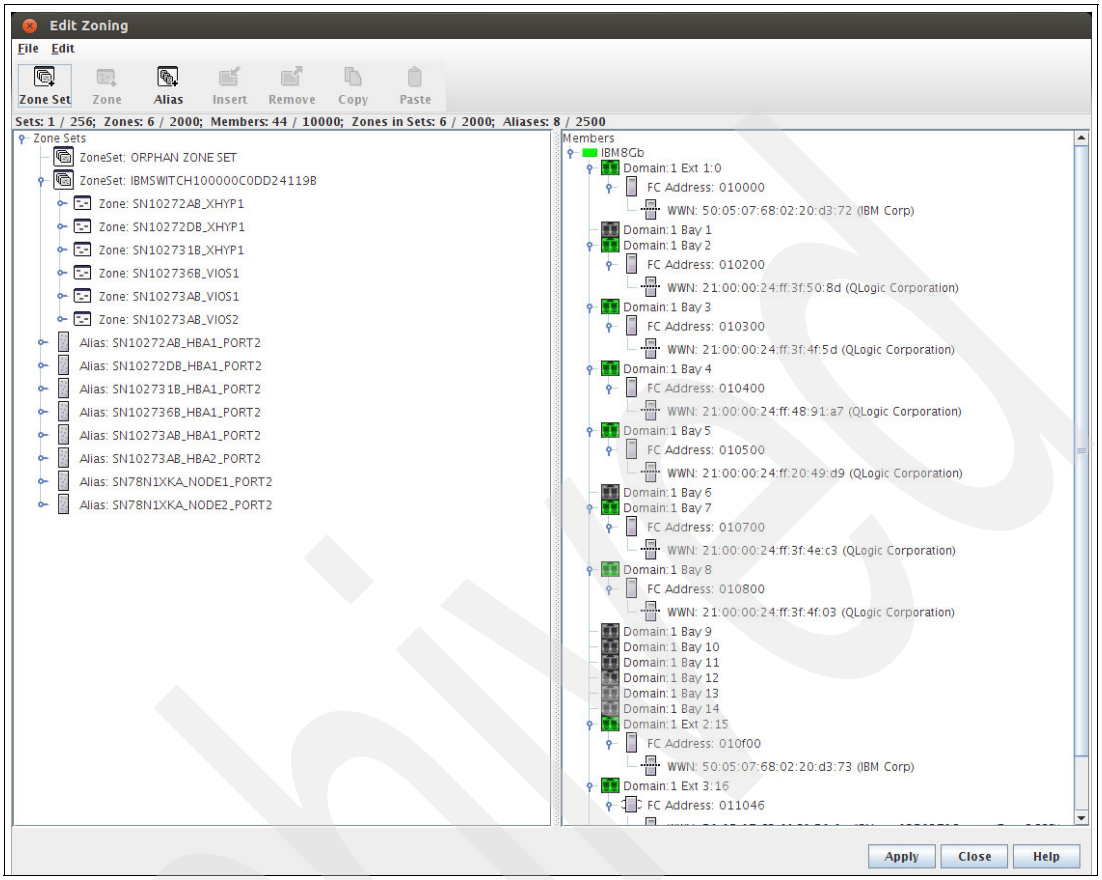

Figure 4-64 Edit Zoning window

This panel show the zones for the example environment. To see the world wide names for each alias in the zone, click on the zone, then the alias, as Figure 4-65 demonstrates.

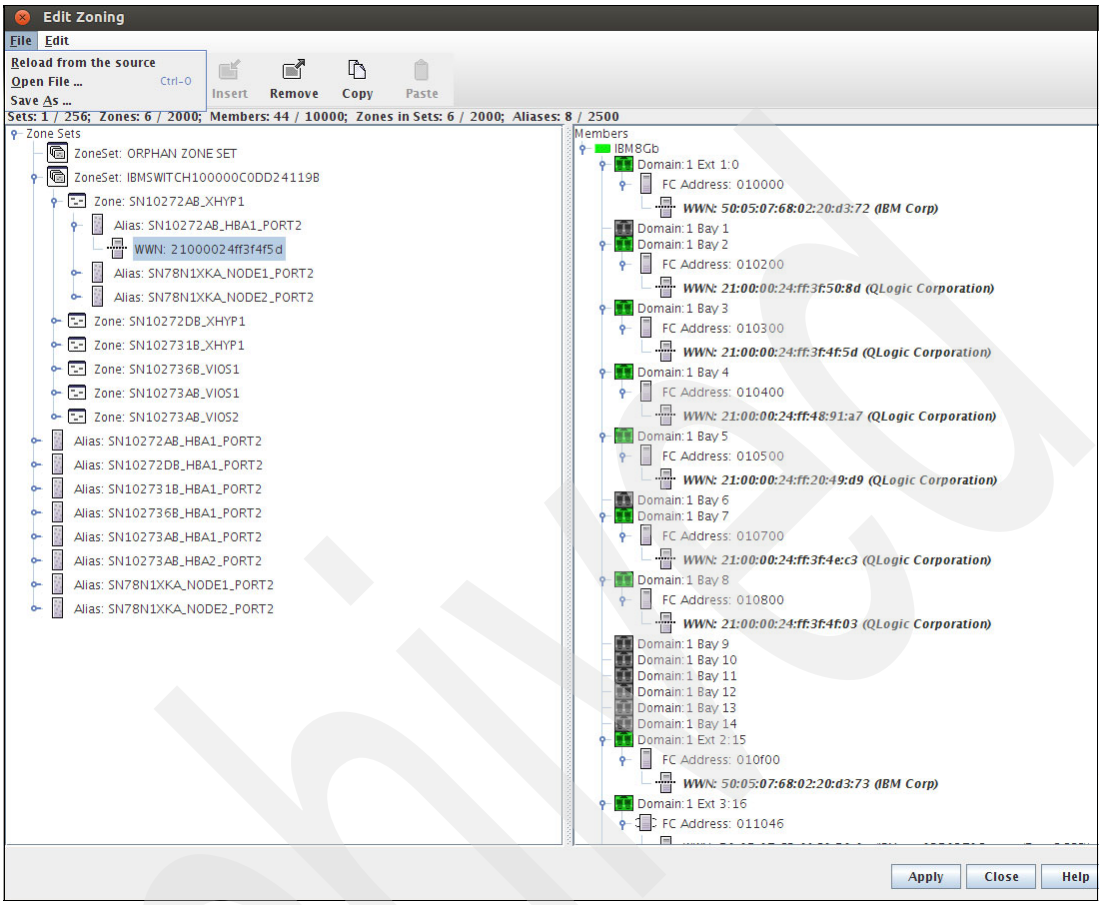

Figure 4-65 Edit zoning

You can also use the command-line interface (CLI) to edit the zoning information. This was discussed at the beginning of this chapter.

The commands used to create the aliases and zones are not covered here. For greater detail, go to the following links:

- http://www.ibm.com/systems/flex/networking/fibrechannel/fc3171_8gb_s/index.html
- http://ibm.co/14AYxmq
- http://www.redbooks.ibm.com/abstracts/tips0866.html

4.5.3 FSM Setup for Storage Management

After the SAN setup is completed on the ScSEs and V7000, the FSM must be associated with the storage. A single command, **`smcli manageV7000`**, is provided, which performs the several necessary steps. An SSH key-pair for USERID must exist on the FSM so that this command can succeed (in the `/home/USERID/.ssh` directory). Issue the **`smcli manageV7000`** command with the two flags listed in the following example (the password for the superuser on the V7000 and the IP address of the subsystem):

```
USERID@f3efsm:~> smcli manageV7000 -p passw0rd -i 129.40.180.68
```

The output of the command is similar to Example 4-5.

Example 4-5 Output of smcli command

```
Pushing SSH public key for admin to V7000 using command:  su superuser -c "scp -o
StrictHostKeyChecking=no -o UserKnownHostsFile=/dev/null
/home/USERID/.ssh/id_rsaV7000.pub superuser@129.40.180.68:/tmp/id_rsaV7000.pub  "
 SSH public key for admin was pushed to the V7000 successfully
Assigned SSH public key file to admin.
Issue mkdatasource command:  /usr/smrshbin/smcli mkdatasource -c svc -f
/home/USERID/.ssh/id_rsaV7000  -v V7000 -i 129.40.180.68
 mkdatasource was successful
 Collect inventory of Farm resource.
 Waiting for storage device mep to be created.
 V7000 with management address: 129.40.180.68 was managed succesfully.
 Issue mksvcsshrsap command:  /usr/smrshbin/smcli mksvcsshrsap -s "Storwize
V7000-2076-f3e7000a-IBM" -u USERID -f /home/USERID/.ssh/id_rsaV7000
 The mksvcsshrsap command was successful.
 V7000 with management address: 129.40.180.68 was managed succesfully.
```

Now, the Storage Array appears in Resource Explorer and the storage setup is complete.

4.5.4 FSM Storage Subagent and Image Repository setup

The Flex Service manage requires that a subagent be installed so that it can correctly manage the V7000. In addition, the FSM must be configured so that the IBM FlashCopy® services of the V7000 can be used as an image repository. This will be used when deploying POWER Cloud based services.

From the home page of the FSM, select the Plugins tab. Scroll to find the VMControl Enterprise Edition link (Figure 4-66).

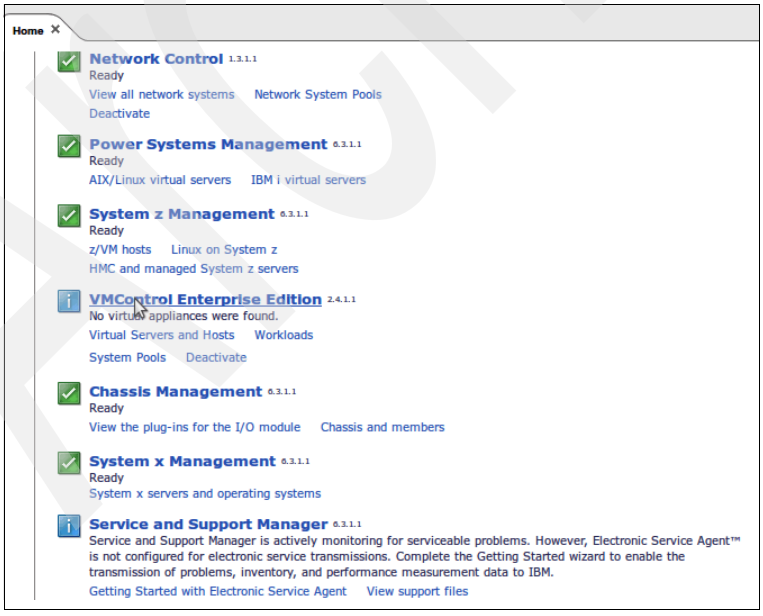

Figure 4-66 Selecting VMControl

The VMControl Enterprise window opens. A warning message indicates that no image repository is detected. Click **Install Agents**, below the warning message (Figure 4-67).

Figure 4-67 Starting the Agent installation wizard

The Agent Installation wizard starts (Figure 4-68).

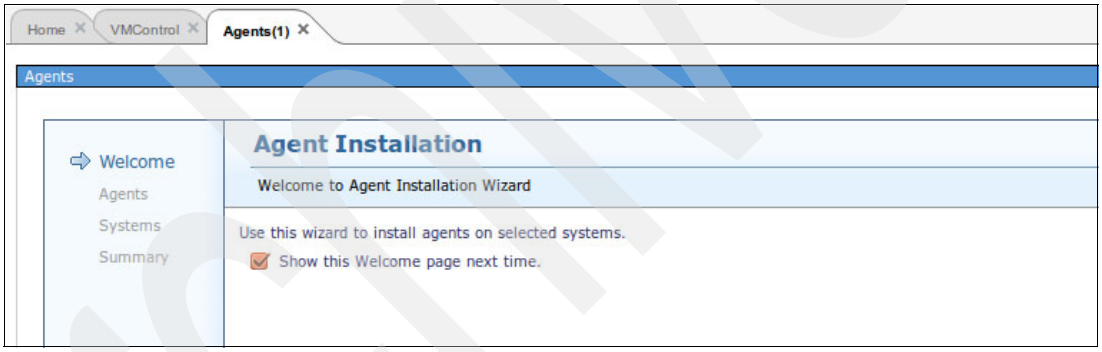

Figure 4-68 Agent Installation Welcome

From the list of agents (Figure 4-69), select all agents.

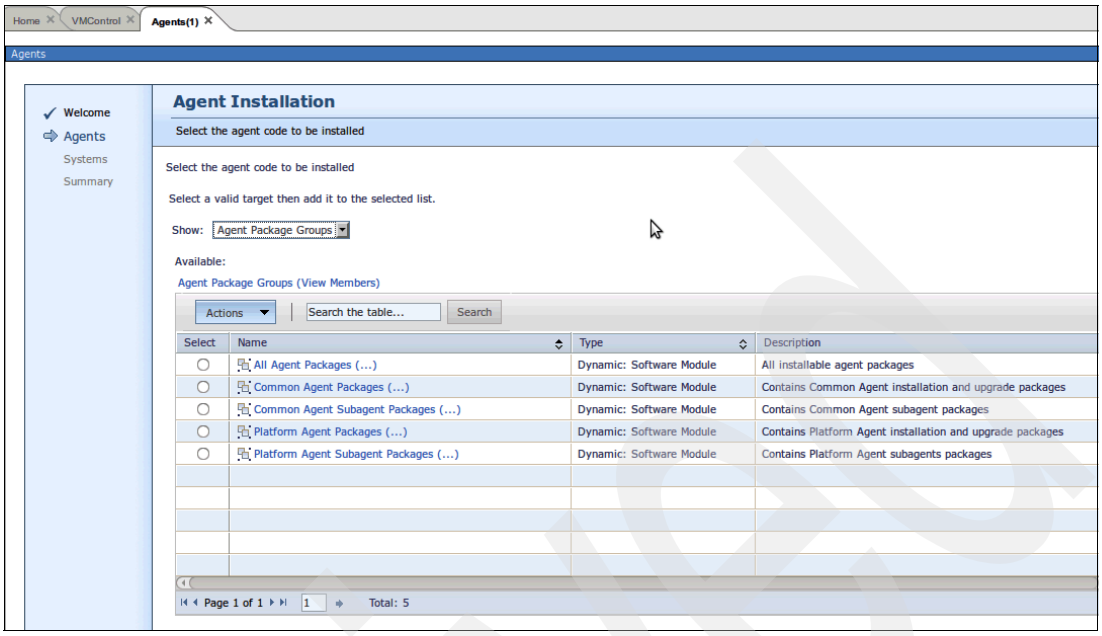

Figure 4-69 Agent Installation

Select **CommonAgentSubagent_VMControl_CommonRepository-2.4.1.1** (Figure 4-70), and then click **Add**.

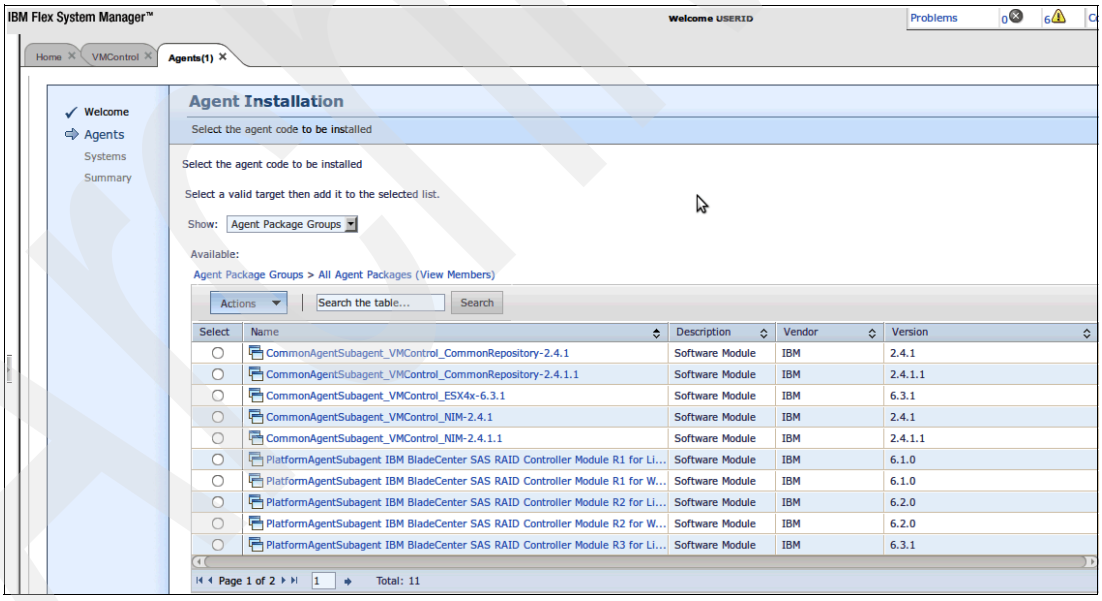

Figure 4-70 Selecting Repository for agent and choosing Add

Select the system on which to install the agent. See Figure 4-71. For the purpose of this document, the host must be a VIO server, because this agent will be used for the POWER compute node cloud services. If there is no VIO server defined to FSM (see also 4.7.1, “VMware preinstall tasks” on page 101), then locate and select the appropriate system and click **Add**.

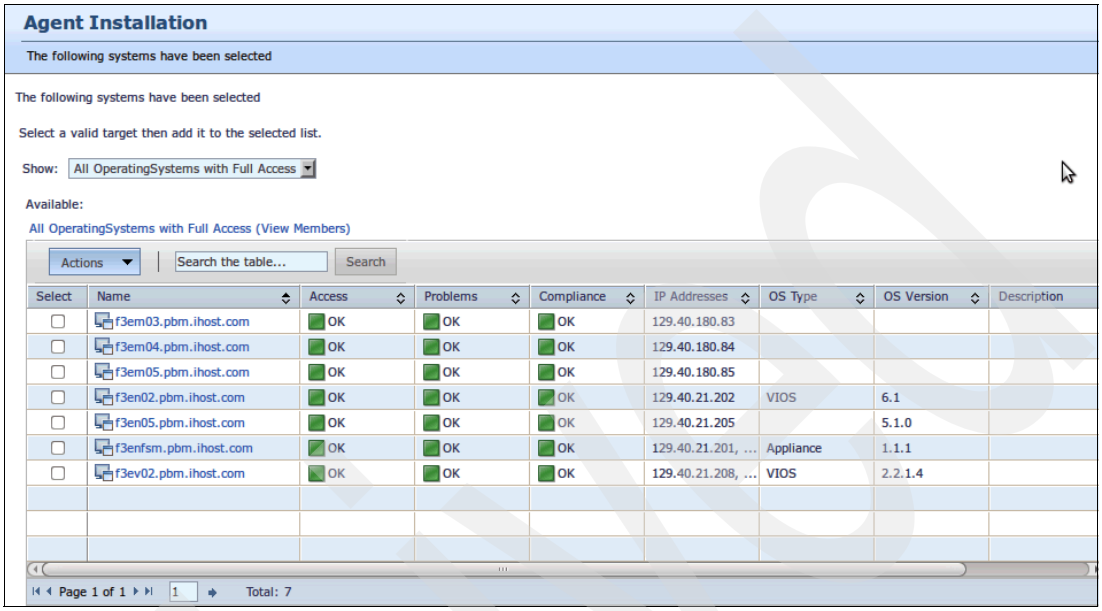

Figure 4-71 Systems that have been selected

Click **Next** to show the summary window (Figure 4-72).

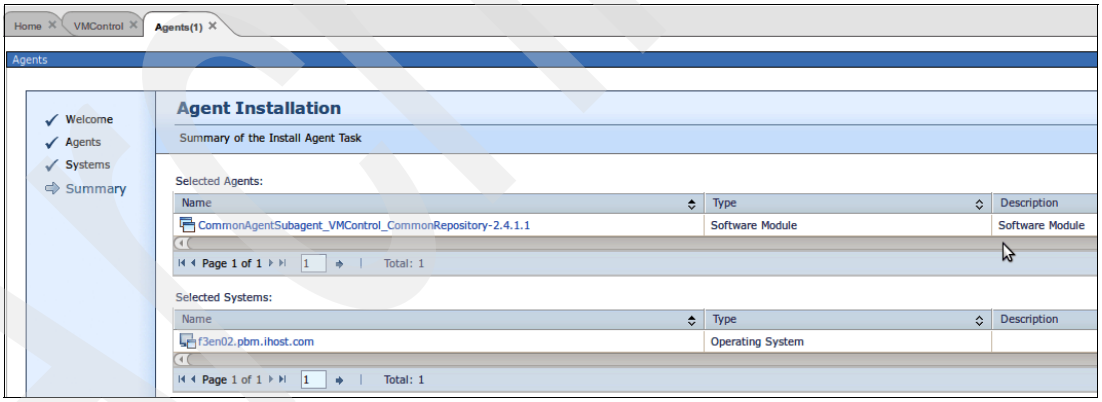

Figure 4-72 Summary of Install Agent Task

Click **Finish** to start the installation and a new job.

After it finishes completely, create an image repository. From the VMControl Enterprise Edition main panel, click the **Create image repository** link, under the virtualization tasks section (Figure 4-73).

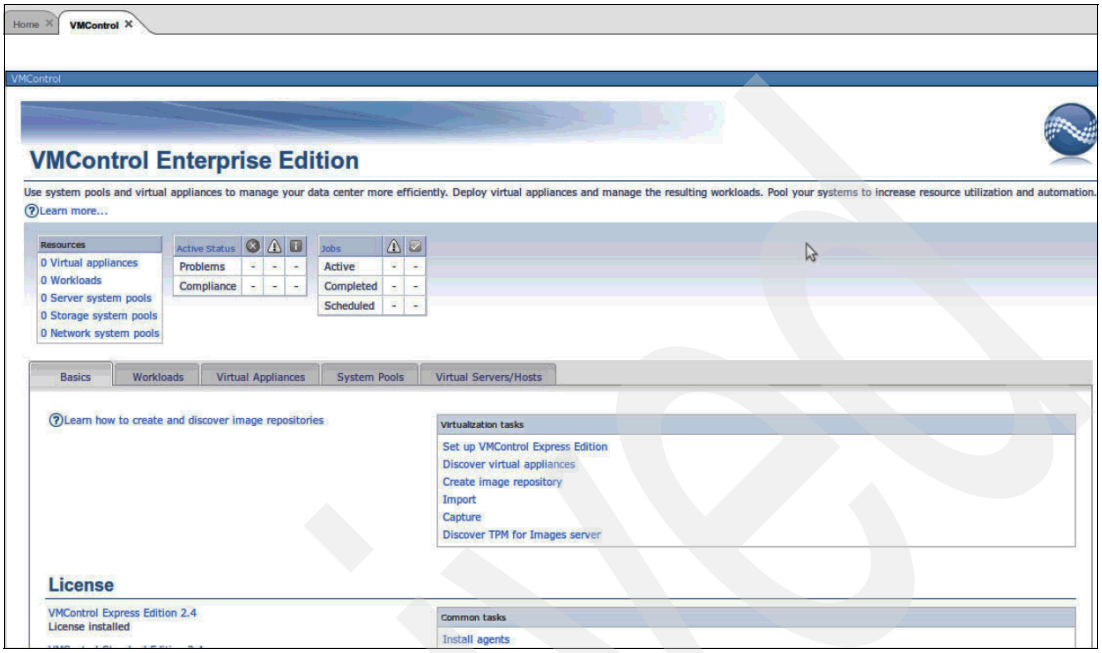

Figure 4-73 Virtualization Tasks: choosing Create image repository

The wizard that will create the image repository starts (Figure 4-74).

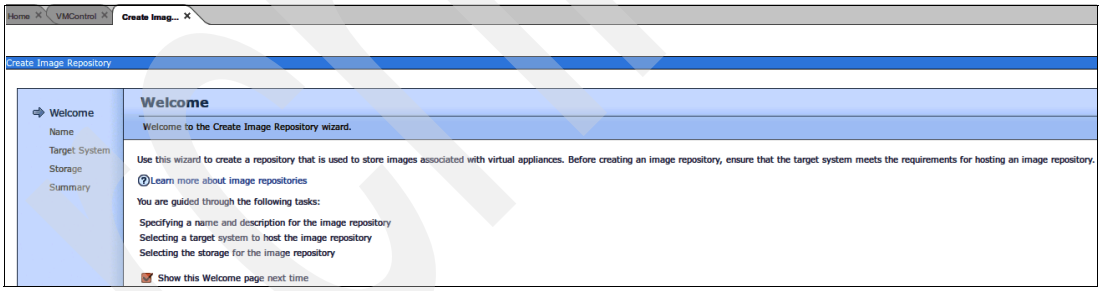

Figure 4-74 Create Image Repository: Welcome

Because this is a creating an image repository that will be located on the V7000, the name `v7000_ImageRepository` is specified (Figure 4-75).

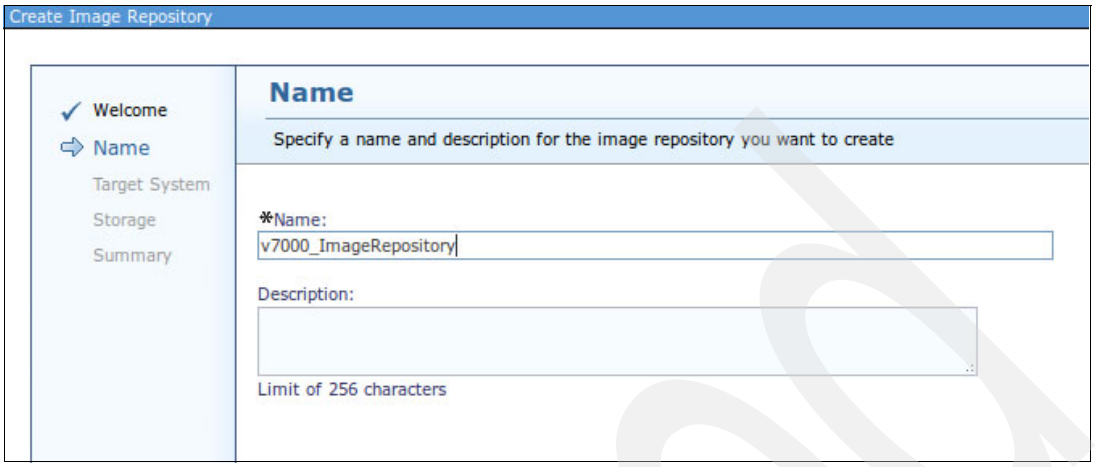

Figure 4-75 Create Image Repository: Specify name

Select the host on which the subagent is installed (Figure 4-76).

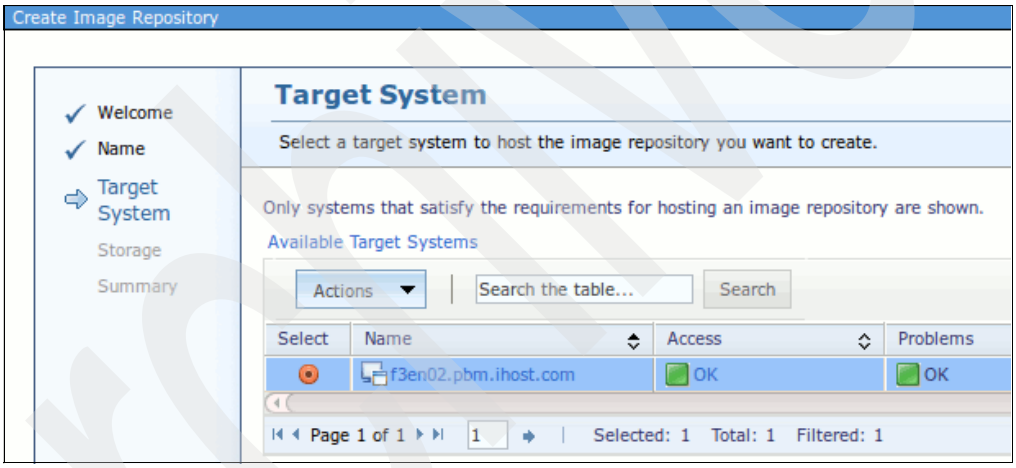

Figure 4-76 Create Image Repository: Select host that had subagent installed

Select the storage subsystem (Figure 4-77) that was set up in 4.5.1, “V7000 setup” on page 51.

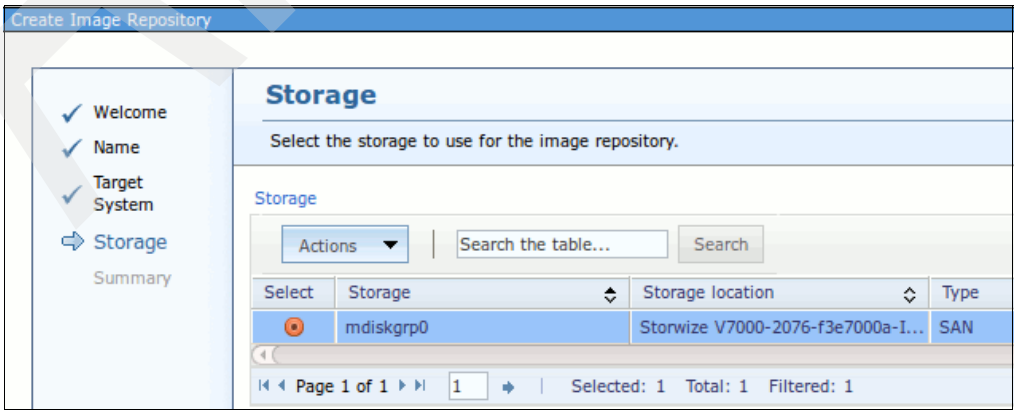

Figure 4-77 Create Image Repository: select storage to use for image

The summary panel opens (Figure 4-78).

Figure 4-78 Create Image Repository: Summary Panel

Click **Finish** to start the installation and start a new job. After it finishes completely, the cloud services can be set up for the Power Compute node

4.6 Power Compute node setup

This section discusses the steps that are required to allow SmartCloud Entry to offer cloud services on POWER compute nodes.

The first step to build these services is to set up one or more VIO instances on all of the POWER compute nodes. After the VIO instances are set up, a new virtual server can be created and captured. This captured virtual server is the basis of the POWER cloud services. The remainder of this chapter covers these steps, which include the following topics:

- Create new volume on V7000
- Map new volume to VIO
- Create virtual server
- Install AIX
- Install VSAE
- Define virtual server to FSM
- Capture virtual server
- Deploy virtual server

4.6.1 Set up VIO

Depending upon how the Flex System components were delivered, the POWER compute nodes might need to have VIO servers installed or configured. There is more than one way to install VIOS on a POWER node. The assumption for this document is that one or more VIOS are installed on all of the POWER compute nodes. The remainder of this chapter uses a VIOS that is preinstalled on a half-wide POWER node. This VIOS must be reconfigured.

If dual VIO servers are being used on a full-wide POWER compute node, and the operating systems are clones of one another, the FSM might not be able to discover them as separate operating systems. See Appendix A.3, “Discovering cloned operating systems” on page 197 for information about using cloned operating systems.

The first step is setting up the network information. In a VIOS environment, a Shared Ethernet Adapter is used to provide virtual networking to multiple virtual servers. Figure 4-79 shows how the various networking adapters detected by the VIOS operating system are related.

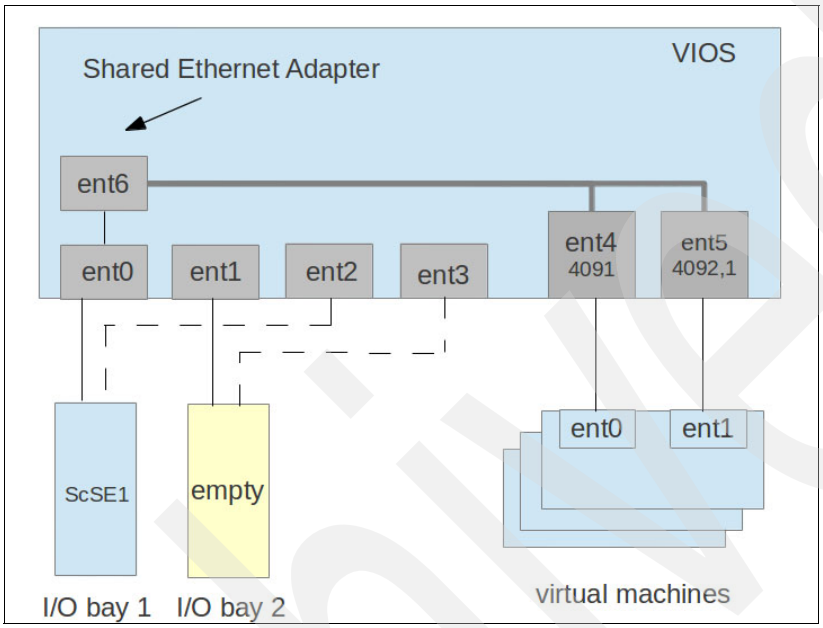

Figure 4-79 networking adapters detected by the VIOS

Adapters ent0, ent1, ent2, and ent3 are physical devices. They represent the four ports on the mezzanine card installed in the half-wide POWER node. Adapters ent4 and ent5 are virtual devices. They represent the VLANs that are configured on the physical devices. Adapter ent6 binds the virtual devices to the physical devices. It also provides the virtualization necessary for more that one virtual machine to use a physical device.

The VIO server LPAR profile must be created or edited to configure the adapters to match. From the FSM's Chassis Manager, select the node on which the VIO server is installed, then select **Manage Power Systems Resources** from the **General Actions** drop-down menu. From the Virtual Machines section, select the appropriate VIOS. Then, from the Actions menu, select **System Configuration** → **Manage Profiles**. From this panel, select the appropriate profile and from the Actions drop-down menu, select **Edit**. Select the Virtual Adapters tab.

There are two virtual adapters (slots 2 and 3), as shown in Figure 4-80.

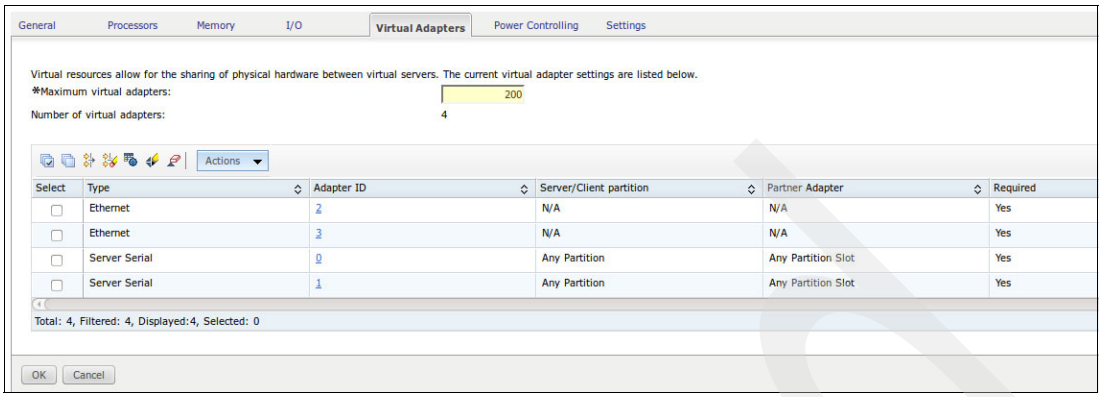

Figure 4-80 VIO Virtual Adapters tab

The virtual adapter in slot 2 must be connected to VLAN ID 4091 (Figure 4-81).

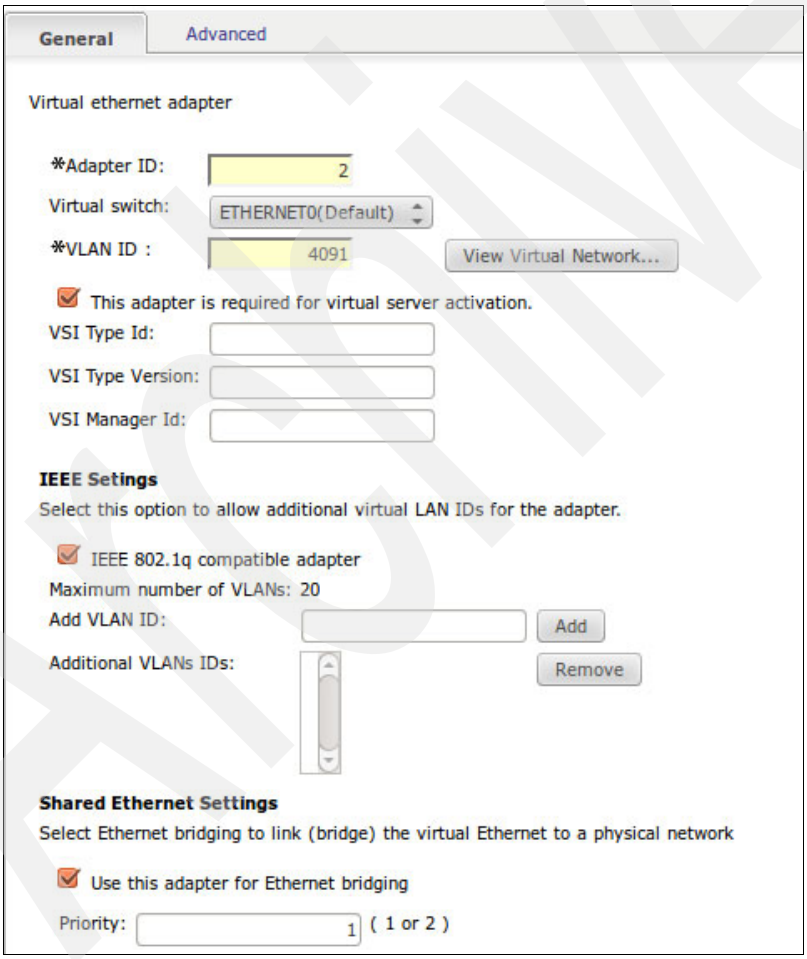

Figure 4-81 Ethernet adapter 2: General tab

The virtual adapter in slot 3 must be connected to VLAN ID 1, with 4092 listed as an additional VLAN (Figure 4-82).

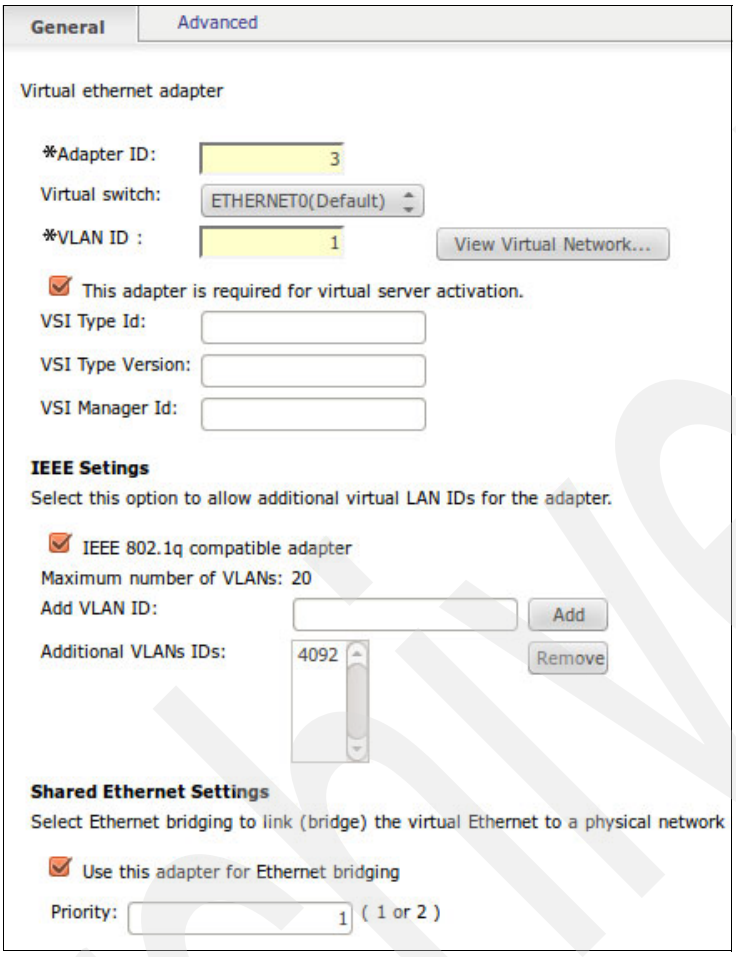

Figure 4-82 Ethernet adapter 3: General tab

To verify these connections, click **View Virtual Network**. From here, each VLAN ID can be selected by way of a radio button. For instance, when VLAN 1 is selected, the result is similar to Figure 4-83.

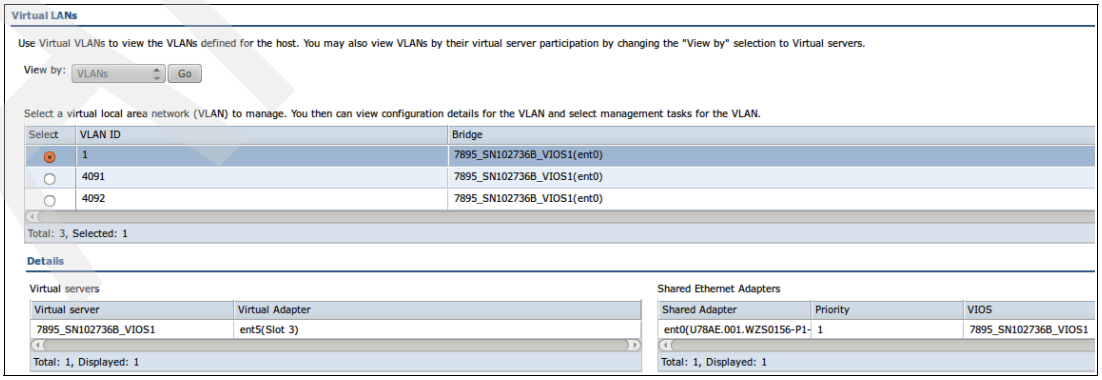

Figure 4-83 Virtual LANs

Notice that the virtual adapter ent5 is slot 3. Figure 4-84 shows 4091 is selected.

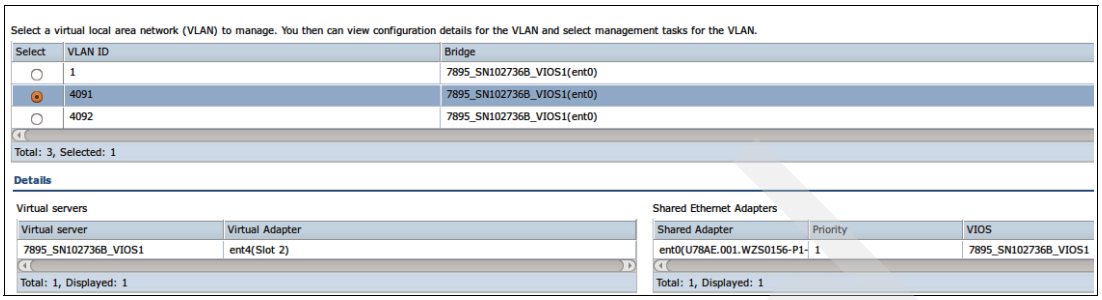

Figure 4-84 Virtual LAN: 4091

Notice that the virtual adapter ent4 is slot 2. Figure 4-85 shows 4092 is selected. Notice that the virtual adapter ent5 is slot 3.

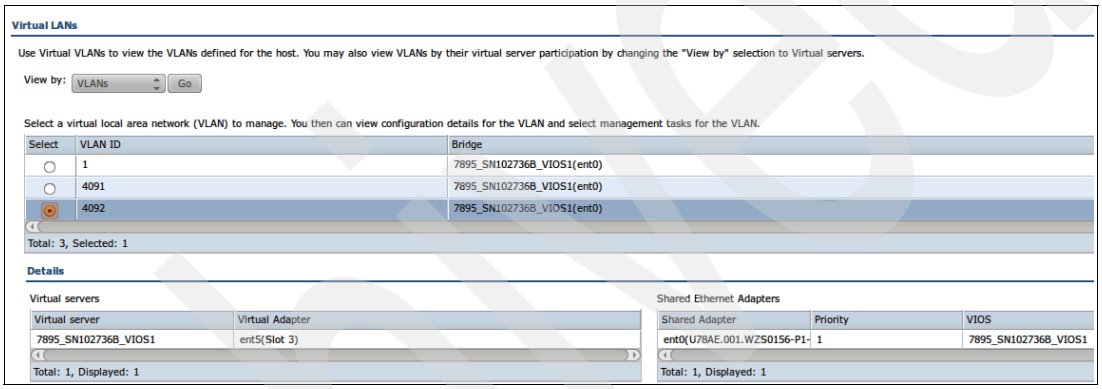

Figure 4-85 Virtual LAN: 4092

After the LPAR profile is saved and activated, you can use the **mkvdev** command to create a Shared Ethernet Adapter:

```
mkvdev -sea ent0 -vadapter ent4,ent5 -default ent4 -defaultid 4091
```

This command creates a Shared Ethernet Adapter ent6. The **-sea** flag signals **mkvdev** to use ent0 as its physical device. The **-vadapter** flag tells ent4 and ent5, created in the profile previously, to use this Shared Ethernet Adapter. The **-default** flag specifies the Ethernet adapter to use for untagged packets. In this case, ent4 is this default adapter. The **-defaultid** flag specifies the VLAN ID of the default adapter.

The VIOS **lsdev** command can be used to show the configuration. Example 4-6 lists all of the Ethernet devices configured for this VIOS.

Example 4-6 List of Ethernet devices

```
$ lsdev -dev ent*
name             status      description
ent0             Available   10GbE 4-port Mezzanine Adapter (a2191007df1033e7)
ent1             Available   10GbE 4-port Mezzanine Adapter (a2191007df1033e7)
ent2             Available   10GbE 4-port Mezzanine Adapter (a2191007df1033e7)
ent3             Available   10GbE 4-port Mezzanine Adapter (a2191007df1033e7)
ent4             Available   Virtual I/O Ethernet Adapter (l-lan)
ent5             Available   Virtual I/O Ethernet Adapter (l-lan)
ent6             Available   Shared Ethernet Adapter
```

Adapter ent0 is the only one of the four physical ports that is actually connected to a ScSE. Adapter ent2 is not used because each compute node is licensed to use only one port on the ScSE. Also, ent1 and ent3 are not used because there is no ScSE installed in I/O Bay 2.

A Shared Ethernet Adapter is a logical construct that is “backed” by one or more physical adapters, or ports. In this case, there is only one active port, so the SEA will be backed only by eth0.

When a new virtual machine is deployed, it can have up to two Ethernet devices. These will appear to the deployed OS as ent0 and ent1. Adapter ent0 on the deployed virtual machine will be backed by ent4 on the VIOS and have access to VLAN 4091. Adapter ent1 will be backed by ent5 on the VIOS and have access to VLAN 4092 and VLAN 1. Physically, of course, all traffic passes through ent0 on the VIOS. The ScSE is configured to only allow external packets that are tagged for VLAN 4091, so the deployed virtual machine must use ent0 as its external based adapter. The ent1 adapter can be used for virtual machine to virtual machine communication.

As Example 4-7 shows, the SEA configuration can be verified by using the **lsdev** command. The output of interest in this discussion is highlighted in bold text.

Example 4-7 SEA configuration verification

```
$ lsdev -dev ent6 -attr
attribute     value     description
user_settable

accounting    disabled  Enable per-client accounting of network statistics                True
ctl_chan                Control Channel adapter for SEA failover                         True
gvrp          no        Enable GARP VLAN Registration Protocol (GVRP)                    True
ha_mode       disabled  High Availability Mode                                           True
jumbo_frames  no        Enable Gigabit Ethernet Jumbo Frames                             True
large_receive no        Enable receive TCP segment aggregation                           True
largesend     0         Enable Hardware Transmit TCP Resegmentation                      True
lldpsvc       no        Enable IEEE 802.1qbg services                                    True
netaddr       0         Address to ping                                                  True
pvid          4091      PVID to use for the SEA device                                   True
pvid_adapter  ent4      Default virtual adapter to use for non-VLAN-tagged packets       True
qos_mode      disabled  N/A                                                              True
real_adapter  ent0      Physical adapter associated with the SEA                         True
thread        1         Thread mode enabled (1) or disabled (0)                          True
virt_adapters ent4,ent5 List of virtual adapters associated with the SEA (comma separated) True
```

4.6.2 Create volume on V7000

The virtual server that forms the basis for the cloud service requires a new volume to store the operating system. To create a new volume click **New Volume** (Figure 4-86).

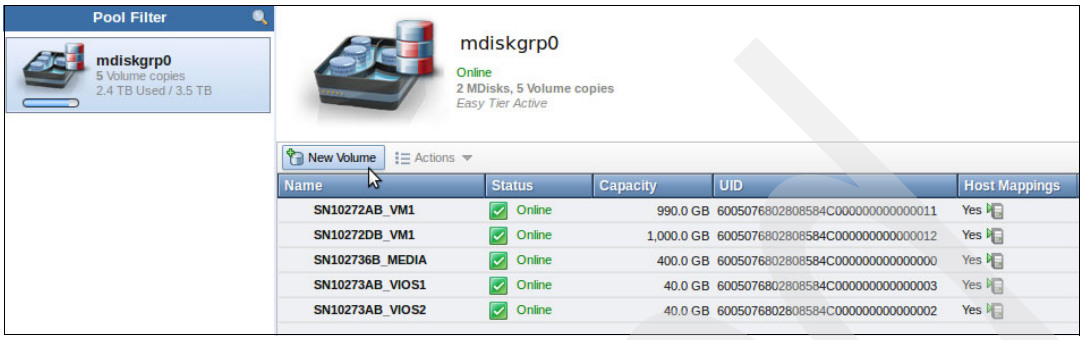

Figure 4-86 Click New Volume to define a new volume in a pool

The details of the new volume definition are presented. If you select the **Generic** option (Figure 4-87), a volume that uses a set amount of capacity from a single pool is created.

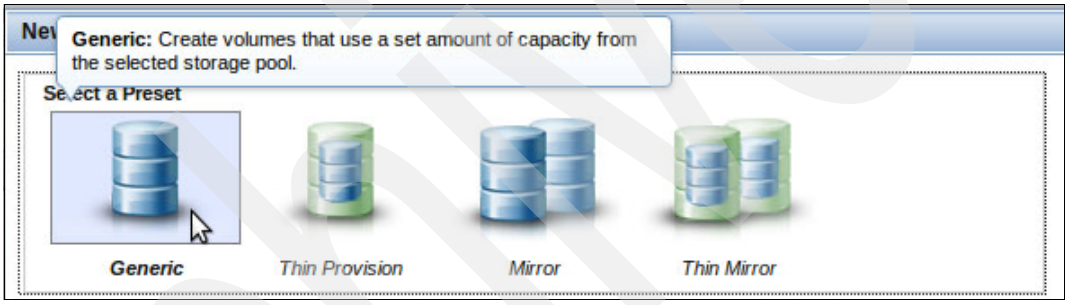

Figure 4-87 Details of New Volume definition

After choosing the Generic option, the pool must be selected. Figure 4-88 shows only one pool listed, which has the name mdiskgrp0.

Figure 4-88 New Volume: Select a Pool

Choosing this pool causes the text input areas shown in Figure 4-89 to be displayed. Enter the volume name and size in the context of the increments selected from the drop-down list. Although this interface accepts blanks as part of the Volume Name field, as noted in A.2, “FlashCopy not used during capture of a virtual appliance” on page 196, the use of blanks in this field can lead to problems later in the process. Click **Create and Map to Host**.

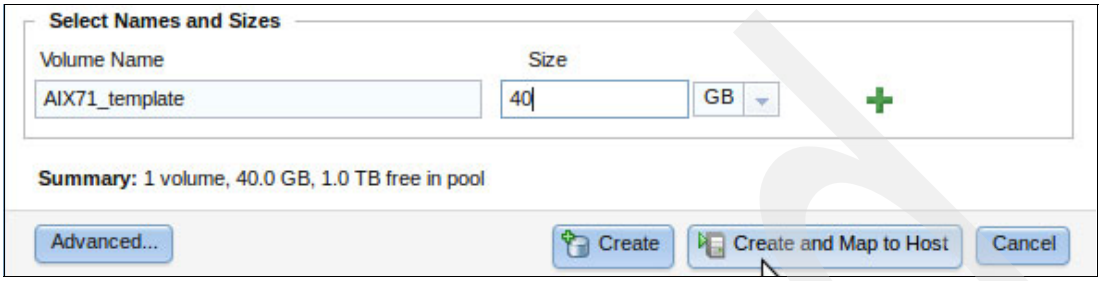

Figure 4-89 Select volume name and size

The result is the execution of commands that include `mkvdisk` to create a volume as shown in Figure 4-90. Click **Continue**.

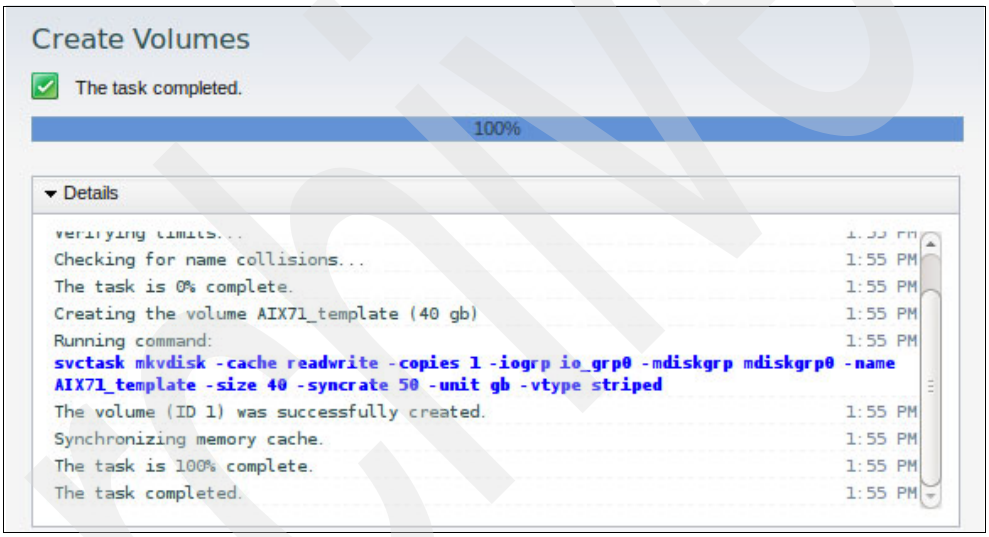

Figure 4-90 Create Volumes: Task Completed

After the requested volume is created it must be mapped to a host, as shown in the next several figures (Figure 4-91 through Figure 4-93 on page 70).

In the Modify Host Mappings section, select a host from the list. In Figure 4-91 the host containing the serial number of the POWER node (SN102736V) is selected.

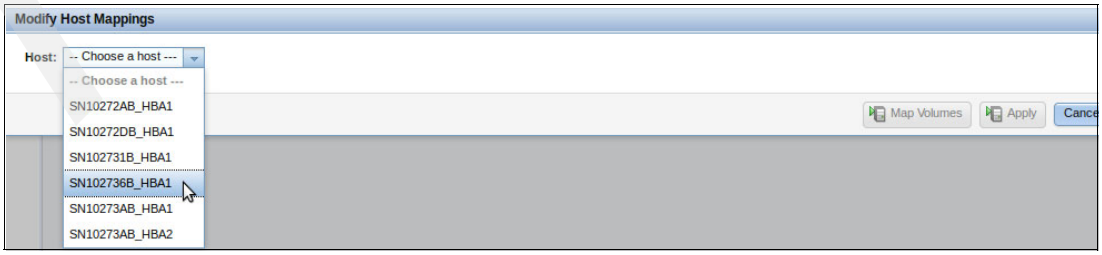

Figure 4-91 Modify Host Mappings

As a result, the volume is added to the Volumes Mapped to the Host panel (Figure 4-92). Click **Apply** to complete the mapping.

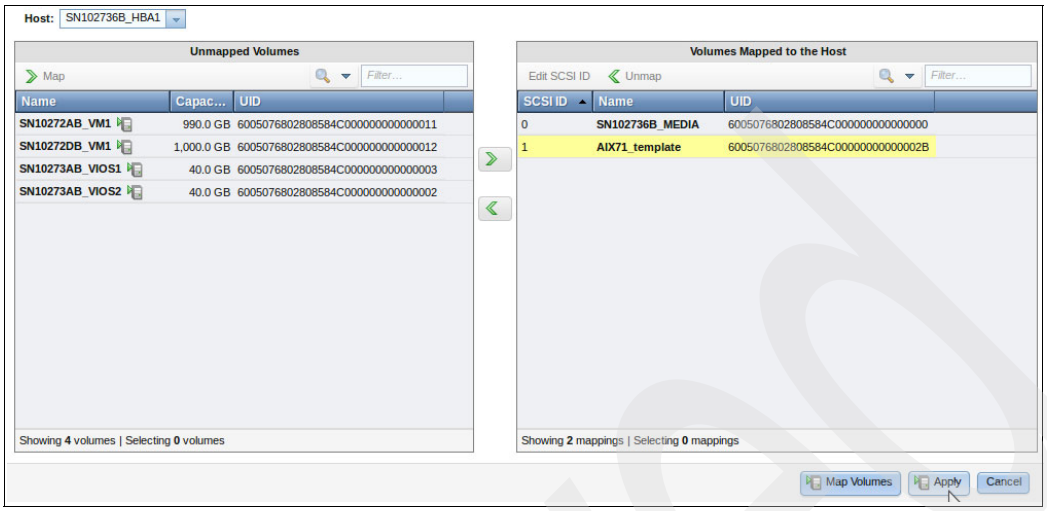

Figure 4-92 Volumes mapped to the host

Click **Close** as shown in Figure 4-93.

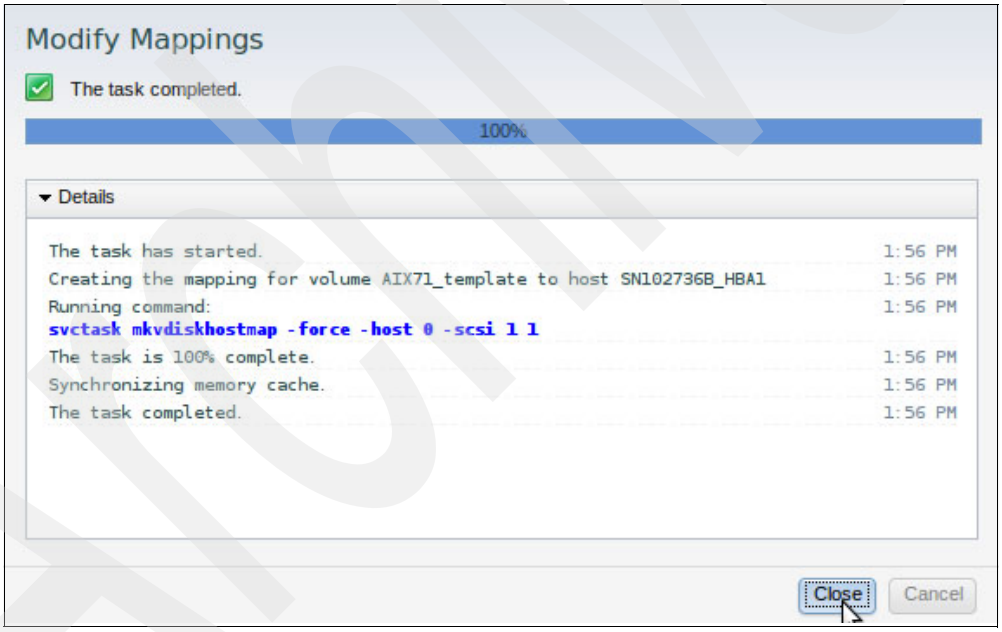

Figure 4-93 Modify Mappings: Task completed

The newly mapped volume appears as a physical disk (hdisk4) on the VIO server. This device will appear in the **lsdev** command output only after running the **cfgdev** command, as shown in Example 4-8.

Example 4-8 Confirming newly mapped volume is available

```
$ cfgdev
$ lsdev -dev hdisk*
name             status      description
hdisk0           Available   MPIO IBM 2076 FC Disk
hdisk1           Defined     MPIO IBM 2076 FC Disk
hdisk2           Available   SAS Disk Drive
hdisk3           Available   SAS Disk Drive
hdisk4           Available   MPIO IBM 2076 FC Disk
```

4.6.3 Create virtual server

From the FSM, select the single-wide POWER Node by clicking its associated graphical representation. Then, to start the Create Virtual Server wizard, select **System Configuration** → **Create Virtual Server**, as shown in Figure 4-94.

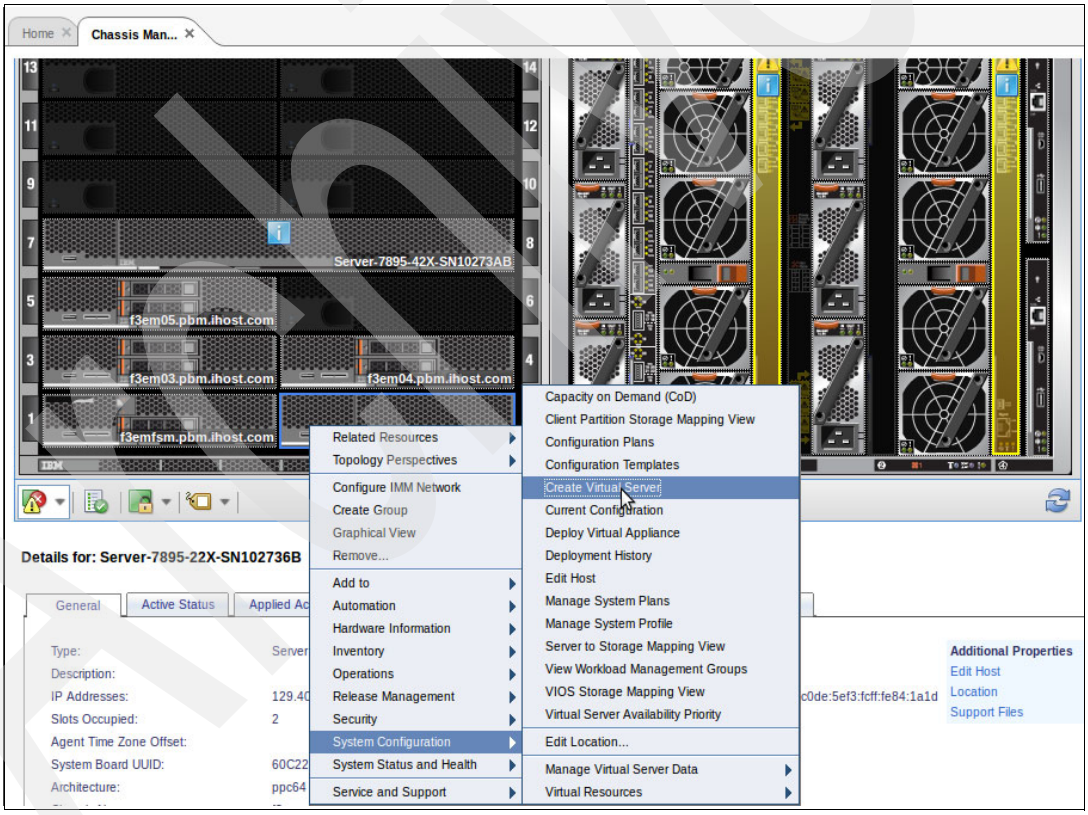

Figure 4-94 Chassis GUI

In the Create Virtual Server Wizard, enter a unique name for the virtual server (Figure 4-95). Keep the default value for Virtual server ID field and be sure all the check boxes are clear. Click **Next**.

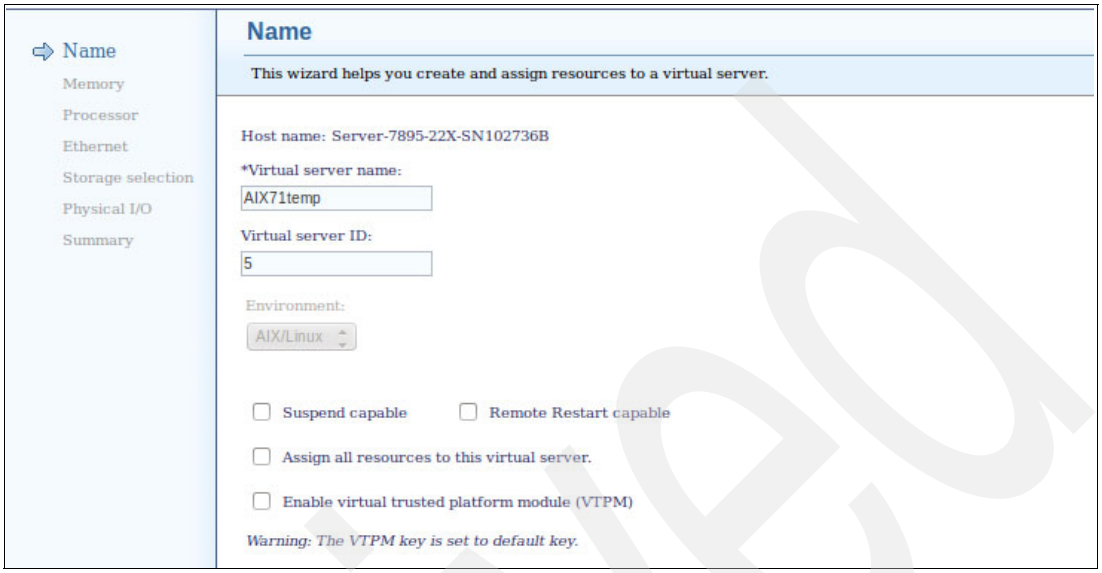

Figure 4-95 Create Virtual Server Wizard: Specifying name

Assign dedicated memory to the virtual server, and assign one or more processors to the virtual server. Also, select which Port VLAN IDs will be associated with the Ethernet adapters.

The scalable switch element is configured so that VLAN 4091 is accessible through the data network from entities external to the chassis.

In Figure 4-96, virtual Ethernet adapter 2 is assigned to Port VLAN ID 4091 and virtual Ethernet adapter 3 is assigned to Port VLAN ID 1.

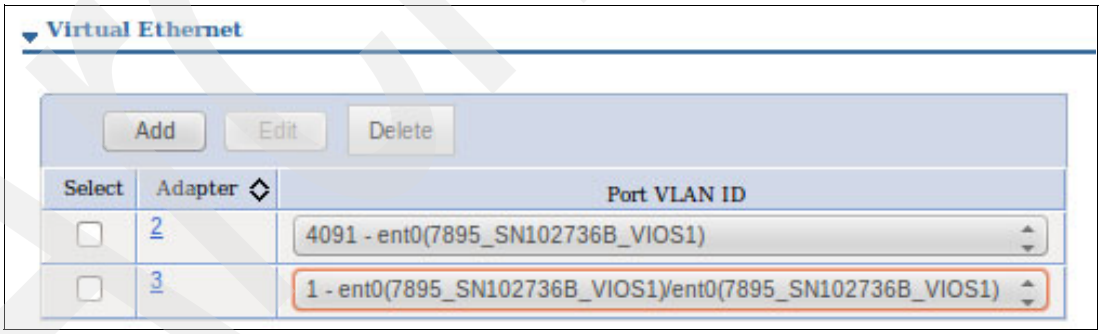

Figure 4-96 Virtual Ethernet assignments

Storage selection is the next step. Select the **Yes, Automatically manage the virtual storage adapters for this Virtual Server** (Figure 4-97), and then select the **Physical Volumes** check box. Click **Next**.

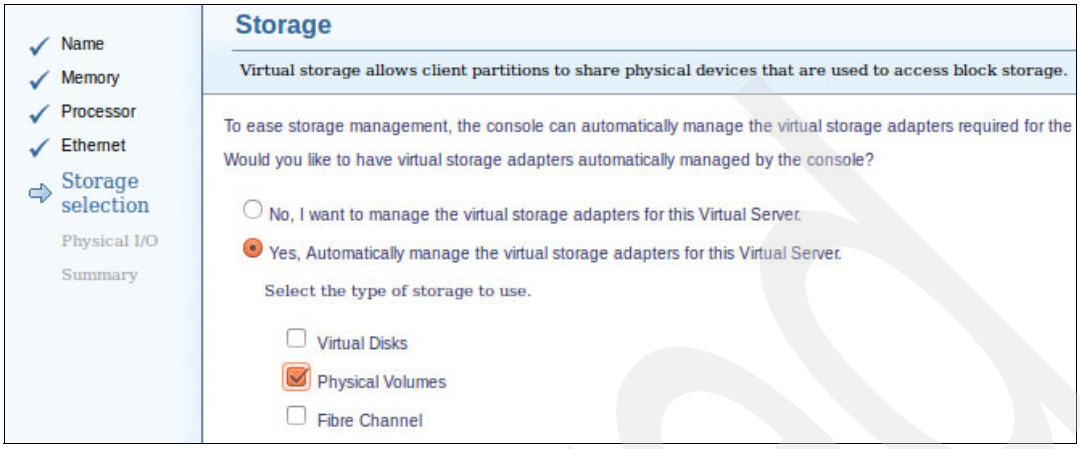

Figure 4-97 Create virtual server: Storage selection

In the Physical Volumes table (Figure 4-98), select the check box next to the disk that was created previously.

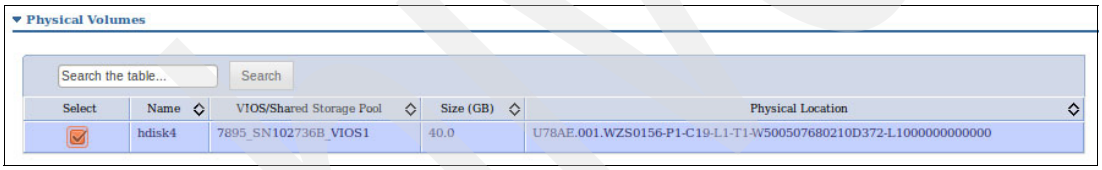

Figure 4-98 Create virtual server: Physical Volume Selection

The next Create Virtual Server Wizard window displayed is associated with optical devices (Figure 4-99). No Physical Optical Devices should be listed. Under the Virtual Optical Media section, select the check box next to the AIX 7.1 install disk, and then click **Next**.

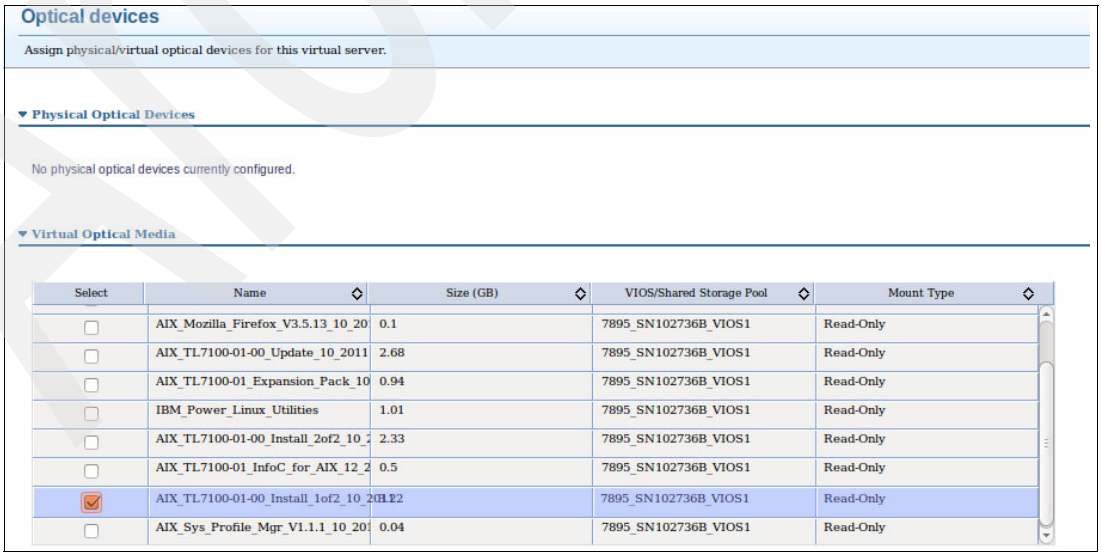

Figure 4-99 Create virtual server: Optical device selection

This environment must be fully virtualized to best suit a cloud environment. Therefore, no physical adapters need to be assigned to this virtual server. However, all of the physical I/O adapters are dedicated to the VIO server, so none appear on the Create Virtual Server Wizard Physical I/O Adapters panel (Figure 4-100). Click **Next**.

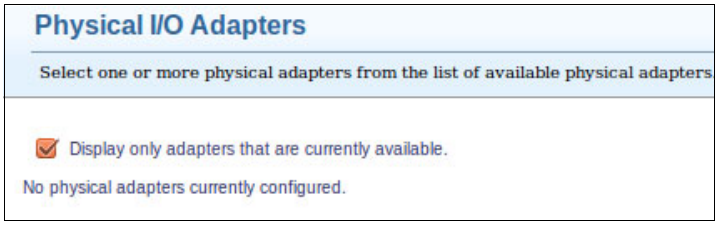

Figure 4-100 Create virtual server: Physical I/O Adapters

The Summary page opens (Figure 4-101). Review the information. If changes must be made to the virtual server specifications, click the **Back** button or click the appropriate topic link in the left-side navigation. Click **Finish** to complete the creation of the virtual server.

Figure 4-101 Create virtual server: Summary

To activate the profile for the virtual server, use the Resource Explorer, which is accessible through FSM. In FSM, right-click the Virtual Server Name, and then select **Operations** → **Activate** → **Profile**, as Figure 4-102 shows.

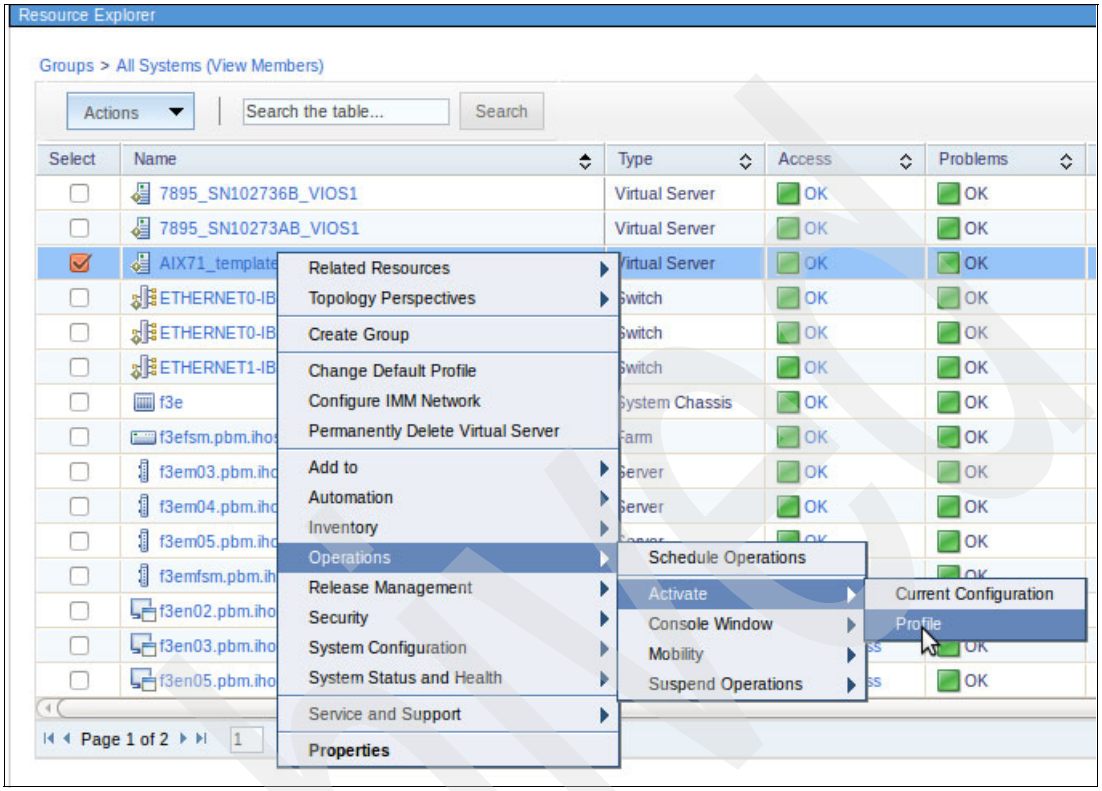

Figure 4-102 Activate virtual server profile

The result is the choices shown in Figure 4-103. At this point select the **Open a terminal window or console session** check box. Click **OK**.

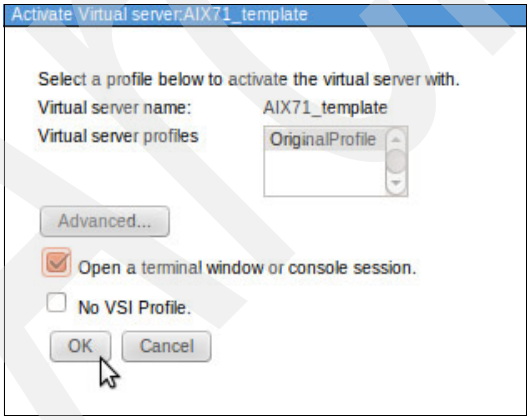

Figure 4-103 Activate virtual server

To continue, select the check box next to **Always trust content for this publisher** in the Security Warning dialog window and then click **Yes**, as shown in Figure 4-104.

Figure 4-104 Certificate cannot be verified

When prompted to authenticate (Figure 4-105), enter the user ID and password of an authorized FSM user. The default user ID is USERID and the password was set, as described in Figure 4-31 on page 36.

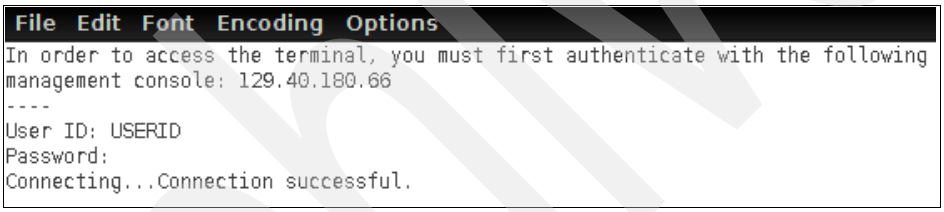

Figure 4-105 Authenticate with the management console

After authenticating, skip the informational screens by pressing Enter until the language selection screen opens. Select **1** for English (Figure 4-106). At this point, the installation of AIX can begin as described in 4.6.4, “Install AIX” on page 77.

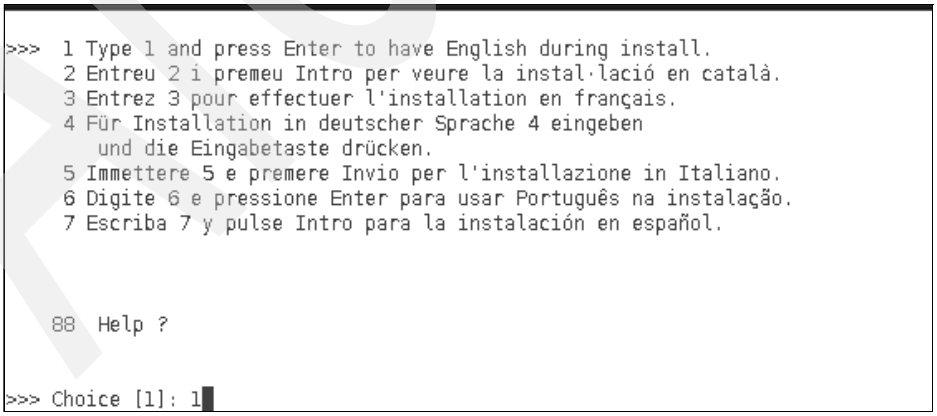

Figure 4-106 Specify English

4.6.4 Install AIX

After selecting the language (Figure 4-106 on page 76), change the installation type from **Express** to **Standard** and select **0, Install with the current settings listed above,** at bottom of the menu. The installation and settings menu is shown in Figure 4-107.

Note that hdisk0 is specified as `Disk Where You Want to Install` as part of `1 System Settings`. This is the first (and only) disk for this LPAR. This disk is the same volume that is created (and referred to as hdisk4) in 4.6.2, “Create volume on V7000” on page 68.

```
                        Installation and Settings

Either type 0 and press Enter to install with current settings, or type the
number of the setting you want to change and press Enter.

    1  System Settings:
         Method of Installation.............New and Complete Overwrite
         Disk Where You Want to Install.....hdisk0

    2  Primary Language Environment Settings (AFTER Install):
         Cultural Convention................English (United States)
         Language ..........................English (United States)
         Keyboard ..........................English (United States)
         Keyboard Type......................Default
    3  Security Model.....................Default
    4  More Options  (Software install options)
    5  Select Edition.....................standard
>>> 0  Install with the current settings listed above.

                       +-----------------------------------------------------
    88  Help ?         |    WARNING: Base Operating System Installation will
    99  Previous Menu  |    destroy or impair recovery of ALL data on the
                       |    destination disk hdisk0.
>>> Choice [0]: 0
```

Figure 4-107 AIX standard installation

The Overwrite Installation Summary screen opens (Figure 4-108). Enter `1` (one) and press Enter to begin the installation.

```
                        Overwrite Installation Summary

Disks:  hdisk0
Cultural Convention:  en_US
Language:  en_US
Keyboard:  en_US
JFS2 File Systems Created:  Yes
Graphics Software:  Yes
System Management Client Software:  Yes
Enable System Backups to install any system:  Yes
Selected Edition:  standard

Optional Software being installed:









>>> 1   Continue with Install
                       +-----------------------------------------------------
    88  Help ?         |    WARNING: Base Operating System Installation will
    99  Previous Menu  |    destroy or impair recovery of ALL data on the
                       |    destination disk hdisk0.
>>> Choice [1]: 1
```

Figure 4-108 Installation Summary

Wait for the non-interactive part of the installation to complete. When prompted (as shown in Figure 4-109) set the terminal type to vt100.

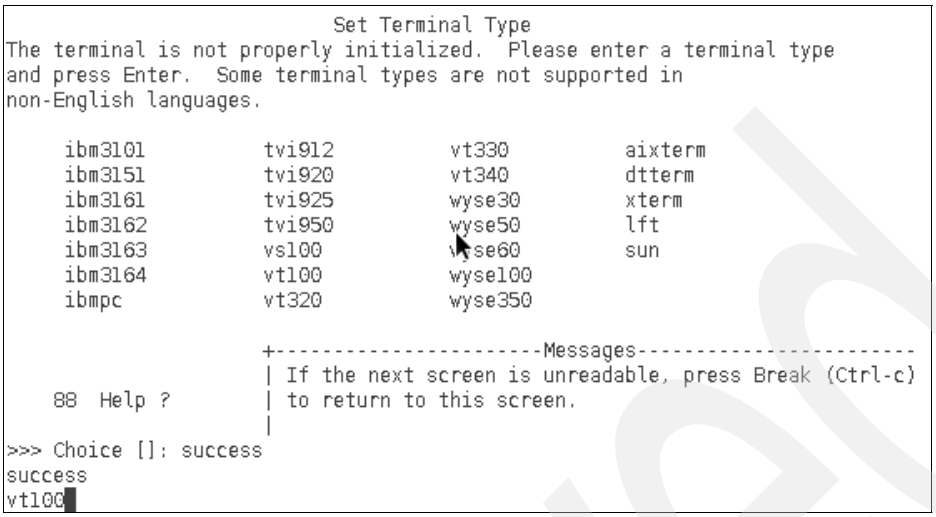

Figure 4-109 Set terminal type to vt100

When the Software License Agreements screen opens (Figure 4-110), select **Accept License Agreement** and press Enter.

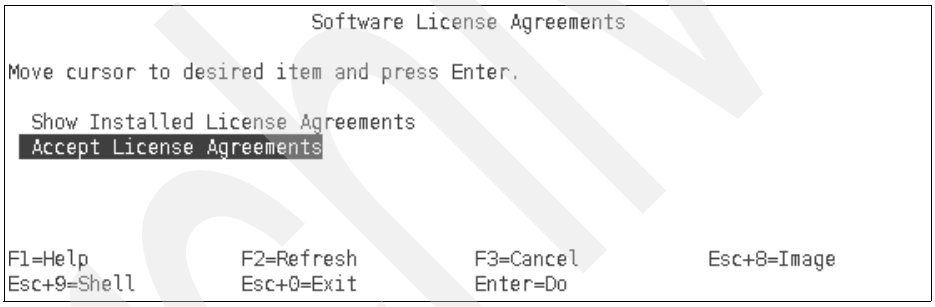

Figure 4-110 AIX install: Software License Agreements

To accept the license agreement, use the Tab key to specify **yes** (Figure 4-111), and then press Enter to continue.

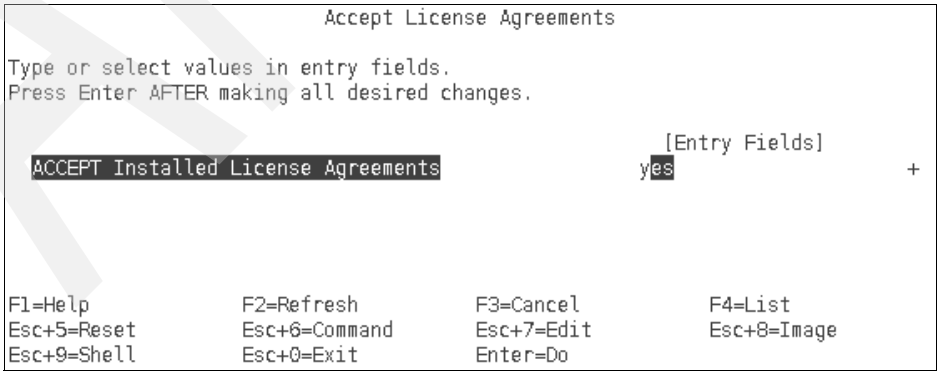

Figure 4-111 AIX install: Accept License Agreements

Wait for the `Command: OK` message to be displayed (Figure 4-112).

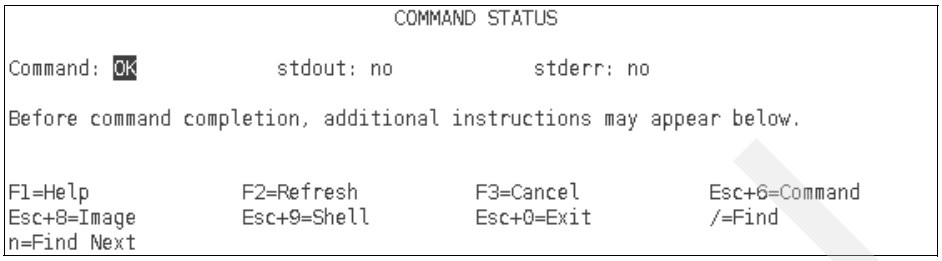

Figure 4-112 AIX install: Wait for Command OK to appear

To continue with the installation, press F3 multiple times until the Installation Assistant menu opens (Figure 4-113).

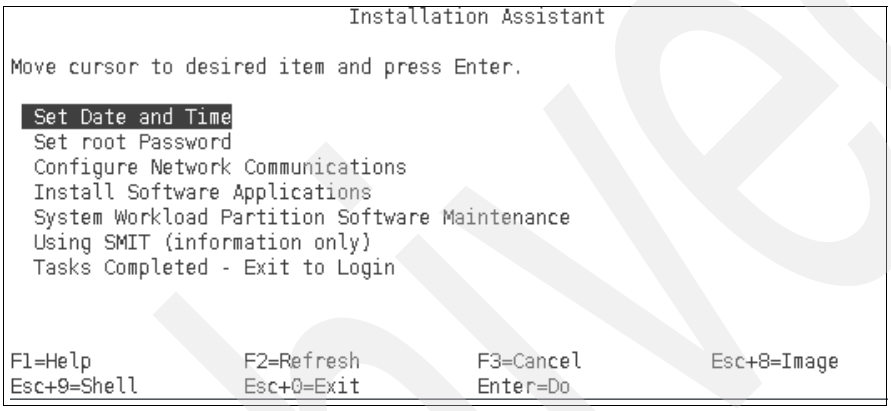

Figure 4-113 AIX Install: Set Date and Time

Use the Installation Assistant menu to select and configure the following settings:

- Set Date and Time
- Set root Password
- Configure Network Communications:
  a. Choose **TCP/IP**.
  b. Respond to the prompt of Available Network Interfaces pop-up menu (Figure 4-114 on page 80). Choose **en0** from the list of available network interfaces. This is the interface that connects to the 4091 VLAN.

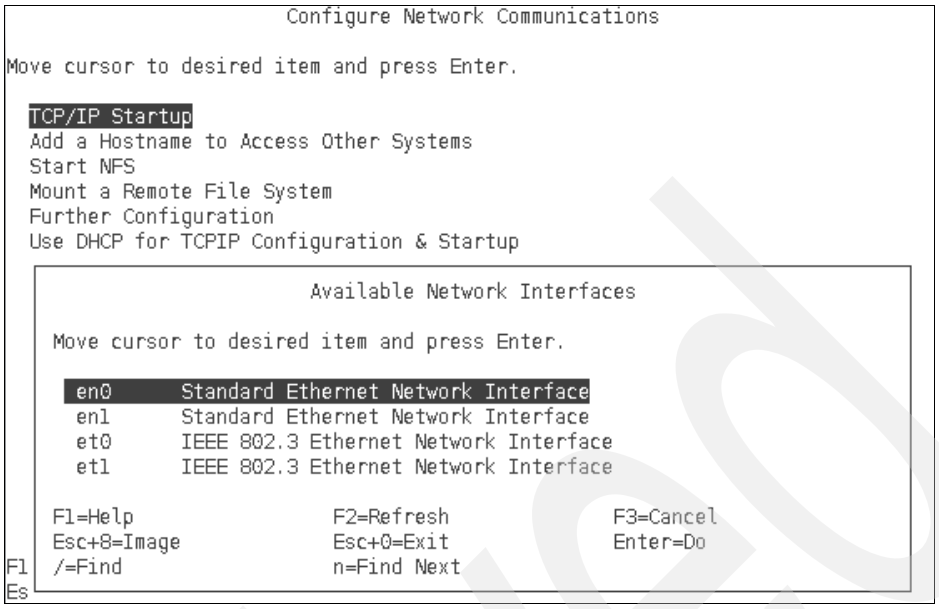

Figure 4-114 Configure Network Communications

c. After selecting **en0**, provide configuration and start up information, as shown in Figure 4-114. Select **Enter** (or **Do**) to proceed.

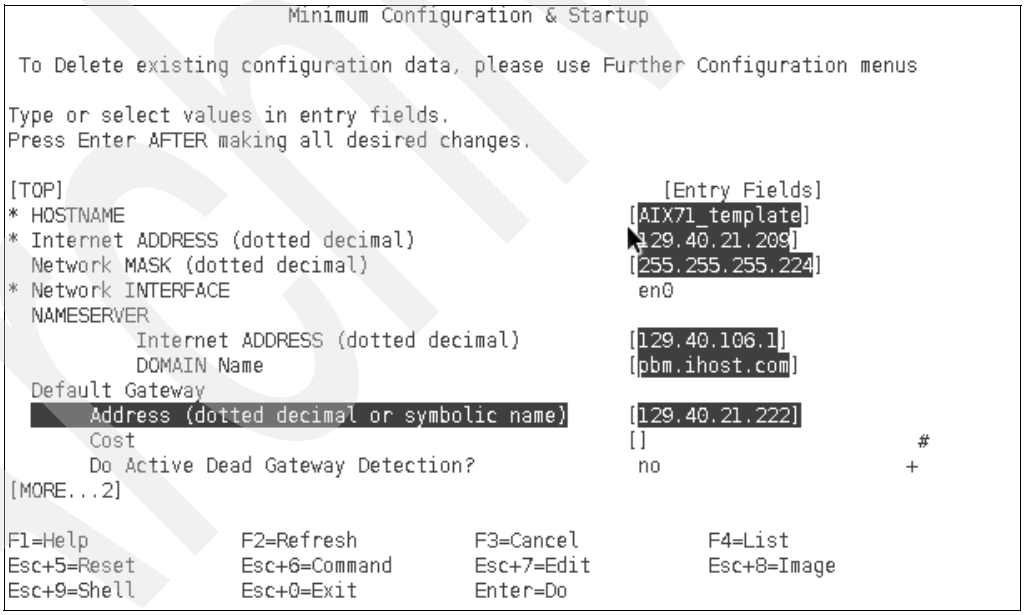

Figure 4-115 Minimum Configuration and Startup

The Install Software menu opens (Figure 4-116). Specify `/dev/cd0` as the device from which to install.

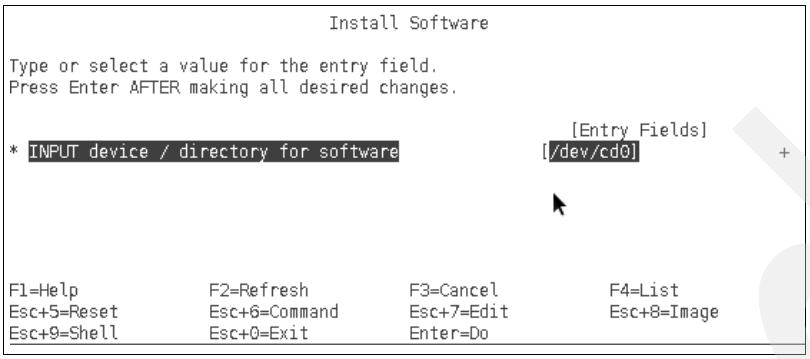

Figure 4-116 AIX Installation: Install Software Menu

For the FSM to access this virtual server, SSH must be installed.

On the subsequent menu (Figure 4-117) enter `openssh` next to SOFTWARE to install. Press Enter to proceed.

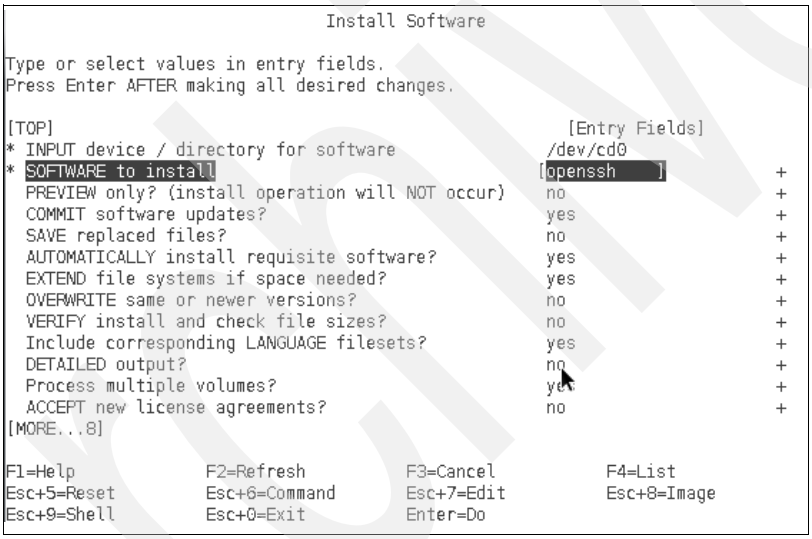

Figure 4-117 AIX Install: Specifying openssh

A confirmation screen opens. Select **ENTER** to proceed, and then press Enter to confirm the the request to install. When the message `command: OK` appears, as shown in Figure 4-118, scroll to bottom to make sure that `Success` is indicated for the installation.

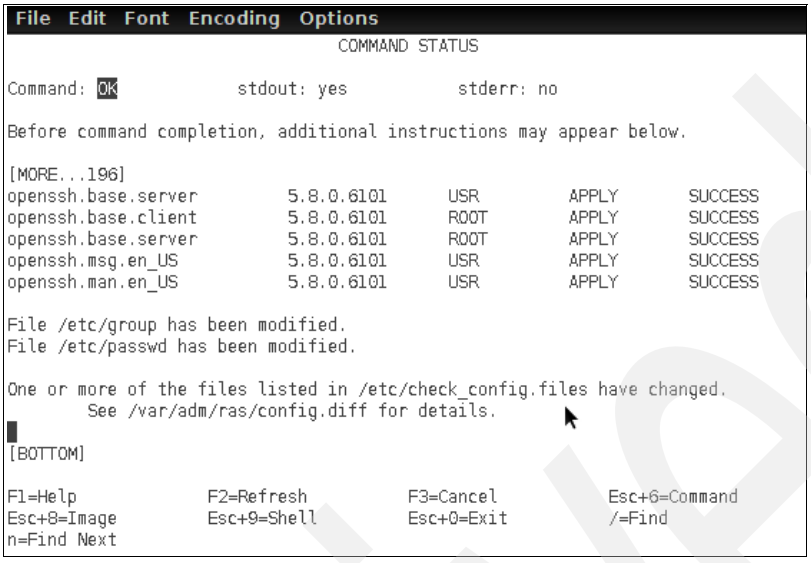

Figure 4-118 AIX Install: Waiting for Command OK

Back on the Installation Assistant menu, select **Tasks Completed - Exit to Login** (Figure 4-119), and then press Enter to get to the login prompt. Login as root, using the password you set previously.

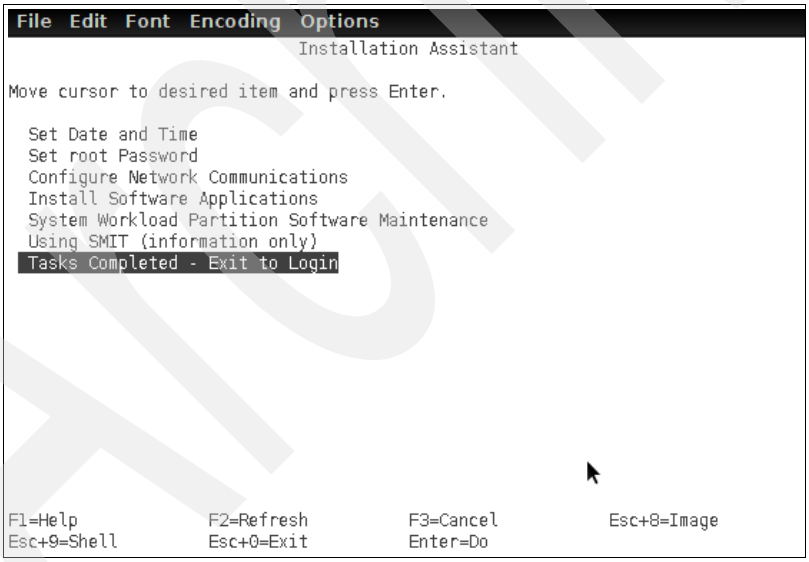

Figure 4-119 AIX Install: Tasks completed

Figure 4-120 shows that login is successful.

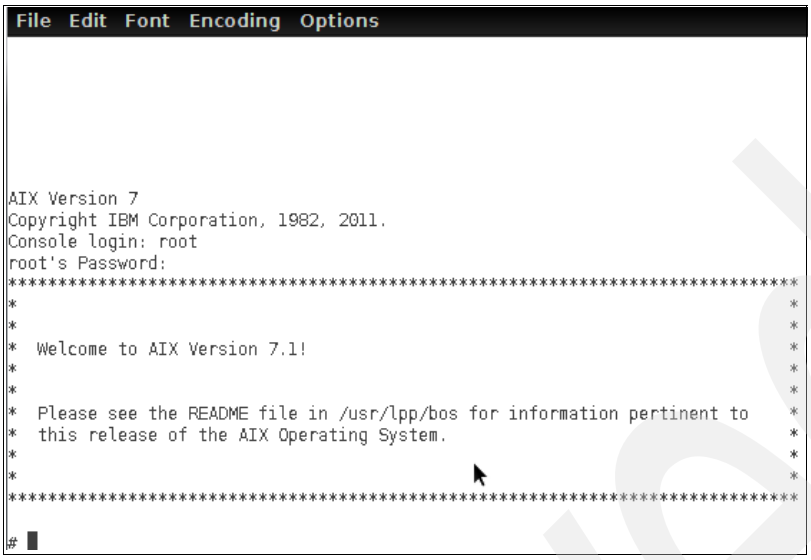

Figure 4-120 AIX Install: Logging in to AIX

To confirm that the network settings are correct, use the **ifconfig -a** command, as shown in Figure 4-121. Then, ping the gateway.

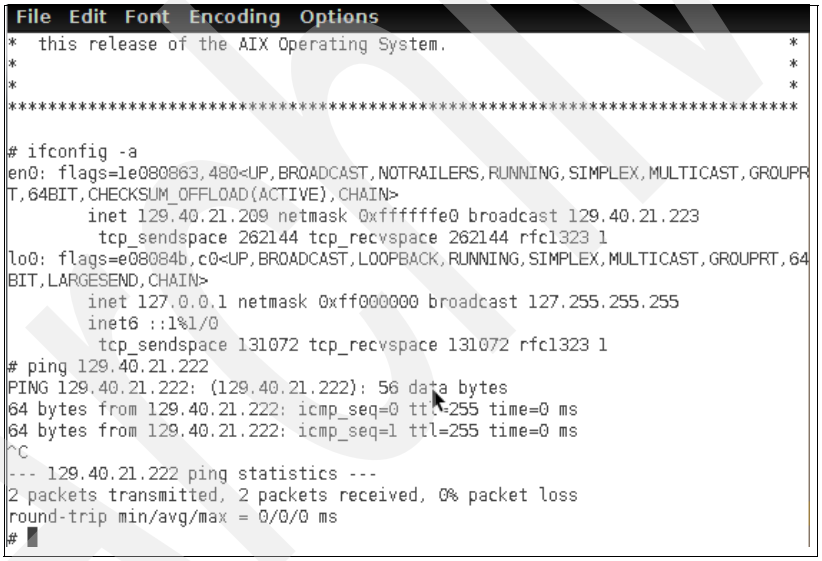

Figure 4-121 AIX Install: Confirm IP config

4.6.5 Define virtual server to FSM

After the AIX operating system is installed on the virtual server, the FSM must be notified that it exists. This is a multistep process:

1. Discover the operating system through the IP address (Figure 4-122). The OS is listed in the Discovered Systems table.

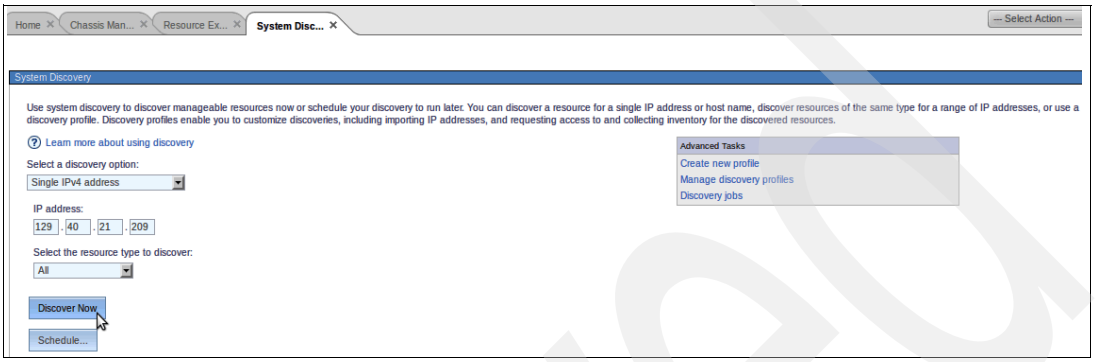

Figure 4-122 Discovery

2. Grant access, as shown in Figure 4-123. Click **No Access** link.

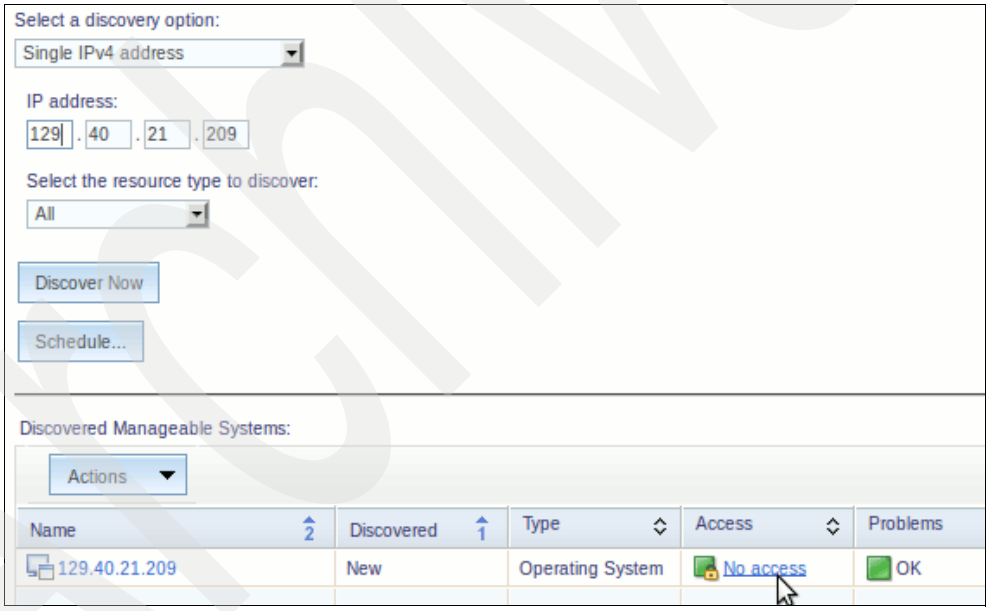

Figure 4-123 Access

3. Click **Request Access** (Figure 4-124). Enter root and root's password.

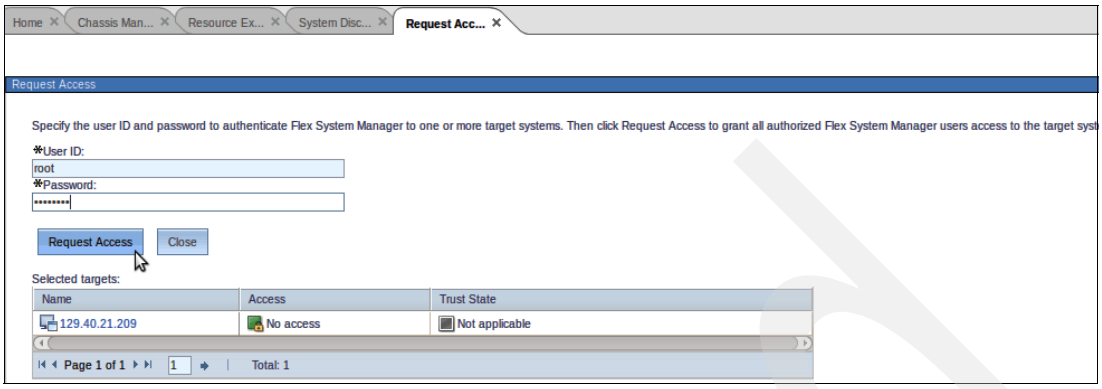

Figure 4-124 Request Access

4. Wait for access to change to OK, as shown in Figure 4-125.

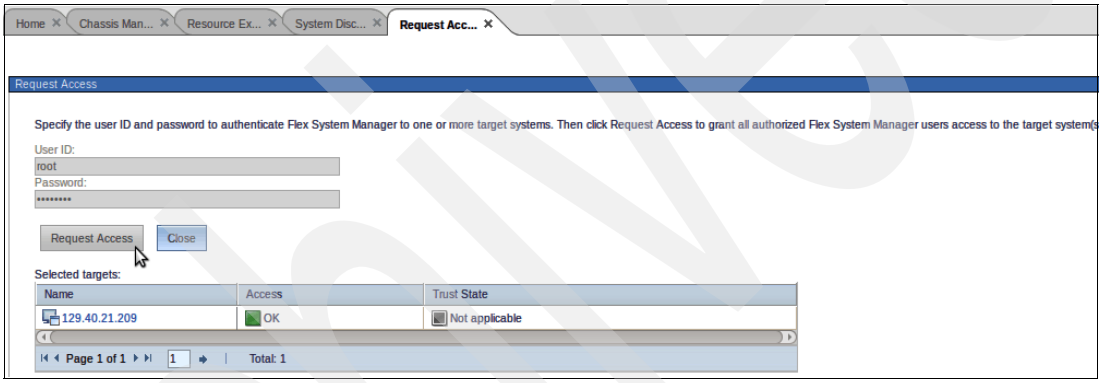

Figure 4-125 Request Access: Access Column changed to OK

5. From the Resource Explorer, select both the new Operating System and Virtual Server, and then collect the inventory by selecting **Actions** → **Inventory** → **Collect Inventory** (Figure 4-126).

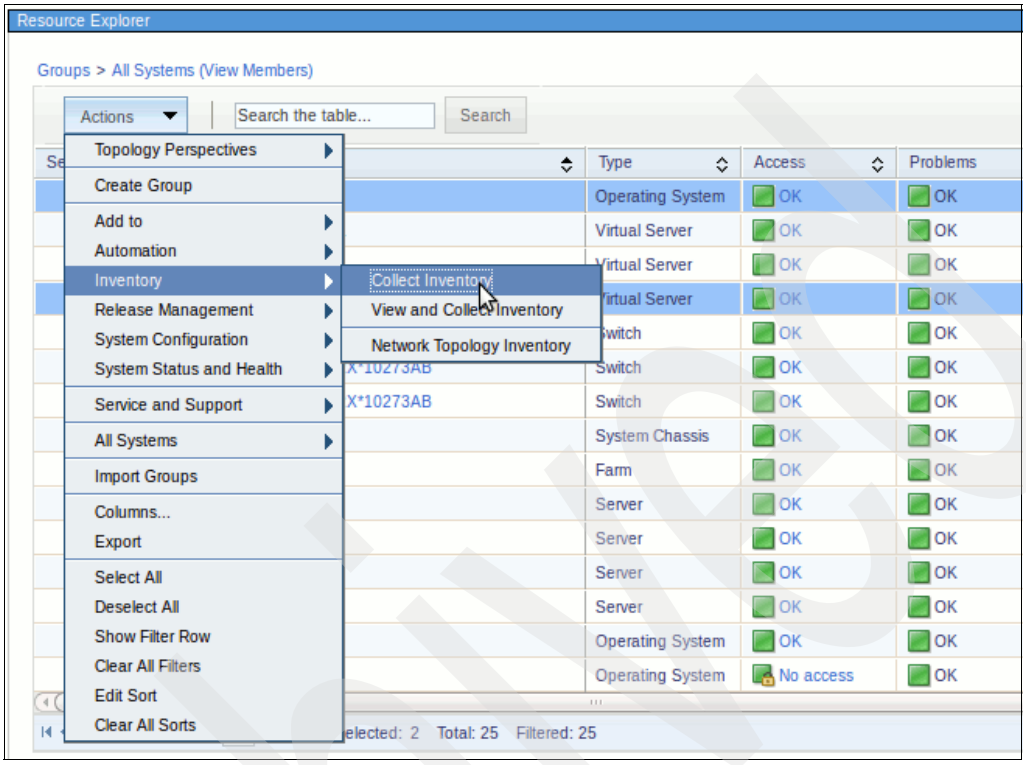

Figure 4-126 Resource Explorer: Collect Inventory

The inventory job starts (Figure 4-127).

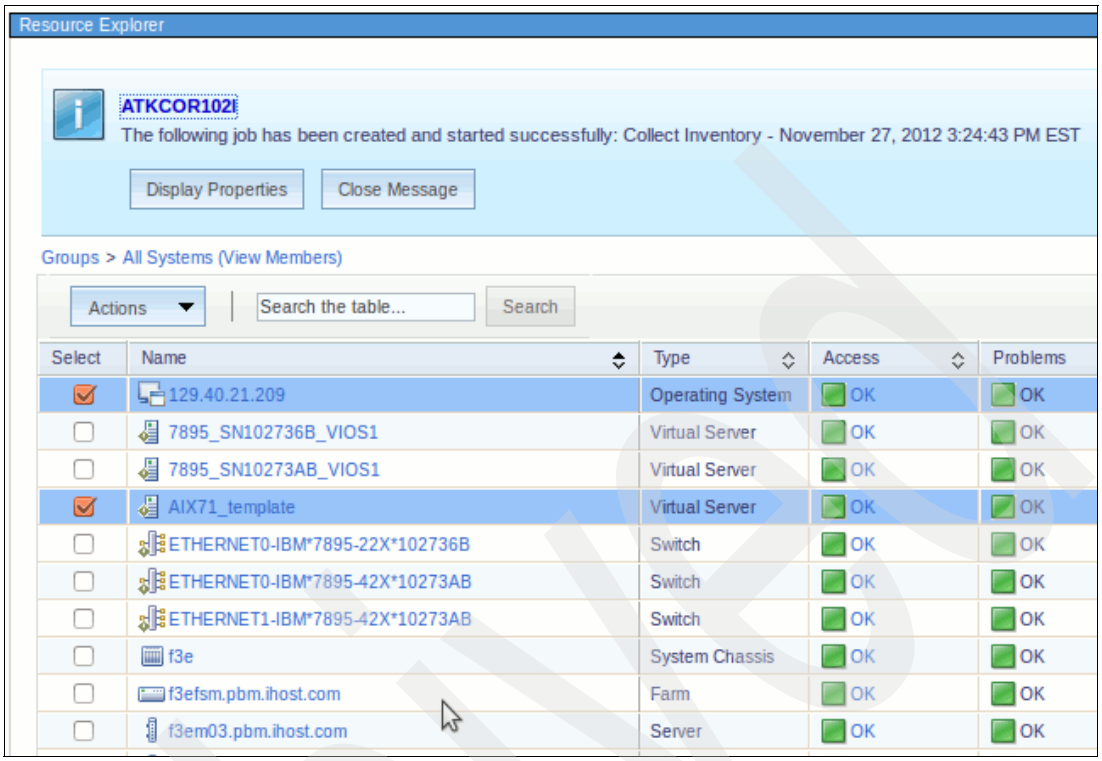

Figure 4-127 Inventory job started

6. In addition, collect inventory on the V7000, FC Switches, the POWER Server itself and the farm.

4.6.6 Install Activation Engine

To enable personalization of each of the workloads that will be deployed from this appliance, a code set must be installed. This code is located on the Flex Service Manager and is called the Virtual Server Activation Engine (VSAE).

To transfer the VSAE package from the FSM to the new AIX use the **scp** command as follows:

```
# scp USERID@129.40.21.201:/opt/ibm/director/proddata/activation-engine /tmp
```

Although many examples in this book use the GUI for the FSM, by way of the management IP address on the management subnet, that IP address is not accessible from the AIX server to which the activation engine is being copied. The data network's IP address that was configured on FSM's eth1 (129.40.21.201) must be used. This is on the same subnet as the address with which the AIX server's en0 interface is configured (129.40.21.209).

Extract the TAR file. Set the JAVA_HOME environment variable, then run the **aix-install.sh** script shown in Figure 4-128.

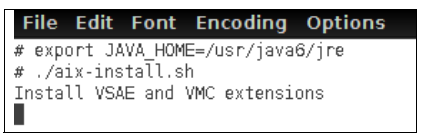

Figure 4-128 Installing VSAE

After the VSAE package is installed, run the following commands:

Before issuing commands: The following commands shut down the virtual server. Be sure that the virtual server was properly discovered and inventoried by the FSM (described in 4.6.5, “Define virtual server to FSM” on page 84) before issuing these commands.

```
rm /opt/ibm/ae/AP/*
cp /opt/ibm/ae/AS/vmc-network-restore/resetenv /opt/ibm/ae/AP/ovf-env.xml
/opt/ibm/ae/AE.sh --reset
```

Verify that the activation engine was activated. In Figure 4-129, the `OVERALL STATUS for 'vs0': Not Activated` message is displayed.

```
File  Edit  Font  Encoding  Options
[2012-11-27 15:42:27,095] INFO: [0@1354052547.1] OVERALL STATUS For 'vs0': Not A
ctivated
[2012-11-27 15:42:27,118] INFO: Resetting operating system
[2012-11-27 15:42:27,131] INFO: Cleaning AR and AP directories
[2012-11-27 15:42:27,141] INFO: Shutting down the system
```

Figure 4-129 VSAE not successfully activated

If this message occurs, restart the OS and run the three commands again. Figure 4-130 shows the result of successful execution.

```
File  Edit  Font  Encoding  Options
[2012-11-27 15:47:16,253] INFO: [0@1354052836.25] OVERALL STATUS For 'vs0': Succ
essful
[2012-11-27 15:47:16,275] INFO: Resetting operating system
[2012-11-27 15:47:16,289] INFO: Cleaning AR and AP directories
[2012-11-27 15:47:16,299] INFO: Shutting down the system


Broadcast message from root@AIX71_template (tty) at 15:47:17 ...

shutdown: PLEASE LOG OFF NOW !!!
System maintenance is in progress.
All processes will be killed now.
```

Figure 4-130 VSAE reset successful

4.6.7 Capture virtual server

With the new virtual server template fully configured, it can now be captured. The capture process involves creating a set of metadata in the FSM and using FlashCopy to clone the virtual servers storage:

1. In the Resource Explorer, scroll until the virtual server is listed (Figure 4-131).

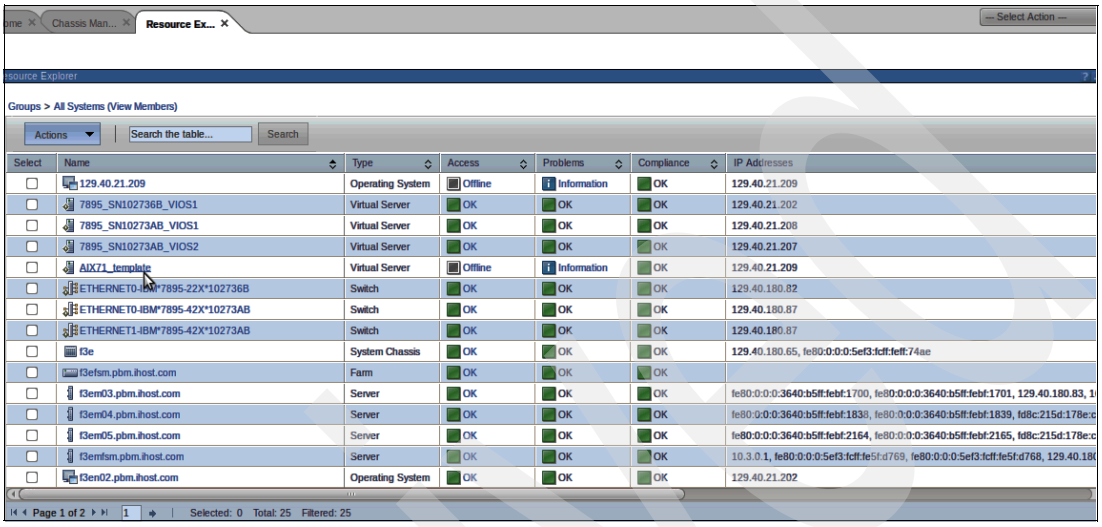

Figure 4-131 List of virtual servers

2. Select this server and click **Actions** → **System Configuration** → **Capture** (Figure 4-132).

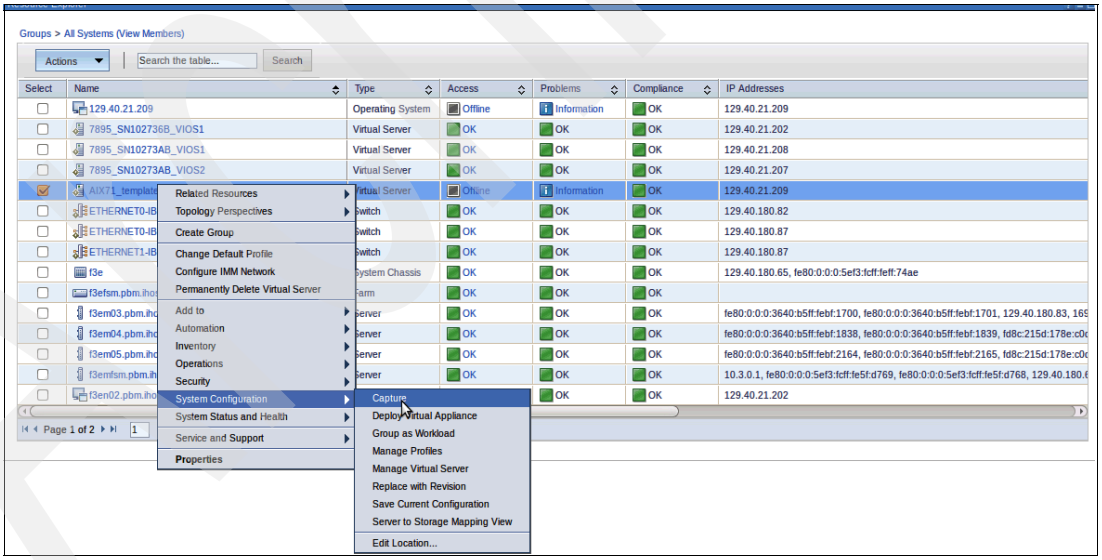

Figure 4-132 Resource Explorer: Select Capture

The Capture wizard starts (Figure 4-133).

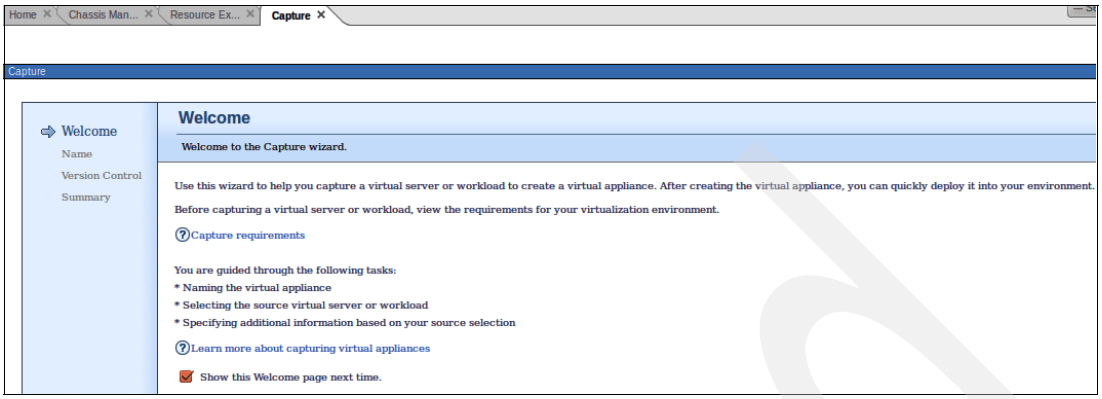

Figure 4-133 Capture wizard: Welcome

3. The first step in capturing the template is to give it an appropriate name (Figure 4-134). This virtual server template will be used as an appliance, and contains AIX version 7.1.

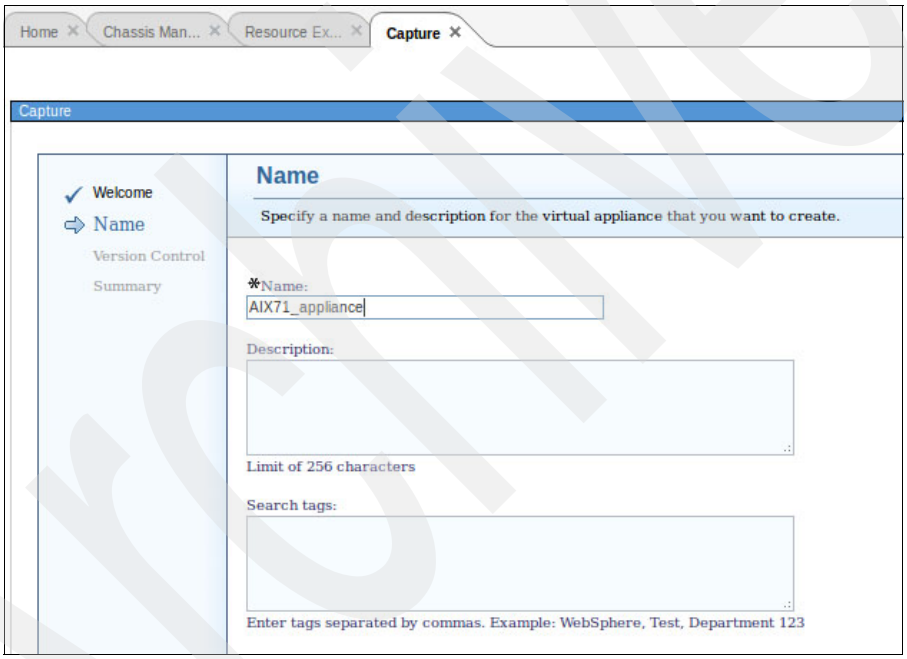

Figure 4-134 Capture Wizard: Specify name

4. The next step is to select a disk to capture (Figure 4-135). Because only one volume was created on the V7000 for this virtual server, only one is listed.

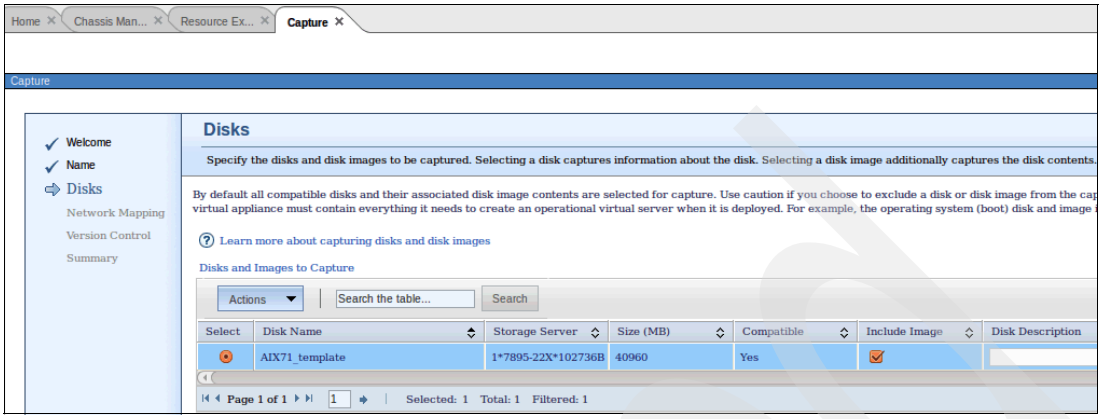

Figure 4-135 Capture Wizard: Specify disks

5. The Network Mapping is next (Figure 4-136). These two networks are the same as configured in 4.3, “Networking setup” on page 27. In the following figures, they are shown in numerical order, and not necessarily the order that will be used for the captured appliance.

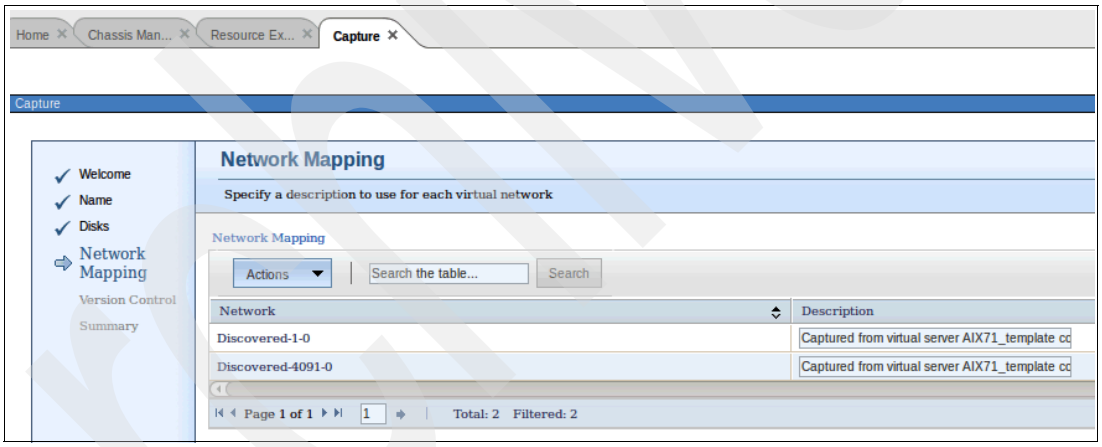

Figure 4-136 Capture Wizard: Network Mapping

This appliance is the first version. If, at a later date, the base virtual server is changed, a new capture must be processed to create a new version of this appliance. This allows for the service management process to be included and for maintenance to be applied. See Figure 4-137.

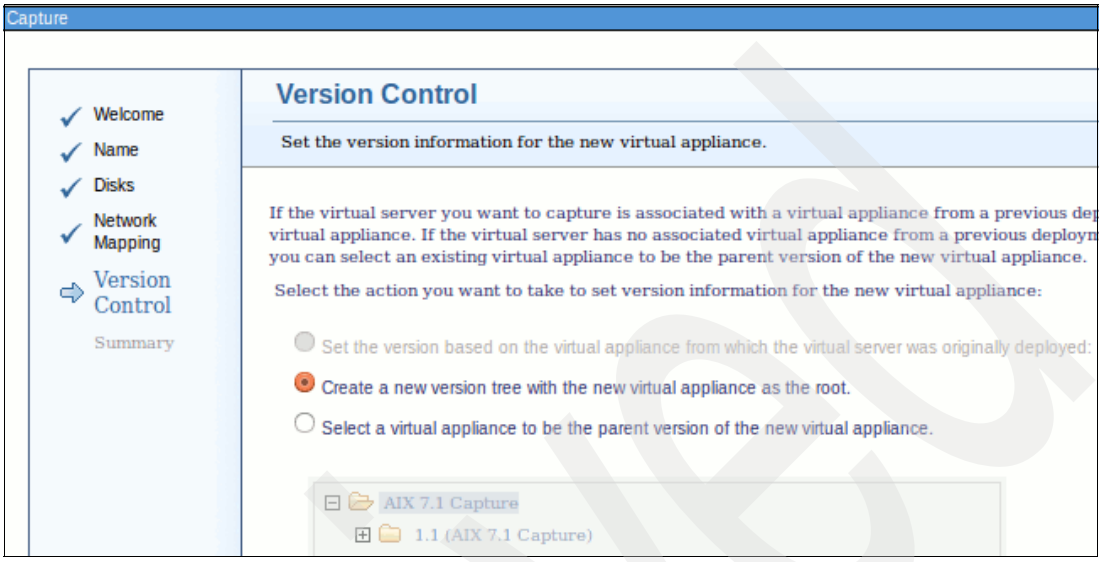

Figure 4-137 Capture Wizard: Version Control

The Summary window opens (Figure 4-138) and a new job is created to capture this virtual server.

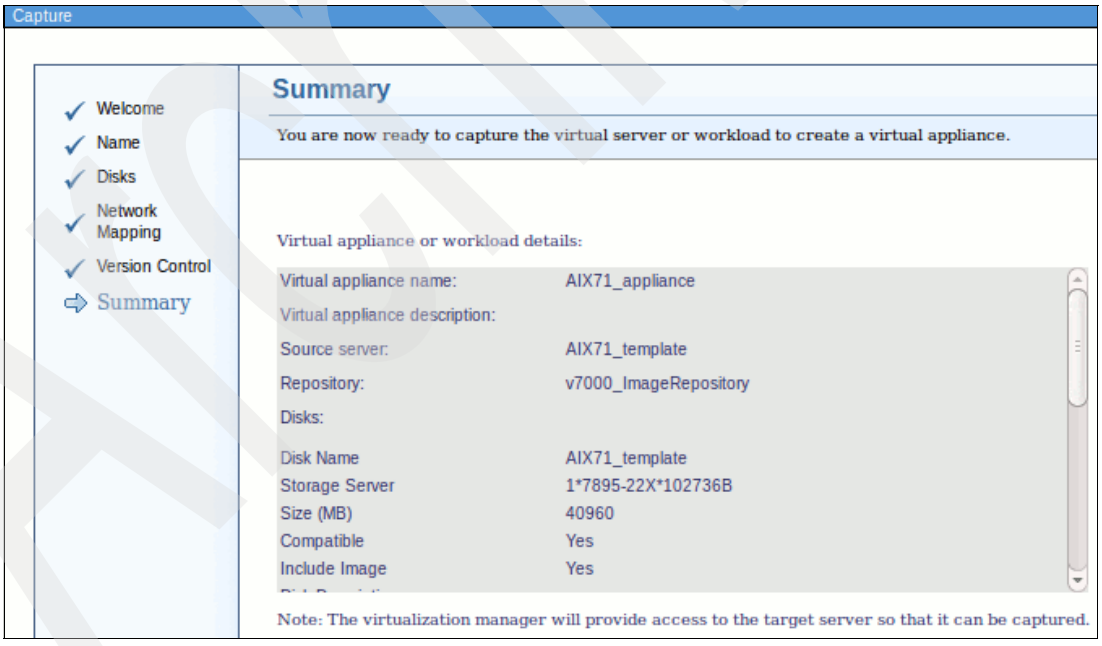

Figure 4-138 Capture Wizard: Summary

The job log (Figure 4-139) shows that the capture is complete and was successful.

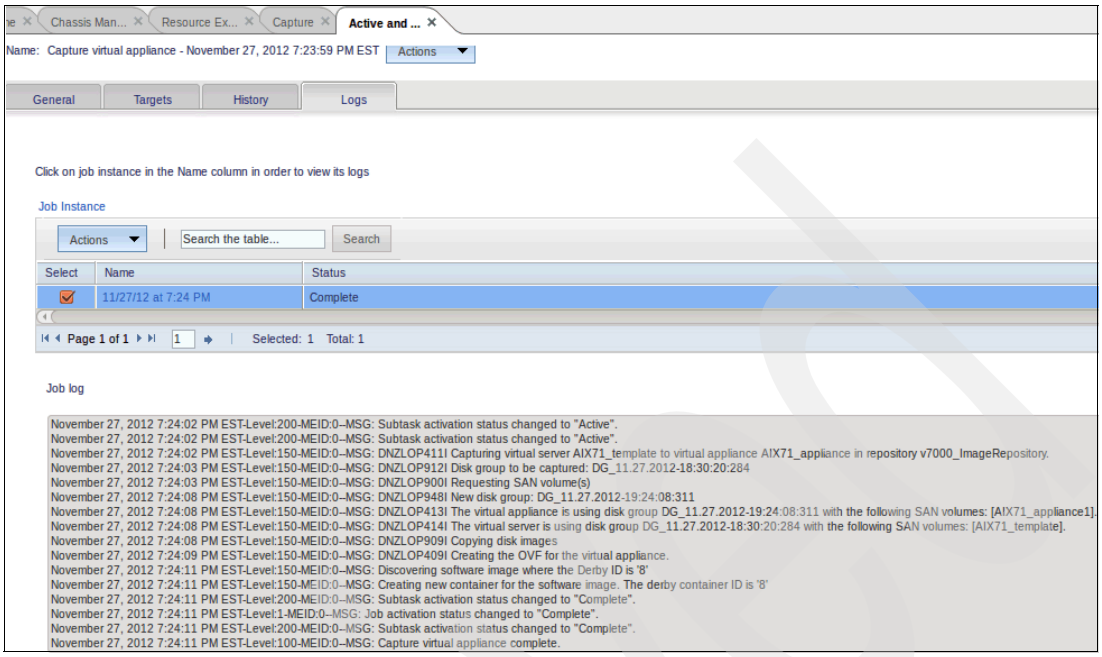

Figure 4-139 Capture virtual appliance complete: Job Log

4.6.8 Deploy virtual server

The previous section created a new appliance for use with SmartCloud Entry. A good approach is to test the deployment of this appliance in the Flex System Manager to ensure that everything is configured properly. From the VMControl panel, select the Virtual Appliances tab (Figure 4-140).

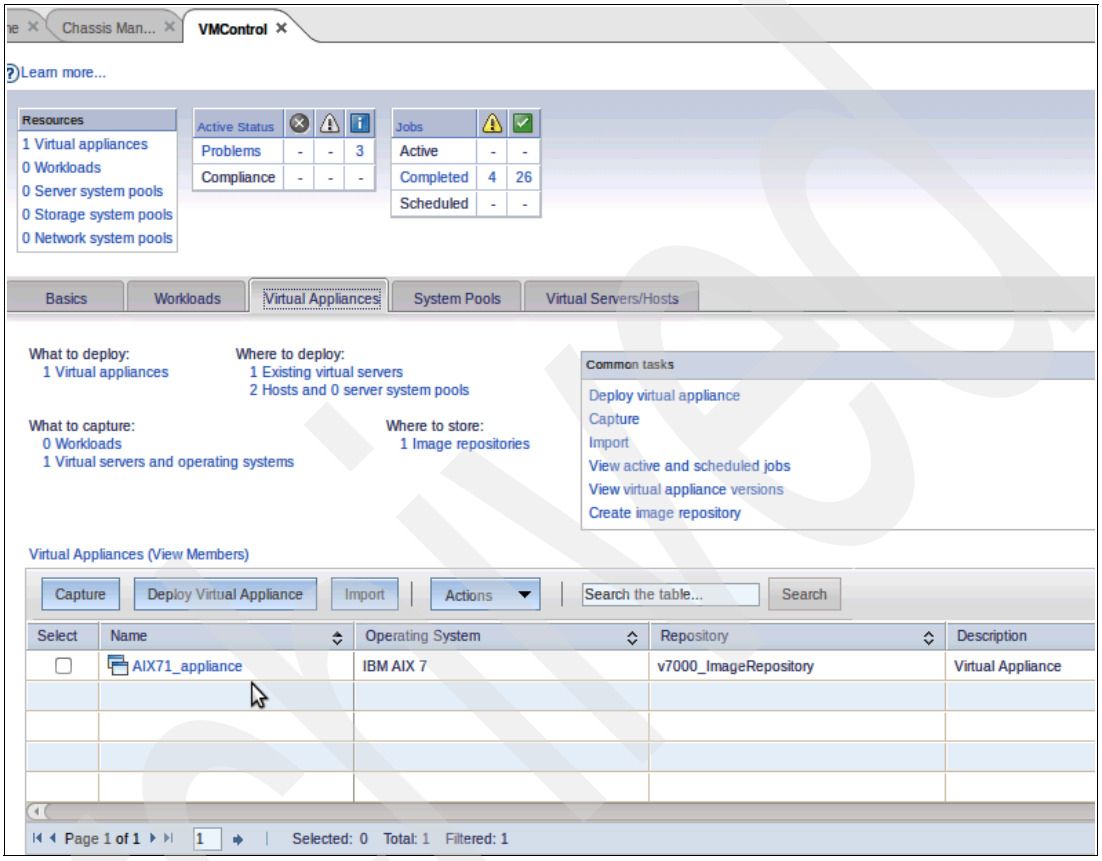

Figure 4-140 VMControl: Virtual appliances tasks

Select the appliance and click **Deploy Virtual Appliance** (Figure 4-141).

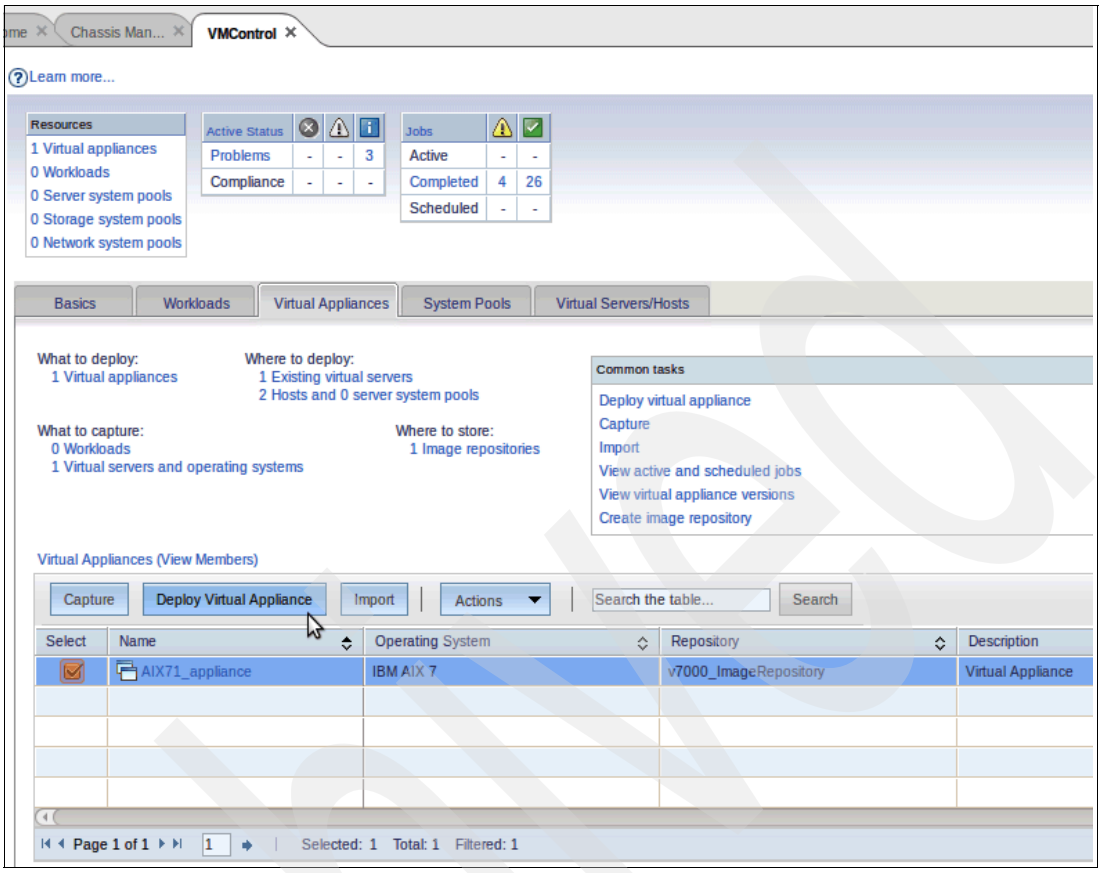

Figure 4-141 VMControl: Choosing Deploy Virtual Appliance

The Deploy Virtual Appliance wizard is shown in Figure 4-142.

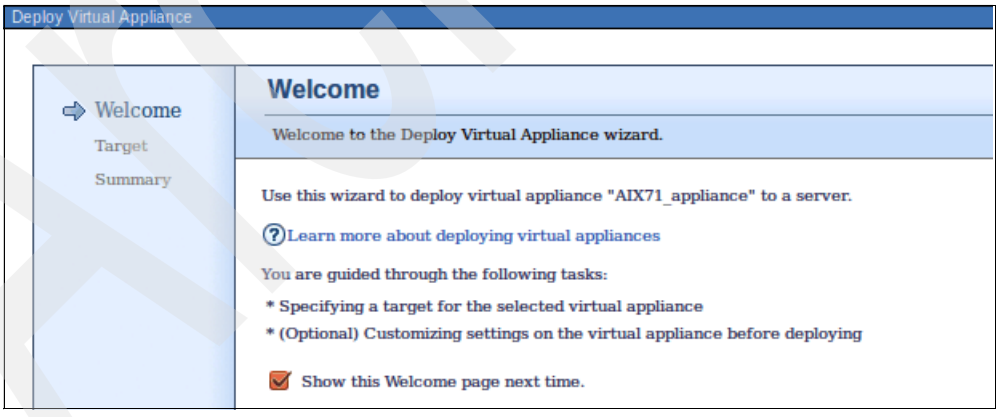

Figure 4-142 Deploy Virtual Appliance wizard: Welcome

Select the target from a list, based on the number of POWER compute nodes installed and configured in the chassis. See Figure 4-143.

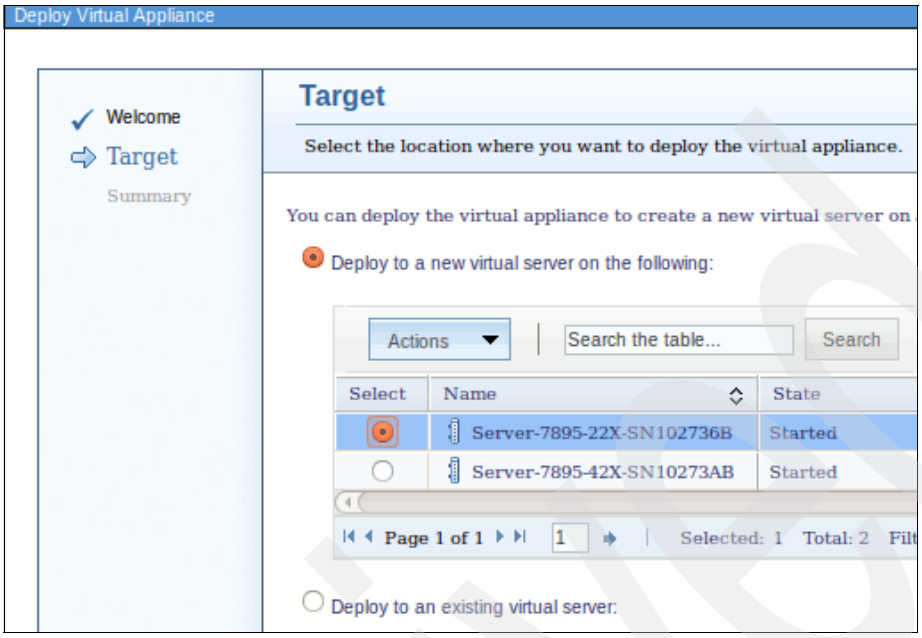

Figure 4-143 Deploy Virtual Appliance wizard: Target

This appliance will be deployed as a workload. Give any appropriate name to this workload (Figure 4-144). As with the name of the appliance, this name helps you more easily locate the workload.

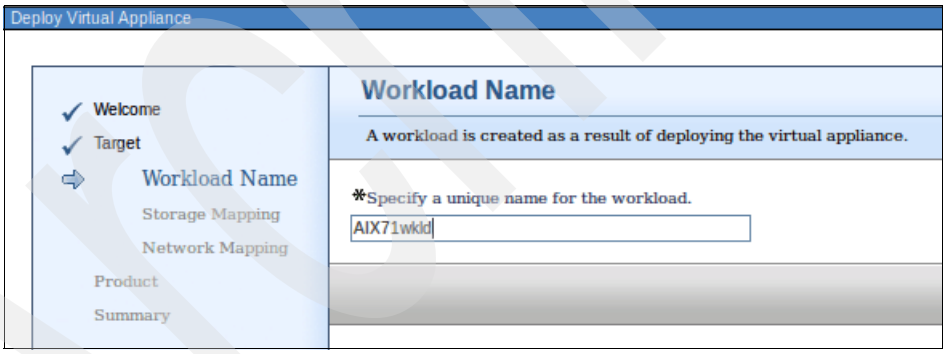

Figure 4-144 Deploy Virtual Appliance wizard: Workload Name

To ensure that this workload has the correct disks, and that these disks are allocated from the proper pool, storage mapping is the next step. The single disk is selected. Figure 4-145 shows that the Assigned Storage column indicates `Not assigned`.

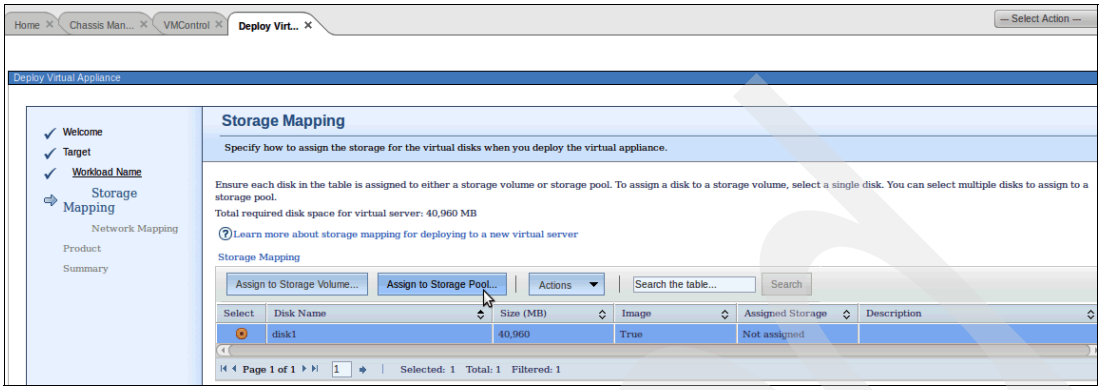

Figure 4-145 Deploy Virtual Appliance Wizard: Storage Mapping

The storage that represents the V7000 configured in 4.5.1, “V7000 setup” on page 51 is selected (Figure 4-146). Click **OK**.

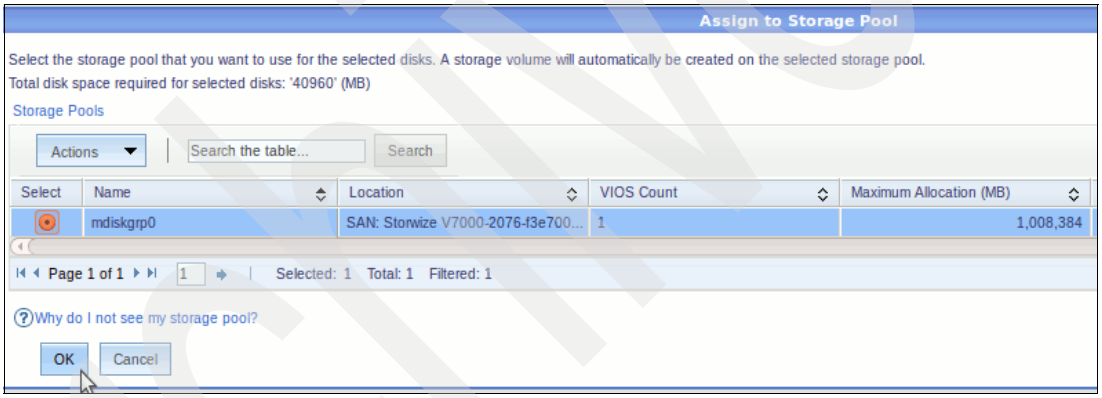

Figure 4-146 Deploy Virtual Appliance wizard: Assign to Storage Pool

Now that the storage is assigned, the Assigned Storage column is populated (Figure 4-147).

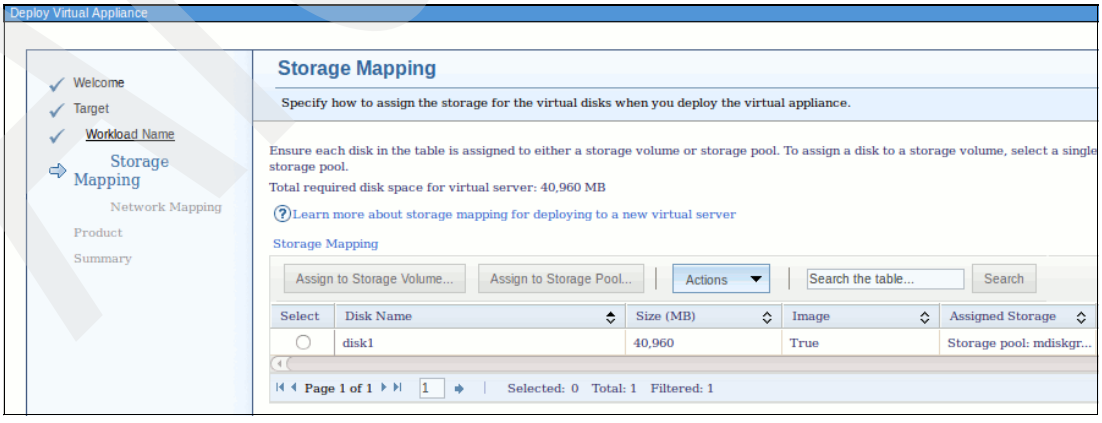

Figure 4-147 Deploy Virtual Appliance wizard: Storage Mapping

Configure Network Mapping next (Figure 4-148). Keep the defaults because they map to the correct VLANs.

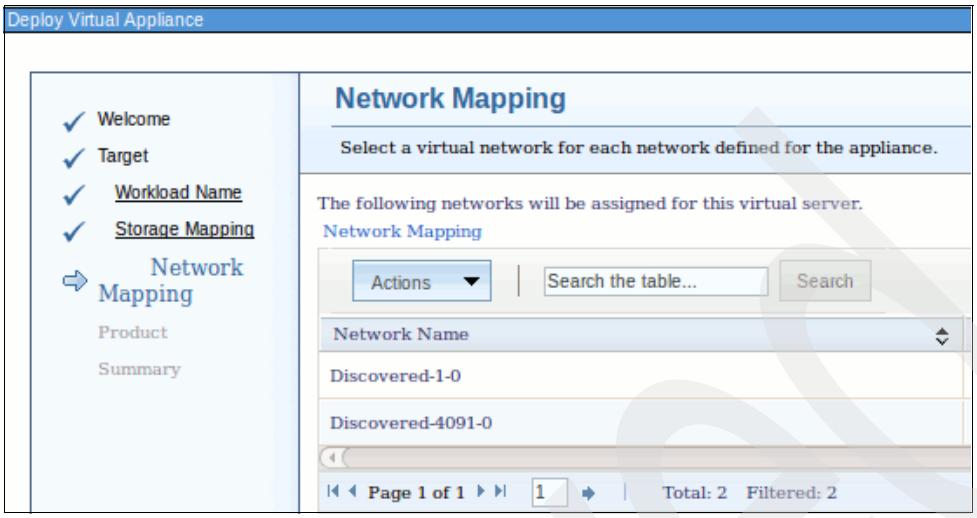

Figure 4-148 Deploy Virtual Appliance wizard: Network Mapping

The configuration for this workload must be manually entered on the panel shown in Figure 4-149 on page 99.

Advantages: One of the advantages of IBM SmartCloud Entry is that the users do not need to know this depth of technical detail to request a new virtual server workload. In addition, the network information in this panel is stored and controlled from a pool of network information, which makes management easier.

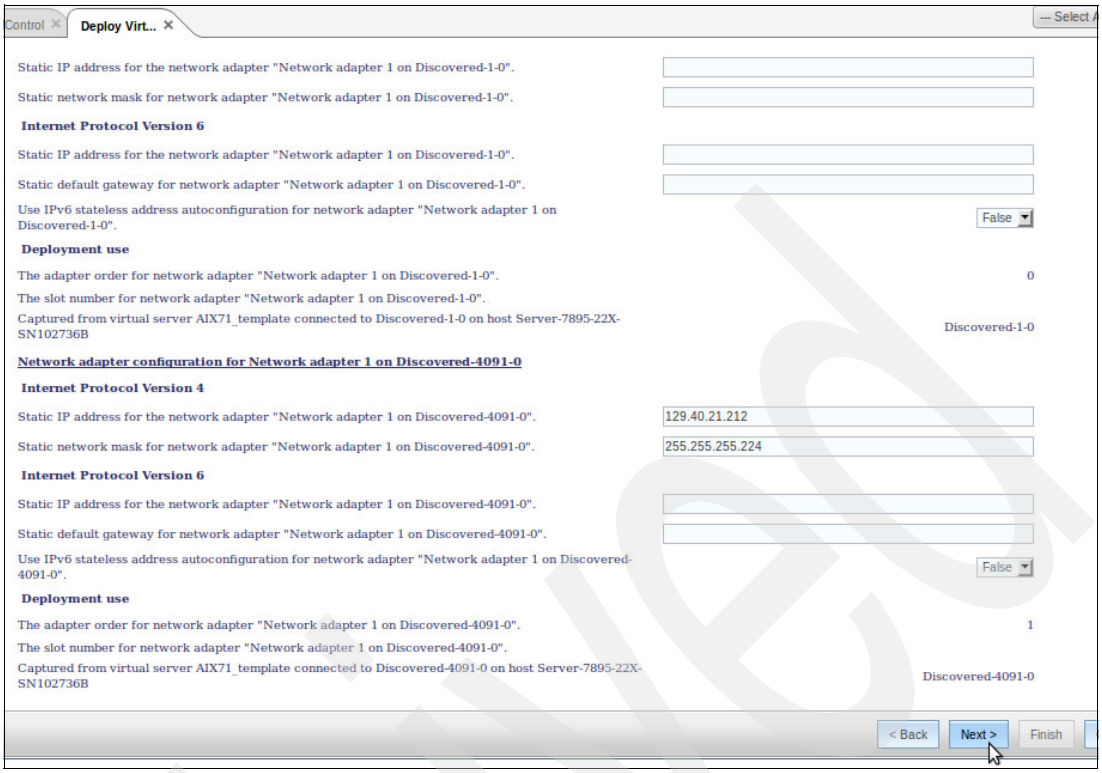

Figure 4-149 Deploy Virtual Appliance Wizard: IPV4 specifications

As with other wizards, the wizard's summary panel opens (Figure 4-150).

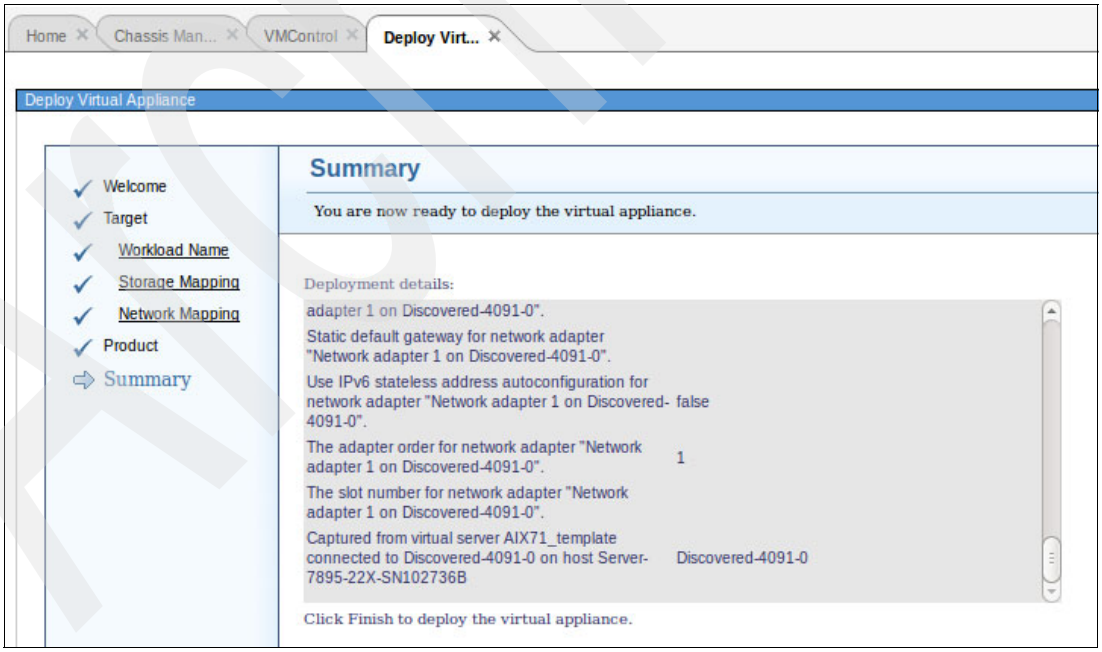

Figure 4-150 Deploy Virtual Appliance wizard: Summary

While watching the logs of this deploy, as shown in Figure 4-151, the output in the log might pause for several minutes. Wait for successful completion.

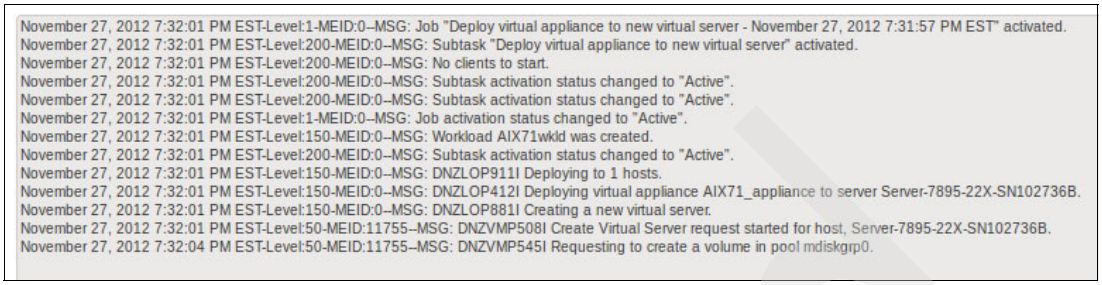

Figure 4-151 Waiting for volume request to complete

The job complete and a message indicates `Deploy virtual server complete` (Figure 4-152).

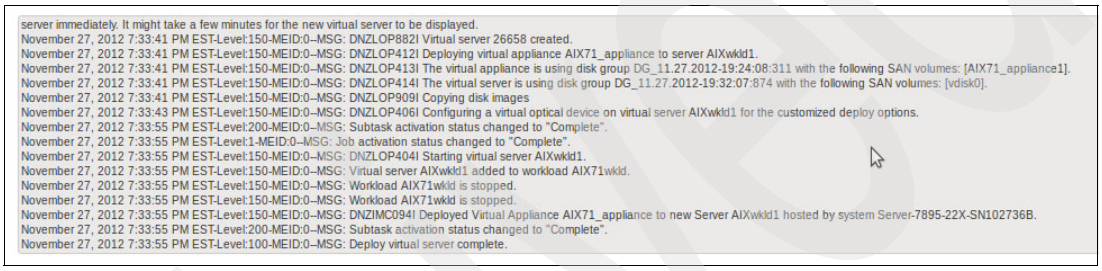

Figure 4-152 Workload deployed

Consider how this process is accelerated by using FlashCopy. As shown in Figure 4-153, the storage subsystem is making a fully asynchronous clone of the initial appliance disk and also the virtual server disk. This process is completely transparent to the workload request and the virtual server is ready for use. The background volume copies (using FlashCopy) continue for a while after the workload is operational.

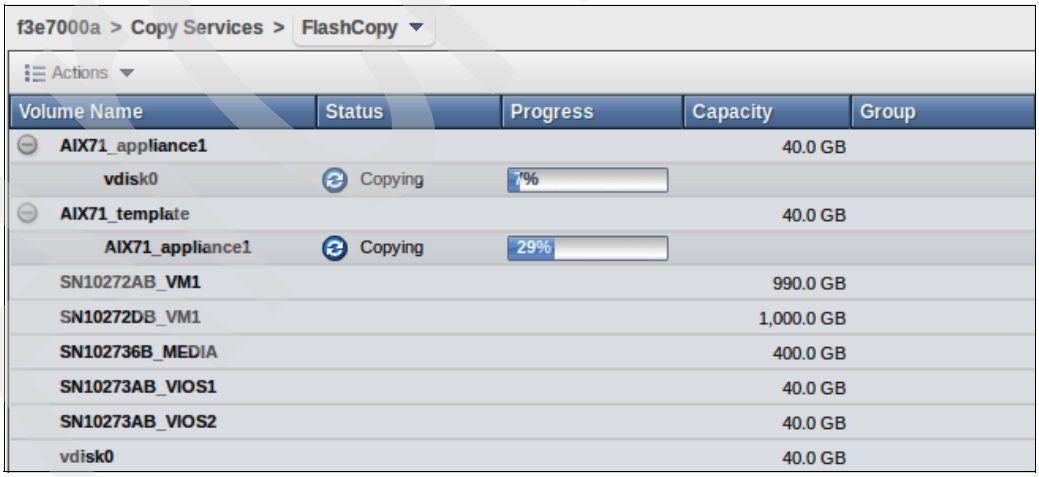

Figure 4-153 Asynchronous cloning using FlashCopy

4.7 The x86 Compute Nodes setup

This section covers the setup of the x86 compute node, the installation of VMware Hypervisor, and the VMware management appliance. To accomplish these goals for each x86 node installed in the chassis, the following tasks must be completed:

- VMware preinstall tasks
- Install VMware ESXi
- Install VMware vSphere client
- Creating the data store
- Deploy and Configure VMware vCenter Server Appliance

4.7.1 VMware preinstall tasks

Similar to the other nodes and modules that plug into the chassis, the first step before managing an x86 compute node it to configure an IP address on the IMM. This node also must be defined in the Flex Service Manager so that it can be properly managed.

From the CMM, select **Chassis Management**, and the **Component IP Configuration**. Click on the node that needs its IP set. On this window, set the proper IP address.

In Figure 4-154, the IMM already has its IP information set.

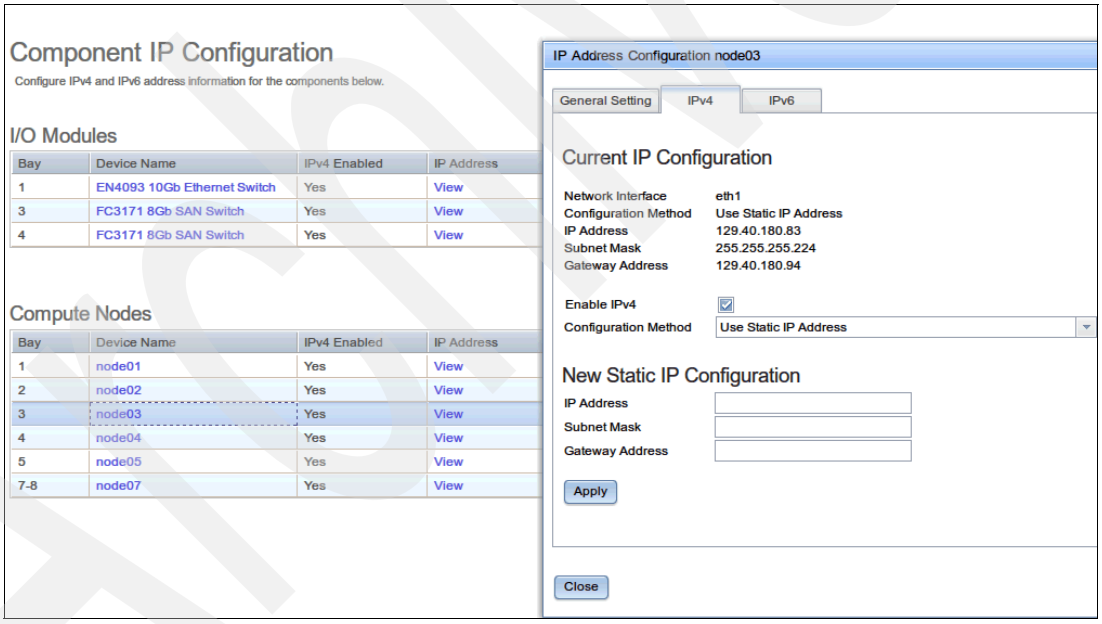

Figure 4-154 IP Address Configuration for node03

This IMM can now be discovered by the Flex Service Manager and inventoried. These steps are not specifically covered in this section, but are similar to 4.6.5, “Define virtual server to FSM” on page 84. The remainder of this section defines the steps to install or upgrade ESXi hypervisor on the internal storage of the node.

A remote console must be started to the compute node from the FSM. This same task can be accomplished from the IMM’s remote control interface, but that is not covered here.

Start the remote console by clicking the node to be controlled in the main chassis window, as shown in Figure 4-155. This is slot 3, the second box from the bottom on the left (outlined) and labeled `f3em03.pbm.ihost.com`.

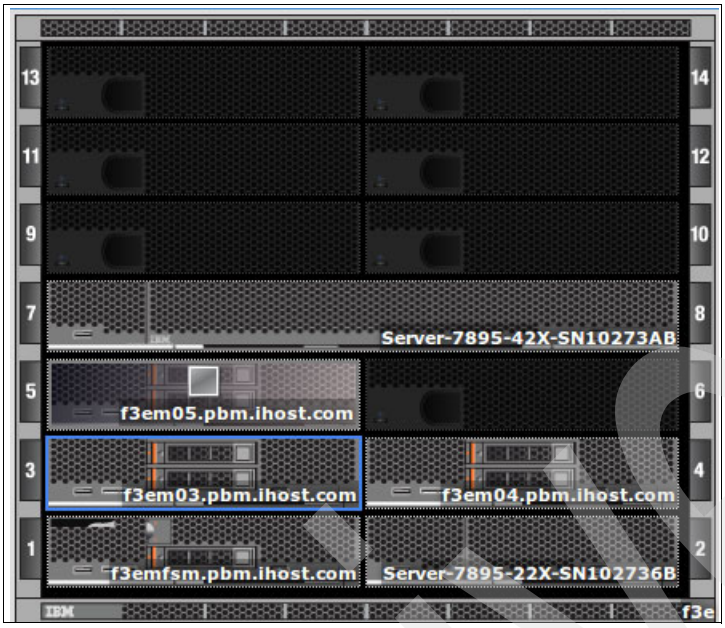

Figure 4-155 Starting remote console using GUI

This action opens a Java-based Remote Control window similar to Figure 4-156.

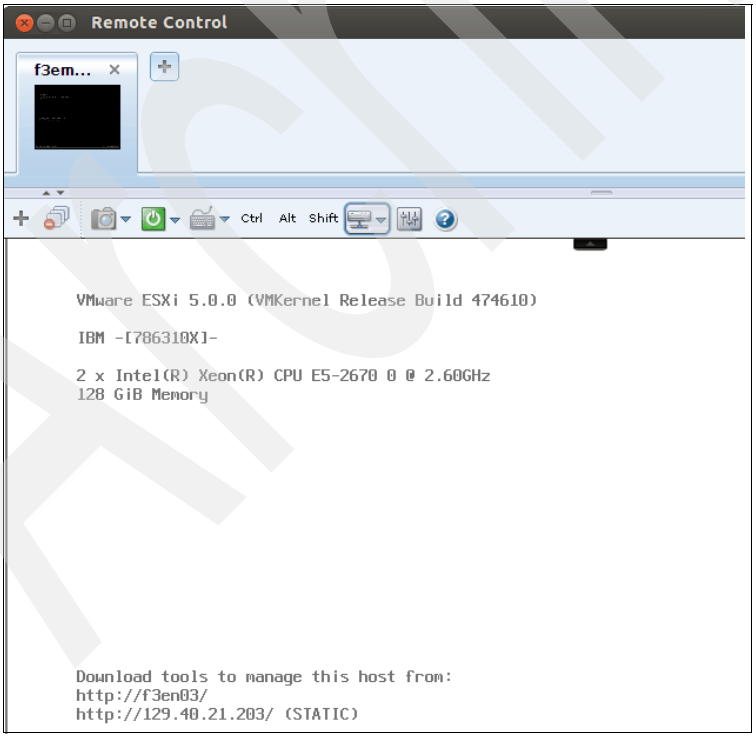

Figure 4-156 Remote console

Notice that ESXi is already installed on this node. For purposes of this section, what is currently installed is irrelevant.

Mount the installation CD by clicking the disk mount icon, and then select **Mount Remote Media** as shown in Figure 4-157.

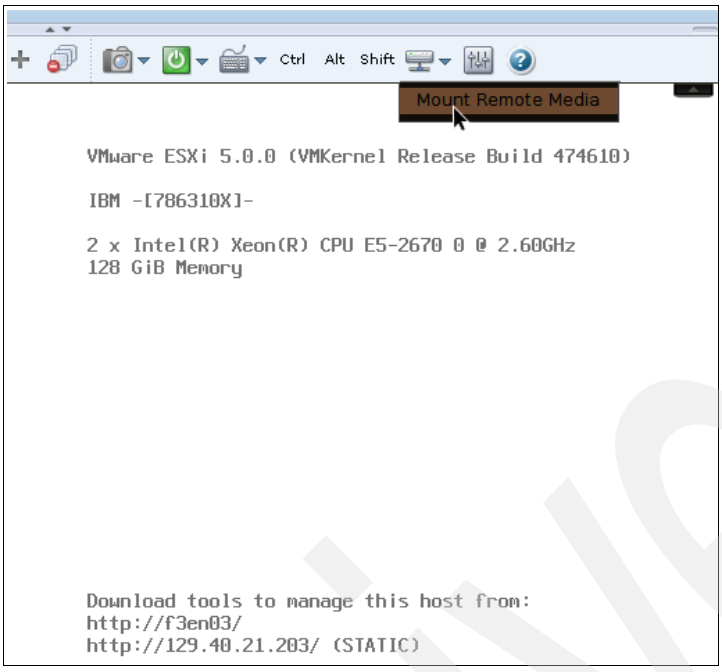

Figure 4-157 Mount Remote Media

The next window (Figure 4-158) indicates the mounting options for this node. It shows there are no Selected Resource, and no drives mapped to the node.

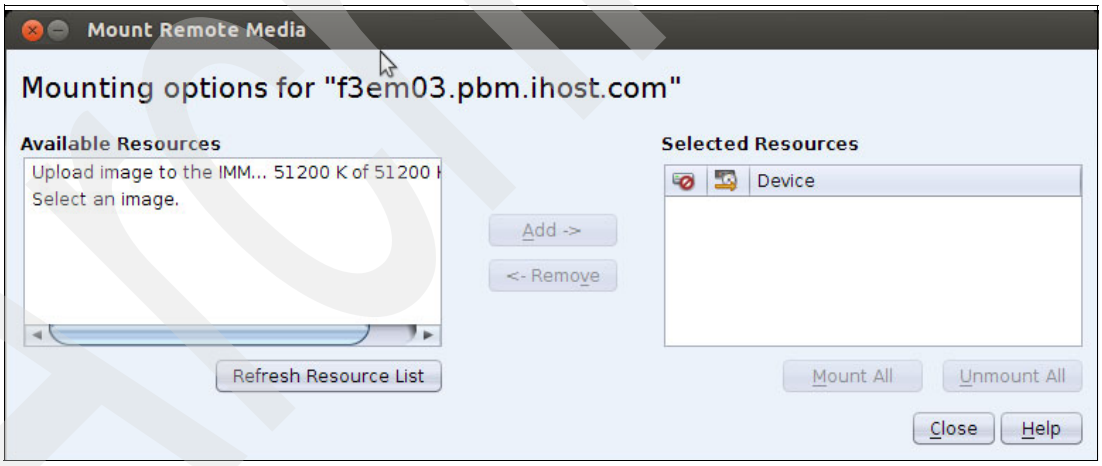

Figure 4-158 Mount Remote Media: Mounting options

To add a new CD based resource, click **Select an image** (Figure 4-159).

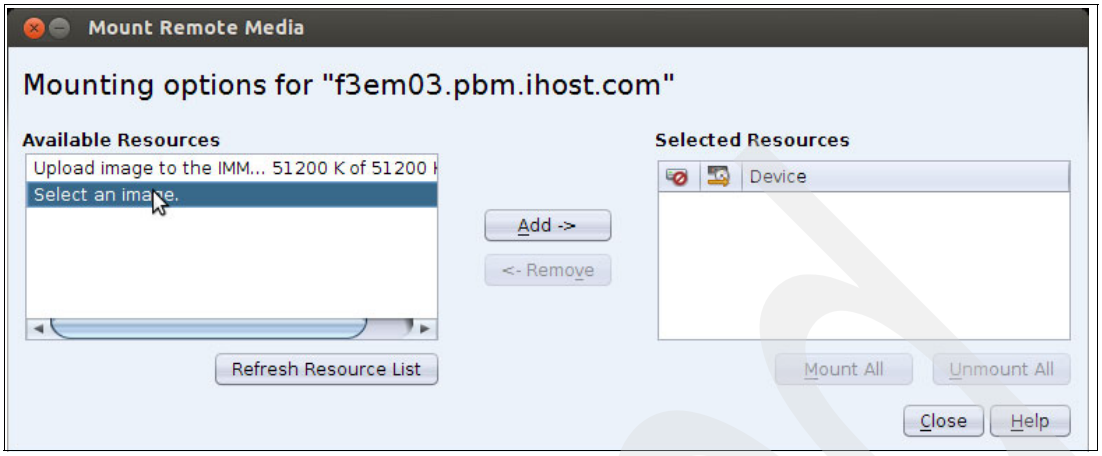

Figure 4-159 Mount Remote Media: Select an image

Locate the CD image of the VMware ESXi to be installed (Figure 4-160) and click **Open**.

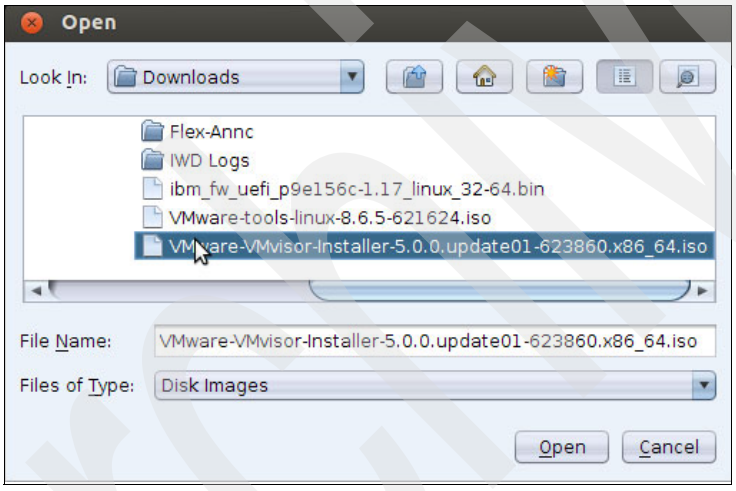

Figure 4-160 Mount Remote Media: Look in downloads

As Figure 4-161 shows, a resource is now listed under Selected Resources. Click **Mount All** to mount this resource to the node.

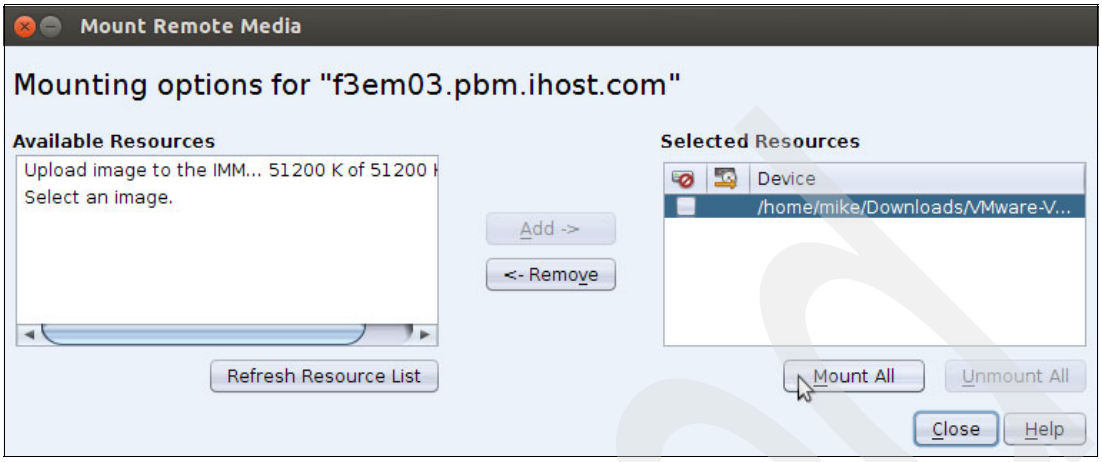

Figure 4-161 Mount Remote Media: Choosing Mount All

The Mount All button is now disabled, but the Unmount All button is enabled (Figure 4-162).

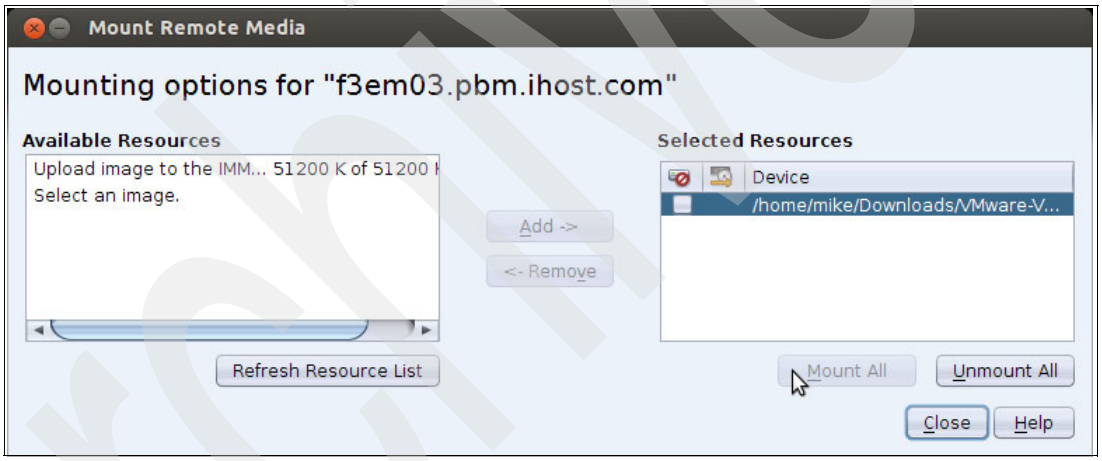

Figure 4-162 Mount Remote Media: Unmount All now available

With the CD mounted, the node can be restarted. Click the power icon and then select **Shut Down OS and Restart** (Figure 4-163).

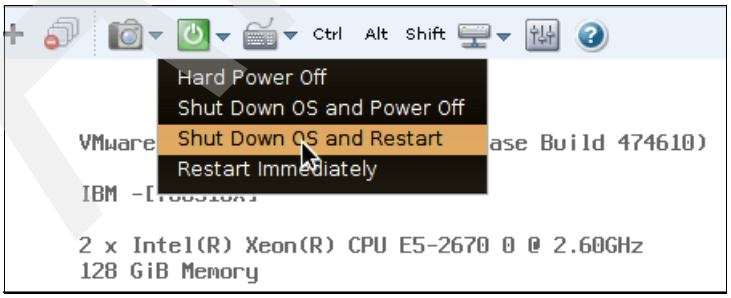

Figure 4-163 Shut Down OS and Restart

4.7.2 Install VMware ESXi

After the node is booted with the CD mounted, the installation programs begin. The VMware ESXi 5.1 welcome screen opens (Figure 4-164). Select **Continue** by pressing the Enter key.

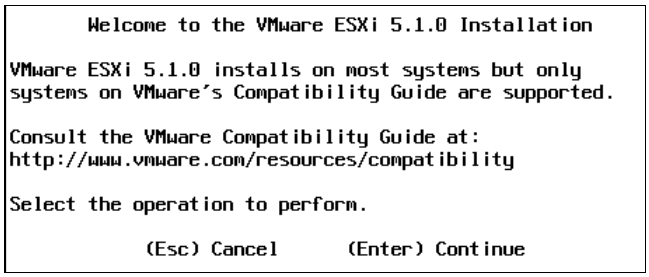

Figure 4-164 VMware ESXi: Installation Welcome

The End User License Agreement (EULA) opens (Figure 4-165). Press F11 to accept the license and continue.

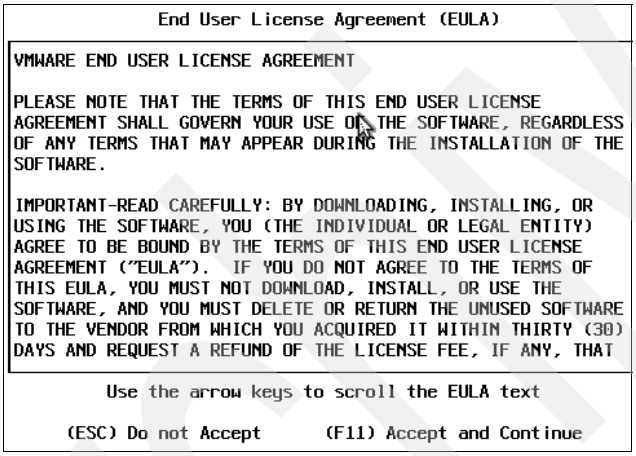

Figure 4-165 VMware ESXi: EULA

The installation program probes the hardware that the compute node has access to. This can take several seconds. See Figure 4-166.

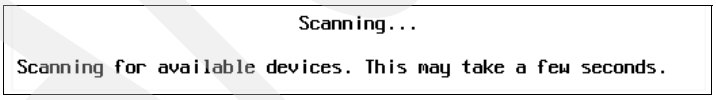

Figure 4-166 VMware ESXi: Scanning for available devices

When the probe is done, the hardware list is displayed. Figure 4-167 shows that three disks were discovered: one local (the USB) and two remote (the storage volumes). This setup might differ for each environment. For this example, the USB disk is chosen.

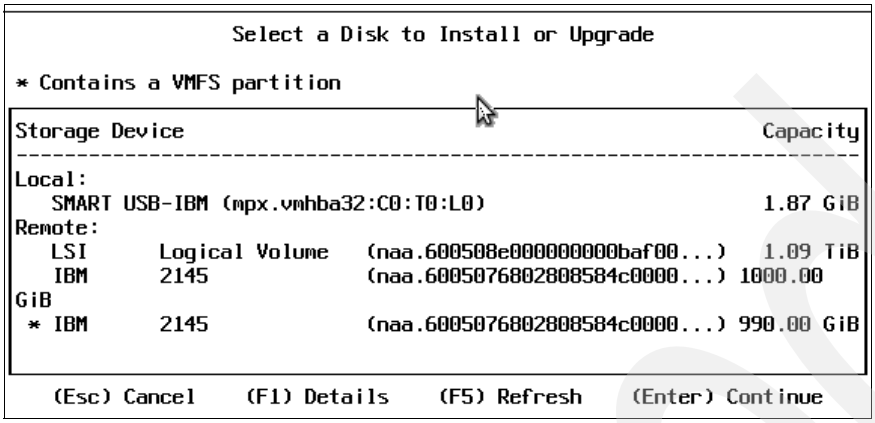

Figure 4-167 VMware ESXi: Select Disk to Install or Upgrade

A new window opens, overlaying the current window. This causes the drive to be scanned to determine whether it meets the requirements for the ESXi installation.

The Select a Disk dialog is overlaid during scanning (Figure 4-168).

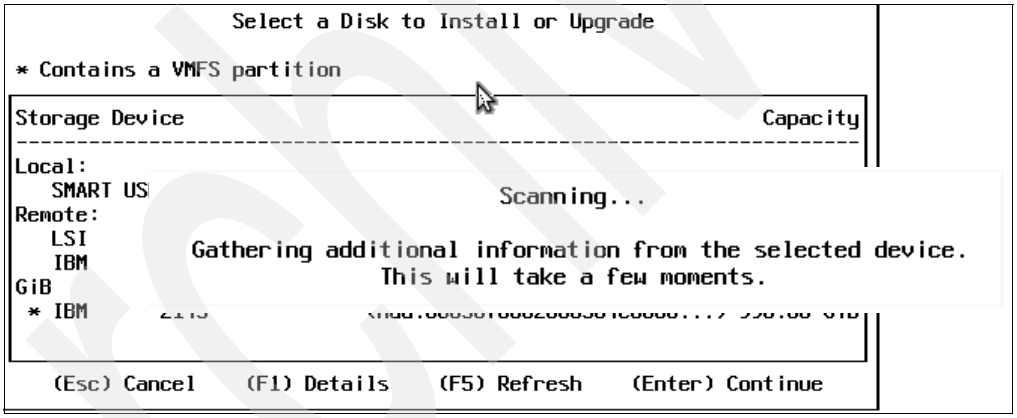

Figure 4-168 VMware ESXi: Gathering information about selected install device

When scanning is complete, a message indicates that the installation found a previously installed version of ESXi (Figure 4-169). The **Install** option was selected to demonstrate the full configuration options.

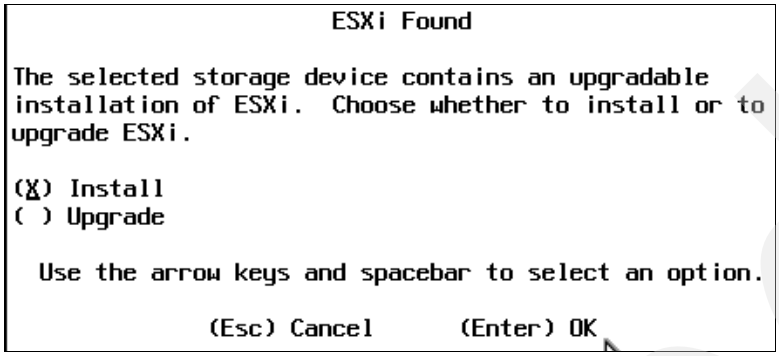

Figure 4-169 VMware ESXi: Storage device Install or Upgrade prompt

Select the keyboard layout that is appropriate. US Default is highlighted in Figure 4-170.

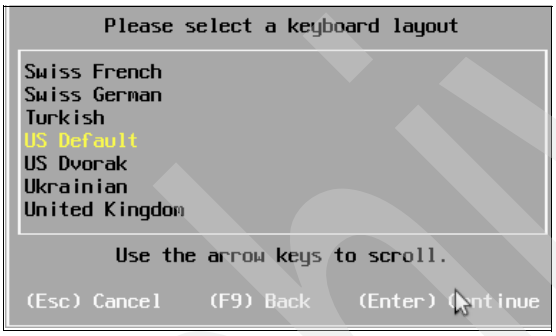

Figure 4-170 VMware ESXi: Choices for keyboard layout (US is default)

Set a root password (Figure 4-171).

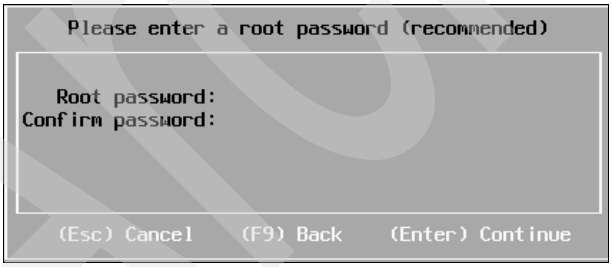

Figure 4-171 VMware ESXi: Root password specification

Figure 4-172 gives a summary of what the installation will do. To prevent unnecessary overwrites, the installation program provides a warning.

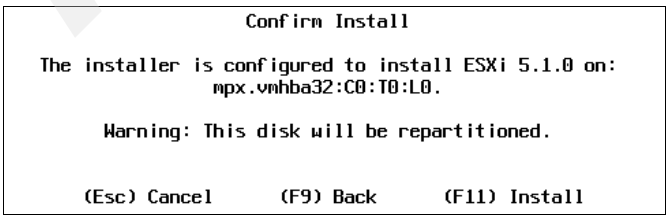

Figure 4-172 VMware Installation: Confirm Install

A progress bar shows the progress of the installation. The Installation Complete screen opens when the installation is successful (Figure 4-173).

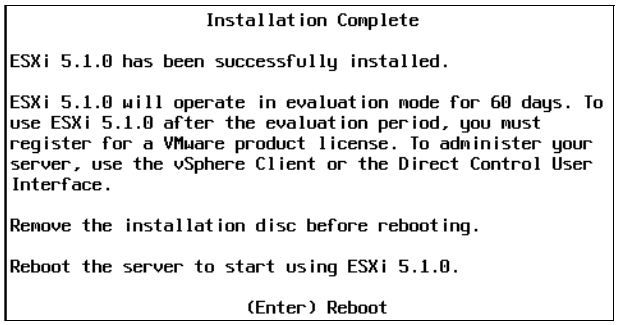

Figure 4-173 VMware Installation: Installation complete (press Enter to reboot)

After the installation completes, click **Unmount All** (Figure 4-174) to unmount the CD.

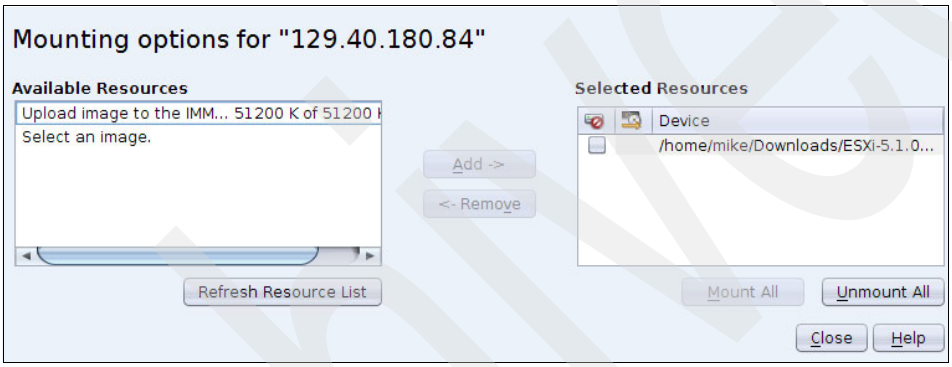

Figure 4-174 Unmounting the CD using Unmount All

After the CD is unmounted, press Enter to reboot. The server reboots (Figure 4-175).

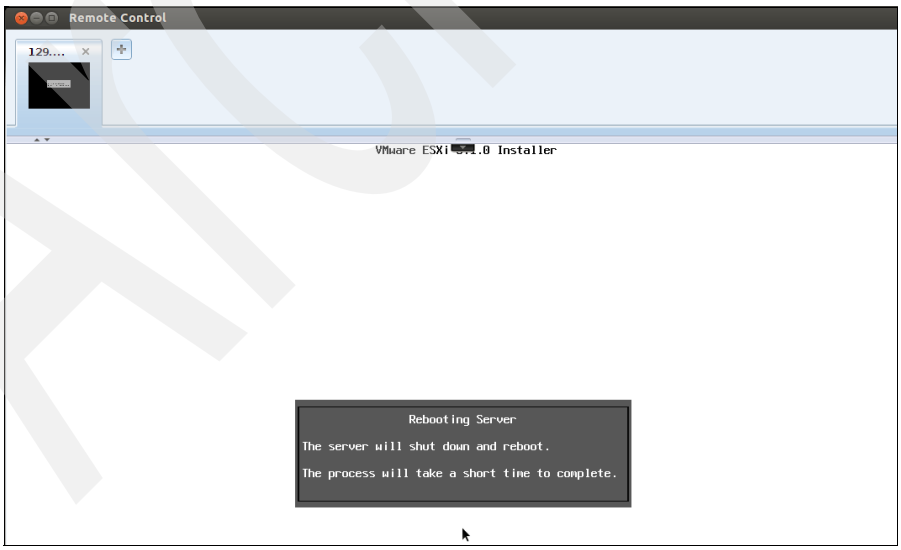

Figure 4-175 VMware ESXi: Rebooting Server

After completing the installation, configure the node so that it can be managed by higher level systems. The main window shows the simple base configuration that is created by default. Press F2 to customize ESXi. See Figure 4-176.

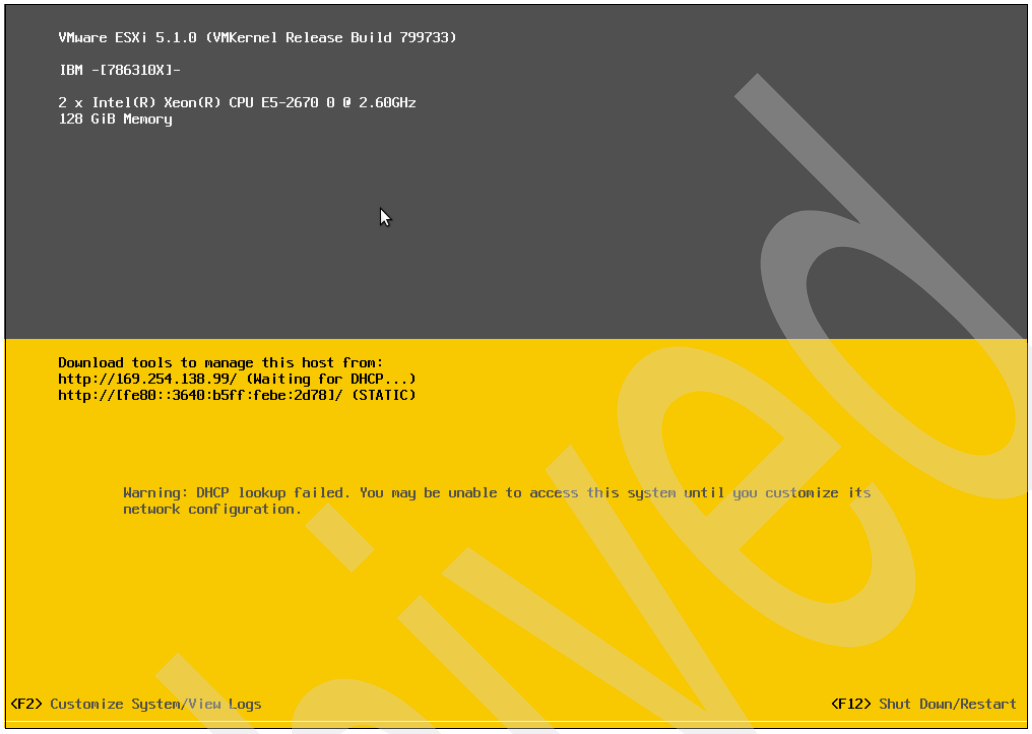

Figure 4-176 Choose F2 to customize system or view logs

Enter the login name (root) and password that was defined earlier (Figure 4-177).

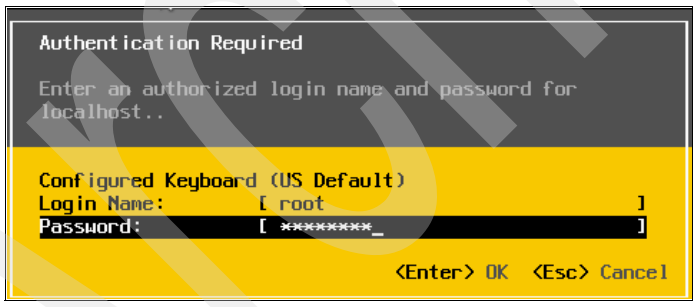

Figure 4-177 Specifying Password to access system configuration

From this screen many items can be configured. Select this network to set an IP address (Figure 4-178). Press Enter.

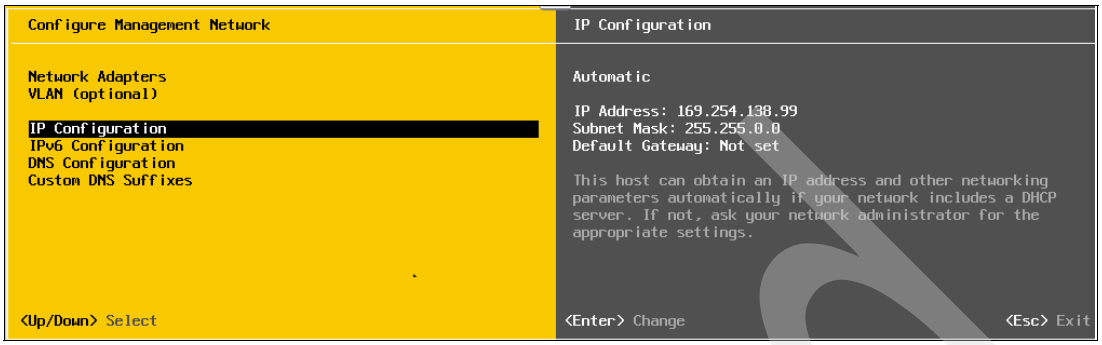

Figure 4-178 Specify IP Configuration

Enter the appropriate networking information and press Enter (Figure 4-179).

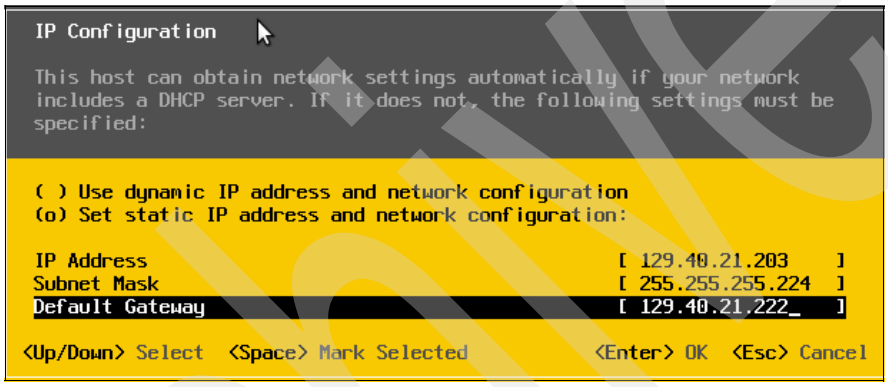

Figure 4-179 Specifying IP static IP addresses

To set up DNS, which is strong suggestion, scroll to and select **DNS Configuration**, and then press Enter (Figure 4-180).

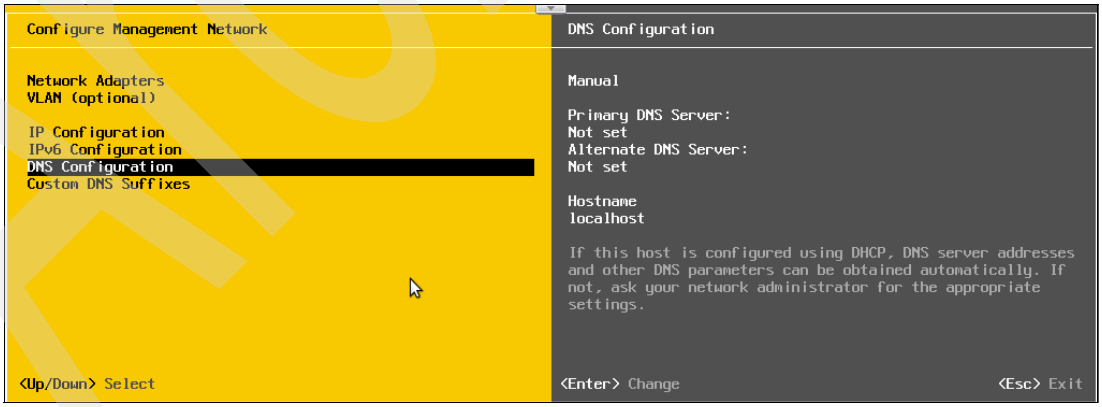

Figure 4-180 Choosing DNS Configuration

Enter the Primary and Secondary DNS servers (Figure 4-181). Also set the host name of this node because it is now on the management interface.

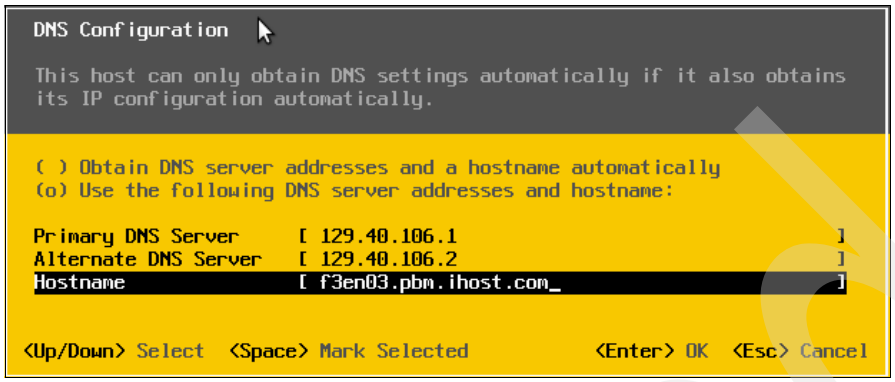

Figure 4-181 Specifying DNS server information

The DNS suffixes can also be set, if required (Figure 4-182).

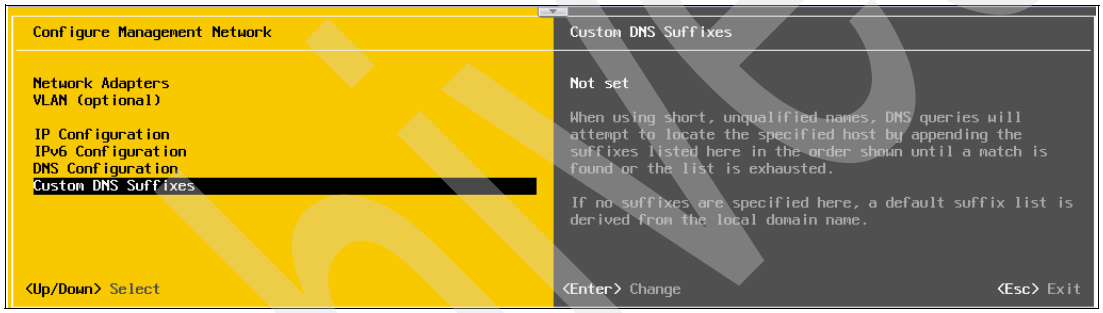

Figure 4-182 Custom DNS suffixes

Enter the information and press Enter (Figure 4-183).

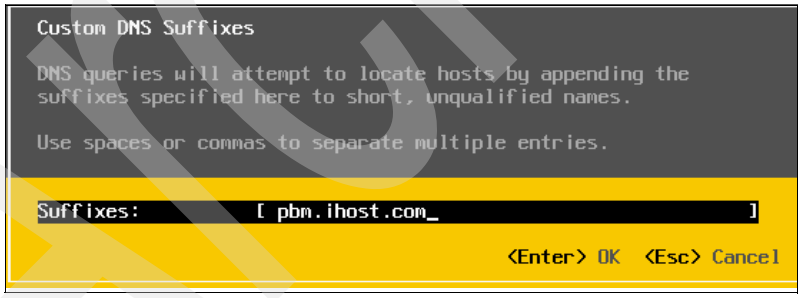

Figure 4-183 Custom DNS Suffixes for short unqualified names

Press ESC to exit the configuration. A confirmation is displayed (Figure 4-184).

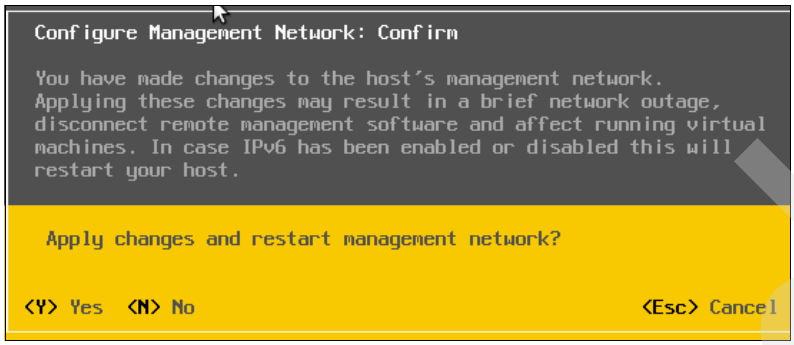

Figure 4-184 Confirm Configuration Changes

Press Y (Yes) to apply the changes and restart the networking stack (Figure 4-185). A good approach is to test the networking changes when the restart operation is complete.

Figure 4-185 Testing the configuration

The test checks the network connection to gateway, to DNS servers, and ensures the node's host name can be resolved. After all of these tests return OK message, press Enter (Figure 4-186).

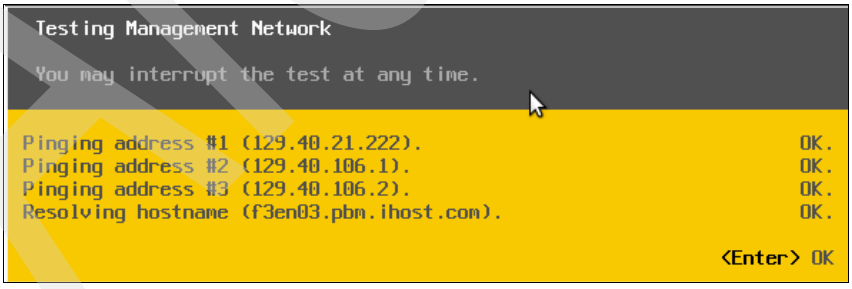

Figure 4-186 Pinging progress is displayed

The configuration process is complete. Press ESC to end the customization session and return to the main screen (Figure 4-187).

Figure 4-187 Main screen after configuration changes have been made

Repeat this process for all x86 nodes in the chassis. This allows for a greater number of cloud services. In the next few sections, the VMware management environment is built.

4.7.3 Install VMware vSphere client

After all of the nodes have the ESXi hypervisor installed, install a VMware client. This installation is covered in this section. The requirements for the client vary for each release of VMware and are not covered.

From a supported browser, open the following URL format, where `<address>` is the IP address of the management network configured in the last section:

```
https://<address>/
```

This URL loads the Getting Started page (Figure 4-188). On this page, click the **Download vSphere Client** link.

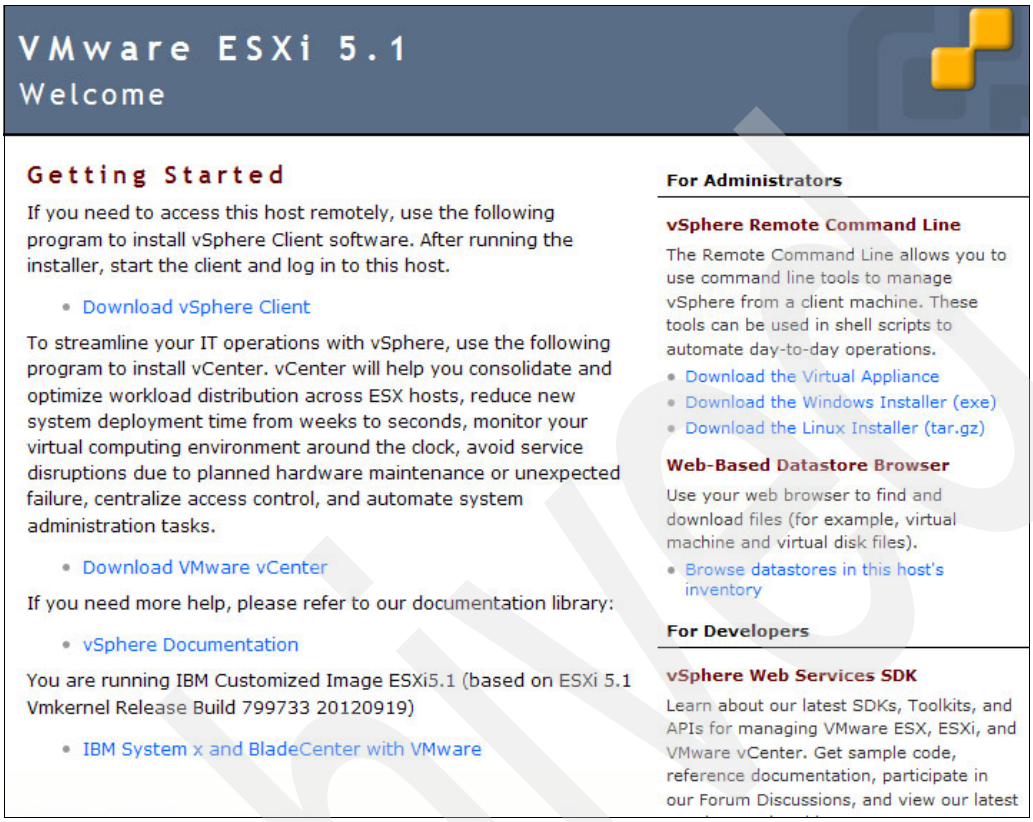

Figure 4-188 VMware vSphere: Welcome

The link points to the executable that installs the client. Click **Run** or **Save** to begin (Figure 4-189). The browser shows a warning that this software might be untrusted. In our case, we clicked **Run**.

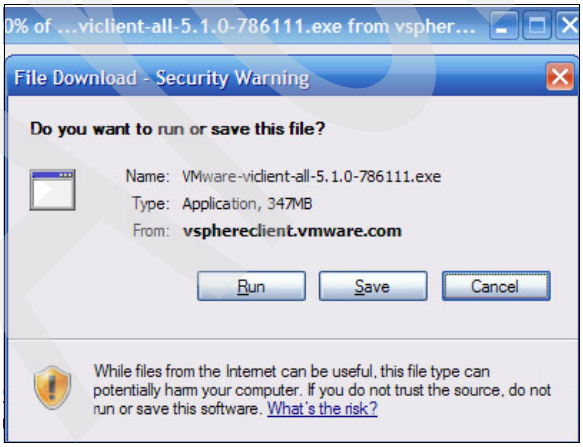

Figure 4-189 VMware vSphere: File Download - Security Warning

The browser shows a progress bar (Figure 4-190) as the code is copied from the node to the machine on which the client is to be installed.

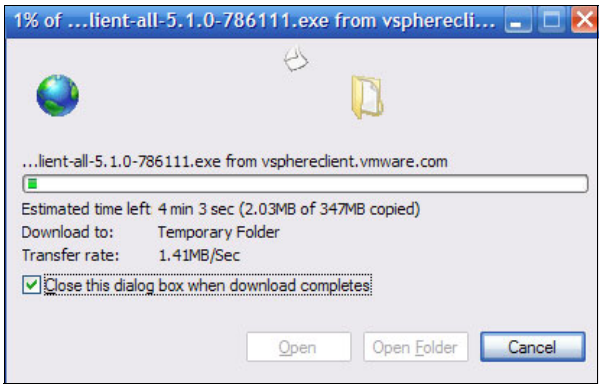

Figure 4-190 VMware vSphere: Downloading VMware vSphere Client

Click **Run** to start the installation (Figure 4-191).

Figure 4-191 VMware vSphere: Security Warning

In this example, the installation is done in English, which is the default. Select the appropriate language and click **OK** (Figure 4-192).

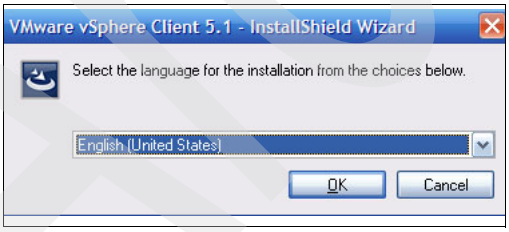

Figure 4-192 VMware vSphere: InstallShield Wizard

Another progress bar indicates that the installation program is proceeding (Figure 4-193).

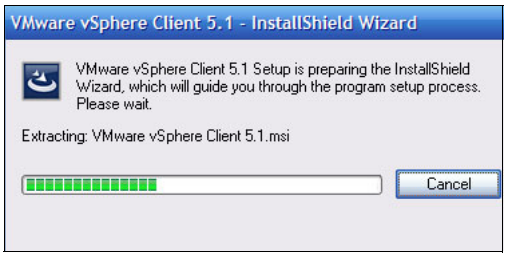

Figure 4-193 VMware vSphere: Progress Bar

The Windows Installer prepares to install (Figure 4-194).

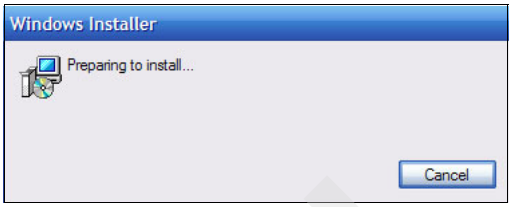

Figure 4-194 VMware vSphere: Preparing to install

Click **Next** to start the installation (Figure 4-195).

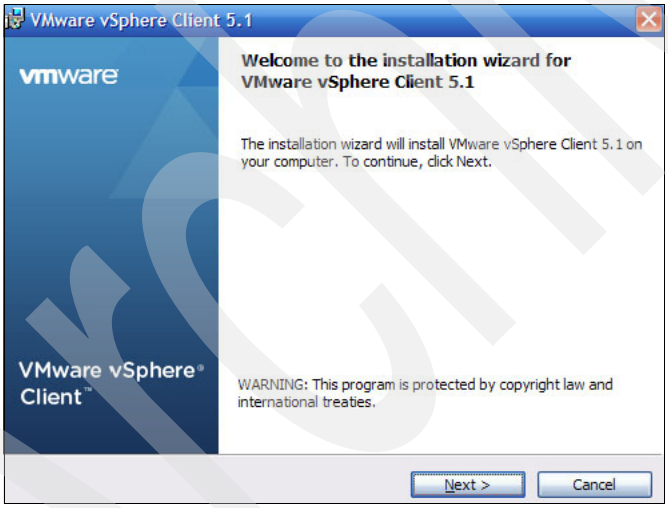

Figure 4-195 VMware vSphere: Starting the installation

Read the End-User Patent Agreement and click **Next**.

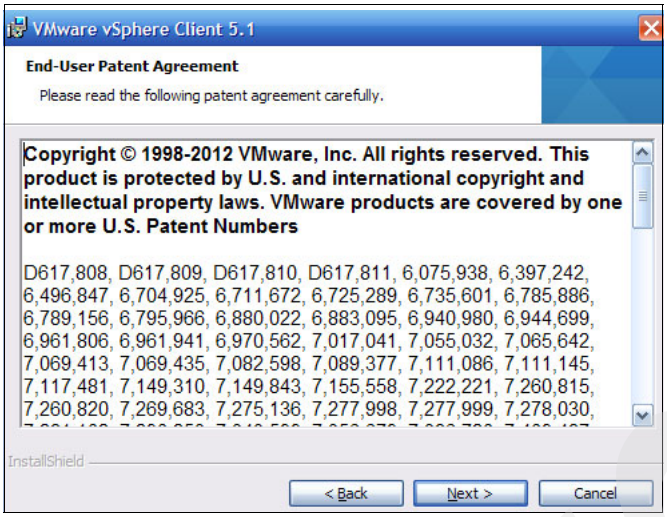

Figure 4-196 VMware vSphere: Patent Agreement

Read the License Agreement for the client, select **I agree to the terms in the license agreement**, and click **Next** (Figure 4-197).

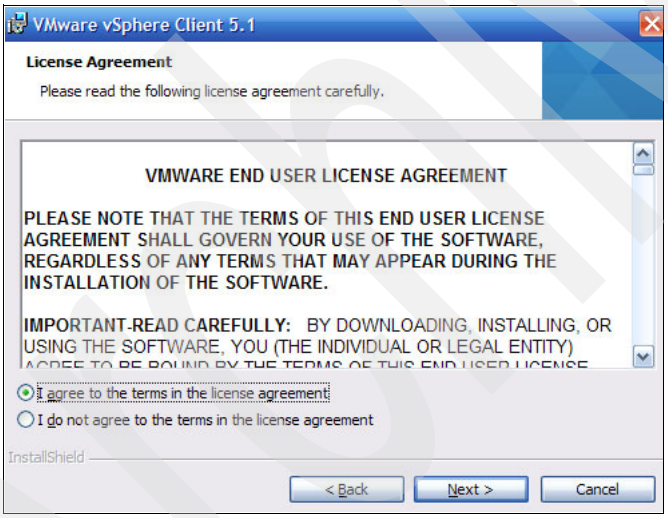

Figure 4-197 VMware vSphere: License Agreement

Select an installation (destination) folder (Figure 4-198) and then click **Next**.

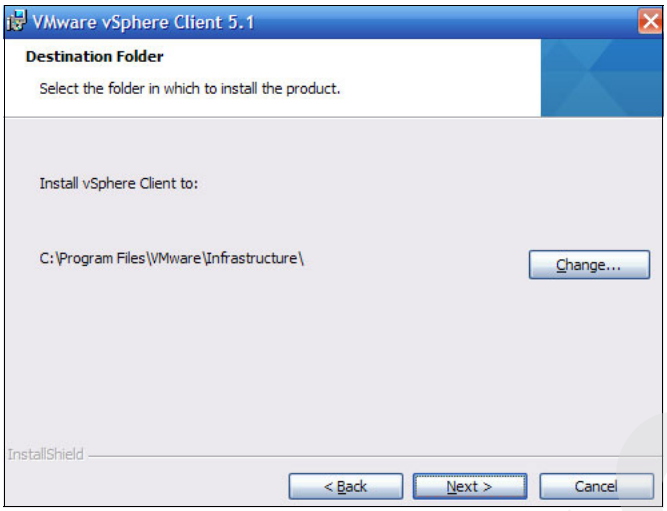

Figure 4-198 VMware vSphere: Specify Destination Folder

The vSphere client installation begins. When it finishes, launch the client from the Start menu. The client opens (Figure 4-199) and can now be used to configure the nodes on which ESXi installed. Enter in the IP address or host name and the credentials that were set up.

Figure 4-199 VMware vSphere: Specify IP address, name, password

The client often has a self-signed certificate. Figure 4-200 shows the certificate warning. Depending on the level of security needed, either ignore or add this certificate.

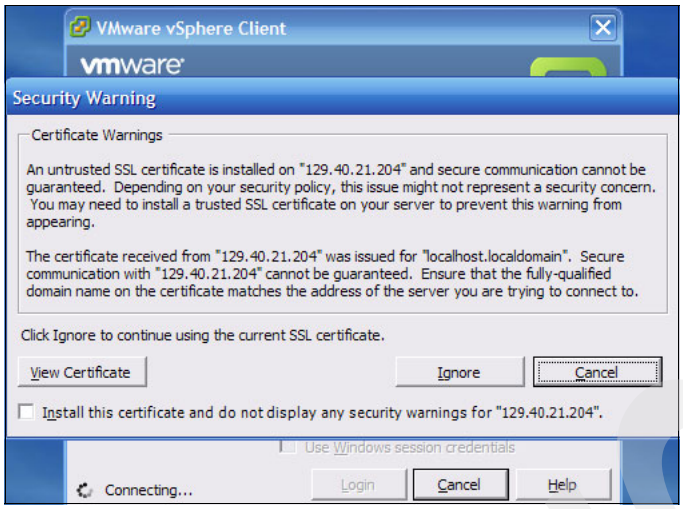

Figure 4-200 VMware vSphere: Security Warning

The client opens the vCenter environment. This installation has an invalid license (Figure 4-201). In this case, an evaluation license can be used, or a site license from VMware can be used.

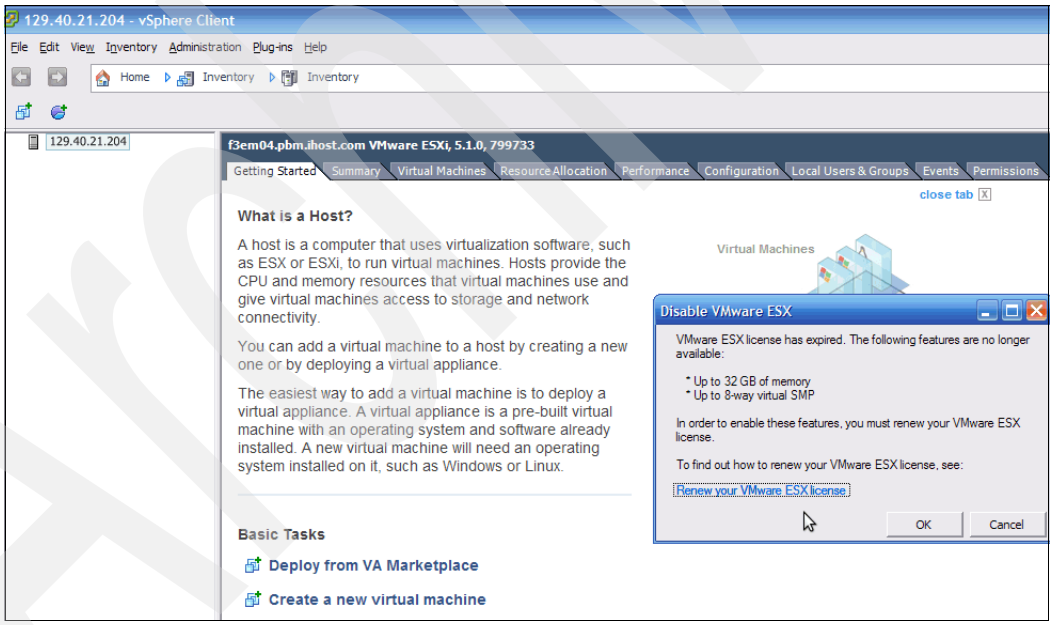

Figure 4-201 VMware ESX: Invalid License

4.7.4 Data store

Similar to the POWER nodes setup (in 4.6, “Power Compute node setup” on page 62), VMware can use the storage attached network to store virtual machines. However, VMware does not create a new volume for each virtual machine. Instead, a single larger pool is created to act as persistent storage. This is known as a data store (or datastore).

The creation of the volume and host mapping are not covered in this section but are similar to the steps in 4.6.2, “Create volume on V7000” on page 68.

This example was completed on the internal USB drive in the previous section. No persistent storage was selected, as shown in the Configuration tab of the main vCenter window (Figure 4-202). Click the link to **create a datastore**.

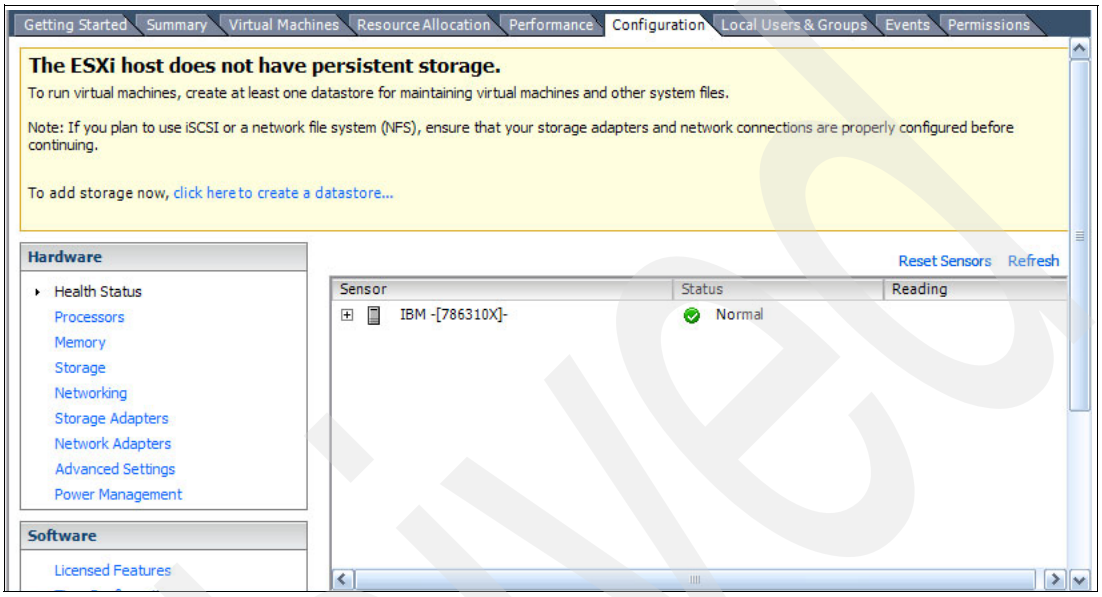

Figure 4-202 Add Storage: Click here to create a datastore

From here, a new storage device can be added. In this case, the first discovered LUN is used (Figure 4-203) and corresponds to the volume created on the V7000.

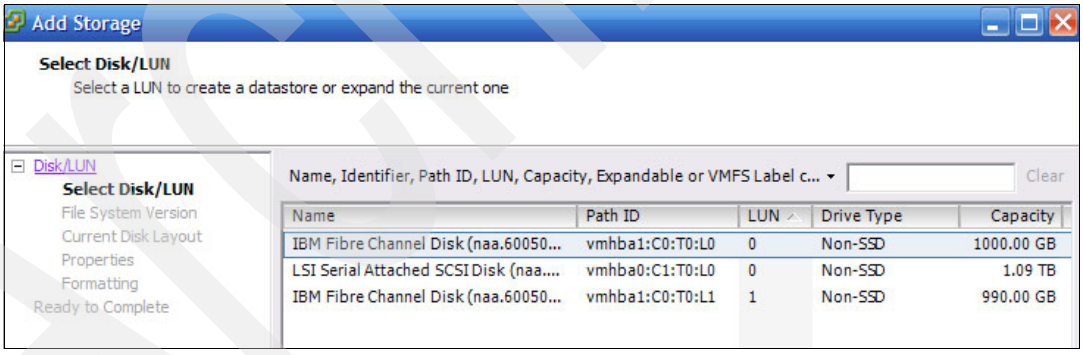

Figure 4-203 Add Storage: Select Disk/LUN

This example is using VMware version 5, so VMFS-5 is selected (Figure 4-204).

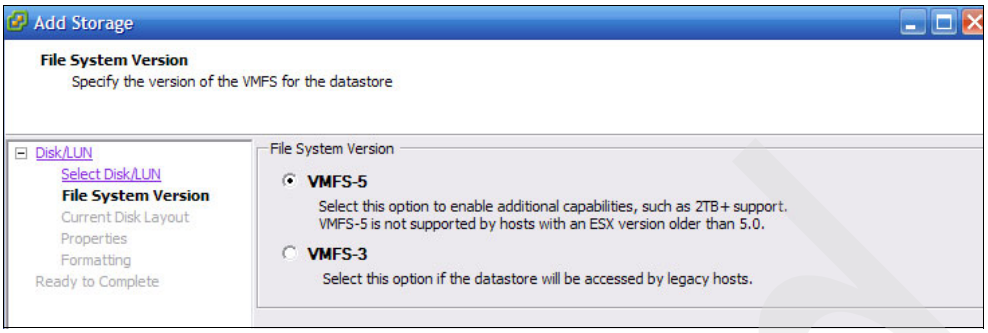

Figure 4-204 Add Storage: Select File System Version

Figure 4-205 shows the full ID of the LUN, indicating that it is blank.

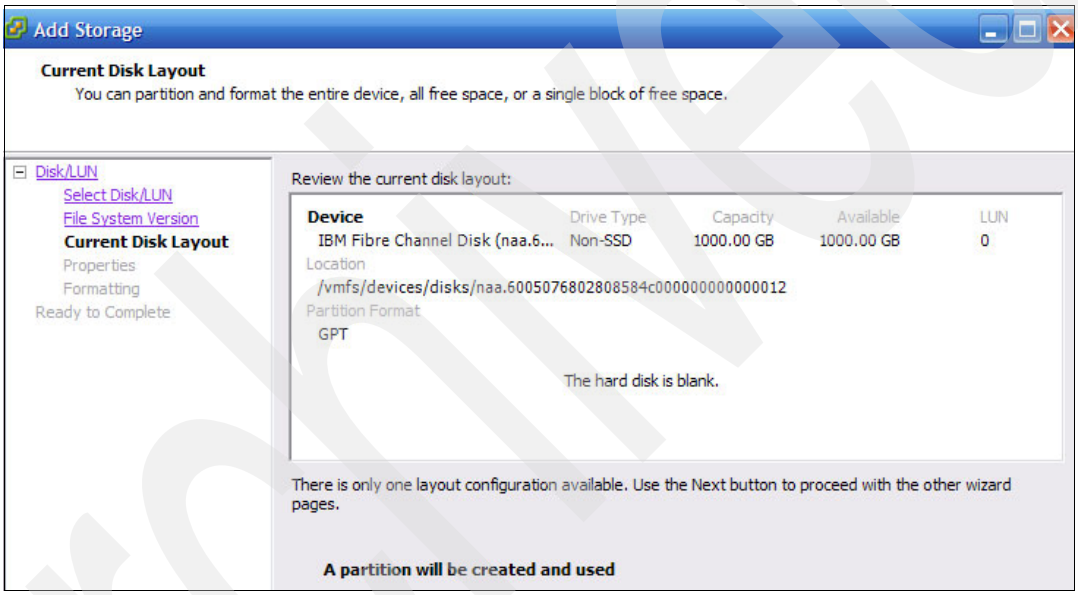

Figure 4-205 Add Storage: Review current disk layout

Enter the name of this data store. This data store will be used for SmartCloud Entry workloads and appliances, so the example uses the name `SCEDataStore` (Figure 4-206).

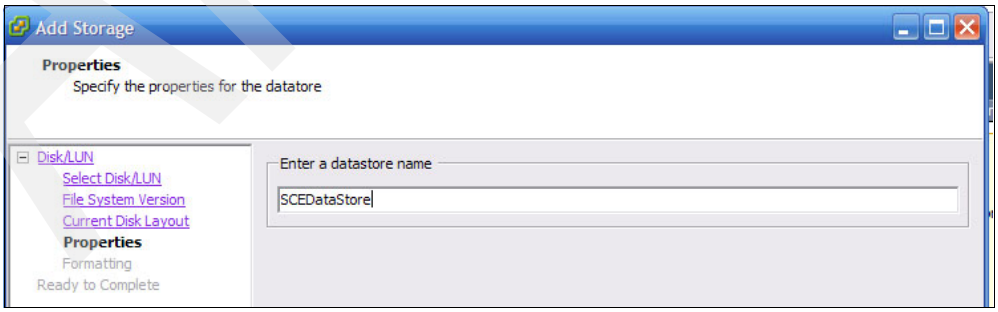

Figure 4-206 Add Storage: Enter datastore name

The maximum amount of space will be used (Figure 4-207).

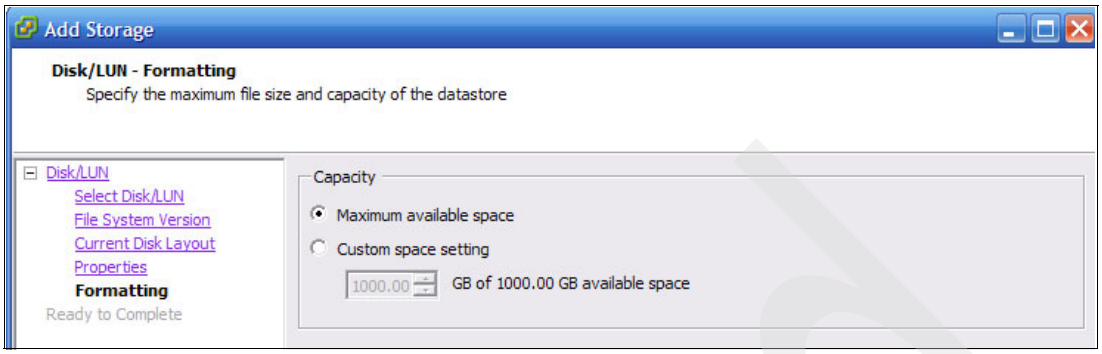

Figure 4-207 Add Storage: Formatting LUN

After the format is complete, the data store is ready to use. The first virtual machine to use this data store is the vCenter Server Appliance, which is described in 4.7.5, “Deploy and configure VMware vCenter Server Appliance” on page 123.

4.7.5 Deploy and configure VMware vCenter Server Appliance

Deploy vCenter and set up the vCenter server (client and web).

Deploy vCenter open virtual appliance (OVA)

Although the vCenter client can control the ESXi server effectively, IBM SmartCloud Entry requires that a vCenter server be installed. This server acts as a gateway to all of the compute nodes. SmartCloud Entry can treat all of the nodes in the exact same way by directing its instructions to the vCenter server. The server also provides high availably, clustering, and ease of management.

From a client session, select **File** → **Deploy OVF Template** (Figure 4-208).

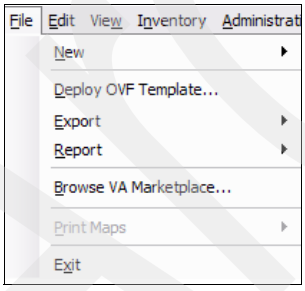

Figure 4-208 Choosing Deploy OVF Template

This wizard first requires a source. Although the selection was for deploying an open virtualization format (OVF), this wizard can deploy an OVA file also. The vCenter must be acquired from VMware and is not covered in the document.

After the OVA is downloaded, specify its location (Figure 4-209).

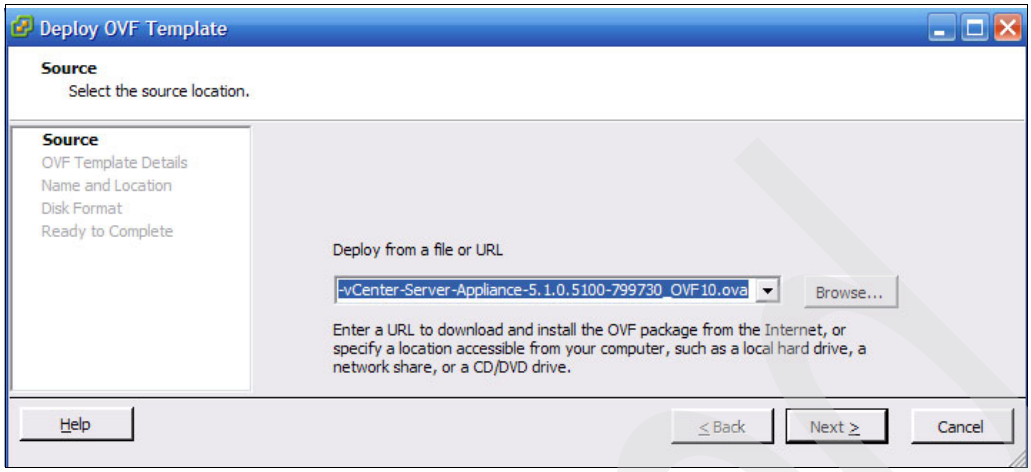

Figure 4-209 Specify deployment file or URL

Figure 4-210 shows the details of the appliance represented by the OVA file.

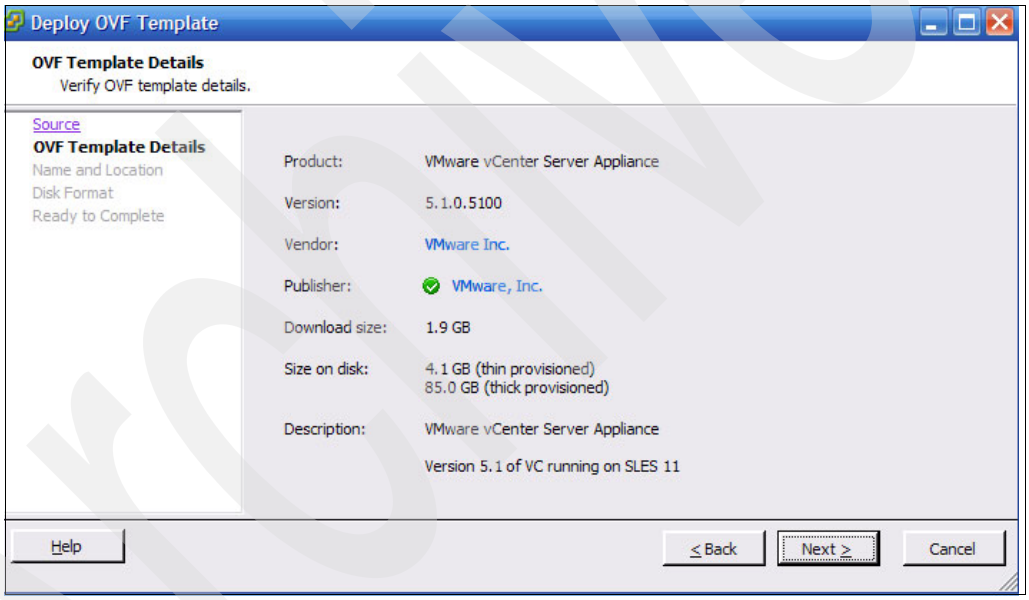

Figure 4-210 OVF Template Details

Enter the name for the virtual machine that this appliance will create and click **Next** (Figure 4-211). The name `VMware vCenter Server Appliance` is used in this example.

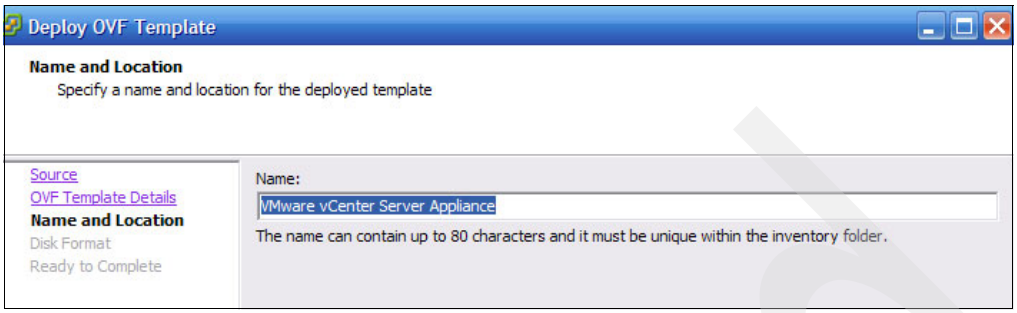

Figure 4-211 OVF Template Name and Location

Figure 4-212 shows the settings. Verify that the information is correct. Select the **Power on after deployment** check box and click **Finish**.

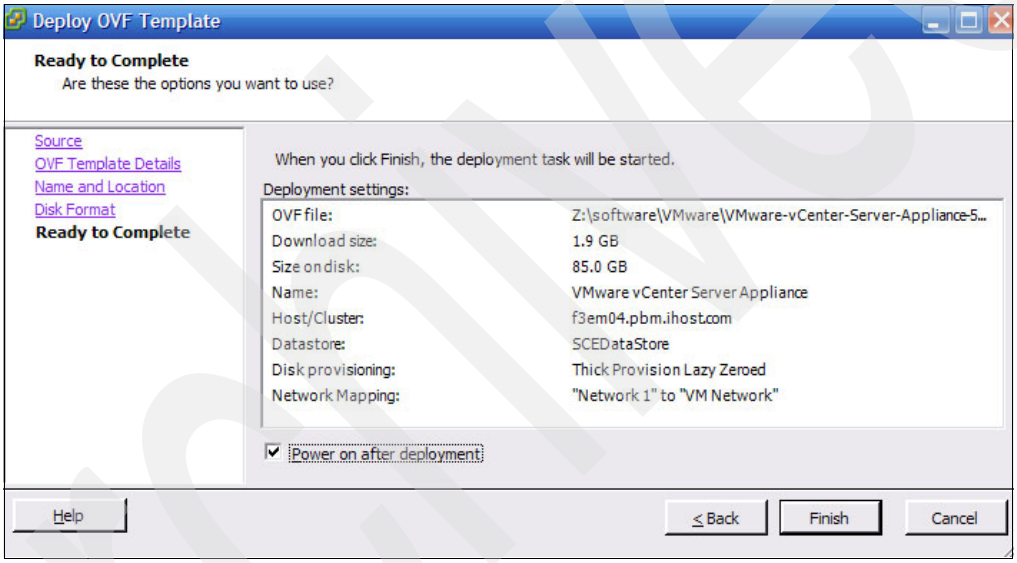

Figure 4-212 Deployment settings

The deployment is in progress (Figure 4-213).

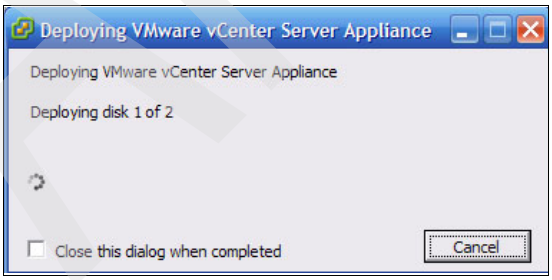

Figure 4-213 Deploying VMware vCenter Server Appliance

This deployment copies the contents of the OVA to the compute node and configures the settings.

The progress window indicates that the appliance is deployed successfully (Figure 4-214).

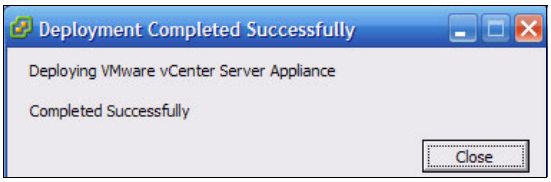

Figure 4-214 VMware vCenter Server Appliance Deployed Successfully

Set up vCenter Server (client-based)

Now that the appliance is deployed as a virtual machine, it must be configured. The first step is to select it from the list of virtual machines on the left (Figure 4-215) and then select the Console tab.

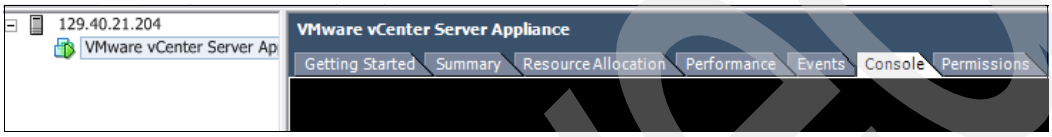

Figure 4-215 Configuring the appliance

In the example environment, DHCP was not enabled. Therefore, as Figure 4-216 shows, a warning message indicates that no network is detected. Select the Login option at the bottom of the window and press Enter.

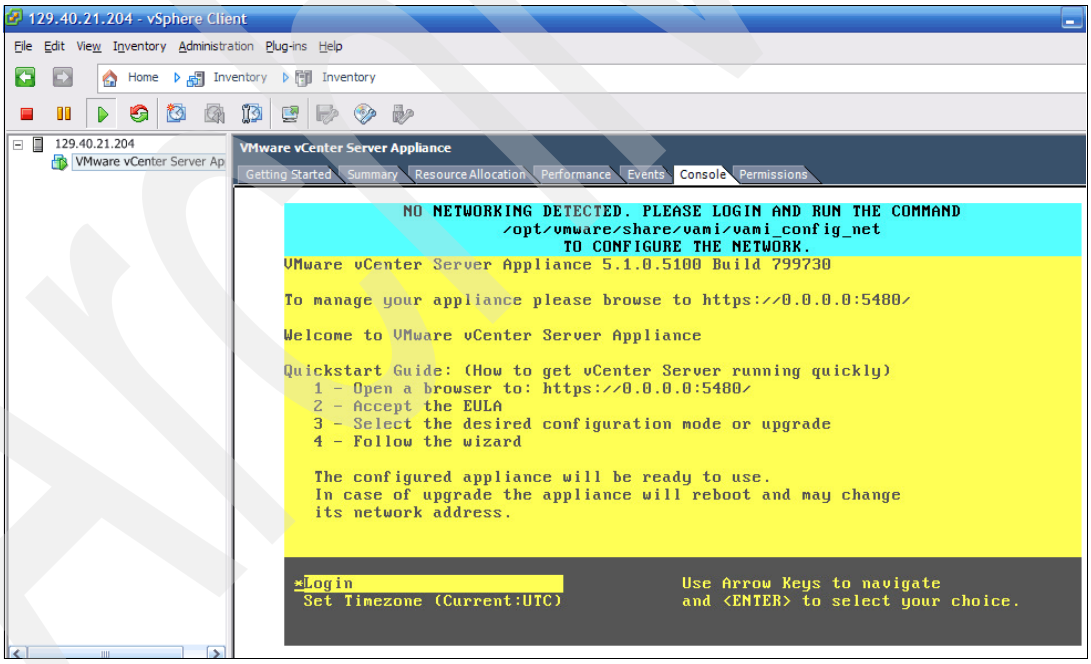

Figure 4-216 Warning that no network was detected

The login prompt opens (Figure 4-217).

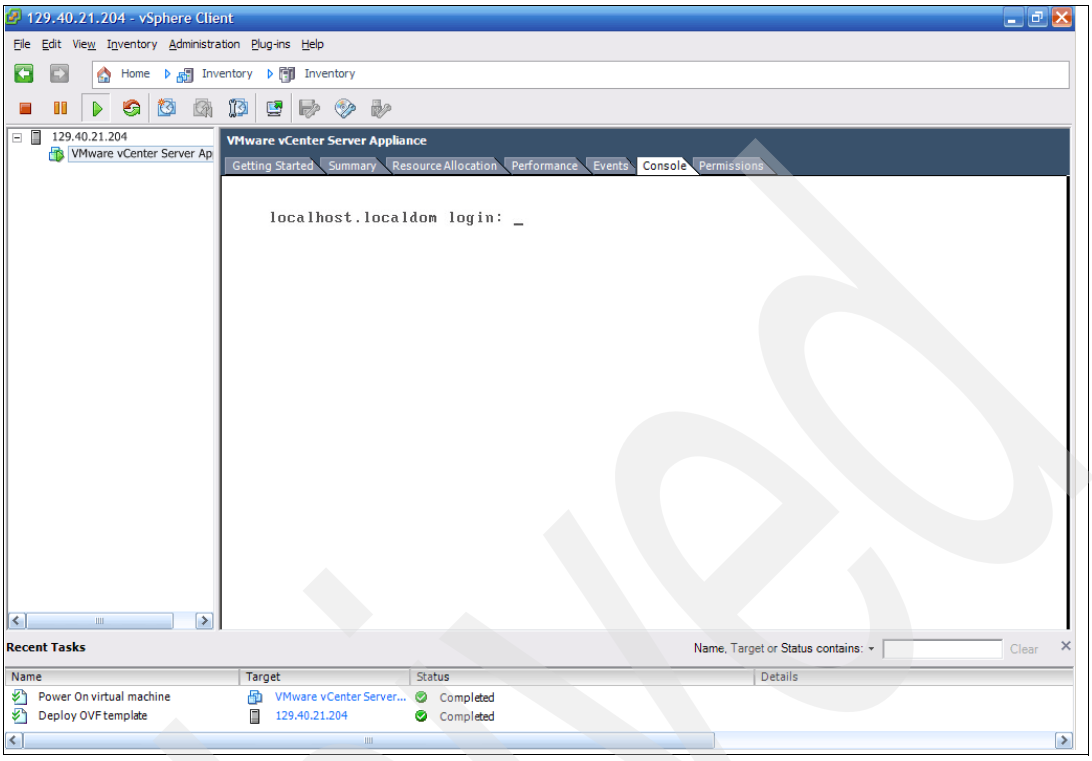

Figure 4-217 Login prompt

Log in with the default user name (root) and password (vmware) as Figure 4-218 shows.

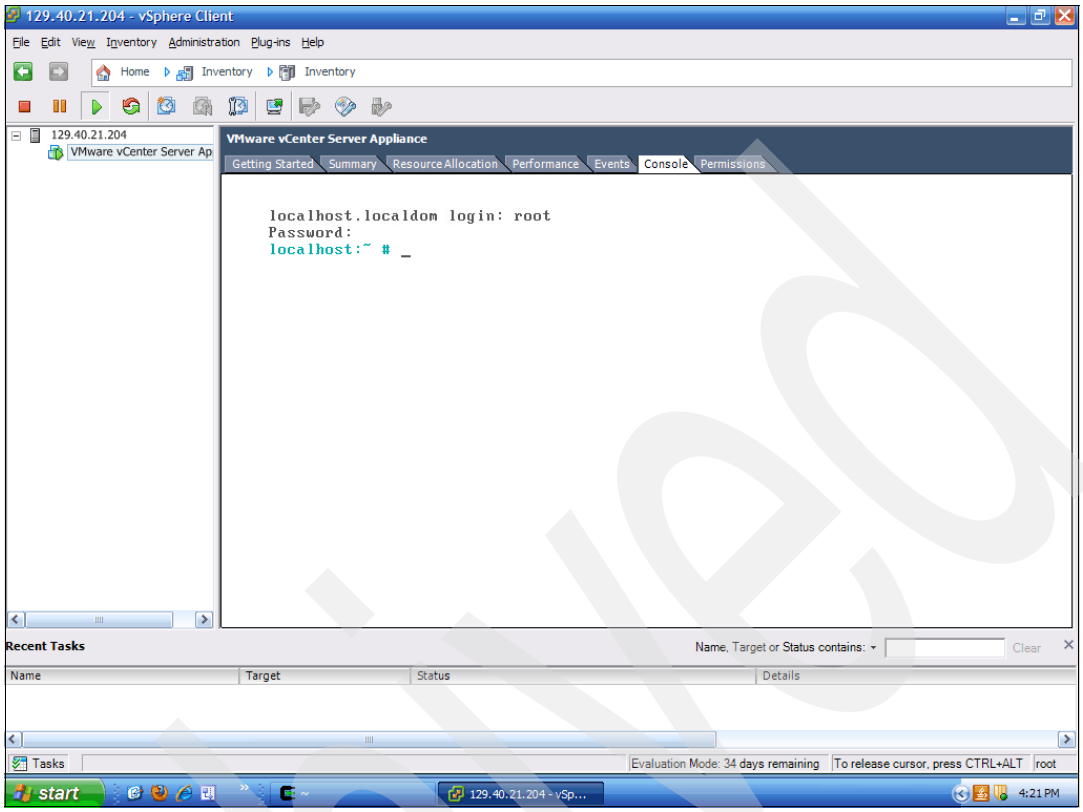

Figure 4-218 Logging in with defaults

YaST, a set up tool, can be use to configure the virtual server. Type `yast` (Figure 4-219) and press Enter.

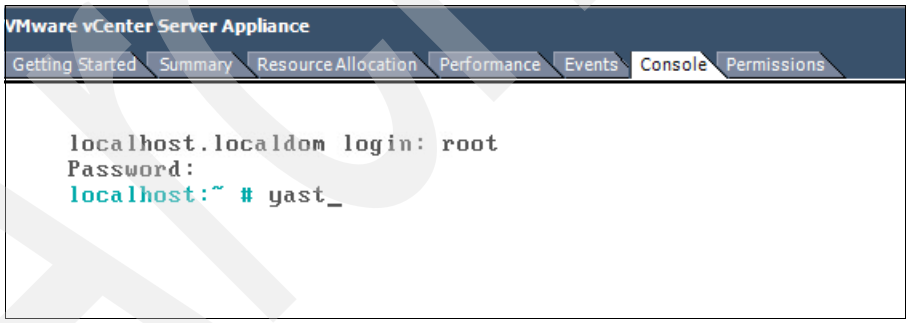

Figure 4-219 Invoking YaST

Page down to select **Network Devices** (Figure 4-220).

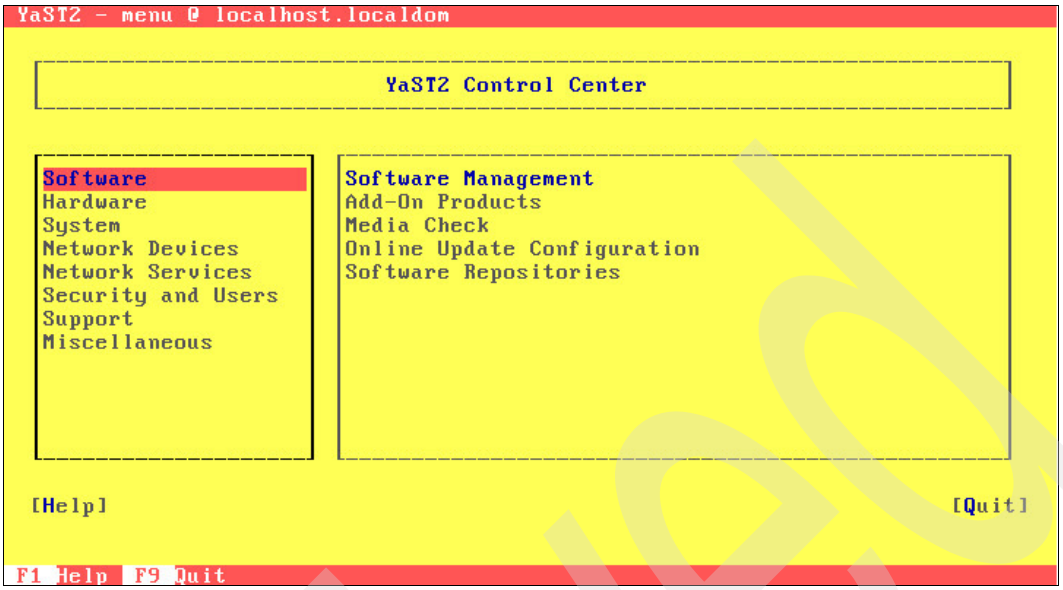

Figure 4-220 Selecting Network Devices

Select the **Network Settings** option and press Enter (Figure 4-221).

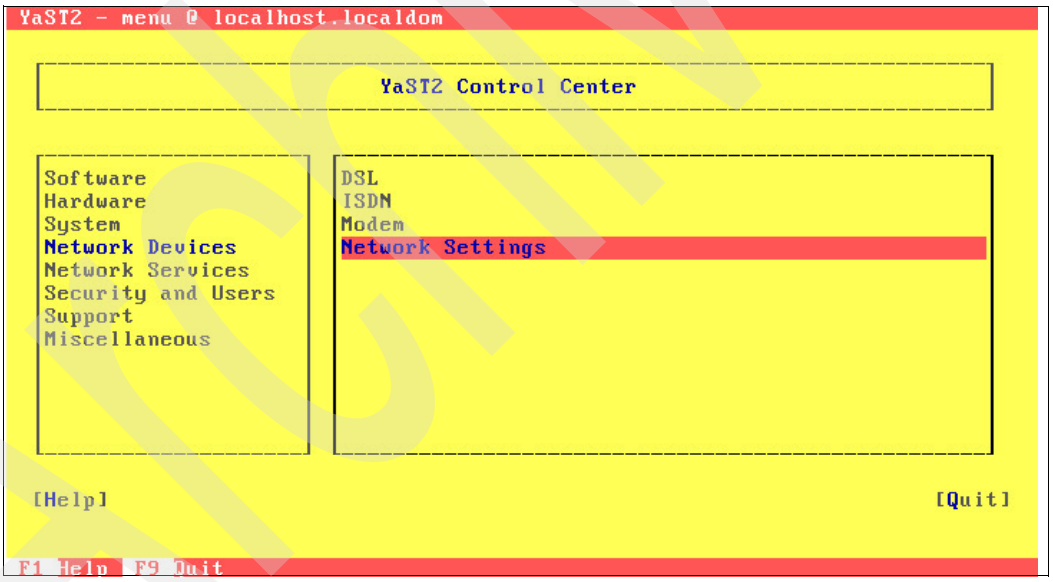

Figure 4-221 Selecting Network Settings

Only a single network interface was allocated for this machine. It is currently configured as DHCP. Select this interface and then select **Edit** (Figure 4-222).

Figure 4-222 Selecting the DHCP interface

Change the networking configuration to match the current environment. In this example (Figure 4-223), static IP addresses are used.

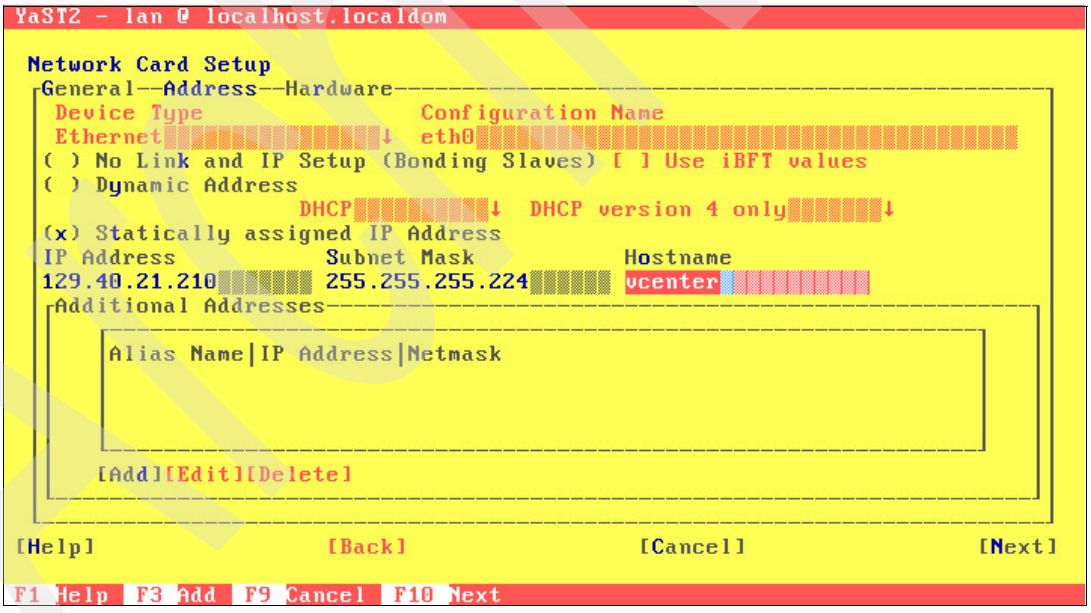

Figure 4-223 Selecting static IP addresses

In the Hostname and Domain Name, and DNS configuration sections (Figure 4-224), enter the appropriate host name, domain name, and other network name servers.

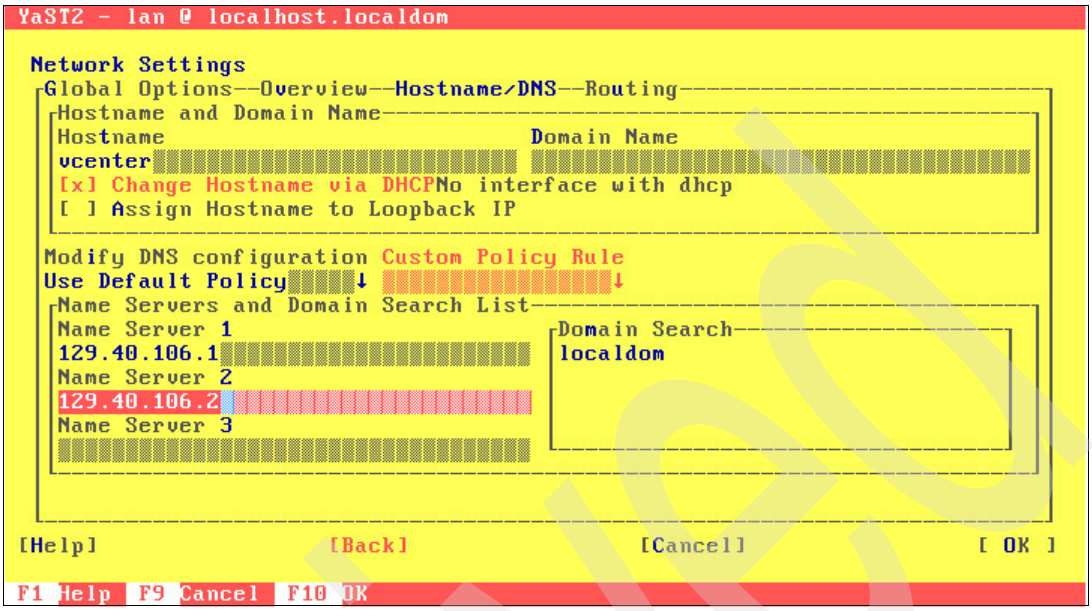

Figure 4-224 Entering host name and DNS information

Select the **Routing** option and set the Default Gateway (Figure 4-225). Press F10 for OK.

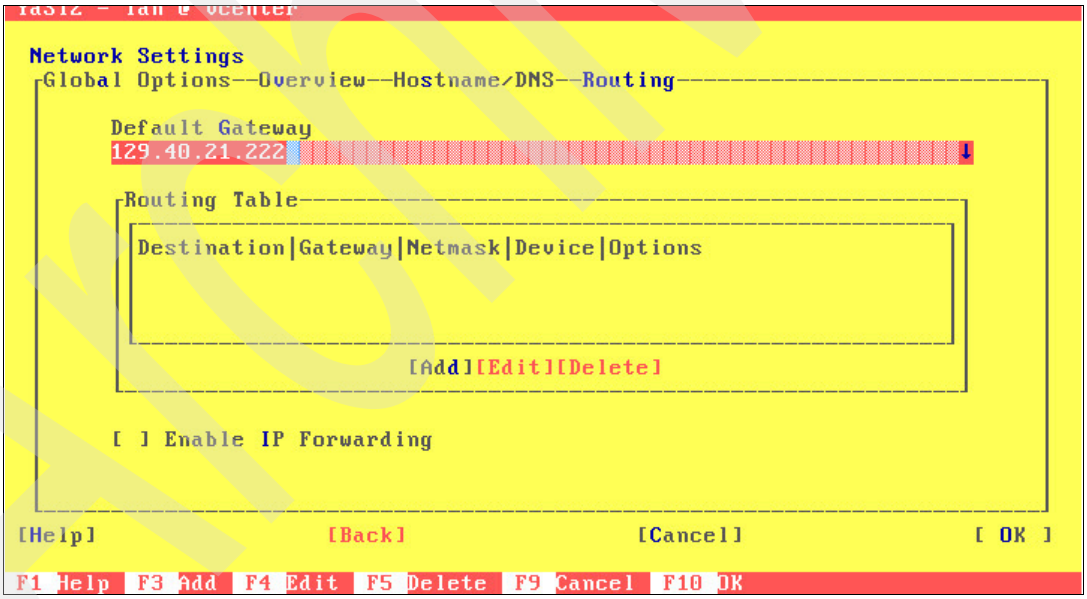

Figure 4-225 Defining routing and gateway information

YaST reconfigures the network. When this is complete, the remainder of the configuration can be done from a browser.

Open a new browser window and enter the following URL, where <address> is the networking information specified previously:

```
https://<address>:5480
```

The login window opens (Figure 4-226).

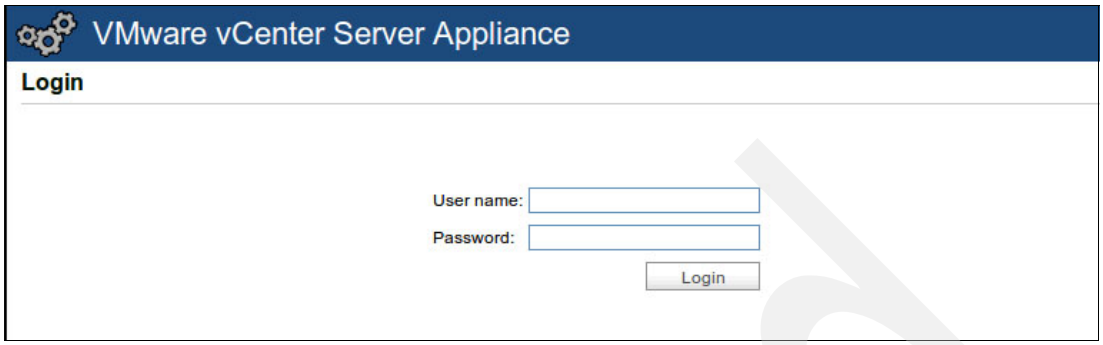

Figure 4-226 VMware vCenter Server Appliance: Login Panel

Set up vCenter Server (web-based)

After logging in with the same user name and password as before, read the user license agreement (Figure 4-227), select the **Accept license agreement** check box, and click **Next**.

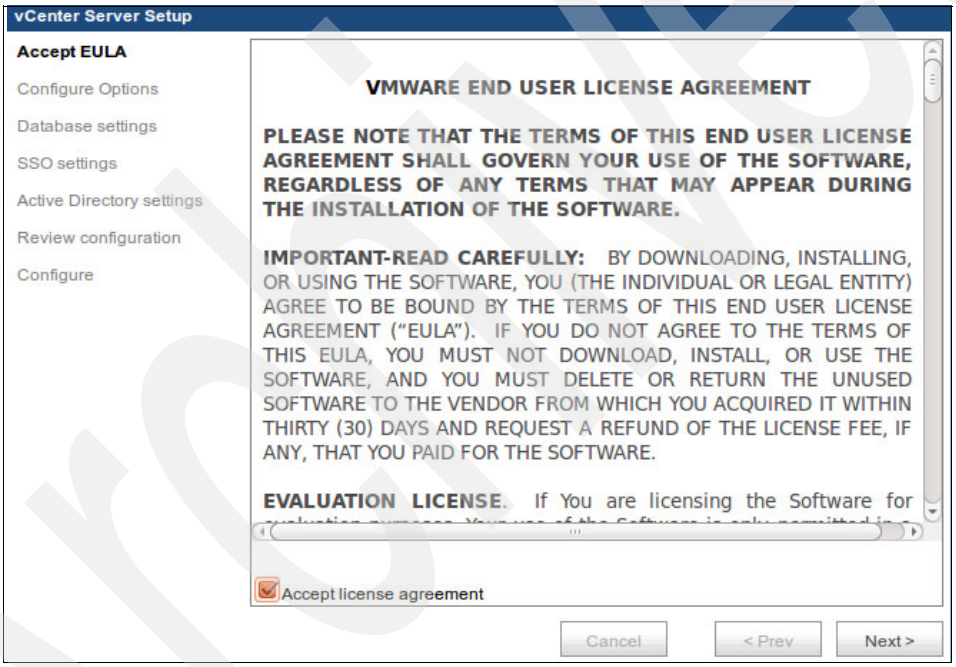

Figure 4-227 VMware vCenter Server Appliance: Accepting License Agreement

Review the options (Figure 4-228). The host name is already configured, so select the **Configure with default settings** option, and then click **Next**.

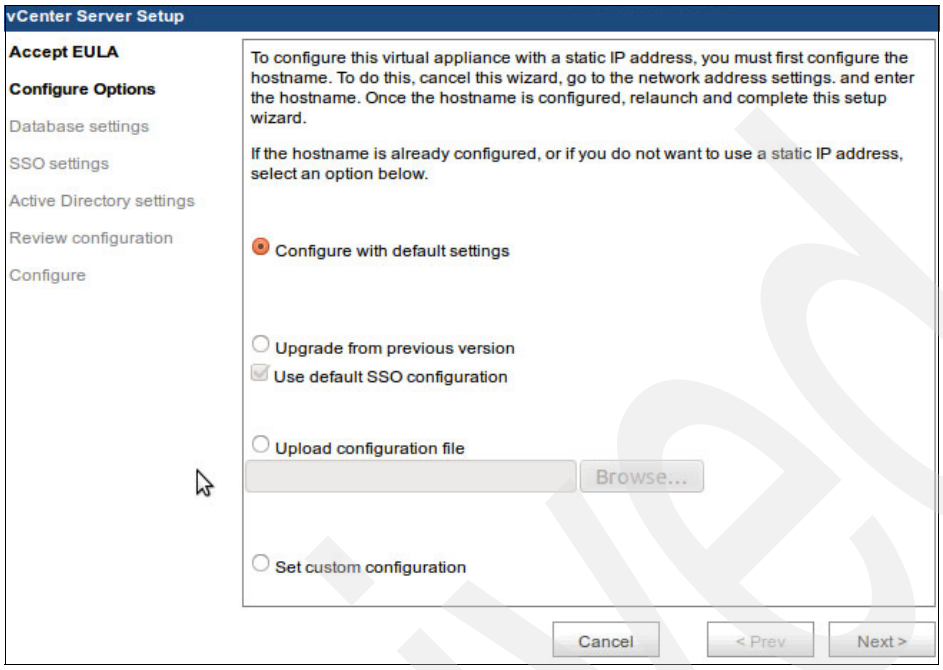

Figure 4-228 VMware vCenter Server Appliance: Configure Options

Review the information (Figure 4-229).

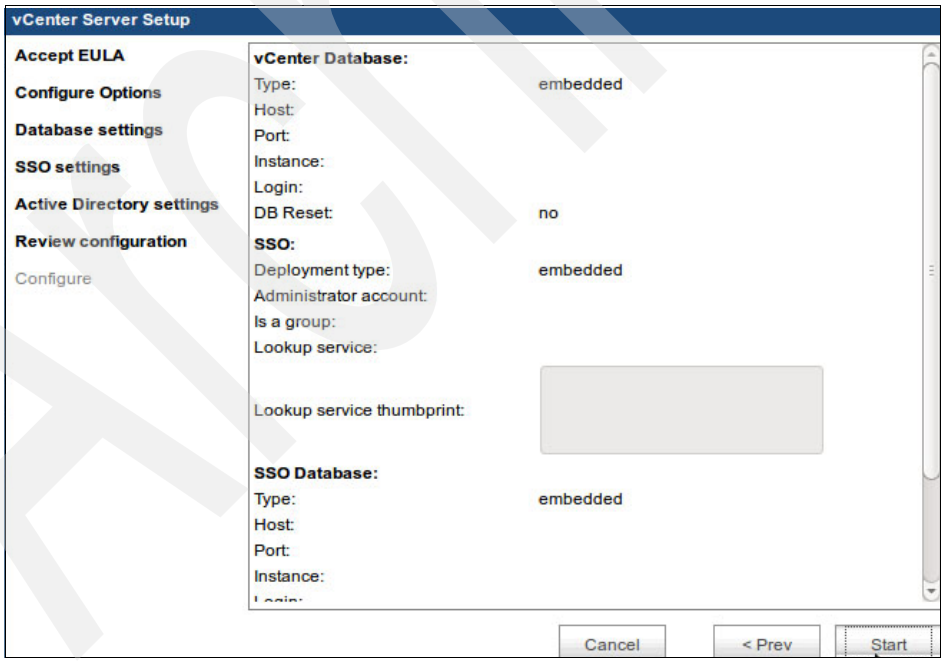

Figure 4-229 VMware vCenter Server Appliance: Ready to review configuration

The internal vCenter server database configures itself (Figure 4-230).

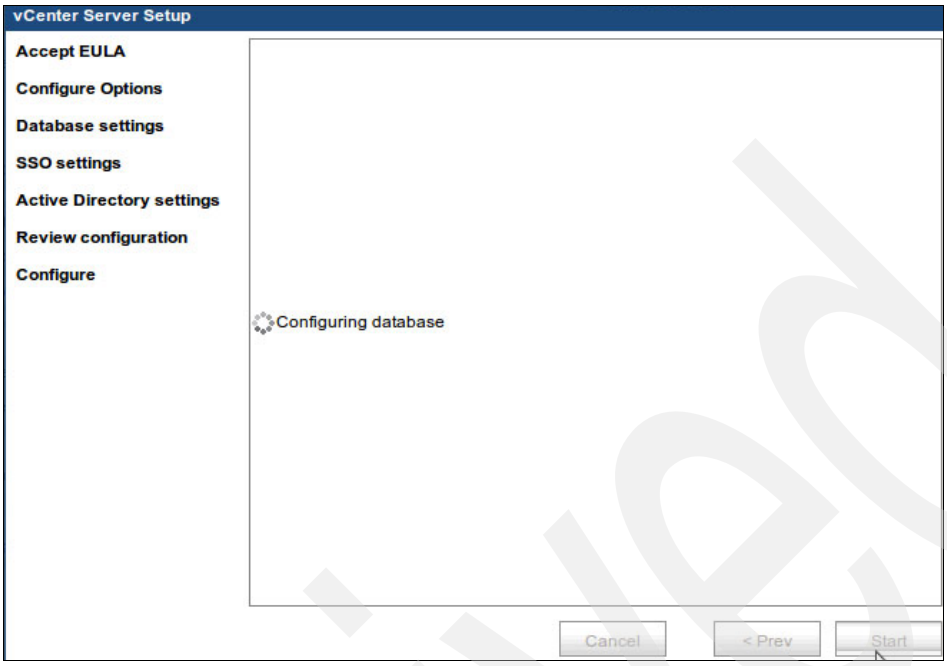

Figure 4-230 VMware vCenter Server Appliance: Database configures itself

After vCenter Server finalizes its internal setup, including the database and setting up of single sign-on (SSO), it starts the actual server. Click **Close** to finish (Figure 4-231).

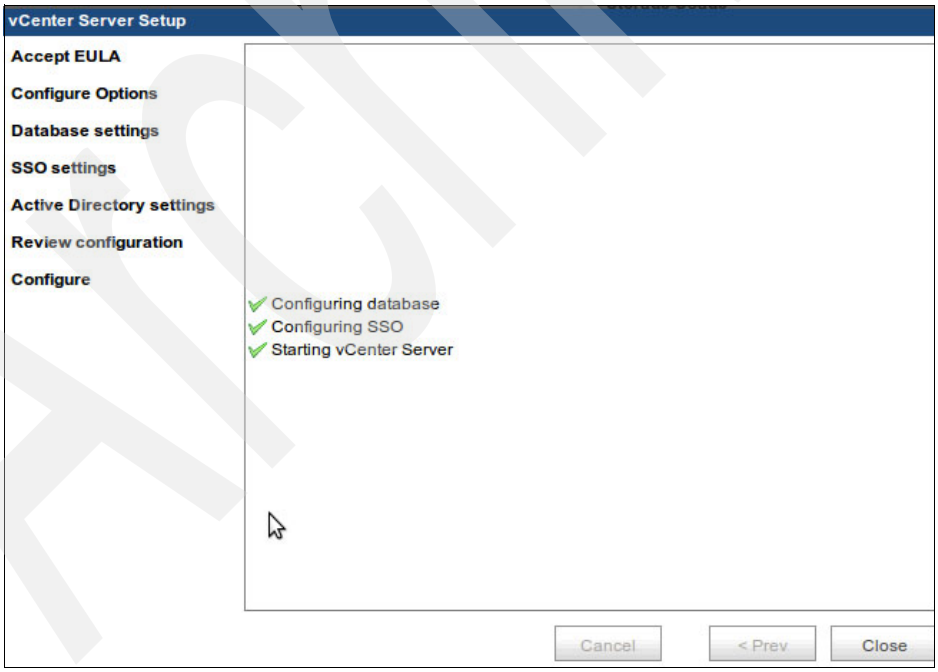

Figure 4-231 VMware vCenter Server Appliance: Database, SSO complete, and server starting

Figure 4-232 shows the main summary window of the appliance. Many of the internal services can be controlled from this panel.

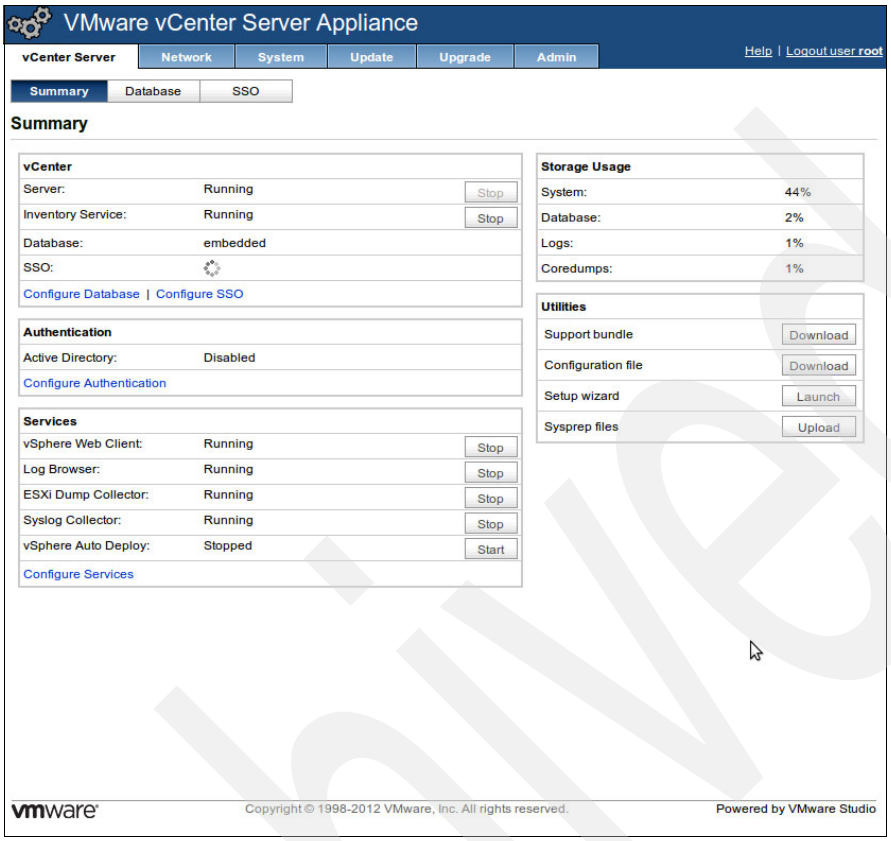

Figure 4-232 VMware vCenter Server Appliance: Summary

4.7.6 Adding hosts to vCenter Server

Now that the vCenter Server is configured, the client that was installed (in 4.7.2, “Install VMware ESXi” on page 106) can point to the server’s address. After logging in, a similar window opens (Figure 4-233). By default, vCenter server uses an evaluation license.

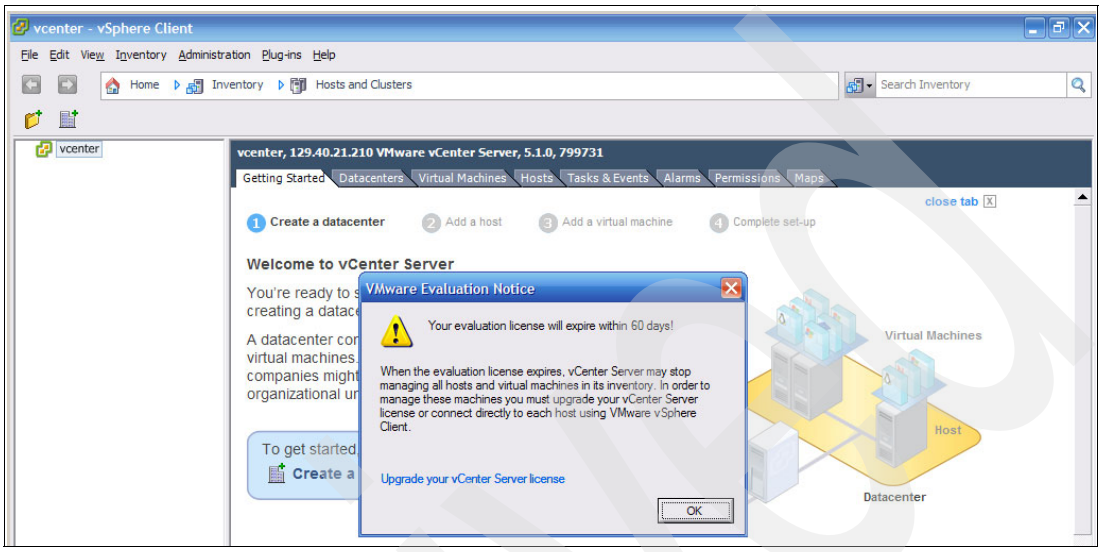

Figure 4-233 VMware Evaluation Notice

To manage all of the nodes as one unit, a new data center must be created. This is a vCenter server logical construct (Figure 4-234). To add a host, click **File** then **Add host**.

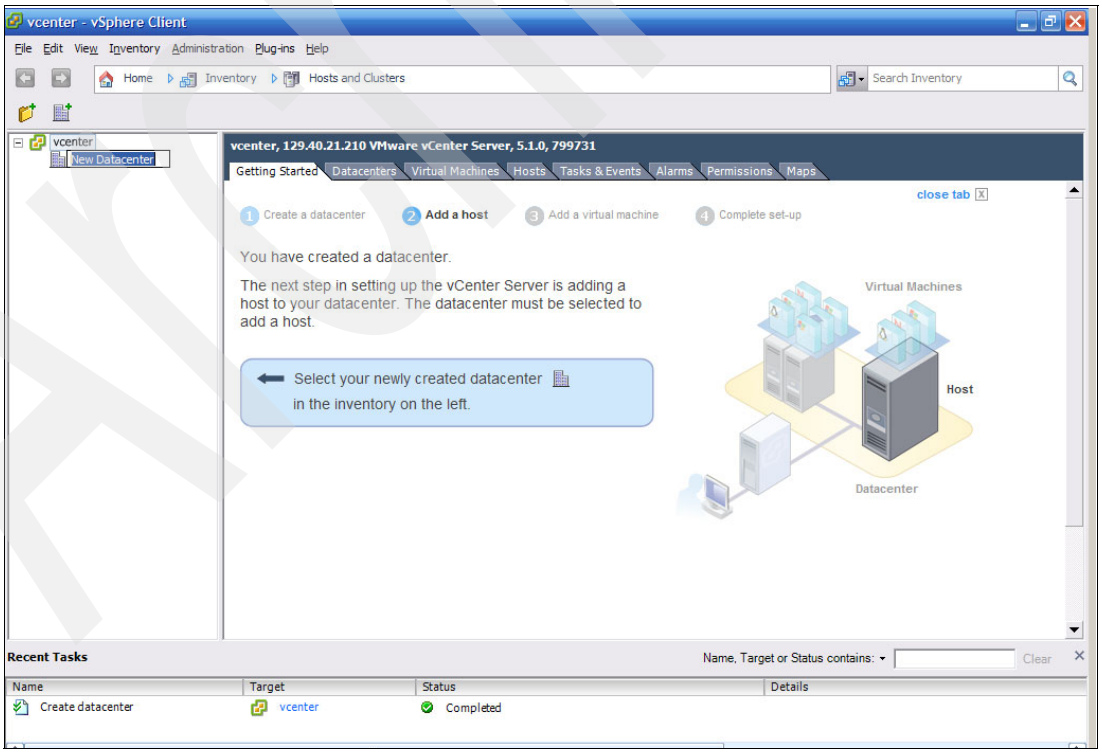

Figure 4-234 Creating a new data center

This initiates the Add Host Wizard (Figure 4-235).

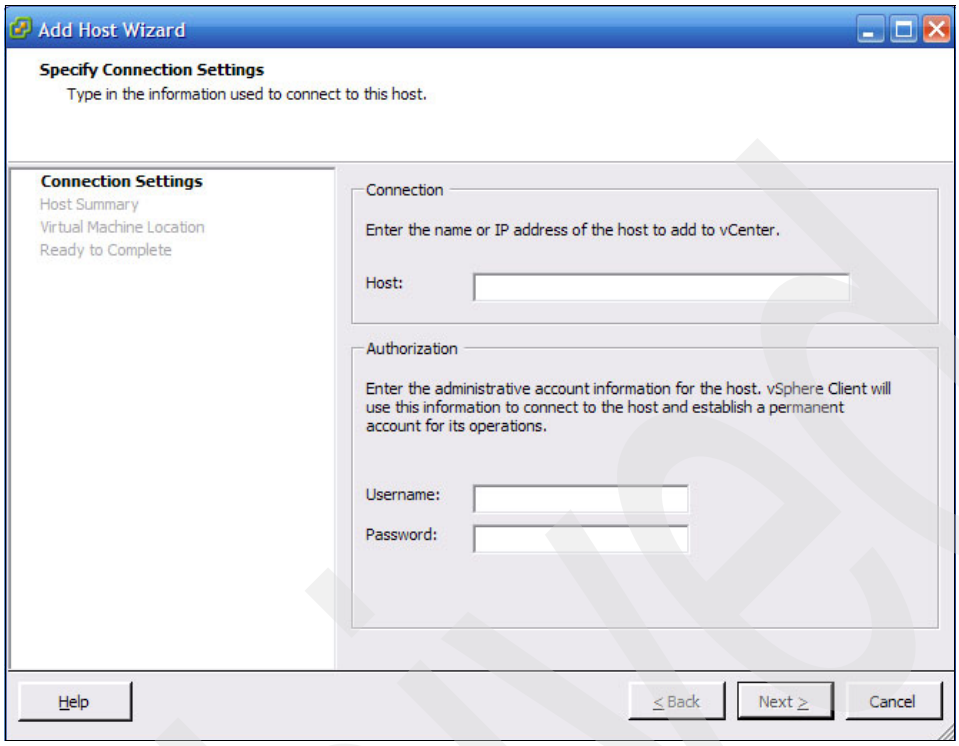

Figure 4-235 Add Host Wizard: Specify Connection Settings

The following steps must be done for every x86 compute node that will be used for cloud services. Specify connection settings and authorization (Figure 4-236).

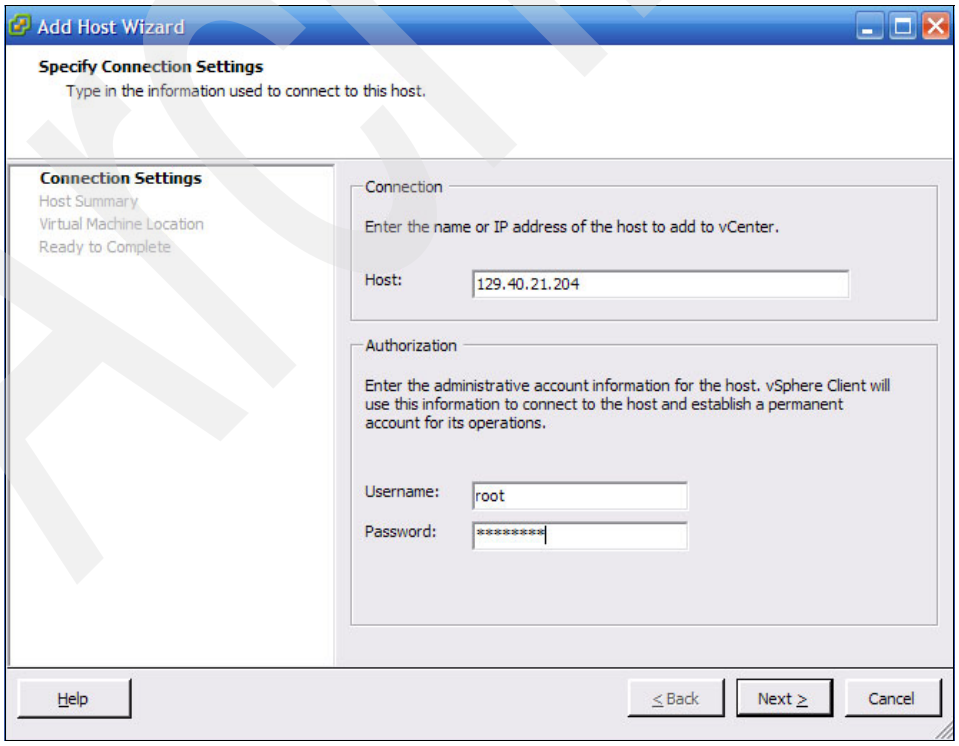

Figure 4-236 Add Host Wizard: Specify Connection Settings

This host's SSL certificate, as with many products, is self-signed. In general it can be trusted. Click **Yes** (Figure 4-237).

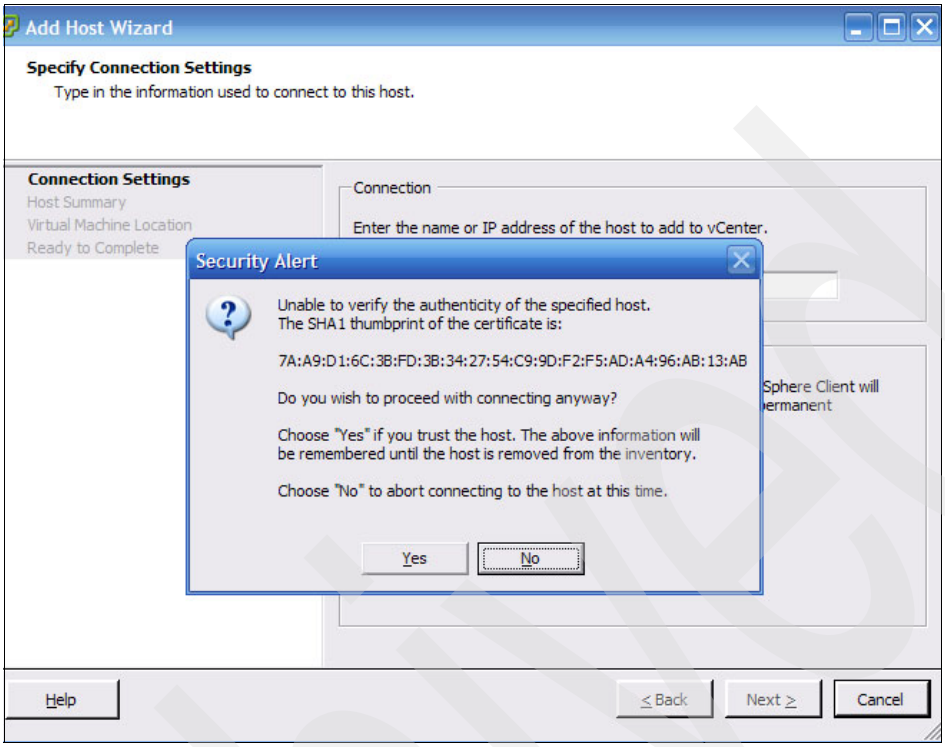

Figure 4-237 Add Host Wizard: Security Alert

The host summary window opens (Figure 4-238). Verify this is the correct host. Click **Next**.

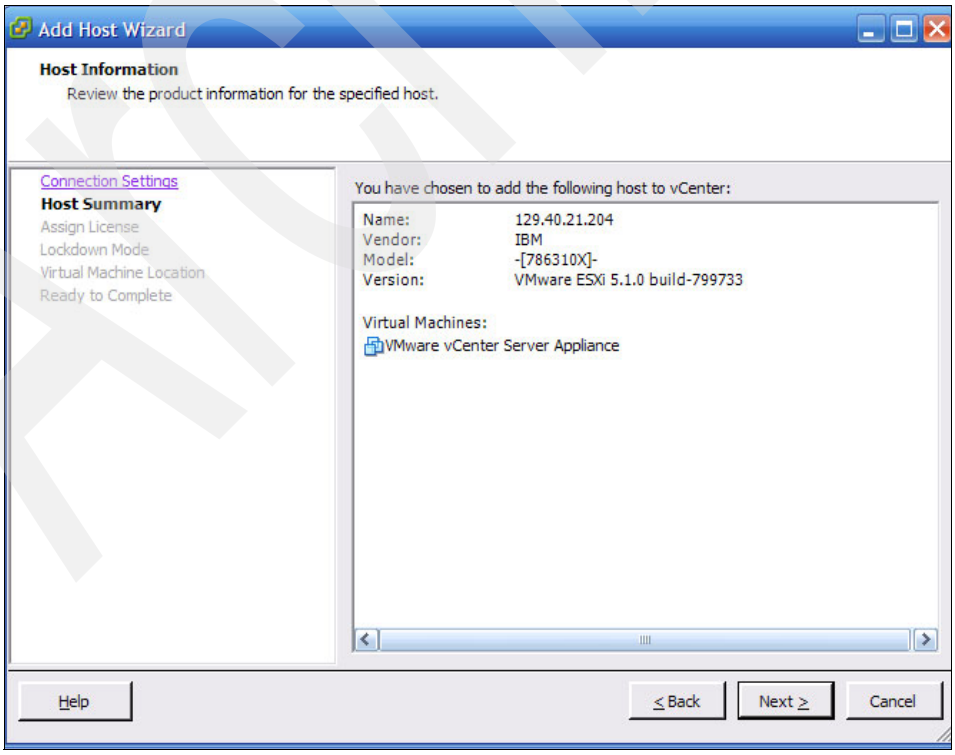

Figure 4-238 Add Host Wizard: Summary

For this example, the evaluation license is used (Figure 4-239).

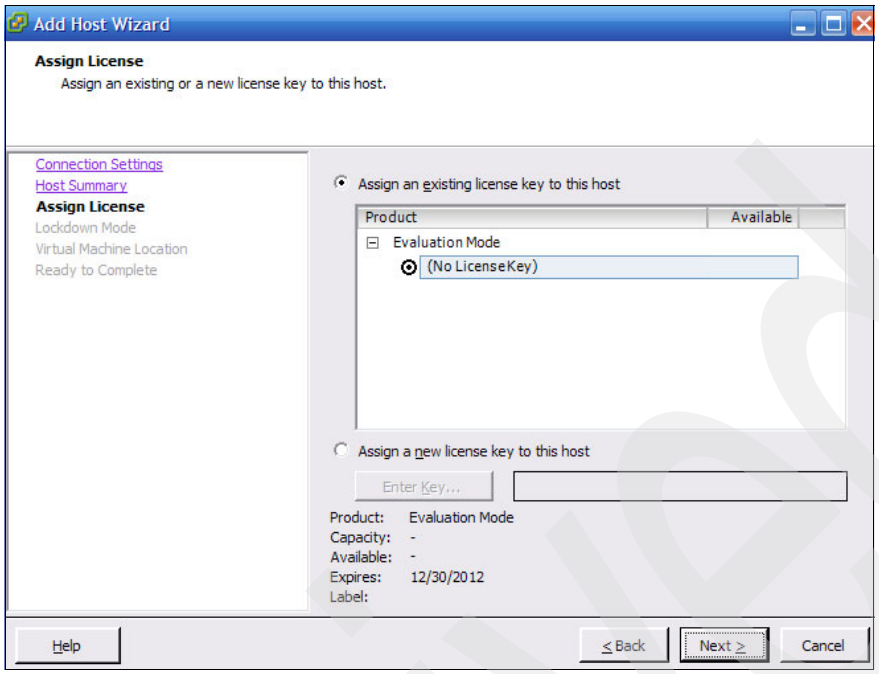

Figure 4-239 Add Host Wizard: Assign a license

Lockdown mode (Figure 4-240) is not required for the security in this environment.

Note: Check with authorized security personnel to understand ramifications of this panel.

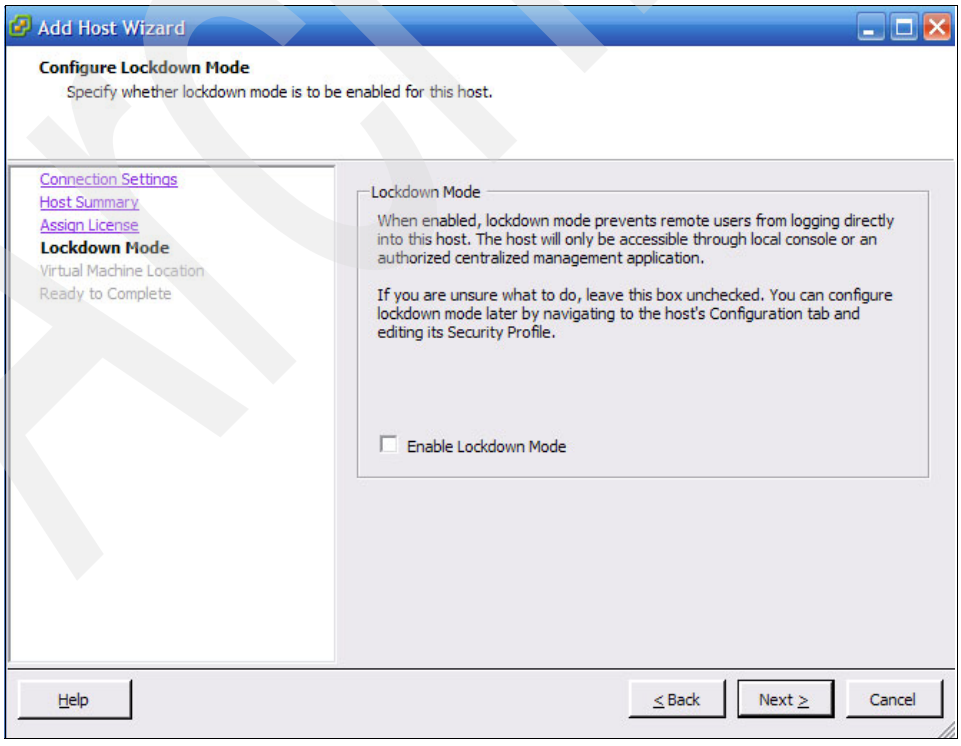

Figure 4-240 Add Host Wizard: Lockdown Mode

On the Virtual Machine Location panel (not shown), select the newly created data store for the virtual machine location, and then click **Next**.

Review the summary (Figure 4-241). Click **Finish** to start the host addition.

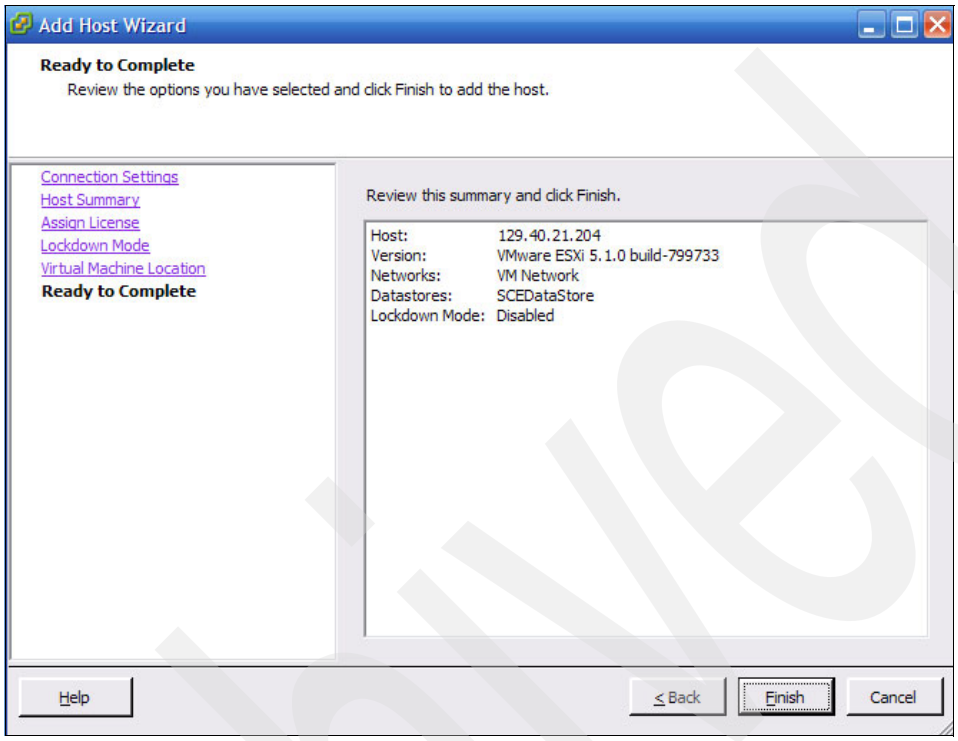

Figure 4-241 Add Host Wizard: Summary

When the process bar is finished (Figure 4-242), the host is added and can now be used for deploying other workloads, appliances, and templates.

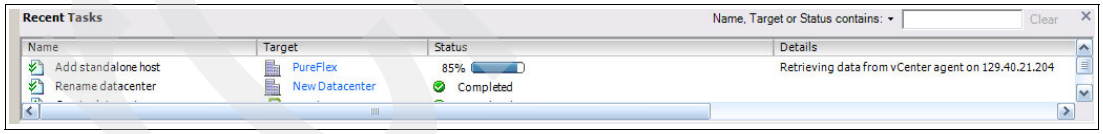

Figure 4-242 New data center is completed

These host addition tasks must be repeated for each ESXi in the chassis. After the vCenter server is fully configured with all of the available resources, IBM SmartCloud Entry can be installed and set up.

5

SmartCloud Entry installation and configuration

This chapter describes the process of IBM SmartCloud Entry installation and initial administrative tasks including the creation of an administrator, enabling billing, and creation of new users. In addition, this chapter covers connecting IBM SmartCloud Entry to the x86 and Power system-based compute nodes configured in Chapter 4, “Hardware setup and configuration” on page 17.

5.1 Installing IBM SmartCloud Entry

IBM SmartCloud Entry can be installed on x86 or Power nodes. For the purposes of this document SmartCloud Entry is installed on an x86 node with a Linux operating system installed. The installation of the operating system is not covered in this book.

After a virtual server is created and the operating system installed, the installation media (**sce240_linux_installer.bin**) can be copied to the operating system.

Example 5-1 shows the command (**./sce240_linux_installer.bin**) executed on the compute node.

Example 5-1 Command and responses to start the installation of SmartCloud Entry

```
[root@sce-intel SmartCloudEntry]# ./sce240_linux_installer.bin
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...
strings: '/lib/libc.so.6': No such file
Launching installer..
```

The installation process requires the specification of language by number as shown in Example 5-2.

Example 5-2 Specifying locale or language

```
Choose Locale...
----------------
1- Deutsch
2- English
3- Español
.
(additional lines of output were deleted for this example)
.
.
CHOOSE LOCALE BY NUMBER: 2
===============================================================================
(created with InstallAnywhere)
```

Example 5-3 shows the beginning of the installation dialog.

Example 5-3 Preparing CONSOLE mode installation

```
Preparing CONSOLE Mode Installation...
===============================================================================
Introduction
------------
InstallAnywhere will guide you through the installation of IBMSmartCloud Entry
version on Linux.
It is strongly recommended that you quit all programs before continuing with
this installation.
Respond to each prompt to proceed to the next step in the installation. If you
want to change something on a previous step, type 'back'.
You may cancel this installation at any time by typing 'quit'.
PRESS <ENTER> TO CONTINUE:
```

Press Enter multiple times to scroll to further licensing details. A final prompt requires a response of **Y** to continue installation (Example 5-4).

Example 5-4 Accepting license agreement

```
--- preceded by many screens of license agreement information ---

PRESS <ENTER> TO CONTINUE:
DO YOU ACCEPT THE TERMS OF THIS LICENSE AGREEMENT? (Y/N):
Y <<<<<<<<<<<<<<<<<<<<<< note: a reply of "Y" is required in order to continue
```

Choose a link location. The link location represents the directory that contains links to the IBM SmartCloud Entry binaries. These binaries are used to start or uninstall SmartCloud Entry. The links are created as a convenience to the administrator. Configure the directory for these links, as shown in Example 5-5.

Example 5-5 Choosing link location

```
Choose Link Location
--------------------
Where would you like to create links?
->1- Default: /root/SmartCloud_Entry_24
2- In your home folder
3- Choose another location...
4- Don't create links
ENTER THE NUMBER OF AN OPTION ABOVE, OR PRESS <ENTER> TO ACCEPT THE DEFAULT
```

Specify the installation folder and property folder by either entering the path in response to the prompts, shown in Example 5-6, or choosing the default values. To accept the default, press Enter.

Example 5-6 Choose installation and property folders

```
Please choose an installation folder.
Where Would You Like to Install? (DEFAULT: /opt/ibm):
Please choose a destination folder for the property files.
Where Would You Like to Install the Property Files? (DEFAULT: /root):
```

Before continuing to the actual installation, a preinstallation summary is provided (Example 5-7). If the settings are satisfactory, press Enter.

Example 5-7 Pre-Installation Summary

```
Pre-Installation Summary
------------------------
Please Review the Following Before Continuing:
Property File Install Folder:
/root/.SCE24/
Install Folder:
/opt/ibm/SCE24
Disk Space:
Free: 11242 MB Required: 505 MB
PRESS <ENTER> TO CONTINUE:
```

The installation progress is indicated in the text-based “progress bar.” A completed progress bar is shown in Example 5-8.

Example 5-8 Installation progress bar

```
Installing...
-------------
[------------------|------------------|------------------|------------------]
Install Finished
----------------
IBMSmartCloud Entry finished installing.
```

The next part of the installation is to configure user data. The installation prompts for importing existing SmartCloud Entry data or configuring new data as shown in Example 5-9.

Example 5-9 Specify I want to configure now

```
You can migrate data from a previous version now or migrate the data later by reading
Chapter 12 of the IBMSmartCloud Entry Admin Guide. You can also configure user data now or
later by editing the properties files in /root/.SCE24/
1- I want to migrate the data now.
2- I want to configure the user data now.
3- I will manually migrate or configure the properties later.
ENTER THE NUMBER FOR YOUR CHOICE, OR PRESS <ENTER> TO ACCEPT THE DEFAULT:: 2
```

Add configuration values (Example 5-10).

Example 5-10 Add configuration values

```
================================================
Add Configuration Values
------------------------
Changing these properties is optional. All options can be changed post install in
/root/.SCE24/authentication.properties. Press <Enter> to accept the default value.

Initial admin user name (DEFAULT: admin):
Initial admin name (DEFAULT: SmartCloud Entry Administrator):
Initial administrator password:
Verify initial administrator password:
================================================
IBMSmartCloud Entry has been successfully installed
------------------------------------------------
```

Example 5-11 shows a successful installation.

Example 5-11 IBM SmartCloud Entry has been successfully installed

```
IBMSmartCloud Entry has been successfully installed to:
/opt/ibm/SCE24

If you choose to create a silent install response file, it will be located in this
directory.

1- Create Silent Install Response File
->2- Do Not Create Silent Install Response File

ENTER THE NUMBER FOR YOUR CHOICE, OR PRESS <ENTER> TO ACCEPT THE DEFAULT:
```

5.1.1 Starting IBM SmartCloud Entry with console in the background

If SmartCloud Entry was installed manually (not deployed as a virtual appliance), the SmartCloud Entry server is started using a terminal program to access the command line, and issuing the **/opt/ibm/SCE24/skc** command. Having an open terminal window dedicated solely to running a server is not preferable. If the terminal that was used to start the server closes, the server also exits. To avoid the server from exiting, the preferred way to run the SmartCloud Entry server is by using the **screen** or **nohup** commands.

The **screen** command allows for the terminal in which a program is running to be detached. In this case, this means that the **/opt/ibm/SCE24/skc** command can be issued in a terminal window and then the window can be closed without aborting the command. To accomplish this, type **screen** in a terminal. Next, run the **/opt/ibm/SCE24/skc** command, then type Ctrl+A, then D. At this point, the terminal running the **/opt/ibm/SCE24/skc** command is detached and all windows can be closed. The SmartCloud Entry server will continue to run. In some situations you might need to re-attach to the detached terminal. From any terminal window running on the SmartCloud Entry virtual server, type **screen -r** to return to the terminal window in which the command to start SmartCloud Entry was originally run.

The **screen** program might not be available in some AIX environments, so the **nohup** command can be used instead. The SmartCloud Entry Administrator's guide gives this example of using the **nohup** command to start IBM SmartCloud Entry:

```
nohup /opt/ibm/SCE24/skc -nosplash < /dev/null > /dev/null &
```

This starts the SmartCloud Entry server in the background and discards any output from the OSGI console. Even if the terminal window is closed, the server continues to run. See Chapter 11 of the *IBM SmartCloud Entry Administrator Guide 2.4* for more information about these commands:

https://www.ibm.com/developerworks/mydeveloperworks/wikis/home?lang=en#/wiki/W21ed5ba0f4a9_46f4_9626_24cbbb86fbb9/page/Documentation

5.2 Deploying IBM SmartCloud Entry virtual appliance

SmartCloud Entry is also provided as a VMware virtual appliance. This virtual appliance is easier to install and manage than manually installing. Instead of installing as done in 5.1, “Installing IBM SmartCloud Entry” on page 142, the virtual appliance can be deployed using the Deployment Wizard:

1. Select the source OVA as shown in Figure 5-1. Click **Next**.

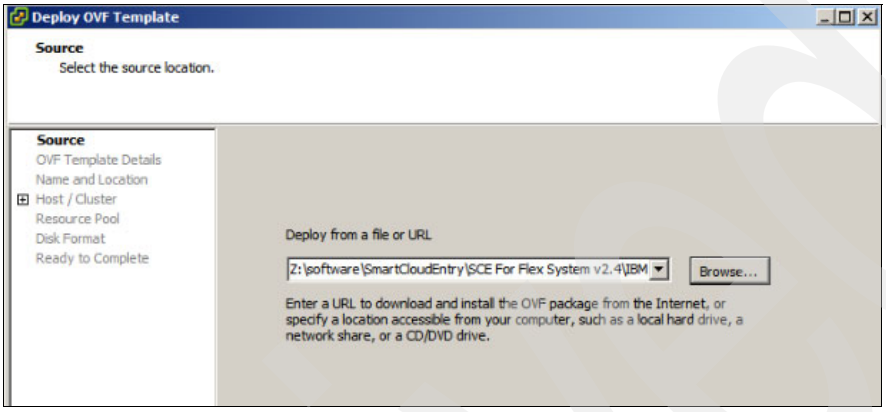

Figure 5-1 IBM SmartCloud Entry Deploy/Configure: Specify file or URL of source

The summary of the virtual appliance is listed (Figure 5-2).

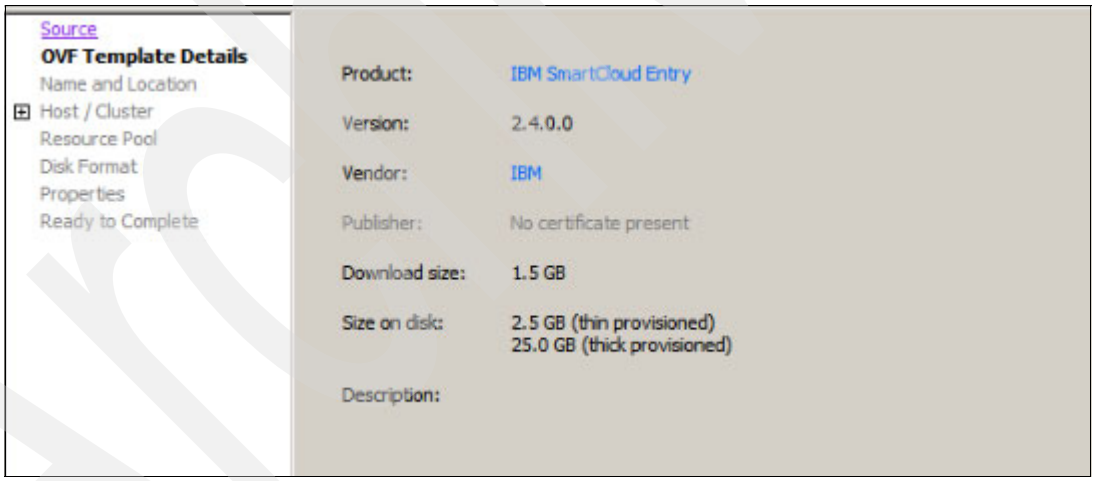

Figure 5-2 IBM SmartCloud Entry Deploy/Configure: OVF template details

2. The appliance has a default name of IBM SmartCloud Entry (Figure 5-3). If necessary, overwrite the name.

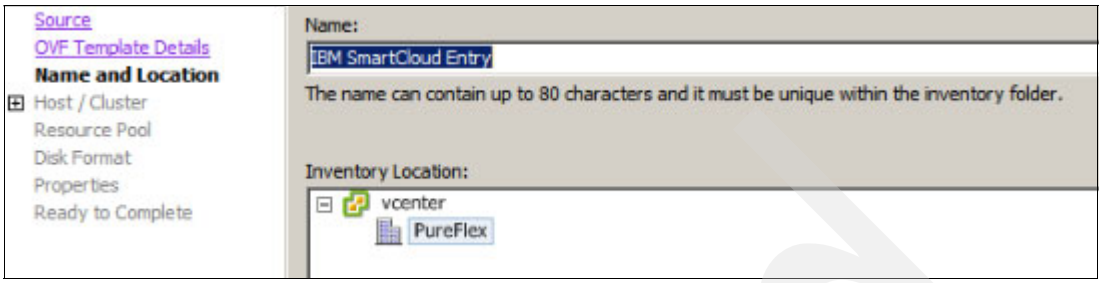

Figure 5-3 IBM SmartCloud Entry Deploy/Configure: Name and Inventory Location

3. The deployment wizard next asks for the data center that is to be used (Figure 5-4).

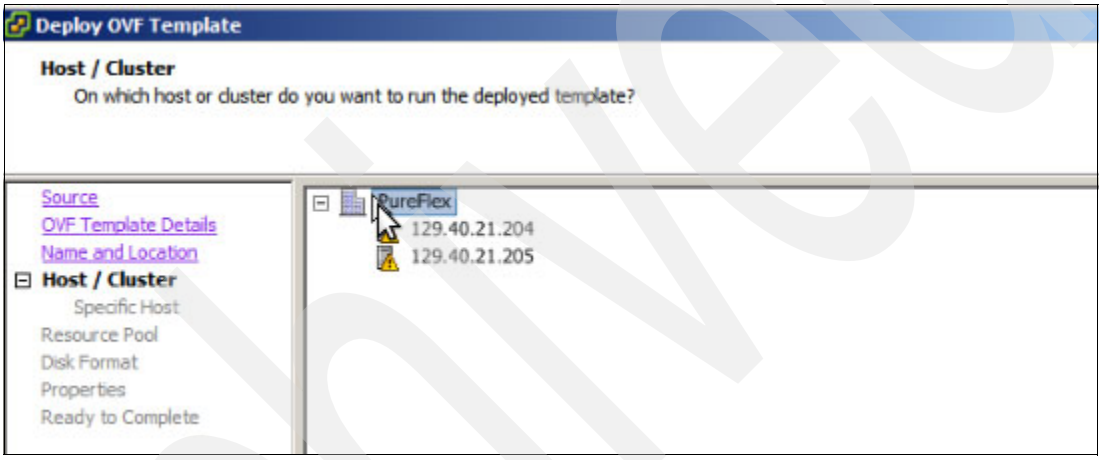

Figure 5-4 IBM SmartCloud Entry Deploy/Configure: Specific data center

4. In the data center, select a preferred host (Figure 5-5).

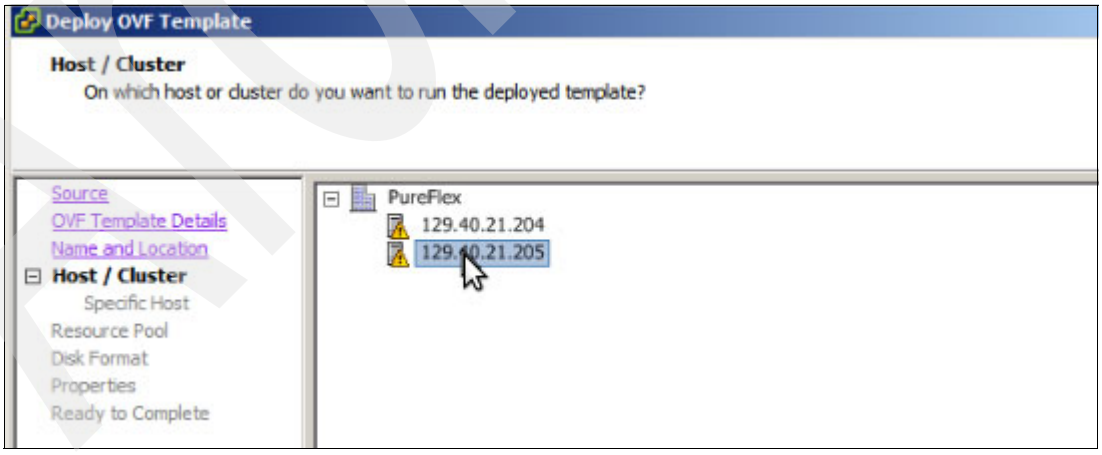

Figure 5-5 IBM SmartCloud Entry Deploy/Configure: Specific Host, IP address selected

5. Select the data store (Figure 5-6) that is created in 4.7.4, “Data store” on page 120.

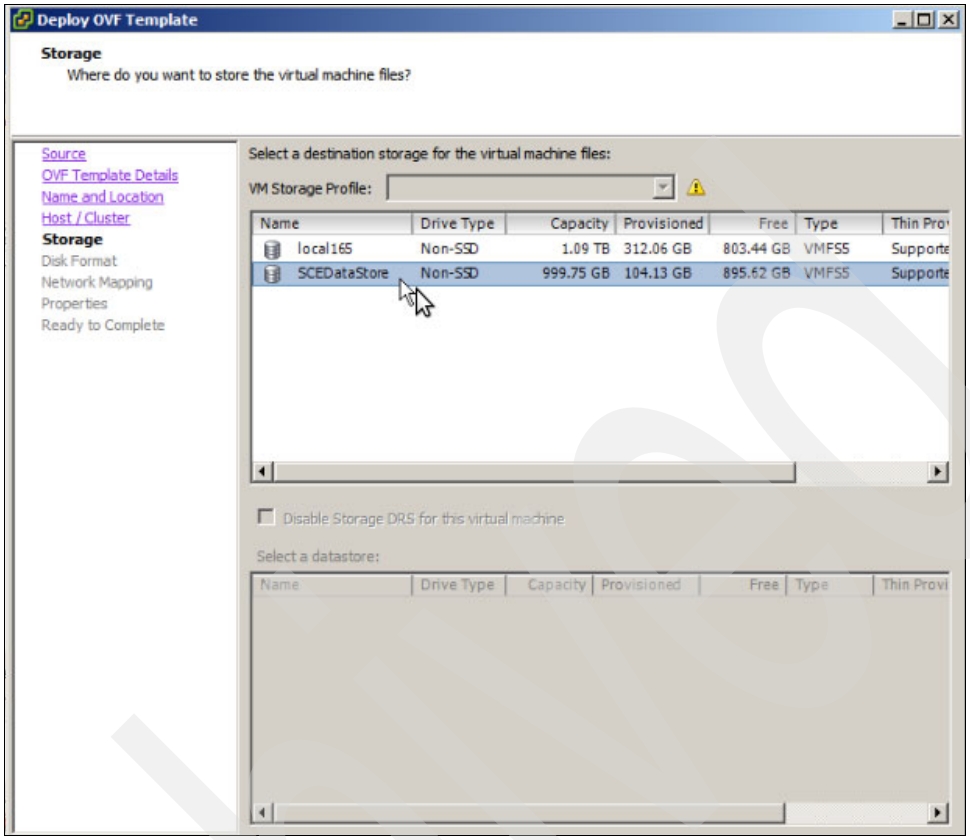

Figure 5-6 IBM SmartCloud Entry Deploy/Configure: Storage

6. For this appliance, the thin provision disk format is applicable (Figure 5-7).

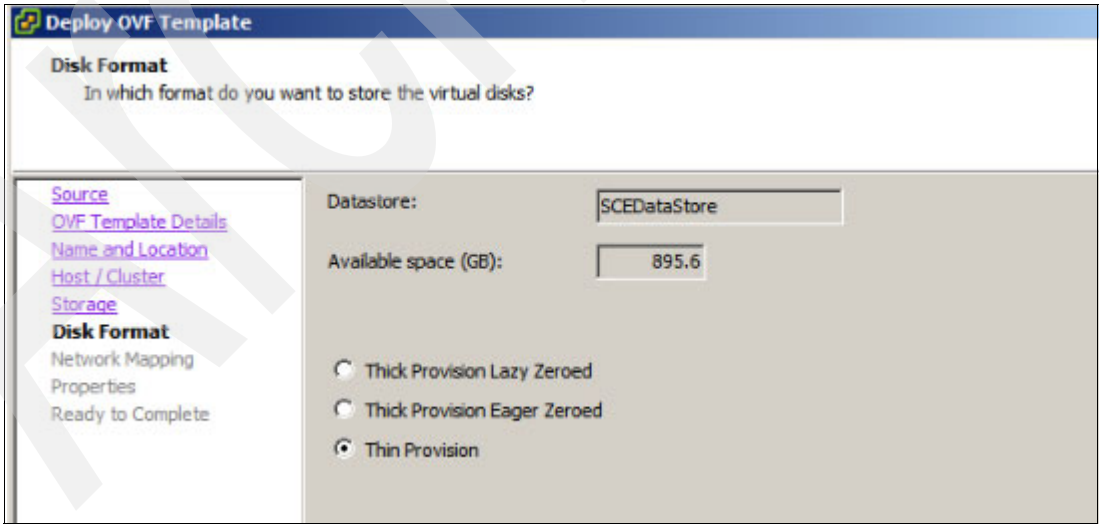

Figure 5-7 IBM SmartCloud Entry Deploy/Configure: Specify Disk Format

7. Select the appropriate VM network (Figure 5-8).

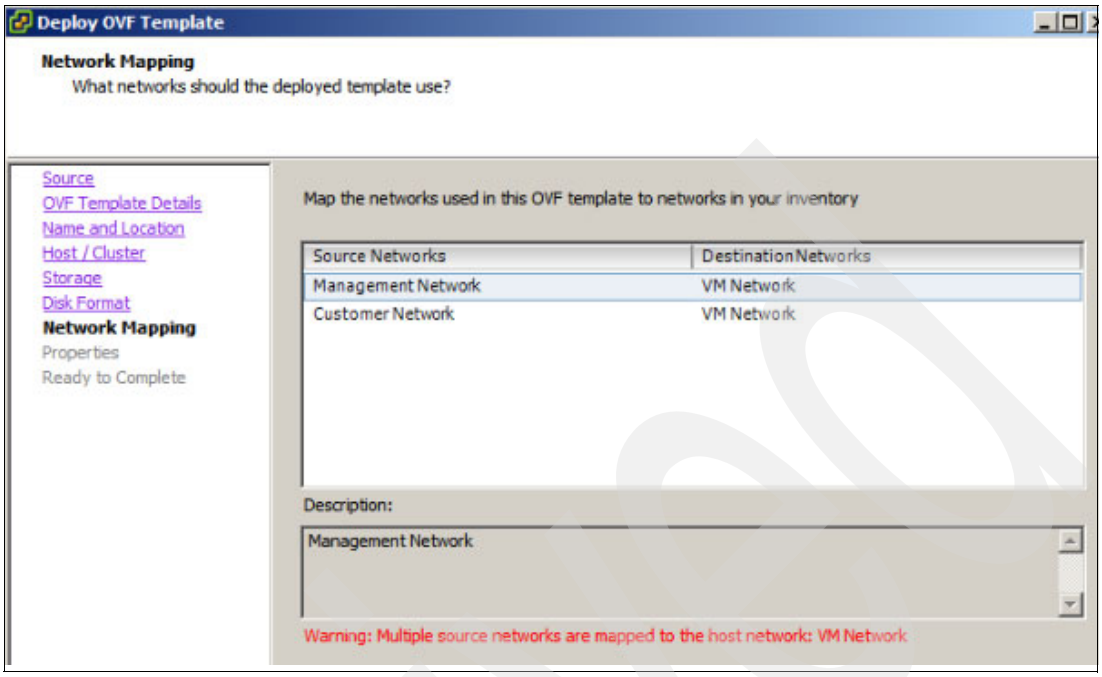

Figure 5-8 IBM SmartCloud Entry Deploy/Configure: Select the network

8. If necessary, modify various aspects of this appliance and click **Next** (Figure 5-9).

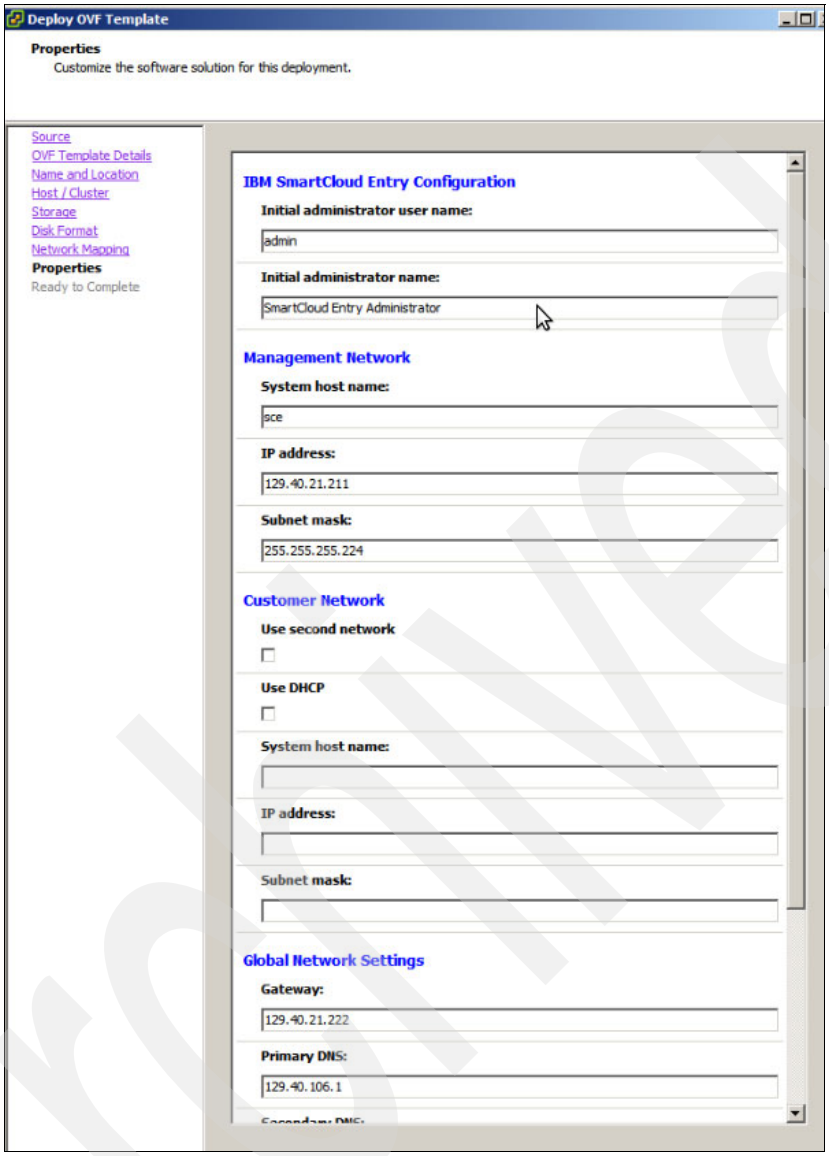

Figure 5-9 IBM SmartCloud Entry Deploy/Configure: Customize properties of a software solution

9. Verify the values in the summary window (Figure 5-10).

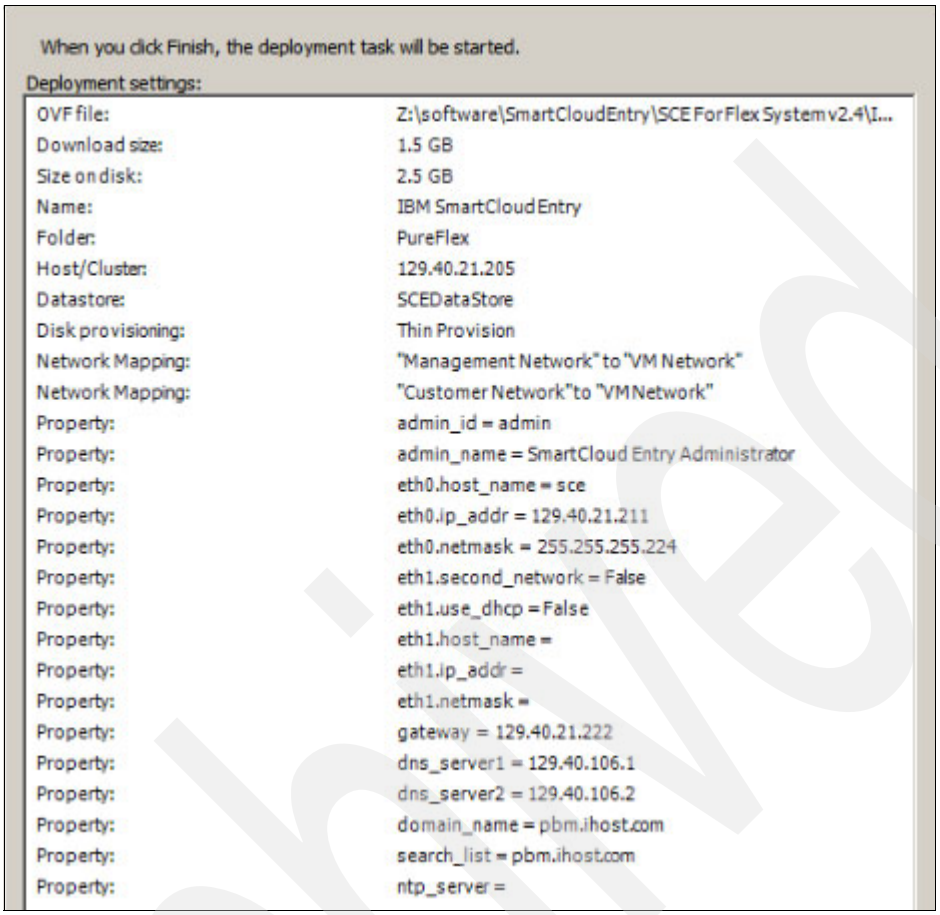

Figure 5-10 IBM SmartCloud Entry Deploy/Configure: Properties summary

5.3 Configuring cloud network pools

In 4.3, "Networking setup" on page 27, three VLANs were configured on the Power nodes. Similar VLANs are configured on the x86 nodes with VMware. SmartCloud Entry must be configured to set the proper information for each of the workload network configurations when deployed.

As the IBM SmartCloud Entry Administrator, log in and select the Network section of the configuration tab.

Figure 5-11 shows a Default Network Configuration option, has a single IP available to it, and is defined by default to enable deployment. To match the VLAN setup for this example, two new network configurations must be created. Click **New Configuration**.

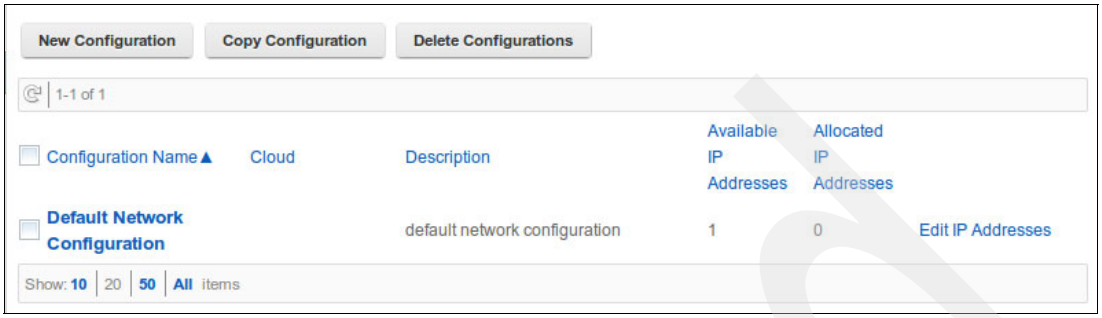

Figure 5-11 IBM SmartCloud Entry: Administrator Network section

The New Network Configuration panel opens. In this panel, several items can be set, as Figure 5-12 shows. Specify the appropriate data, and save the network configuration.

Note: Cloud must be set to None so that more than one pool can use this configuration. The same is true for the Network ID being set to blank.

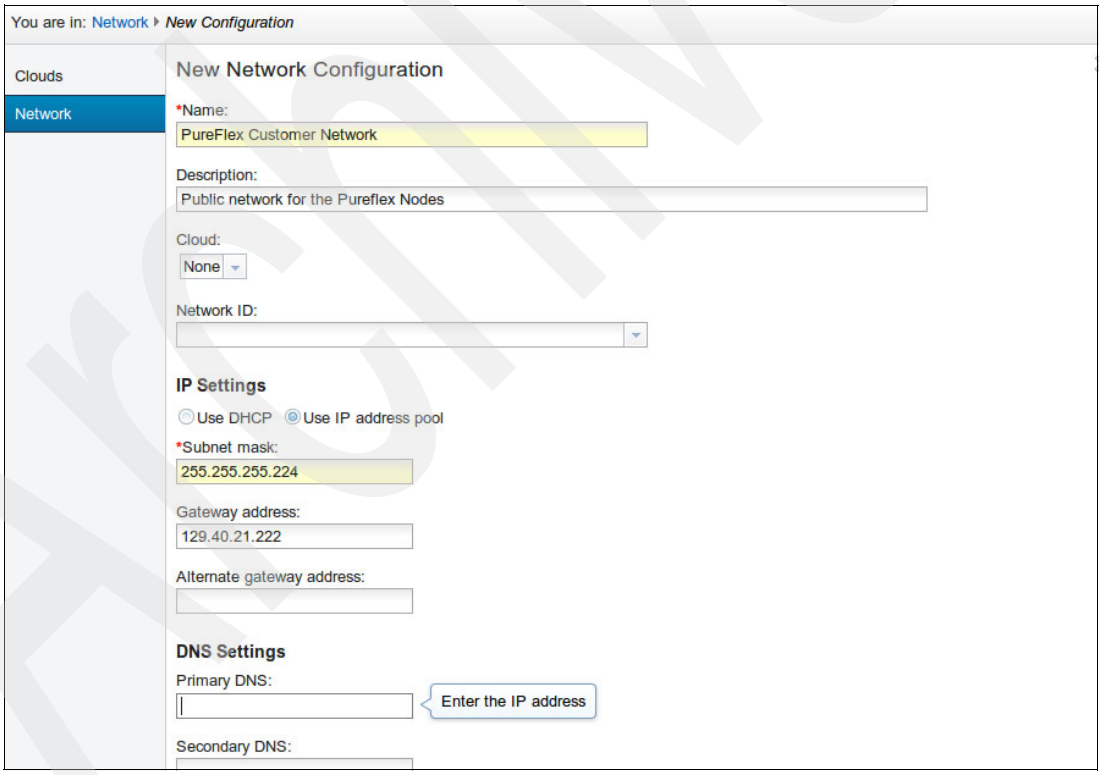

Figure 5-12 New Network Configuration

IBM SmartCloud Entry prompts you about whether the IP address should be added to the configuration (Figure 5-13). Click **Yes**.

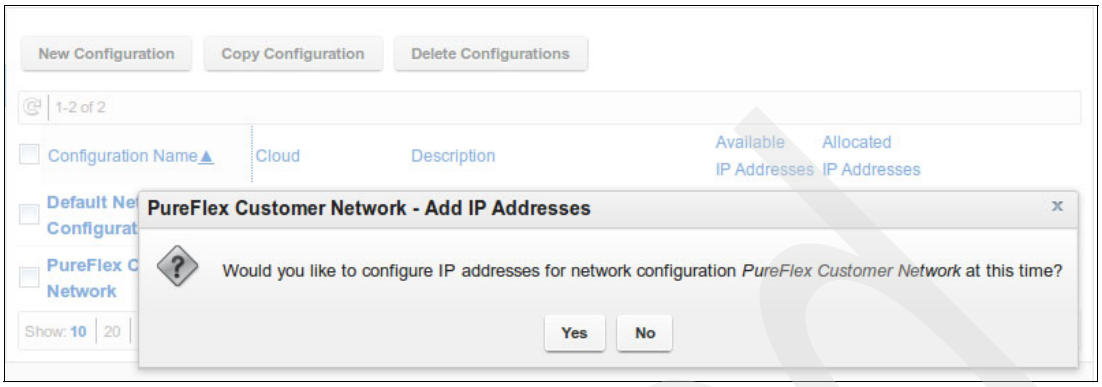

Figure 5-13 Confirmation Prompt: Would you like to configure IP addresses

The IP address panel is displayed (Figure 5-14).

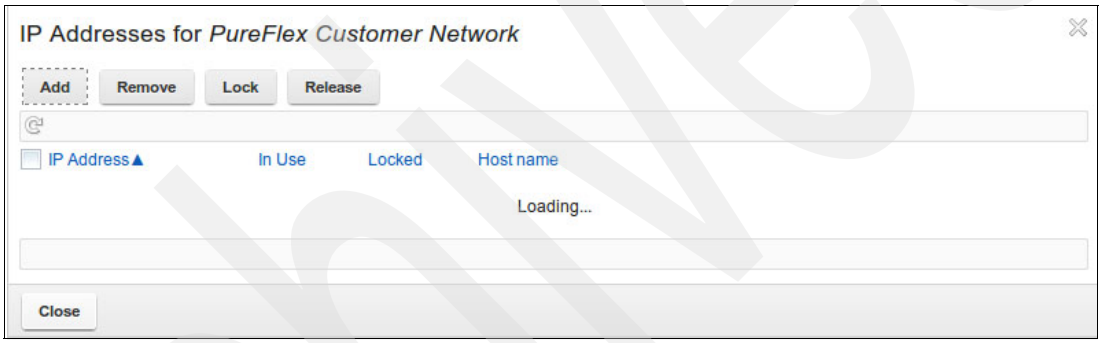

Figure 5-14 Requesting: Add IP Addresses for PureFlex Customer Network

New IPs can be added by clicking **Add**. A single IP can be added to the configuration or a range. When choosing to add a range, enter the appropriate start address and end address as Figure 5-15 shows. Click **OK**.

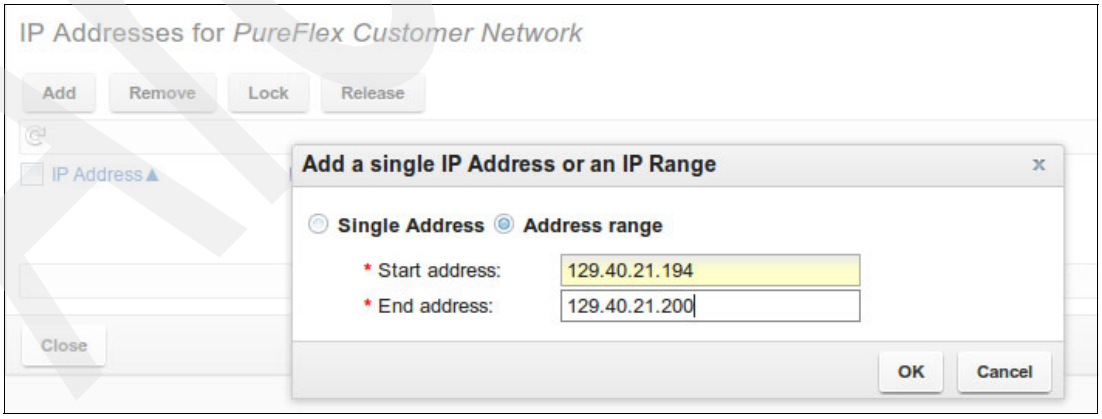

Figure 5-15 Add single IP Address or Range of Addresses

A message indicates that the IP address was added.

The network setup that is being used for this environment also has a private network that is used for the virtual machines that form the basis of the cloud services to communicate with each other. The network configuration setup for this private network is similar to the one described previously. Enter the appropriate data (Figure 5-16).

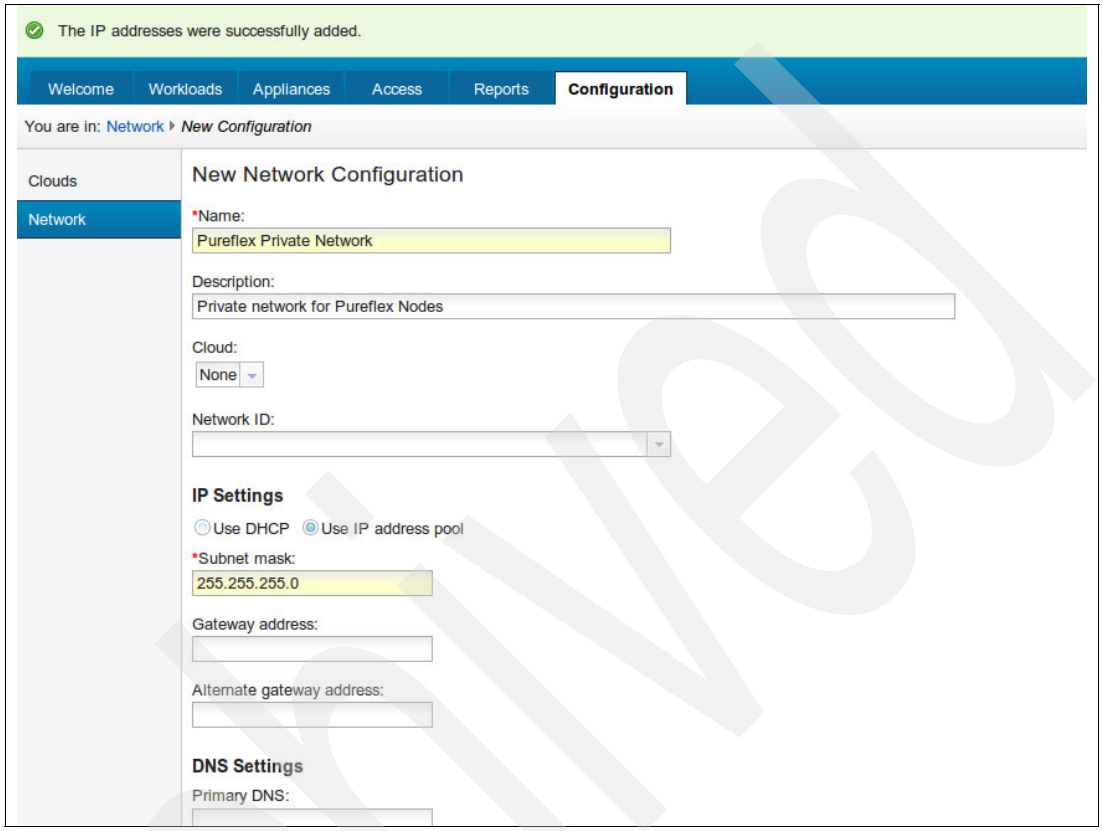

Figure 5-16 New network configuration

Add an IP address range (Figure 5-17),

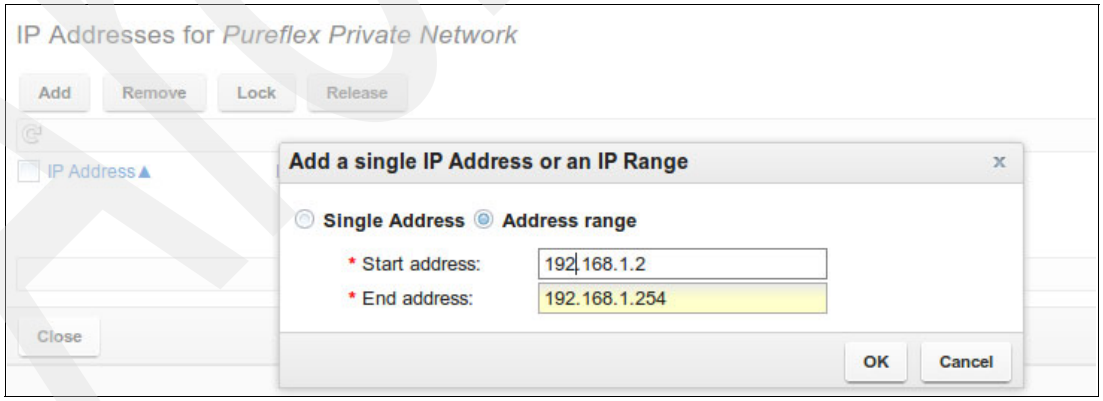

Figure 5-17 Add an IP address range

Verify that the addresses were added correctly (Figure 5-18).

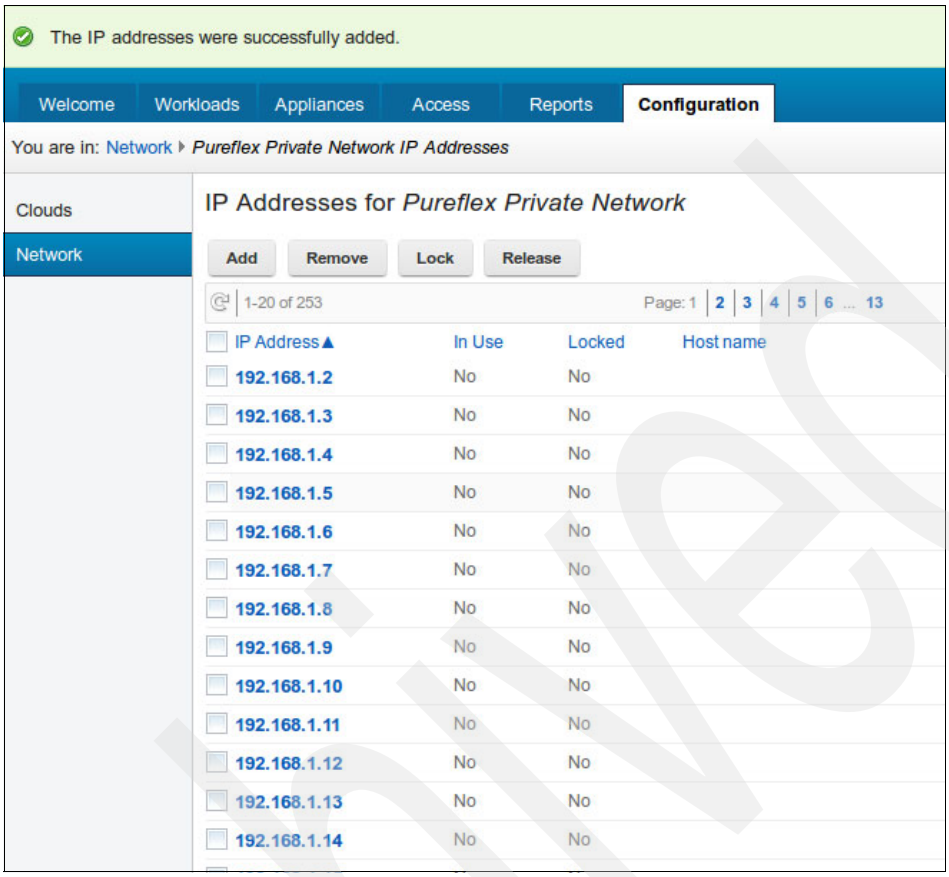

Figure 5-18 Summary of IP addresses

Now that both the private and the public network configurations are defined to IBM SmartCloud Entry, cloud pools can be added and the networks can be assigned to them.

5.4 Adding an x86 node based cloud

For IBM SmartCloud Entry to use the x86 nodes that have VMware installed on them, a new cloud configuration must be added. As the IBM SmartCloud Entry Administrator, click **Clouds** on the Configuration tab (Figure 5-19). Notice that no cloud configurations are defined by default. Click **Add Cloud**.

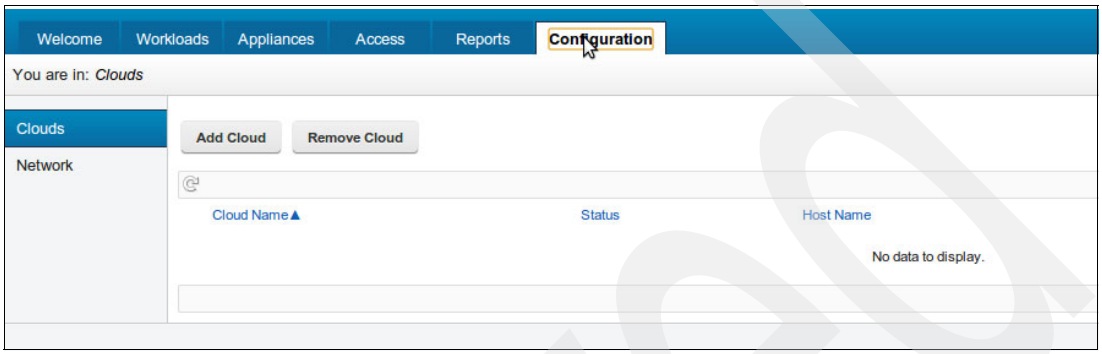

Figure 5-19 Selecting Clouds on Configuration tab

The Add Cloud Configuration panel opens (Figure 5-20).

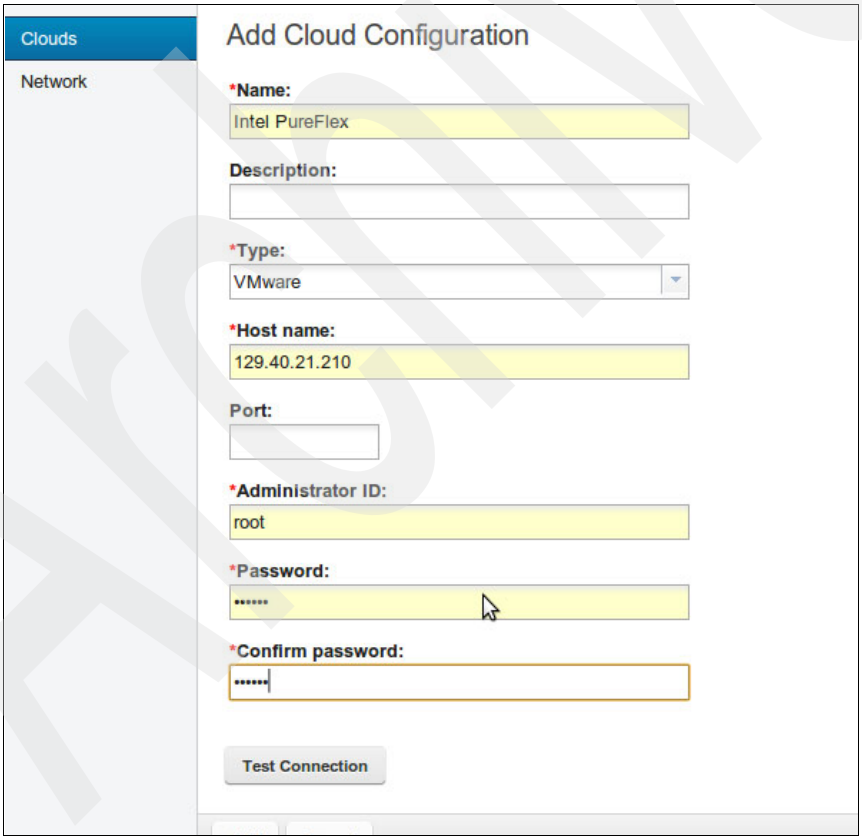

Figure 5-20 Panel used to Add Cloud Configuration

For the VMware environment, use the vCenter Server that was installed in 4.7.5, “Deploy and configure VMware vCenter Server Appliance” on page 123. On this same panel set the following values:

- Type: Set to VMware.
- Host name: Set to the IP address or host name of the vCenter Server.
- Administrator ID: Set to the ID for vCenter Server administrator.
- Password: Set for the vCenter Server administrator's password.

After entering this information, a good approach is to test the connection and the parameters that were entered by clicking **Test Connection** (Figure 5-21).

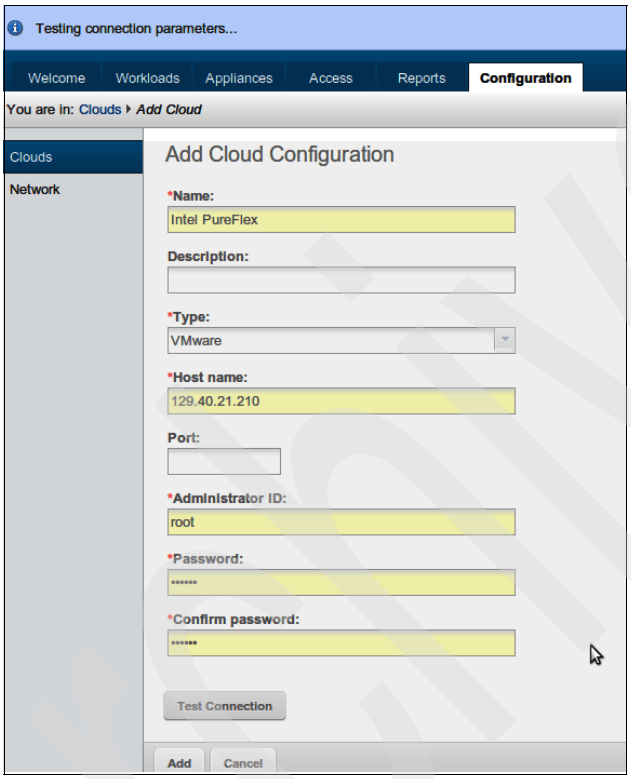

Figure 5-21 Test connection

The notification to the administrator indicates that the test was successful (Figure 5-22). Click **Add**.

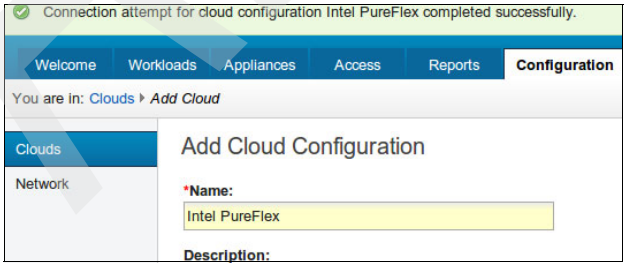

Figure 5-22 Test connection successful

The Certificate Information window opens (Figure 5-23). If the information is correct, click **Accept Certificate**.

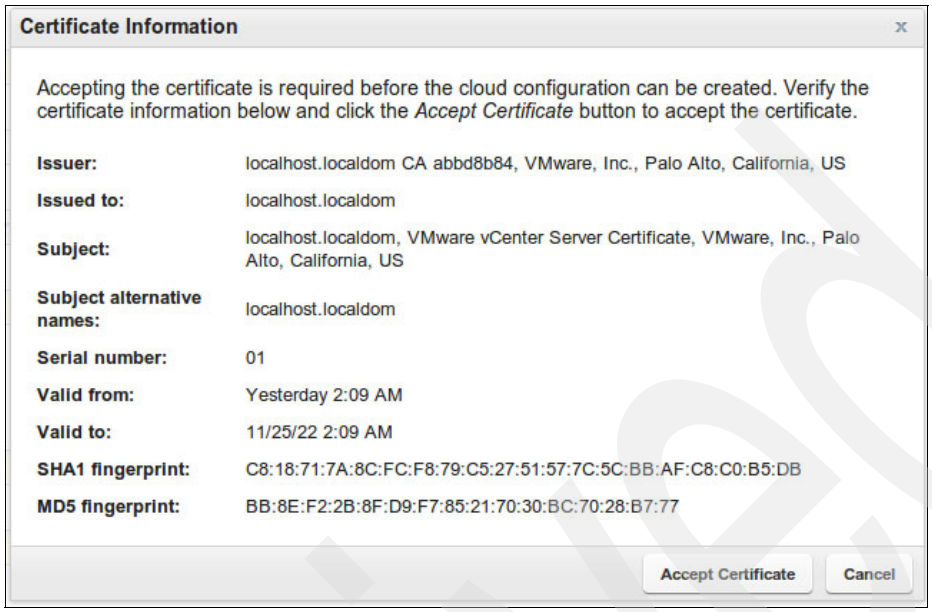

Figure 5-23 Certificate Information

The cloud configuration is added successfully, as the message indicates and the you are returned to the Clouds sections of the Configuration tab (Figure 5-24).

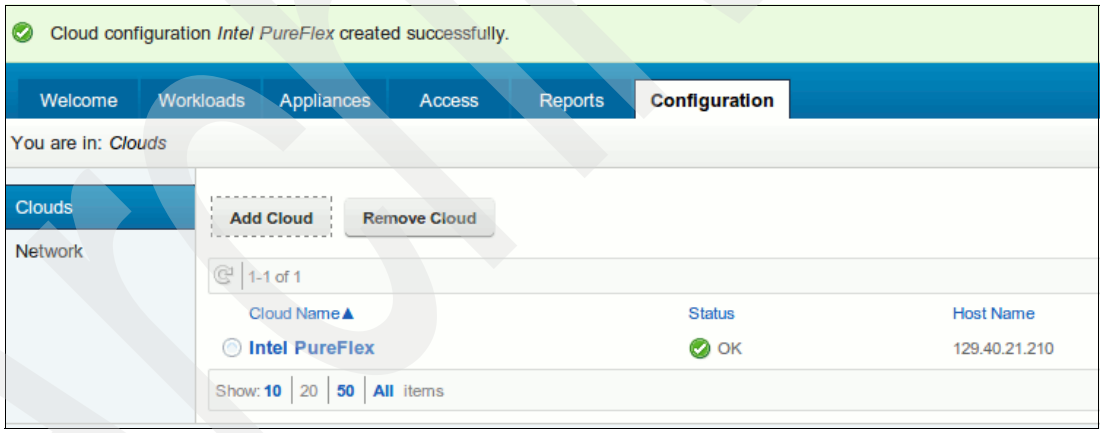

Figure 5-24 Cloud Configuration Created Successfully

Notice that there are workloads automatically discovered (Figure 5-25). These workloads correspond to the virtual machines that were deployed earlier.

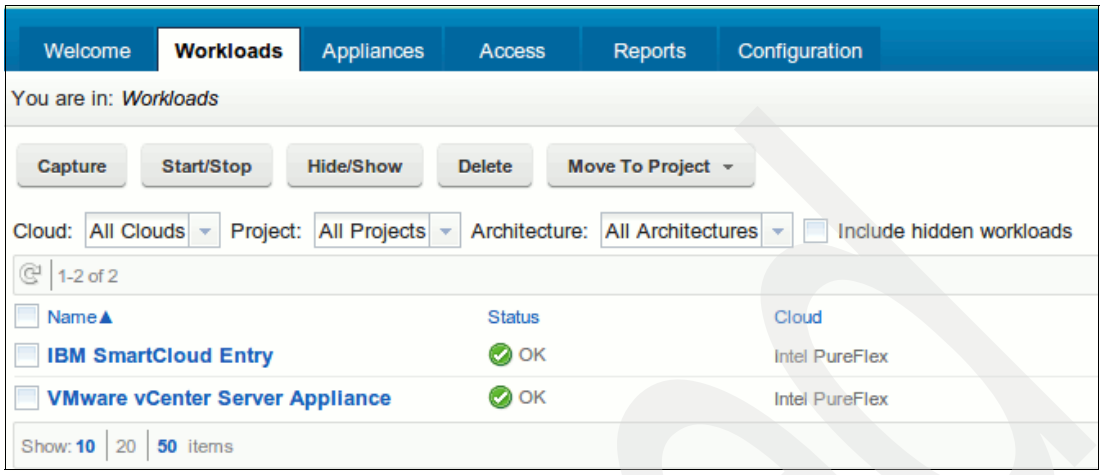

Figure 5-25 Workloads that correspond to deployed VM

5.5 Adding a Power node-based cloud

Similar to 5.4, “Adding an x86 node based cloud” on page 156, SmartCloud Entry must be configured to “talk to” the Flex Service Manager to deploy cloud services on Power nodes. The first step is to define the following values:

- Type: Set to VMControl.
- Version: Set to 2.4.1.x.
- Host name: Set to the IP address or host name of the FSM.
- Administrator ID: Set to the ID for FSM.
- Password: Set for the FSM.

After entering this information, a good approach is to test the parameters and the connection by clicking **Test Connection** (Figure 5-26). As with the x86, click **Add**.

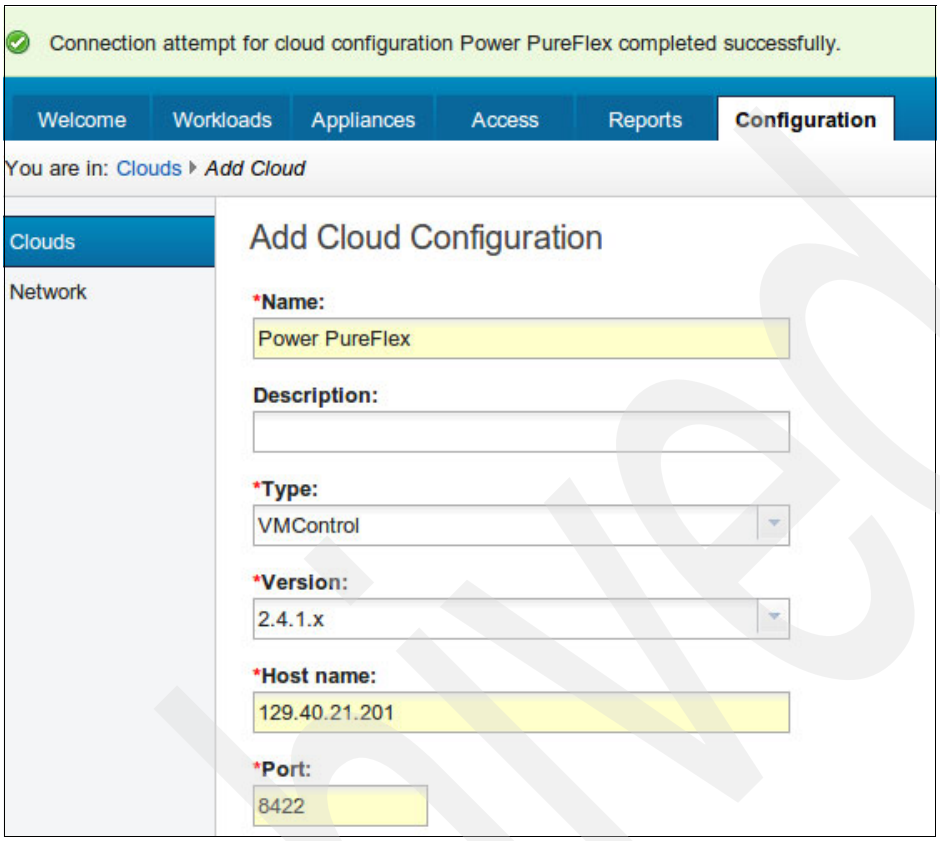

Figure 5-26 Add Cloud Configuration

Click **Accept Certificate** (Figure 5-27).

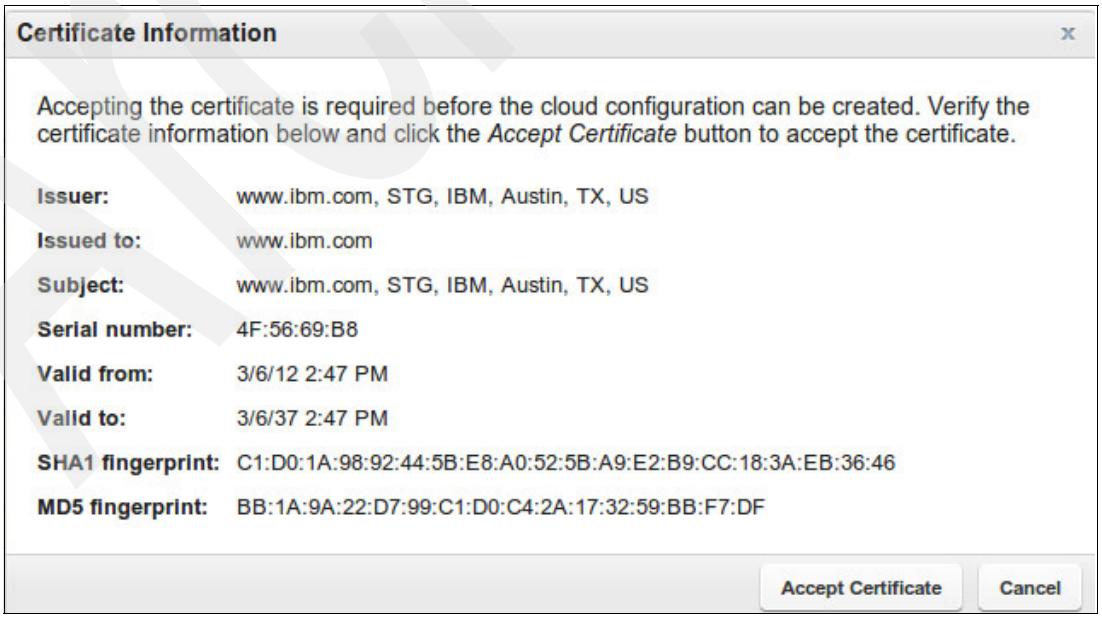

Figure 5-27 Certificate Information

The new workload (Figure 5-28) that is automatically discovered is the AIX workload that was captured in 4.6.7, “Capture virtual server” on page 89.

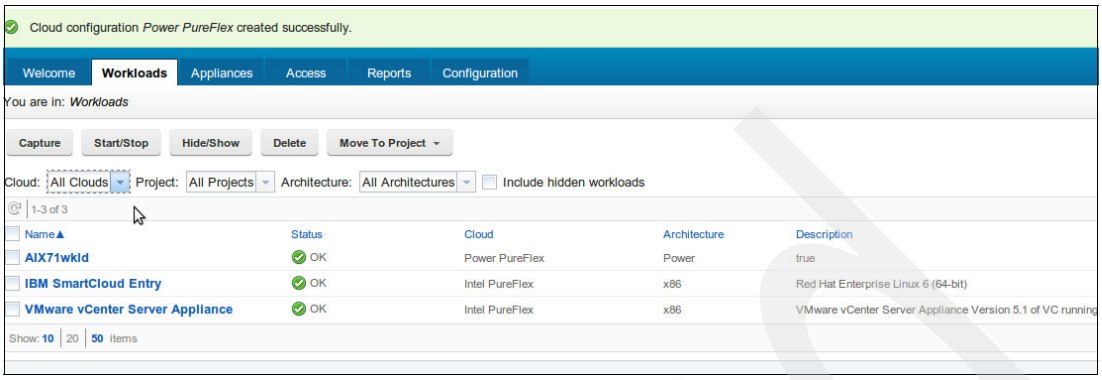

Figure 5-28 New workload: Automatically discovered

5.6 Enable email notifications

The notifications that SmartCloud Entry provides for various events can be sent to users and administrators as email messages. This allows administrators to monitor asynchronous events as they happen and not have to constantly poll the system to see what events warrant their attention.

To enable this functionality, a SMTP relay host must be set. Setting up a relay host is outside the scope of this book.

Complete the following steps:

1. Log in to the command line interface of the host that has IBM SmartCloud Entry installed, or the virtual appliance:

   ```
   ssh sysadmin@129.40.21.211
   sysadmin@129.40.21.211's password:
   ```

2. Edit the `email.properties` file:

   ```
   IBM-SCE> vi ~/.SCE24/email.properties
   ```

3. Look for the following commented line:

   ```
   #com.ibm.cfs.email.relay.host=email.server.com
   ```

4. Edit the line as follows:

   a. Remove the pound sign (#) so that the line is now uncommented.

   b. Update with the appropriate server information. The line now reads as follows:

      ```
      com.ibm.cfs.email.relay.host=smtpserver.cloud.private.com
      ```

5. Save the file and restart IBM SmartCloud Entry; the output with the command is as follows:

   ```
   IBM-SCE> /etc/init.d/sce restart
   Stopping IBMSmartCloud Entry: [ OK ]
   Starting IBMSmartCloud Entry: [ OK ]
   IBM-SCE>
   ```

5.7 Enable and configure billing

IBM SmartCloud Entry supports a metering system that allows to assign billing information to the resources that are provided by the Flex based environment. Currently, the following characteristics of the cloud services can be associated with the billing system.

- Processor
- Memory
- Disks

By default, SmartCloud Entry billing is disabled. This means that no references to Accounts are listed on the Access Menu when the Access tab is selected, as shown in Figure 5-29.

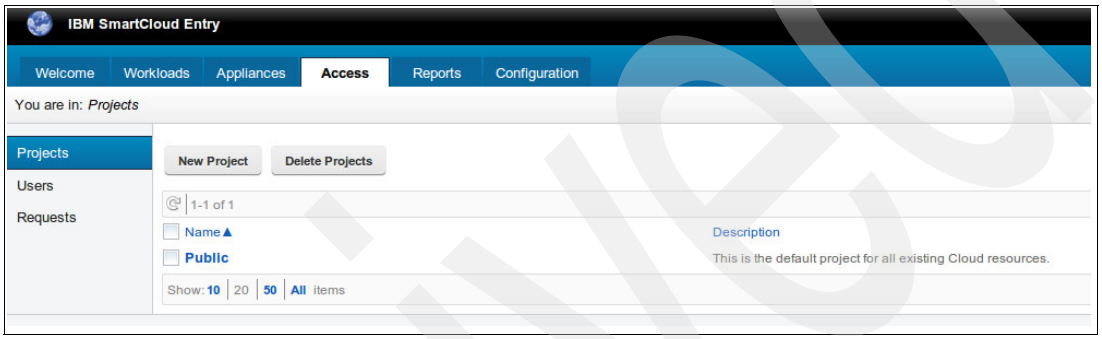

Figure 5-29 Access menu with billing disabled

If billing functionality is required, it must be enabled in the properties file, which is in the following location in the configuration directory:

`~/.SCE24/billing.properties`

To enable billing, edit the `billing.properties` file as follows:

1. Open the file and look for the following line:

   `com.ibm.cfs.billing.enabled=false`

2. Edit the line so the value is set to true:

   `com.ibm.cfs.billing.enabled=true`

After enabling billing, a new subtab (Accounts) is now available on the left side of the Access tab of the IBM SmartCloud Entry portal as shown in Figure 5-30.

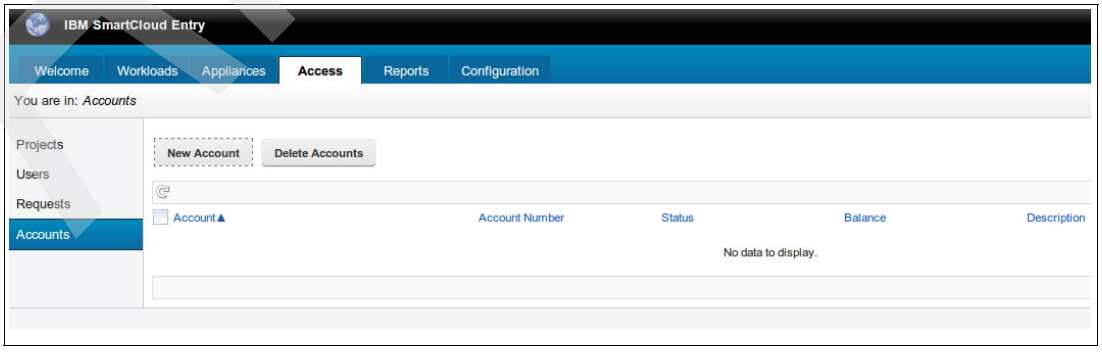

Figure 5-30 Access menu with billing enabled

Click the new **Accounts** subtab in the Access Menu. Two choices are now available in the form of New Account and Delete Account. To add a new account, click **New Account**.

The New Account dialog panel opens (Figure 5-31).

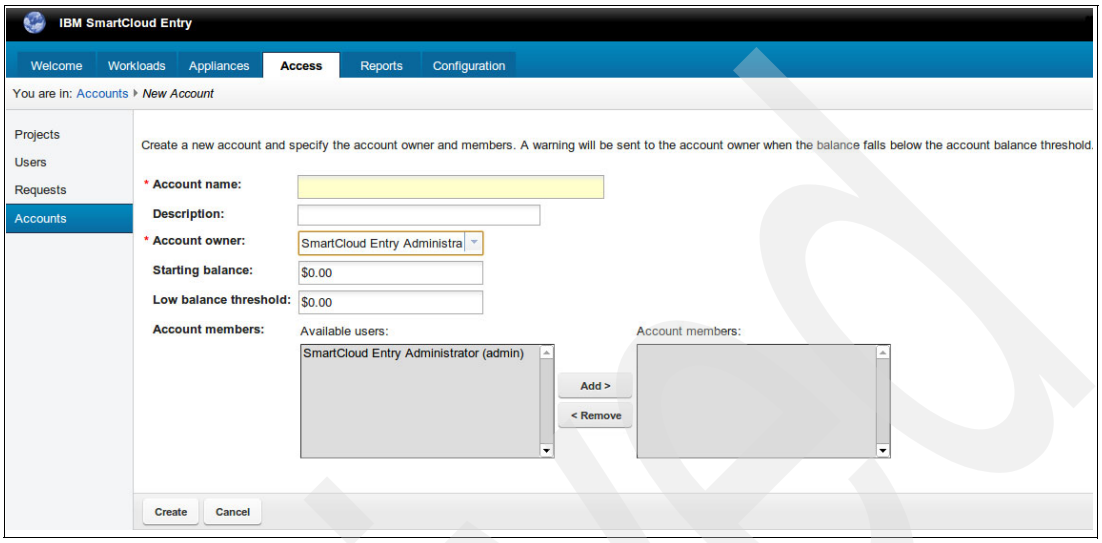

Figure 5-31 New Account dialog

Use this New Account dialog to set the default information for this account. This information includes the account owner and the account members. When a bill becomes due, the account owner will be emailed with the bill contents.

The starting amount of credits and the low balance threshold are also set here. Figure 5-32 shows an example of creating a new account with a starting balance of $100, a low balance threshold of $10, and only one account member. Click **Create**.

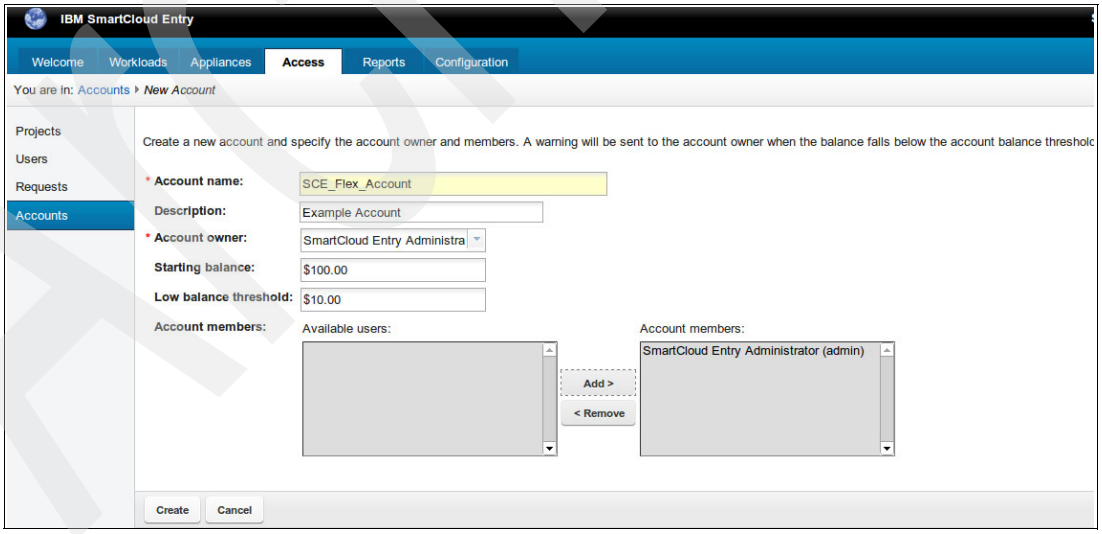

Figure 5-32 New Account starting balance and low threshold

The account is established and you, as the administrator, are notified, (Figure 5-33).

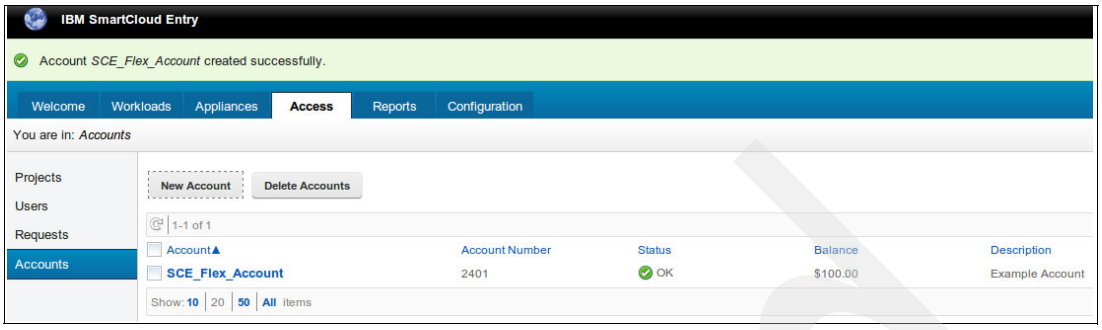

Figure 5-33 New Account has been created

When modifying billing information for these accounts becomes necessary, a similar set of panels can be used. Simply click the name of the account to be modified.

When an account balance reaches the threshold set, the account members are notified by email that their accounts are delinquent. Figure 5-34 shows that account number 2551 is overdrawn. The default delinquency policy shuts down all workloads associated with this account.

```
SmartCloud Entry Administrator,

A delinquency policy has been run on your SmartCloud Entry billing account.

Account Number: "2551"
Name: "SCE_Flex_account"
Description: "Example Account"
Current Balance: -USD 0.10
Threshold: USD 10.00
Is Delinquent: true
Delinquency Policy: Shutdown deployments

If you need to update any of your account details, you can do so from the SmartCloud Entry Web interface once you are logged in.

Thank you.

**This is a machine generated message, please do not reply.
```

Figure 5-34 Delinquent account email

5.8 Creating users

To create a new user, completing the following steps:

1. Log in as SmartCloud Entry Administrator.

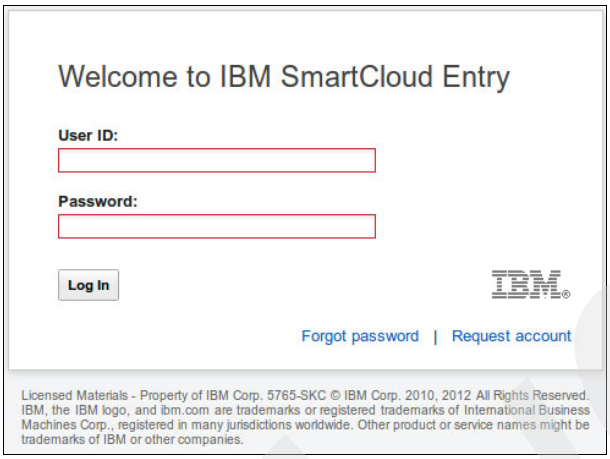

Figure 5-35 Log in to IBM SmartCloud Entry

2. Navigate to the Access menu (using Access tab) in the IBM SmartCloud Entry interface, click **Users**, and then click **New User** (Figure 5-36).

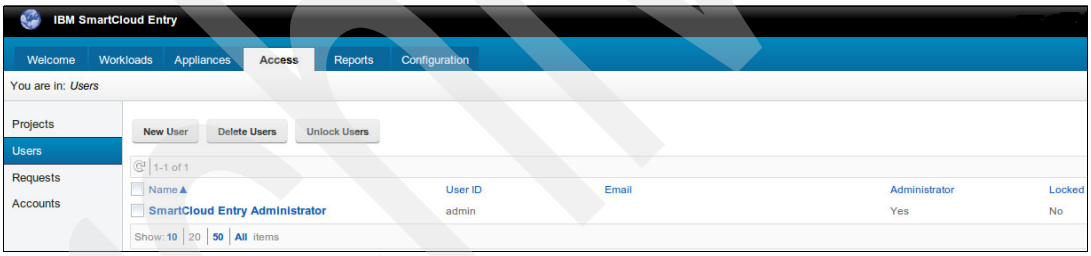

Figure 5-36 Specifying new user

3. Provide requested user information. Figure 5-37 and Figure 5-38 on page 167 provide before-and-after views of the New User panel. The information provided indicates a request for the creation of a user with the following attributes:
   - User ID: sce_user02@example.com
   - Name: SCE_User02
   - Email: sce_user02@example.com
   - Password: passW0rd

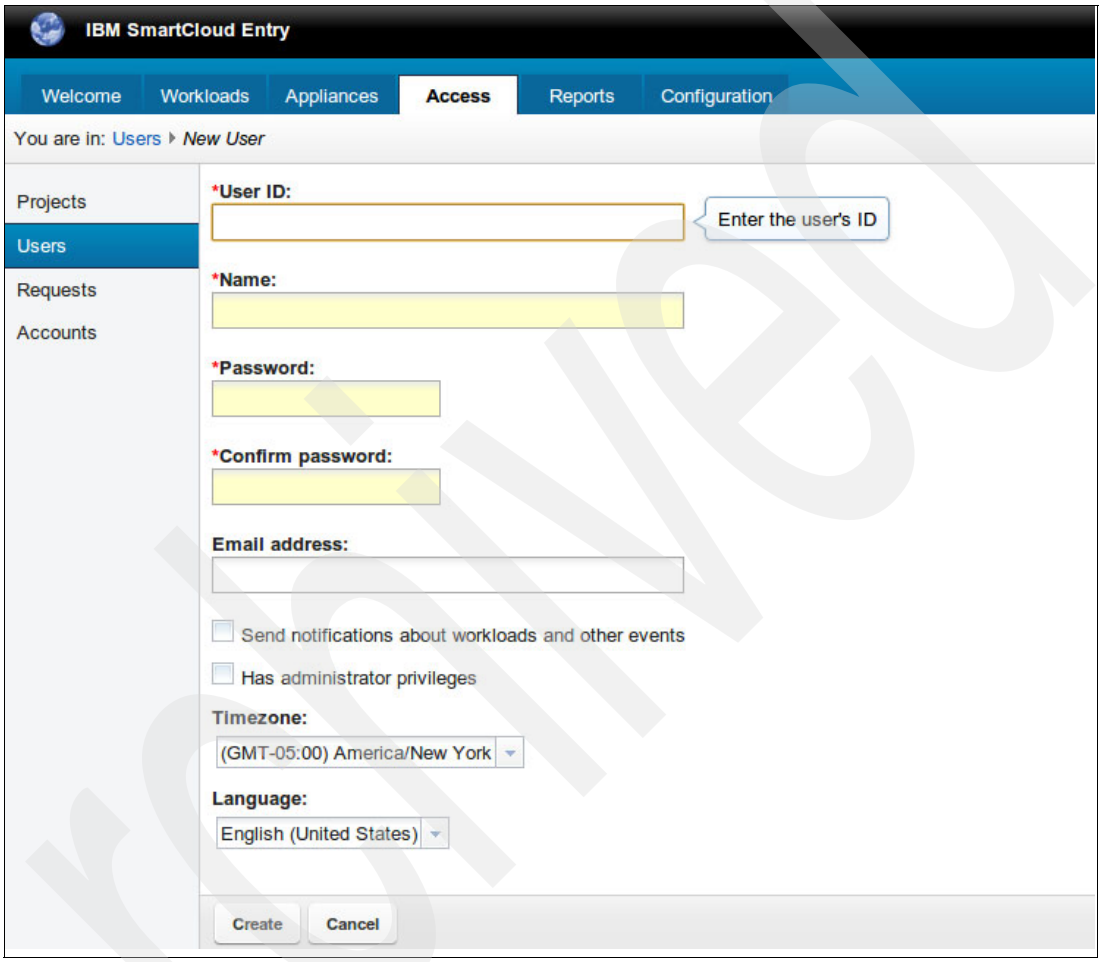

Figure 5-37 Enter new user specifications

4. After completing the information, click **Create** (Figure 5-38).

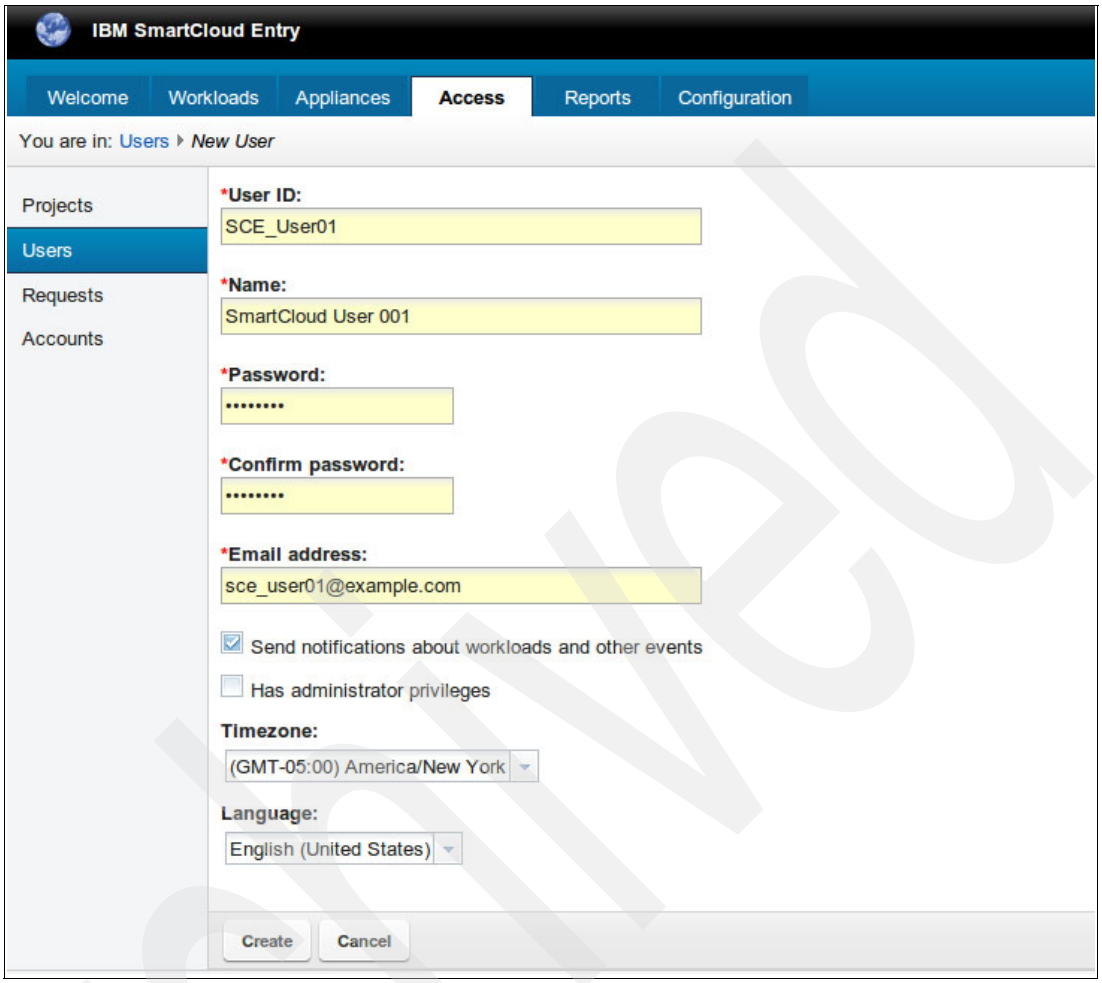

Figure 5-38 New user information provided

A summary panel shows the list of currently defined users (Figure 5-39). Additional new users may be added by clicking **New User** and repeating the previous steps.

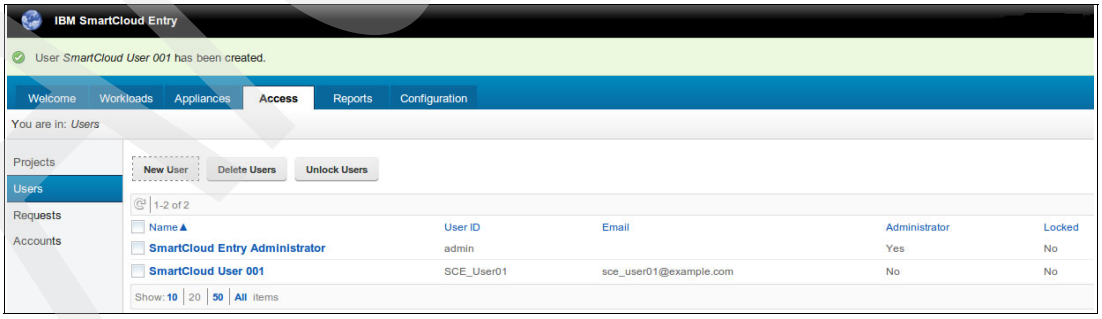

Figure 5-39 New user has been created

5.9 Creating projects

To create a new project, the basic steps are as follows:

1. Log in as SmartCloud Entry Administrator.
2. Click **Projects** and then click **New Project** (Figure 5-40).

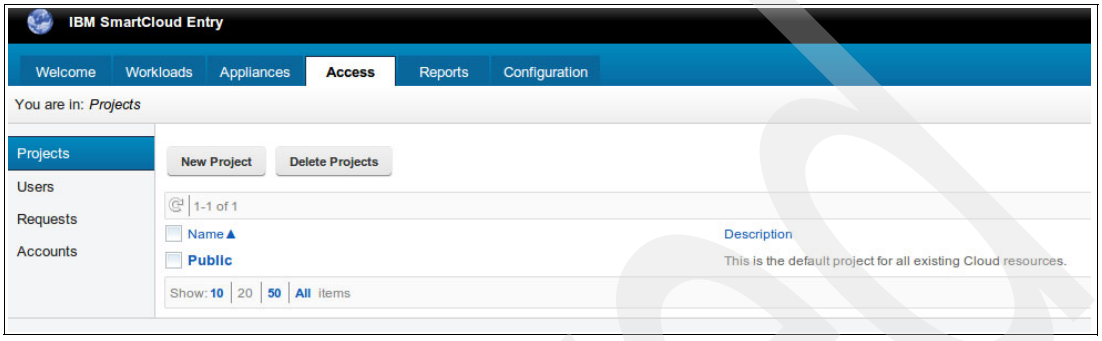

Figure 5-40 Projects Panel with New Project as a choice

3. The New Project panel opens (Figure 5-41). Enter the name and description of the project to be created. Click **Create**.

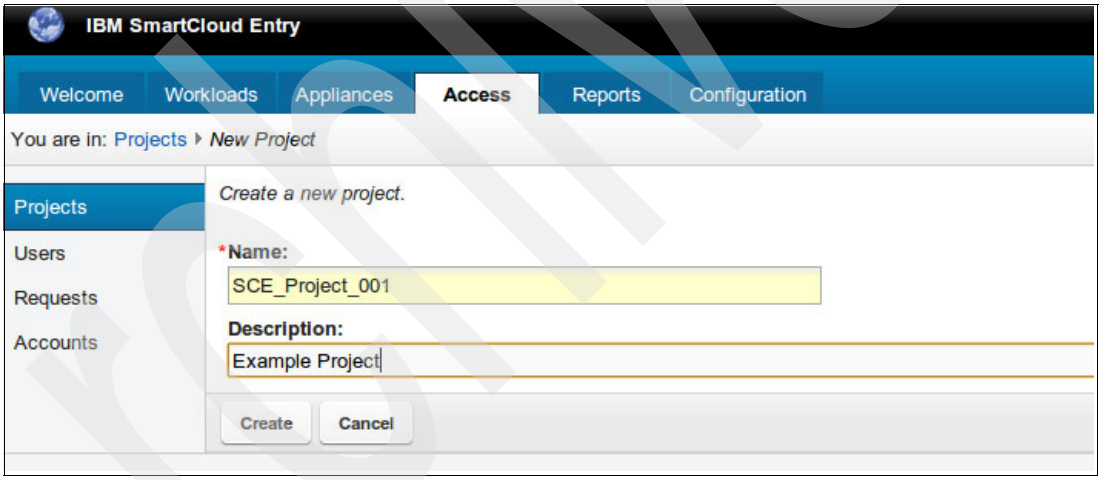

Figure 5-41 Provide new project description

A new project is created; it is listed on the Projects panel (Figure 5-42). Note the `Created project SCE_Project_001` message. On the actual panel, a green icon with a white check mark is displayed to the left of the message near the top of the panel.

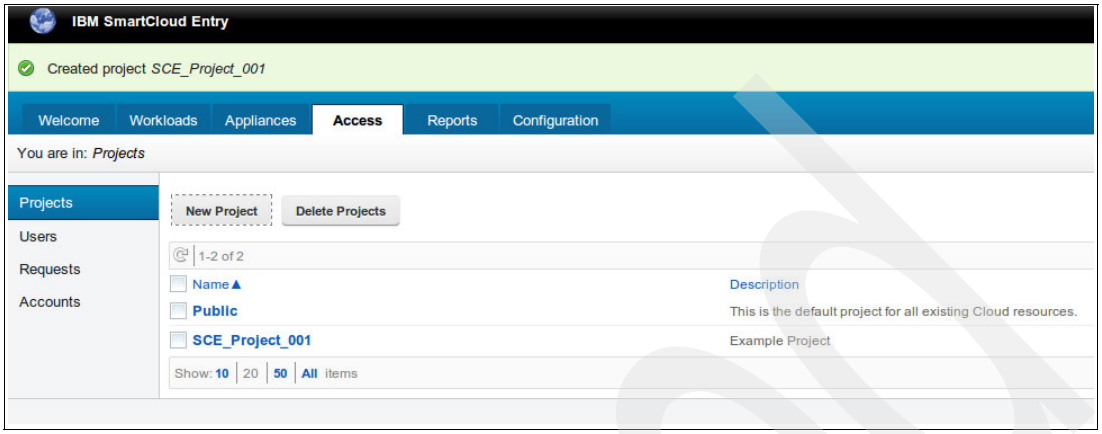

Figure 5-42 New project created

5.10 Adding users to projects

To add a user to a project, complete the following steps:

1. Log in as SmartCloud Entry Administrator.
2. As shown in Figure 5-43, select **Access Tab** → **Projects** and then select the **SCE_Project_001** project name (the project created in 5.9, “Creating projects” on page 168) to which new users are to be added (or continuing on from Figure 5-42, select the project that new members are to be added to). Then, expand the Project Members section (if not already expanded) and click **Add**.

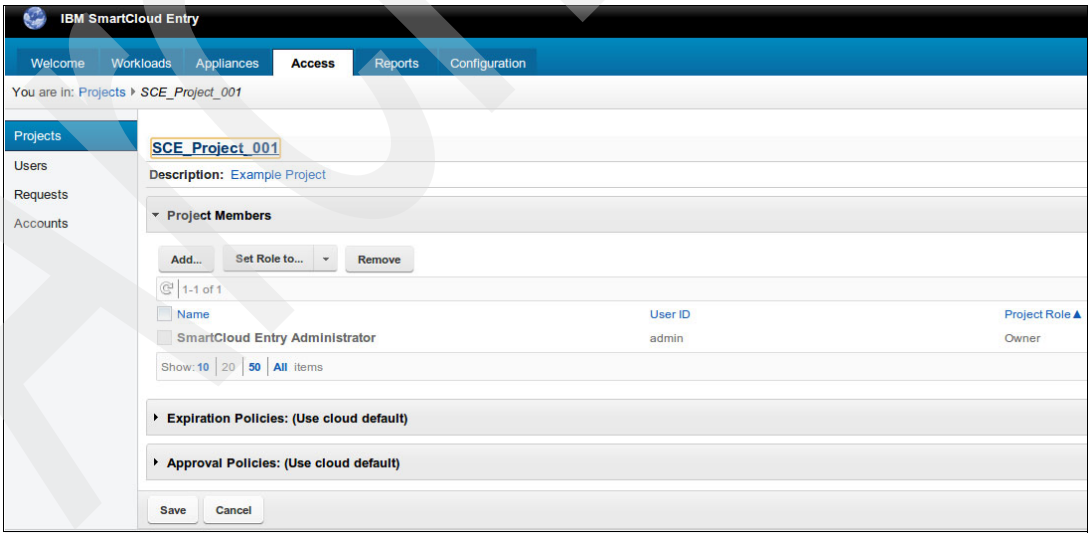

Figure 5-43 Select project

3. The Add Project Members - SCE_Project_001 panel is displayed (Figure 5-44). On this panel, in the Project role, select **User** from the drop-down list and select one or more names of the users to be added as Project Members from the prepopulated list. Click **OK**.

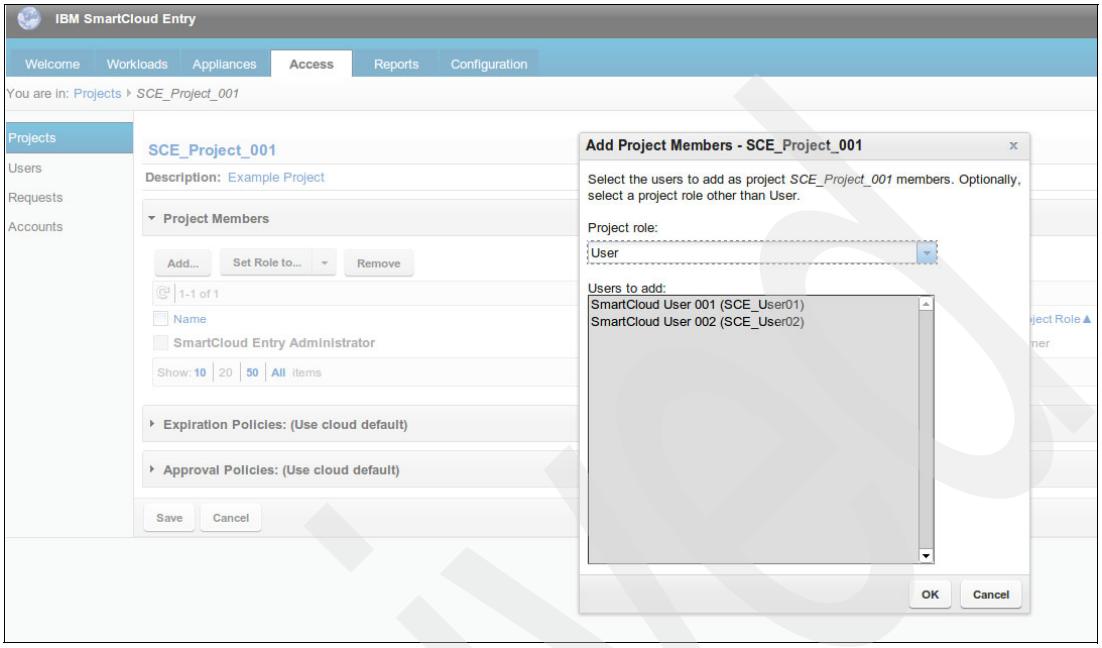

Figure 5-44 Add Project Members pop-up

The newly added member (or members) of the project is listed in the Project Members section of the project. Review the information. Note the message `Project members updated successfully for project SCE_Project_001`, near the top of the Access panel (Figure 5-45).

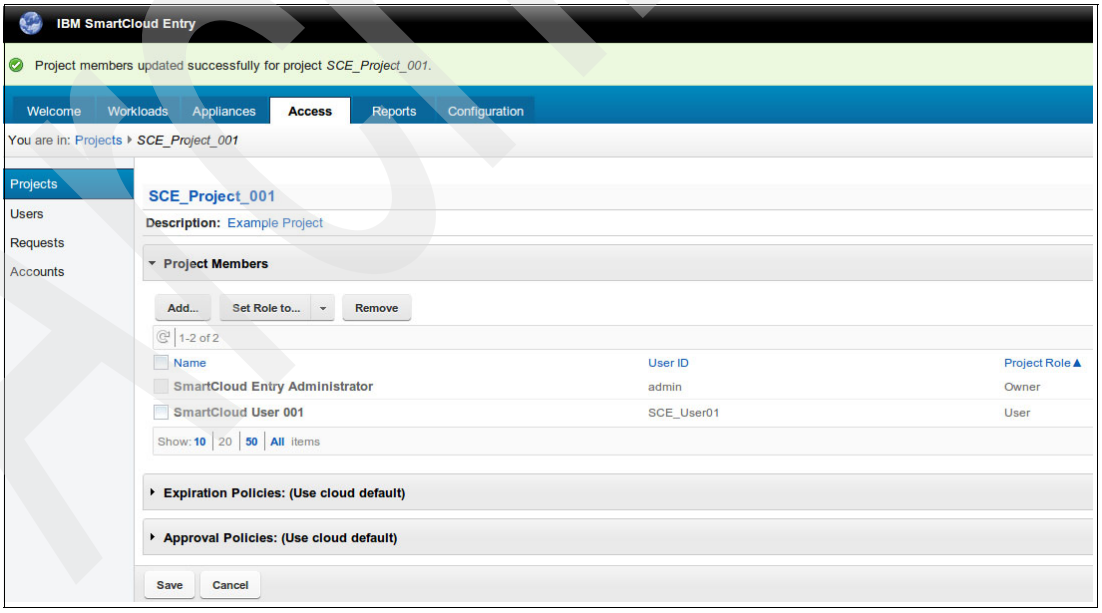

Figure 5-45 Newly added members

5.11 Adding members to accounts

To add users as project members, complete the following steps:

1. Log in as SmartCloud Entry Administrator
2. Select the **Access** tab, click **Accounts**, and select **SCE_Flex_Account (**Figure 5-46).

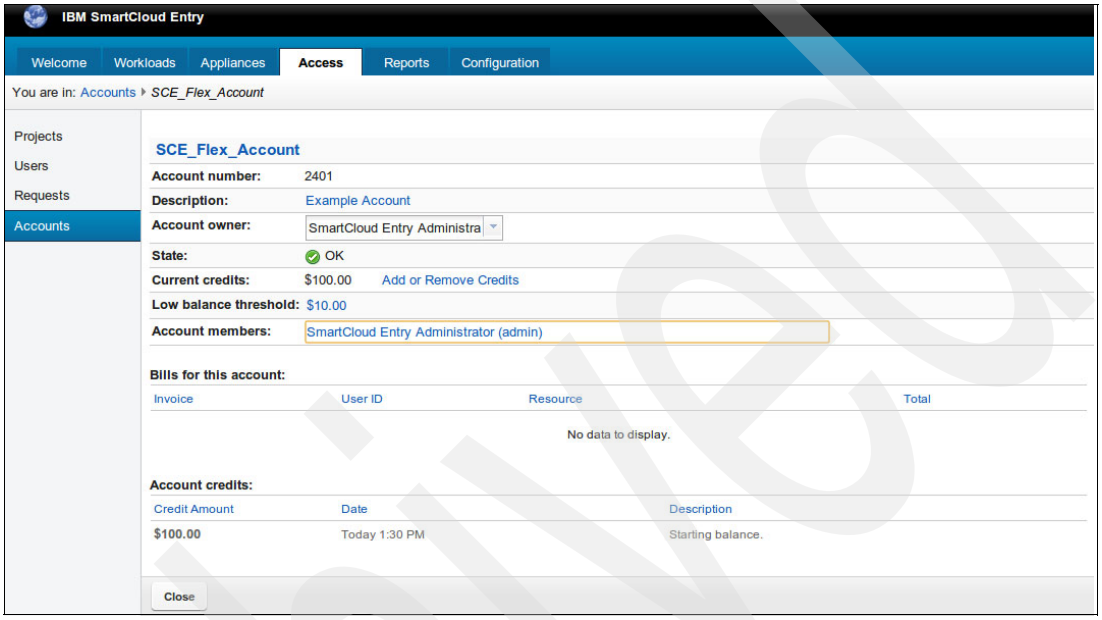

Figure 5-46 Selected SCE_Flex_Account

3. To add users to the Account Members list, click in the **Account members** field, which will then display Add and Remove buttons with a list of Available users to the left and current Account Members listed to the right (Figure 5-47).

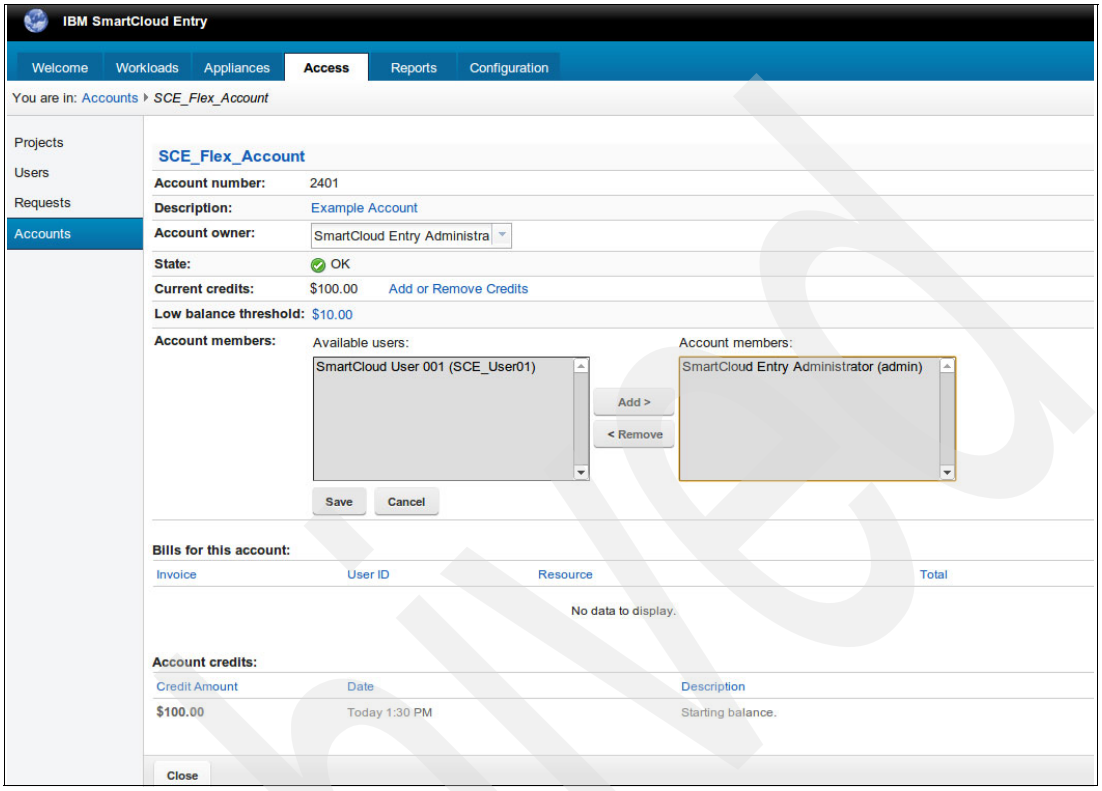

Figure 5-47 Display available users and account members

4. To add an available user, click the user name from the list of available users, and then click **Add**. The selected name moves into the list of account members (Figure 5-48). To remove a user from the list of account members, simply select a user from that list and click **Remove**. The selected name moves back to the Available users.

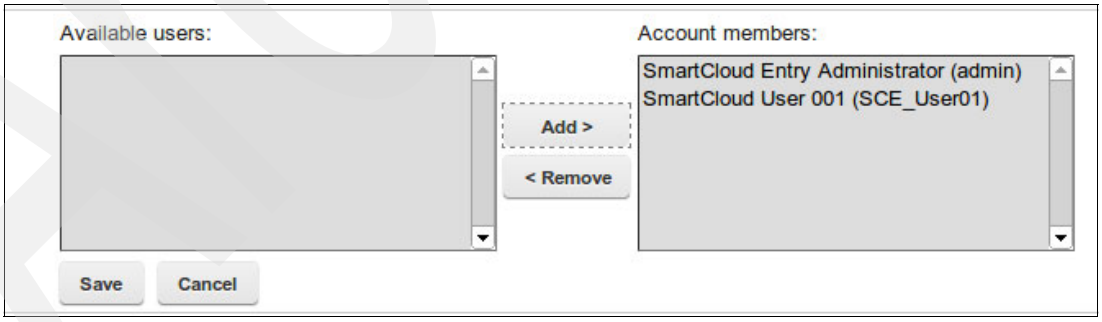

Figure 5-48 Add user to Account members list

5. When all of the appropriate users are added to (or removed from) the Account members list, click **Save**. The account summary view is displayed (Figure 5-49); it includes the resulting account members list.

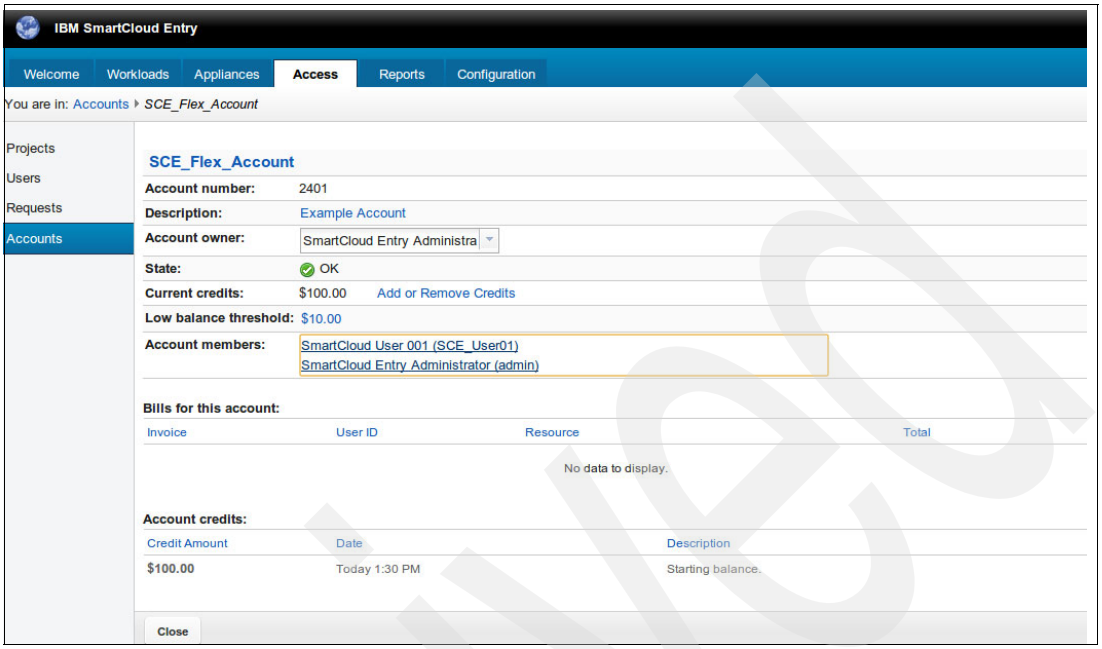

Figure 5-49 Account summary view

5.12 Configure approvals

Approvals allow a SmartCloud Entry Administrator to control how the services in a cloud are used. To enable approvals for a project, complete the following steps:

1. Select the project in the Access tab of the SmartCloud Entry user interface (Figure 5-50).

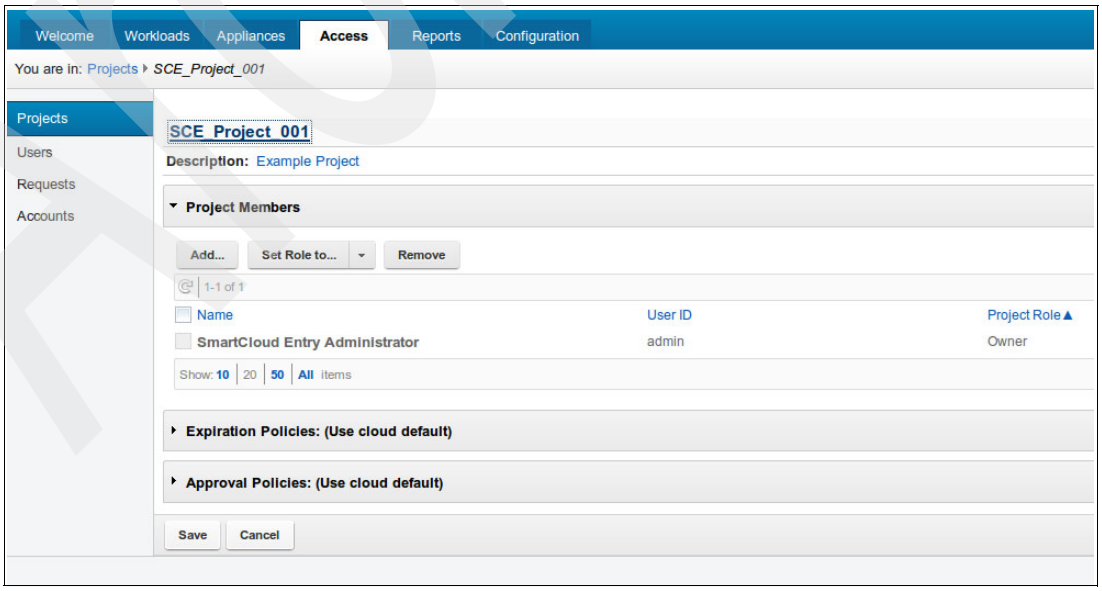

Figure 5-50 Select a project

2. After the panel loads, expand the Approval Policies widget. Notice that the project will use the default approval policies associated with the cloud.

   As the SmartCloud Entry administrator, you can override the default cloud approval settings by using the check boxes. To do this, select **Customize Settings**.

   For this example, all events require approvals, as selected in Figure 5-51. After making the selections, click **Save** to finalize the changes to the project.

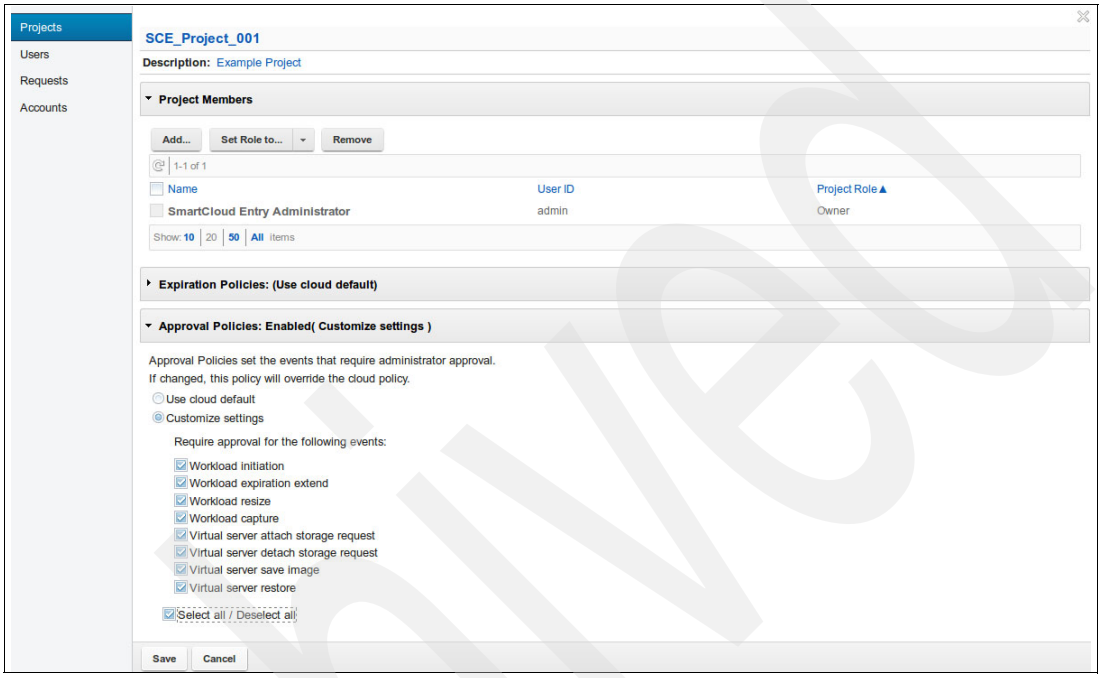

Figure 5-51 Customizing settings

The policies have been updated (Figure 5-52).

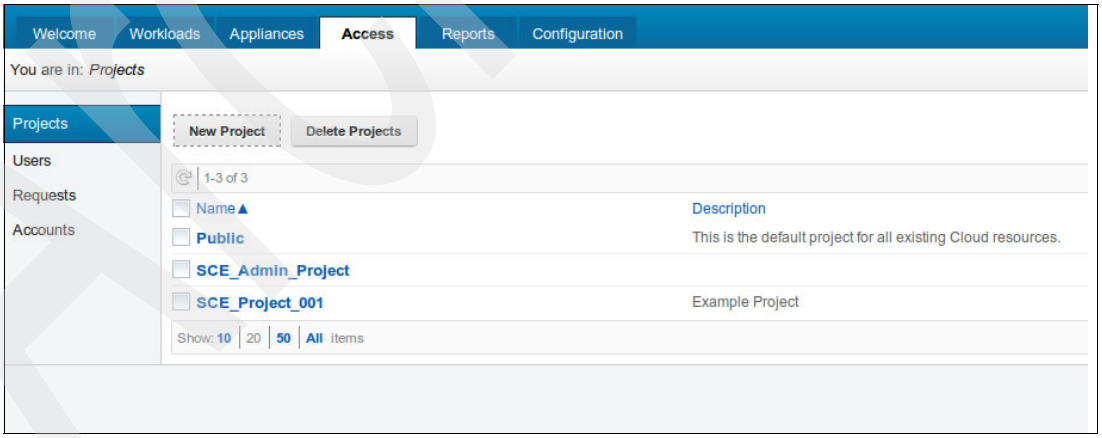

Figure 5-52 List of stored projects

3. Now that the project has the appropriate policies attached, an appliance can be assigned to the project. In the Appliances tab (Figure 5-53), select the appliance to be added to the project with the new approval policies.

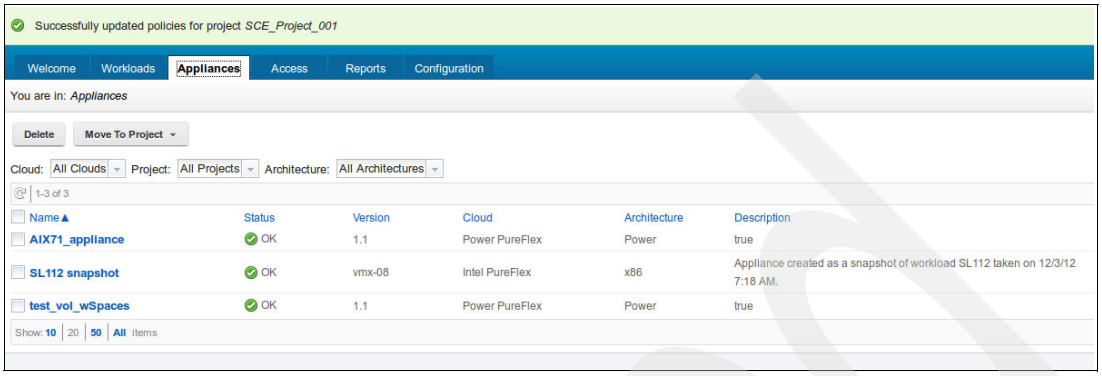

Figure 5-53 Viewing list of appliances on appliance tab

For instance, in Figure 5-54, the AIX71_appliance is moved to project SCE_Project_001. Recall (from 5.8, “Creating users” on page 165) that SCE_Project_001 has two users assigned to it: SCE_User01 and SCE_User02.

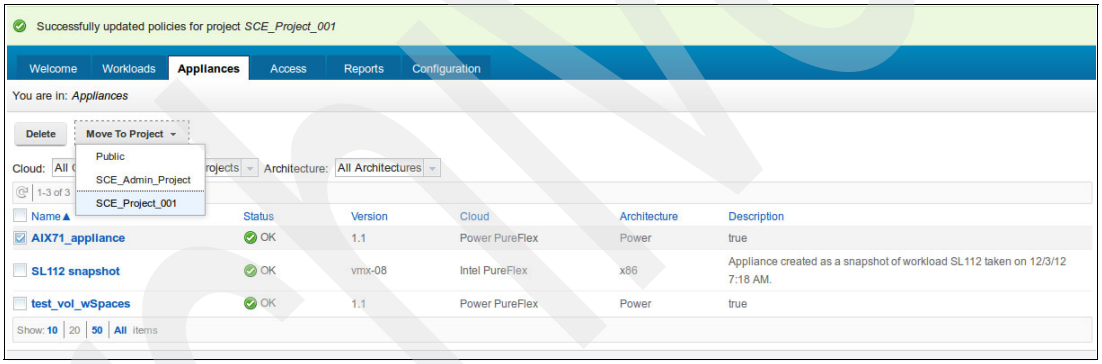

Figure 5-54 Choosing Move Appliances to Project

To show the request and approval mechanisms, Figure 5-55 shows user SCE_User01's attempt to request a workload. After *user one* logs in, the main welcome page opens.

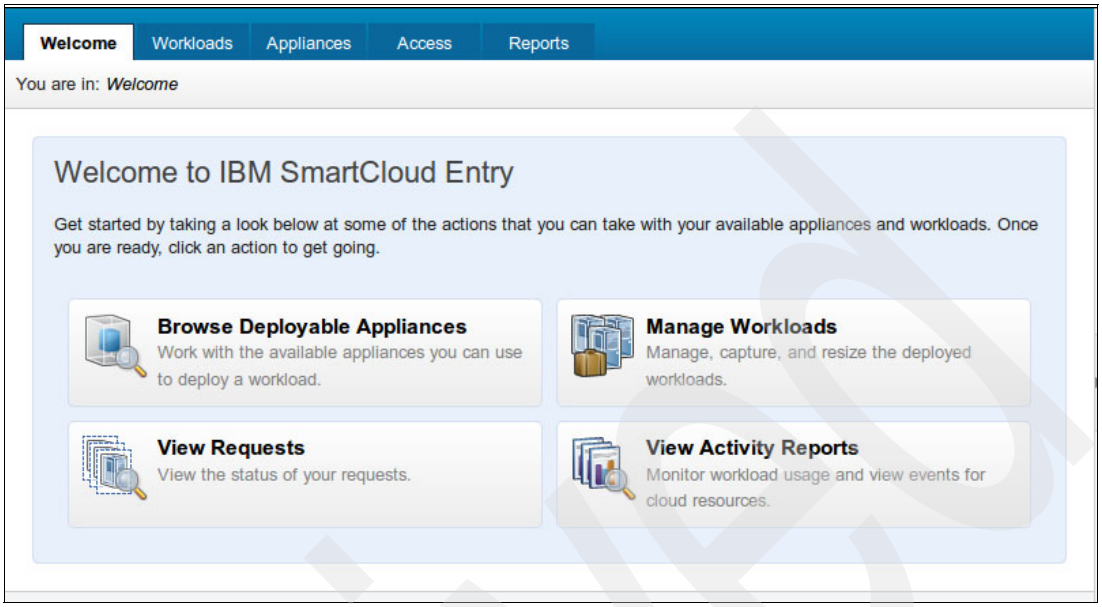

Figure 5-55 Welcome to IBM SmartCloud Entry

4. This window might differ from previous windows in this chapter because SCE_User01 has the role of a regular user rather than an administrator. Again select the Appliances tab (Figure 5-56).

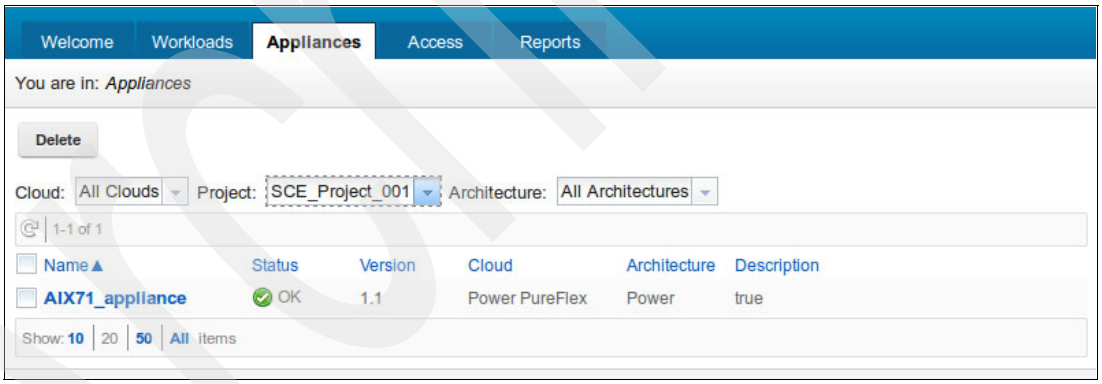

Figure 5-56 Appliance tab

The appliance deployment can be requested. Figure 5-57 shows the requested settings for a new appliance deployment named `AIX Deploy 003`.

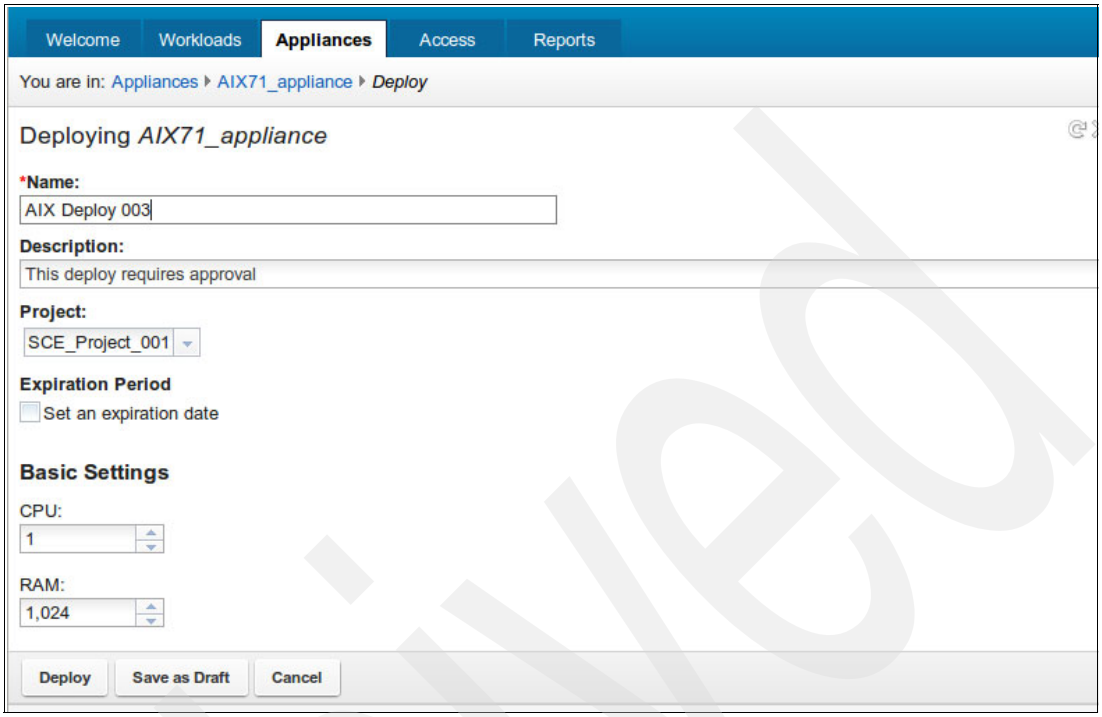

Figure 5-57 Appliance tab: Deploy Appliance

Figure 5-58 shows the notification that the request is being held for approval.

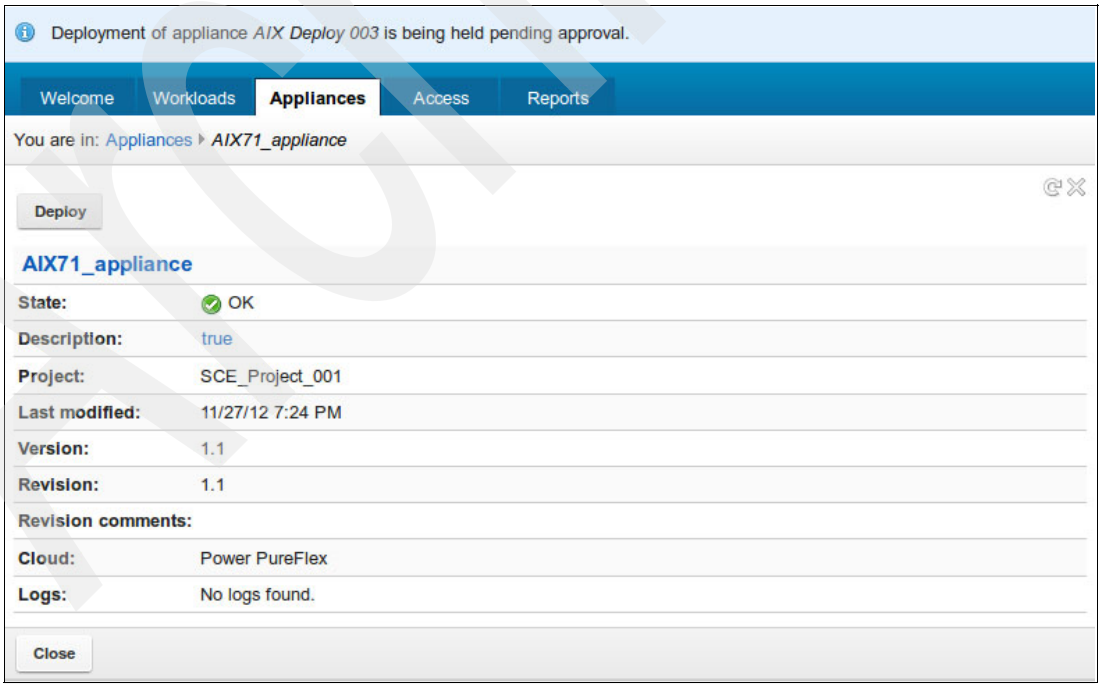

Figure 5-58 Appliance tab

The SmartCloud Entry Administrator will be notified that there is a pending request. Figure 5-59 shows the email that was sent for the `AIX Deploy 003` request.

```
SmartCloud Entry Administrator,

The SmartCloud Entry request has been modified. Request details:

Request number: "3901"
Name: "Workload initiation"
Description: ""
State: "Pending"
Comments:


Thank you.

**This is a machine generated message, please do not reply.
```

Figure 5-59 Generated email example: Notifies SmartCloud Entry Administrator

The SmartCloud Entry Administrator can then log into the user interface and go the Request Section of the Access tab. Figure 5-60 shows SCE_User01's request (3901) for the `AIX Deploy 003` resource. The administrator can approve or reject the request.

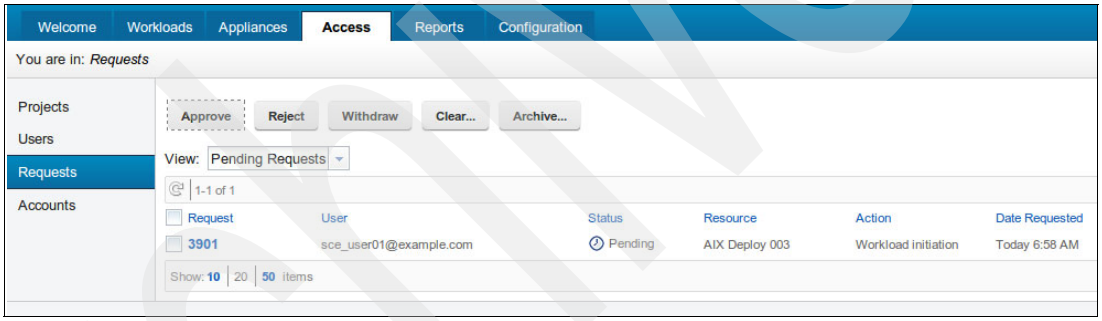

Figure 5-60 Access Tab: Approve button

In Figure 5-61, the administrator approves the request.

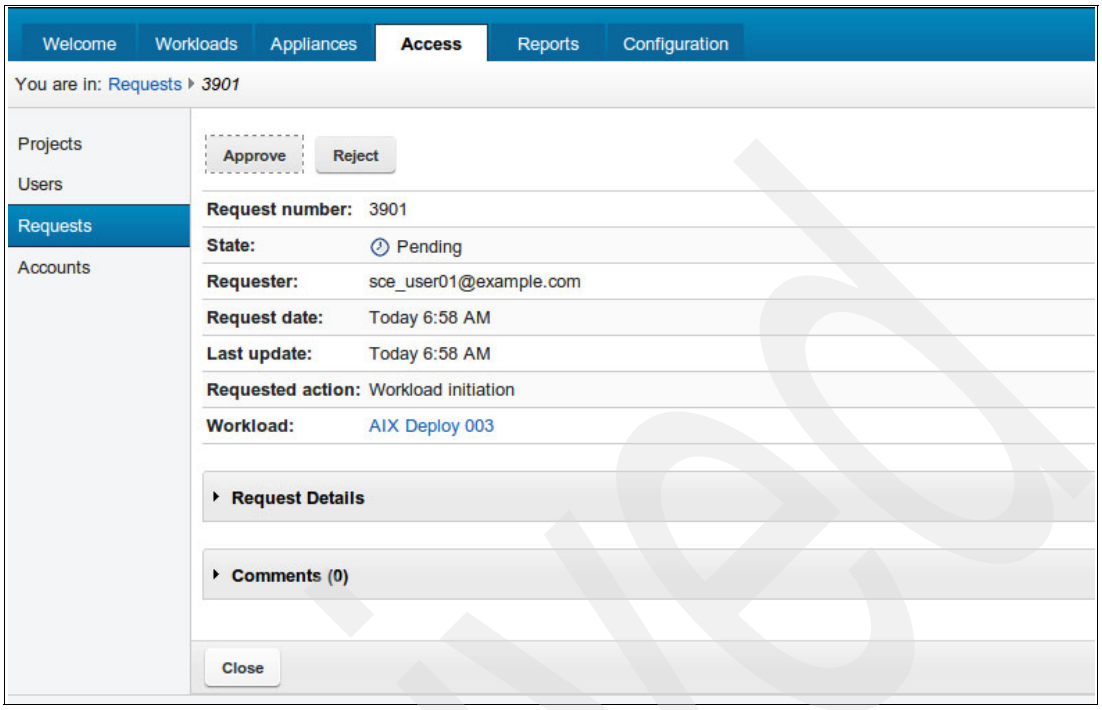

Figure 5-61 A request ready to be approved

The administrator is notified that the request approval was successful (Figure 5-62).

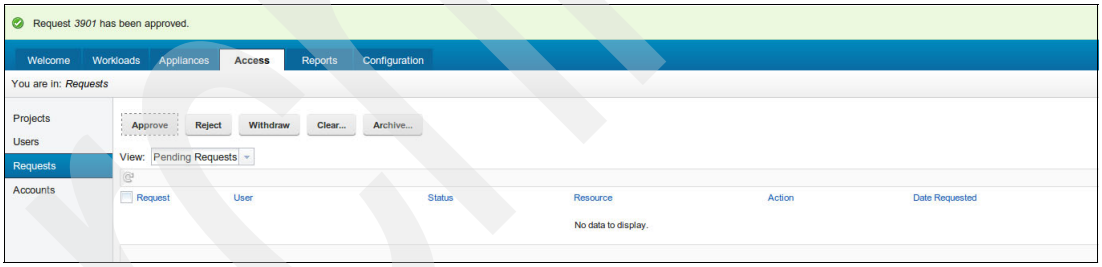

Figure 5-62 Request has been approved

SCE_User01 is then sent an email to indicate the request was approved (Figure 5-63).

```
SCE_User01,

The SmartCloud Entry request has been modified. Request details:

Request number: "3901"
Name: "Workload initiation"
Description: ""
State: "Approved"
Comments:
[admin|Wednesday, December 5, 2012] Approved request.


Thank you.

**This is a machine generated message, please do not reply.
```

Figure 5-63 User received mail when request has been approved.

SmartCloud Entry begins the deployment after the administrator approves the request. When the deployment is successful, the user is notified and given the important information about the requested appliance (Figure 5-64).

```
SCE_User01,

Your deployment "AIX Deploy 003" has been successfully provisioned on the Cloud.

Here are your new VM details:

Name: 129-40-21-197
IP: [129.40.21.197, 192.168.1.4]
root password: [Unknown. Please contact your Cloud administrator]

Please log into the SmartCloud Entry Web interface to view additional details.

Note: It is recommended that you change your root password.

You will need to update SmartCloud Entry with the new root password in order to take snapshots of your deployment.

Thank you.

**This is a machine generated message, please do not reply.
```

Figure 5-64 User receives email when deployment has been successfully provisioned on the Cloud

The other events that require approval have a similar flow.

5.13 Configure metering

By default, IBM SmartCloud Entry metering is disabled. To enable metering, edit the `~/.SCE24/metering.properties` file as follows:

1. Look for the following line:

   `com.ibm.cfs.metering.enabled=false`

2. Change the value from false to true:

   `com.ibm.cfs.metering.enabled=true`

After the changes are made, IBM SmartCloud Entry must be restarted. For an example of a usage metering report, see Figure 3-3 on page 16.

6

IBM SmartCloud Entry usage

This chapter covers usage of the cloud system that is configured in the previous chapter. Two main topics are covered for both the Power nodes and x86 nodes.

- How to deploy appliances. These actions allow users to create workload based on the appliances create by the administrators.
- How to capture deployed workloads into new appliances. This allows users to back up their active workloads. It also provides administrators a standard interface for creating appliances for both platforms.

6.1 Deploying appliances

This section tells how to deploy appliances on POWER and x86 to create workloads.

6.1.1 Deploying a Power appliance

The discovered appliances that IBM SmartCloud Entry lists are ready to deploy. However, in the following example, the appliance must be slightly modified because the networking configuration that IBM SmartCloud Entry will use for the appliance as it is deployed might not be correct. Recall in 4.3, “Networking setup” on page 27 that there is a complex network inside the IBM Flex Chassis. Also, recall that two VLANs are used (4091 and 1).

In 5.3, “Configuring cloud network pools” on page 151, two network pools were created to correspond to the two VLANs.

Again, log in to the user interface. In the main panel, select the Appliances tab. Several appliances are listed (Figure 6-1).

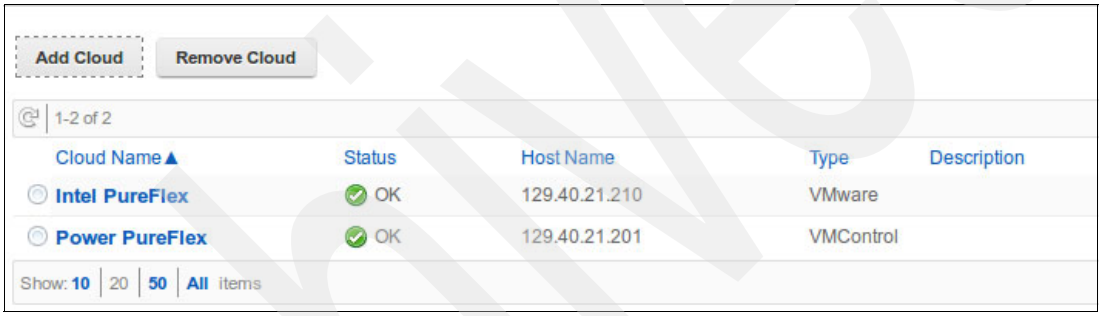

Figure 6-1 List of appliances ready to be deployed

Select the AIX71_appliance link. This appliance corresponds to the AIX image that was created in 4.6.4, “Install AIX” on page 77. Click **Configure** (Figure 6-2).

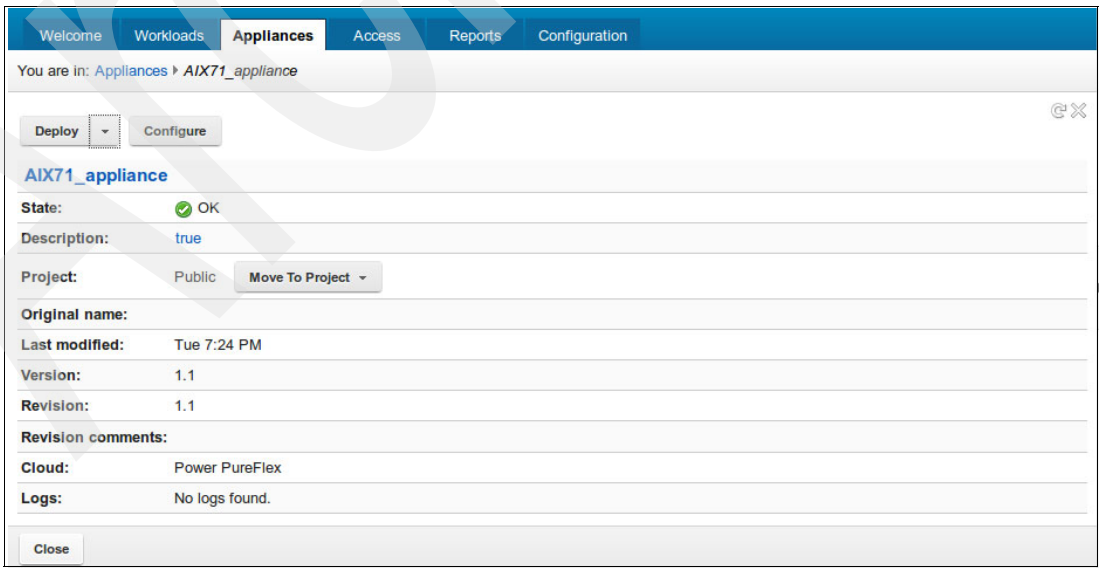

Figure 6-2 An appliance is selected to be configured

The configuring appliance panel opens. For ease of reading, this panel is spread across several figures (Figure 6-3 through Figure 6-6 on page 184).

Figure 6-3 shows the overall settings of the appliance. It includes shared virtual processors.

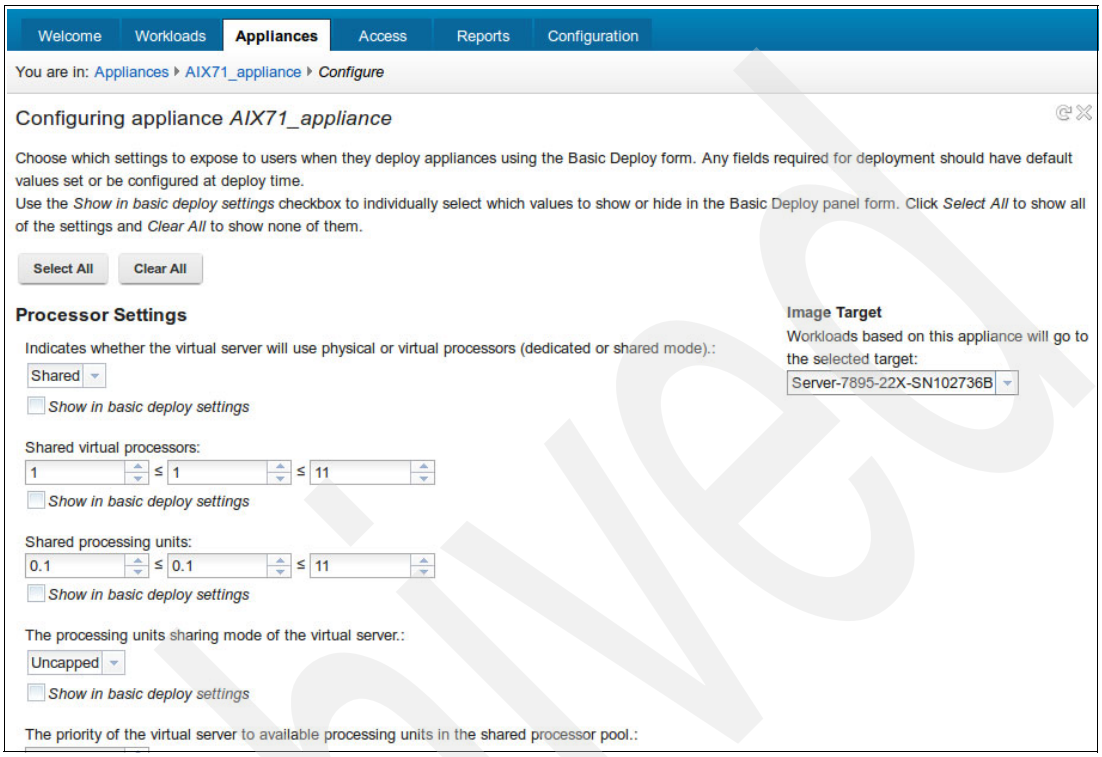

Figure 6-3 Overall settings of the appliance

Scroll down to see the network settings (Figure 6-4). This information is pre-populated from the information set in 5.3, “Configuring cloud network pools” on page 151.

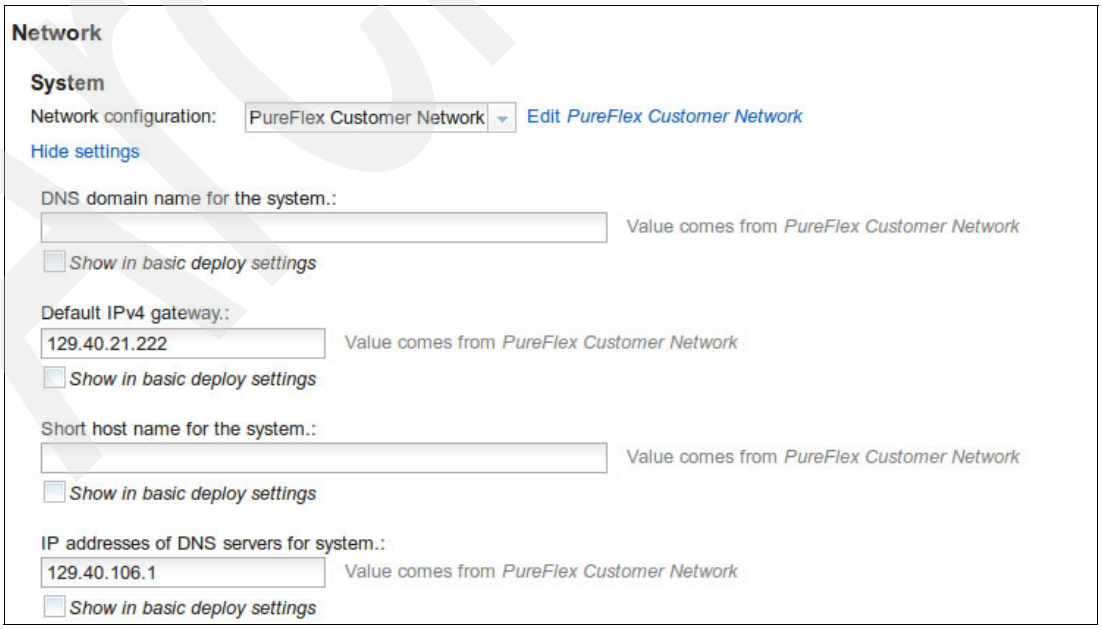

Figure 6-4 Network settings of the appliance

The first of the two adapters is shown in Figure 6-5. It is important to select the correct Network ID for this adapter. These adapters represent the setup configured in 4.3.1, “Chassis and compute node setup” on page 27. In Figure 6-5, the Network ID for VLAN 4091 is selected.

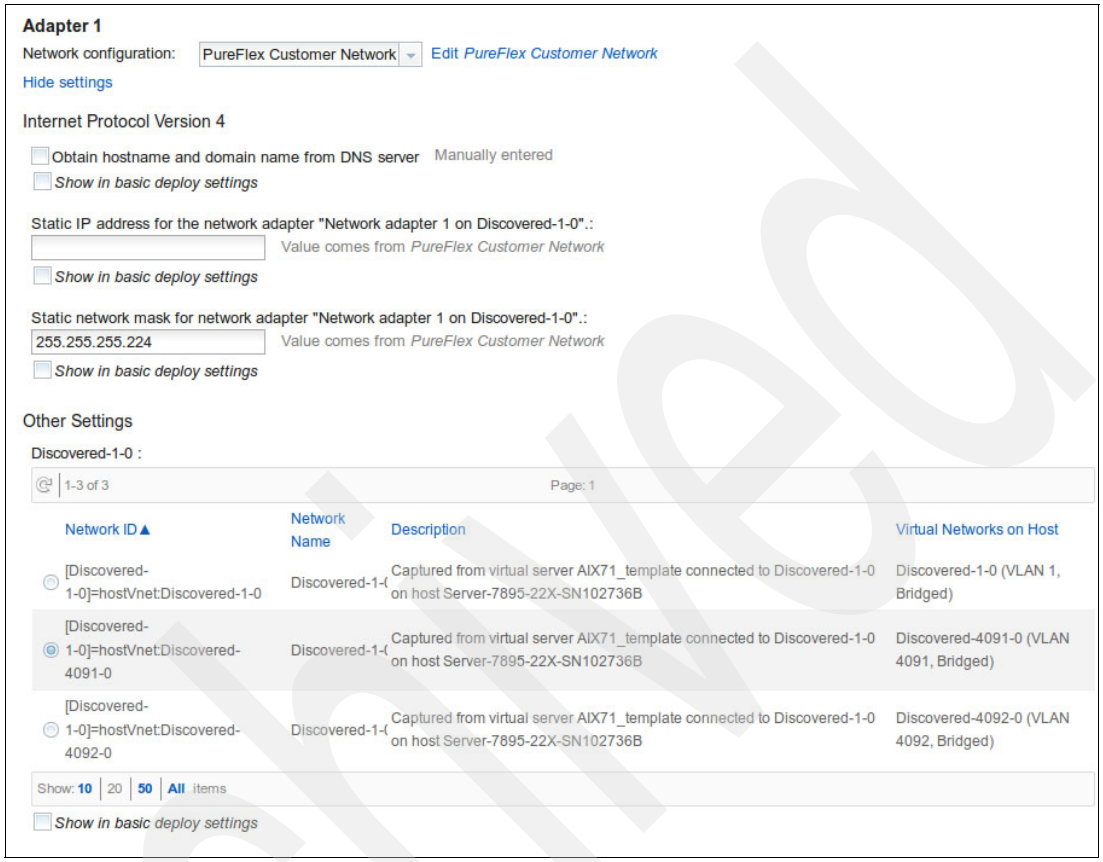

Figure 6-5 Adapter Specifications

Figure 6-6 shows the second adapter. This is the adapter for VLAN 1.

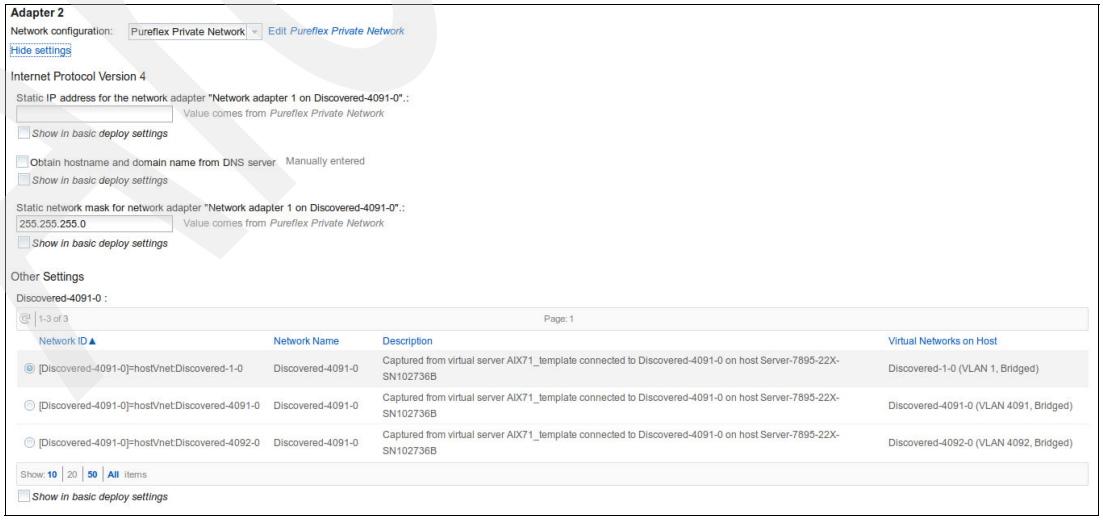

Figure 6-6 Second adapter specification

Save the appliance. You may then deploy it by clicking **Deploy**. Set a name for this deployment (Figure 6-7).

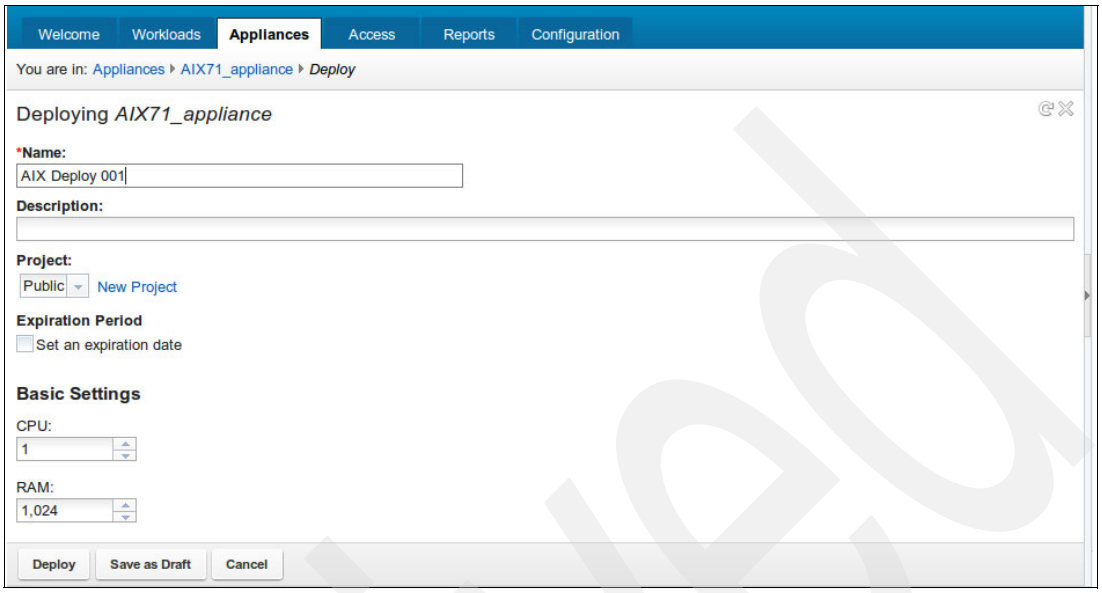

Figure 6-7 Providing a name for the deployment

The request for the deployment is sent as a new workload (Figure 6-8). The new workload has the same settings as configured in the previous step.

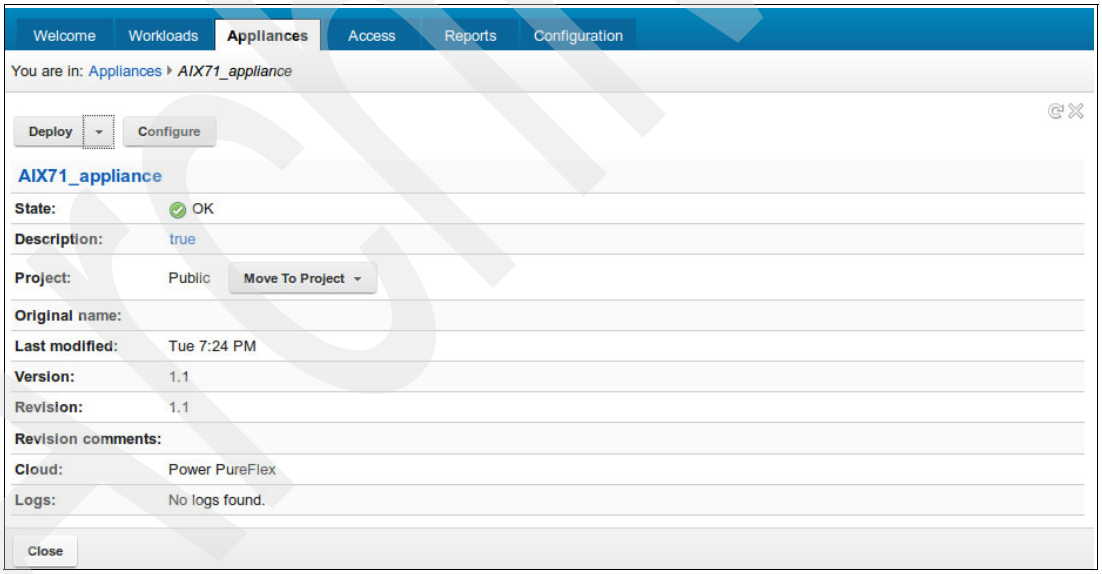

Figure 6-8 Request for deployment sent as a new workload

The Workloads tab shows the workload as being deployed (Figure 6-9).

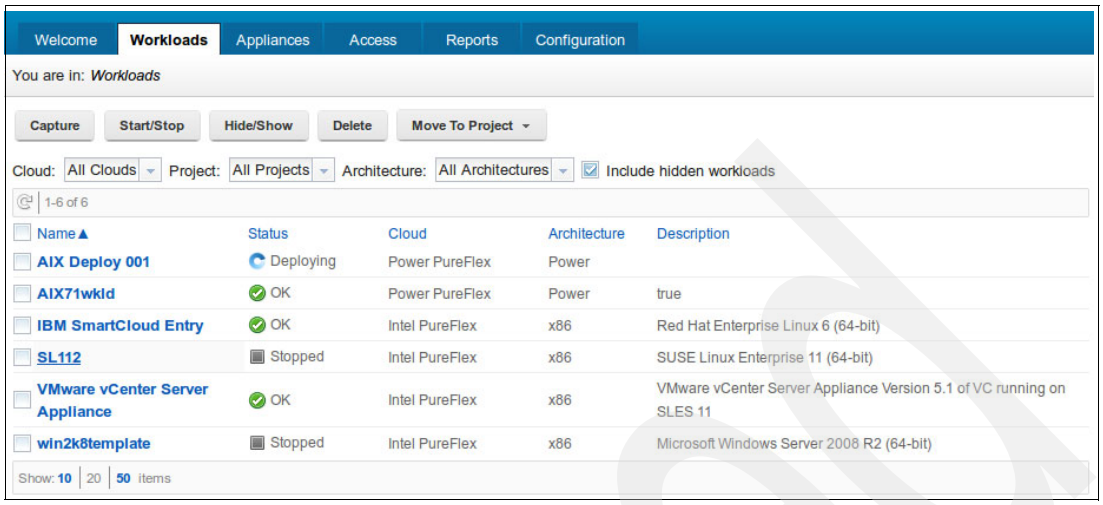

Figure 6-9 Workloads Status

When the workload is finished, an email notifies (Figure 6-10) the appropriate user that the workload was deployed. This email also contains the information to log in to the newly created operating system.

```
SmartCloud Entry Administrator,

Your deployment "SmartCloud Entry Administrator's AIX71_appliance on 12/5/12 1:48 PM" has been successfully provisioned on the
Cloud.

Here are your new VM details:

Name: 129-40-21-198
IP: [129.40.21.198, 192.168.1.5]
root password: [Unknown. Please contact your Cloud administrator]

Please log into the SmartCloud Entry Web interface to view additional details.

Note: It is recommended that you change your root password.

You will need to update SmartCloud Entry with the new root password in order to take snapshots of your deployment.

Thank you.

**This is a machine generated message, please do not reply.
```

Figure 6-10 Workload completed email

Using the information provided in the email shown in Figure 6-10 on page 186, the user requesting the workload can log in, as shown in Example 6-1.

Example 6-1 Logging in

```
#:~$ ssh root@129.40.21.198
The authenticity of host '129.40.21.198 (129.40.21.198)' can't be established.
RSA key fingerprint is 12:b3:0b:2c:2e:4c:38:ba:48:10:4c:4b:5b:41:97:4c.
Are you sure you want to continue connecting (yes/no)? yes
Warning: Permanently added '129.40.21.198' (RSA) to the list of known hosts.
root@129.40.21.198's password:
X11 forwarding request failed on channel 0
Last login: Tue Nov 27 18:43:01 CDT 2012 on ssh from f3enfsm.pbm.ihost.com
*******************************************************************************
*                                                                             *
*                                                                             *
* Welcome to AIX Version 7.1!                                                 *
*                                                                             *
*                                                                             *
* Please see the README file in /usr/lpp/bos for information pertinent to     *
* this release of the AIX Operating System.                                   *
*                                                                             *
*                                                                             *
*******************************************************************************
# ifconfig -a
en0:
flags=1e080863,480<UP,BROADCAST,NOTRAILERS,RUNNING,SIMPLEX,MULTICAST,GROUPRT,64BIT
,CHECKSUM_OFFLOAD(ACTIVE),CHAIN>
inet 129.40.21.198 netmask 0xfffffe0 broadcast 129.40.21.223
tcp_sendspace 262144 tcp_recvspace 262144 rfc1323 1
en1:
flags=1e080863,480<UP,BROADCAST,NOTRAILERS,RUNNING,SIMPLEX,MULTICAST,GROUPRT,64BIT
,CHECKSUM_OFFLOAD(ACTIVE),CHAIN>
inet 192.168.1.5 netmask 0xffffff00 broadcast 192.168.1.255
tcp_sendspace 262144 tcp_recvspace 262144 rfc1323 1
lo0:
flags=e08084b,c0<UP,BROADCAST,LOOPBACK,RUNNING,SIMPLEX,MULTICAST,GROUPRT,64BIT,LAR
GESEND,CHAIN>
inet 127.0.0.1 netmask 0xff000000 broadcast 127.255.255.255
inet6 ::1%1/0
tcp_sendspace 131072 tcp_recvspace 131072 rfc1323 1
#
```

6.1.2 Deploying an x86 appliance

VMware appliances must be configured in a similar manner as in the previous section. The first panel differs slightly from that section to show that VMware requires different information to successfully be deployed (Figure 6-11).

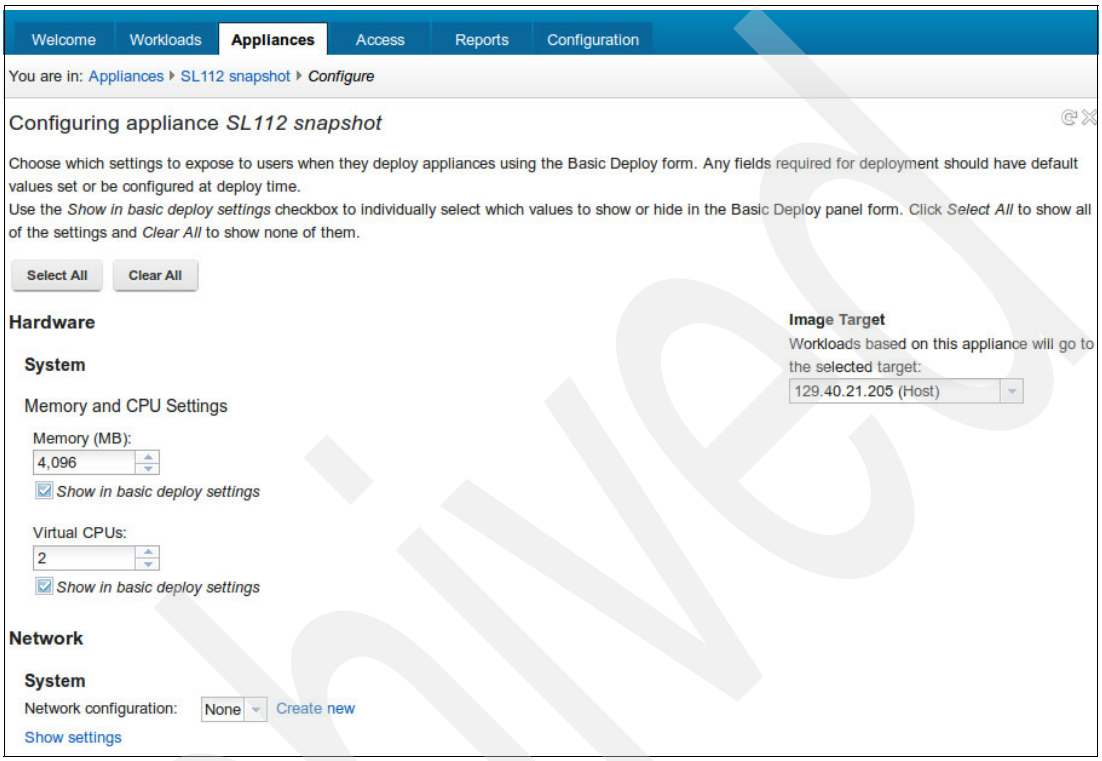

Figure 6-11 Configuration using VMware

The network setup is not as complex, because the VMware switch on the ESXi already defines the network ID.

Click **Save**. The appliance is saved (Figure 6-12).

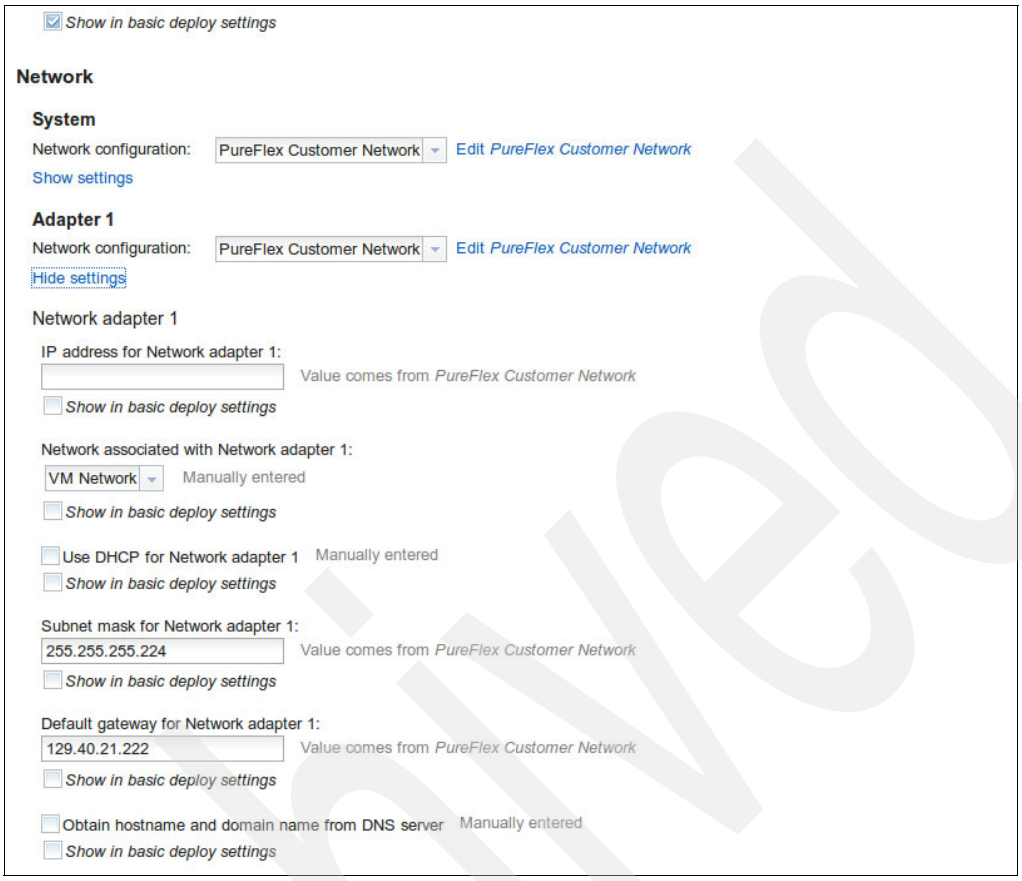

Figure 6-12 Saved appliance

The saved appliance can be deployed (Figure 6-13).

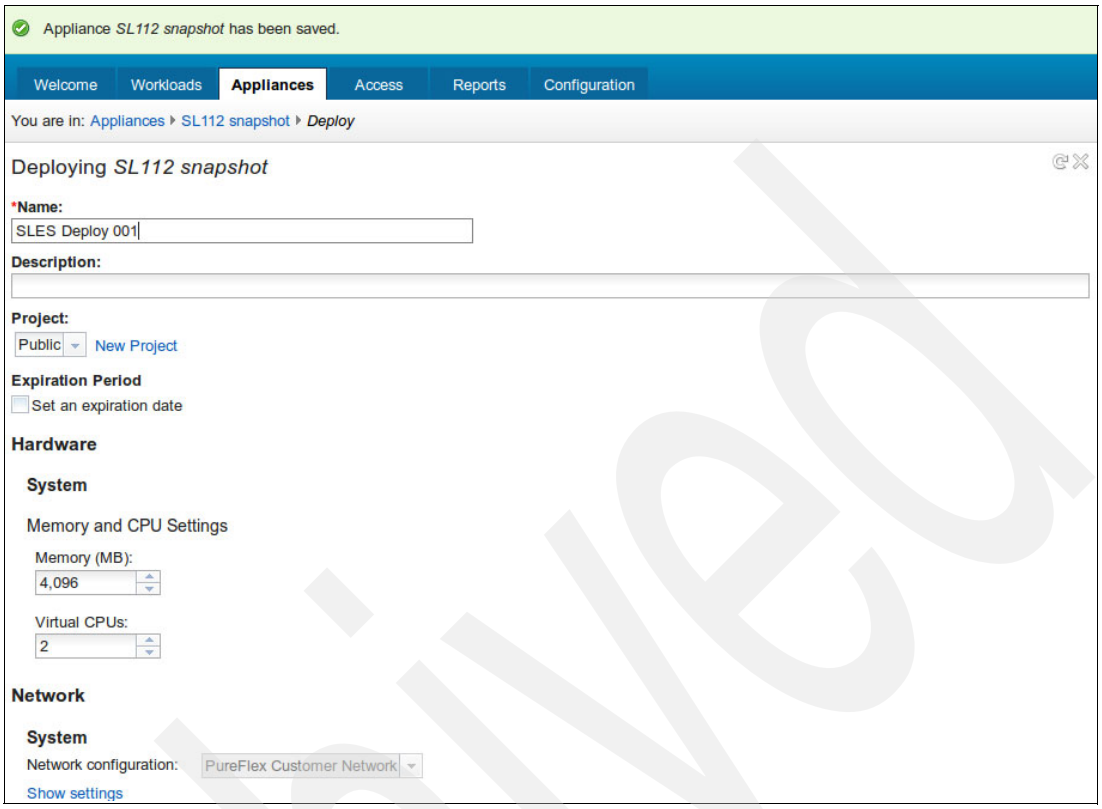

Figure 6-13 Deployment snapshot

A notification is sent for the x86 appliance (Figure 6-14).

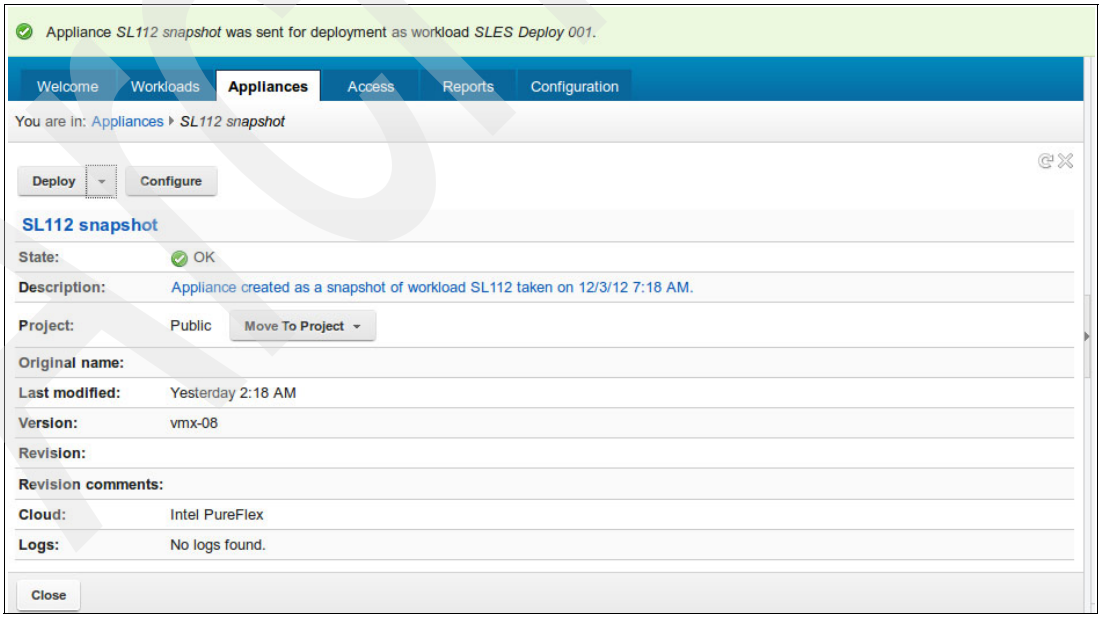

Figure 6-14 Deployment snapshot has been sent for deployment as workload

6.2 Capturing a workload

In this section, workloads are captured on POWER and on x86.

6.2.1 Capturing a Power workload

Capturing a Power workload in SmartCloud Entry is nearly identical to capturing a workload in VMControl on the FSM (covered in 4.6.7, “Capture virtual server” on page 89). The first step is to shut down the workload so that it is eligible for capture (Figure 6-15).

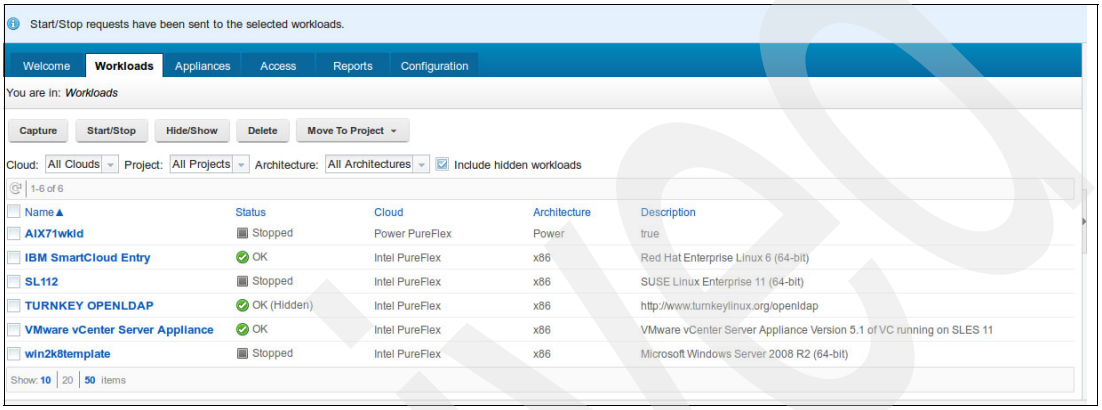

Figure 6-15 Shut down workload so it is eligible for capture

SmartCloud Entry sets an on-screen notification that the workload has been stopped. Click **Capture**. Figure 6-16 shows the notification as the request is sent to the cloud.

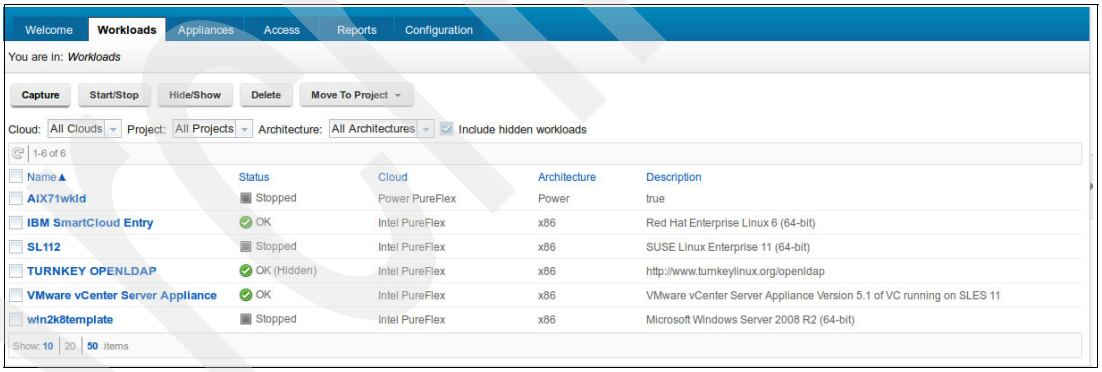

Figure 6-16 Workload has been stopped

An email is sent to the requesting user when the capture is completed (Figure 6-17).

```
SmartCloud Entry Administrator,

Capture of your deployment "AIX71wkld" has completed successfully.

Please log into the SmartCloud Entry Web interface to view additional details.

Thank you.

**This is a machine generated message, please do not reply.
```

Figure 6-17 Confirmation email that deployment has completed successfully

The IBM SmartCloud Entry Appliances tab is updated with the new appliance (Figure 6-18).

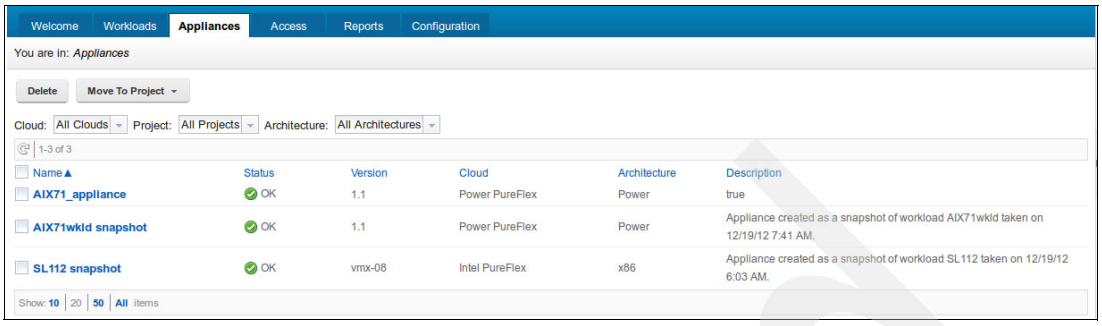

Figure 6-18 Appliance tab is updated with new appliance information and status

6.2.2 Capturing an x86 Workload

To capture an x86 workload, its status must appear as stopped (Figure 6-19).

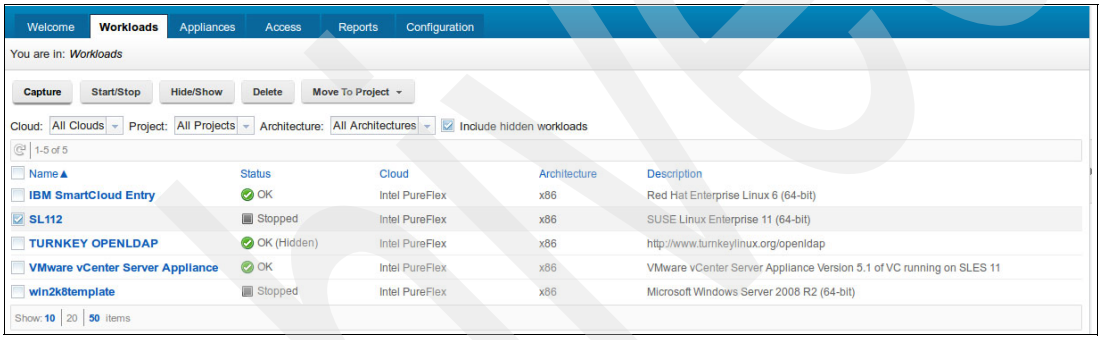

Figure 6-19 On screen notification that a workload has stopped

IBM SmartCloud Entry issues an on-screen notification that the workload was stopped. Now you can click **Capture**. A notification that is displayed (Figure 6-20) as the request is sent to the cloud as a result. An information icon is displayed in the upper left of window.

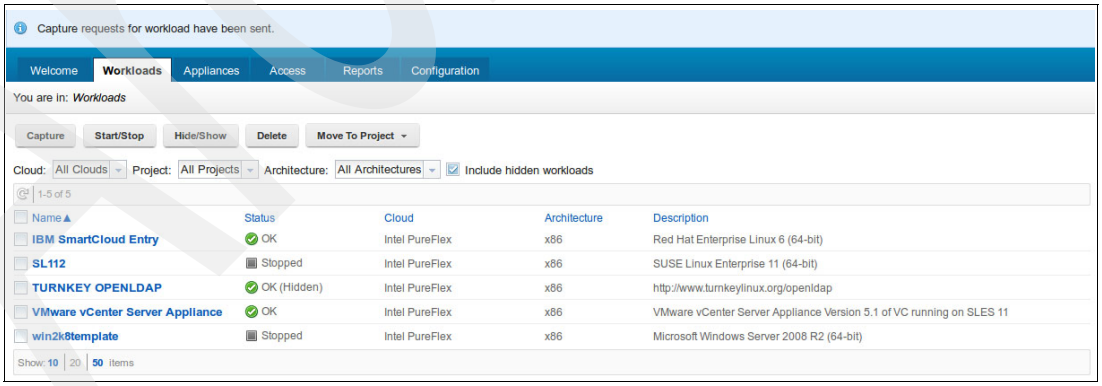

Figure 6-20 Information icon and message: Capture request has been sent

An email is sent to the requesting user when the capture is completed (Figure 6-21).

```
SmartCloud Entry Administrator,

Capture of your deployment "SL112" has completed successfully.

Please log into the SmartCloud Entry Web interface to view additional details.

Thank you.

**This is a machine generated message, please do not reply.
```

Figure 6-21 An email notification: capture of deployment has completed successfully

The IBM SmartCloud Entry Appliances tab is updated with the new appliance (Figure 6-22).

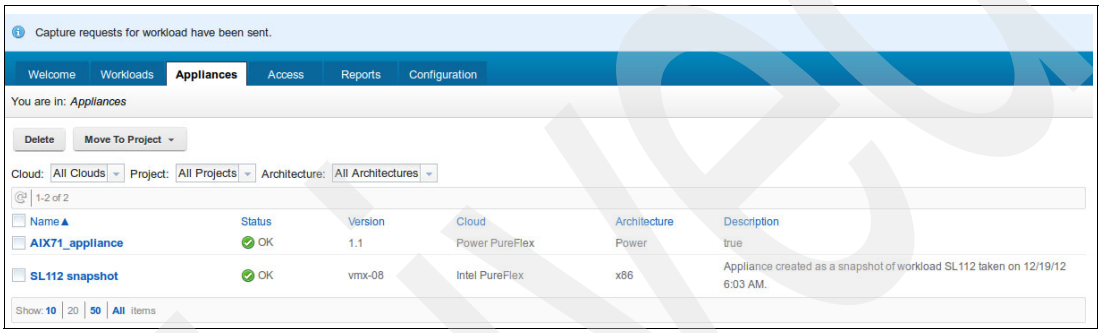

Figure 6-22 Appliance tab is updated with new appliance information

A

Known problems and workarounds

Various problems were encountered while researching and completing some of the tasks described in this document. This appendix describes them and offers workaround solutions for each problem encountered:

- DNZVMP016E: Internal application error

  This error occurs when an LPAR is missing a default profile or a current profile. This causes some VMControl operations to fail and prevents the removal of the LPAR itself.

- FlashCopy not used during capture of a virtual appliance

  This problem occurs during capture of a virtual appliance. If the OS of the virtual server being captured is installed on a volume on the V7000, and that volume has spaces in its name, FlashCopy will not be used to copy the volume.

- Discovering cloned operating systems

  If two operating systems are clones of one another, the FSM will not be able to differentiate between them. When discovering two operating systems that are clones, the FSM will correctly discover the first operating system, but the second will be seen as the same operating system as the first.

A.1 DNZVMP016E: Internal application error

This error occurs when an LPAR is missing a default profile or a current profile. Some VMControl operations can fail, preventing the removal of the LPAR itself. To determine if a particular LPAR has these profiles defined, issue the following command from the FSM command line:

```
lssyscfg -r prof -m <LPAR name> -Flpar_id:name
```

In Example A-1, LPARs 1 and 2 are missing values for both `default_profile` and `curr_profile`.

Example A-1 Missing values

```
USERID@f3efsm:~> lssyscfg -m Server-7895-22X-SN102736B -r lpar -F
lpar_id,name,default_profile,curr_profile

2,TemplateOS,,

1,7895_SN102736B_VIOS1,,

3,AIX71template,OriginalProfile,OriginalProfile
```

The workaround to this problem is to assign values to the empty variables using the **chsyscfg** command on the FSM's command line. The commands are at the following site:

http://www.ibm.com/support/docview.wss?uid=nas7d16d5ce0f7a5484b86257a3f00603d94

A.2 FlashCopy not used during capture of a virtual appliance

This problem occurs during capture of a virtual appliance. If the OS of the virtual server being captured is installed on a volume on the V7000, and that volume has spaces in its name, FlashCopy will not be used to copy the volume. If the image repository is hosted by a VIO server, the **dd** command will be used by the capture (or copy) operation instead. The capture operation succeeds but takes significantly longer than it does when FlashCopy is used.

The "Capture virtual appliance" job log shows the following message:

```
DNZLOP899I Storage copy services is unable to copy one or more disks using the
FlashCopy function, because FlashCopy is not configured. The server that controls
the image repository will copy the disks.
```

To prevent this problem, do not use spaces when naming volumes on the V7000 (for example, use `AIX_bootdisk` instead of `AIX boot disk`).

A.3 Discovering cloned operating systems

If two operating systems are clones of each other, the FSM is unable to differentiate between them. When discovering two operating system that are clones, the FSM correctly discovers the first. The FSM will create a new object to represent this operating system and its discovery status will be set to new. Discovery of the second operating system succeeds. However, instead of creating a new object to represent the second operating system, the FSM lists the second operating system as “previously discovered” and the IP address of the second operating system will simply be added to the list of IP addresses associated with the first.

Because the operating systems’ file systems are clones, they contain identical UIDs, Tivoli GUIDs, and SSH keys. To correct this issue, new UIDs, GUIDs, and SSH keys must be generated for one of the two operating systems, and the Common Agent Services agent must be restarted.

The information center provides the procedure:

http://pic.dhe.ibm.com/infocenter/director/pubs/topic/com.ibm.director.discovery.helps.doc/fqm0_t_discovering_systems_mirrored_image.html

Related publications

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

IBM Redbooks

The following IBM Redbooks publications provide additional information about the topic in this document. Note that some publications in this list might be available in softcopy only.

- *Choosing IBM Flex System for Your Private Cloud Infrastructure*, REDP-4920
- *Connecting an IBM PureFlex System to the Network*, TIPS0941
- *Creating Smart Virtual Appliances with the IBM Image Construction and Composition Tool*, SG24-8042
- *IBM Flex System Interoperability Guide*, REDP-FSIG
- *IBM Flex System Manager*, TIPS0862
- *IBM Flex System Networking in an Enterprise Data Center*, REDP-4834
- *IBM Flex System p24L, p260 and p460 Compute Nodes,* TIPS0880
- *IBM Flex System p260 and p460 Planning and Implementation Guide*, SG24-7989
- *IBM Flex System V7000 Storage Node Introduction and Implementation Guide,* SG24-8068
- *IBM PureFlex System and IBM Flex System Products and Technology*, SG24-7984
- *IBM PureFlex System Solutions for Managed Service Providers*, REDP-4994
- *IBM SmartCloud Entry 2.4 Deployment Use Cases*, REDP-4908
- *IBM Storwize V7000 Unified Disk System*, TIPS0837
- *IBM System Storage SAN Volume Controller and Storwize V7000 Replication Family Services,* SG24-7574
- *Implementing Systems Management of IBM PureFlex System*, SG24-8060
- *Implementing the IBM Storwize V7000 Unified*, SG24-8010
- *Implementing the IBM Storwize V7000 V6.3*, SG24-7938
- *Implementing the Storwize V7000 and the IBM System Storage SAN32B-E4 Encryption Switch*, SG24-7977
- *Moving to IBM PureFlex System x86-to-x86 Migration*, REDP-4887
- *Overview of IBM PureSystems*, TIPS0892
- *Positioning IBM Flex System 16 Gb Fibre Channel Fabric for Storage-Intensive Enterprise Workloads*, REDP-4921
- *Real-time Compression in SAN Volume Controller and Storwize V7000*, REDP-4859
- *Simple Configuration Example for Storwize V7000 FlashCopy and PowerHA SystemMirror for i*, REDP-4923

You can search for, view, download or order these documents and other Redbooks, Redpapers, Web Docs, drafts, and additional materials at the following website:

ibm.com/redbooks

Online resources

These websites are also relevant as further information sources:

- IBM SmartCloud Entry documentation on developerWorks:

  https://www.ibm.com/developerworks/mydeveloperworks/wikis/home?lang=en#/wiki/W21ed5ba0f4a9_46f4_9626_24cbbb86fbb9/page/Documentation

- IBM Storwize V7000:

  http://pic.dhe.ibm.com/infocenter/storwize/ic/index.jsp

- IBM PureSystems:

  http://www.ibm.com/developerworks/puresystems/index.html

- Workaround information:

  - http://www.ibm.com/support/docview.wss?uid=nas7d16d5ce0f7a5484b86257a3f00603d94
  - http://pic.dhe.ibm.com/infocenter/director/pubs/topic/com.ibm.director.discovery.helps.doc/fqm0_t_discovering_systems_mirrored_image.html

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services

(0.2”spine)
0.17”<->0.473”
90<->249 pages

Implementing IBM SmartCloud Entry on IBM PureFlex System

Introducing IBM PureFlex System and cloud

Installing IBM SmartCloud Entry hardware and software

Configuring and using IBM SmartCloud Entry

Distributed computing has been transformed with the introduction of virtualization technology. This has driven a re-architecture of traditional data center workload placement. In 2012, IBM announced IBM PureSystems, a new offering based on preconfigured software, servers, and storage that form an expert integrated system.

Expert integrated systems now combine traditional IT resources into a single optimized solution, with prepackaged components including servers, storage devices, networking equipment, and software. With this evolution of technology, we move from discrete, siloed, and underutilized IT resources to shared resource pools.

This IBM Redbooks publication can help you install, tailor, and configure IBM SmartCloud Entry on the IBM PureFlex System offering. This book is intended for anyone who wants to learn more about cloud computing with IBM SmartCloud Entry and offerings based on IBM Flex System elements.

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

SG24-8102-00 ISBN 0738438189