WebSphere software

IBM

WebSphere Cloudburst Appliance and PowerVM

- Discover PowerVM WebSphere Cloudburst Environment
- Install and administer WebSphere Cloudburst Appliance
- Learn by example and practical scenaios

Shivaji D Bhosale
Martin Cernicky
Paul Reynolds
Margaret Ticknor

Redbooks

ibm.com/redbooks

International Technical Support Organization

WebSphere Cloudburst Appliance and PowerVM

October 2010

SG24-7806-00

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

First Edition (October 2010)

This edition applies to WebSphere CloudBurst Appliance Version 2.0.0.

Contents

Notices

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

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

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

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

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

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

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

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

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

COPYRIGHT LICENSE:

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

Trademarks

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

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

AIX®
CloudBurst™
DB2®
IBM®
Passport Advantage®
Power Systems™
POWER6®
PowerHA™
PowerVM™
POWER®
pSeries®
Redbooks®
Redpaper™
Redbooks (logo) ®
System p5®
System p®
Systems Director VMControl™
Tivoli®
WebSphere®
z/VM®

The following terms are trademarks of other companies:

ITIL is a registered trademark, and a registered community trademark of the Office of Government Commerce, and is registered in the U.S. Patent and Trademark Office.

Java, and all Java-based trademarks are trademarks of Sun Microsystems, Inc. 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.

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

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

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

Preface

This IBM® Redbooks® publication discusses the concepts and implementation of PowerVM™ and the WebSphere® CloudBurst™ appliance. This book is aimed at administrators and developers who have little knowledge of PowerVM, but in-depth knowledge of WebSphere software.

Cloud computing is the pooling of computing resources to provide a single source of computing power to multiple users. A cloud manager provides a self-service portal that maintains permissions and information about cloud objects such as virtual images, patterns, and resources. The WebSphere CloudBurst Appliance represents a cloud manager. It is a secure hardware appliance that optimizes the configuration, deployment, and management of WebSphere Application Server environments in a cloud. It can also be used by service providers providing hosted public clouds and software-as-a-service environments to simplify and standardize repeated deployments of their software.

This book includes an introduction to cloud computing and how the WebSphere CloudBurst appliance fits into business today. It presents the features and benefits of using the WebSphere CloudBurst Appliance and the advantages of using PowerVM. It provides the steps required to implement WebSphere CloudBurst appliance with PowerVM.

The team who wrote this book

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

Shivaji D Bhosale is a Team Lead with the India Software Lab, Systems and Technology group, Pune, India. He has 14 years of experience in UNIX® and middleware. He has a master's degree in Computer Management from Shivaji University Kolhapure. His areas of expertise include UNIX, pSeries®, and related technologies. He is an IBM Certified Specialist in pSeries AIX® System Administration, Tivoli® Monitoring, and EXIN ITIL®. His areas of interest include cloud computing, pSeries virtualization and systems management automation.

Martin Cernicky is a Certified IT Specialist from IBM STG Lab Services in the Czech Republic. He has 20 years of experience within the IT sector and has been working with IBM since 1995. Martin is currently working as a System p® Specialist and has extensive knowledge of AIX/pSeries systems and solutions. Additional responsibilities over the last eight years have included supporting WebSphere MQ middleware messaging and WebSphere Message Broker. He has designed and implemented solutions using the WebSphere MQ family products and has substantial experience implementing messaging and broker solutions. Martin is a co-author of the IBM Redbooks publications *Migrating to WebSphere Message Broker Version 6.0*, SG24-7198; *Managing WebSphere Message Broker Resources in a Production Environment*,SG24-7283; and *WebSphere MQ V7.0 Features and Enhancements*, SG24-7583. Martin has a degree in Automated Technology Systems at the Czech Institute of Technology.

Paul Reynolds is a Virtualization Specialist for IBM Australia. Paul has worked in UNIX and virtualization-based solutions for the past seven years. In 2006 he began supporting customers using PowerVM and PowerHA™ solutions, and his areas of expertise are PowerVM and x86 virtualization. Paul recently moved to the Emerging Technologies team,

where he focuses on cloud computing infrastructure. Paul is an IBM Certified Systems Expert - IBM System p5® Virtualization Technical Support.

Margaret Ticknor is an IT Support Specialist at the IBM ITSO Center in Raleigh. She manages the production and testing environments for the development of IBM Redbooks publications in Raleigh. Prior to joining the ITSO in 1997, Margaret worked in Endicott on the WW VM platform, supporting internal VM customers. Margaret attended the Computer Science program at State University of New York at Binghamton.

Residency team (from left to right): Martin Keen, Shivaji Bhosale, Martin Cernicky, Paul Reynolds, and Margaret Ticknor

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

Martin Keen and Tamikia Barrow
International Technical Support Organization, Raleigh Center

Matthew Sheard
IBM US

Dustin Amrhein
IBM US

Peter Van Sickel
IBM US

Brian Stelzer
IBM US

Larry E Williamson
IBM US

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

Part 1

Overview

This part introduces the concept of cloud computing and the WebSphere CloudBurst Appliance as a self-service cloud management device. It discusses the advantages of PowerVM and the PowerVM solution for WebSphere CloudBurst Appliance. This part comprises the following chapters:

1

Introduction to WebSphere CloudBurst Appliance

Cloud computing is a concept that describes the pooling of computing resources in a *cloud* to provide a single source of computing power to multiple users. The cloud manager provides a self-service portal that maintains permissions and information about cloud objects such as virtual images, patterns, and resources.

The WebSphere CloudBurst Appliance is a secure hardware appliance that represents such a cloud manager. It optimizes the configuration, deployment, and management of WebSphere Application Server environments in a cloud. Whereas the WebSphere CloudBurst Appliance is targeted for private or on-premise cloud computing environments, it can also be used by service providers that provide hosted public clouds and software-as-a-service environments to simplify and standardize repeated deployments of their software.

In this chapter we introduce:

1.1 Today's business issues

IT costs are continually increasing, often due to higher administration and operation costs, and from low utilization of servers that are consuming valuable space and power. These trends have increased the interest in the IT industry of using improved management and server consolidation capabilities as the means to reduce these costs.

The need for a consolidated infrastructure extends to the IBM WebSphere environments within the data center. Like other applications, current WebSphere environments, particularly development and test, often span many different servers, all running at low utilization. This inefficient use of resources can lead to higher power and space costs than necessary.

Reducing operational cost is perhaps an even more significant concern. Much of the administrative cost in an IT environment is driven by the amount of work involved in designing, installing, configuring, and maintaining solutions comprising different software components. For example, installation of a middleware solution can take weeks or months to complete, depending on staff resources, staff skills, problems encountered during installation/configuration, change authorization, and availability of hardware.

1.2 What is WebSphere CloudBurst

WebSphere CloudBurst (Figure 1-1) is a self-service cloud management device that decreases labor costs of IT operations and increases hardware utilization. It is a new class of hardware appliance that sits in a data center, dispensing WebSphere Application Server topologies into an external pool (cloud) of virtualized hardware.

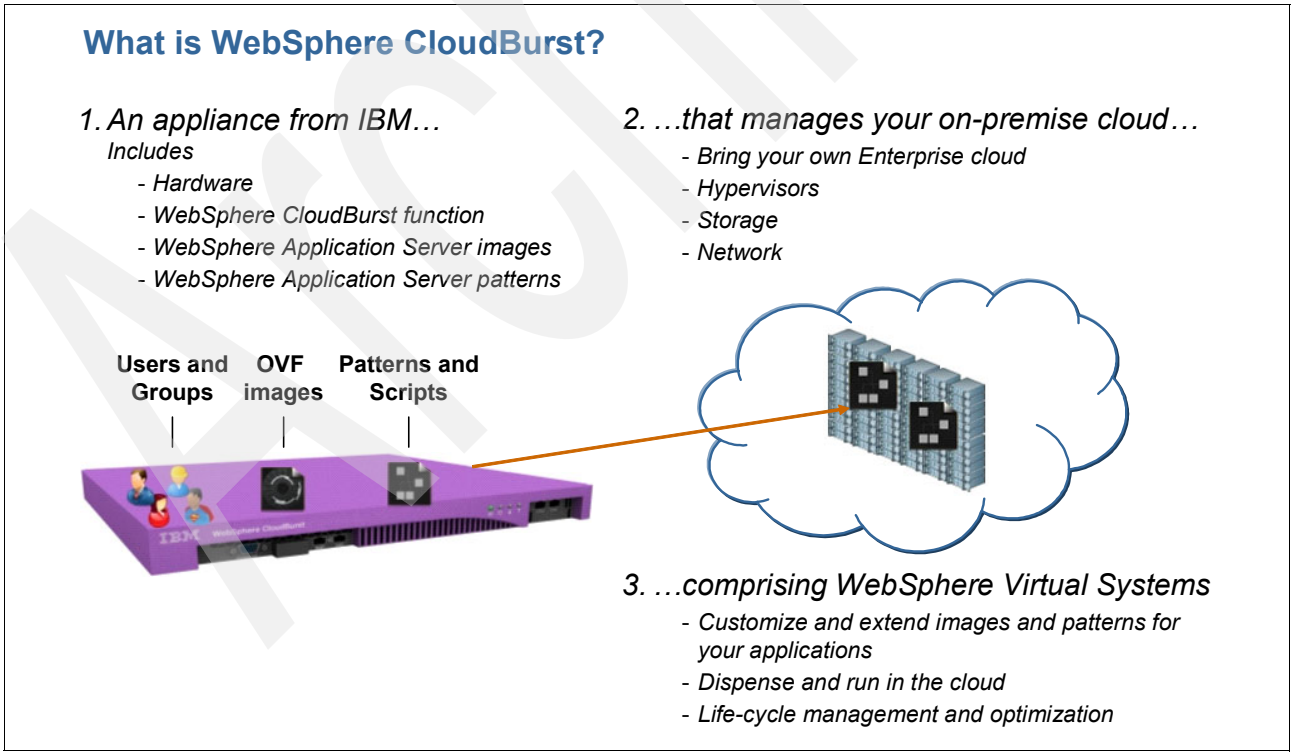

Figure 1-1 WebSphere CloudBurst

WebSphere CloudBurst Appliance helps you easily and quickly develop, test, and deploy business application environments, ending the use of manual, complex, and time-intensive processes associated with creating application environments. Its role is to deploy and monitor runtime environments. After the topology is deployed using WebSphere CloudBurst, the users and administrators connect directly to the deployed WebSphere environments.

The WebSphere CloudBurst Appliance does not play any role in the operations and running of the deployed WebSphere Application Server environments. As such, it is not in the critical data path for a particular environment after that environment is running. At the end of a reservation, resources are automatically returned to the shared resource pool and logged for internal charge-back purposes.

WebSphere CloudBurst Appliance brings together preconfigured WebSphere Application Server Hypervisor Edition virtual images and patterns of WebSphere Application Server environments that are based on more than 10 years of experience. WebSphere CloudBurst understands the software that it manages and, as such, does not treat it as a black box. This means that more can be done toward the goal of removing manual labor and administrative costs from IT operations, when compared with more general-purpose competition. For example, WebSphere CloudBurst not only stands up individual instances of IBM software products, but it configures them to communicate and work together. It is also able to apply IBM-provided fixpacks to operating systems or middleware that have been dispensed into the cloud.

WebSphere CloudBurst allows you to use existing hardware in your datacenter as the cloud into which WebSphere CloudBurst deploys its patterns. This hardware must be running a supported hypervisor (for example, PowerVM, z/VM®, or VMware ESX). A hypervisor is a software or firmware virtualization program that provides a layer of abstraction between operating systems and physical resources on a machine. This abstraction enables multiple operating systems and application stacks to run on a single physical resource, enabling higher levels of resource utilization.

1.3 Features of WebSphere CloudBurst Appliance

The next sections provide an overview of several abilities and features of WebSphere CloudBurst Appliance.

1.3.1 Consistency and compliance

Without central management of runtime environments, you often find that within an organization each person or department supplies its own unique environment and configuration. This can lead to excessive time spent in building, deploying, and maintaining systems across the departments. It can also lead to unpredictability in runtime performance and compliance.

WebSphere CloudBurst allows you to simplify the maintenance and management of complex environments. It allows you to establish a catalog of virtual images and patterns that provide the basis for the systems. Users can use artifacts in this catalog either as the precise environment they need or as the basis for further customization (depending on configured permissions). This ensures consistency across deployments by reducing the number of configurations and by driving compliance to company standards. It also reduces the number of configurations that have to be maintained.

1.3.2 Deployment and updating

Building new systems can be involved and, when done manually, can be a slow process that is prone to error. The steps include installing and customizing the operating systems, installing the middleware, bringing it up to the current maintenance levels, and building and customizing the WebSphere topology. This can take days, weeks, or months, depending on the number of systems and the complexity of the configuration. Deploying applications to those systems can also take significant time.

With WebSphere CloudBurst, new runtime environments can be defined quickly by using existing virtual images and patterns. The virtual images can be used as is, or can be extended and customized for your unique deployment needs. WebSphere CloudBurst dispenses WebSphere virtual image software that is pre-installed, configured, optimized, and ready to run. Dispensing the systems to the cloud can be done quickly and efficiently, with the appliance ensuring that hardware resource usage is balanced and optimized. Each organization has control over customizing and dispensing the systems, leading to a more efficient process with less human resources required.

Not only does WebSphere CloudBurst allow organizations to build and deploy systems faster, but the repeatability of the process can significantly reduce the risk of errors in the configuration and deployment.

Making it possible to consolidate and use existing hardware resources efficiently can reduce the number of servers required. Servers that traditionally ran at low utilization rates now run at higher rates. This has advantages in both the cost of hardware and in the cost of the space, energy, and cooling requirements to run that hardware.

1.3.3 Maximization of resources

IT organizations realize the advantages of reducing hardware expenditures by allowing multiple organizations to consolidate resources. This can mean moving individual development and test machines into a common pool. Until they see the benefits of this new model, organizations are often reluctant to give up exclusive control of their hardware. Security, user/group permissions, and usage tracking mechanisms for charge-back purposes exist in WebSphere CloudBurst to help ease this initial reluctance. Next we provide details about these mechanisms

All the credentials for shared resources in the cloud are kept on encrypted flash inside the appliance. The appliance is tamper-resistant. If someone tries to get inside the appliance, it does not work anymore. The images are all encrypted on the appliance. The management capabilities allow for segmenting the cloud by organization, and also provide fine-grained access control to images, patterns, and scripts so that organizations have full control over who can see and use assets within or across organization.

WebSphere CloudBurst allows a self-service model. The WebSphere CloudBurst roles and permissions model facilitates self-service access while controlling what is dispensed and by whom. Roles restrict who can create and modify images, patterns, and scripts. For example, you can have users who can dispense environments, but they cannot change the image that is dispensed. Permissions allow fine-grained control to specific images, patterns, and scripts. For example, production HR patterns could not be accessed by anyone other than the HR operations team. This combination of security mechanisms allow the company to lock down WebSphere CloudBurst to only dispense approved, compliant images and patterns.

1.3.4 Optimization of performance, hardware, and infrastructure

WebSphere CloudBurst Appliance uses real-time performance data for CPU, memory, and storage to make placement decisions to optimize and balance the overall performance of the cloud. In addition, WebSphere CloudBurst is aware of both the virtual machines and the underlying physical servers, and places WebSphere cells to maximize both performance and availability. For example, custom nodes will be spread across physical servers to avoid single points of failure from the underlying physical server.

1.3.5 User chargeback

Cloud computing is often associated with pay-for-use pricing structures. To achieve this, of course, cloud resource usage must be tracked and reported. WebSphere CloudBurst provides statistics about cloud usage that support charge back in the enterprise. For each user within WebSphere CloudBurst, administrators can retrieve information about his virtual machine usage and CPU, memory, and IP utilization rates. All of this information can be viewed within the WebSphere CloudBurst console, or it can be downloaded into a spreadsheet.

In addition to resource usage by user, WebSphere CloudBurst also provides information about the usage of the cloud's resources. The utilization rates of physical and virtual CPU and memory usage, IP usage, and storage usage are all available in the WebSphere CloudBurst console.

1.3.6 Virtual images

The WebSphere CloudBurst Appliance provides virtual images that are immediately available for use or that can be customized. The images are WebSphere Application Server Hypervisor Edition images, a packaging of IBM WebSphere Application Server that is optimized for virtual environments. With the WebSphere Application Server Hypervisor Edition, users get a pre-installed operating system, a WebSphere Application Server, and an HTTP server. These virtual images are used by the appliance to create virtual machines for deployment into the cloud.

Virtual images follow the Open Virtualization Format (OVF) specification, which is an industry standard specification for packaging and distribution of virtual appliances containing one or more virtual machines. Using OVF provides benefits in providing a standard mechanism to communicate virtual machine resource requirements to different hypervisors. OVF standardizes the mechanism to pass parameter information into a virtual machine.

1.3.7 Patterns

WebSphere CloudBurst uses patterns to define the WebSphere Application Server topology for deployment. WebSphere CloudBurst includes preconfigured patterns based on years of WebSphere leading practices. These patterns have been hardened and tuned to ensure high-performing WebSphere environments. These patterns can be used as is or they can be cloned and customized.

Patterns are similar to templates that can be used to deploy on the cloud. Each virtual image part in the pattern is deployed as its own virtual machine into the cloud. A pattern that has been deployed is called a virtual system.

WebSphere CloudBurst is said to be middleware aware, as it has knowledge of how to dispense each pattern, including how to properly bring infrastructure online and initialized. It

understands the relationship between the clustered nodes and the Websphere Deployment Manager (DMgr).

After a pattern has been created in the WebSphere CloudBurst Appliance, the pattern can be reused to create multiple identical WebSphere environments in the cloud. Just as with the custom images, these custom patterns are stored on the appliance and can be reused as necessary to ensure consistent, repeatable WebSphere environments. Custom patterns can also be reused between different cloud group platforms.

1.3.8 Deployment

An administrator dedicates a portion of the data center servers to the cloud. This cloud of servers is used by WebSphere CloudBurst as the pool of machines into which it dispenses WebSphere Application Server. Instead of installing and configuring the WebSphere Application Servers on each machine, a WebSphere CloudBurst administrator creates the WebSphere topology as a pattern and deploys that pattern into the cloud.

1.3.9 Customization

The software delivered to the cloud is highly customizable. This customization can occur at multiple layers, corresponding to typical IT roles:

- Virtual images can be extended to include additional software, for example, Java™ Database Connectivity (JDBC) drivers, adapters, firewalls, and anti-virus programs.
- Patterns can be customized to create WebSphere topologies that meet the requirements for each type of deployment.
- The WebSphere installations can be customized with scripts that specify configuration actions to be taken at deployment, for example, installing applications, configuring resources, and customizing security settings.

These customization options can be saved for repeated deployments. In addition, the operations team can provide customizations that are unique to a deployment, for example, passwords.

1.4 IBM WebSphere CloudBurst Appliance 2.0

IBM WebSphere CloudBurst Appliance 2.0 is the second version of IBM WebSphere-based cloud manager. It supports the three leading cloud platforms across three hardware architectures:

- VMware
- PowerVM
- zVM

The WebSphere CloudBurst Appliance is built on 10 years of WebSphere Application Server best practices, so you know that any predefined deployments will be built and tuned to bring the best out of any environment.

New features for the WebSphere CloudBurst Appliance 2.0 include:

- Updated virtual images of WebSphere Application Server 6.1.31 and 7.0.0.11.
- Introduction of the Intelligent Management Pack, which has features similar to WebSphere Virtual Enterprise and enables the creation of complex, highly available WebSphere Application Server cluster patterns.
- New advanced cluster predefined patterns, which allows the customer to quickly deploy and modify best practices designed WebSphere Application Server clusters.
- New licensing tool, which provides the ability track and enforce WebSphere Application Server licenses rules. It also has the ability to alert administrators of high license usage.
- New placement algorithm, which attempts to decrease license usage, while optimizing and balancing workloads on physical hardware.
- Monitoring of the appliance is now possible via Simple Network Management Protocol (SNMP). It is now possible to share SNMP trap information with your preferred management device.
- The user and group permissions have been improved to allow a read-only setting to cloud and appliance administration tasks.
- The new version of the appliance now allows you to import virtual images from a secure shell (SSH) enabled server.
- Multi-image pattern composition and deployment, single pattern now can be built using different products and versions of WebSphere Application Server.
- Image built with RedHat (VMware only): No longer is a services engagement required.

1.5 Summary

In summary, WebSphere CloudBurst provides high business value through increased efficiency, cost-effectiveness, and usability of WebSphere topologies in a data center by taking advantage of cloud computing principles.

The IBM WebSphere CloudBurst Appliance lowers costs through reduced labor costs, higher hardware utilization, fewer errors, and less power consumption. It can provide faster time-to-value by automating repetitive, time-consuming manual tasks. It empowers individual application managers, developers, or testers to perform business without lengthy approval processes.

By packaging WebSphere virtualization and WebSphere intelligence along with easy-to-use interfaces, WebSphere CloudBurst brings private WebSphere clouds to the enterprise. WebSphere CloudBurst ships with environment patterns that codify 10 years of WebSphere management leading practices. WebSphere environments can be customized and captured once, then are able to be dispensed at the push of a button. Its increase agility and rapid setup and tear down of WebSphere Application Server environments means less time spent managing WebSphere environments, and it integrates well into your enterprise.

One of the strengths of WebSphere CloudBurst is the ease of dispensing an optimized solution, more than just provisioning a few virtual machines or deploying and installing middleware. WebSphere CloudBurst supports dispensing complex and loosely connected middleware with inherent knowledge of that infrastructure. WebSphere CloudBurst is thus uniquely positioned to dispense secure and optimized solutions into a cloud.

2

WebSphere CloudBurst Appliance with PowerVM Overview

This chapter discusses the details about the PowerVM solution for WebSphere CloudBurst Appliance. It provides an overview from the business, architecture, and configuration perspectives for the reader to understand the PowerVM solution.

The following topics are discussed:

2.1 PowerVM advantages

PowerVM virtualization is the stable and powerful IBM solution to consolidate multiple workloads onto fewer systems, increasing server utilization and reducing cost. PowerVM provides a secure and scalable virtualization environment for AIX, IBM i, and Linux® applications built upon the advanced reliability, availability, and serviceability (RAS) features and leading performance of the Power Systems™ platform. Highlights are:

- Consolidate diverse sets of applications built for multiple operating systems on a single server:
  - AIX
  - IBM i
  - Linux

  Note: Only PowerVM on AIX is supported for WebSphere CloudBurst Appliance at this time. The PowerVM resources must be managed by a Hardware Management Console (HMC). Resources managed by Integrated Virtualization Manager (IVM) are not supported.

- Virtualize processor, memory, and I/O resources to increase asset utilization and reduce infrastructure costs.
- Dynamically adjust server capability to meet changing workload demands.
- Move running workloads between servers to maximize availability and avoid planned downtime.

The IBM Power Systems family of servers includes proven server consolidation platforms that help you control costs while improving overall performance, availability, and energy efficiency. With these servers and IBM PowerVM virtualization solutions, your business can consolidate applications and servers, virtualize its system resources, and provide a more flexible, dynamic IT infrastructure.

PowerVM delivers industrial strength virtualization for AIX, IBM i, and Linux environments on IBM POWER® processor-based systems. IBM Power Systems servers coupled with PowerVM technology are designed to help clients build a dynamic infrastructure, reducing costs, managing risk, and improving service levels.

PowerVM offers a secure virtualization environment, built on the advanced RAS features and leadership performance of the Power Systems platform.

Using PowerVM consolidation and virtualization for WebSphere cloud is the efficient way to meet specific WebSphere CloudBurst Appliance requirements in a more dynamic fashion. The latest IBM advances in virtualization break out of traditional solutions and offer better platform support and scalability than ever before.

2.2 PowerVM architecture for WebSphere CloudBurst Appliance

The WebSphere CloudBurst Appliance solution contains several components (Figure 2-1).

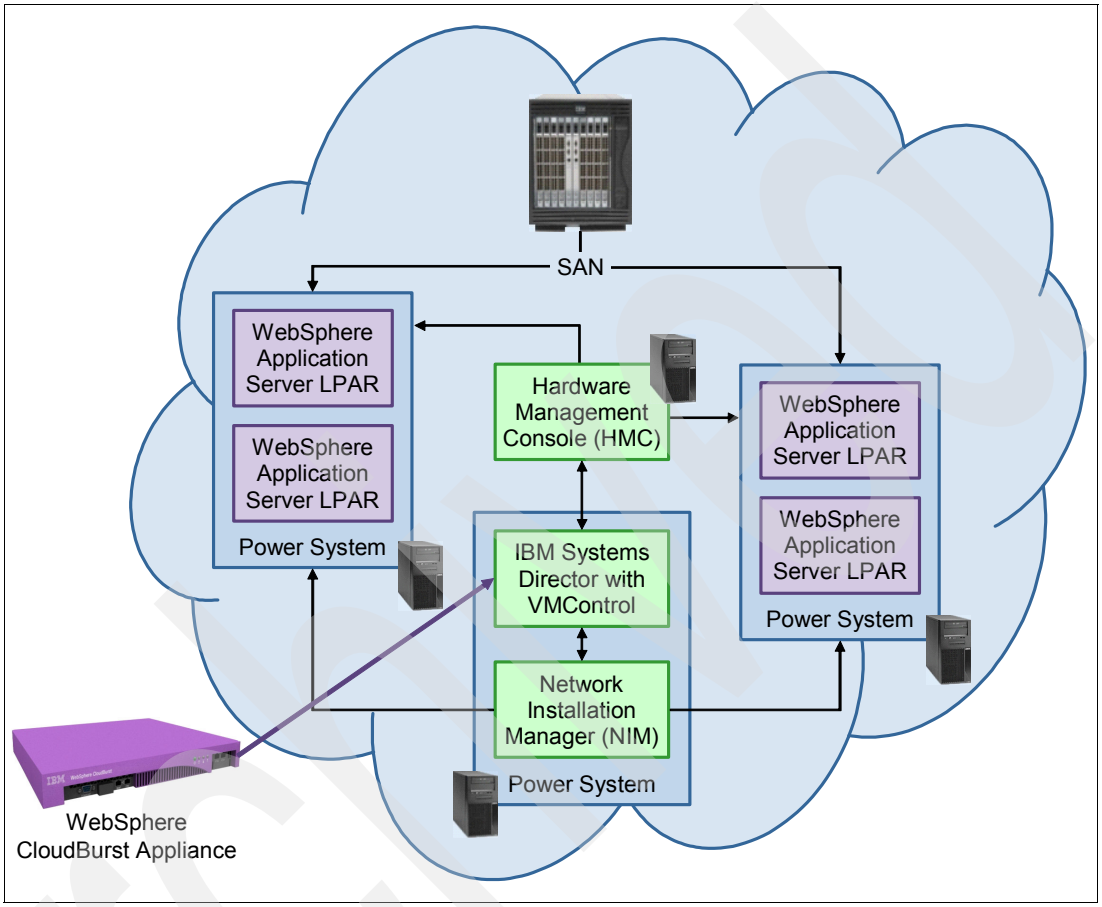

Figure 2-1 Components of WebSphere Cloud

The basic components are:

- Power Systems
- Storage systems
- Management systems
- Networks
- WebSphere CloudBurst Appliance

The following sections describe these components and their roles in the solution. The last section in this chapter describes the communication flow of the entire solution.

2.2.1 Power Systems

The processing core of a WebSphere cloud is Power Systems. Power Systems run the application workload. Power Systems are divided into logical partitions (LPARs). Each LPAR is a logically independent AIX system that hosts WebSphere Application Server. The AIX and Websphere application server are bundled together into WebSphere Application Server Hypervisor Edition images. The images are stored in WebSphere CloudBurst Appliance.

LPARs are totally virtualized via Virtual I/O Server (VIOS). This means that the disks and Ethernet adapters are not dedicated to LPARs, but are shared via VIOS, which owns these resources. VIOS itself runs within its own LPAR.

2.2.2 Storage systems

VIOS provides virtualized storage for LPARs. Therefore, each VIOS needs enough storage space for LPARs. A common solution is a disk storage system, typically in a Storage Area Network (SAN) environment. However, you can use any other storage system or bunch of local disks supported by VIOS.

The sizing of appropriate storage systems must be done in accordance with required capacity and I/O performance. This sizing is out of the scope of this book.

VIOS itself can use the same storage as for LPAR virtualization, but dedicated local disks are suggested.

2.2.3 Management systems

Several management systems are needed to manage and operate the cloud. The basic ones are:

- Hardware Management Console
- Network Installation Manager (NIM)
- IBM Systems Director with VMControl plug-in

Hardware Management Console

Power Systems and their resources are controlled from the HMC. However, administrators use WebSphere CloudBurst Appliance for managing virtual resources in a cloud solution.

Though you can have up to two HMCs to manage a central electronics complex, VMControl only supports a single HMC.

Network Installation Manager server

NIM is a free AIX subsystem that allows you to install, upgrade, patch, or back up individual AIX systems. The NIM environment consists of a NIM server and one or more NIM clients. The NIM server must be installed and configured manually on an AIX LPAR with sufficient resources to work with image sizes ranging up to 16 GB. NIM clients are installed automatically during standard AIX installation.

Note: Consider using a separate AIX system for the NIM server.

IBM Systems Director with VMControl plug-in

IBM Systems Director is a platform manager tool. It can monitor and manage several IBM server systems, storage systems, uninterruptible power supplies, and so on. In the case of Power Systems, IBM Systems Director works with HMC, not directly with hypervisor.

The VMControl plug-in for Director is a tool for virtual resource management. It can:

- Create, modify, and delete LPARs and their resources
- Initiate LPAR operations as activate, shutdown, restart, or dynamic operations.
- Manage VMware ESX and z/VM hypervisor environments.

IBM Systems Director and the VMControl plug-in can be installed on various operating systems. Consider running IBM Systems Director on an AIX platform when using the WebSphere CloudBurst Appliance PowerVM cloud solution.

Further management and utility systems can be used with the cloud, for example:

- Backup and archiving systems (IBM Tivoli Storage Manager)
- Monitoring systems (IBM Tivoli Enterprise Monitoring)
- License management systems (IBM License Metric Tool)
- User management systems (LDAP).

However, these systems are not required to build a functional WebSphere cloud solution.

2.2.4 Networks

The cloud requires several Ethernet networks for management or application communications. Networks that can be used are:

- A separate private HMC network for Power Systems
- A private HMC network for RMC/deploy communications
- A management network
- An application network

Separate private HMC network for Power Systems

This network is required for Power Systems hardware management. It is totally separated and uses dynamic IP addresses. The HMC acts as a Dynamic Host Configuration Protocol (DHCP) server. HMC controls Power Systems via this network through a Flexible Service Procesor (FSP) embedded in each Power System.

Private HMC network for RMC/deploy communications

This network is used for Resource Monitoring and Control (RMC) and deploy communication. RMC is an AIX subsystem for system control and in a virtualized environment is used for communication between HMC and LPARs. Consider setting up a separate RMC/deploy network for cases of a higher number of LPARs and deployment. The management network can be use instead for small solutions.

Management network

The management network connects all management systems (WebSphere CloudBurst Appliance, HMC, NIM, and IBM Systems Director) and is also used for administrators access.

Application network

The network for application traffic is only connected to each VIOS and to the appropriate part of the customer's application network. There can be more than one application network in the cloud solution, such as the customer core network, internal demilitarized zone, external demilitarized zone, and so on.

2.2.5 WebSphere CloudBurst Appliance

WebSphere CloudBurst Appliance is the main management component of the WebSphere cloud. It contains specific management software and WebSphere Application Server Hypervisor Edition images. The cloud administrator uses the appliance user interface to set up, deploy, and monitor running systems and administer the cloud.

The appliance communicates with the cloud through VMControl. VMControl is responsible for communication with NIM and HMC.

The types of cloud resources are:

- IP groups

  An IP group is a collection of IP addresses that are available for virtual systems to use.

- Hypervisors

  Examples of hypervisors that are supported by WebSphere CloudBurst Appliance are VMware ESX, PowerVM, and z/VM. Hypervisors are what make up a cloud group.

- Cloud groups

  A cloud group is a collection of related hypervisors. When deploying patterns to create virtual systems, you use a cloud group as the deployment target. One or more hypervisors of the same type make up a cloud group.

2.2.6 Communication flow

This section describes the communication among the cloud components to show how the cloud works.

The WebSphere CloudBurst Appliance is on the top of the cloud. As mentioned in the previous section, the appliance communicates only with VMControl by way of the management network. Administrators connect the appliance by way of the management network and uses one of the appliance user interfaces (web-based GUI, command line) to work with the cloud, and then the appliance sends the command to VMControl (Figure 2-2).

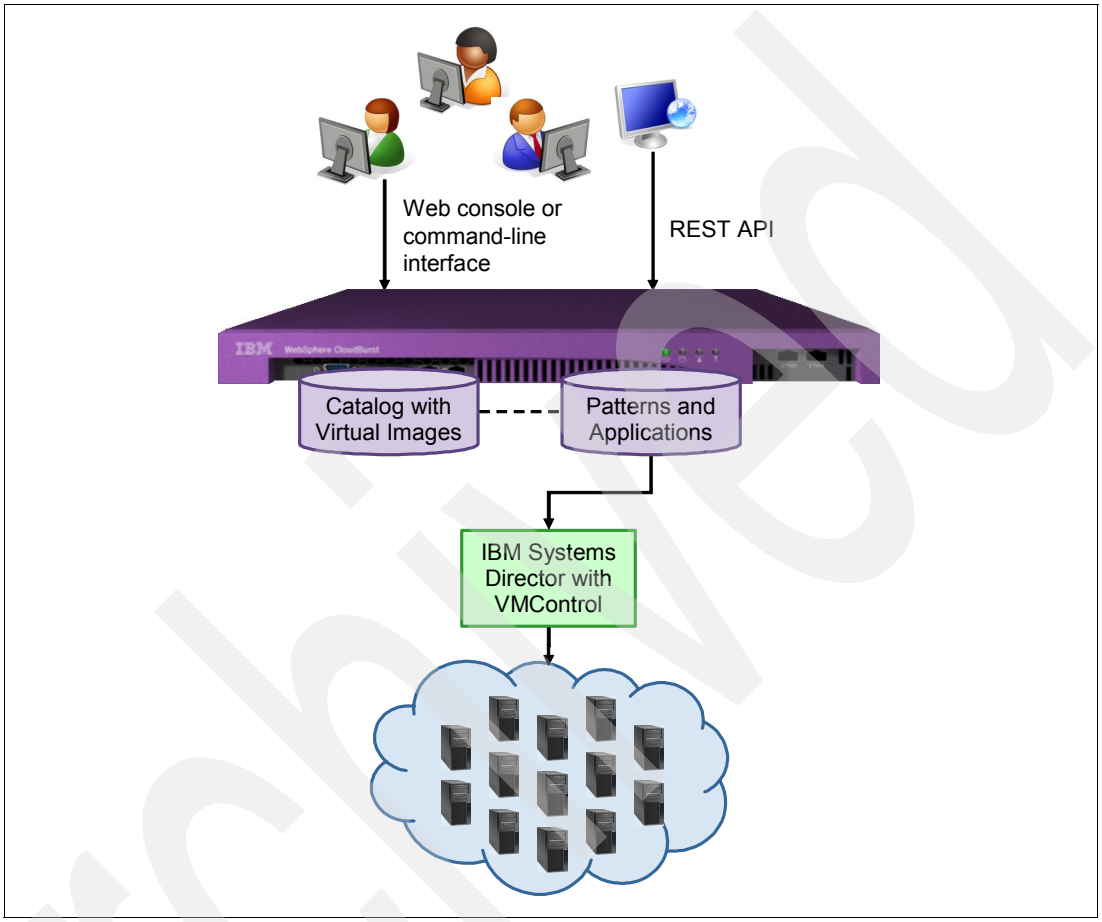

Figure 2-2 WebSphere CloudBurst Appliance communication flow

For LPAR-oriented operations, VMControl, by way of the IBM Systems Director, sends commands to the HMC, by way of the management network. The operations, which require RMC communication, use the RMC/deploy network if it exists. On a lower level, the HMC sends commands from VMControl to the Power Systems.

In the case of image operations, such as deploy and capture, VMControl sends commands to the NIM, which in turn prepares the required resources, initiates the boot, and installs the software on the LPAR.

During the first installation of a virtual image the image is transferred from the appliance to VMControl. A virtual appliance is created on the VMControl server and the image is transferred from the VMcontrol to the NIM. A Shared Product Object Tree (SPOT) image is created on the NIM server and an installation from the SPOT is performed. The SPOT resource is retained on the NIM after the initial successful deployment to speed up the next deploy of the same image. One virtual image is used to deploy several patterns.

2.3 Implementation overview

This chapter describes basic installation and configuration steps needed to build the cloud and suggestions for cloud architecture. For detailed steps for installing and setting up specific components, refer to the appropriate product documentation or IBM Redbooks publication.

Basic configuration steps

This section lists a high-level overview of what must be done for the configuration. The reader is expected to discuss these steps with the appropriate systems specialist if further explanation and help are required.

The process comprises the following steps:

1. Hardware and network installation
2. HMC configuration
3. Disk storage installation and configuration
4. VIOS installation and configuration
5. NIM server installation and configuration
6. IBM Systems Director installation and configuration
7. VMControl plug-in installation and configuration
8. NIM integration into VMControl
9. WebSphere CloudBurst Appliance configuration

The suggestions and rules for each step are covered in the following sections.

2.3.1 Leading practices and rules

All cloud LPARs are fully virtualized. All required I/O resources are virtualized and made available to the LPARs in the cloud by the VIO server. The following circulation techniques are used:

- Shared Ethernet Adapter (SEA) for network virtualization
- Virtual SCSI for storage virtualization

The main implementation rule is to *always* read the product `readme` files and use proper versions for software and firmware. Especially in case of any component upgrade, check for ties to other components.

Hardware and network installation

The special hardware requirements from a WebSphere CloudBurst Appliance point of view are:

- Appropriate firmware version for Power Systems (central electrical complex)
- Appropriate firmware version for HMC console
- Minimum of 3 GB of physical memory in HMC console

For the network setup, note the following important information:

- For the Websphere CloudBurst Appliance Solution, time synchronization is helpful for better orientation in the logs for problem determination.
- Correct name resolution is required for the WebSphere CloudBurst Appliance solution. Define appropriate TCP/IP settings and Domain Name Server (DNS) records. Both the forward and reverse resolutions have to work properly.

A separate network can be created for use with RMC/deploy communications.

HMC configuration

Use the proper version of the HMC software and then apply the latest service pack and required fixes. Check that the Power System firmware version meets the VMControl requirements and update it if required.

Disk storage installation and configuration

Configure the disk storage system and create appropriate LUNs to use in the VIOS. The capacity of the LUNs must be designed in accordance with WebSphere CloudBurst Appliance requirements and the expected number of LPARs. See “Storage pool” on page 35 for the storage capacity calculation. Use standard documentation for storage products. There are no special requirements, except that the storage has to be supported in VIOS.

VIOS installation and configuration

The special requirements from the WebSphere CloudBurst Appliance point of view for a VIOS configuration are:

- At least one configured Shared Ethernet Adapter (SEA) with external access.
- At least one volume group to use as a storage pool.

  A new volume group has to be created. Consider setting up a new volume group to be used so that you are not using rootvg.
- Only a single VIOS configuration can be used with WebSphere CloudBurst Appliance Version 2.0.

The SEA and storage pool have to be configured for application LPARs. For more information see 3.3, “VIOS logical partition configuration” on page 33.

NIM server installation and configuration

Use a separate NIM server for WebSphere CloudBurst. Only a basic configuration of the NIM server is needed. For more information see 3.5.2, “NIM Master configuration” on page 38.

IBM Systems Director installation and configuration

Use the correct version of the product. See 4.2, “Installing IBM System Director” on page 43.

VMControl plug-in installation and configuration

Use the required version of the product. See 4.3, “Installing VMControl” on page 58.

NIM integration into VMControl

Common Agent Services (CAS) and VMControl subagents must be installed on the NIM server. See 4.4, “Installing agents to NIM” on page 61.

WebSphere CloudBurst Appliance configuration

The configuration of WebSphere CloudBurst Appliance is the last, but not least, part of WebSphere cloud initialization. It consists of the following steps:

1. Set up the appliance (network, time, security, and so on.).
2. Set up the cloud (IP group, cloud groups).
3. Create virtual systems (working with patterns).

See Chapter 6, “WebSphere CloudBurst Appliance Administration” on page 99, for configuration details.

Part 2

Building the cloud

This part describes the steps required to build a PowerVM CloudBurst environment. It discusses the physical and logical components of the cloud, the required software, and basic configuration steps.

We assume that the reader has working knowledge of the PowerVM components, such as hypervisors, Hardware Management Console (HMC), Virtual I/O Server (VIOS), and Network Installation Manager (NIM). Only the tasks that are required or are specific to the WebSphere CloudBurst Appliance and the PowerVM solution are covered in this section.

This part includes the following chapters:

3

Power System components

This chapter describes the configuration of the Hardware Management Console (HMC), central electronics complex, Virtual I/O Servers (VIOS), and Network Installation Manager (NIM) server based on our lab scenario. The installation of the operating systems, AIX, and VIOS is not covered in this chapter.

Note: We assume that the reader has an understanding of HMC, central electrical complex, AIX, VIO, and NIM concepts and practical knowledge.

This chapter covers the following topics:

3.1 Test lab scenario

In this test scenario, the solution of WebSphere CloudBurst Appliance with PowerVM is shown, while staying within the framework of a realistic business situation. Like most test environments, physical resource restrictions limited the design of our test lab. Our network and hardware topology is not necessarily intended as a guide for leading practices, but rather as an example for demonstrating particular features of WebSphere CloudBurst and its cloud environment.

3.1.1 Logical topology

Figure 3-1 shows the logical topology of our test environment, WebSphere cloud topology.

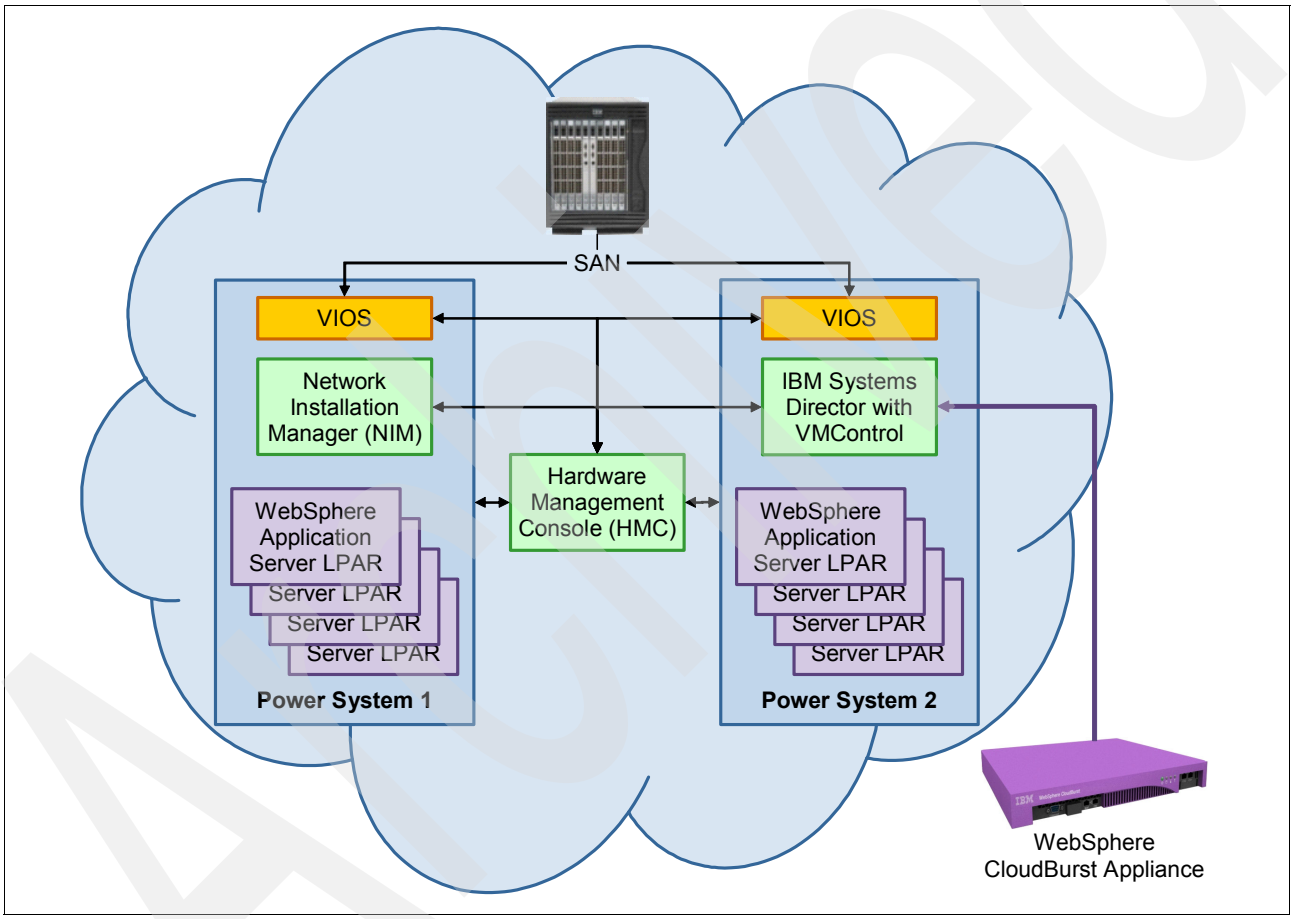

Figure 3-1 Logical topology of scenario environment

The environment consists of several kinds of logical components. Table 3-1 list all logical components used and their properties.

Table 3-1 Logical components in scenario

Component	Usage	Comments
SAN	Provides disk storage for VIOS	The SAN configuration is out of the scope of this book.
Power System 1	Hosts NIM, VIOS, and application LPARs	
Power System 2	Hosts IBM Systems Director, VIOS, and application LPARs	
VIOS	Provides network and storage virtualization for LPARs	Shared Ethernet Adapter (SEA) is used for network virtualization. vSCSI adapters/disk are used for storage virtualization.
NIM	Provides network installation capabilities	Only a basic, simple configuration for IBM Systems Director.
IBM Systems Director	Provides hardware management and monitoring	It acts also as a framework for VMControl.
VMControl	Provides complete cloud management	Runs as IBM Systems Director plug-in and is integrated with NIM.
HMC	Controls Power Systems	
Application LPARs	Runs WebSphere Application Server Hypervisor Edition	It has to be installed from WebSphere CloudBurst Appliance.
WebSphere CloudBurst Appliance	Provides complete WebSphere cloud management	

3.1.2 Physical topology of the test environment

Figure 3-2 shows the physical topology of the test environments.

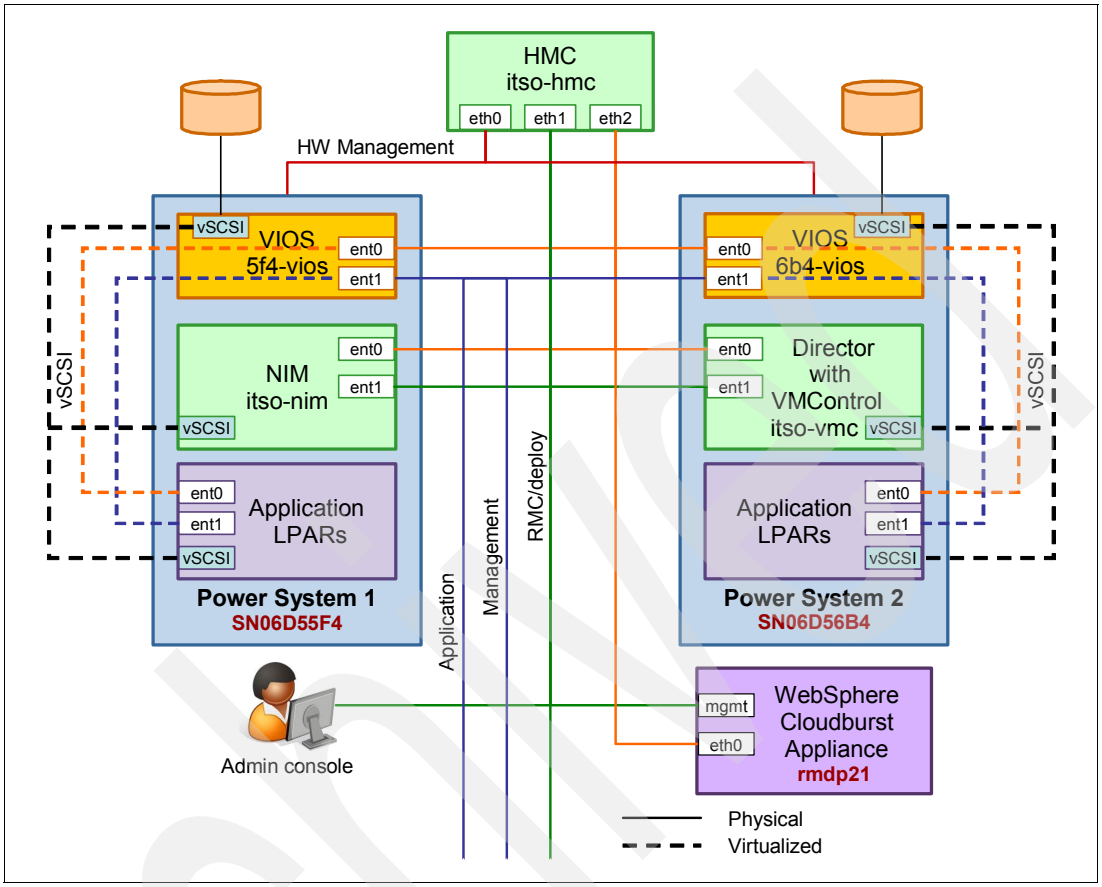

Figure 3-2 Physical topology of scenario environment

The physical environment consists of:

- Two Power Systems that run VIOS
- NIM
- IBM Systems Director server
- WebSphere CloudBurst Appliance
- Several networks

The following section list the products and versions used.

Physical hardware

Table 3-2 to Table 3-5 on page 27 provide information about the physical hardware used in this scenario. Table 3-2 provides information about the HMC machine.

Table 3-2 itso-hmc - HMC

Machine parameter	Value
Machine type	7310-CR4
Code level	7.3.5 SP, MH01217, MH01227

Table 3-3 provides information about the PowerVM managed node 1 machine.

Table 3-3 SN06D55F4 - Power System 1

Machine parameter	Value
Machine type	8204-E8A
System firmware	EL340_075
Number of CPU	8 POWER6® 4.2 GHz
Size of memory	128 GB RAM

Table 3-4 provides information about the PowerVM managed node 2 machine.

Table 3-4 SN06D56B4 - Power System 2

Machine parameter	Value
Machine type	8204-E8A
System firmware	EL340_075
Number of CPU	8 POWER6 4.2 GHz
Size of memory	128 GB RAM

Table 3-5 provides information about WebSphere CloudBurst Appliance.

Table 3-5 rmdp21 - WebSphere CloudBurst Appliance

Machine parameter	Value
Firmware level	2.0.0.0-24024
I/O	One Ethernet port for management: mgmt One Ethernet port for deploy: eth0

Installed LPARs

Table 3-6 to Table 3-9 on page 29 provide information about installed LPARs used in the scenario.

Table 3-6 provides information about the VIO Server on Power VM managed node1.

Table 3-6 itso-vio5f4 - VIOS on system 1

Machine parameter	Value
VIOS level	2.1.1.10-FP21
Number of CPU	Desired: 0.5 Maximum: 2 Uncapped Weight: 192
Size of memory	Desired and maximum: 4 GB

Machine parameter	Value
I/O	One HEA port for management/RMC One 4-port 10/100/1000 Ethernet for SEA One 2-port 4 GBit FC adapter for vSCSI disks One SAS controller for local disks Two local disks for VIOS rootvg Two local disks for NIM vSCSI disks SAN storage disks for cloud vSCSI disks

Table 3-7 provides information about VIO Server on Power VM managed node 2.

Table 3-7 itso-vio6b4 - VIOS on system 2

Machine parameter	Value
VIOS level	2.1.1.10-FP21
Number of CPU	Desired: 0.5 Maximum: 2 Uncapped Weight: 192
Size of memory	Desired and maximum: 4 GB
I/O	One HEA port for management/RMC One 4-port 10/100/1000 Ethernet for SEA One 2-port 4 GBit FC adapter for vSCSI disks One SAS controller for local disks Two local disks for VIOS rootvg Two local disks for IBM Systems Director vSCSI disks SAN storage disks for cloud vSCSI disks

Table 3-8 provides information about NIM server on Power VM managed node 1.

Table 3-8 itso-nim - NIM on system 1

Machine parameter	Value
AIX level	6100-03-01-0921
dsm.core	6.1.3.1[1]
DirectorCommonAgent	6.1.0.3[1]
VMControl subagent level	2.1 + WebSphere CloudBurst Appliance interim fix
Number of CPU	Desired: 1 Maximum: 2 Uncapped Weight: 128
Size of memory	Desired and maximum: 12 GB
I/O	Two HEA ports for management/RMC Two vSCSI disks for rootvg

1. Filesets dsm.core and DirectorCommonAgent are not a part of basic AIX installation and have to be installed manually.

Table 3-9 provides information about IBM Systems Director server on Power VM managed node 2.

Table 3-9 itso-vmc - IBM Systems Director on system 2

Machine parameter	Value
AIX level	6100-03-01-0921
IBM Systems Director level	6.1.1.2
CAS/DCA agent level	6.1.0.3
VMControl level	2.1 + WebSphere CloudBurst Appliance interim fix
Number of CPU	Desired: 1 Maximum: 2 Uncapped Weight: 128
Size of memory	Desired: 4 GB Maximum: 8 GB
I/O	Two HEA ports for management/RMC Two vSCSI disks for rootvg

Networks

Table 3-10 provides information about the network used in this scenario.

Table 3-10 Scenario networks

Network	Usage	Comments
HW mgmt	Power Systems hardware management from HMC	Separate network, configured as private at HMC console.
Management	Component administration and management communication among hardware	NIM and IBM Systems Director LPARs use Host Ethernet Adapter (HEA).
RMC/deploy	Separate network for RMC and deploy traffic	VIOS, NIM, and IBM Systems Director LPARs use HEA. Application LPARs use virtualized network via SEA on VIOS.
Application	Network for application communication	Application LPARs use virtualized network via SEA on VIOS.
vSCSI	Framework for storage virtualization	Not really a network. It is virtualization technique uses hypervisor and memory sharing.

Leading practices

Leading practices for configuring a PowerVM environment for use with WebSphere CloudBurst Appliance 2.0 are:

- Keep the NIM and IBM System Director servers on separate physical servers, as they both consume high amounts of CPU and disk resource during a deployment.
- Use dedicated physical machines for cloud infrastructure so that they do not contend for other production system resources.

- Give any infrastructure servers (such as VIOS, NIM, and IBM System Director servers) high ID numbers (80 or higher) on the HMC. This way you can quickly differentiate cloud-related servers from infrastructure servers.
- Use as few volume groups as possible for storage pools.
- If using a SAN or iSCSI backend storage pool, use large LUNs to reduce management overhead. When active, the deployed WebSphere Application Server LPARs preform little I/O.
- Do not use rootvg as a volume group for storage pools.
- Give your VIO servers a weighting value of192, as all deployed servers in the cloud have a weight value of 129. It is important that your VIO servers have the highest weight value.
- Use a dedicated HMC that only manages cloud resources. This reduces any possibility of affecting running production servers.
- Use a separate network interface for RMC communication. This network can also be used for pattern deployment.
- Install and configure the software stack.

The remaining sections of this chapter and all of Chapter 4, “Installing IBM Director and VMControl” on page 41, and Chapter 5, “WebSphere CloudBurst Appliance” on page 69, discuss configuring and installing the software required for WebSphere CloudBurst Appliance on PowerVM.

All of the following installation configuration procedures were completed by our team to create a working PowerVM cloud. The version levels we used are listed. We tested only the versions listed in this document.

3.2 Hardware Management Console configuration

The Hardware Management Console (HMC) is a system that controls managed systems, such as servers and logical partitions.

3.2.1 Prerequisites for HMC

The HMC used in the WebSphere CloudBurst Appliance software stack performs many tasks and continually collects data for each LPAR, which it feeds back to the VMControl and WebSphere CloudBurst Appliance.

The HMC requirements are:

- Model 7042-CR5 or 7310-CR4 or 7042-CR4.
- Minimum of 3 GB of physical memory, but 4 GB physical memory is suggested.
- HMC minimum code level V7R3.5.SP1.
- eFIX MH01227 must be installed on HMC.
- Network connection between HMC and client LPARs for RMC network traffic.
- Suggestion: additional Ethernet cards for connecting to multiple networks.

HMC physical memory check

To check the size of HMC physical memory perform the following steps:

1. Log in to HMC as `hscroot` via SSH.
2. Issue the command:

   ```
   cat /proc/meminfo | grep MemTotal
   ```

   This command shows the size of memory in KB.

If the physical memory is less than 3 GB, contact your IBM sales representative to upgrade HMC physical memory. A specific request for a 2 GB Miscellaneous Equipment Specification (MES) might be required.

Note: Examples of the physical memory size for certain HMC models are:

- 7042-C08: 4 GB RAM standard
- 7042-C07: 1 GB RAM standard

  2 GB RAM can be add via feature code 1935.
- 7042-C06: 1 GB RAM standard

  2 GB RAM can be add via feature code 1924 or 1935.
- 7042-CR5: 4 GB RAM standard
- 7042-CR4: 1 GB RAM standard

  2 GB RAM can be add via feature code 1924 or 1935.

3.2.2 HMC user configuration

In a shared environment where an HMC can manage many IBM POWER servers for production and cloud, it is a good idea to create a separate HMC user who only has the administrative authority over cloud hardware devices. This ensures that actions that are instigated by the WebSphere CloudBurst Appliance have no effect on any non-cloud resources connected on the HMC.

In our test scenario we created a user on the HMC called wcaadmin who had full access to the cloud servers.

HMC role

Roles are a way of allowing a user to have privileged access to certain physical machines and tasks. An HMC user can be a member of multiple roles.

We advise creating a dedicated role on the HMC for the WebSphere CloudBurst Appliance, especially in a shared environment where a single HMC manages production and cloud servers. This isolates access to the physical servers that will make up the cloud. We created a new super admin role to perform, which only allowed access to the two physical servers that make up our cloud.

To create a new role that only allows super admin access to your cloud CECs:

1. Open an HMC window.
2. Log in as hscroot or any superadmin privileged user.
3. In left pane click **HMC Management** → **Administration** → **Manage Task and Resource Roles**.

4. Click **Edit** → **Add**, and then:
   a. Enter a role name (for example, Cloudresouces).
   b. Select each available central electrical complex that is part of the cloud setup and click **Add** (Figure 3-3).

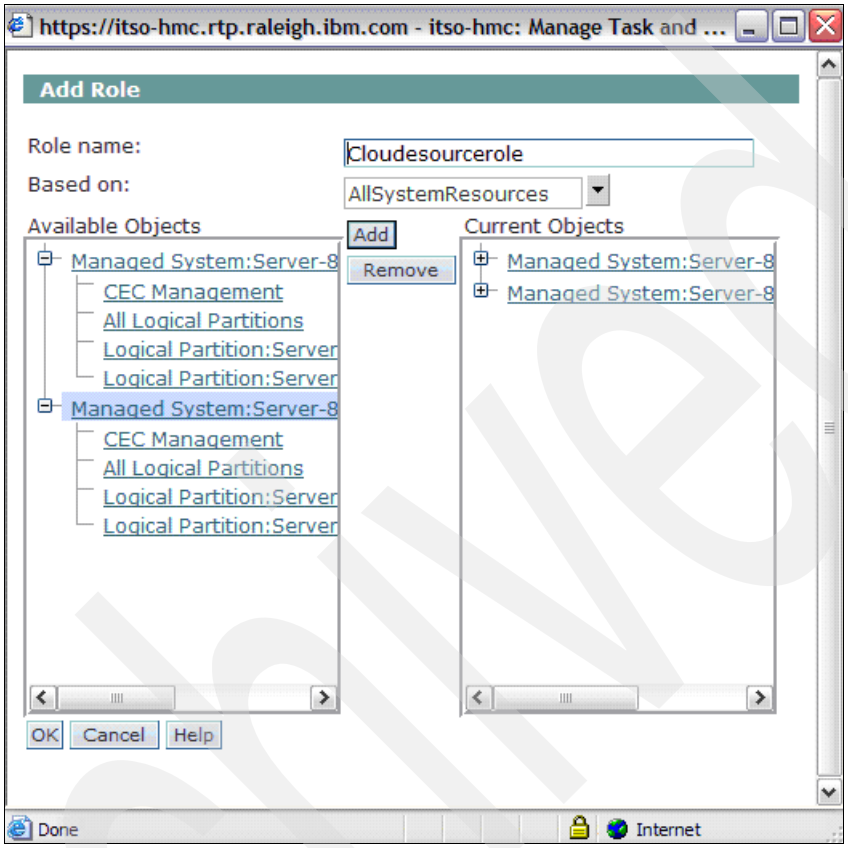

Figure 3-3 Creating HMC resource roles step GUI

   c. Click **OK**.

The role creation step is complete.

HMC user

Users are another way of restricting access to an HMC. Using the role created in the previous step gives a user full access to the cloud resources, but does not allow access to any other physical servers.

To create a dedicated user on HMC:

1. In the left pane click **HMC Management** → **Manage User Profiles and Access.** Click **User** in User profiles window.
2. Click **Add**, and then:
   a. Enter the user ID (for example, wcaadmin).
   b. Enter a description.
   c. For authentication, select **Local Authentication**.
   d. Enter the password and confirm it.

e. Select the role that has access to the CECs that are part of the cloud resource.

f. For the tasks roles, select **hmcsuperadmin**. Click **OK**.

The user creation step is complete.

HMC network configuration

The HMC needs to be able to communicate with the managed systems. The following networks are required for clear communication between the user and HMC, the LPARs and the HMC, and the CECs and the HMC:

- HMC management for accessing the HMC web and command-line console.
- RMC/deploy network for RMC communication between HMC and logical partitions. Without an RMC network, the WebSphere CloudBurst Appliance solution does not work.
- System Management Network, for communication between the HMC and the physical Power servers (central electrical complexes).

3.2.3 Central Electronics Complex configuration

The central electrical complex is a building block of CPUs, physical memory, and PCI backplane linked together by cables. The CPU and physical memory interconnect at the front and I/O interconnect at the back. At the basic level, the central electrical complex is the physical server.

WebSphere CloudBurst Appliance 2.0 supports all Power 5 and Power 6 architectures that are enabled on host a Virtual I/O Server and that are managed by an HMC. All LPARs deployed in the cloud use virtual I/O devices for disk and network and are able to dynamically use and share resources as required, which is why it is essential that the hardware is capable of supporting a VIOS server. IVM is not supported by WebSphere CloudBurst Appliance 2.0.

Below are the minimum firmware levels for Power 5 and 6 hardware in a WebSphere CloudBurst Appliance 2.0 managed environment:

- IBM Power 5: 01SF240_381
- IBM Power 6: 01EL340_075

3.3 VIOS logical partition configuration

When using a PowerVM cloud with the WebSphere CloudBurst Appliance, only one VIO server is supported per physical machine. This single VIO server supplies all of the cloud LPAR's network and disk. Thus, it is advised if the physical server is dedicated to the cloud to assign all physical I/O resources to the VIO server.

The requirements for the VIOS partition are:

- Minimum of two shared virtual processors, suggested minimum CPU 0.5. This value can change depending on network traffic.
- Uncapped weight of 192, as this needs to be higher than the cloud LPARs, which have an uncapped weight of 129.
- Minimum memory of 4 GB.
- A network device with RMC communication with the HMC.

- At least one Shared Ethernet Adapter (SEA) with access to the external network. For WebSphere CloudBurst Appliance to view available networks the SEA must have external access. More than one SEA can be configured.
- Minimum of 100 virtual adapters. This sizing can change depending on your environment. With each pattern deployment a new virtual SCSI mapping is dynamically added to the Virtual I/O Server. When the virtual system is removed, the virtual SCSI adapter is removed.

  Figure 3-4 shows a summary of our partition profile.

> **Note:** Do not create empty virtual SCSI adapters in the profile, as they serve no purpose. Upon a new deployment, the WebSphere CloudBurst Appliance causes a new virtual SCSI adapter to be dynamically added to the appliance.

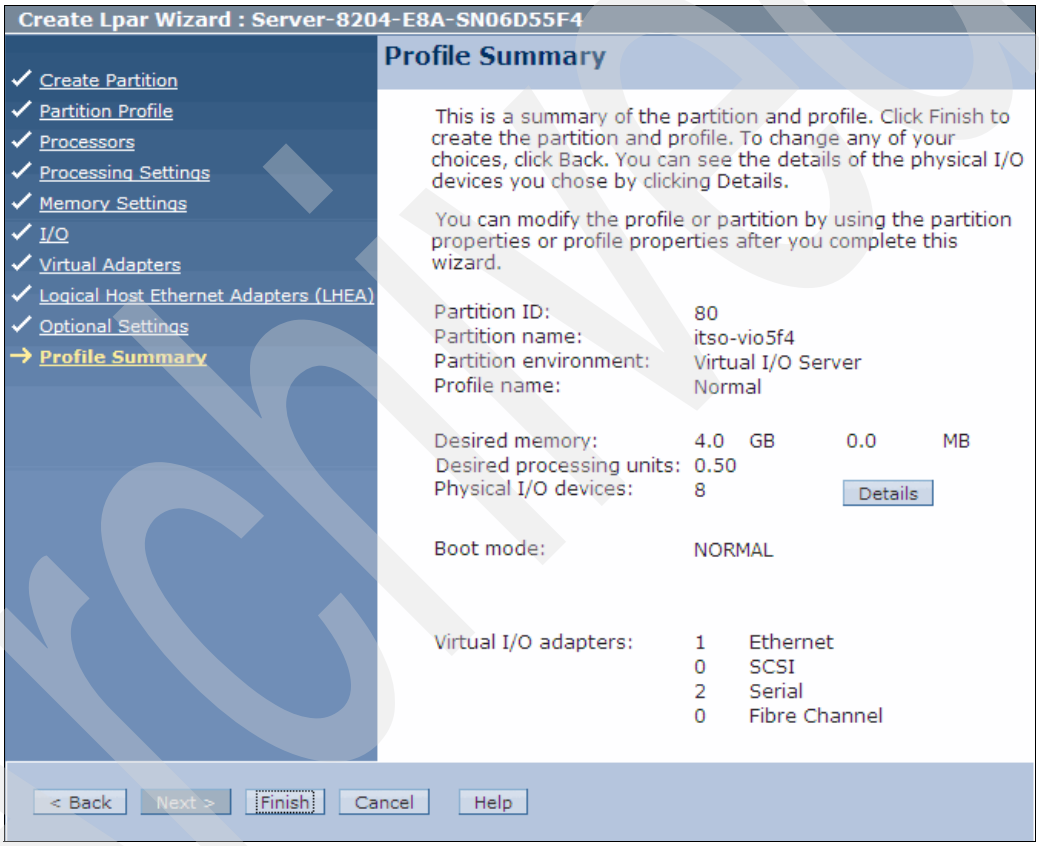

Figure 3-4 Summary of partition profile used for VIOS

3.4 Virtual I/O Server configuration

When installing the VIO server for use with WebSphere CloudBurst Appliance, the minimum supported level is 2.1.1.10-fp21.

Configuring the VIOS server after installation for use with the WebSphere CloudBurst Appliance is simple, as only two components need to be configured:

- Configure SEA for external access.
- Create a storage pool.

We assume that the system administrator will configure other common VIOS operating systems items such as IP address, host name, and monitoring integration.

Shared Ethernet Adapter configuration

To configure the SEA for external access:

1. To view and discover the SEA in the VIOS, issue the following command:

   ```
   lsdev -virtual
   ```

   In our scenario our command and output looked like:

   ```
   $lsdev -virtual
   name             status      description
   [...]
   ent9             Available   Virtual I/O Ethernet Adapter (l-lan)
   [...]
   ```

2. To create the SEA mapping from the virtual adapter to the physical adapter, use the following command:

   ```
   mkvdev -sea <physical adapter> -vadapter <virtual adapter> -default <virtual
   adapter)> -default <vlan>
   ```

   In our example we ran the command:

   ```
   $mkvdev -sea ent1 -vadapter ent9 -default ent9 -defaultid 1
   ent11 Available
   ```

   In our scenario it was previously identified that ent0 is the physical adapter and ent9 is the virtual adapter that has been granted external access. The defaultID parameter defines the VLAN ID to use. A new shared adapter was created with the name ent11.

Storage pool

WebSphere CloudBurst Appliance uses storage pools that are part of cloud groups to store deployed LPARs. The storage pools are created on the VIOS server and are presented up the software stack to the WebSphere CloudBurst Appliance. At the VIOS level a storage pool is a volume group. When the cloud group is added to the WebSphere CloudBurst Appliance, any volume groups discovered on the VIO server (including rootvg) are identified in the WebSphere CloudBurst Appliance as storage groups.

A volume group can be made of multiple disks of various types, but it is a best practice to create volume groups using the same type of disk (that is, separate volume groups for SAN, iSCSI, and internal disks). It is also a best practice to create as few volume groups as possible to reduce complexity. Remember that a running LPAR that has been deployed by the WebSphere CloudBurst Appliance has very little I/O.

Estimate storage space for PowerVM cloud

To estimate the required storage for the PowerVM cloud use the following formula:

```
Storage pool size (in MB) = size of deployed virtual image (MB) x (deployed LPARs
+ stored LPARs + extended LPARs)
```

Where:

- Size of deployed virtual image: In V2.0 of the WebSphere CloudBurst Application the deployed virtual image size is 27,386 MB.
- Deployed LPARs: the estimated maximum number of logical paritions that you expect to be deployed at any one time.

- Stored LPARs: the estimated maximum number of logical partitions that you expect to have stored on the system.
- Extended LPARs: the number of logical partitions to extend (usually 1 or 2).

Configuring storage pool on Virtual I/O Server

To set up storage pools for WebSphere CloudBurst Appliance in the VIOS partition follow these steps:

1. Log in as *padmin* to the VIOS.
2. Enter the **lspv** command to find which physical volumes are free (Example 3-1). In Example 3-1, hdisk0 and hdisk1 are part of the rootvg volume group. We already know that the physical volumes hdisk4 through hdisk8 are of the same type. If you are unsure of your storage type, consult with your storage administrator.

   Example 3-1 Listing of physical volumes

   ```
   # lspv
   hdisk0          00cd55f4e93fada3                    rootvg          active
   hdisk1          00cd55f496e66d03                    rootvg          active
   hdisk4          00cd55f49acc3dce                    None
   hdisk5          00cd55f49f1ca52a                    None
   hdisk6          00cd55f49f1ca62a                    None
   hdisk7          00cd55f49f1ca702                    None
   hdisk8          00cd55f49f1ca7a4                    None
   ```

3. Enter the following command:

   ```
   mkvg -vg <volume group name> <hdiskX> .....<hdiskN>
   ```

   In Example 3-1 we used the following command:

   ```
   mkvg -vg cloudvg hdisk4 hdisk5 hdisk6 hdisk7 hdisk8
   ```

Storage pool setup is now complete.

3.5 Network Installation Manager configuration

This section describes the required steps to create the NIM server for use with the WebSphere CloudBurst Appliance. The tasks are:

- The required filesets installation
- The NIM master configuration

3.5.1 Required filesets for installation

This section describes the filesets required for the NIM server.

AIX prerequisites check

Take the following steps before proceeding with the NIM Master installation:

1. Log in as root and check the operating system version. Issue the **`oslevel -s`** command.

   The correct level is 6100-03-01-0921, which is required for WebSphere CloudBurst Appliance.

   > **Note:** AIX 6.1 TL 3-SP1 is the only supported AIX version for the NIM and IBM Systems Director when used in the WebSphere CloudBurst 2.0 software stack.

2. Verify that the host resolution works properly (forward and reverse DNS host resolution). The NTP time synchronization protocol setup is required.
3. Verify the version of OpenSSL and OpenSSH filesets. Always install the version recommended by the IBM Systems Director installation. A back-level or incorrectly installed OpenSSL or OpenSSH fileset can create deployment failures.

   > **Note:** We used these versions in our scenario:
   >
   > - OpenSSL 0.9.8.803
   > - OpenSSH 5.0.0.5302
   >
   > For detailed information about OpenSSL and OpenSSH versions, see the topic "Preparing to install IBM Systems Director Server on AIX" in the IBM Systems Director 6.1.x Information Center at:
   >
   > http://publib.boulder.ibm.com/infocenter/director/v6r1x/index.jsp?topic=/director.install_6.1/fqm0_t_preparing_to_install_ibm_director_on_aix.html

WebSphere CloudBurst Appliance related software

The following filesets have to be installed:

- `dsm.core`

  Install this fileset manually from AIX 6.1 Technology Level 03, Service Pack 01 Base media.

  > **Note:** The fileset `dsm.core` is not on the base media below AIX 6.1.3.0.

- `DirectorCommonAgent`

  Install this fileset from the IBM Systems Director 6.1.0.3 Common Agent for AIX and VIOS (Installp format) package. Follow these steps to download it:

  a. Open the IBM Systems Director Agents web page, available at:

     http://www.ibm.com/systems/software/director/downloads/agents.html

  b. Under the section **IBM Systems Director Common or Platform Agents** → **Manual Installation**, select **Download Version 6.1** and continue.

  c. Select **IBM Systems Director 6.1.0.3 Common Agent for AIX and VIOS (Installp format)** on the next page and download it.

  d. Store the DCA installation package, previously downloaded, to the temporary file system (for example, `/inst/dca`) and extract it:

     ```
     #cd /inst/dca
     #gzip -cd package_name | tar -xvf -
     ```

e. Run SMIT installation:

```
#smitty install_latest
```

f. Go through the SMIT screens:

i. On the INPUT device/directory for software window, between the brackets, type `/inst/dca`, then press Enter.

ii. On the SOFTWARE to Install window, press the F4 key. This provides a list of software to install. Using the F7 key, select DirectCommonAgent, then press Enter.

iii. For Accept new license agreements?, select **Yes**, and then press Enter twice.

- Wait for the installation to complete, then exit SMIT.

NIM Master related software

To create the NIM Master server the following filesets have to be installed:

- `bos.sysmgt.nim.master`
- `bos.sysmgt.nim.spot`

Install these filesets manually from AIX 6.1 Base media, then update it to Technology Level 03, Service Pack 01.

3.5.2 NIM Master configuration

Now you can set up your server as NIM Master. Set up the required operating system parameters and then set up NIM Master itself, as described in this section.

Operating system setting

Change the below network parameters using the **no** command:

```
#no -p -o tcp_recvspace=524288
#no -p -o tcp_sendspace=524288
#no -p -o udp_recvspace=655360
#no -p -o udp_sendspace=65536
#no -p -o arpqsize=128
#no -p -o rfc1323=1
#no -p -o sb_max=8388608
```

Note: The network parameter arpsize should be greater than 64. We used 128 in our scenario.

Change the maximum number of user processes using **chdev** command:

```
#chdev -l sys0 -a maxuproc=1024
```

Note: The maximum number of processes per user should be greater than 1024.

Create a separate file system for the `/export/nim` path:

```
#/usr/sbin/crfs -v jfs2 -g rootvg -a size=55000 -m /export/nim
```

Note: The minimum size of `/export/nim` should be at least 55 GB. `/export/nim` is used to store the virtual appliances used for deployment. It might be possible to have several virtual appliances stored in `/export/nim` at any one time.

NIM Master simple setup

This section describes the setup of the NIM Master, as required for the WebSphere CloudBurst Appliance. The NIM Master configuration is very basic. Creating a more advanced NIM configuration does not work with the IBM Systems Director VMControl™ components of the software stack.

To create the NIM master setup:

1. Log in as the root user.
2. Issue the SMIT command **`smit nimconfig`**.
3. Enter the following parameters:
   - For the network name, enter `master_net`.
   - For the primary network install interface, select the appropriate Ethernet adapter that belongs to RMC/deploy private network.
4. Press Enter.

Example 3-2 illustrates the NIM master configuration menu. The private network is en0 in this example.

Example 3-2 smit nimconfig command menu

```
Configure Network Installation Management Master Fileset

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

                                                        [Entry Fields]
* Network Name                                         [master_net]
* Primary Network Install Interface                     [en0]                   +

  Allow Machines to Register Themselves as Clients?  [yes]                   +
  Alternate Port Numbers for Network Communications
       (reserved values will be used if left blank)
    Client Registration                                 []                       #
    Client Communications                               []                       #

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

Important: Do not use NIM Easy Setup or similar tasks. The steps described above are sufficient for the WebSphere CloudBurst Appliance solution.

In addition, the selection of the correct Ethernet interface for the NIM Master is important in case you are using more than one Ethernet interface on the NIM server. The interface for the RMC/deploy network must be used, not the interface for system management.

4

Installing IBM Director and VMControl

This chapter describes the installation and configuration of IBM Systems Director Version 6.1.1.2 and VMControl plug-in Version 2.1. This chapter discuses all important steps, from software installation to VMControl configuration.

The installation of the operating system is not covered in this chapter. Appropriate AIX installation at the supported level is a prerequisite for these tasks. For the AIX level used in this scenario, refer to Table 3-9 on page 29.

Important: The reader is strictly advised to follow the steps as written to avoid potential installation and configuration problems. The written steps apply to the WebSphere CloudBurst Appliance Version 2.0 only.

The following topics are discussed:

4.1 Downloading installation packages

This section describes how to obtain appropriate installation packages for IBM Systems Director Version 6.1 and VMControl plug-in Version 2.1, plus the required updates for them. You can download them either directly to the temporary installation directory on the director server or to a workstation and then upload it to the server.

To download IBM Systems Director and VMControl plug-in base filesets follow these steps:

1. Open the IBM Systems Director Plug-ins web page, available at:

   http://www.ibm.com/systems/software/director/downloads/plugins.html

2. Under the IBM Systems Director VMControl section, select **Download Version 2.2** and continue.
3. Select **IBM Systems Director VMControl 2.1** and **AIX and IBM Systems Director 6.1 Management Server for AIX (must be updated to 6.1.1.2 to support VMControl 2.1 for AIX)** on the next page and download them.

To download DirectorCommonAgent 6.1.0.3 follow these steps:

1. Open the IBM Systems Director Agents web page, available at:

   http://www.ibm.com/systems/software/director/downloads/agents.html

2. Under the IBM Systems Director Common or Platform Agents → Manual Installation section, select **Download Version 6.1** and continue.
3. Select **IBM Systems Director 6.1.0.3 Common Agent for AIX and VIOS (Installp format)** on the next page and download it.

To download IBM Systems Director update follow these steps:

1. Open the IBM Fix Central web page, available at:

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

2. Select **IBM Systems Director**, version **6.1** and continue.
3. Find and select **fix pack: SysDir6_1_1_2** on the next page and download it.

> **Important:** There is also a 2.1.0.1 patch collection for VMControl (group: com.ibm.director.vmcontrol.collection_2.1.0.1) on the download page. Do *not* download it or install it. These fixes interfere with WebSphere CloudBurst Appliance interim fix and the cloud deploy does not work.

To download WebSphere CloudBurst Appliance interim fix for VMControl follow these steps:

1. Log in to the FTP server at `testcase.boulder.ibm.com`. Use `anonymous` for the user ID and your email address as the password.
2. Go to the `/fromibm/websphere` directory.
3. Download these files:
   - `WebSphereCloudBurstVMC2.1fix.tar`
   - `WebSphereCloudBurstVMC2.1readme.txt`

4.2 Installing IBM System Director

We perform the IBM Systems Director product installation in accordance with the product installation guides, which are available online at the IBM Systems Director 6.1.x Information Center at IBM web page:

http://publib.boulder.ibm.com/infocenter/director/v6r1x/index.jsp

The steps in this section are only for the installation flow required for WebSphere CloudBurst Appliance. They are not a replacement of product standard documentation.

4.2.1 Base product installation and configuration

The product installation and configuration consist of the following parts:

1. AIX prerequisites check
2. Product filesets installation
3. Basic configuration

All mentioned previously, parts are covered in detail in the next three sections.

AIX prerequisites check

Provide the following checks before the IBM Systems Director installation:

1. Log in as root and check the operating system version. Issue the **`oslevel -s`** command.

   The correct level is 6100-03-01-0921, which is required for WebSphere CloudBurst Appliance.

2. Verify that the host resolution works properly (forward and reverse DNS host resolution). Set up the NTP time synchronization protocol.
3. Verify the file system size. A minimum of 10 GB of free space in the `/tmp` file system and 1 GB of free space in the `/` file system are required with WebSphere CloudBurst Appliance.

   Note: For detailed information about file system requirements, see the “Recommended hardware requirements for IBM Systems Director Server running on AIX” topic in the IBM Systems Director 6.1.x Information Center at:

   http://publib.boulder.ibm.com/infocenter/director/v6r1x/index.jsp?topic=/director.plan_6.1/fqm0_r_hardware_requirements_servers_running_aix.html

4. Verify the version of the OpenSSL and OpenSSH filesets. We strongly suggest having the correct versions installed before IBM Systems Director installation.

   Note: We use these versions in our scenario:

   - OpenSSL 0.9.8.803
   - OpenSSH 5.0.0.5302

   For detailed information about OpenSSL and OpenSSH versions, see the “Preparing to install IBM Systems Director Server on AIX” topic in the IBM Systems Director 6.1.x Information Center at:

   http://publib.boulder.ibm.com/infocenter/director/v6r1x/index.jsp?topic=/director.install_6.1/fqm0_t_preparing_to_install_ibm_director_on_aix.html

Product filesets installation

The product filesets installation can be done either from DVD media or from a downloaded installation package. Follow these instructions to install the product from the downloaded package:

1. Log in as root to the director server operating system.
2. Store the director base installation package, previously downloaded, on the temporary file system (for example, `/inst/isd`) and extract it:

   ```
   #cd /inst/isd
   #gzip -cd package_name | tar -xvf -
   ```

3. Copy the `DirServer.bnd` bundle file to `/usr/sys/inst.data/user_bundles`:

   ```
   #cp /inst/isd/DirServer.bnd /usr/sys/inst.data/user_bundles
   ```

4. Run the **`smitty install`** command and:
   a. Select **Install and Update Software**.
   b. Select **Install Software Bundle**.
   c. Select **INPUT device / directory for software**, enter `/inst/isd` into the bracket, and then press Enter.
   d. Select a fileset bundle, select **DirServer**, and then press Enter.
   e. For Accept new license agreements?, select **Yes**, and then press Enter twice.
5. Wait until the installation completes, then exit SMIT.

> **Note:** Be patient. The installation process can take 40 minutes.

To perform a DirectorCommonAgent (CAS/DCA) update, follow these steps:

1. Store the DCA installation package, previously downloaded, to the temporary file system (for example, `/inst/dcaupd`) and extract it:

   ```
   #cd /inst/dcaupd
   #gzip -cd package_name | tar -xvf -
   ```

2. Run a SMIT update:

   ```
   # smitty update_all
   ```

3. Go through the SMIT screens:
   a. On the INPUT device/directory for Software window, between the brackets enter /inst/dcaupd , then press Enter.
   b. For Accept new license agreements?, select **Yes**, and then press Enter twice.
4. Wait until the installation completes, then exit SMIT.

> **Note:** Be patient. The update process can take 30 minutes.

Basic configuration

This section describes the steps needed for the basic product configuration. First, the configuration script must be run as follows:

1. Issue the `/opt/ibm/director/bin/configAgtMgr.sh` script.
2. Answer the script questions as follows:
   a. Enter the Resource Manager user ID that you want to set for your Agent Manager: `usmi`
   b. Enter the Resource Manager password to set for your Agent Manager: `usmi`
   c. Verify the Resource Manager password to set for your Agent Manager: `usmi`
   d. Enter the Agent Registration password to set for your Agent Manager: `usmi`
   e. Verify the Agent Registration password to set for your Agent Manager: `usmi`
   f. Would you like to use an existing Agent Manager (yes or no): `no`

   Note: Use your own user ID and password. usmi is only the example we used in the scenario.

3. Wait until the script completes. Example 4-1 shows the end of the successful script output.

Example 4-1 End of the successful script output

```
Configuring Resource Manager and Common Agent.
Starting Agent Manager... OK
Retrieving Agent Manager Instance ID... OK
Registering Agent Manager toolkit... OK
Removing Agent Manager user manager... OK
Adding Agent Manager user usmi... OK
Configuring Common Agent... OK
Waiting for Common Agent certificates... OK
Waiting for Common Agent SLP advertisement... OK
Waiting for Common Agent status report... OK
Stopping Agent Manager... OK
Agent Manager configuration completed successfully.
#
```

After the successful configuration, you can start the director server following these steps:

1. Log out from your SSH session on director operating system and then log in again.

   Note: The new login process refreshes the shell search path. Otherwise, the IBM Systems Director **`smstart`**, **`smstop`**, and **`smstatus`** commands have to be preceded by the path information (`/opt/ibm/director/bin`).

2. Start the IBM Systems Director server, issue the **`smstart`** command, and then issue the **`smstatus -r`** command. Example 4-2 illustrates the command output.

Example 4-2 smstart and smstatus command output

```
# smstart
Starting IBM Director...
The starting process can take a while. Use smstatus to check if the server is
active.
#
# smstatus -r
Inactive
```

```
Starting
Active
#
```

3. Wait until the director is in the active state, which can take a while. Press Ctrl+C to escape from the command operation after `Active` appears.

> **Note:** You can also use the **`smstatus`** command without the -r parameter. This shows the actual status and returns you to the command line immediately.
>
> Before the inactive status you might briefly see the error status. The error status changes within 30 seconds to the inactive statu, unless there is a permanent startup error.

Now you can log into the IBM Systems Director server and provide the initial tasks:

1. In the operating system, create a new user, which will be used as the main admin user for VMControl (for example, vmcadmin). The primary group for this user has to be smadmin. Create new password for the user and log into it.
2. Open a web browser and enter the following URL:

   `https://x.x.x.x:8422/ibm/console`

   Where *x.x.x.x* is your director server IP address.

   > **Note:** A website certification dialog box displays. Follow the users' guide of your browser to accept the certificate.
   >
   > For supported browsers, see the "Supported Web browsers AIX" topic in the IBM Systems Director 6.1.x Information Center at:
   >
   > http://publib.boulder.ibm.com/infocenter/director/v6r1x/index.jsp?topic=/director.plan_6.1/fqm0_r_supported_web_browsers.html

3. Log in using the user ID and password created in the step 1.

The Welcome window of IBM Systems Director now displays (Figure 4-1).

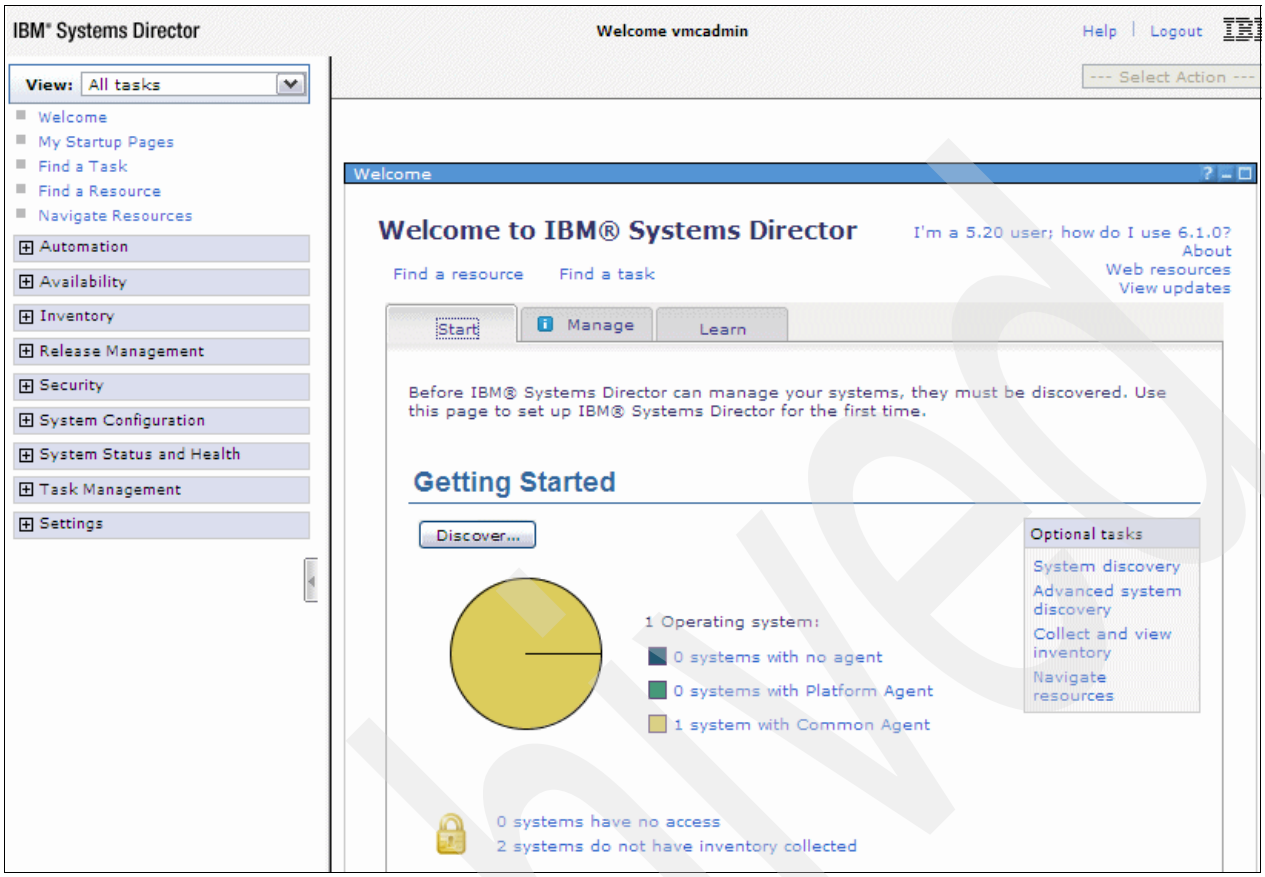

Figure 4-1 IBM Systems Director Welcome window

4. Within the Welcome window, click **Navigate Resources** in the left pane, and then **All Operating Systems** in the main window (Figure 4-2).

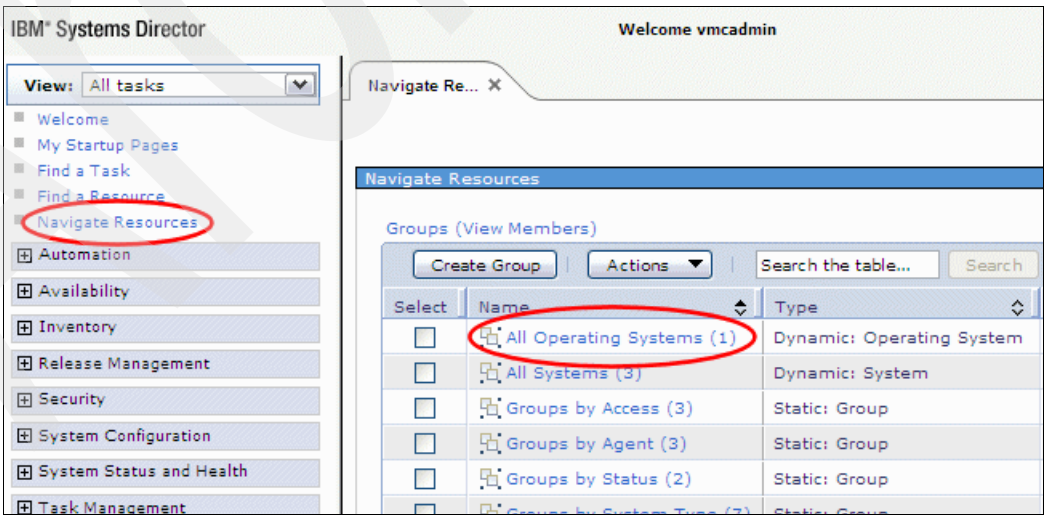

Figure 4-2 Navigate Resources option in Welcome window

The director server system displays and the Access column shows `OK` (Figure 4-3).

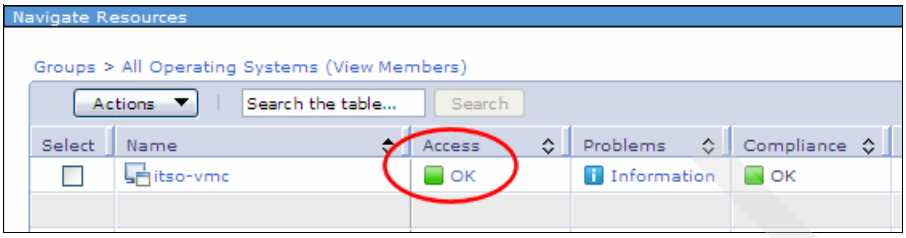

Figure 4-3 Discovered director server system

Note: If the Access column shows `No access` for the server name, click the **No access** status, enter the user ID and password, and click **Request Access**. Now the Access column entry indicates `OK`. Click **Close** to return to the list.

If no system appears in the Navigate Resources pane, you have to discover it:

a. Go to the Welcome window and click the **System Discovery** option under Discovery Manager task (Figure 4-4).

Figure 4-4 System Discovery at Discovery Manager

b. Select **Single system (IP address)**, enter the IP address of the director server operating system, and click **Discover** (Figure 4-5).

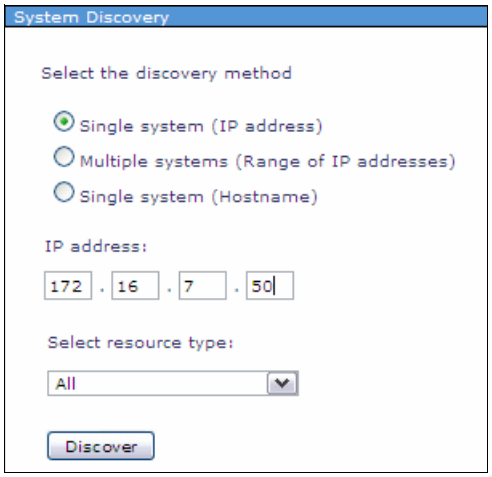

Figure 4-5 System discovery, single IP address

> **Note:** Allow the discovery process enough time to complete while the discovered systems list is being built.
>
> Do not initiate discovery of any other system except the IBM Systems Director server itself.

5. Within the discovered systems list, right-click the name of the director server, and then select **Inventory** → **View and Collect Inventory** (Figure 4-6).

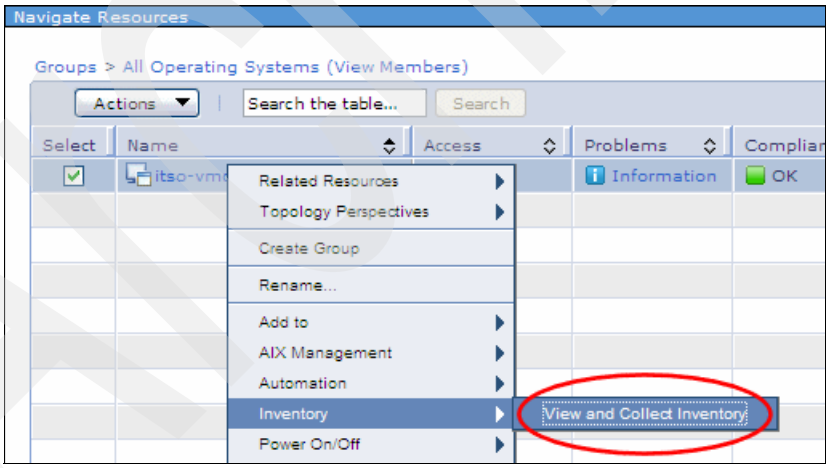

Figure 4-6 View and Collect Inventory

6. Click **Collect Inventory** (Figure 4-7). Then ensure that Run Now is selected and click **OK**:

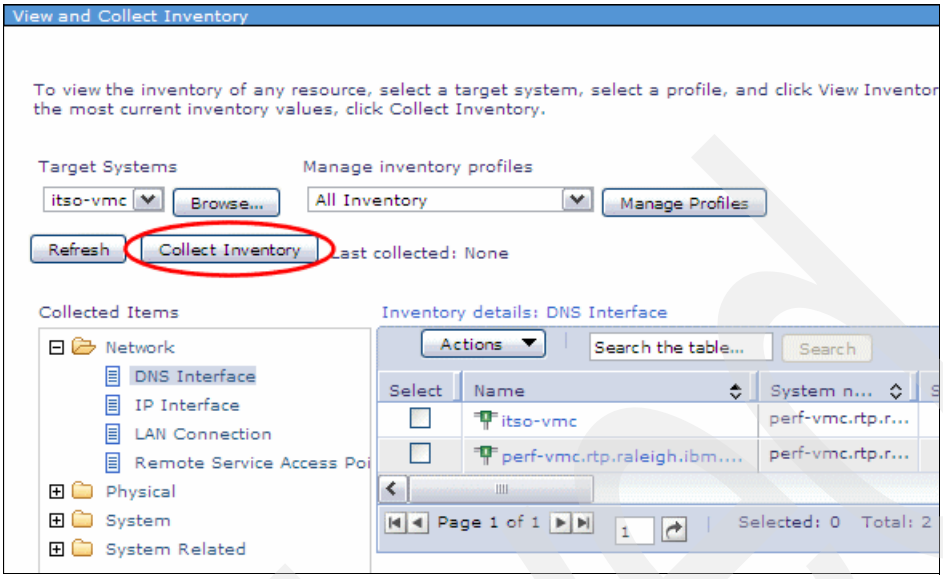

Figure 4-7 Collect Inventory

7. On the left, expand **Task Management** and click **Active and Scheduled Jobs**. You will see the collect inventory job complete with no errors (Figure 4-8). The IBM Systems Director now knows about its own resources.

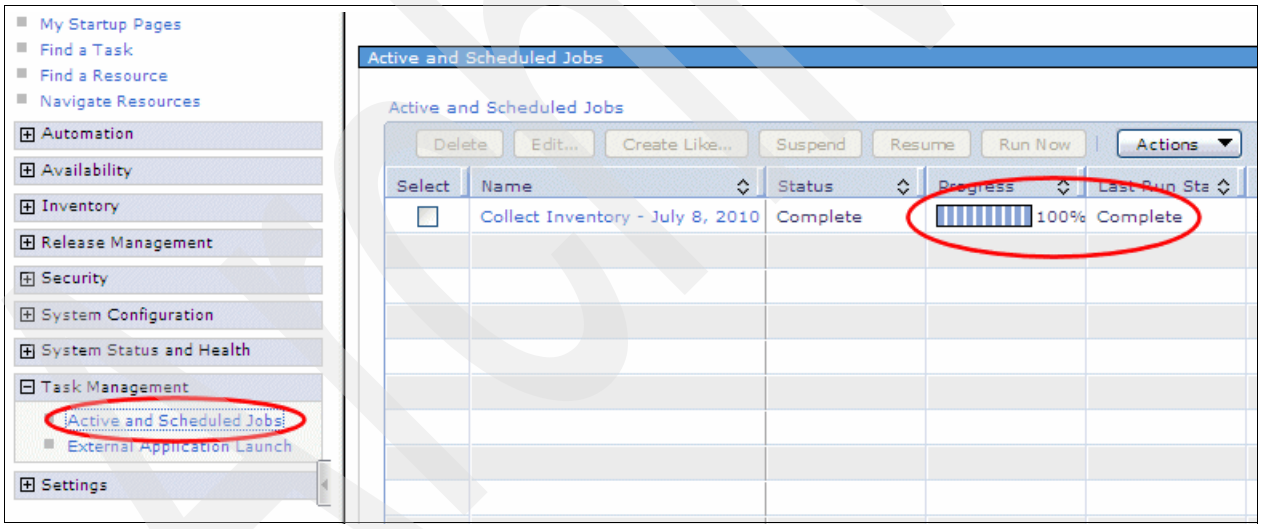

Figure 4-8 Successful collected inventory for the director server

> **Important:** At this point, do not make additional discoveries or changes to your system until maintenance is applied to the IBM System Director.

8. Log out of the director web GUI.

4.2.2 Product updates installation

The IBM Systems Director update consists of two parts:

- The command-line part
- The GUI part

Both are covered in this section.

Command-line part

Now the updates for IBM Systems Director have to be installed. Follow these instructions to install updates from the downloaded package:

1. Log in as root to the director server operating system.
2. Store the director base installation package, previously downloaded, to the temporary file system (for example, `/inst/isdupd`) and extract it:

   ```
   #cd /inst/isdupd
   #unzip package_name
   ```

   Note: The unzip tool can be downloaded from the AIX Toolbox for Linux Applications web page, available at:

   ftp://public.dhe.ibm.com/aix/freeSoftware/aixtoolbox/RPMS/ppc/unzip/unzip-5.51-1.aix5.1.ppc.rpm

   Install it with command **`rpm -ivh unzip-5.51-1.aix5.1.ppc.rpm`**.

3. Import the updates into the IBM Systems Director using the following **`smcli`** command:

   ```
   #smcli importupd –vr /inst/isdupd
   ```

4. From the command prompt list, you will see system director updates shown in Example 4-3.

Example 4-3 smcli lsupd command

```
# smcli lsupd | more
cfgdbcmd.feature_1.0.1
com.ibm.director.agent.helps.doc.nl1_6.1.1
com.ibm.director.agent.helps.doc_6.1.1.1
com.ibm.director.automation.collection.feature_6.1.1.1
com.ibm.director.automation.helps.doc.nl1_6.1.1
com.ibm.director.automation.helps.doc_6.1.1.1
com.ibm.director.base.feature_6.1.1.1
com.ibm.director.bcxconfigmanager_6.1.1.1
com.ibm.director.configmanager_1.0.1.1
com.ibm.director.configuration.helps.doc.nl1_6.1.1
com.ibm.director.configuration.helps.doc_6.1.1.1
com.ibm.director.configurationmanager_6.1.1.1
com.ibm.director.console.automation.feature_6.1.1.1
com.ibm.director.console.configmanager_1.0.1.1
com.ibm.director.console.helps.doc.nl1_6.1.1
com.ibm.director.console.helps.doc_6.1.1.1
com.ibm.director.console.lic_6.1.1.1
com.ibm.director.console.welcomeExt_6.1.1
com.ibm.director.core.manager_6.1.1.1
com.ibm.director.core.manager_6.1.1.2
```

```
com.ibm.director.discovery.base.feature_6.1.1.1
com.ibm.director.discovery.helps.doc.nl1_6.1.1
com.ibm.director.discovery.helps.doc_6.1.1.1
Standard input
```

5. Install two main director updates. Issue the listed **smcli** commands:

   ```
   #smcli installupd –i x.x.x.x –u com.ibm.director.core.manager_6.1.1.1
   #smcli installupd –i x.x.x.x –u com.ibm.director.core.manager_6.1.1.2
   ```

   Where *x.x.x.x* represents the IP address of the director server.

   > **Note:** We obtain the following warning message during the second **installupd** issued:
   >
   > `ATKUPD767W – Update "com.ibm.tivoli.tpm.common.feature_5.1.1" installed failed for system` *`your system name`*
   >
   > This error does not have any negative effects and you can ignore it.

6. Stop and restart the IBM Systems Director server:

   ```
   #smstop
   #smstart
   #smstatus –r
   ```

   Wait until `Active` appears, then press Ctrl+C to exit from the command.

Now you are ready to go to the next part of the update.

Director GUI part

Finish updating IBM System Director as described in these steps:

1. Log in to the director web GUI again using your user ID and password.
2. Select **Release Management** → **Updates** in the left pane (Figure 4-9).

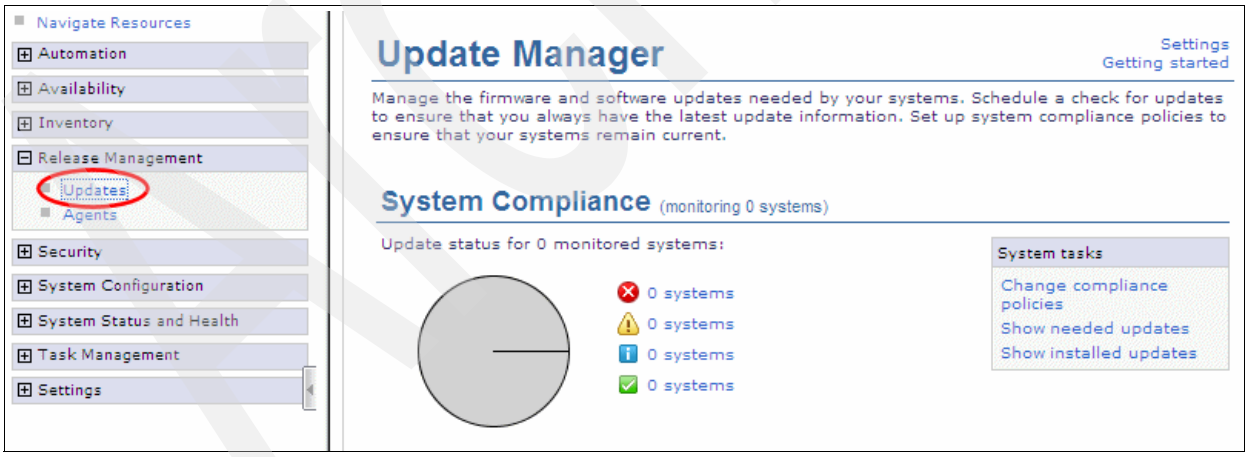

Figure 4-9 Update Manager window

The pie chart under Systems Compliance shows that `0 systems` are compliant.

3. Under the System Compliance group, select **Change compliance policies** (Figure 4-10).

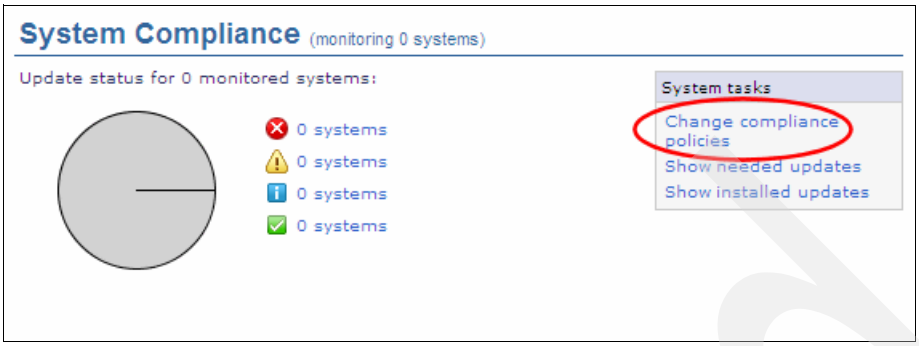

Figure 4-10 Change compliance policies

a. Next to the Select system pull-down menu, click **Browse**.

b. On the Context Chooser window, click **Browse** to select your director server. The 0Type must be Operating System. Click **Add** (Figure 4-11). The itso-vmc system is the director server in our scenario.

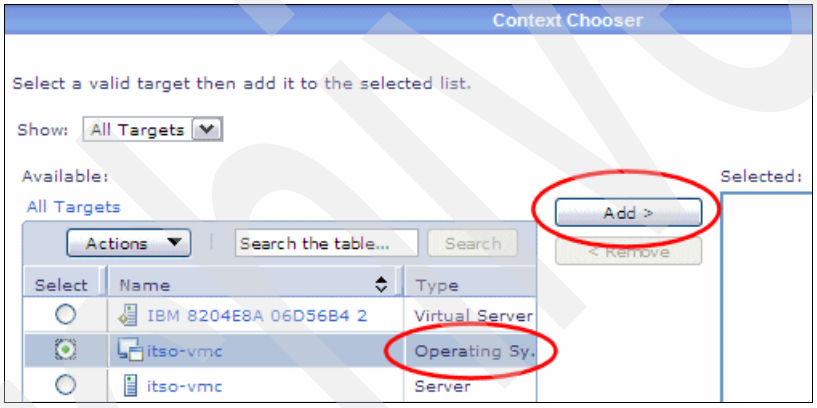

Figure 4-11 Select the director operating system

c. Return to the previous menu by clicking **OK**.

d. Under the Update compliance policy section, click **Show Compliance Policy**, and then click **Add** (Figure 4-12).

Figure 4-12 Add system to policy

e. Locate All IBM Systems Director 6.1 Updates under the Update Groups listing, check the selection box, and click **Add** (Figure 4-13).

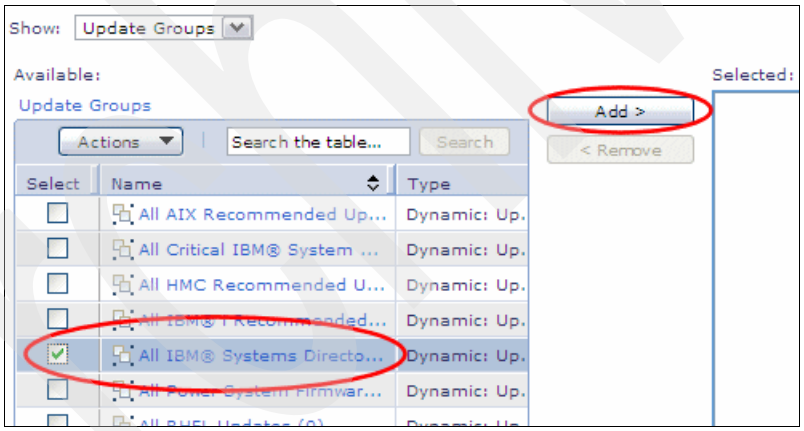

Figure 4-13 Add directors updates

f. Verify that the **Ensure that all applicable updates are installed** radio button is selected and click **OK** (Figure 4-14).

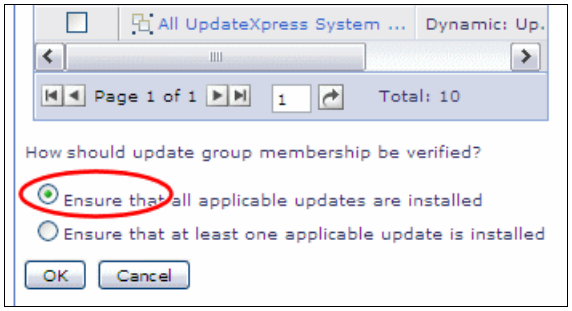

Figure 4-14 Commit the update selection

g. Select the **IBM Systems Director** update group and click **Save** (Figure 4-15).

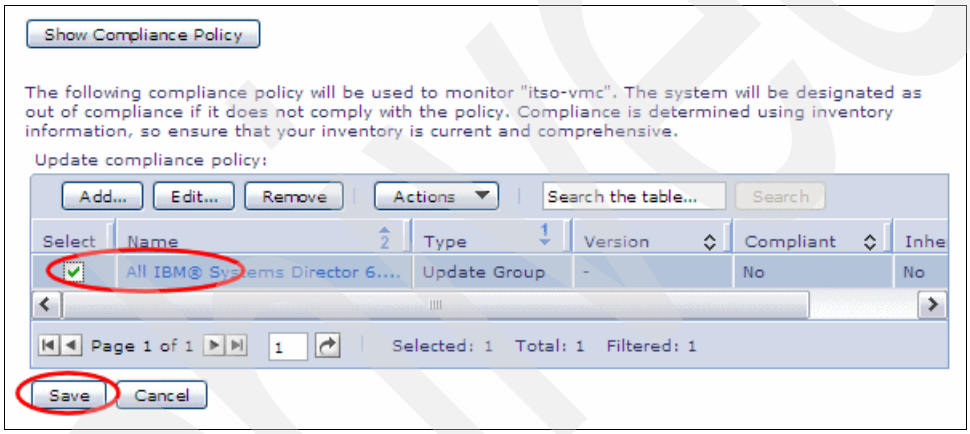

Figure 4-15 Select the director update group

4. From the Task menu, again select **Release Management** → **Updates**. The pie chart under Systems Compliance is colored red. This is because the monitored system is out of compliance with the update policy.

> **Note:** If the pie chart happens to be colored green, wait a couple of minutes, close the window, and click **Release Management** → **Updates** again. If the pie chart remains green, consider logging out of the IBM System Director, and stopping (**`smstop`**) and restarting (**`smstart`**) the director.
>
> If it is still green, re-import the update packages. Click **Import updates** under the Common tasks shaded box, then select the directory (for example `/inst/isdupd`).

5. From the Update Manager window, select **Show needed updates** (Figure 4-16).

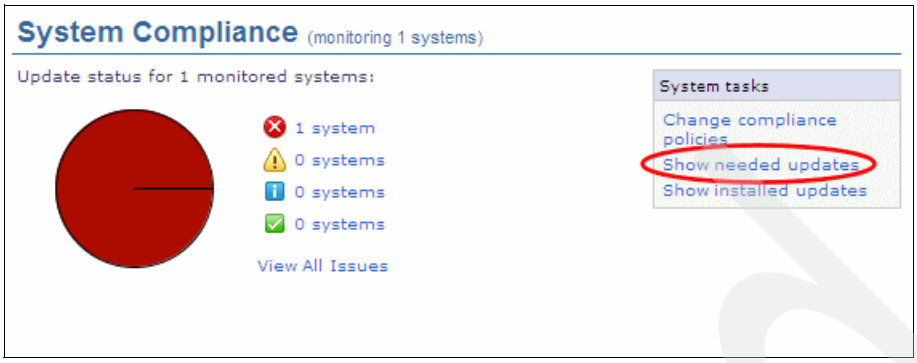

Figure 4-16 Show needed update selection

6. Select your IBM Systems Director server (Operating System), and click **Show Needed Updates** (Figure 4-17).

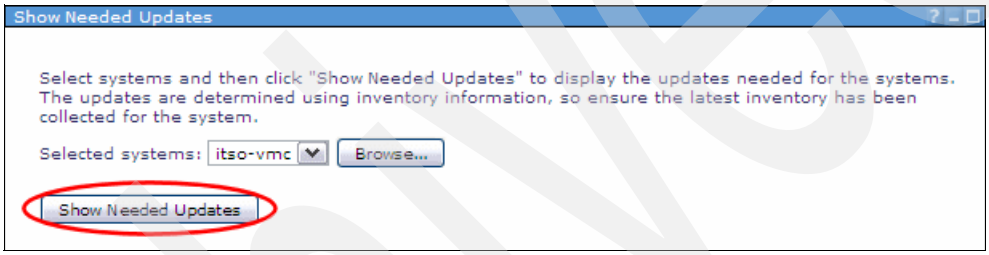

Figure 4-17 Show needed update button

7. From the Show Needed Updates window, select the **Actions** pull-down menu and choose **Select All**. From the **Actions** pull-down menu, click the **Install** menu (Figure 4-18) to initiate the installation wizard.

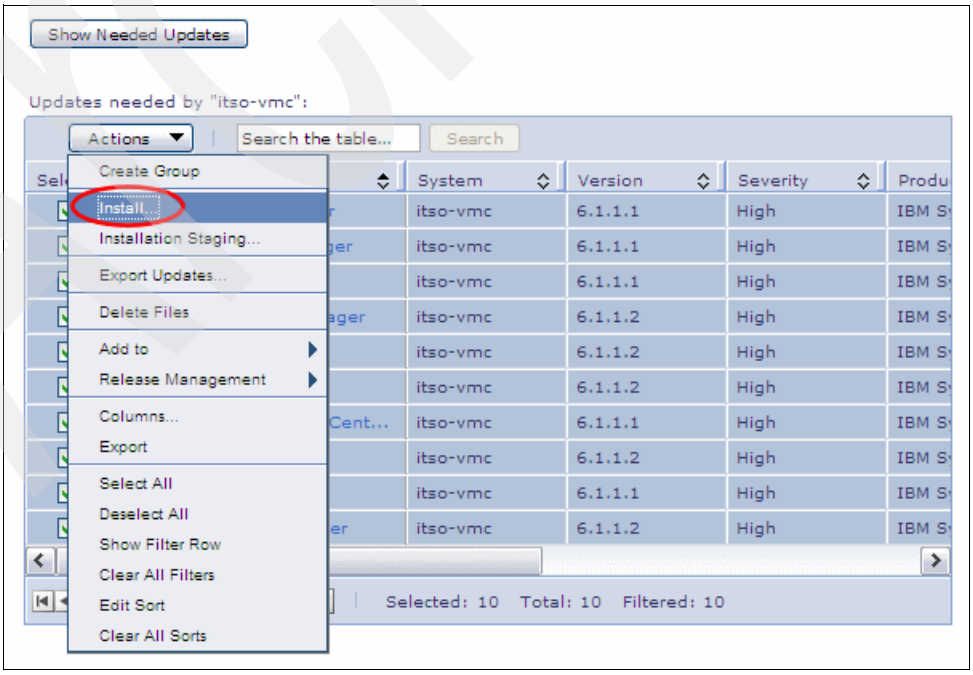

Figure 4-18 Updates install menu

8. Follow the installation wizard through these options:
   a. At the Welcome step, click **Next**.
   b. At the Systems step, select your director server (Operating System) from the list, click **Add**, and then click **Next**.
   c. At the Options step, verify that the **Automatically install missing update requirements** box is selected and click **Next** (Figure 4-19).

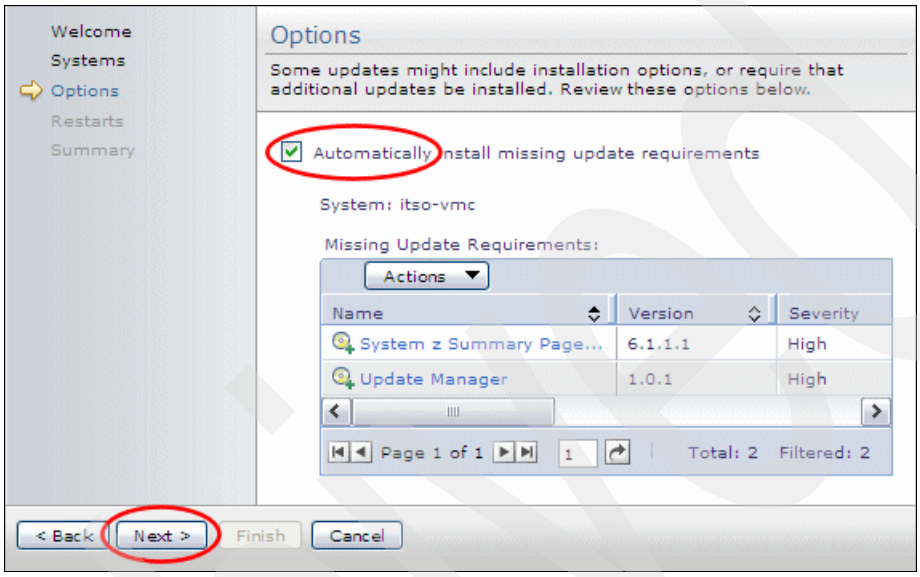

Figure 4-19 Automatically install missing update requirement selection

   d. At the Restarts step, verify that the **Automatically restart as needed during installation** box is selected and click **Next** (Figure 4-20).

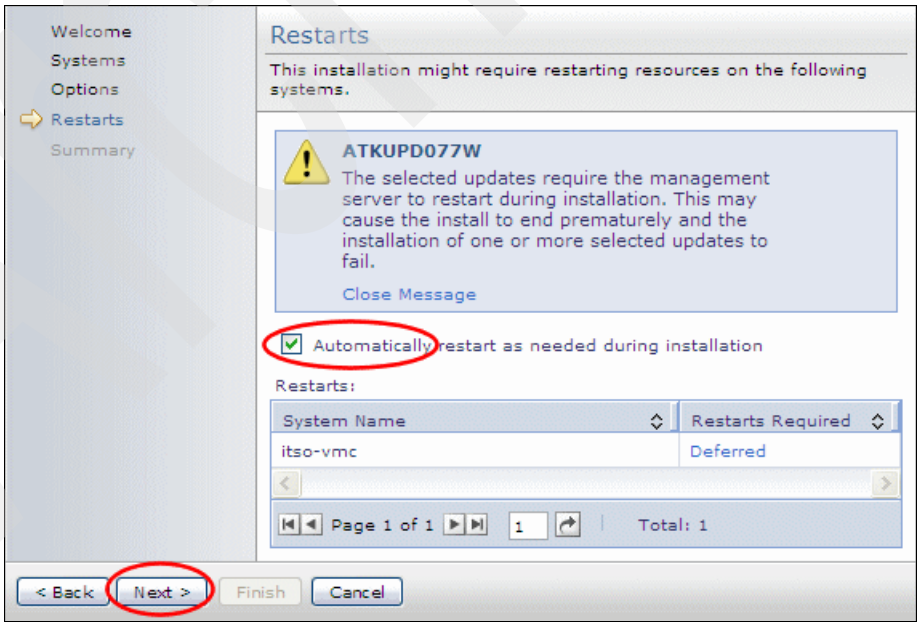

Figure 4-20 Automatically restart as needed during installation selection

   e. At the Summary step, click **Finish**.

f. On the next window, select the **Run now** option and click **OK** (Figure 4-21).

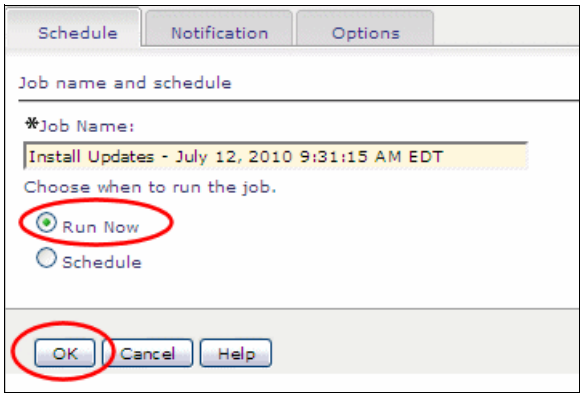

Figure 4-21 Run now option

9. Expand **Task Management** in the left pane and click **Active and Scheduled Jobs** (Figure 4-8 on page 50) to watch the status of the update job. The update step takes approximately 20 to 30 minutes, so be patient.

> **Note:** The update typically ends with a status of `Complete with Errors`, which does not indicate that the updates failed to install. Continue to the next step.

10. After the job completes, restart the IBM Systems Director Server:

```
#smstop
#smstart
#smstatus -r
```

Wait until `Active` appears, then press Ctrl+C to exit from the command.

11. From the Task menu, again select **Release Management** → **Updates**. The pie chart under Systems Compliance is colored green.

> **Note:** If the pie chart is colored red, wait a couple of minutes, close the window, and relaunch by clicking **Release Management** → **Updates** again. If the pie chart is still red after 5 to 10 minutes, consider running an inventory collection of the IBM Systems Director server. If it is still red, consider logging out of the IBM System Director and stopping (**`smstop`**) and restarting (**`smstart`**) the director.

12. Return to the Update Manager menu and select **Show Installed Updates**. You will see a list of updates installed on the system director.

Now you can continue to install the VMControl plug-in, as described in the next section.

4.3 Installing VMControl

We perform the VMControl plug-in product installation in accordance with the product installation guides, available online at the IBM Systems Director 6.1.x Information Center at the IBM web page:

http://publib.boulder.ibm.com/infocenter/director/v6r1x/index.jsp?topic=/vim_220/fsd0_vim_c_vmc_210.html

The following section provides a summary of important steps and is not a replacement to the product standard documentation.

4.3.1 Base product installation

We used the trial version of the product. Installation of the licensed product can be performed in the same manner.

To install the product from the trial downloaded package:

1. Log in as root to the director server operating system.
2. Store the VMControl installation package, previously downloaded, on the temporary file system (for example, `/inst/vmc`) and extract it:

   ```
   #cd /inst/vmc
   #gzip -cd package_name | tar -xvf -
   ```

3. Edit the following lines of the `/inst/vmc/installer.properties` file to enable the installation to run silently. Set up these parameters:

   ```
   INSTALLER_UI=silent
   LICENSE_ACCEPTED=true
   START_SERVER=false
   ```

4. Stop the director server:

   ```
   #smstop
   ```

5. Issue the VMControl installation script:

   ```
   #chmod 755 /inst/vmc/Systems_Director_VMControl_2_1_AIX.sh
   #/inst/vmc/Systems_Director_VMControl_2_1_AIX.sh
   ```

   Example 4-4 shows successful output from the VMControl installation.

   Example 4-4 Successful output from VMControl installation

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




   ===============================================================================
   Installing...
   -------------

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

   Installation Complete.
   #
   ```

6. Change to `/opt/ibm/director/VMControlManager`:

   ```
   #cd /opt/ibm/director/VMControlManager
   ```

7. Verify that the `installLog.txt` file exists:
   - If the `installLog.txt` file does not exist, then the installation is not yet complete.
   - After the file appears, review the file and check whether all the install actions are successful.
8. Start the IBM Systems Director:

   ```
   #smstart
   #smstatus –r
   ```

   Wait until `Active` appears, and then press Ctrl+C to exit from the command.
9. After the Systems Director shows an active state, log in the director web GUI using your user ID and password.
10. MControl functions can be seen quickly in two places within the console (Figure 4-22):
    - On the director Welcome window, by the plug-in named IBM Systems Director VMControl.
    - On the left pane, under the expansion of System Configuration, by the link IBM Systems Director VMControl.

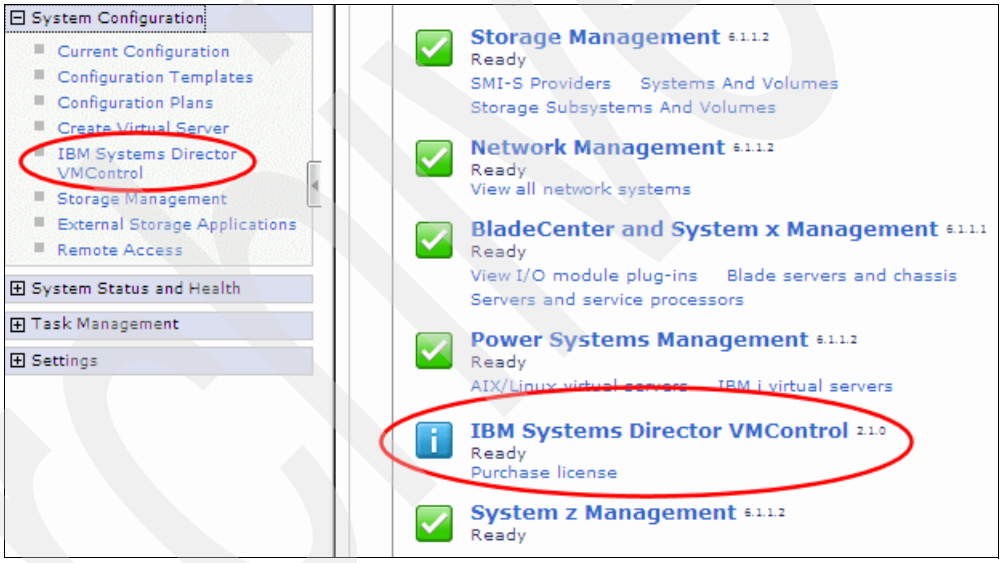

Figure 4-22 IBM Systems Director VMControl plug-in

Now continue to the VMControl updates installation, described in the next section.

4.3.2 WebSphere CloudBurst Appliance interim fix installation

We performed the WebSphere CloudBurst Appliance interim fix installation in accordance with the update `readme` file, downloaded as a part of update package.

Follow these steps to install the WebSphere CloudBurst Appliance interim fix:

1. Log in as root to the director server operating system.
2. Store the VMControl update packages, downloaded earlier, to the temporary file system (for example, `/inst/vmcwca`) and extract it:

   ```
   #cd /inst/vmcwca
   #tar -xvf package_name
   ```

Note: Keep this directory. You will need the uninstall script if you want to uninstall the interim fix.

3. Stop the director server:

   ```
   #smstop
   ```

4. Issue the VMControl installation script:

   ```
   #./hotfix_install_aix.sh
   ```

   The `EFIX-WCA-06-11-10-Version-1.0.TXT` file in the `/opt/ibm/director` directory appears in case of successful installation.

5. Look at the `/opt/ibm/director/lwi/conf/overrides/USMIKernel.properties` file and check whether the following lines are in the file:

   ```
   max.cli.threads=17
   com.ibm.director.im.rf.nim.deploynew.outputOID=true
   ```

   If not, add them at the end of the file.

6. Start the IBM Systems Director:

   ```
   #smstart
   #smstatus –r
   ```

   Wait until `Active` appears, and then press Ctrl+C to exit from the command.

Now the IBM Systems Director and VMControl are installed and ready for operation. Continue to the next section to integrate the NIM server into the director environment.

4.4 Installing agents to NIM

The VMControl software needs to communicate with the NIM server to provide, deploy, or recapture operating system images. Two agents are required for this communication, and they must be installed on the NIM server.

The WebSphere CloudBurst Appliance interim fix has to be installed on the NIM server. The installation of these agents/fixes is covered in this section.

Note: If the installation or update of any agents fails, the best way to correct it is:

1. Remove representation of the NIM server from the director.
2. Delete the `DirectorCommonAgent` fileset on the NIM server and install it again.
3. Perform all steps again from 4.4.1, “CAS/DCA agent update” on page 61.

4.4.1 CAS/DCA agent update

The DirectorCommonAgent (CAS/DCA) is installed as a part of standard AIX installation. However, the proper version of this agent is required for VMControl operation.

Perform a CAS/DCA update via the IBM System Director GUI as described in these steps:

1. Log in to the director web GUI using your user ID and password and discover the NIM server. Select **Inventory** → **System discovery** in the left pane, then enter the private IP address of your NIM server and click **Discover** (Figure 4-23).

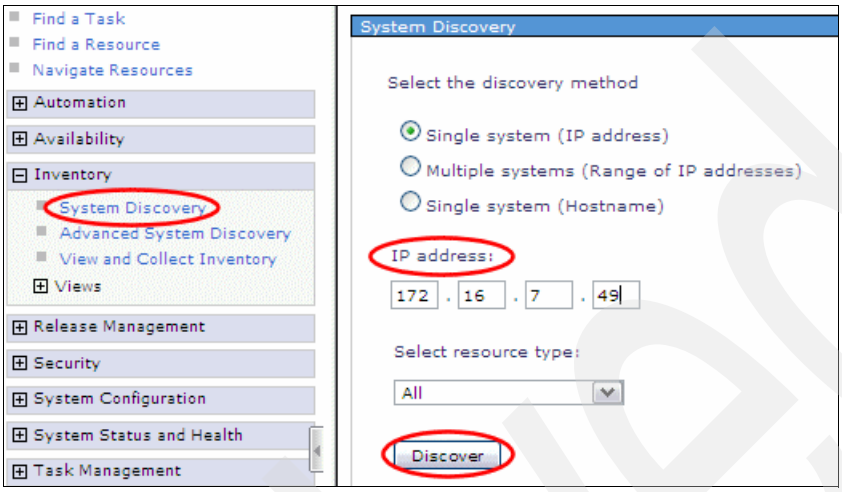

Figure 4-23 NIM server discovery

Important: The private address of your NIM server must be used for the discovery.

2. Within the discovered systems list, right-click the name of the NIM server, and then select **Inventory** → **View and Collect Inventory**.

Note: If the Access column shows `No access` for the server name, click the **No access** status, enter the user ID and password, and click **Request Access**. Now the Access column entry indicates OK. Click **Close** to return to the list.

3. Click **Navigate Resources**, and then click **All Systems** (Figure 4-24).

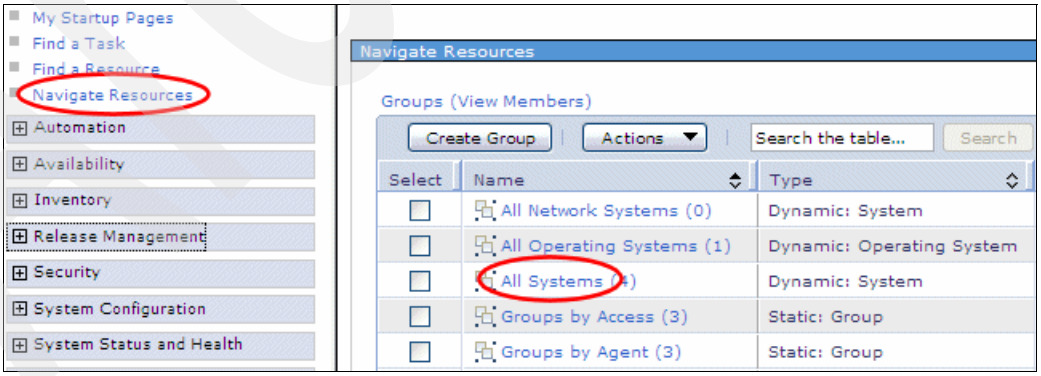

Figure 4-24 All System selection

4. Select the resource name that represents the NIM server and from the **Actions** pull-down menu choose **Release Management** → **Show needed updates** (Figure 4-25).

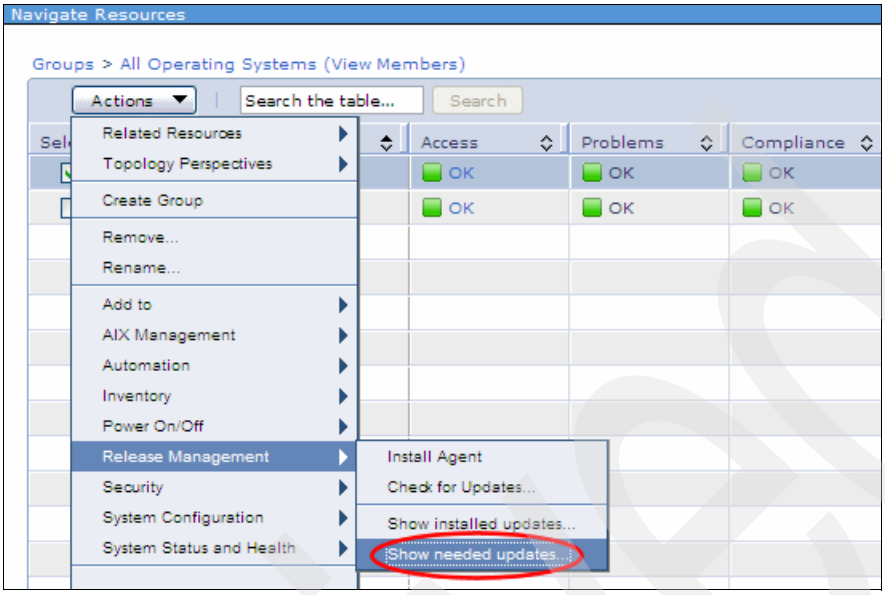

Figure 4-25 Show needed updates selection

5. From the Show Needed Updates window, select the **Actions** pull-down menu and choose **Select All**. From the Actions pull-down menu, click the **Install** menu item to initiate the installation wizard.
6. Follow the installation wizard in the same manner as described in steps 6 and 7 in "Director GUI part" on page 52.

Continue on to the VMControl subagent installation, described in the next section.

4.4.2 VMControl subagent

Perform the VMControl subagent installation via the IBM System Director GUI, as described in these steps:

1. Log in to the director web GUI using your user ID and password, select **Release Management** → **Agents** in the left pane, and click **Common Agent Subagent Packages** (Figure 4-26).

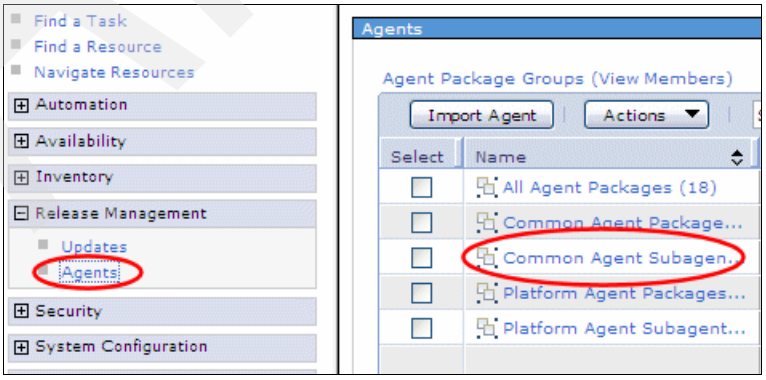

Figure 4-26 Release management agents

2. Select the **CommonAgentSubagentVMControl_NIM-2.1.0** resource name, and from the **Actions** pull-down menu choose **Release Management** → **Install Agent** (Figure 4-27).

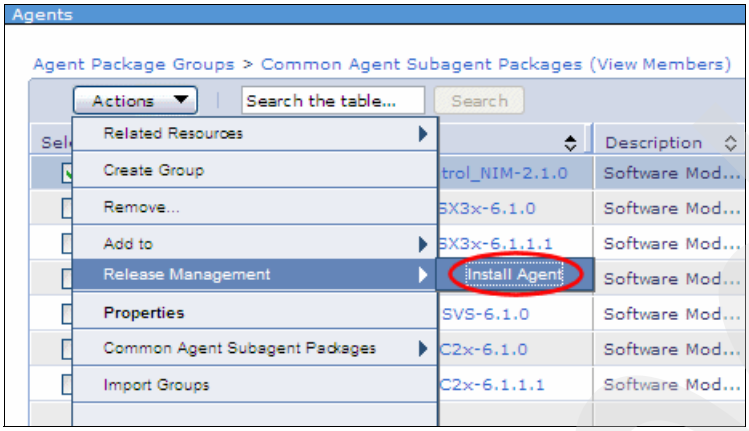

Figure 4-27 NIM subagent installation selection

3. Follow the installation wizard through these options:
   a. At the Welcome step, click **Next**.
   b. At the Agents step, the CommonAgentSubagent_VMControl_NIM-2.1.0 agent is already within the Selected window. If it is not, select that agent in the left list and click **Add** → **Next**.
   c. At the Systems step, select your NIM server from the list, and then click **Add** → **Next**.
   d. At the Summary step, click **Finish**.
   e. On the next window, select the **Run now** option and click **OK**.
4. Expand **Task Management** in the left pane, and then click **Active and Scheduled Jobs** (Figure 4-8 on page 50) to watch the status of the update job. The update step takes approximately 20 minutes, so be patient. There are two jobs to review for this activity. Both jobs must show `Complete` with no errors indicated in the Last Run Status column:
   – Install Agent
   – Agent Installation

Note: If you see any errors in the Last Run Status column, click the associated task entry, click the **Logs** tab, and review the error messages.

5. Discover the repository on your NIM server. Select **System Configuration** → **IBM Systems Director VMControl** in the left pane (Figure 4-28).

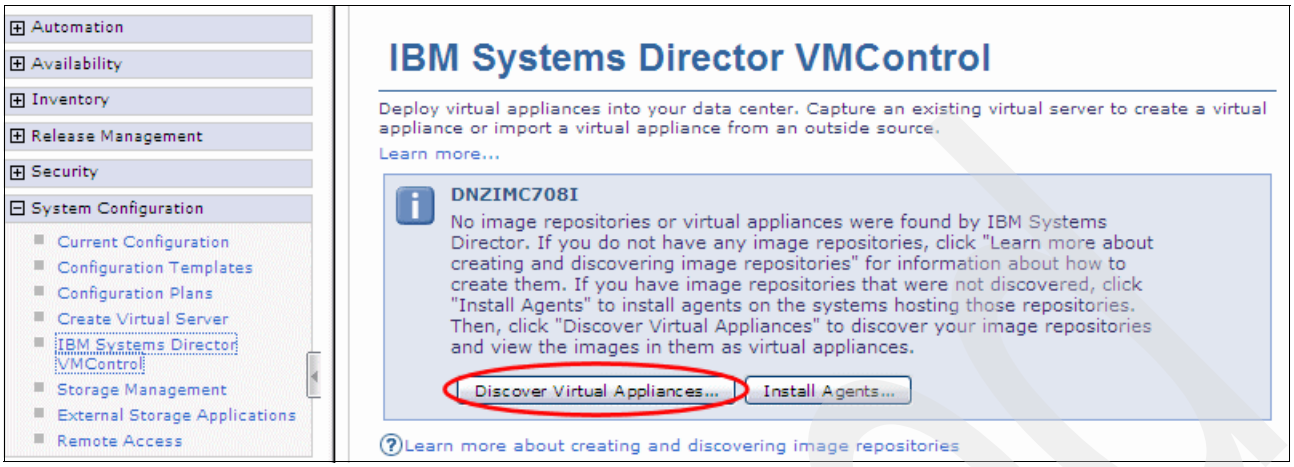

Figure 4-28 Discover virtual appliances selection

6. Locate your NIM server in the list, check the selection box, and click **Add** → **OK** button. You will see the message:

   `The following job has been created and started successfully: Discover virtual appliances – (date/time)`

7. Expand **Task Management** in the left pane and click **Active and Scheduled Jobs** (Figure 4-8 on page 50) to watch the status of the update job.

8. Return to IBM Systems Director VMControl and check the number of image repositories (it should be 1), and then click it (Figure 4-29).

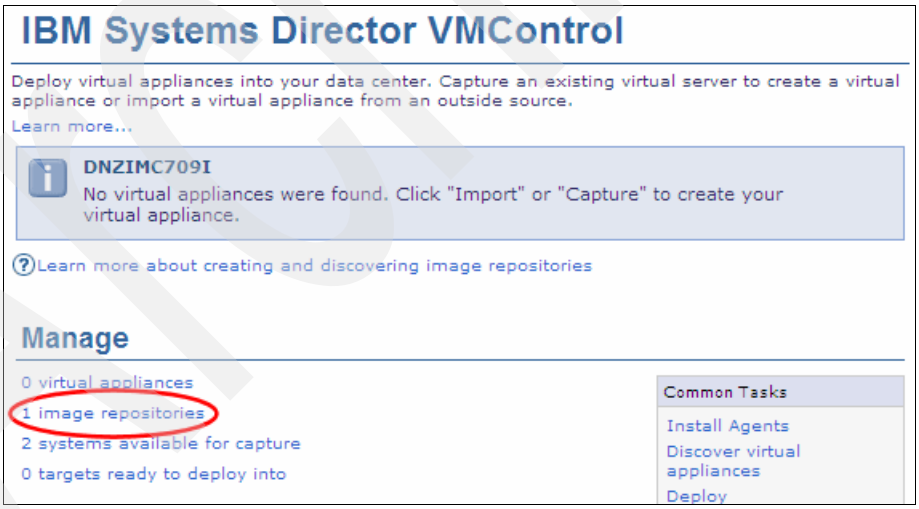

Figure 4-29 VMControl image repositories

9. If the detail of the repository contains your NIM server, the installation of the VMControl plug-in was successful (Figure 4-30).

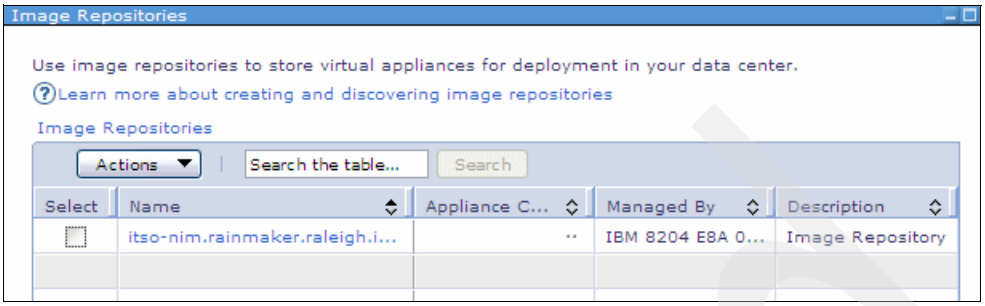

Figure 4-30 VMControl image repositories detail

Continue to the WebSphere CloudBurst Appliance interim fix installation in the next section.

4.4.3 WebSphere CloudBurst Appliance interim fix

We performed the WebSphere CloudBurst Appliance interim fix installation in accordance with the update `readme` file, which as downloaded as a part of update package.

Follow these steps to install the WebSphere CloudBurst Appliance interim fix:

1. Log in as root to the NIM server operating system.
2. Store the VMControl update packages, downloaded earlier, to the temporary file system (for example, `/inst/wcaupd`) and extract it:

   ```
   #cd /inst/vmcwca
   #tar -xvf package_name
   ```

   Note: Keep this directory. You will need the uninstall script if you want to uninstall the inerim fix.

3. Stop the director agent:

   ```
   #/opt/ibm/director/agent/runtime/agent/bin/endpoint.sh stop
   ```

4. Issue the VMControl installation script:

   ```
   #./hotfix_agent_install_aix.sh
   ```

5. Start the director agent:

   ```
   #/opt/ibm/director/agent/runtime/agent/bin/endpoint.sh start
   ```

You have completed a successful installation of the IBM Systems Director and VMControl plug-in.

4.5 Discovering systems

The last step, which has to be done before the setup will work with WebSphere CloudBurst Appliance, is to discover HMC console and Power servers.

Important: Discover only the systems mentioned in the steps below.

Follow these steps to complete the discovery process:

1. Log in to the director web GUI using your user ID and password and discover the HMC console. Select **Inventory** → **System discovery** in the left pane, then enter the private IP address of your HMC console and click **Discover**.

   Important: The private address of your NIM server must be used for the discovery.

2. Select **Navigate Resources** in the left pane, then right-click your HMC console and select **Inventory** → **View and Collect Inventory**. Click **Collect Inventory** in the next window to issue the operation.

   Note: The Access column shows `No access` for the HMC console. Click the **No access** status and enter the user ID and password for the HMC console. The user with the hscsuperadmin role is required (for example, hscroot). Click **Request Access**. Now the Access column entry indicates `OK`. Click **Close** to return to the list.

3. Expand **Task Management** in the left pane and click **Active and Scheduled Jobs** to watch the status of the update job. Wait until the operation finishes.
4. Remove all Power servers with an access status of unknown, if any. Select these servers from the Actions pull-down menu and click the **Remove** option.
5. Collect the inventory of the newly discovered resources:
   - Power servers
   - Virtual servers
6. Figure 4-31 shows how your Navigate Resources window now looks.

Figure 4-31 Successfully discovered systems

Note: You might see the HMC console twice, as the Hardware Management Console type and the operating system type. The operating system type appears if the CIM agent is set up and running on the HMC console. However, this is not required from the WebSphere CloudBurst Appliance point of view, so if you can see only the Hardware Management Console type, it is enough.

You successfully finished the IBM Systems Director and VMControl plug-in implementation. You are ready to work with the WebSphere CloudBurst Appliance. Continue on to the next chapter.

5

WebSphere CloudBurst Appliance

This chapter describes how to configure a new WebSphere CloudBurst Appliance, create cloud resources, and successfully connect our newly created PowerVM cloud to the WebSphere CloudBurst Appliance.

This chapter describes the following topics:

5.1 The hardware

The WebSphere CloudBurst Appliance is a 1U mountable appliance that installs into a standard 19-inch rack mount (Figure 5-1).

Figure 5-1 WebSphere CloudBurst Appliance

The appliance includes:

- Two power cables
- A serial interface cable
- Rack rails for installation into a standard 19-inch rack

On the front of the WebSphere CloudBurst Appliance is a serial cable port and four 1 Gb Ethernet ports. The Ethernet ports are divided into one management port (labeled MGMT) for managing the appliance using the web console and command-line interface. The remaining three ports are used for deploying the virtual systems (Figure 5-2).

Figure 5-2 WebSphere CloudBurst Appliance interface ports

5.2 Initial setup

Initialize the WebSphere CloudBurst Appliance by performing the following operations:

1. Place the WebSphere CloudBurst Appliance in a rack or a suitable position in your data center.
2. Connect the power cables.
3. Connect the management port labeled MGMT to the appropriate network.
4. Connect the serial cable to the serial connection port on the front of the WebSphere CloudBurst Appliance to a computer that is running a serial emulator program. In our test environment the serial connection was connected to a serial console server.
5. Power on the WebSphere Cloudbust Appliance.

6. When the appliance starts you will see initialization messages (Figure 5-3), followed by a welcome message, requesting to continue. Press C and press the Enter key to continue.

```
cbannex.rtp.raleigh.ibm.com - PuTTY
Thu Jul 08 2010 09:47:22 [supervisor] Spawn: validator state=1 pid=0 stop=0 newp
id=1181
Thu Jul 08 2010 09:47:22 [supervisor] Spawn: ensemble state=1 pid=0 stop=0 newpi
d=1184
Thu Jul 08 2010 09:47:22 [supervisor] Spawn: lightpd state=1 pid=0 stop=0 newpid
=1188
Thu Jul 08 2010 09:47:22 [supervisor] Spawn: cbserial state=0 pid=0 stop=0 newpi
d=1189
Thu Jul 08 2010 09:47:23 [supervisor] Spawn: setupfs state=0 pid=0 stop=0 newpid
=1192
Thu Jul 08 2010 09:47:23 [supervisor] Spawn: dpinfo state=0 pid=0 stop=0 newpid=
1198
Thu Jul 08 2010 09:47:23 [supervisor] Spawn: drouter state=1 pid=0 stop=0 newpid
=1199
Thu Jul 08 2010 09:47:25 [supervisor] Spawn: validator state=1 pid=1181 stop=0 n
ewpid=1210
Initializing hard disks...


Running firmware version 2.0.0.0-24024.


Welcome to the IBM WebSphere CloudBurst Appliance serial console.
Please press the "c" key followed by the [Enter] key to continue: Thu Jul 08 201
```

Figure 5-3 The WebSphere CloudBurst Appliance initalization messages

7. Choose the language that you want WebSphere CloudBurst Appliance to use. In our scenario we selected English by pressing the number 1 and then pressing Enter (Figure 5-4).

```
cbannex.rtp.raleigh.ibm.com - PuTTY
18
Attached to port 18



Welcome to the IBM WebSphere CloudBurst Appliance serial console.
Please press the "c" key followed by the [Enter] key to continue: c

Running firmware version 2.0.0.0-24024.


Select the language for use with Cloudburst:
1) English (en)
2) French (fr)
3) Italian (it)
4) Spanish (es)
5) German (de)
6) Japanese (ja)
7) Portuguese (Brazil) (pt_BR)
8) Chinese (China) (zh_CN)
9) Chinese (Taiwan) (zh_TW)
10) Korean (ko)
Select the language [1-10] (default: 1): 1
```

Figure 5-4 Language selection

8. Read the license agreement. There are three license agreements that you need to agree to during this step. Press Enter to read the next page of the license agreement. When you

have finished reading the license agreement press `Y` and Enter key to accept the license agreement. Perform this step for each license agreement (Figure 5-5).

```
cbannex.rtp.raleigh.ibm.com - PuTTY
MS AND CONDITIONS FOR SEPARATELY LICENSED CODE



WebSphere CloudBurst Appliance V2.0 Managed Beta



The IBM license agreement and any applicable information on the web

download page for IBM products refers Licensee to this file for details

concerning terms and conditions applicable to code identified as

Separately Licensed Code in the License Information document and

included in the products listed above ("the Program").



The "Separately Licensed Code" identified in the License Information

--> License 1 of 3, page 1 of 164.
--> Press [Enter] to go to next page, [Y/y] to accept, [N/n] to reject: y
```

Figure 5-5 The license agreement

9. The cbadmin login user name is the super user for the WebSphere CloudBurst Appliance. Enter your superuser password and press Enter (Figure 5-6). The cbadmin password is shown on the widnow while it is being typed.

```
cbannex.rtp.raleigh.ibm.com - PuTTY
management of such entity, whether by contract or otherwise, or (ii)
ownership of fifty percent (50%) or more of the outstanding shares, or
(iii) beneficial ownership of such entity.

"You" (or "Your") shall mean an individual or Legal Entity exercising
permissions granted by this License.

"Source" form shall mean the preferred form for making modifications,
including but not limited to software source code, documentation
source, and configuration files.

"Object" form shall mean any form resulting from mechanical
transformation or translation of a Source form, including but not
limited to compiled object code, generated documentation, and
conversions to other media types.
--> License 3 of 3, page 5 of 711.
--> Press [Enter] to go to next page, [Y/y] to accept, [N/n] to reject: y

The following entry will be the superuser password and should not be forgotten.
The password WILL BE VISIBLE on this console, so verify
that no untrusted parties can see this console before continuing.
Please read your password as you type it to verify
that your keyboard is working properly.
Enter the new password for the superuser "cbadmin":
```

Figure 5-6 Change the cbadmin user password

10. Enter the IP address of your management port and press Enter (Figure 5-7 on page 73). Only use the management interface should for administering the WebSphere CloudBurst

Appliance via the web console and the command-line tools. The management interface requires a static IP address. In our scenario we use IP address 9.42.89.38.

Figure 5-7 Management interface IP address configuration

11. Enter the subnet mask for the management interface in Classless Inter-Domain Routing (CIDR) notation and press Enter (Figure 5-8).

Figure 5-8 Enter the subnet mask

CIDR notation represents a mask bit from number 8 to 30. Table 5-1 provides examples.

Table 5-1 Examples of subnet mask and equivalent CIDR notation

Subnet mask	CIDR notation
255.255.255.128	25
255.255.255.0	24
255.255.252.0	22
255.255.0.0	16

12. Enter the default gateway for your location and press Enter (Figure 5-9).

Figure 5-9 Default gateway configuration

13. Select the network type of the management interface and press Enter (Figure 5-10). If you are unsure of your network, select option 1 for auto to allow the appliance to determine the correct settings. In our scenario we selected option 1.

```
cbannex.rtp.raleigh.ibm.com - PuTTY

Please wait while the cbadmin password is set...  The password is successfully s
et.


Configure ethernet port MGMT...
Enter the IP address of the interface (default: 9.42.89.38): 9.42.89.38
Enter the subnet mask in CIDR format (default: 24): 24
Enter the default gateway IP address (default: 9.42.89.1): 9.42.89.1
Select the physical mode of the ethernet interface (default: Auto)
  1) Auto
  2) 100baseTx-FD
  3) 100baseTx-HD
  4) 10baseT-FD
  5) 10baseT-HD
  6) 1000baseTx-FD
Enter 1-6: 1
```

Figure 5-10 Select the management network interface

14. Enter the IP address of you primary DNS server and press Enter (Figure 5-11). The DNS server must be reachable from the management network. You must enter a functional DNS server that is able to perform forward and reverse IP address and host name lookups. The WebSphere CloudBurst Appliance does not function correctly without a properly working DNS server.

```
cbannex.rtp.raleigh.ibm.com - PuTTY

Please wait while the cbadmin password is set...  The password is successfully s
et.


Configure ethernet port MGMT...
Enter the IP address of the interface (default: 9.42.89.38): 9.42.89.38
Enter the subnet mask in CIDR format (default: 24): 24
Enter the default gateway IP address (default: 9.42.89.1): 9.42.89.1
Select the physical mode of the ethernet interface (default: Auto)
  1) Auto
  2) 100baseTx-FD
  3) 100baseTx-HD
  4) 10baseT-FD
  5) 10baseT-HD
  6) 1000baseTx-FD
Enter 1-6: 1
Enter the IP address of the DNS server (default: 172.16.15.253): 172.16.15.253
```

Figure 5-11 Primary DNS server configuration

15. The initial configuration is now complete. You can now manage the WebSphere CloudBurst Appliance through the web console and the command-line tools. Review your network settings to make sure that they are correct. If the network settings need to be changed press C, then Enter (Figure 5-12).

```
cbannex.rtp.raleigh.ibm.com - PuTTY
  5) 10baseT-HD
  6) 1000baseTx-FD
Enter 1-6: 1
Enter the IP address of the DNS server (default: 172.16.15.253): 172.16.15.253

No changes to the Network Interface.
Network setup is complete.

Configuration summary for MGMT interface:
  IP address: 9.42.89.38
  Subnet mask: 24
  Default gateway: 9.42.89.1
  Ethernet physical mode: Auto
  DNS server: 172.16.15.253

This completes the first boot process.
Point your web browser at https://9.42.89.38 for the appliance web console
and login as "cbadmin" with the password you entered above.
You may login here again with any superuser userid if you need to reconfigure
the IP interface without the web console.


Welcome to the IBM WebSphere CloudBurst Appliance serial console.
Please press the "c" key followed by the [Enter] key to continue:
```

Figure 5-12 Summary of initialization

16. Log in to the web console of the WebSphere CloudBurst Appliance by opening a web browser and entering the address of your management interface. In our scenario we used `https://9.42.89.38` (Figure 5-13). The web console only supports Mozilla Firefox 3.0, 2.5 and Internet Explorer 7.0 and 8.0. Using the incorrect web browser displays an error message.

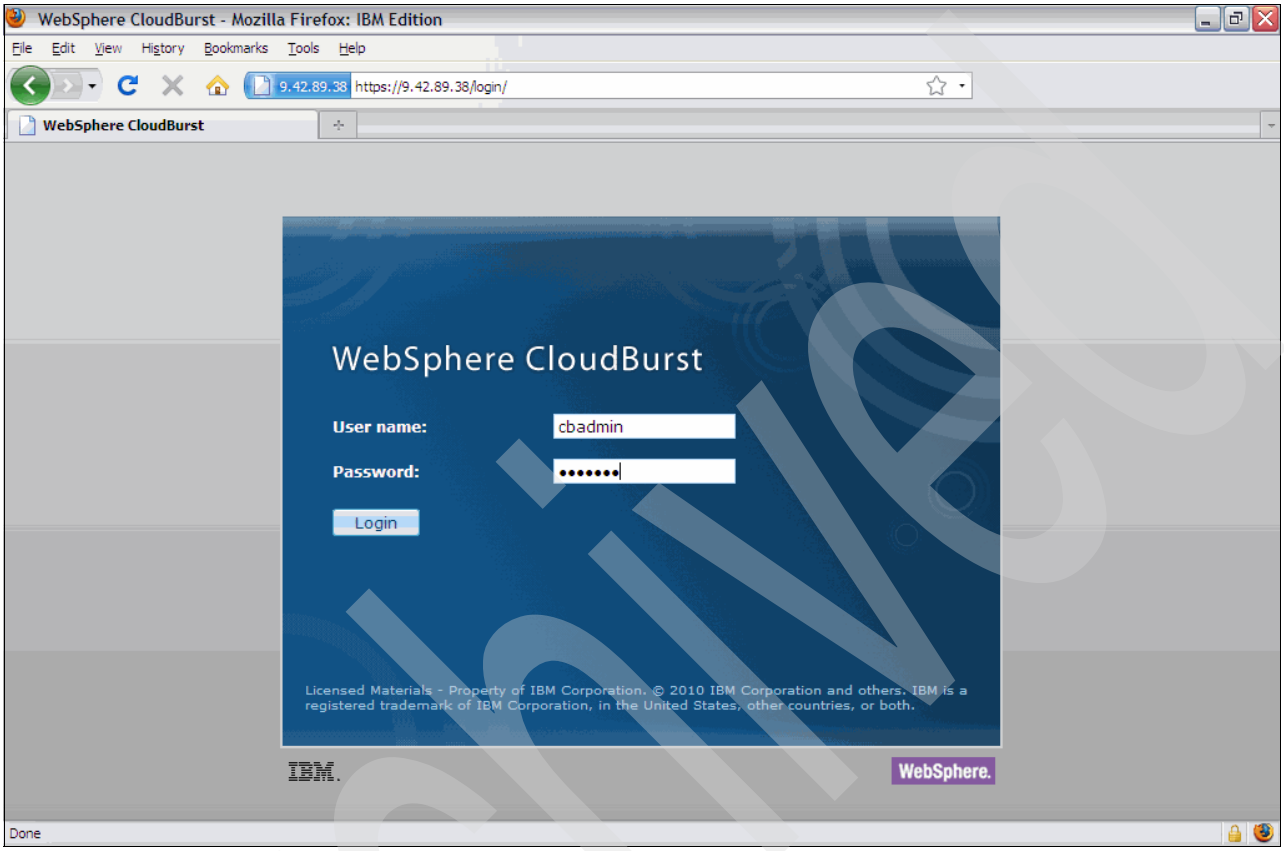

Figure 5-13 WebSphere CloudBurst Appliance web console

17. Enter the user name `cbadmin` and the superadmin password that you specified in step 9, and then click **Login**.

Upon successful login, you are presented with the welcome window (Figure 5-14). Congratulations on a successful initialization of the WebSphere CloudBurst Appliance.

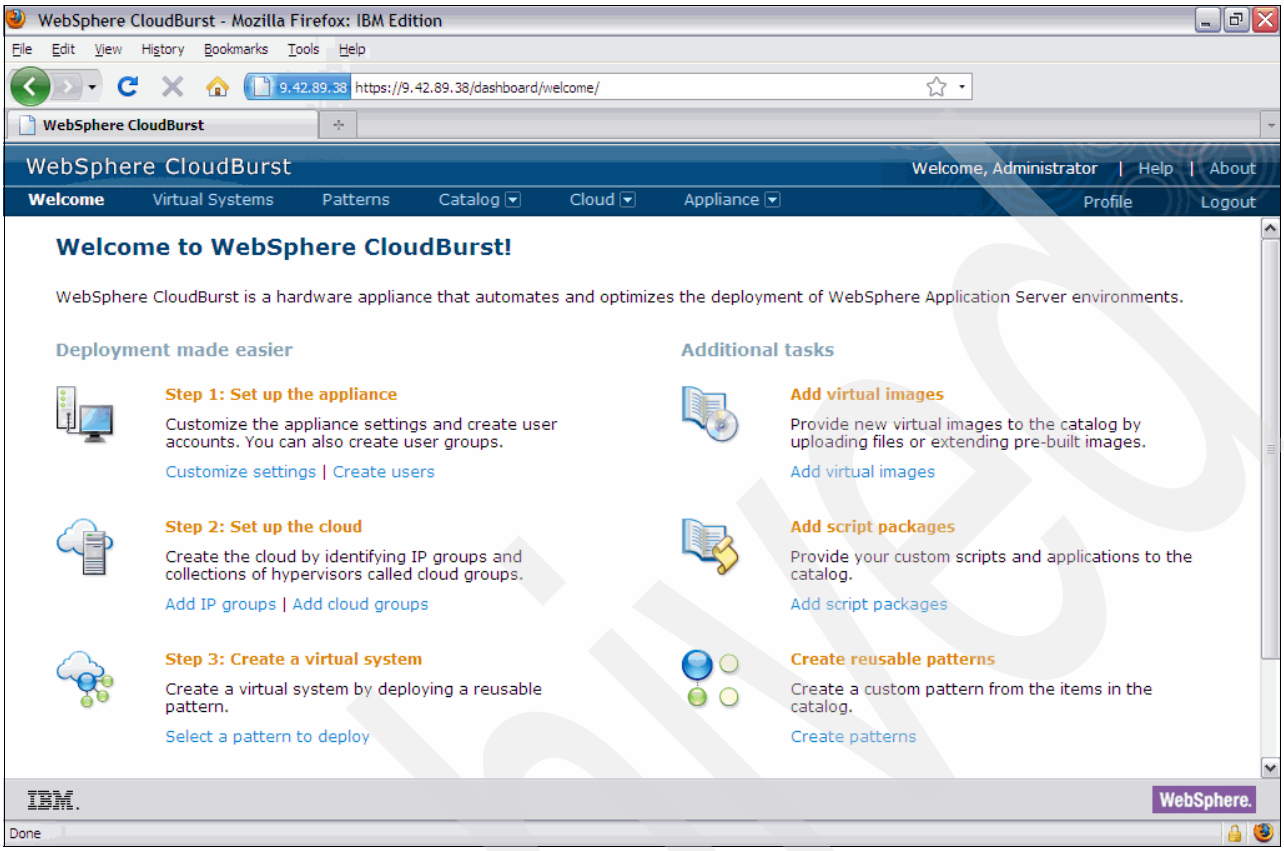

Figure 5-14 WebSphere CloudBurst Appliance web console welcome window

5.3 Basic cloud resources scenario

In this section, we describe how to create the basic resources for a cloud environment on the WebSphere CloudBurst Appliance.

Although there are many useful configuration options that can be configured on the WebSphere CloudBurst Appliance, there are only a few steps required to enable you to quickly setup and start deploying a preconfigured WebSphere Application. In this section we describe how to create the basic resources required to get your WebSphere cloud up and running so you can start deploying images.

5.3.1 Configuring the deployment networks

Before deployment of virtual systems into the cloud, it is a best practice to configure at least one other Ethernet interface on the WebSphere CloudBurst Appliance dedicated to deployment of virtual systems. Using the management network to deploy the images can cause high network latency to the web console and loss of session. Keeping the management and deployment networks separate ensures that this situation does not occur, as the data physically travels along separate paths.

To set up a deployment network, perform the following steps:

1. Log in to the WebSphere CloudBurst Appliance using the `cbadmin` user.
2. Click the **Customize Settings** link (Figure 5-15).

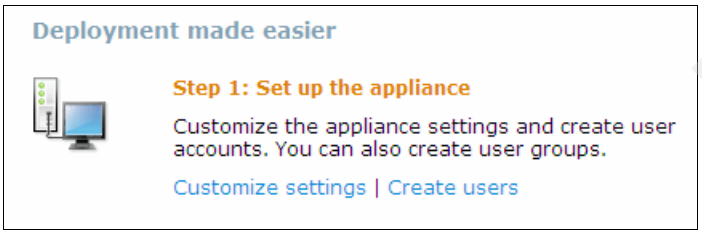

Figure 5-15 The customize settings link

3. Click the **Ethernet Interfaces** link.
4. Configure your network interfaces as desired by clicking the relevant boxes related to IP address/mask and the default gateway (Figure 5-16). In our scenario, our management interface is only used to access the web console and command line of the WebSphere CloudBurst. The other two interfaces, eth0 and eth2, are used for deploying the virtual systems.

Figure 5-16 Network interface configuration

5.3.2 Setting up the cloud resources

To initially define the cloud there are only two steps required:

- Define the IP groups.
- Define the cloud groups.

Add IP group

The IP groups are lists of IP addresses that are reserved for use by virtual systems. The IP addresses need to be fully resolvable by DNS, both forward and reverse, and should be defined in the DNS before creating the IP group. When a new pattern is deployed into the cloud it will use IP addresses from the IP group that has been defined. To create the IP groups, perform the following steps:

1. Log in to WebSphere CloudBurst Appliance web console or, if already logged in, click **Welcome** on the top left of the web console.
2. On the welcome page select the **Add IP groups** link (Figure 5-17).

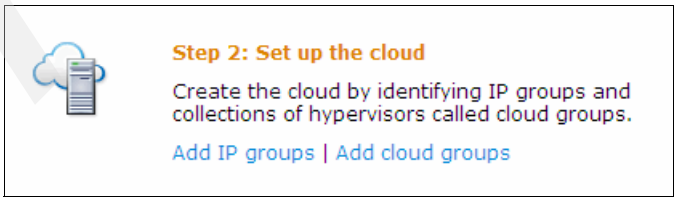

Figure 5-17 The Add IP Groups hyperlink

3. Click the green plus icon (**+**) to create a new IP group (Figure 5-18).

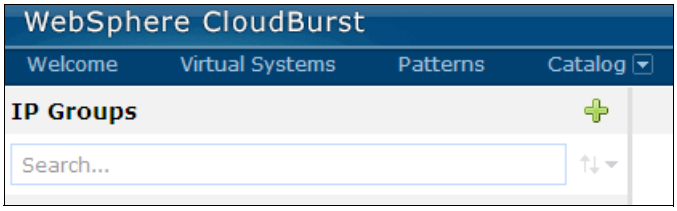

Figure 5-18 Add new IP groups icon

4. Enter the details of the IP group that you want to add and click **OK** (Figure 5-19). The network details used to define the IP group are the network details that are passed to the virtual system for use with its networking configuration. We used the settings shown in Example 5-1. The naming convention ITSO is designed to represent a false department within our scenario. The concept is related to ITSO being an individual department.

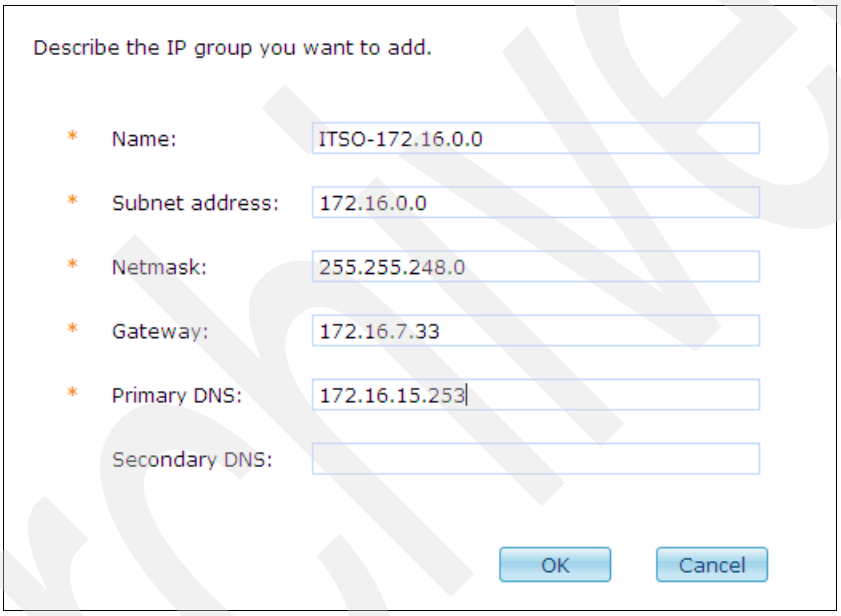

Figure 5-19 Add new IP group details

Example 5-1 Details used to create IP group

```
Name: ITSO-172.16.0.0
Subnet Address:172.16.0.0
Netmask:255.255.248.0
Gateway:172.16.7.33
Primary DNS:172.16.15.253
```

5. Click the newly created IP group. We use a small subset of the available IP addresses in the subnet for use by our cloud group. That is, we are going to define a list of available IP addresses that can be used in our cloud. This is done by entering details in the Add range fields on the right side of the window (Figure 5-20). In our scenario we use all of the 254 addresses in the 172.16.0.0 subnet IP addresses for our cloud deployment.

Figure 5-20 How to add an IP address range

6. A new list of IP address and host names (if IP addresses are defined in the DNS) are displayed from the IP address range defined (Figure 5-21). When an IP address has been assigned to a virtual system by the WebSphere CloudBurst Appliance, the icon next to the IP address changes to a solid green square. This helps users quickly identify which IP addresses are in use.

   In our scenario the host names and IP addresses have been reserved for use by the ITSO department cloud environment. The administrator of the DNS server has updated the DNS details prior to adding of the IP group.

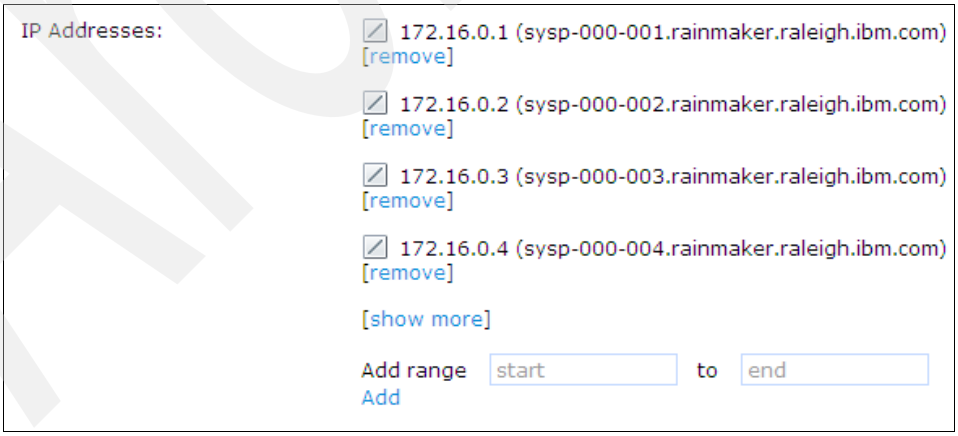

Figure 5-21 IP address and host names from newly defined IP address range

Add a cloud group

To complete the cloud configuration, the WebSphere CloudBurst Appliance needs to connect to a hypervisor that it can control to create and remove LPARs. In our scenario, we connect to

our recently created IBM Director/VMControl server, itso-vmc. All CECs defined in VMControl are used as a cloud resource.

To define a new cloud group, perform the following tasks:

1. From the menu bar across the top of the window select **Cloud** → **Cloud Groups** (Figure 5-22).

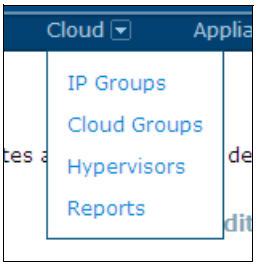

Figure 5-22 The Add cloud groups button on the welcome page

2. Click the green **plus** icon (**+**) to create a new cloud group.
3. Enter the details for the cloud group and click **Create** (Figure 5-23).

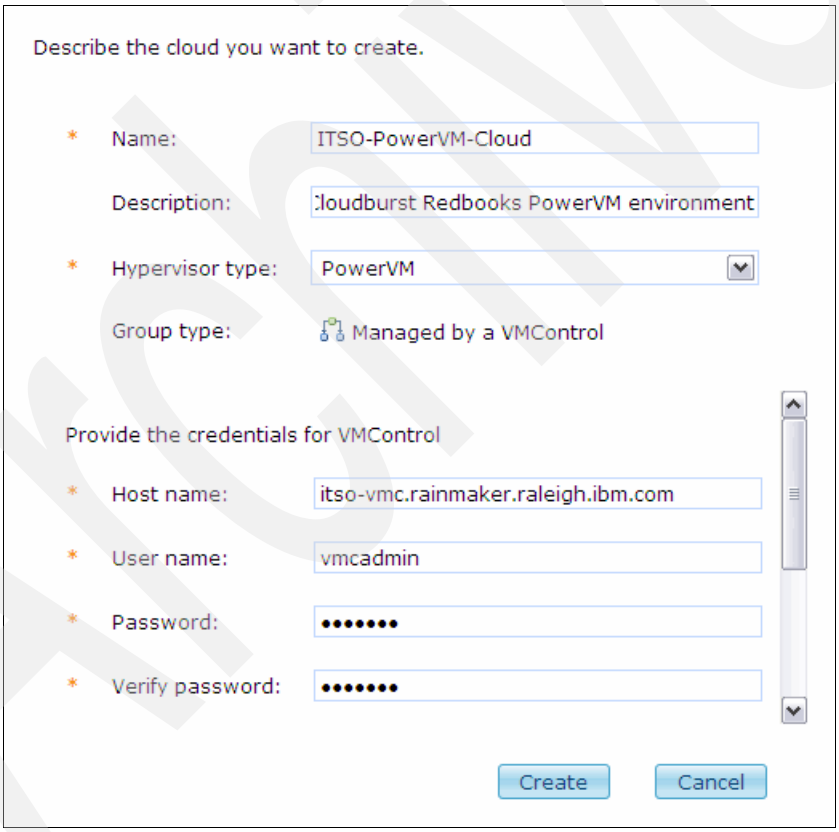

Figure 5-23 Create cloud group panel

In our scenario we entered the details shown in Example 5-2.

Example 5-2 Details used to create cloud group

```
Name:                ITSO-PowerVM-Cloud
Description:         ITSO WebSphere CloudBurst Redbooks PowerVM environment
Hypervisor Type:     PowerVM

Credentials for VMControl

Host name:           itso-vmc.rainmaker.raleigh.ibm.com
User name:           vmcadmin
Password:            **********

Operating System Credentials

User name:           root
Password:            **********
```

4. Upon successful connection to the hypervisor, click **Accept** to accept the connection certificate to the server (in our scenario, this is the itso-vmc server) (Figure 5-24).

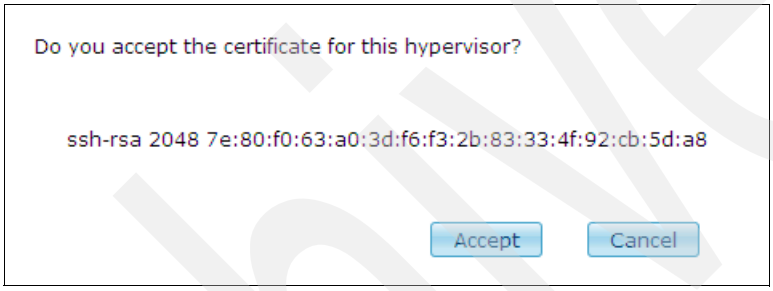

Figure 5-24 Accept certificate

5. A new cloud group has been created and will run a discovery process. The discovery process runs for one or two minutes, discovering the objects in the cloud environment, such as the CECs (known as hypervisors in WebSphere CloudBurst Appliance), the shared Ethernet adapters, and the volume groups on the VIO servers, which will become the storage devices. When discovery has completed, the hypervisors are visible in the hypervisors section and the current status of the cloud states *You must start at least one hypervisor to create a virtual system* (Figure 5-25). Click the hypervisor name (in our scenario, IBM 8204 E8A 06D55F4) to prepare the hypervisor to start.

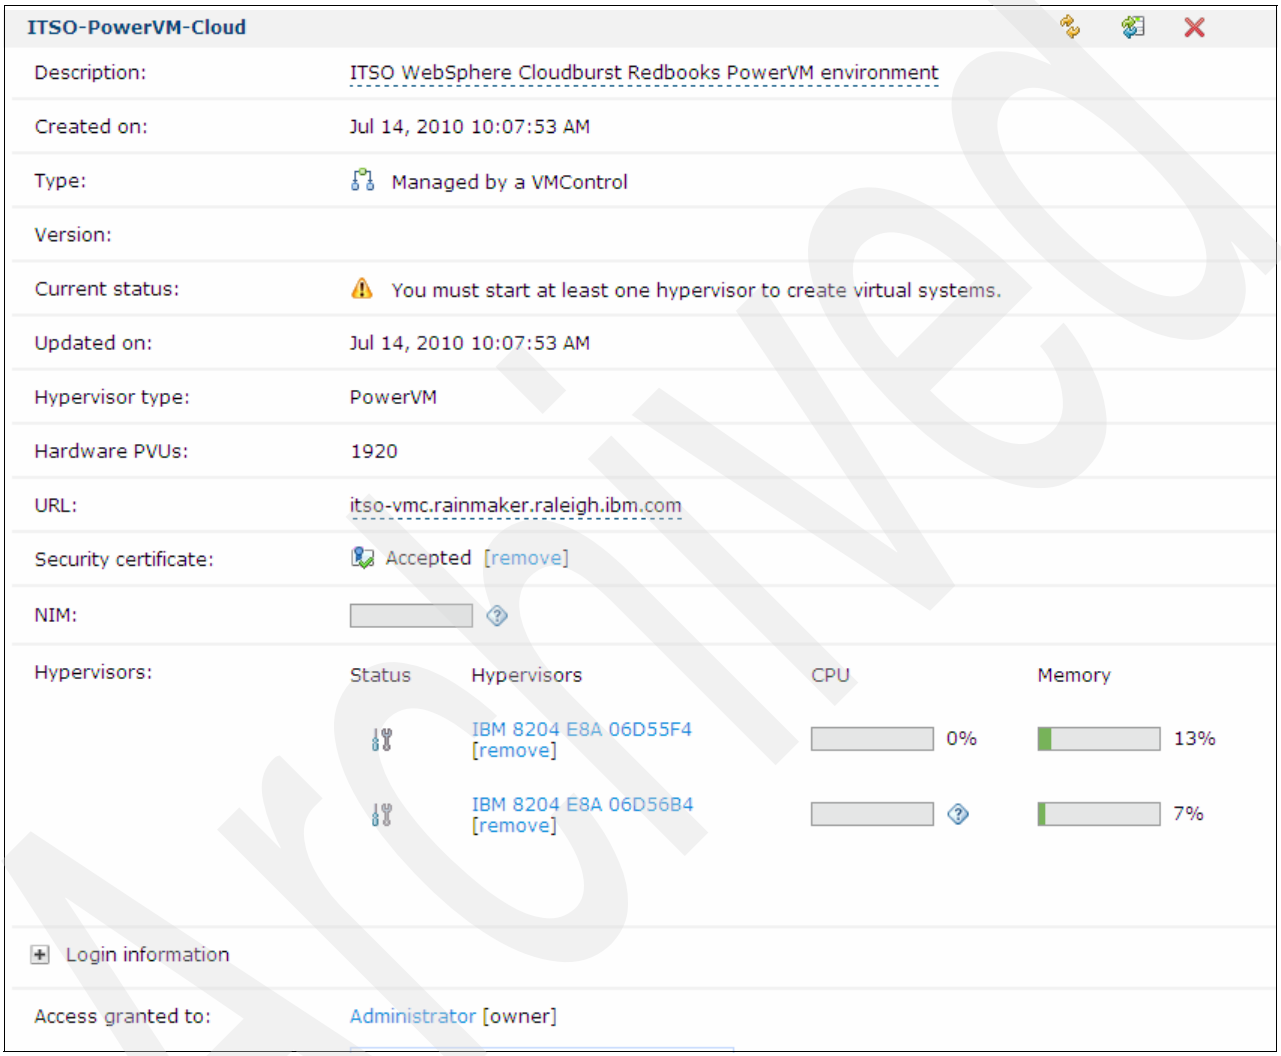

Figure 5-25 The newly created cloud group

6. The hypervisors page opens, displaying the properties of the selected hypervisor. The hypervisor is currently in maintenance mode, as displayed by the icon of the screwdriver and wrench in the current status section. When a hypervisor is in maintenance mode, the actual server is operating normally, and any deployed virtual systems continue to run as normal, but changes to the hypervisor can only be made, and no new deployments are possible on that particular hypervisor. To configure the network of the hypervisor click the **Networks** link (Figure 5-26).

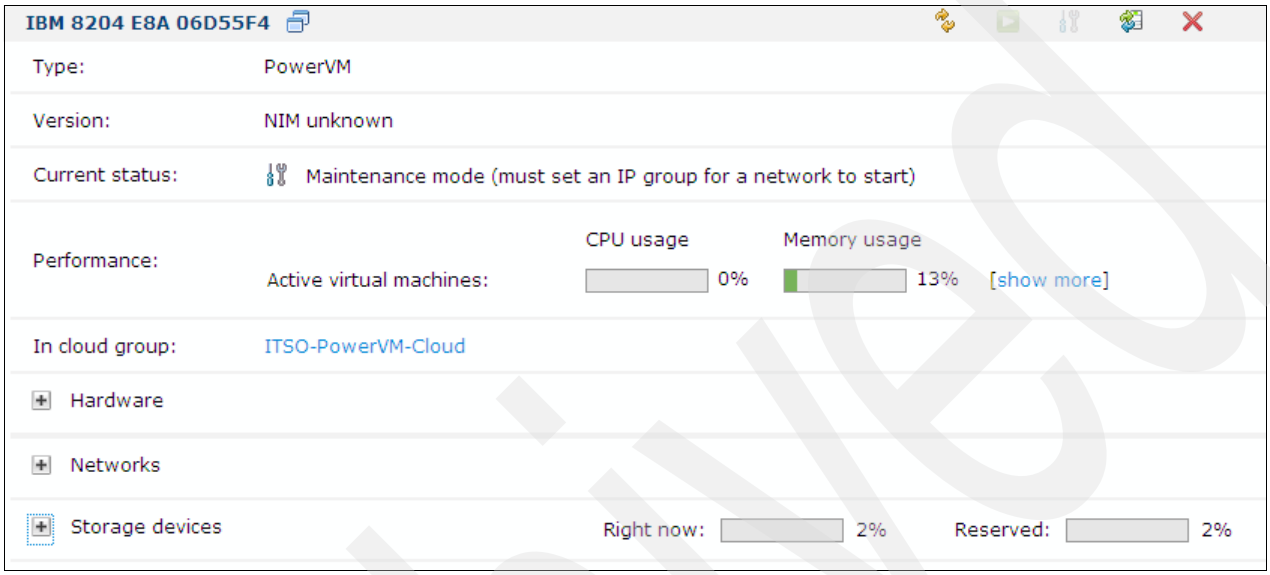

Figure 5-26 A PowerVM hypervisor

7. Click the **Virtual Network 1** link, click the drop-down box **IP group, and select the appropriate IP group (Figure 5-27). We associate a** shared Ethernet adapter on the hypervisor with an IP group that can be used for deployment. In our scenario we selected the ITSO-172.16.0.0 IP group that we created earlier.

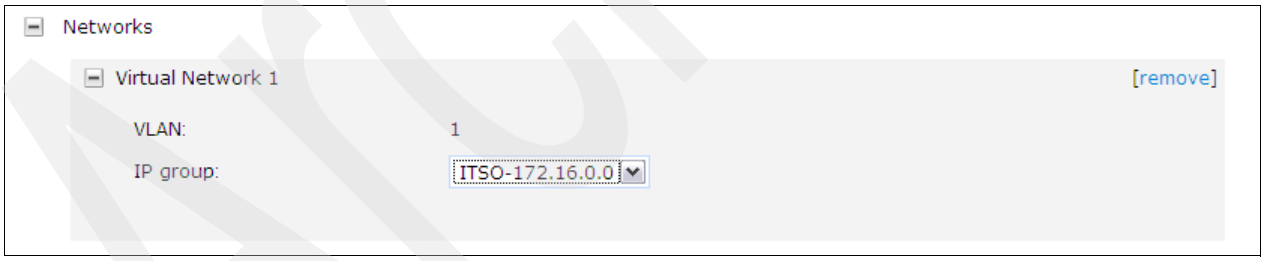

Figure 5-27 Associate an IP group with an SEA

8. Click the **Storage Devices** link to display the available volume groups on the VIO server (Figure 5-28).

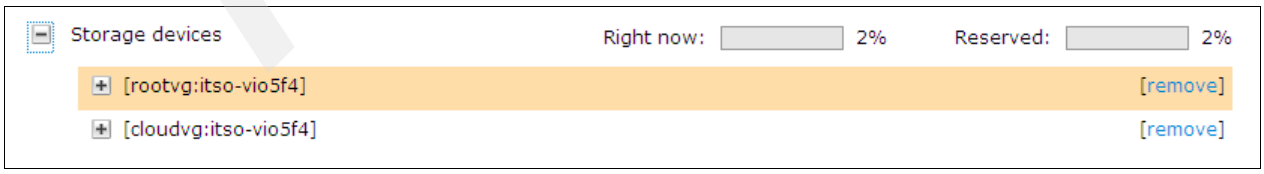

Figure 5-28 Hypervisor storage devices

9. Click the **Remove** button next to `rootvg:<lparname>`, where <*lparname*> is the name of your VIOS LPAR. New deployments of patterns only use the volume groups shown as storage devices in the hypervisor. By removing the root volume group, the rootvg is not used as a storage device for any deployments of patterns. In our scenario we removed rootvg:vio5f4, where vio5f4 is the name of the VIOS server as displayed on the HMC.

> **Note:** Consider removing the root volume group (rootvg) from the available storage devices. In the opinion of the authors, it is important to keep rootvg for operating system functions only. Using the rootvg as a storage device can cause undesirable effects.

10. Click the Start icon on the hypervisor in the top right area of the window (Figure 5-29). This makes the hypervisor available as a cloud resource.

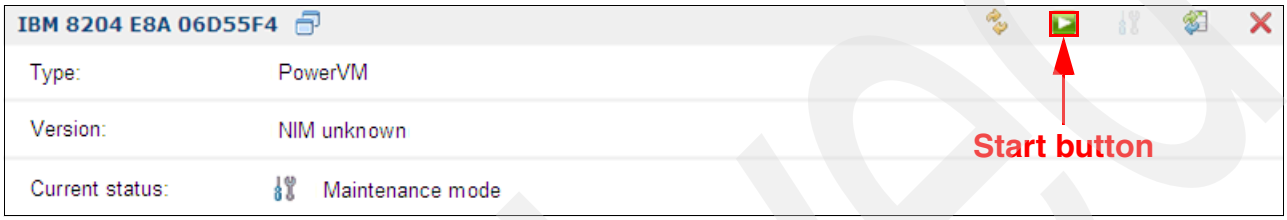

Figure 5-29 Start the hypervisor

11. Repeat steps 6 to 9 for any other hypervisors in your cloud group.

Test deployment

To confirm that all components are set up correctly and are working together, we perform a deployment of a pattern to the PowerVM cloud, ITSO-PowerVM-Cloud. The pattern used in this scenario is a default WebSphere single server pattern, which comes standard with the appliance.

WebSphere CloudBurst Appliance is designed as a self-service portal. Deploying patterns is normally performed by predefined individual users who have access to a limited set of resources. This is only a test deployment to confirm that the system is working. Thus, we decided to use the cbadmin user to perform the test deployment.

By default, WebSphere CloudBurst Appliance comes preconfigured with a set of WebSphere patterns designed for the VMware hypervisors. To use one of the default patterns on PowerVM, all that needs to be done is to clone the existing pattern and change the target platform to PowerVM. This is an example of how the patterns are essentially platform agnostic, and highlight how easy it is to develop on one platform and then quickly deploy the same patterns on production platforms.

In the following scenario we copy the existing WebSphere single server pattern, change the image to PowerVM hypervisor, and deploy in the ITSO-PowerVM cloud.

Accept virtual image license agreements

Before we can copy an existing pattern, we must accept the license agreements for the virtual images of WebSphere Application Server for Novell SUSE Linux Enterprise Server 11 on the VMware hypervisor and WebSphere Application Server for PowerVM.

To accept the license agreements, perform the following tasks:

1. From the menu bar at the top of the window select **Catalog** → **Virtual Images** (Figure 5-30).

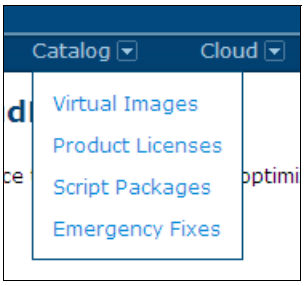

Figure 5-30 The catalog menu

2. Choose the **WebSphere Application Server 7.0.0.11 (PowerVM)** virtual image from the list of images (Figure 5-31). The X symbol in the red box indicates that the license agreements have not been accepted.

Figure 5-31 WebSphere Application Server for PowerVM virtual image

3. Click the **accept** hyperlink next to the accept category from the virtual image properties window (Figure 5-32). (The license agreement currently displays `Not accepted`.)

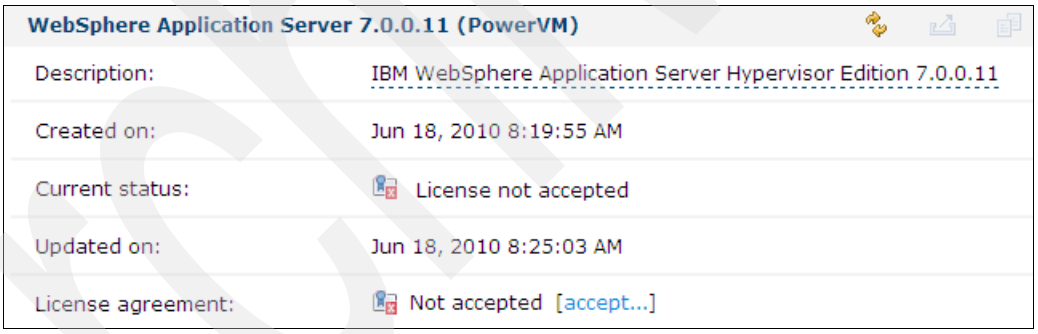

Figure 5-32 License status of virtual image

4. Two licenses need to be accepted for this license agreement. Click the **WebSphere Application Server Hypervisor Edition License** link (Figure 5-33).

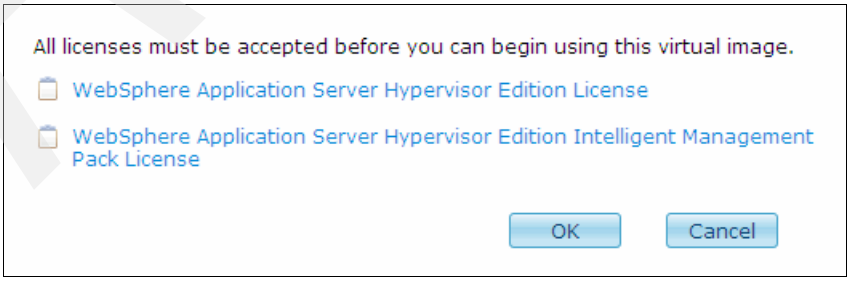

Figure 5-33 License status panel

5. Read the WebSphere Application Server Hypervisor Edition licence and click **Accept** (Figure 5-34).

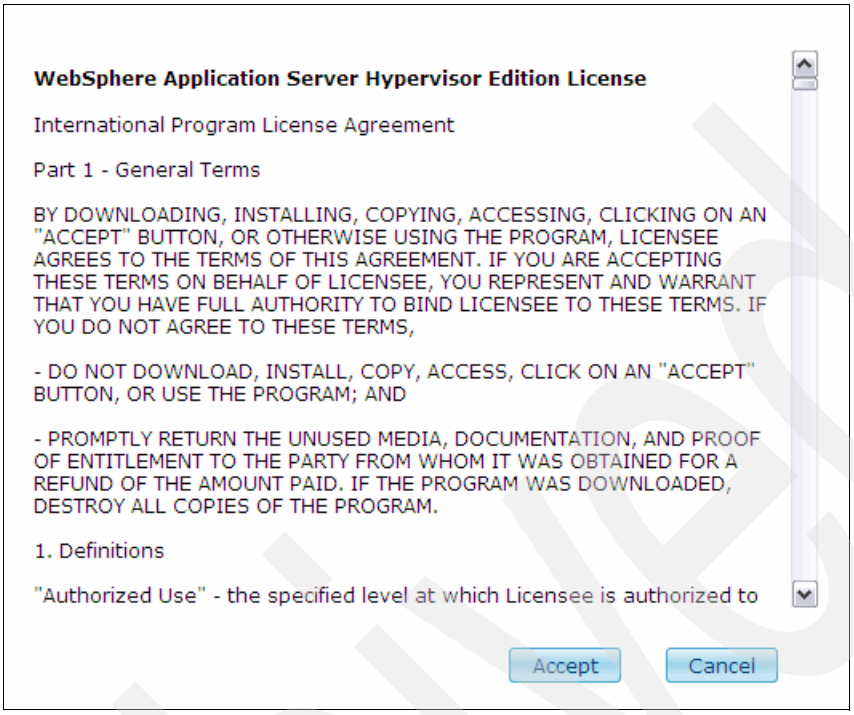

Figure 5-34 The WebSphere Application Server Hypervisor Edition license

6. Click the **WebSphere Application Server Hypervisor Edition Intelligent Management Pack License** link.

7. Read the WebSphere Application Server Hypervisor Edition Intelligent Management Pack license and click **Accept** (Figure 5-35).

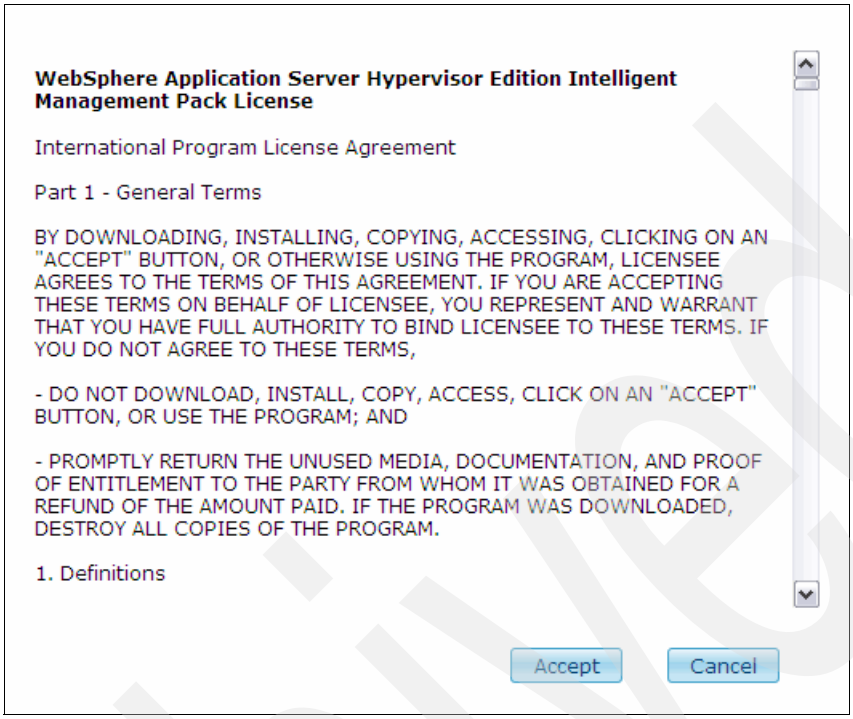

Figure 5-35 The WebSphere Application Server Hypervisor Edition Intelligent Mgmt Pack license

8. Click **OK** on the license summary window. The WebSphere license for PowerVM has been accepted (Figure 5-36).

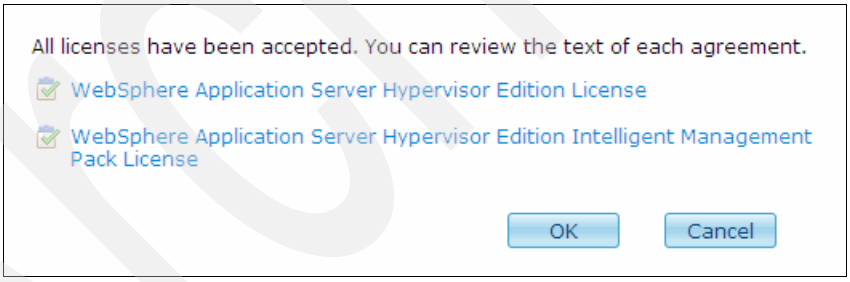

Figure 5-36 The accepted licenses

9. Click the **WebSphere Application Server 7.0.0.11** virtual image from the list of images.

10. There are two WebSphere Application Server 7.0.0.11 VMware images with the same name, one is built on the Novell SUSE Linux Enterprise Server 11 (SLES11) operating system and the other is built on the RedHat Enterprise Linux 5 (RHEL5) operating system. To ensure that we can copy the existing pattern, we need to accept the licenses for the SLES11 platform. Check the hypervisor platform and the operating system by checking the hypervisor type and operating system section of the properties in the properties of the virtual image (Figure 5-37).

Hypervisor type:	ESX
Operating system:	SLES 11 (Novell SUSE Linux Enterprise Server 11)

Figure 5-37 Check the hypervisor and operating system

11. On the VMware/SLES11 image, click the **accept** hyperlink next to the accept category from the virtual image properties window.
12. Accept each of the four licence agreements (Figure 5-38) using the same method as outlined in steps 4 to 8.

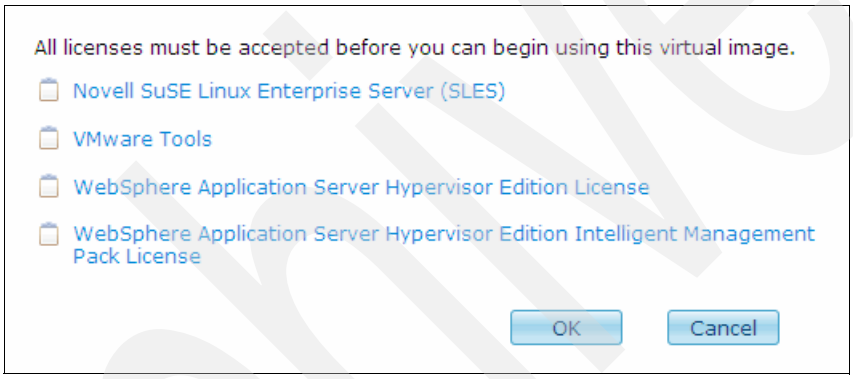

Figure 5-38 The four licenses for the SLES/VMware combination

13. Confirm that all licence agreements have been accepted for WebSphere Application Server 7.0.0.11 on SLES11 and PowerVM. This is done by checking that the virtual image has a green check mark with a padlock icon next to the virtual image name (Figure 5-39).

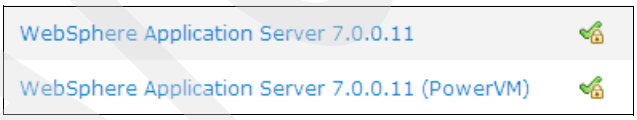

Figure 5-39 Confirmation that the licenses have been accepted

Clone pattern

The licenses for both of the required virtual images has been accepted. The next step in the test deployment scenario is to clone the standard WebSphere single server pattern for the VMware/SLES11 environment and define it to use the PowerVM hypervisors. To do this:

1. Select **Patterns** from the menu bar at the top of the window (Figure 5-40).

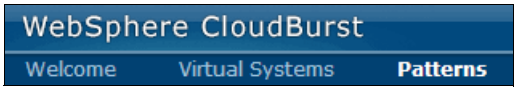

Figure 5-40 The Patterns menu button

2. Select the WebSphere single server pattern from the patterns list (Figure 5-41). The lock icon next to the pattern name indicates that the pattern is a locked, read-only pattern. Modifications cannot be made to this pattern once locked. To make a modification to the pattern, the pattern must be cloned, and then changes can be made to the new copy of the pattern.

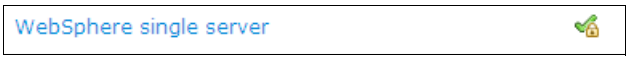

Figure 5-41 The WebSphere single server pattern

3. Clone the pattern by clicking the Clone icon (Figure 5-42) on the patterns properties window.

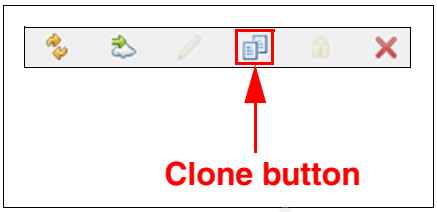

Figure 5-42 The clone pattern button

4. Rename the pattern and change the virtual image to a PowerVM WebSphere Application Server image. In our scenario we changed the a name to WebSphere single server (PowerVM) (Figure 5-43). Click **OK** to clone the new pattern to a PowerVM hypervisor.

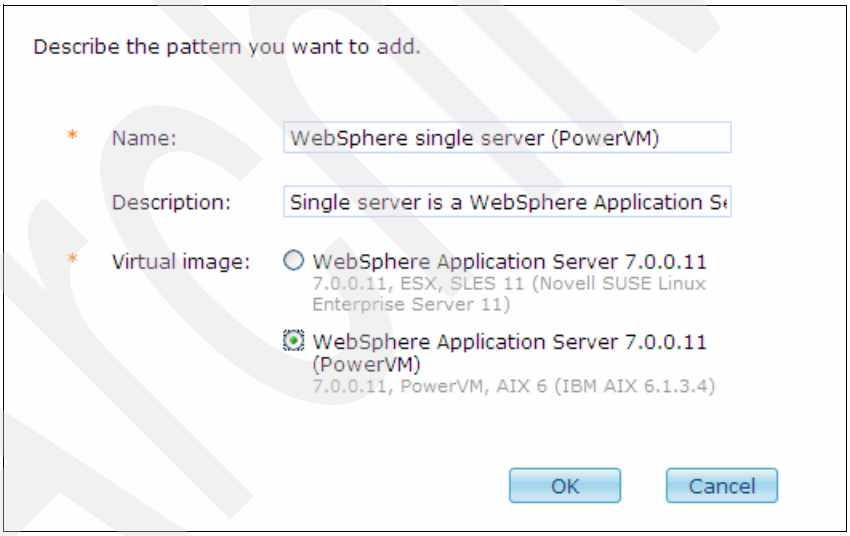

Figure 5-43 Description of the new cloned server

A new WebSphere single server for PowerVM pattern has been created.

Deployment

The last step of the test deployment process is to deploy the pattern into the cloud. To do this:

1. Select the newly created pattern called WebSphere single server (PowerVM) by clicking its name in the pattern list window.
2. Click the Deploy into cloud icon on the patterns properties window (Figure 5-44).

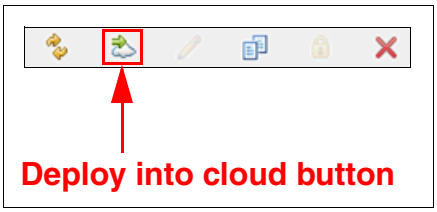

Figure 5-44 The deploy into cloud button

3. Give the new virtual system a name and choose a cloud group to deploy the server on. In our scenario we named our test server itso-test1 and choose the cloud ITSO-PowerVM-Cloud (Figure 5-45). When a pattern is deployed to a cloud, the new objects that are created are part of virtual image.

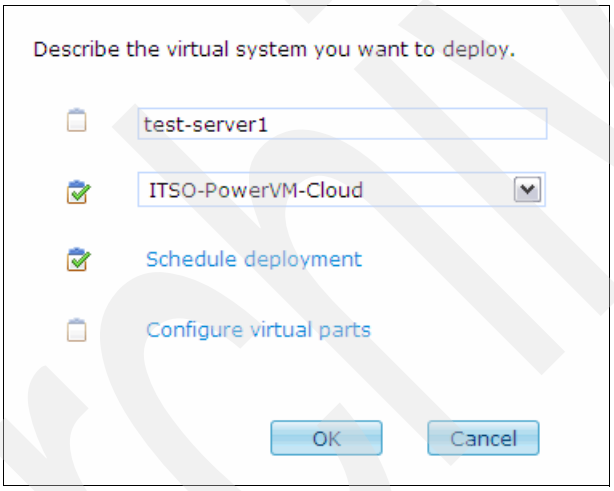

Figure 5-45 Creation of new virtual system

4. Click the **Configure virtual parts** → **Standalone server** link (Figure 5-46).

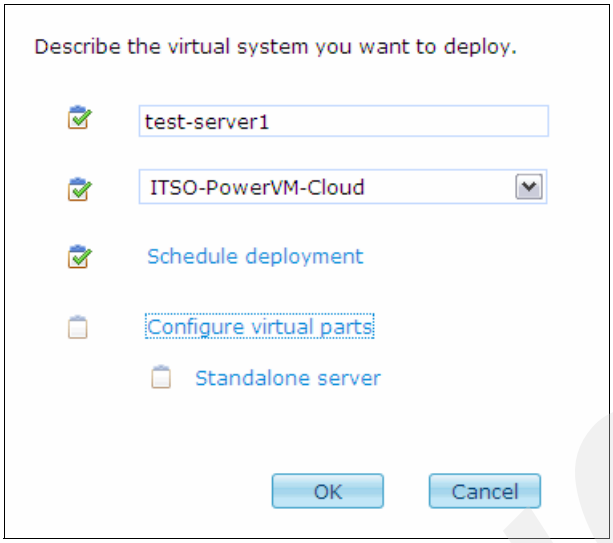

Figure 5-46 Configure standalone server

5. A new window opens that allows us to change the parameters of the server, such as virtual CPU, memory, and cell name (Figure 5-47). In our scenario, we left all settings as the defaults. Use the scroll bar to scroll to the bottom of the pattern properties to the password section. Type a password for the root and the virtuser (Figure 5-48 on page 94) and click **OK**. Passwords are required items before the deployment can proceed.

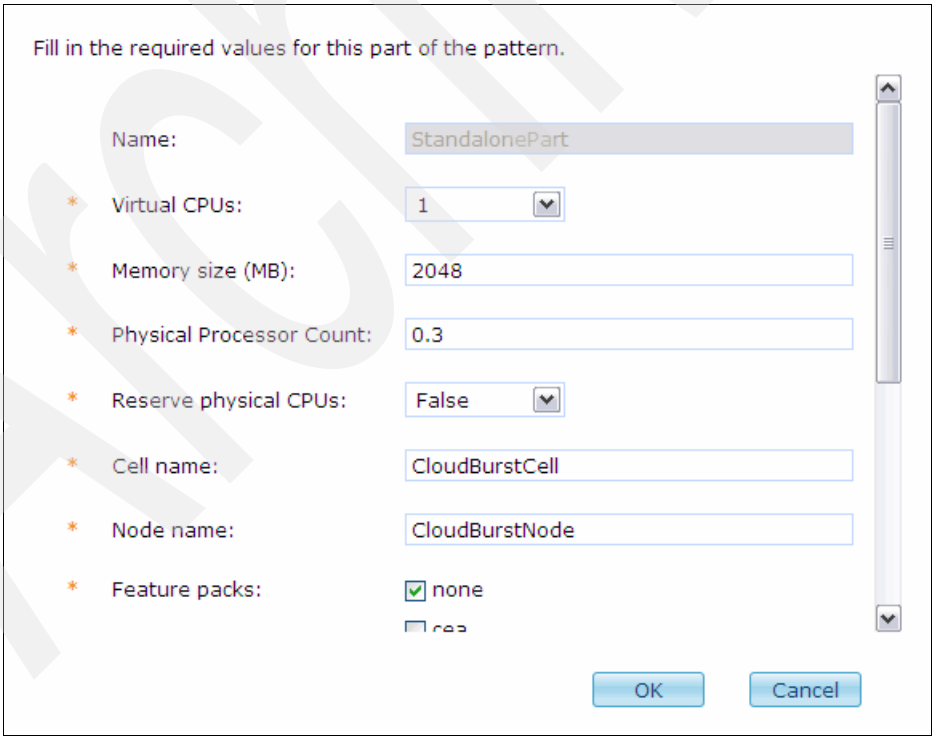

Figure 5-47 Virtual system properties

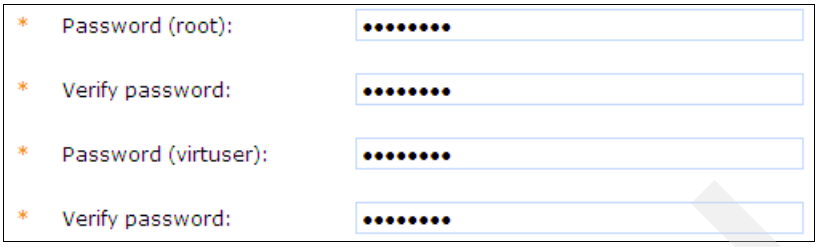

Figure 5-48 Required passwords for root and virtuser

6. Click **OK** on the virtual system description window to start the deployment.
7. After clicking **OK**, you are taken to the Virtual Systems view, where you can see the status of your virtual systems deployment (Figure 5-49). The pattern is being deployed as a virtual system. It takes approximately one hour to perform the first deployment on AIX (results were from our test system and times are highly dependant on network resources). This is due to the OVA image being transferred to the IBM Director/VMControl server from WebSphere, creation of a virtual appliance, transferring the image from IBM Systems Director/VMControl to the NIM server, and creation of the SPOT image. Subsequent installations of this PowerVM pattern do not include these steps and thus take considerably less time to complete.

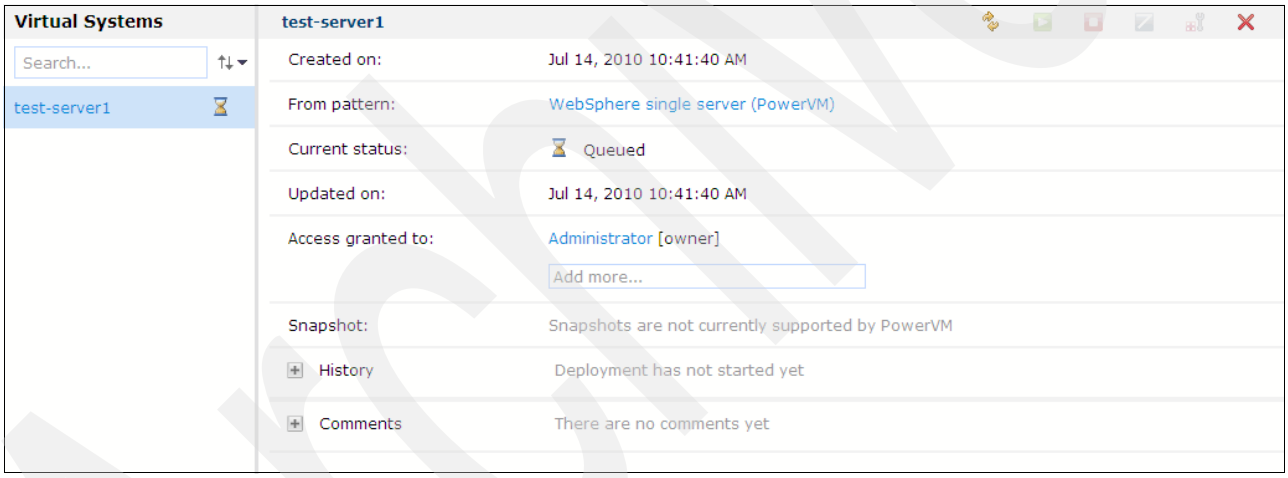

Figure 5-49 New virtual systems

8. Progress of the deployment of the virtual system can also be tracked from the task queue. To open the task queue, click **Appliance** → **Task Queue** on the menu bar (Figure 5-50).

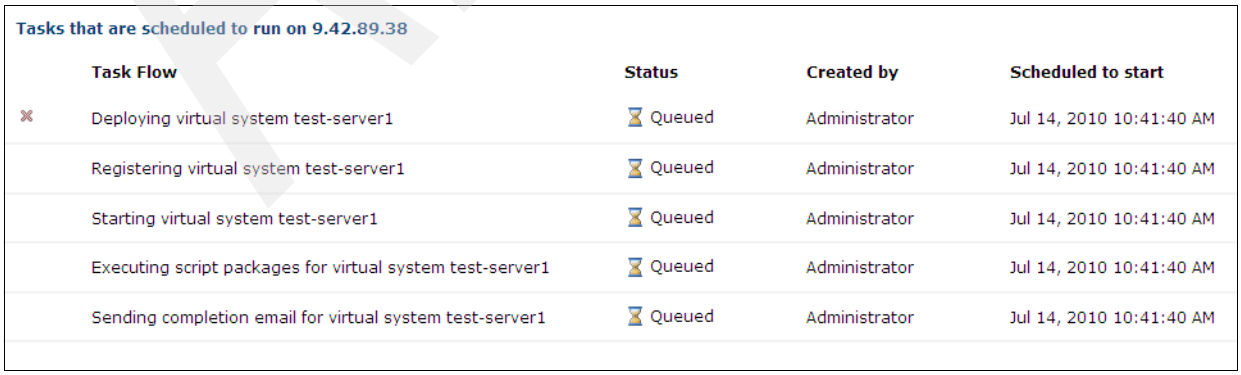

Figure 5-50 Task queue

9. When the deployment is complete there should be no more tasks in the task queue. Click **Virtual Systems** on the menu bar. A start symbol (a green square with a white arrow) signifies a successful deployment.

5.4 Summary

The configuration of the WebSphere CloudBurst Appliance is the last step in the software stack cloud. Although configuration of the entire stack can seem difficult, breaking down the installation steps into components makes the installation a manageable and simpler process.

Now that installation of the software stack is complete, primary management of cloud environment is provided by the WebSphere CloudBurst Appliance.

Part 3

Using WebSphere CloudBurst Appliance on PowerVM

This part describes how to customize and administer a WebSphere CloudBust Appliance. It provides an overview of WebSphere CloudBurst Patterns and IBM Hypervisor Edition Virtual images.

This part includes the following chapters:

6

WebSphere CloudBurst Appliance Administration

The true power in the WebSphere CloudBurst Appliance is its ability to act as a self-service portal, which allows users to deploy their own WebSphere images on demand. In this chapter we discuss how to customize and administer a WebSphere Cloudbust Appliance.

This chapter describes:

6.1 Customization of the WebSphere CloudBurst Appliance

This section covers the common WebSphere CloudBurst Appliance administration tasks. These tasks can also be completed from the command line, but that is not covered in this chapter.

The tasks in this section are performed from the web console. For the following instructions we assume that the user is logged in as the cbadmin user. The settings are changed from the settings page. To access the settings, select **Appliance** → **Settings** on the menu bar (Figure 6-1).

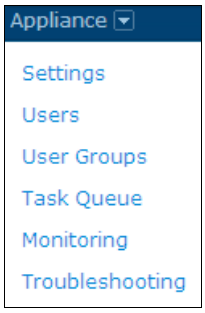

Figure 6-1 The Settings menu item

6.1.1 Application identification

Application identification is the way that the appliance can be named so that it can easily be identified. To change the name of the appliance:

1. Click **Appliance Identification** on the Settings page (Figure 6-2).

Figure 6-2 The appliance identification link

2. Click the underlined boxes next to Appliance name and Appliance group to change the name and group identification (Figure 6-3). In our scenario we change the name to itso-cbapp and the group to IBM.

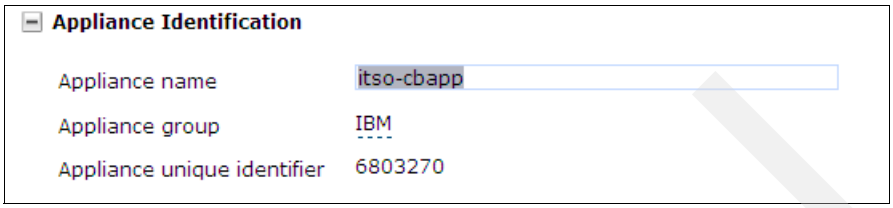

Figure 6-3 Change the appliance identification

6.1.2 Security

Security is an important issue in today's IT environment. Each company has its own security requirements, which vary depending on whether they are external or internal IT systems. The WebSphere CloudBurst Appliance can be configured to conform with a company's security requirements by applying certain security settings (Figure 6-4).

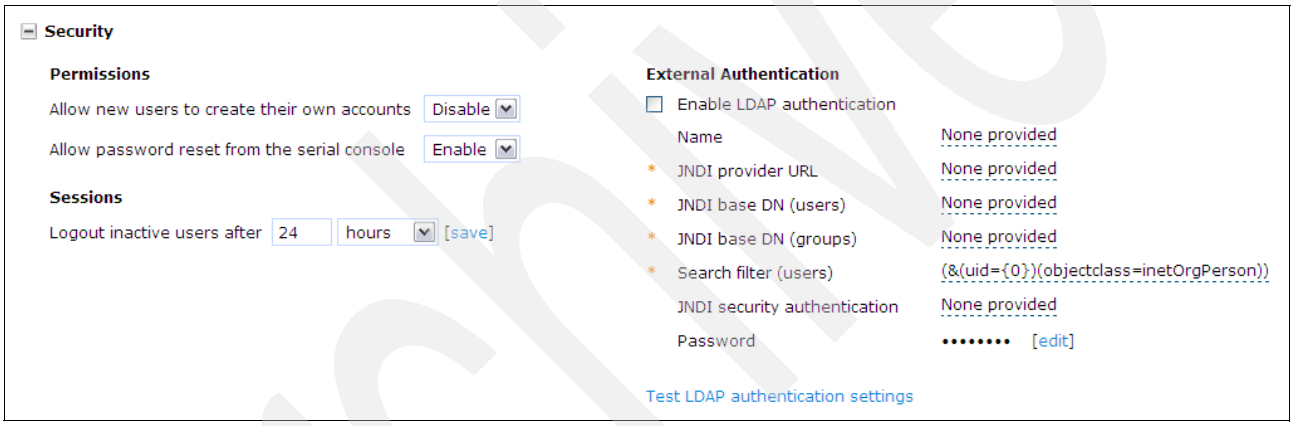

Figure 6-4 Available security settings

The WebSphere CloudBurst Appliance has three security settings:

- Permissions

  Permissions allows users to create their own accounts and reset their passwords from the serial console. Each of these items has an enable/disable setting. In our scenario we use the default settings.

- Sessions

  The session time out can be changed from as low as 1 minute, to many hours. When making any changes of the session time out, the WebSphere CloudBurst Appliance needs to be restarted for the change to take effect. In our scenario we keep the default time out of 24 hours.

- External Authentication

  External authentication of users is allowed against an LDAP server. In our scenario we do not utilize LDAP authentication. We utilize authentication using local users on the WebSphere CloudBurst Appliance.

6.1.3 Ethernet interfaces

The Ethernet interface configuration was explained in 5.3.1, “Configuring the deployment networks” on page 78. From this menu, the status of each interface can be reviewed (Figure 6-5).

Ethernet Interfaces

Interface	IP address/mask	Default gateway	MTU	Mode
eth0	172.16.0.227/21	None provided	1500	Auto
eth1	None provided	None provided	1500	Auto
eth2	172.16.15.208/21	None provided	1500	Auto
mgmt	9.42.89.38/24	9.42.89.1	1500	Auto

Ethernet interfaces status

Click to expand

					Received				Transmitted			
Interface	MAC address	Link status	Mode status	Collisions	kilobytes	packets	errors	drops	kilobytes	packets	errors	drops
eth0	00:1a:64:88:8f:9d	ok	1000baseTx-FD	0	297455	4862578	0	0	0	4	0	0
eth1	00:14:5e:f1:bf:88	no-link	10baseT-HD	0	0	0	0	0	0	0	0	0
eth2	00:14:5e:f1:bf:8a	ok	1000baseTx-FD	0	320768	5180784	0	0	500	6824	0	0
mgmt	00:1a:64:88:8f:9c	ok	100baseTx-FD	0	129419	1913330	0	0	2126848	1534323	0	0

Figure 6-5 The Ethernet Interfaces settings

6.1.4 Domain Name Services

Although Domain Name Server (DNS) settings were applied during the initialization procedure, additional DNS servers can be added and removed from this setting menu. The menu also provides the ability to perform a forward and reverse DNS lookup to check the validity of a host name or IP address (Figure 6-5).

Domain Name Servers

172.16.15.253

Click to add

Lookup host name or IP address

itso-vmc.rtp.raleigh.ibm.com Lookup

Host name: itso-vmc.rtp.raleigh.ibm.com

IP address: 9.42.132.37

Figure 6-6 The DNS server settings and DNS lookup

6.1.5 Date and time

The date and time configuration allows you to configure the desired time zone and a Network Time Protocol (NTP) server. In our scenario, we use EST (US Eastern Time) and NTP server 172.16.15.253 (Figure 6-7).

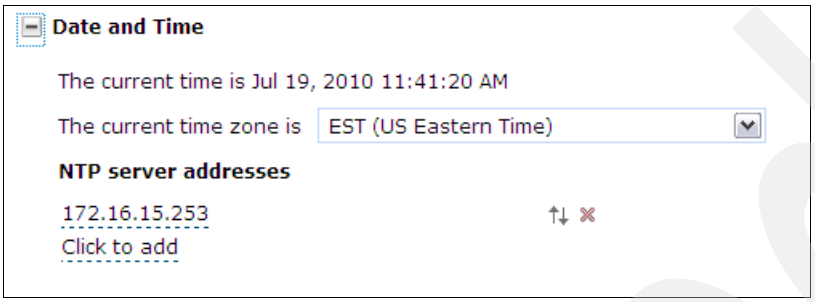

Figure 6-7 The data and time configuration

6.1.6 Mail configuration

Configuring a Simple Mail Transfer Protocol (SMTP) server allows notification of the following events:

- Password resets.
- Virtual system is created.
- Virtual system has successfully started.
- Virtual system has failed to start.
- Virtual image is exported.
- Virtual image is imported.

To configure a SMTP server for mail notification:

1. Click **Mail Delivery**.
2. Add SMTP server information, as either a host name or IP address. In our scenario we use dcertpmail.raleigh.ibm.com.
3. Add a reply-to address. Enter the email address for the WebSphere CloudBurst Appliance here. In our scenario we enter cbadmin@itso-app.raleigh.ibm.com (Figure 6-8).

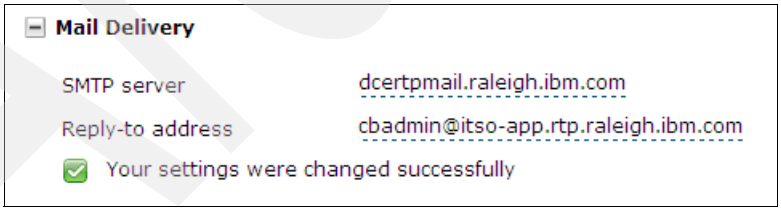

Figure 6-8 Mail delivery configuration

6.2 Administration of WebSphere CloudBurst Appliance

The WebSphere CloudBurst Appliance administration covers a wide range of tasks. Only certain tasks that we believe to be the most common administrative tasks are covered in this section.

6.2.1 Backup and restore

The backup and restore functionality of the WebSphere CloudBurst Appliance allows you to back up the entire appliance and all of its settings (users, Ethernet interface, custom patterns, IP groups, cloud groups, and so on).

All backups are encrypted using RSA encryption to protect any sensitive information. Decryption of the backup package can only be done using the private key. It also must be decrypted on another appliance, so you must own an appliance and have the key to decrypt it. Backups are performed using the SSH protocol for security.

Backup

When a backup occurs, all currently running tasks are completed before the backup begins. No new tasks can be requested until the backup is completed.

A backup of a CloudBurst appliance creates multiple file segments (Figure 6-9). A backup of an appliance containing only default data is approximately 28 GB. that you have enough backup capacity on the target file system and that the backup user has the ability to make large files.

```
vmcontrol.rtp.raleigh.ibm.com - PuTTY
$ ls -l
total 59105416
-rw-r--r--    1 cbbackup staff        343200 Jul 12 23:59 cbRestore.manifest
-rw-r--r--    1 cbbackup staff           383 Jul 12 23:59 cbRestore.metadata
-rw-r--r--    1 cbbackup staff    2863053312 Jul 12 16:51 cbRestore.seg1
-rw-r--r--    1 cbbackup staff            40 Jul 12 17:21 cbRestore.seg1.sha1
-rw-r--r--    1 cbbackup staff     769376368 Jul 12 20:09 cbRestore.seg10
-rw-r--r--    1 cbbackup staff            40 Jul 12 20:17 cbRestore.seg10.sha1
-rw-r--r--    1 cbbackup staff    1055119648 Jul 12 20:21 cbRestore.seg11
-rw-r--r--    1 cbbackup staff            40 Jul 12 20:32 cbRestore.seg11.sha1
-rw-r--r--    1 cbbackup staff    1154046192 Jul 12 20:36 cbRestore.seg12
-rw-r--r--    1 cbbackup staff            40 Jul 12 20:48 cbRestore.seg12.sha1
-rw-r--r--    1 cbbackup staff     608808832 Jul 12 20:52 cbRestore.seg13
-rw-r--r--    1 cbbackup staff            40 Jul 12 20:58 cbRestore.seg13.sha1
-rw-r--r--    1 cbbackup staff    1036728944 Jul 12 21:03 cbRestore.seg14
-rw-r--r--    1 cbbackup staff            40 Jul 12 21:13 cbRestore.seg14.sha1
-rw-r--r--    1 cbbackup staff    1272459360 Jul 12 21:18 cbRestore.seg15
-rw-r--r--    1 cbbackup staff            40 Jul 12 21:31 cbRestore.seg15.sha1
-rw-r--r--    1 cbbackup staff    1184399856 Jul 12 21:36 cbRestore.seg16
-rw-r--r--    1 cbbackup staff            40 Jul 12 21:48 cbRestore.seg16.sha1
-rw-r--r--    1 cbbackup staff    1103493600 Jul 12 21:52 cbRestore.seg17
-rw-r--r--    1 cbbackup staff            40 Jul 12 22:04 cbRestore.seg17.sha1
-rw-r--r--    1 cbbackup staff     839587680 Jul 12 22:08 cbRestore.seg18
-rw-r--r--    1 cbbackup staff            40 Jul 12 22:16 cbRestore.seg18.sha1
```

Figure 6-9 List of backup files on the target server

In our scenario we took a backup of the appliance to our NIM server itso-nim. We used a dedicated user for the backups called cbbackup that can create unlimited file sizes and also had ownership of a 45 GB logical volume mounted on `/backups`. In our test scenario with a 1 GB Ethernet link between the WebSphere CloudBurst Appliance and the target server, the backup took 7.5 hours. The final file size of the backup was 28.2 GB.

In the following instructions we assume that the user is logged in as the cbadmin user. To access the Settings page select **Appliance** → **Settings** on the menu bar.

To perform a backup perform the following tasks:

1. Click **Backup and Restore**.
2. Click **Step 1: Find a place to store your backups**. Backups use the SSH protocol. The destination server on which you choose to store your backup must be a SSH server.
   a. Enter the host name or IP address of the SSH server, where the image will be stored.
   b. Enter the path of the SSH server where the image will be stored.
   c. Enter the user name of the user on the target server used to make the connection.
   d. Enter the password of the user on the target server used in making the connection.
   e. Click the **Submit** button to save the password (Figure 6-10).

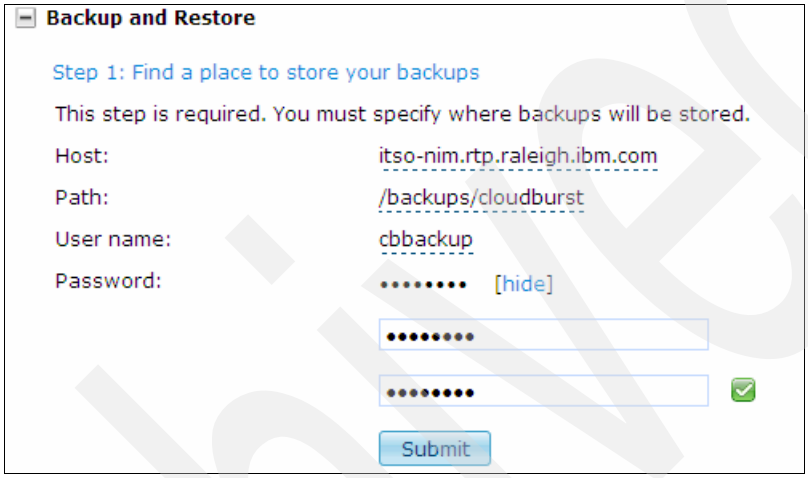

Figure 6-10 Backup user information

3. Click **Step 2: Secure your backups**. Backups are secured using an RSA key. To generate your own key pair:
   a. Enter a password in the New password and Verify password fields.
   b. Click **Generate**.
   c. Click the **Download the backup private key** link.
   d. Save the private key to your local computer and store it in a secure location. The private key is required to restore the backup. The WebSphere CloudBurst Appliance does not store SSH keys. Thus, they must be stored in another secure location (Figure 6-11).

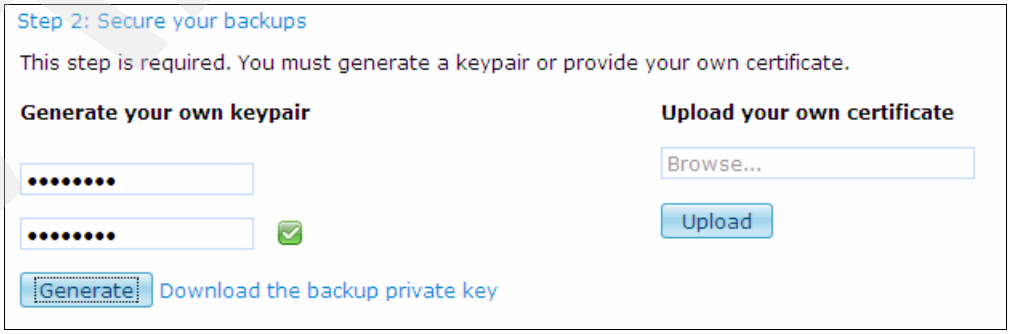

Figure 6-11 Secure the backup

4. Click **Step 3: Schedule a new backup**. A backup can be scheduled to run immediately, which is the default date and time, or they can be scheduled for a future date. When using a future date the date format needs to be *mm/dd/yyyy* (for example, to represent the 26th of July 2010, the date format is 07/26/2010). In our scenario, we used the default settings and did not change the date, so the backup was scheduled to run when all current tasks were complete. To do this click **Schedule** (Figure 6-12).

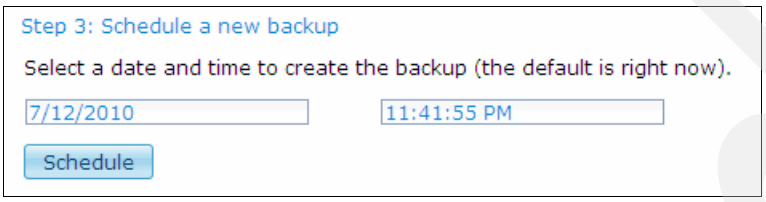

Figure 6-12 Schedule the backup

5. To view the status of the backup task, click **Appliance** → **Task Queue**, which shows the backup task in the task queue (Figure 6-12).

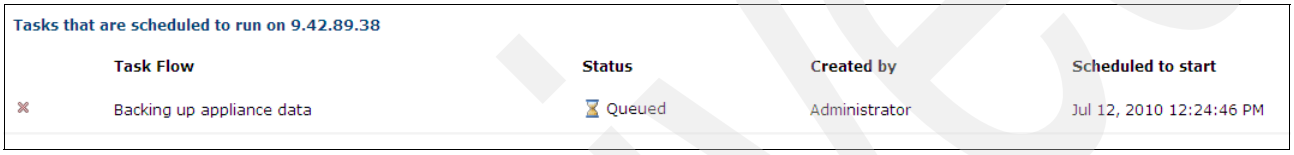

Figure 6-13 Backup is queued

Figure 6-14 shows the active backup.

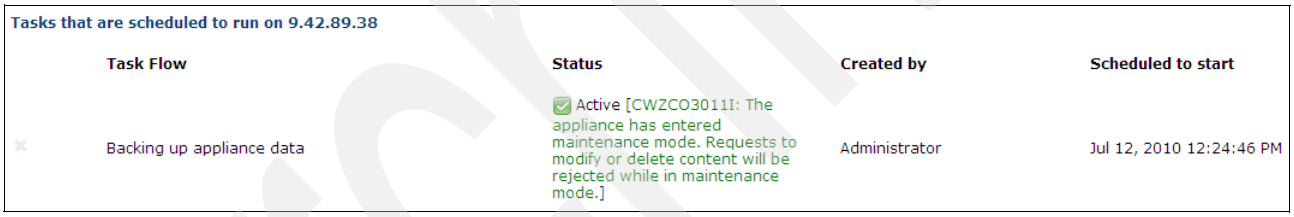

Figure 6-14 Backup is active

Figure 6-15 shows which backup is running.

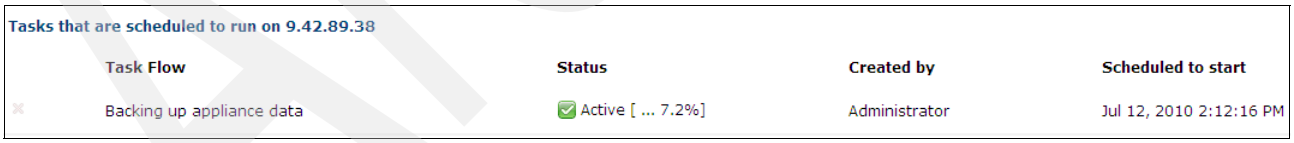

Figure 6-15 Backup is running

Restore

Restores are performed in a similar manner to backups, as they rely upon the SSH credentials stored in step 1 of the backup procedure. The restore facility checks the SSH server and folder location for any backups and displays them as available to users.

Before performing any restore, confirm that the firmware version of the backup is the same as the currently running firmware of the appliance. If you want a backup from an earlier running firmware you must change the firmware level to match.

To perform a restore of WebSphere CloudBurst Appliance:

1. Click **Step 1: Find a place to store your backups**. Enter the credentials for the server where the backup is stored, if it is not already populated with the correct details. In our scenario we use the itso-nim server and the cbbackup user specified in the backup section earlier.
2. Click **Step 4: Restore from backup**.
3. Select a backup to restore, and select the **Restore** link associated with that backup. In our scenario we select the backup with an archive file date of `16:34 Mon, Jul 12, 2010` (Figure 6-16).

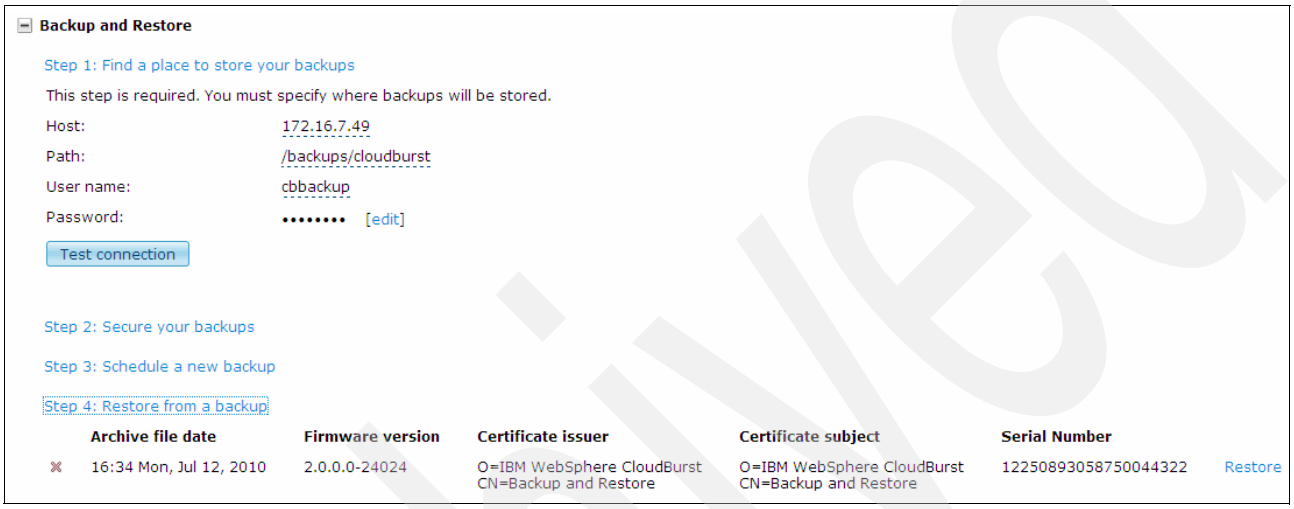

Figure 6-16 Restore from backup

4. Select the private key from your local file system that is associated with the backup. Without the correct private key the backup cannot proceed. In our example we used the private key that was generated in the backup example, backup-071210 (Figure 6-11 on page 105).

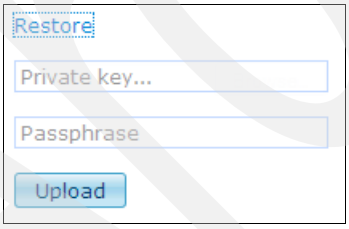

Figure 6-17 Specify the private key on local system

5. Enter the password for the private key (Figure 6-18). This was specified in Step 2: Secure your backups, when the backup was originally taken.
6. Click **Upload** (Figure 6-18). This process checks with the private key and the associated password are correct. If either the private key or the password is incorrect, an error message displays.

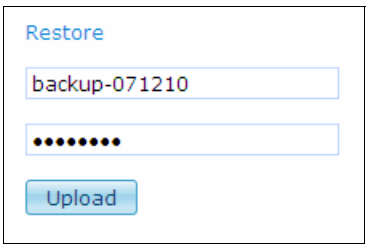

Figure 6-18 Enter password for private key

7. Click **Restore** (Figure 6-19).

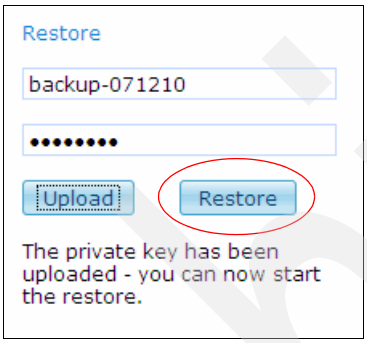

Figure 6-19 The Restore button

In our test scenario the restore procedure took approximately 8.5 hours. This was for restoration of a default WebSphere CloudBurst Appliance with no customization.

6.2.2 Firmware updates

Firmware updates can be downloaded from IBM Fix Central:

`http://www.ibm.com/support/fixcentral`

The update file is a single file with a .script2 extension. To perform an update:

In the following instructions we assume that the user is logged in as the cbadmin user and the required firmware update has been downloaded from IBM Fix Central. Access the Settings page by selecting **Appliance** → **Settings** on the menu bar.

1. Click **Firmware** to expand the firmware settings and display the current firmware level (Figure 6-20).
2. Click in the text box labelled **Browse** (Figure 6-20).

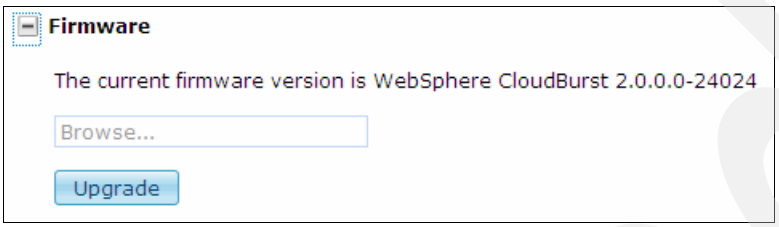

Figure 6-20 The expanded firmware update section

3. From your local machine select the `.script2` file that was downloaded from IBM Fix Central website and click **Open**. In our scenario the file is called `WebSphere_Cloudburst_Appliance-2.0.0.0.scrypt2` (Figure 6-21).

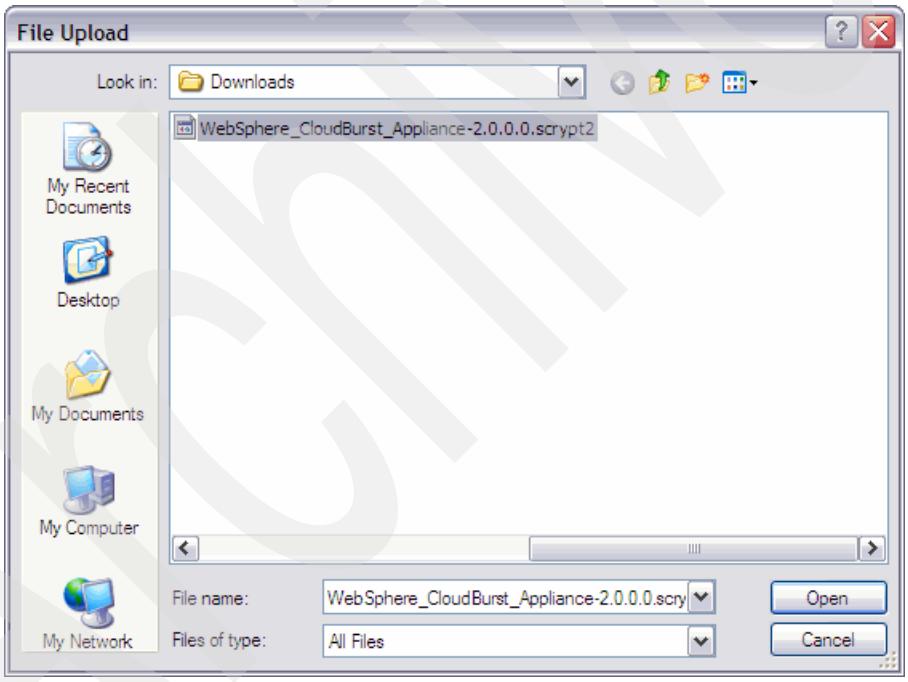

Figure 6-21 Upload the updated file located on the local machine

4. Click **Upgrade** (Figure 6-22). The upload speed of the update depends on you network connection. A message displays when the upload completes.

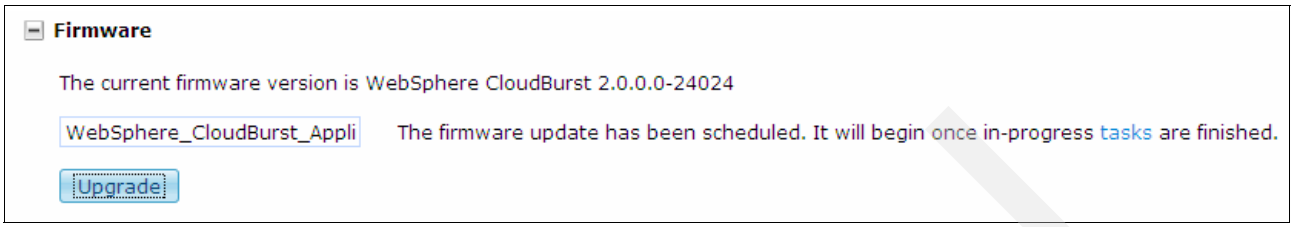

Figure 6-22 Upgrade the appliance

5. The update has started. As the update starts, the WebSphere CloudBurst Appliance restarts. The web console does not indicate the progress of the update because the session has terminated. The web console is not available while the update is preceding. In our scenario, our update and reboot took 8 minutes.
6. Log into the console as the cbadmin user when the update has completed and confirm the new firmware level by navigating to **Appliance** → **Settings** on the menu bar and expanding the **Firmware** section. In our scenario we updated the firmware to the same level that was previously running, 2.0.0.0-24024 (Figure 6-23).

Figure 6-23 Current firmware level

6.2.3 Shutdown/restart

There can be certain times during the normal operation of a WebSphere CloudBurst Appliance when a manual restart is required (for example, changing the session time-out value for users requires the appliance to be restarted).

To restart the appliance:

1. Access the Settings page by selecting **Appliance** → **Settings** on the menu bar.
2. Click the **Power** link to expand the options.
3. Click **Restart the appliance** (Figure 6-24).

Figure 6-24 The power options

4. The appliance can be restarted immediately or you can wait until all the currently active tasks have been completed. In our scenario we choose to wait until active tasks are completed. Click the **Wait until active tasks have completed** radio button and click **OK** (Figure 6-25).

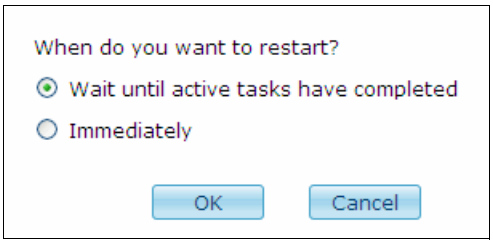

Figure 6-25 Confirmation of a restart

5. The WebSphere CloudBurst Appliance restarts, when any currently tasks are completed. You must log in to the appliance again after this step is complete, as your session expires upon the restart.

6.2.4 Monitoring

Monitoring of a WebSphere CloudBurst Appliance is performed using SNMP traps. This gives the WebSphere CloudBurst Appliance the flexibility to be monitored quickly and easily in enterprise environments. You can download the MIB packages from the appliance for use with an SNMP client so that it can easily understand any traps. The WebSphere CloudBurst Appliance has the ability to change the severity of trap codes, so companies can customize for their environments. To access the monitoring settings, select **Appliance** → **Monitoring** on the menu bar.

SNMP traps were not configured in our test scenario, as there were no monitoring servers in the environment (Figure 6-26).

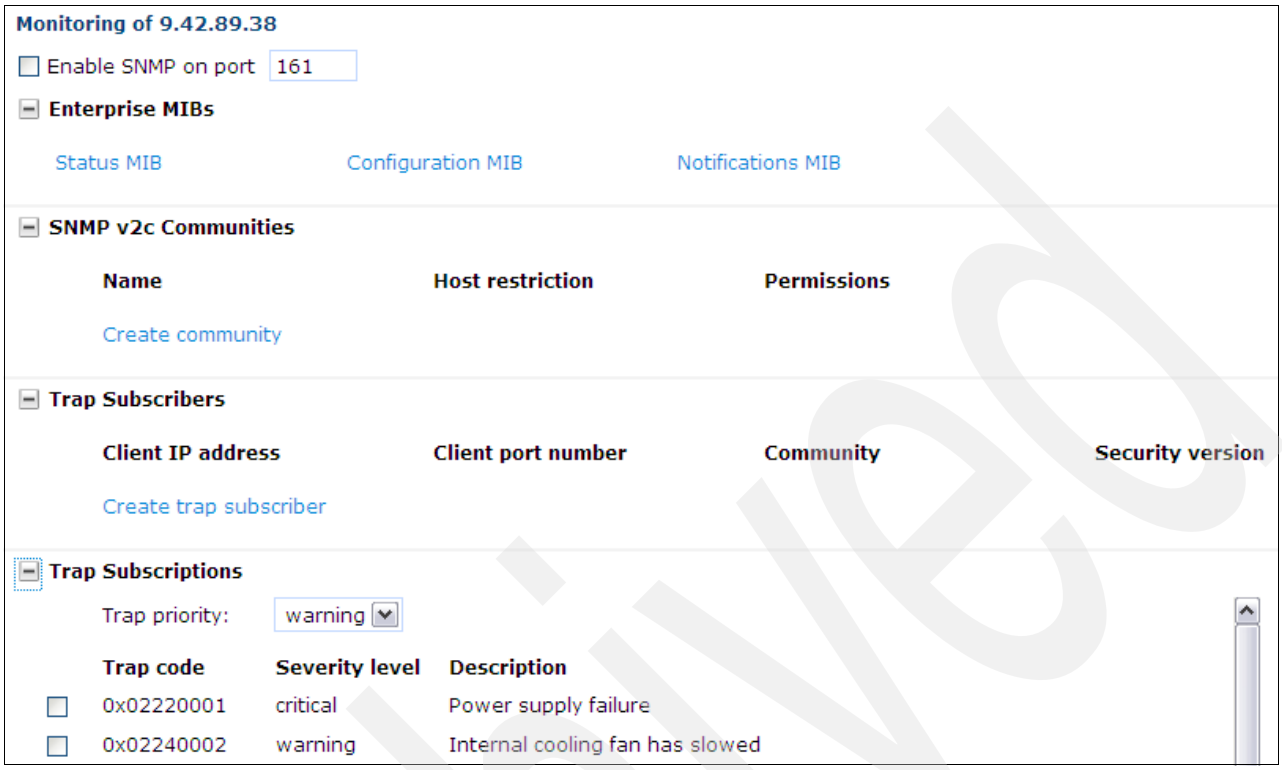

Figure 6-26 The monitoring settings in WebSphere CloudBurst Appliance

6.2.5 Logs

Using the log file functionality of WebSphere CloudBurst Appliance allows users a quick way to check errors and help troubleshoot basic problems. A complete set of logs can be downloaded from the appliance and sent to IBM support in the instance of a major issue.

There are two standard logs stored on the appliance that are easily viewable by users, `error.log` and `trace.log`. Each log file has a maximum size limit and a rotation policy.

- `error.log` has a maximum of 2 MB in size and have a maximum of five versions.
- `trace.log` has a maximum file size of 100 MB and has a maximum of 10 versions.

Any log file greater than 30 days old is removed.

To view the log files:

1. Select **Appliance** → **Troubleshooting** on the menu bar.
2. Click **Logging** to expand the log options.

3. Click the **View current error file** link to view `error.log` or **View current trace file** to view `trace.log` (Figure 6-27).

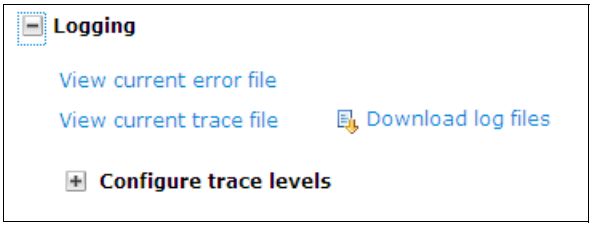

Figure 6-27 The log file options

A new browser tab opens displaying the log file (Figure 6-28). Clear and Pause links display at the top of the window. You can click Pause to stop log entries from being appended to the log file. You can click **Clear** to delete all data from the log viewer.

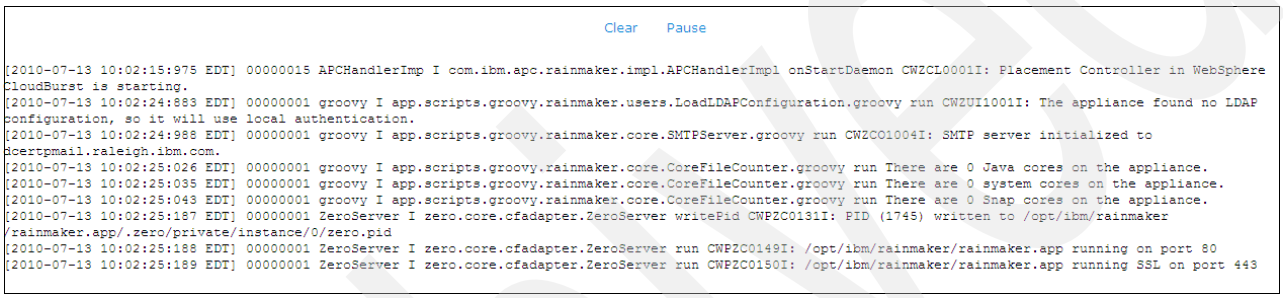

Figure 6-28 Error log file

In the event that all log files are needed or when lodging a support call with IBM, click the **Download log files** link. This link allows you to download all log files in a compressed format to your local machine. When lodging a call with IBM, also include the machine type and model number of the IBM POWER system used for the PowerVM hypervisor.

6.2.6 Auditing

The auditing information stored on the WebSphere CloudBurst Appliance contains records for user activity. This includes creation/updates/deletes of virtual systems, cloud groups, and patterns and who performed the action.

Table 6-1 shows all auditable objects that are available to the user.

Table 6-1 Auditable objects

Object	Add/create	Delete	Update	Other
Virtual system	Y	Y	Y	
Virtual machine	Y	Y	Y	
Task	Y	Y	Y	
Hypervisor	Y	Y	Y	
IP	Y	Y	Y	
IP group	Y	Y	Y	
Storage	Y	Y	Y	

Object	Add/create	Delete	Update	Other
Cloud	Y	Y	Y	
Network	Y	Y	Y	
Script package	Y	Y	Y	
Emergency fix	Y	N/A	N/A	Add part, add script, clone
Pattern	Y	Y	Y	
Virtual image	Y	Y	Y	
Disk image	Y	Y	Y	
User	Y	Y	Y	
Group	Y	Y	Y	
Session	N/A	N/A	N/A	Timeout
Backup profile	Y	Y	Y	
Physical machine CPU utilization	N/A	N/A	Y	
Physical machine memory utilizations	N/A	N/A	Y	
Virtual machine CPU utilization	N/A	N/A	Y	
Virtual machine memory utilization	N/A	N/A	Y	
Storage utilization	N/A	N/A	Y	
IP utilization	N/A	N/A	Y	
Daily summary of hypervisors	N/A	N/A	Y	
Daily summary of virtual systems	N/A	N/A	Y	
Daily summary of virtual machines	N/A	N/A	Y	
Daily summary of patterns	N/A	N/A	Y	
Daily summary of script packages	N/A	N/A	Y	
Daily summary of virtual images	N/A	N/A	Y	
Daily summary of users	N/A	N/A	Y	

There are two choices when downloading auditable objects:

- Download all data.
- Download via date and time range.

Selecting either choice creates a file called `audit.zip`, which downloads to your local PC. The `audit.zip` file contains three comma-separated value (CSV) files (`appliance-audit.csv`, `license-audit.csv`, and `pvu-audit.csv`) for ease of use with a spreadsheet application.

To download the audit files perform the following tasks:

1. Select **Appliance** → **Troubleshooting** on the menu bar.
2. Click the **Auditing** link to expand the auditing options:
   a. Click the **Download all data** hyperlink to download all the data (Figure 6-29).
   b. Adjust the appropriate data range and time zone settings and click the **Download filtered data** hyperlink.

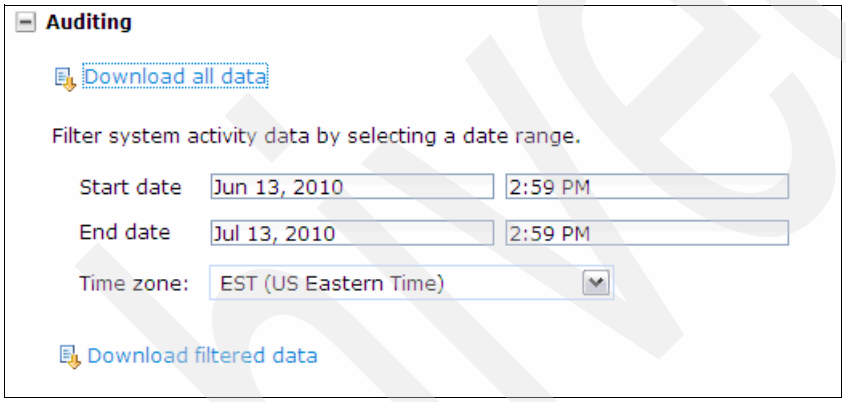

Figure 6-29 The auditing options

3. Save the `audit.zip` file to your local machine so that you can analyze the data (Figure 6-30).

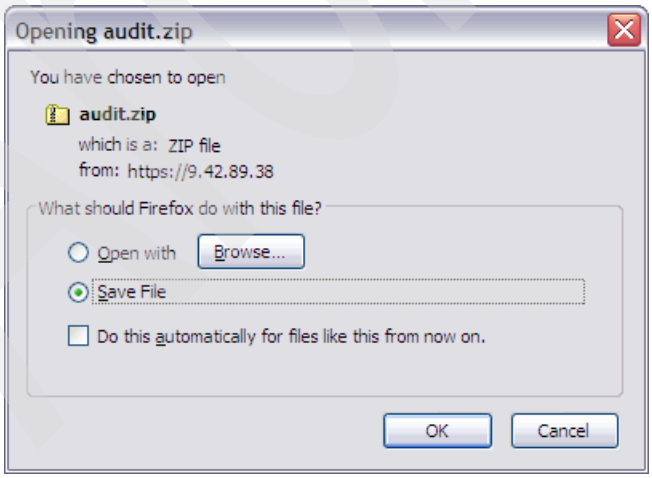

Figure 6-30 Audit.zip file

6.2.7 Hardware information

To quickly view the hardware status of a WebSphere CloudBurst, the web console contains a quick summary of the state of the hardware capacity and the internal hardware temperatures.

The hardware capacity section contains a set of color-coded bar graphs (green for healthy and red for danger) on the following items:

- Memory usage
- Encrypted disk space
- Temporary disk space
- Hard disk space

To view the hardware capacity of the WebSphere CloudBurst Appliance:

1. Select **Appliance** → **Troubleshooting** on the menu bar.
2. Click the **Hardware Capacity** link to view the hardware capacity (Figure 6-31).

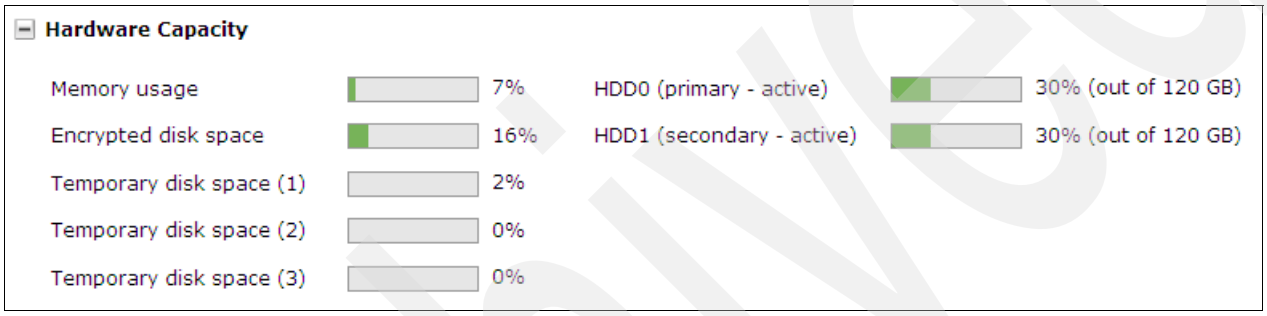

Figure 6-31 Current hardware capacity

The hardware temperature section contains temperature information, in fahrenheit, and color-coded icons to represent the temperature of the system. If the temperature is in a normal operating range, green icons displays. If the temperature is outside normal operating range, yellow icons display. The temperature sensors measure the temperature of the following internal components:

- System
- CPUs
- Memory modules

To view the internal hardware temperature of the WebSphere CloudBurst Appliance:

1. Select **Appliance** → **Troubleshooting** on the menu bar.
2. Click the **Hardware Temperatures** link to view the hardware capacity (Figure 6-32).

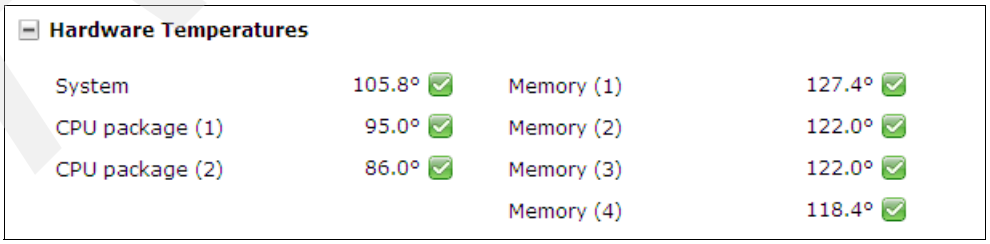

Figure 6-32 The internal hardware temperature

6.2.8 Ping hosts

To help confirm network connectivity, troubleshoot network connection issues and isolate a network problem, the WebSphere CloudBurst Appliance includes a ping facility.

To use the ping tool:

1. Select **Appliance** → **Troubleshooting** on the menu bar.
2. Click **Outbound Connections** to expand the tool.
3. Type the host name or IP address in the text box and click **Ping** (Figure 6-33).

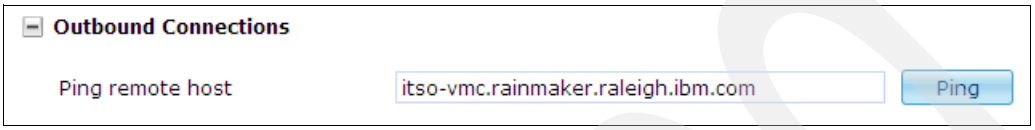

Figure 6-33 The ping tool

4. A successful ping displays a green box with a check mark on the bottom right corner of the panel (Figure 6-34).

Figure 6-34 A successful ping

An unsuccessful ping displays a red circle with an X (Figure 6-35).

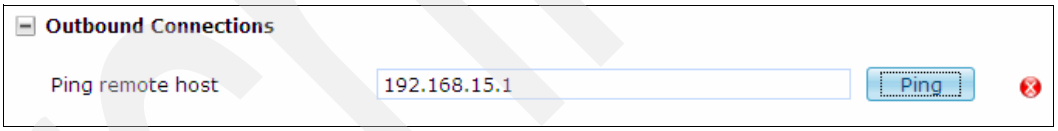

Figure 6-35 An unsuccessful ping

6.2.9 User and group administration

One of the most powerful aspects of a WebSphere CloudBurst Appliance is its ability to behave as a self-service portal, which allows users to deploy complex WebSphere Application Servers quickly and efficiently using repeatable patterns.

This ability is controlled through the user administration interface of WebSphere CloudBurst Appliance. The user administration allows the creation of users and groups and the ability to allow or restrict access to WebSphere CloudBurst Appliance objects.

To create a new user you must be assigned the appliance administration role with full permission. In our scenario we used the cbadmin user to create the new users.

To create a new user:

1. Select **Appliance** → **Users** on the menu bar.
2. Click the green + icon to create a new user.

3. Enter the user name, full name, password, and email address and click **OK** (Figure 6-36). If security is managed by an LDAP server, the password field does not display, as the password in the LDAP directory is used for authentication.

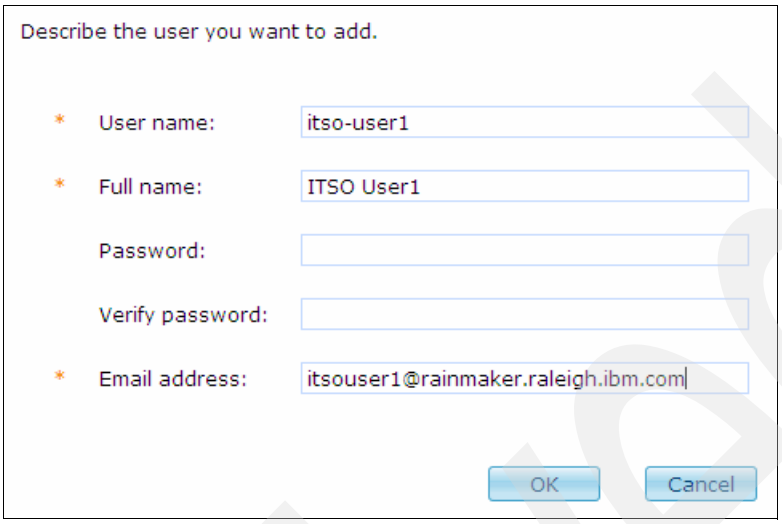

Figure 6-36 New user details

If the password field is left empty, a password is automatically generated and emailed to the user (Figure 6-37). (The SMTP server must be configured for this to work.)

```
An account has been created for you on the WebSphere CloudBurst
Appliance. Your user name and password are below.

User name: itso-user1
Password: wCU1JdxT

To login to the appliance, please visit https://9.42.89.38.
```

Figure 6-37 Email from the CloudBurst Appliance

A new user has been created with the most restrictive security policy. This policy only allows the user to deploy patterns into a cloud.

In our test scenario we created five basic users (Figure 6-38):

- itso-user1
- itso-user2
- itso-user3
- itso-user4
- itso-user5

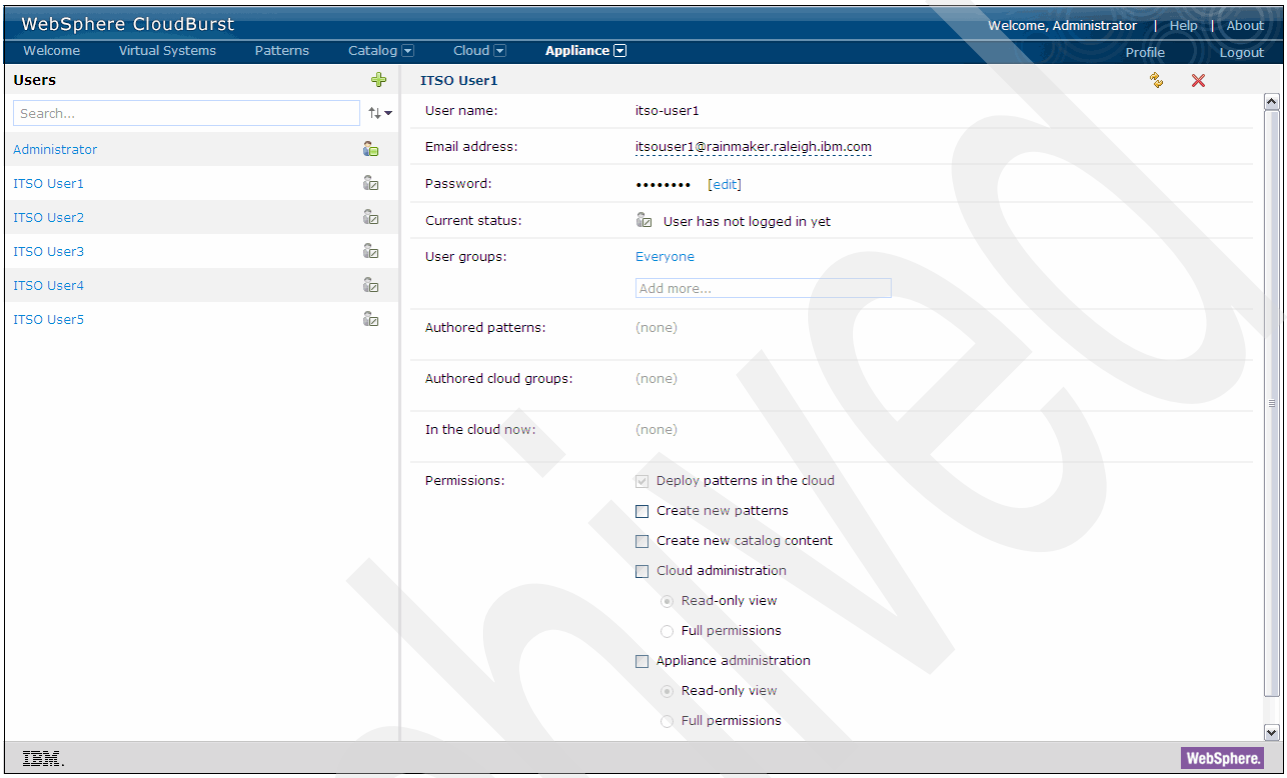

Figure 6-38 itso-user1 properties and list of all users

Extra permissions to objects in the WebSphere CloudBurst Appliance can be given to an individual user. A better way to manage users with privileges is to add them to a privileged group. All permissions granted to a group are inherited by the group members. This makes management easier to manage, especially when there are multiple users in a group. WebSphere CloudBurst Appliance groups contain exactly the same permissions as individual users. The permissions are:

- Create new patterns.
- Create new catalog content.
- Cloud administration:
  - Read-only view
  - Full permissions
- Appliance administration:
  - Read-only view
  - Full permissions
- IBM License Metric Tool.

In our test scenario we create a privileged group called ITSO Administrators that can create new patterns, add new catalog content, access the IBM License Metric Tool (ILMT), and have read-only access to cloud administration. ITSO User1 is the full name of our administrative user in this scenario and will be added to the ITSO Administrators group.

To create a group:

1. Select **Appliance** → **User Groups** on the menu bar.
2. Click the green plus icon (**+**) to create a new group.
3. Enter the group name and description in the text box and click OK (Figure 6-39). In our scenario the group name is ITSO Administrators and the description is The ITSO Administrators.

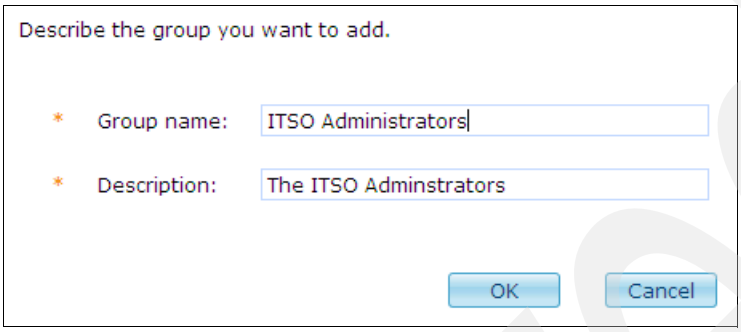

Figure 6-39 Create a new group

4. A new groups is created. Click the new group **ITSO Administrators**.
5. Click the check box next to the permissions that we require. In our scenario, we check the boxes for the following hardware:
   - Create new patterns
   - Add new catalog content
   - Cloud Administration: Read-only view
   - IBM License Metric Tool (Figure 6-40)

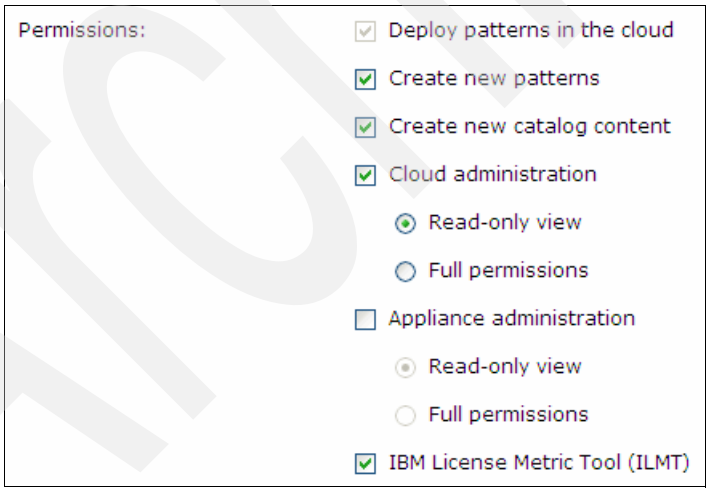

Figure 6-40 Extra permissions given to the group ITSO Administrators

6. Click in the Group members text box, select the user to add to the group. In our scenario, we selected ITSO User1 (which is the full name of user ID itso-user1) (Figure 6-41).

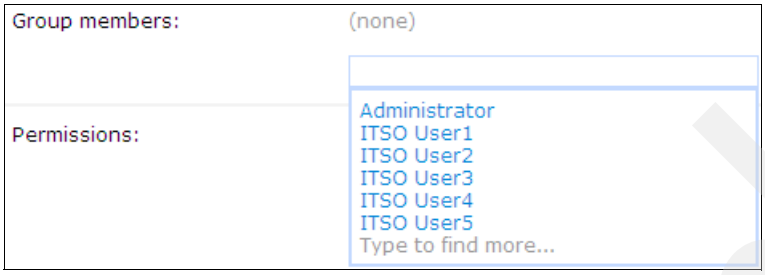

Figure 6-41 Add user ITSO User1 to group ITSO Administrators

The itso-user1 user has been added to the ITSO Administrators group and has inherited all of the permissions granted to the group (Figure 6-42).

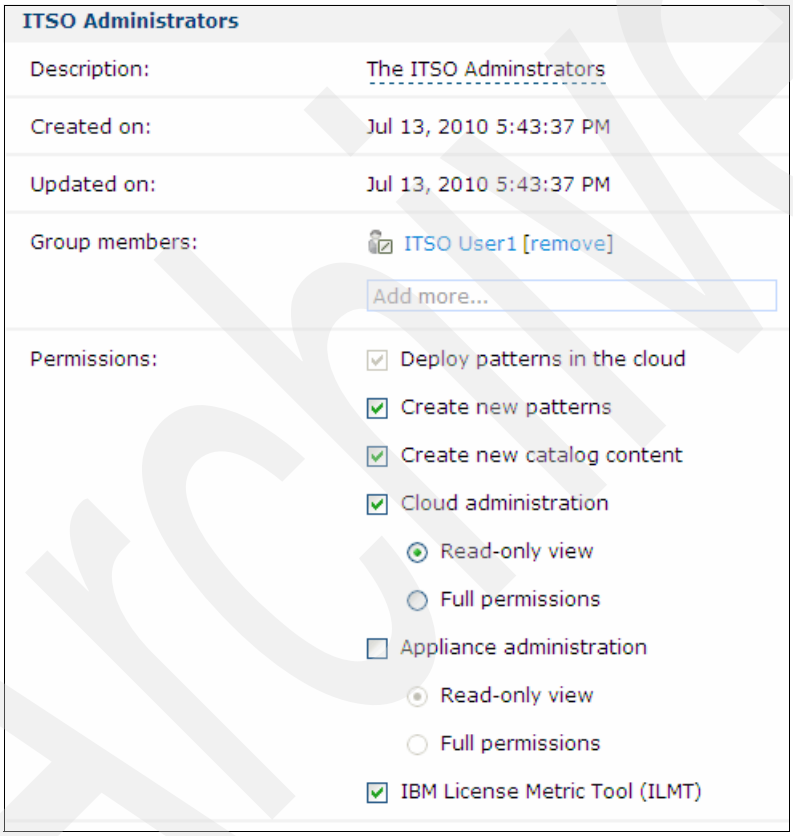

Figure 6-42 The completed ITSO Administrators group

6.2.10 Virtual systems

A virtual system is a deployed pattern made up of one or more LPARs that contain a component of the overall system. Figure 6-43 shows examples of different virtual systems. In the example, virtual system 1 is made up of one LPAR called server1. Virtual system 2 is made up of three LPARs (server2, server3, and server4), each containing different components of the entire virtual system 2.

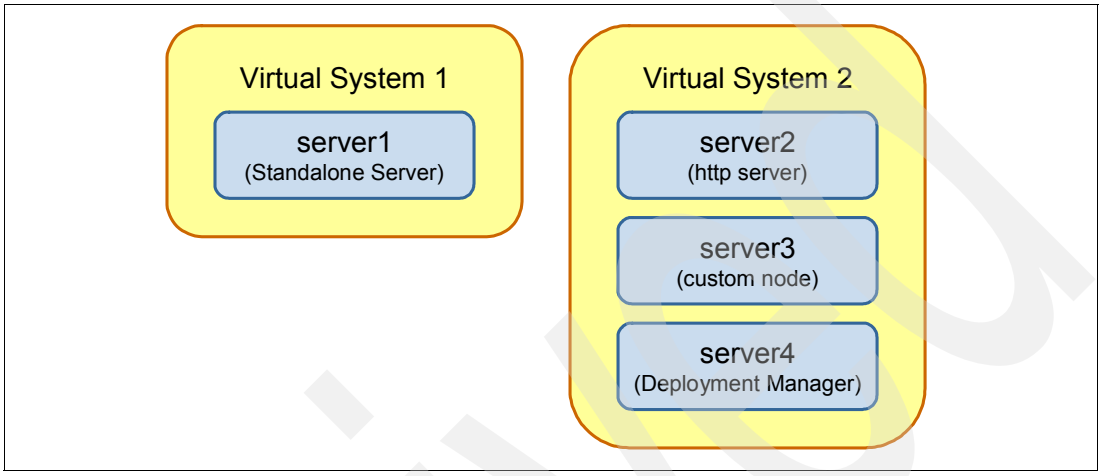

Figure 6-43 Examples of virtual systems

To view and administer virtual systems click **Virtual Systems** on the menu bar. A summary view of virtual systems displays. The state of the virtual system displays next to the virtual system name via an icon, green for started and red for stopped. You are only able to view systems to which you have been given access.

Click a virtual system to view the properties of the virtual system (Figure 6-44). In our scenario we chose to view test system itso-cluster-dev, which is a virtual system of a predefined pattern *WebSphere cluster (development)*, which comes with the WebSphere CloudBurst Appliance (The pattern has been cloned and changed the target platform to PowerVM. For further details about how to do this refer to Chapter 7, “Working with patterns” on page 157.

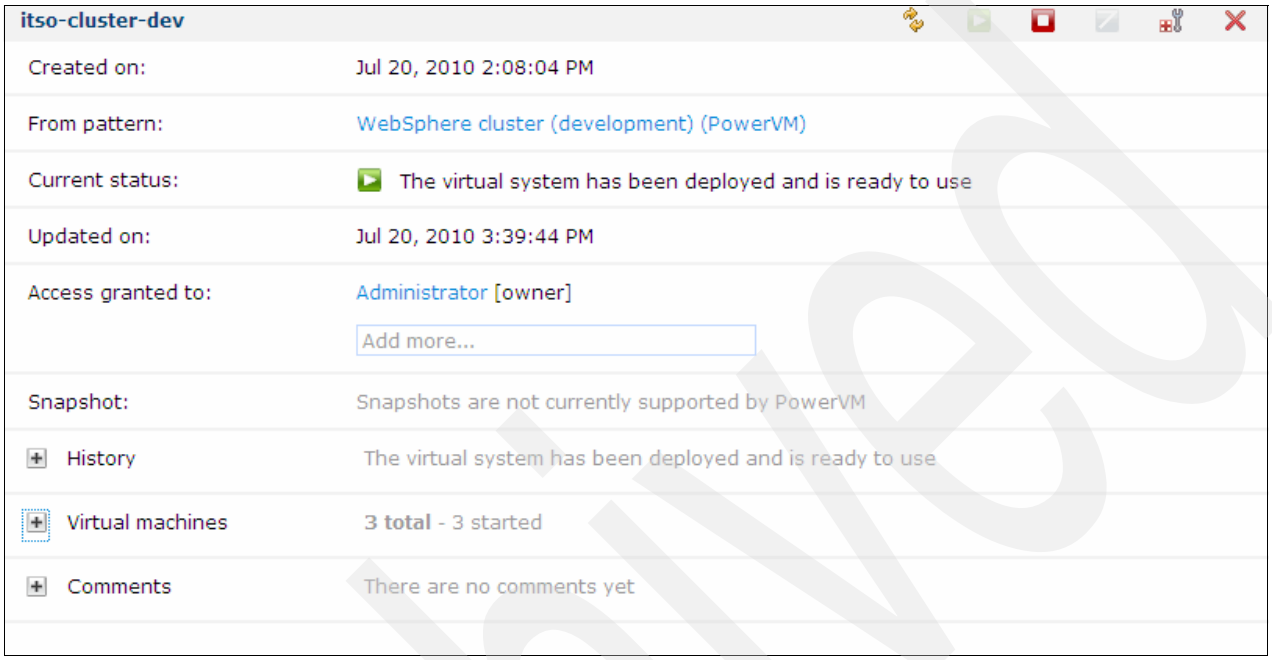

Figure 6-44 Virtual systems properties

The top right corner of the virtual systems view provides functions for start, stop, store apply service, and delete virtual systems (Figure 6-45):

- Start virtual system: starts all the LPARs that make up the virtual system.
- Stop virtual system: stops all the LPARs that make up the virtual system.
- Store virtual system: releases all of the hypervisor resources and keeps the assigned IP addresses and the virtual system remains on the hypervisor. All other resources (that is, memore, CPU, PVUs) are released.
- Applying a service: allows you to apply an emergency fix or service pack to a system.
- Delete virtual system: removes the virtual system from WebSphere CloudBurst and deletes all virtual system resources including disk images and IP addresses, and returns the resources to the pool.

Figure 6-45 The basic virtual system functions

Other basic information is provided by the virtual system view, which includes the creation date, the pattern the virtual system was deployed from, the status, and access rights (Figure 6-46). Snapshots are not currently supported for PowerVM on WebSphere CloudBurst Appliance 2.0.

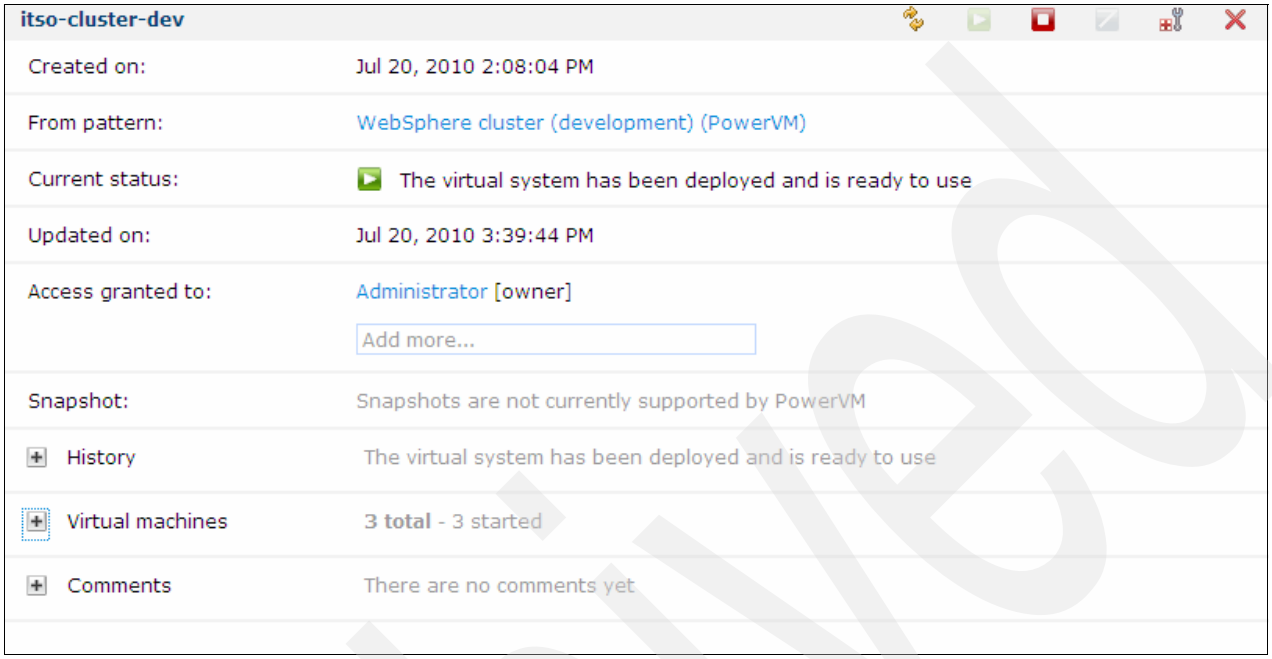

Figure 6-46 Basic virtual system information

The Access granted to field on the information page allows access privileges to be added to or removed from users and groups. Access privileges include read, write, and all privileges.

- Read access only allows users to view the virtual system in the catalog.
- Write access allows users to start, stop and store the virtual image and apply fixes.
- All access grants the same permission as the owner of the application.

Unlike most other link-style buttons in WebSphere CloudBurst Appliance, the permission listed next to the user's name (read, write, or all), is the current permission granted to the user (Figure 6-47). To change the permission click the link and it changes to the next permission. Clicking the link again rotates the permission to its third permutation.

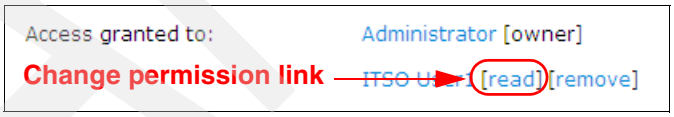

Figure 6-47 Change permission link

To allow a user or group access to a virtual system:

1. Click the **Add more** text box (Figure 6-48)

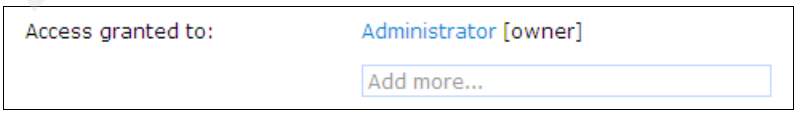

Figure 6-48 Add more text box

2. Select the user or group. In our scenario we chose ITSO User2 (itso-user2) (Figure 6-49).

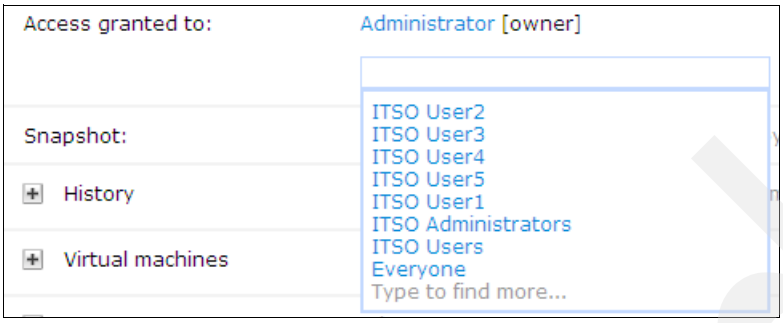

Figure 6-49 Select user or group

3. By default the user or group only has read access, as can be seen by the permission link displaying read. In our scenario we want to give the ITSO User2 write access to the virtual system. Click the permission link **read** to change the permission, so the link says `write`.

The user now has write access to the virtual system (Figure 6-50).

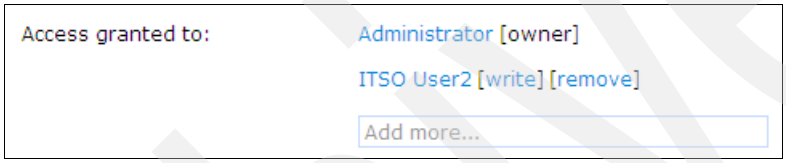

Figure 6-50 Write permission has been granted to user

Each virtual system keeps a history of the significant events since it was created. The events recorded include virtual systems startup, shutdown, certain errors and resource allocation. Click the **History** section to expand and view the history panel (Figure 6-51).

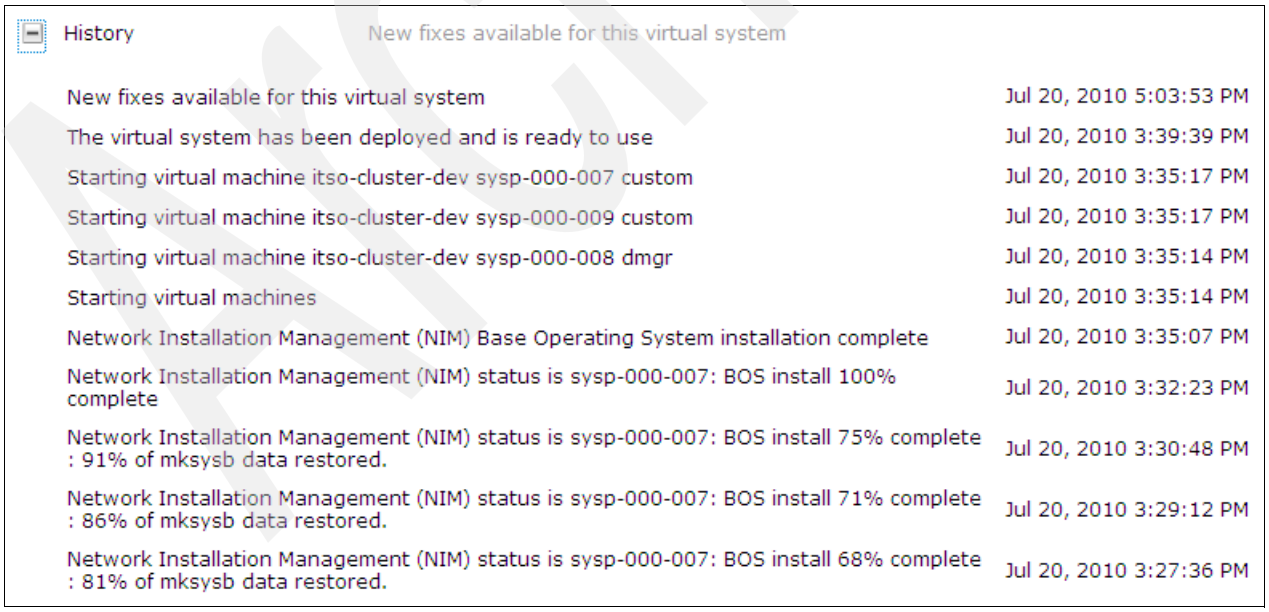

Figure 6-51 Virtual system history

Note: The WebSphere CloudBurst UI uses the words *virtual machine* to identity Logical Partitions (LPARs). All references in this publication to the words *virtual machine* have the same meaning as LPAR.

Each virtual system is made up of one or more LPARs/virtual machines. The number of virtual machines that are in the virtual systems is outlined in the pattern from which the virtual system was deployed. Click the **Virtual machines** section to expand and list the virtual machines that make up the virtual system. In our example, our virtual system was deployed from the predefined WebSphere cluster (development) pattern, which is made up of three virtual machines (Figure 6-52).

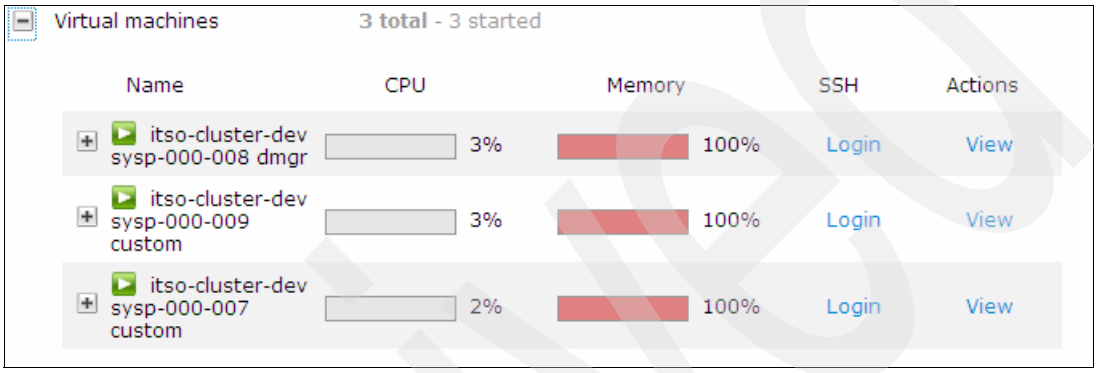

Figure 6-52 Virtual machines in the virtual system

Each virtual machine in the system is listed by host name, its role in the virtual system, and an indicator for CPU and memory consumption.

It is possible to perform start, stop, clone, and delete actions on the individual virtual machines that make up the virtual system from this view. To perform an action click the **View** link for the associated virtual machine and select the icon associated with the action (Figure 6-53). Certain actions cannot be performed on a virtual machine unless it is in a particular state (that is, to clone or delete a virtual machine, the virtual machine must be in a stopped state). It also depends on they type of the part (for example, it does not allow you to clone a deployment manager part).

Figure 6-53 Perform an action on a virtual machine

It is also possible to open an SSH console session to each individual virtual machine by clicking the **Login** link next to the associated virtual machine. When connecting via SSH to a virtual machine this way, the WebSphere CloudBurst Appliance behaves as a proxy between your web browser and the virtual machine. The console does not connect your computer to the virtual machine.

To log in via ssh:

1. Click the **Login** link next to the associated virtual machine.
2. Provide the user name and password details and click **Login** (Figure 6-54).

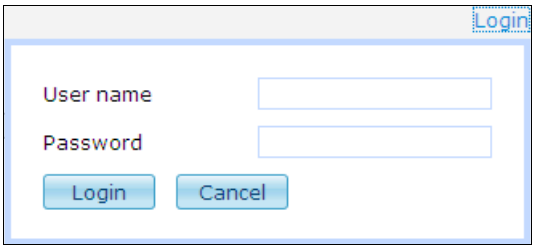

Figure 6-54 Log in via ssh

3. A new browser window opens, starting a console connection to the target virtual machine (Figure 6-55).

```
Last login: Tue Jul 20 15:38:49 2010 on ssh from 172.16.0.227
IBM WebSphere Application Server Hypervisor Edition for IBM AIX
# #
```

Figure 6-55 SSH console connection to target virtual machine

Each virtual machine that makes up the virtual system can be expanded to display properties about the individual virtual machine. Click the + icon next to the virtual machine to view the properties of the virtual machine. The information available in the display window is:

- General Information: where the hypervisor virtual machine is located, storage information, virtual machine status
- IBM Products: the installed IBM products and license usage
- Hardware and Network: IP addresses, host name, CPU, and memory
- Operating System: name type and version of the operating system
- WebSphere Configuration: cell/node/profile name, environment variables
- Script packages: download log files
- Consoles: direct link to the WebSphere web console

Figure 6-56 displays a summary of a virtual machine.

Figure 6-56 Virtual machine summary

6.2.11 License management

A cloud infrastructure can be a volatile environment, where application servers are being built and removed all the time. This can lead to license infringements or the opposite situation, low utilization of licenses.

WebSphere CloudBurst Appliance is a license-aware appliance and provides the capability to track license usage, generate reports, enforce licensing rules, and integrate with the IBM License Metric Tool. The IBM License Metric Tool is a no-charge product, which is installed on its own server, for tracking licenses of IBM software products. In our test environment we did not use the IBM License Metric Tool.

To view the license management, click **Catalog** → **Product Licenses** on the menu. The product licenses window that opens is made up of two components (Figure 6-57):

- The IBM License Metric Tool configuration
- Native WebSphere CloudBurst Appliance license tracking tools

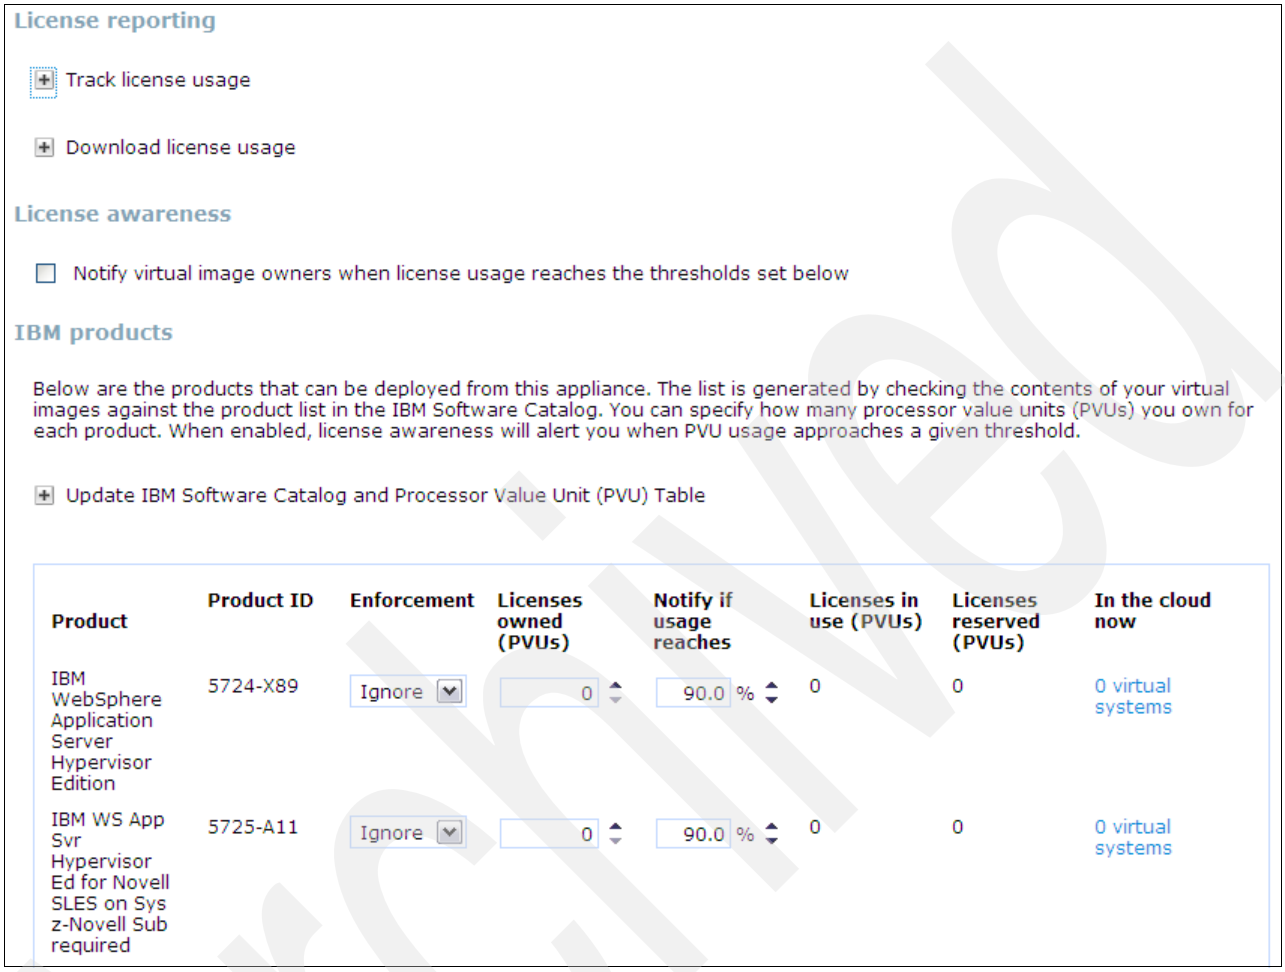

Figure 6-57 WebSphere CloudBurst Appliance license usage catalog

Understanding licensing in the cloud

To fully understand how license usage is calculated in a virtualized environment. IBM uses a unit called a processor value unit (PVU). The PVU is a pricing structure for software that takes into account the type of processor. Different processors have different PVU values.

In a virtualized cloud environment that is managed by WebSphere CloudBurst Appliance, two pieces of information are required to calculate the number of PVUs being used:

- The per core PVU score
- The number of processor cores that a product can use

The numbers are multiplied together to come up with how many PVUs are being used.

As an example, we have two physical machines in the cloud, each with two dual-core processors, making a total of four cores per physical machine (Figure 6-58). In the example, assume that each core has a PVU value of 120.

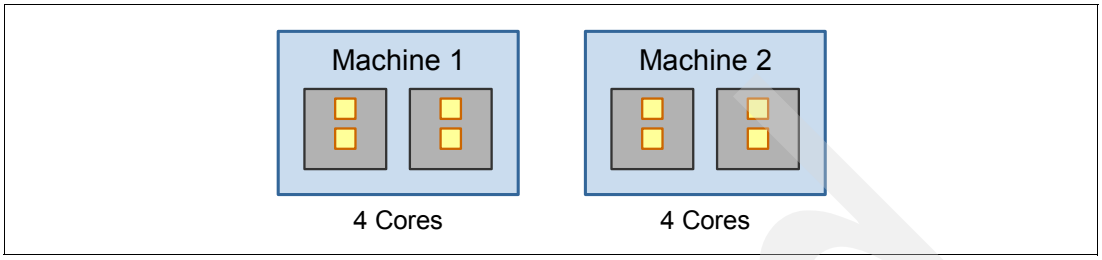

Figure 6-58 Physical servers in cloud

Four virtual machines are placed in the cloud, two with two virtual processors and one with a single virtual processor (or core) (Figure 6-59).

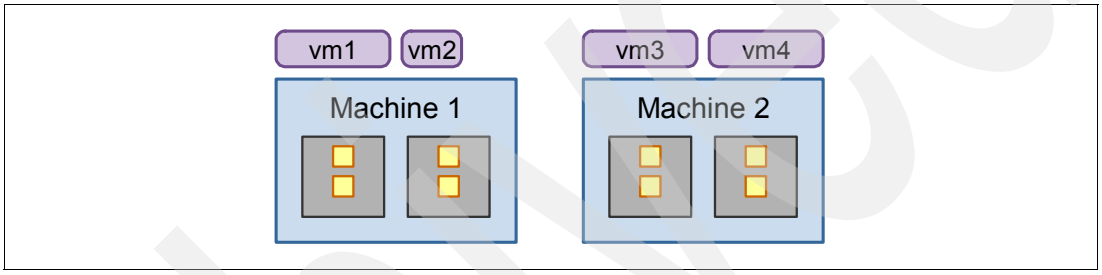

Figure 6-59 Virtual machines in the cloud

The total PVU usage is PVU score x number of cores in use. In our situation these values are 120 x 5 = 600.

Another four virtual machines are deployed, three with two virtual processors and one with a single virtual processor (Figure 6-60).

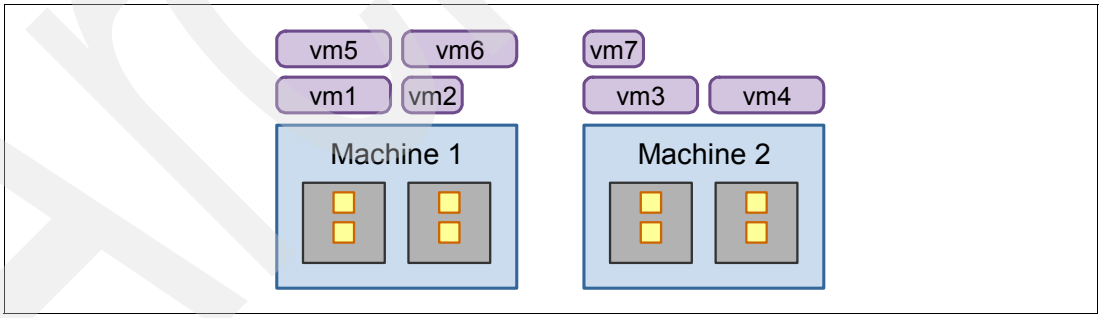

Figure 6-60 More virtual machines in the cloud

The PVU usage for this scenario is 120 x 8 = 960. The PVU license usage value cannot exceed the number of cores on the physical hardware, no matter how many virtual machines are deployed.

To help better understand, let us add another physical machine to the cloud, which has two dual-core processors, four cores in total (Figure 6-61).

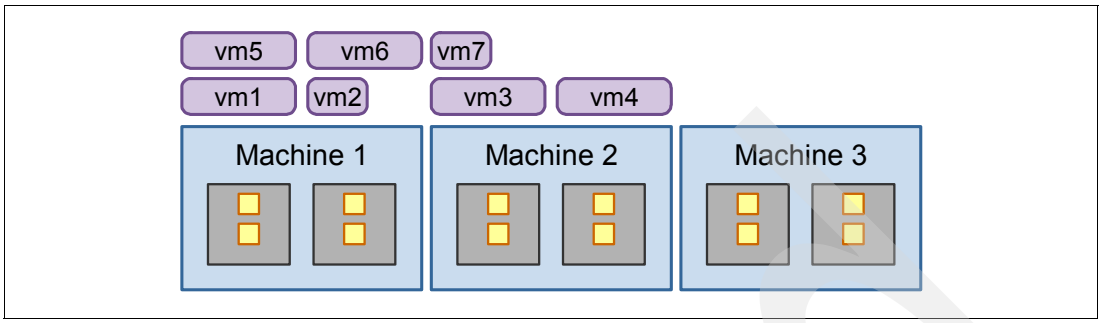

Figure 6-61 Extra physical machine in the cloud

The PVU usage depends on where the virtual machines are located. In the scenario depicted in Figure 6-62, the PVU license usage is still 960. If vm7, which contains a single virtual processor, is moved to machine3 (Figure 6-62), the PVU license usage changes and becomes 120 x 9 = 1080. This is because all virtual machines running in the cloud now have access to a total of nine cores.

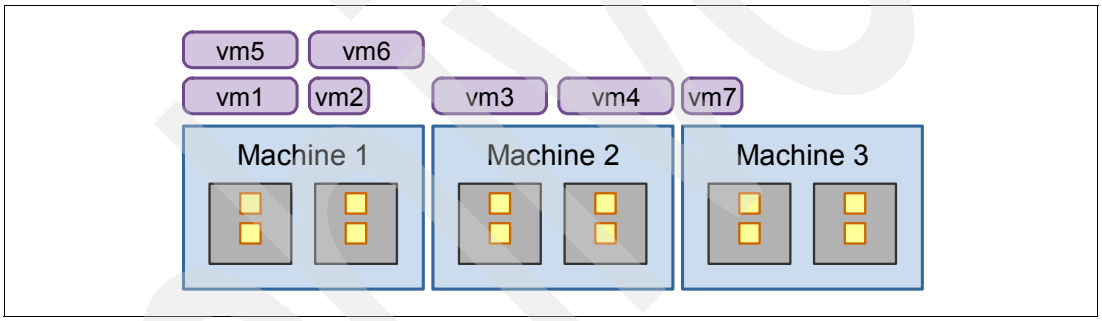

Figure 6-62 Virtual machine move

During deployment of a new virtual system, WebSphere CloudBurst Appliance analyzes the running systems in the cloud and uses an algorithm that attempts to minimise PVU usage and balance workloads across physical machines in the cloud.

Note: For more information about PVU licensing refer to:

http://www-01.ibm.com/software/lotus/passportadvantage/pvu_licensing_for_customers.html

http://www-01.ibm.com/software/lotus/passportadvantage/Counting_Software_licenses_using_specific_virtualization_technologies.html

IBM License Metric Tool Integration

The IBM License Metric Tool section can be viewed by clicking the **Track license usage** link (Figure 6-63). The section contains configuration options for integrating the WebSphere CloudBurst Appliance with the IBM License Metric Tool for license tracking. As this task was not performed by our team due to no IBM License Metric Tool server in our environment, see the WebSphere CloudBurst Appliance information center for more information:

http://publib.boulder.ibm.com/infocenter/wscloudb/v2r0/index.jsp

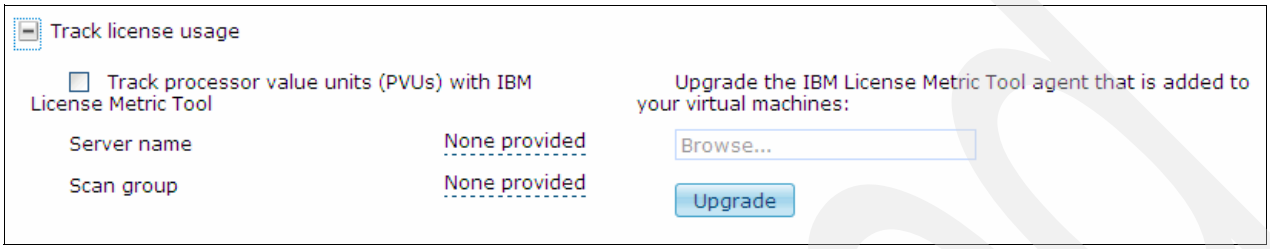

Figure 6-63 IBM License Metric Tool integration configuration

Native license tools

The native license tools that come with WebSphere CloudBurst Appliance allow you to generate audit reports, update software catalogs, update PVU tables, and create alerts based on license usage.

To make license auditing easier, the WebSphere CloudBurst Appliance has the ability to download the license usage data for cloud managed resources. To view the options available for recording and downloading license usage, click the **Download license usage** link to expand the section (Figure 6-64).

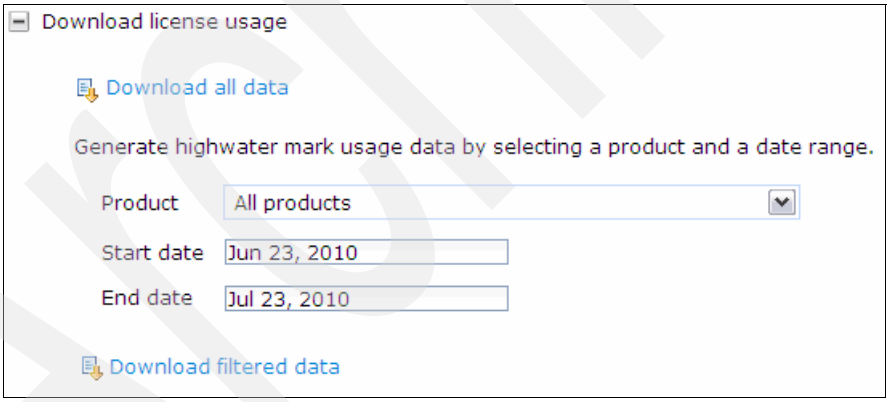

Figure 6-64 The download license usage section

The WebSphere CloudBurst Appliance tracks the maximum usage, known as the high water mark, for each product on the WebSphere CloudBurst Appliance and maintains three months' worth of data.

There are options that can be used to download the license usage data:

- Download all data: downloads the last three months of data for every product
- Download filtered data: allows you to customize the downloaded data via product and data range

When the file is downloaded, the default name for the file is `highwatermarklicense.csv`. This is a comma-separated value formatted file that can be imported into third-party software for viewing and formatting.

Predefined thresholds can be set as a specified percentage to alert administrators when licenses have surpassed the threshold. To enable this function, as it is turned off by default, select the **Notify virtual image owners when license usage reaches the thresholds set below** check box (Figure 6-65).

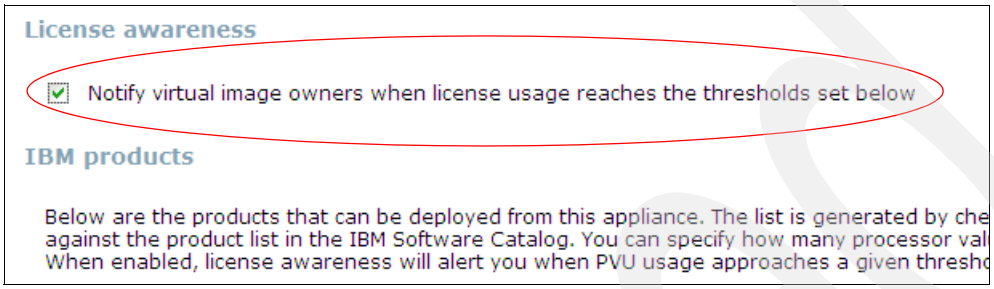

Figure 6-65 Enable email notification when license threshold reached

To change the thresholds for each product, set the Notify if usage reaches field to the required percentage that you believe is required, which is found in the license awareness section (Figure 6-66).

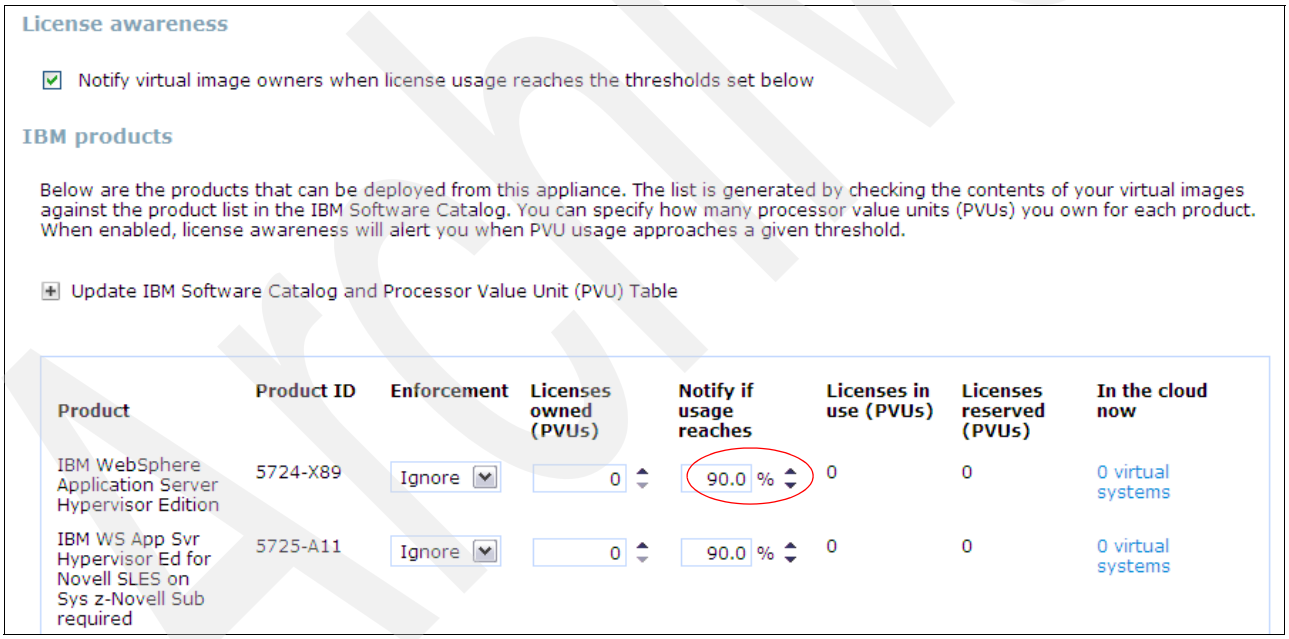

Figure 6-66 Email notification thresholds settings

By default, no license information is provided to the WebSphere CloudBurst Appliance. The license entitlement and any enforcement rules need to be completed by you. To add the number of PVU licenses owned, enter the number in the licenses owned section applicable to the product. The licence tracking also displays the current PVU usage of each product and reservation of any licenses. Enforcement rules are actions to be taken when the maximum license entitlement is reached. The actions are:

- Ignore: no enforcement action is taken and deployments can continue unhindered.
- Warn: an error is logged in the audit logs and a warning message is included in the virtual system history. Email notifications are sent to communicate the warning. Deployments can still continue unhindered.
- Enforce: Deployments of virtual systems will fail. Emails are sent to notify of unsuccessful deployments.

Figure 6-67 highlights where to find the information about licenses.

Product	Product ID	Enforcement	Licenses owned (PVUs)	Notify if usage reaches	Licenses in use (PVUs)	Licenses reserved (PVUs)	In the cloud now
IBM WebSphere Application Server Hypervisor Edition	5724-X89	Ignore	0	90.0 %	0	0	0 virtual systems
IBM WS App Svr Hypervisor Ed for Novell SLES on Sys z-Novell Sub required	5725-A11	Ignore	0	90.0 %	0	0	0 virtual systems
IBM WS App Svr Hyper Ed for Novell SLES on Sys z-Novell Sub Not required	5725-A12	Ignore	0	90.0 %	0	0	0 virtual systems
IBM WebSphere Application Server Hypervisor Edition on AIX	5725-A25	Ignore	1920	90.0 %	720	720	4 virtual systems
IBM WebSphere App Svr Hypervisor Edition for Red Hat Enterprise Linux Svr	5725-A26	Ignore	0	90.0 %	0	0	0 virtual systems
IBM WebSphere Appl Server Hypervisor Edition Intelligent Management Pack	5725-A27	Ignore	1920	90.0 %	0	0	0 virtual systems

Figure 6-67 Product license entitlement and enforcement

The software catalog and PVU table

The IBM Software Catalog and the PVU table are used together to track PVU usage of software that is deployed and in use by the WebSphere CloudBurst Appliance. Update these files regularly, as new updates are released periodically and can effect your license usage. To view the software catalog and PVU table information click **Update IBM Software Catalog and Processor Value Unit (PVU) Table** (Figure 6-68).

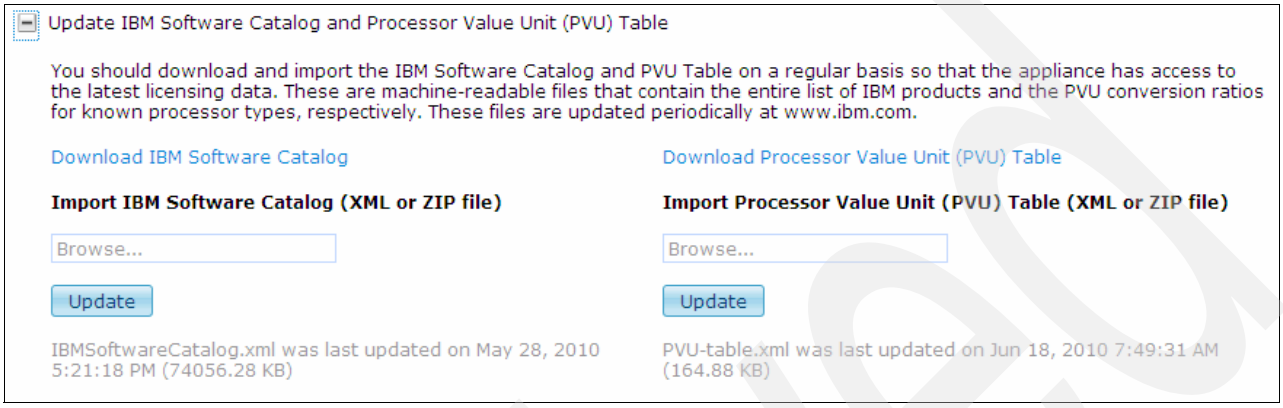

Figure 6-68 Software catalog and PVU table

To get and update the software catalog:

1. Click **Download IBM Software Catalog.** This takes you to a web page from which you can download the latest software catalog. Download the latest catalog image and save it to your local computer. In our example we used the IBM use-only catalog, called `IBMUseOnlySoftwareCatalog.zip`. It is advised for the customer to use the customer catalog. In our example in Figure 6-69, the site it is called `IBMSoftwareCatalog_canonical_form_20100630.zip`.

Figure 6-69 Software catalog download site

2. Click **Browse** and select the catalog file that you just downloaded.
3. Click **Upload** (Figure 6-70).

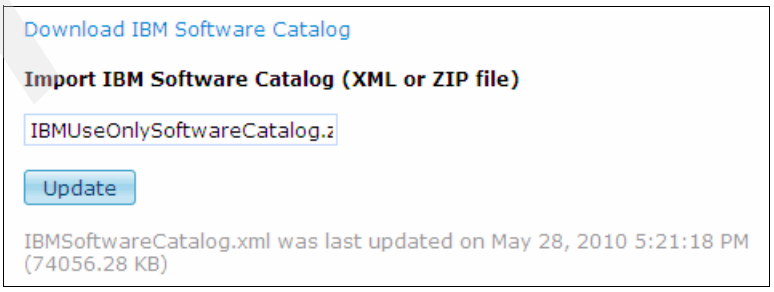

Figure 6-70 Upload the latest catalog file

This updates the catalog. A message displays stating that the catalog updated successfully (Figure 6-71).

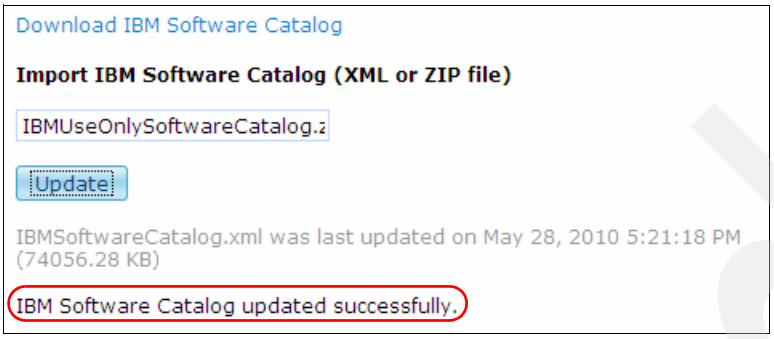

Figure 6-71 Successful update of the catalog

4. Refreshing the page displays a new date on which the catalog was updated (Figure 6-72).

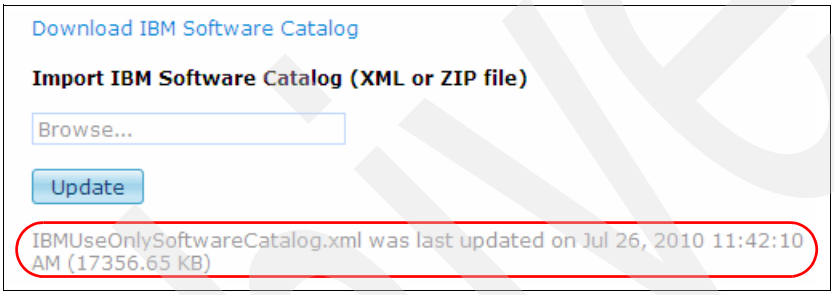

Figure 6-72 The updated date in the catalog

To get and update the PVU table:

1. Click **Download Processor Value Unit (PVU) Table.** This takes you to a web page from which you can download the latest PVU table (Figure 6-73).

United States [change]

Search

Home | Solutions | Services | Products | Support & downloads | My IBM

Welcome [IBM Sign in] [Register]

Support & downloads >

Support & downloads
IBM Support Portal
Feedback

PVU table for IBM License Metric Tool

Downloadable files

Abstract
Processor value unit table download page for IBM License Metric Tool.

Download Description
This page allows you to download the latest version of the processor value unit definition file. A processor value unit is a unit of measure that can be used to differentiate licensing of middleware on different processors. A processor value unit table is a machine-readable file that contains information about the conversion ratio used for available processor types. The conversion ratio indicates value units needed to run a specific processor. Consequently, it is necessary to ensure that the table is updated to reflect the ratios for all processors in use.

Prerequisites
IBM License Metric Tool 7.1
IBM License Metric Tool 7.2

Installation Instructions
For instructions on how to import an IBM processor value unit table update, refer to "Importing a processor value unit table" in the Licence Metric Tool information center.

Download package
If your browser opens XML when you left click to download via "FTP", please right click on "FTP" and use "Save Target As..."

What is DD?

Download	RELEASE DATE	LANGUAGE	SIZE(Bytes)	Download Options	
PVU Table	6/10/2010	US English	168860	FTP	DD

Copyright and trademark information

Document information
Product categories:
Software
Systems and Asset Management
Change & Configuration
IBM License Metric Tool
Operating system(s):
AIX, HP-UX, Linux, Solaris, Windows, i5/OS
Software version:
7.1, 7.2
Reference #:
4020292
IBM Group:
Software Group
Modified date:
2010-06-11

Translate My Page
Select Language

Rate this page
Help us improve this page. Your response will be used to

Figure 6-73 PVU table download website

2. Click **Browse** and select the catalog file that you just downloaded.
3. Click **Upload** (Figure 6-74).

Download Processor Value Unit (PVU) Table

Import Processor Value Unit (PVU) Table (XML or ZIP file)

IBM_ProcessorValueUnitTable.

Update

PVU-table.xml was last updated on Jun 18, 2010 7:49:31 AM (164.88 KB)

Figure 6-74 Upload the PVU table

This updates the PVU table. A message stating that the PVU table has been updated displays (Figure 6-75).

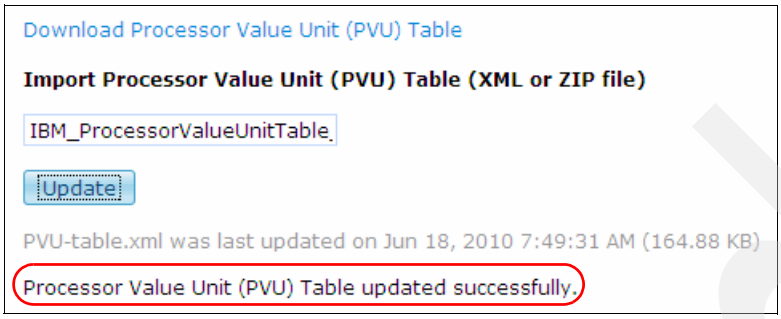

Figure 6-75 Successful update message

4. By refreshing the web page, the date that the PVU table was last updated changes to the current date (Figure 6-76).

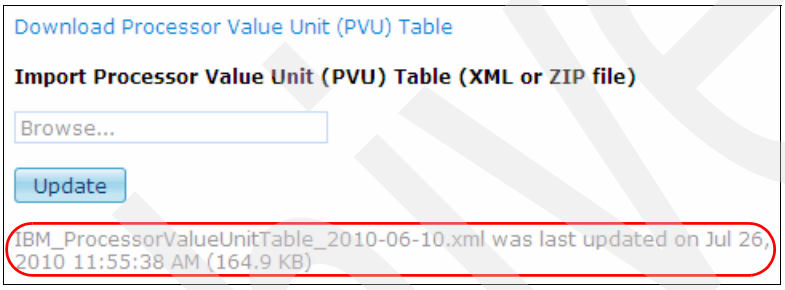

Figure 6-76 PVU table updated with latest date

6.2.12 Cloud usage reports

WebSphere CloudBurst Appliance comes with a diverse set of graphical reports, which provides administrators with an overview of the resource usage. The resource usage can be broken down by resource type and date and can be used for planning or charge back.

WebSphere CloudBurst Appliance offers a variety of reports:

- CPU Usage by Hypervisor: displays the percent of CPU used by each hypervisor.
- CPU Usage by Virtual Machine: displays the percent of CPU used by each virtual machine.
- Memory Usage by Hypervisor: displays the percent of memory used by each hypervisor.
- Memory Usage by Virtual Machine: displays the percent of memory used by each virtual machine.
- Storage Usage by Device: displays the percent of storage used. Storage is defined for each hypervisor.
- IP Usage in the Cloud: displays the percent of IP addresses that are being used.

The default for each report is to show the average usage. To view the maximum CPU and to display by name use the maximum toggle button in the upper left portion of the widow.

Each of the reports can be customized for specific date ranges.

To view the cloud usage reports:

1. Click **Cloud** → **Reports** from the menu (Figure 6-77).

Figure 6-77 The reports menu option

2. Select the appropriate cloud report. In our example we select **Memory Usage by Hypervisor** (Figure 6-78).

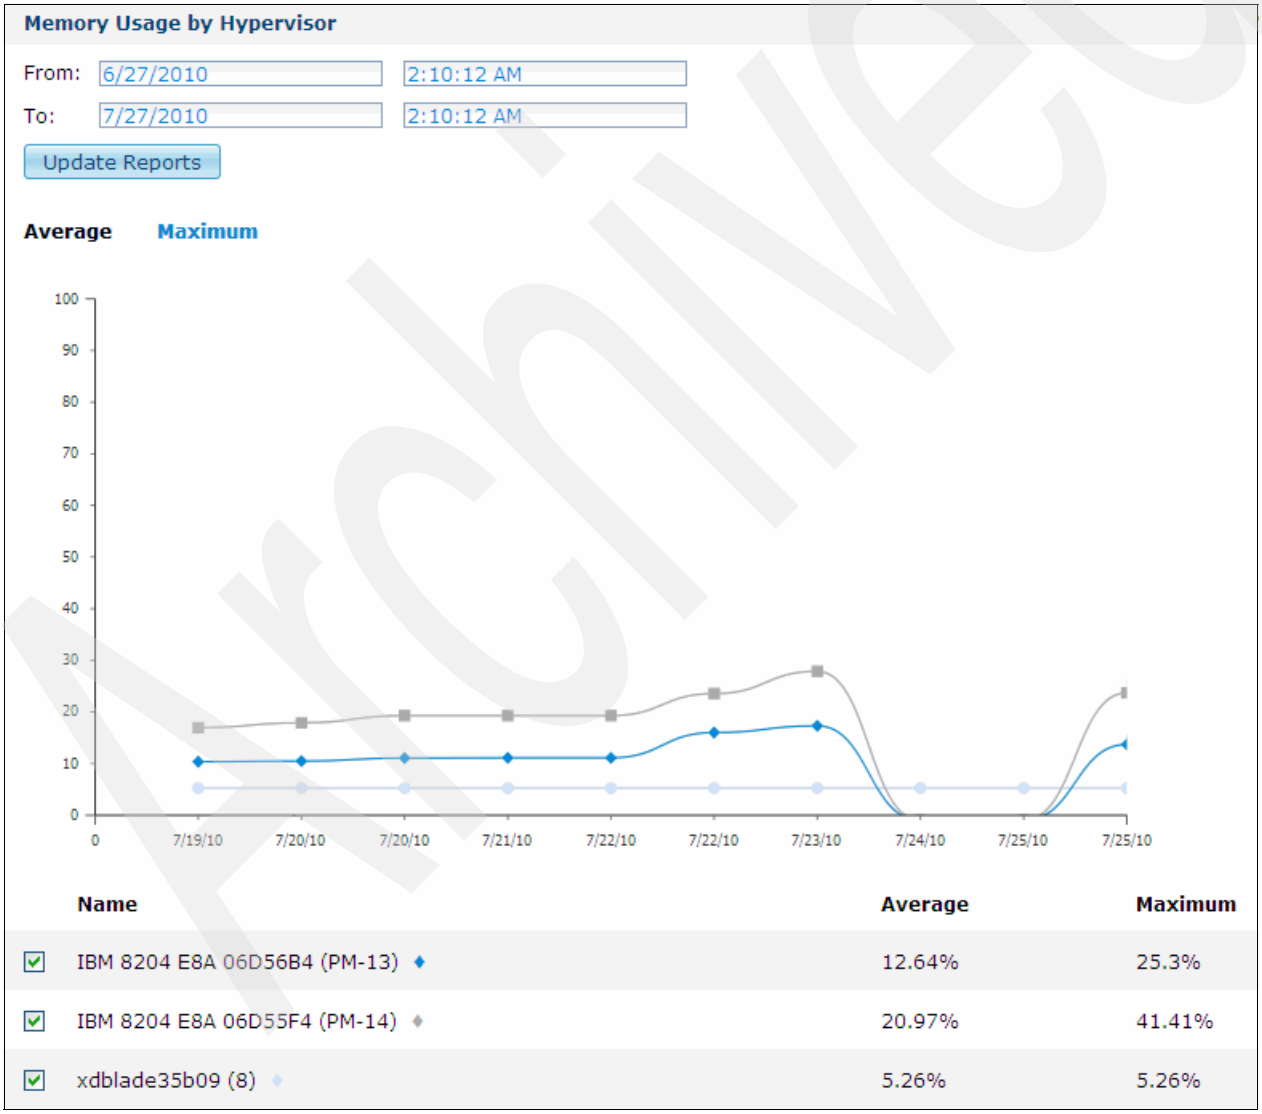

Figure 6-78 The memory usage by hypervisor graph

3. Optional: Modify the date range by selecting a start date and an end date.
4. Optional: Change the view by selecting the view box.

6.3 Installation of fix packs

As with all large applications and operating systems, mistakes or undesirable effects can occur when the application is being programmed. Fix packs are released to help rectify any mistakes by improving performance, fixing problems, and preventing possible security issues. Fix packs are generally rolled out on a routine basis unless a high-severity problem occurs and there is a need for an immediate fix.

WebSphere CloudBurst gives you the ability to update your deployed environments and your stored environments (that is, virtual images residing on the appliance) either using a method called extend and capture, which allows you to apply a permanent fix to an image, or by applying an emergency fix to running virtual systems.

There are several ways to apply fixes:

- Extend the virtual image, apply a fix pack, and re-capture the virtual image.

  This is the suggested method, as all subsequent deployments of the extended image will have the fix pack installed.
- Apply a fix pack to the running virtual system.
- Apply a fix pack to the virtual machine using native operating system tools.

  This is not suggested, as you loose the tracking and snapshot capabilities of WebSphere CloudBurst.

  This does not use the WebSphere CloudBurst Appliance.

  This can be done as an extend and capture or applied directly to each running virtual machine.

As well as having several methods of applying fixes, there are also a couple of types of fixes that can be installed:

- Emergency fixes: short-term fixes to fix urgent issues
- Service packs: upgrades the product version levels

The following section describes how to apply an emergency fix using the extend and capture method and applying an emergency fix to a running virtual system. Applying a fix using the native tools is out of scope of this document, as we focus on using the WebSphere CloudBurst Appliance. To complete the section we describe how to import an image and how it can be used to apply a service pack,

Note: Emergency fixes are WebSphere Application Server and WebSphere CloudBurst Appliance specific. AIX fixes cannot be applied to a deployed virtual system using the WebSphere CloudBurst Appliance fix tool. It is possible to package your own emergency fix. For more information refer to:

`http://publib.boulder.ibm.com/infocenter/wscloudb/v2r0/topic/com.ibm.websphere.CloudBurst.doc/pc/pct_package_ef.html`

6.3.1 Creating an emergency fix

The first step in applying any emergency fix is to create the emergency fix in the catalog. To create the emergency fix:

1. Download the emergency fix from fix central:

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

   In our scenario, we downloaded the fix for APAR PM12973. (There is no reason that we chose this particular fix other than we were looking for an example.) The file we saved to our local system was `7.0.0.11-WS-WAS-IFPM12973.pak`.

2. Click **Catalog** → **Emergency Fixes** on the menu.
3. Click the plus icon (**+**) to create a new emergency fix.
4. Enter the emergency fix name and a description for the emergency fix. In our scenario we called our emergency fix IFPM12793. Click **OK** (Figure 6-79).

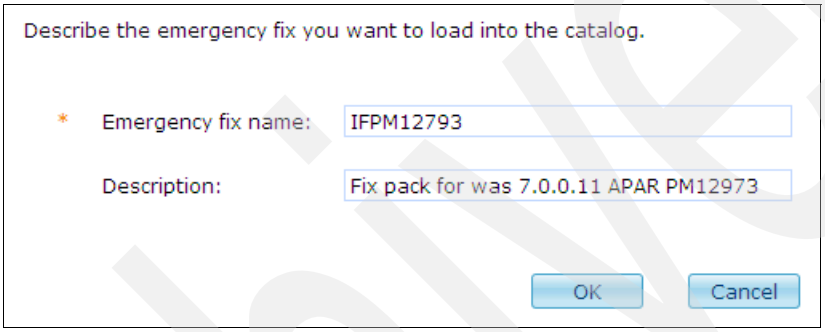

Figure 6-79 Create a fix pack

5. Click **Browse to** select the emergency fix file that was downloaded from fix central. Emergency fixes can only be .zip, .pak, and .tgz files. In our scenario we selected the file that we downloaded earlier, `7.0.0.11-WS-WAS-IFPM12973.pak`. Click **Upload** (Figure 6-80).

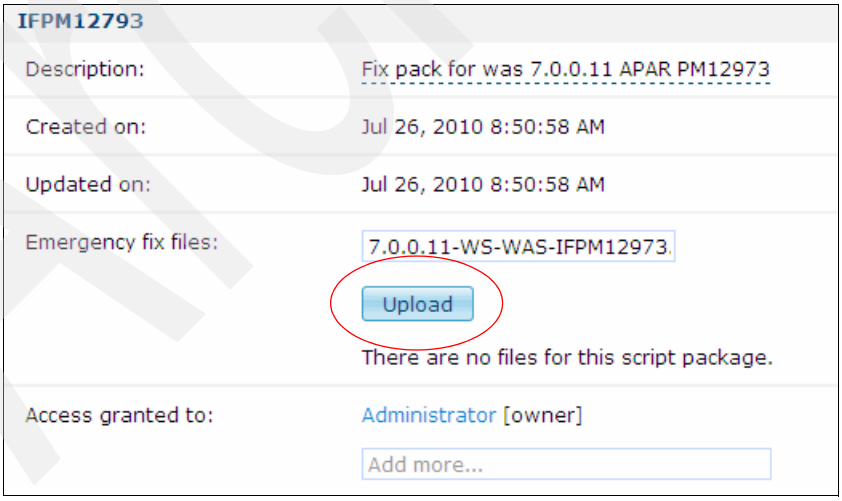

Figure 6-80 Upload the fix pack

6. The new fix has been added to the catalog. It is possible to further customize the properties of the fix. Click the fix name in the catalog to further customize the emergency fix.

7. Click **Add more** and select other users for whom to grant access to the emergency fix.
8. Select the severity of your emergency fix from the drop-down menu. This is a good way to help classify the serverity of the fix. In our scenario we left the fix at the default setting (Figure 6-81).

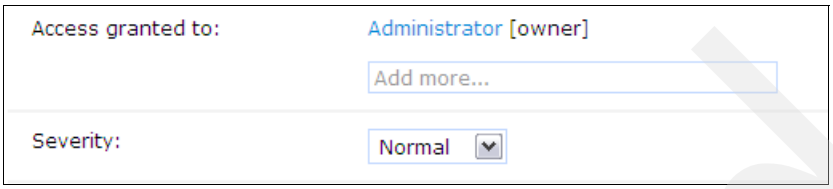

Figure 6-81 Change severity of the fix

9. Select the virtual images that the virtual image is applicable to by clicking **Add more** (Figure 6-82). In our scenario we chose the WebSphere Application Server V7.0.0.11 (PowerVM) image. The fix can only be applied to the images selected. By default it makes the change applicable to all subsequent virtual images that are cloned or extended from the selected virtual image.

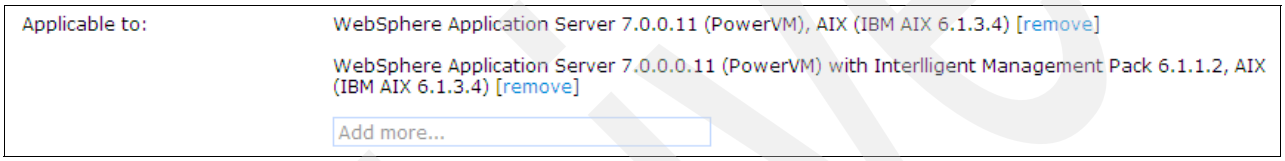

Figure 6-82 Select virtual images to apply emergency fix

10. The emergency fix can be deployed on all virtual systems created with the virtual image selected in the previous step.

6.3.2 Extending and capturing

Applying a fix pack by extending a virtual image and recapturing it is the suggested method for applying emergency fixes because it means that all subsequent deployments of the virtual image automatically have the fix pack installed.

When an extension of a virtual image is performed, the WebSphere CloudBurst Appliance:

- Creates a new virtual image in the catalog
- Deploys the raw virtual image, using the default pattern, to the cloud

This allows customization to be made to the virtual image without modifying the original image.

To extend an existing virtual image:

1. From the menu bar select **Catalog** → **Virtual Images**.
2. Click the virtual image that you want to extend from the catalog. In our scenario we update the WebSphere Application Server V7.0.0.11 (PowerVM) image.
3. Click the Clone and extend icon in the right corner of the image properties (Figure 6-83).

Figure 6-83 Clone and extend icon

4. Enter the values for each of the relevant fields:
   - General Information:
     - Name: name of the new virtual image. We named our image WebSphere Application Server V7.0.0.11 IFPM12973 (PowerVM).
     - Description: (optional) a description of the new virtual image.
     - Version: This must be a valid version of WebSphere Application Server. In our scenario we name the Version 7.0.0.11.
   - Deployment configuration
     - In cloud group: used as the cloud group to deploy the virtual image
     - Password: the password for the virtuser
   - Hardware configuration
     - Network interfaces: The number of network interfaces required by the new image. We use the default of 1.
     - image1.mksysb: the size on the new volume. By default an extra 15 GB is added. This space is used to allow space for a new mksysb to be generated. We used the default of 26 GB.
5. Click **OK** (Figure 6-84). This starts the process of creating the new image in the catalog and deploying the raw virtual image to the cloud for modification (extension). This process takes around 40 minutes but is dependent on your network speed.

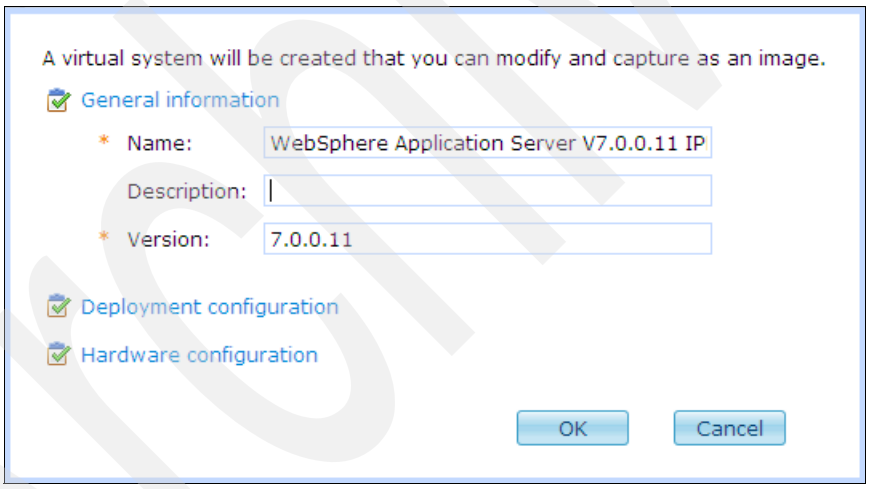

Figure 6-84 Creating new image

6. On the menu bar click **Virtual Systems**. A new virtual system has been created, and is based on the naming and version number generated from the previous step. The virtual image is ready to use and modify when a icon next the name is green. In our scenario our image is called `WebSphere Application Server V7.0.0.11 IFPM12793 (PowerVM) 7.0.0.11`.

6.3.3 Applying the fix to a deployed virtual image

To apply an emergency fix to the deployed virtual image:

1. On the menu bar click **Virtual Systems**.
2. Click the virtual system to which to apply the fix. In our scenario we select WebSphere Application Server V7.0.0.11 IFPM12793 (PowerVM) 7.0.0.11, which is the virtual system directly related to the virtual image that we previously extended.
3. Click **Apply service** (Figure 6-85).

Figure 6-85 Apply service pack

4. Optional: Schedule when you want to apply the fix. The default schedule is to apply the fix immediately. In our scenario we use the default to schedule immediately (Figure 6-86).

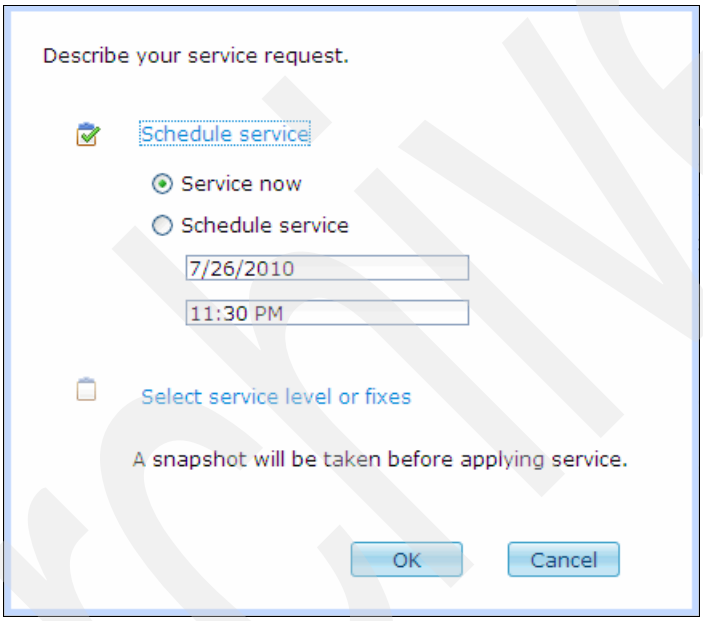

Figure 6-86 Schedule when to apply the fix

5. Select the **Apply emergency fixes** radio button because we are installing an emergency fix.

6. Select the emergency fix from the catalog and click **OK**. You can select more than one emergency fix if more than one is applicable. We selected the emergency fix created earlier called IFPM12973 (Figure 6-87).

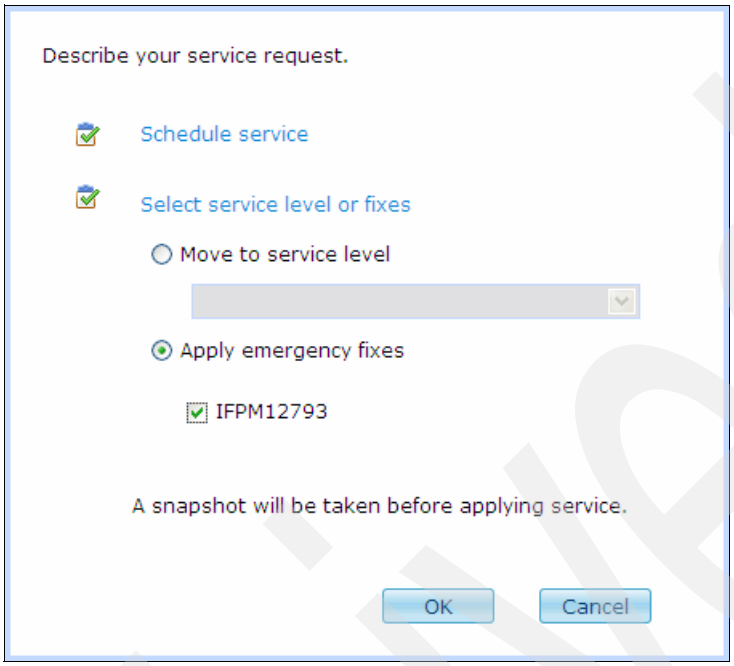

Figure 6-87 Select the correct fix packs

7. The emergency fix runs at the specified schedule. In our scenario it occurred immediately. Upon successful completion of emergency fix installation a successful message displays in the virtual systems history, and you receive an email about the successful installation (Figure 6-88).

```
Service applied on the virtual system

To login to the appliance, please visit https://9.42.89.38.
```

Figure 6-88 Successful email message

6.3.4 Applying a fix pack to a running virtual system

There are certain circumstances in which you might want to apply a fix to a running virtual system, such as when a deployed virtual system is set to run indefinitely and there is a critical fix that must be applied immediately.

The process is to deploy a fix to a running virtual system:

1. Create the fix in the catalog.
2. Apply the fix to the running virtual system.

Section 6.3.1, “Creating an emergency fix” on page 141, describes how to create the fix.

To apply the fix to a running virtual system:

1. On the menu bar click **Virtual Systems**.
2. Click the virtual system to apply the fix to. In our scenario we select `itso-was-standalone`, which is the virtual system directly related to the virtual image that we previously extended.
3. Click **Apply service** (Figure 6-89).

Figure 6-89 Apply service pack

4. Optional: Schedule when you want to apply the fix. The default schedule is to apply the fix immediately. In our scenario we use the default to schedule immediately.
5. Select the **Apply emergency fixes** radio button because we are installing an emergency fix.
6. Select the emergency fix from the catalog and click **OK**. We select the emergency fix created earlier called SR12973 (Figure 6-90).

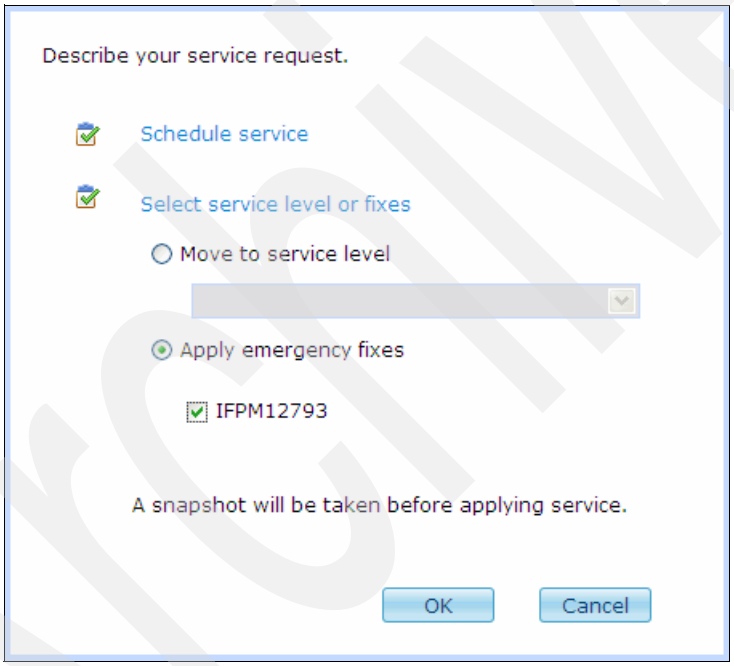

Figure 6-90 Select the fix pack

The emergency fix runs at the specified schedule. In our scenario it occurred immediately. Upon successful completion of emergency fix installation, a successful message displays in the virtual systems history and you receive an email about the successful installation.

6.3.5 Capture the updated system

After successful installation of the emergency fix, the image can be captured and placed in the catalog for deployments. To capture the updated image:

1. On the menu click **Catalog** → **Virtual Images**.
2. Select the virtual image that is in extension mode. In our scenario we select WebSphere Application Server V7.0.0.11 IFPM12973 (PowerVM).

3. In the right corner of the properties window click **Capture** (Figure 6-91). This process starts a mksysb, which creates a new virtual image that will be transferred to the WebSphere CloudBurst Appliance.

Figure 6-91 Capture virtual system

4. Optional: After capturing the modified image, permanently lock the virtual image to make it read-only to prevent further modification. To lock the virtual image, click the lock icon in the right corner of the virtual image properties window (Figure 6-92). If further modification is required, another extend and capture of this image can be performed.

Figure 6-92 Make virtual image read-only

To complete the entire process and ensure that users are using the new virtual image with the applied fix, clone all patterns that are using the previous virtual image (which does not contain the fix) and choose the deploy on the new virtual image.

If this process is not performed all future deployments continue on the unpatched virtual image.

6.3.6 Applying service packs

A service pack is a new virtual image that contains a more recent level of the software, and often contains multiple enhancements over previous versions. When the new virtual image is imported into the catalog, it automatically creates a service pack available for related virtual systems.

The steps to install a service pack are:

1. Install the service pack image. The service pack will be a new virtual image. See 6.4.1, “Importing virtual images” on page 149, for the steps for importing a virtual image.
2. On the menu bar click **Virtual Systems**.
3. Click the virtual system to apply the fix to. In our scenario we select our extended virtual image WebSphere Application Server 6.1.0.29 Update (PowerVM), which is the virtual system directly related to the virtual image that we previously extended.
4. Click **Apply service** (Figure 6-93).

Figure 6-93 Apply the service pack

5. Optional: Schedule when you want to apply the fix. The default schedule is to apply the fix immediately. In our scenario we use the default to schedule the fix immediately.
6. Select the **Move to service** level radio button because we are updating to a new server level.

7. Select the service level from the from list of service levels and click **OK**. In our scenario we selected WebSphere Application Server 6.1.0.31 (Figure 6-94).

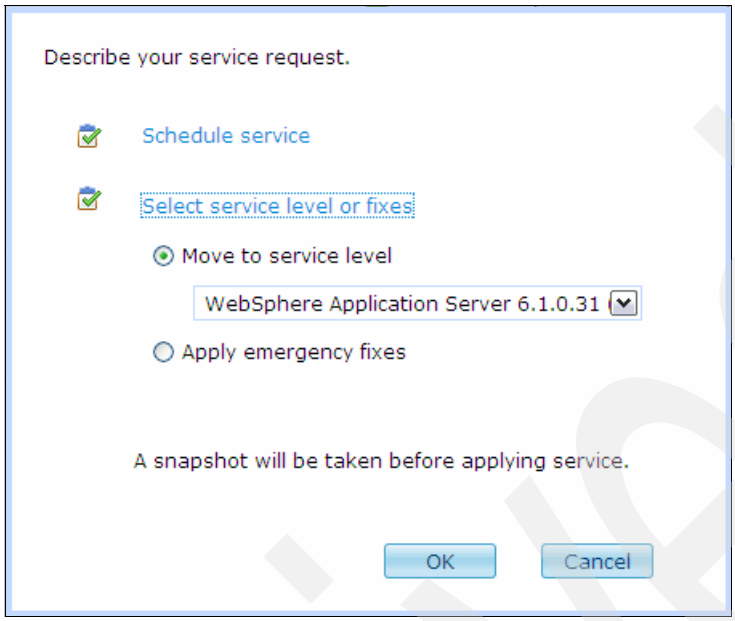

Figure 6-94 Select the new service level

8. The service level starts updating. When completed, you receive an email stating that the update is complete. The virtual system also stops when completed and requires a manual restart. To start the virtual system, click **Start** on the virtual system panel (Figure 6-95).

Figure 6-95 Start the virtual system

9. You can confirm that the virtual system has been updated by checking the Service History section on the virtual systems properties panel (Figure 6-96).

Figure 6-96 The service history of the virtual system

6.4 Import and export of virtual images

IBM periodically releases new virtual images for the WebSphere CloudBurst Appliance that contain new functionality and fixes, and are updated versions of the software products. To

place any new or customized virtual images on the WebSphere CloudBurst Appliance, we use the import function.

6.4.1 Importing virtual images

The web console can import Open Virtual Appliance (.ova) files stored on either an SSH or an http server. An Open Virtual Appliance file contains the files used to describe the virtual appliance (for PowerVM a mksysb).

In our scenario we chose to import the Open Virtual Appliance file for an earlier version of WebSphere Application Server Hypervisor Edition. In our scenario we used WebSphere Application Server Hypervisor Editon V6.1.0.29.

> **Note:** The virtual image for WebSphere Intelligent Management Pack can be downloaded from Passport Advantage®. Download directions can be found at:
>
> http://www.ibm.com/support/docview.wss?uid=27016051

The following steps are how we imported the virtual image:

1. Download the ova image from Passport Advantage. We downloaded `CZC8HML.ova`.
1. We copied the image to our local machine and then copied it to an SSH enabled server. We copied the image `CZC8HML.ova` to the `/tmp` directory of the NIM server itso-nim.
2. Log in to the WebSphere CloudBurst Appliance as a user who has permission to add catalog content. In our scenario we used the cbadmin user.
3. Select **Catalog** → **Virtual Images** on the menu bar.
4. Click the green plus icon (**+**) to create a new virtual image.
5. Enter the OVA file location. The file location can either be an http URL or the image can be transferred using Secure Copy Protocol (SCP) using the fomat *<hostname>:/path*. In our scenario we use SCP to copy the image from the NIM server itso-nim. Our file location looked like:

   ```
   itso-nim.rainmaker.raleigh.ibm.com:/tmp/CZC8HML.ova
   ```

6. Enter the user name and password for the user who has read access to the ova files stored on the SSH server and click **OK**. In our scneario we use the root user of the NIM server itso-nim (Figure 6-97).

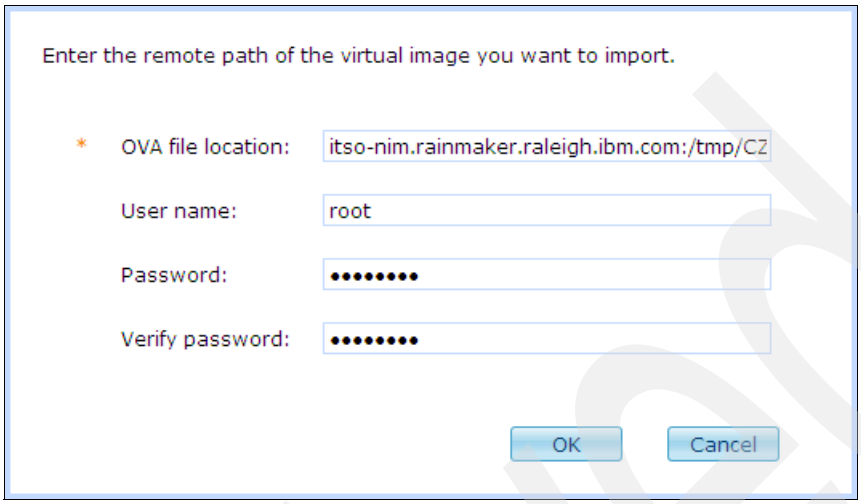

Figure 6-97 The import information

7. The import starts and is placed in the task queue. The name of the new imported virtual image will be something incoherent but will start with `New virtual image <numbers>`. Upon completion and a successful import, the name of the virtual image changes to the correct name, which is stored in one of the files inside the ova file. In our scenario the new server name was WebSphere Application Server 6.1.0.29 (PowerVM) (Figure 6-98).

Figure 6-98 The completion of a successful import

6.4.2 Exporting

Exporting a virtual image for use on other WebSphere CloudBurst Appliances can be particularly useful when you have more than one WebSphere CloudBurst Appliances in your environment. You can share customized versions of virtual images or you can back up and store individual virtual images separate from the backup of the WebSphere CloudBurst Appliance.

Virtual images exported from the catalog are stored in the OVA file format and can only be transferred to an external server that is SSH enabled. Ensure that you have the ability to create files larger than 2 GB, which can be a restriction on certain operating system platforms. The exported image can 10 GB or larger depending on the customization. Coordinate with your system administrator before performing this function.

To export a virtual image:

1. Log in to the WebSphere CloudBurst Appliance web console as a user who has create new catalog permission. In our scenario we used the cbadmin administrative user.
2. On the menu bar click **Catalog** → **Virtual Images**.
3. From the image catalog click the virtual image that you want to export. In our scenario we chose the image WebSphere Application Server V7.0.0.11 IPFM12973 (PowerVM), which has a fix pack applied to it.

4. Click the Export icon on the top right side of the properties window (Figure 6-99).

Figure 6-99 The export icon

5. A new window opens prompting for details of the destination SSH server. Complete each field and click **OK** (Figure 6-100).
   - Remote host: Specify the host name or IP address of the destination server. In our scenario we chose our NIM server itso-nim.rainmaker.raleigh.ibm.com.
   - Remote path: Specify the path to where you want to place the image on the destination. In our scenario we chose the `/tmp` directory.
   - Username: Specify the user name of a user on the destination server who will be used to transfer the image. In our scenario we used root.
   - Password: The password of the user on the destination server.

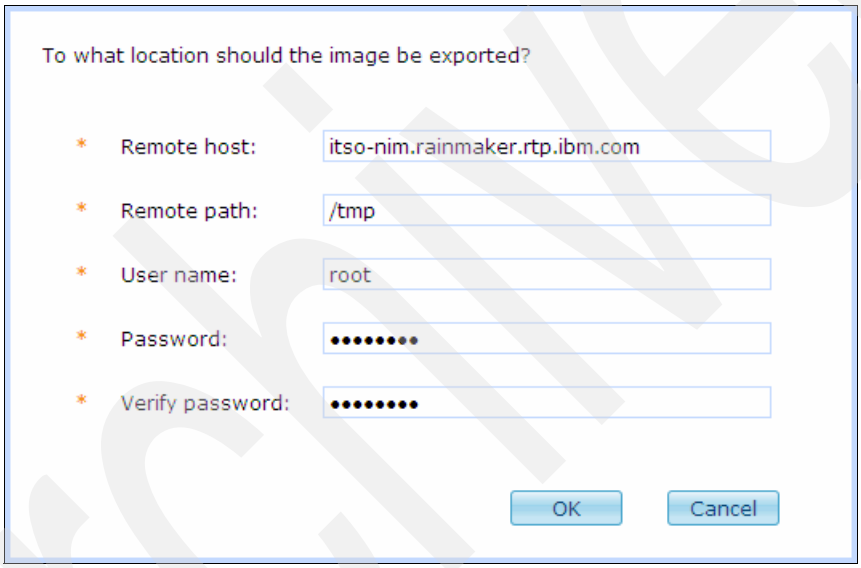

Figure 6-100 Export image options

Your image starts transferring, and a task is created in the task queue.

6. When the task is completed you receive an email notifying you that the task is completed (Figure 6-101). The image, which is in an ova file format, can be imported to another WebSphere CloudBurst Appliance.

```
WebSphere CloudBurst has successfully exported WebSphere Application Server 7.0.0.11 IPFM12973
(PowerVM) to itso-nim.rainmaker.raleigh.ibm.com:/tmp. The OVA file can be imported to another
appliance or used by any tool that implements the OVF standard.

To login to the appliance, please visit https://9.42.89.38.
```

Figure 6-101 Email advising the export was successful

6.5 Installation of new virtual image

Adding extra virtual images to WebSphere CloudBurst Appliance extends the default capabilities of WebSphere CloudBurst Appliance from just being a WebSphere Application Server deployment platform to becoming a general cloud management tool. In this scenario we look at extending the capabilities by installing a PowerVM virtual image of WebSphere Process Server Hypervisor Edition and WebSphere Intelligent Management Pack.

In our scenario we use two methods of installing a virtual image. The WebSphere Process Server package comes with an image loader Java utility. The WebSphere Application Server Hypervisor Edition is loaded using the native WebSphere CloudBurst Appliance tools.

Note: The virtual image for WebSphere Process Server Hypervisor edition can be downloaded from Passport Advantage at:

http://www.ibm.com/support/docview.wss?uid=24026231

6.5.1 WebSphere Process Server

WebSphere Process Server is a standards-based business process integration server underpinned by the robust Java EE infrastructure of WebSphere Application Server, on which it is based. It allows deployment and execution of business processes that orchestrate services within an service-oriented architecture (SOA) environment. It also includes the complete functionality of WebSphere Enterprise Service Bus, enabling the implementation of SOA-based integration solutions.

WebSphere Process Server serves as the execution engine of business modules designed in WebSphere Business Modeler or WebSphere Integration Developer.

New WebSphere Process Server V7 features include enhanced support for open standards such as Java EE 5 and EJB 3.0, simplified system installation, easier cluster configuration, richer capabilities for process administrators, and single-step migration. WebSphere Process Server now allows you to change process instances in-flight to react dynamically for changing business conditions.

6.5.2 Import of image and patterns

The method for installing WebSphere Process Server Hypervisor Edition can be found in the quick start guide, which can be downloaded with the product from Passport Advantage. One purpose for using the image loader Java utility is that it also installs predefined patterns for WebSphere Process Server.

We used the following method to install the WebSphere Process Server Hypervisor Edition virtual image using the quick start instructions:

1. Download part numbers CZFG0ML and CZFG2ML from the Passport Advantage website:

   http://www-01.ibm.com/software/howtobuy/passportadvantage/pao_customers.htm

2. Decompress `wpshv.quickstart.v7.zip` and read through the quick start guide.
3. Extract the `wpshv.v7.aix.ppc64.tgz` image. In our scenario we opened the archive on a Microsoft® Windows® XP desktop using an open source application environment called 7zip (http://www.7-zip.org).

4. The extraction provides the load file. Run `loader.bat` from the directory of the extracted package. This opens the graphical loader user interface (Figure 6-102).

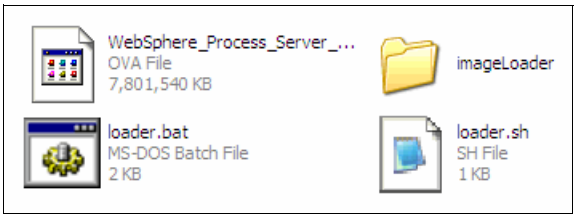

Figure 6-102 Run the loader script

5. Enter the host name or IP address of the WebSphere CloudBurst Appliance where you want to install the image. Enter the user name and password to access the appliance. Click the **Load a Virtual Image** link. In our scenario we entered the IP address of the WebSphere CloudBurst Appliance and used the cbadmin user for the connection (Figure 6-103).

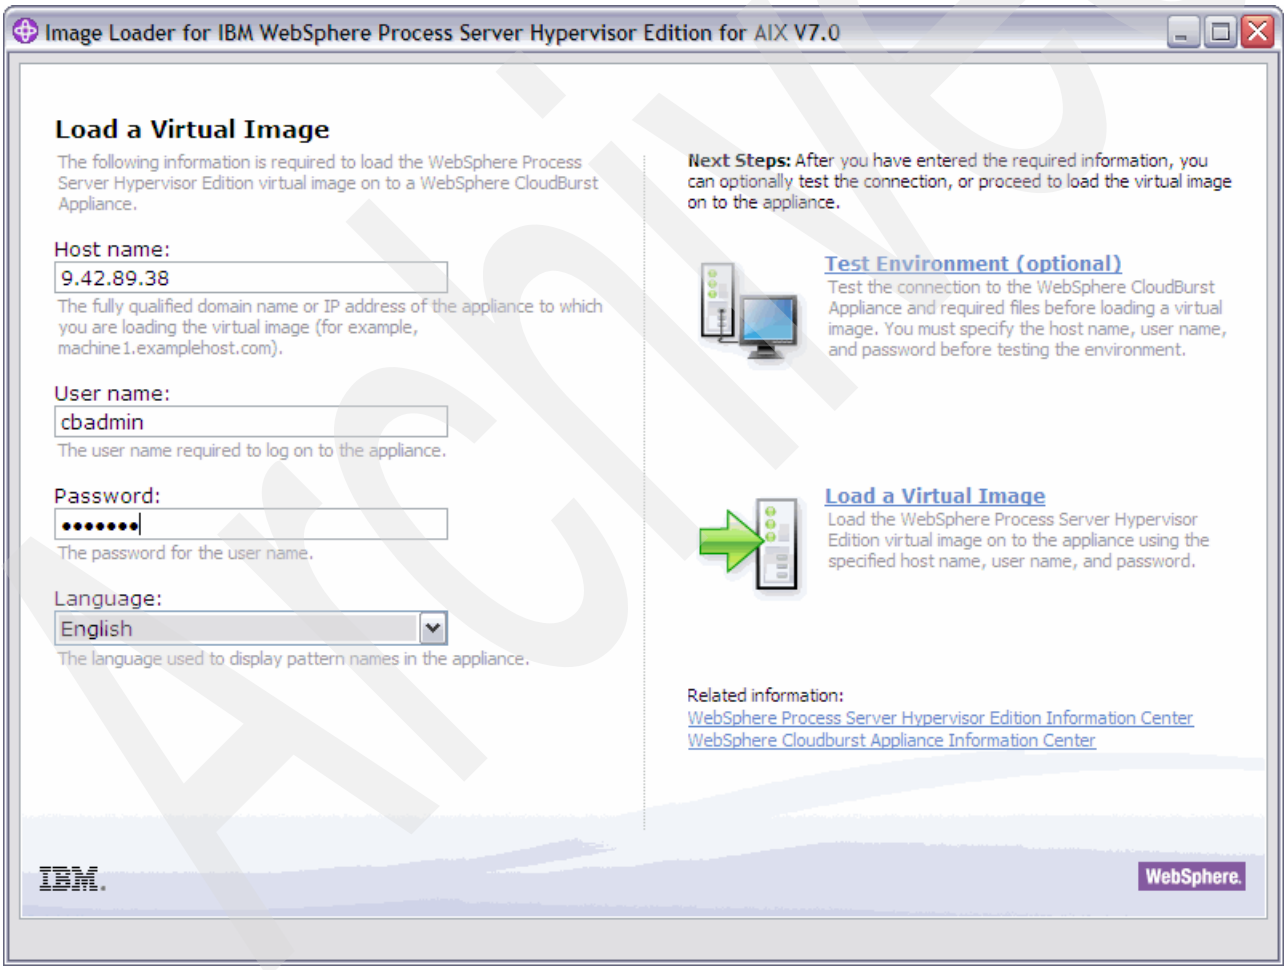

Figure 6-103 Enter the host name, user name, and password of the WebSphere CloudBurst Appliance

You can watch the progress of the WebSphere Process Server Hypervisor Edition virtual image as it is uploaded to the WebSphere CloudBurst Appliance. This step can take a long time and depends on your network speed (Figure 6-104). In our test scenario the installation of the virtual image upload took around 2.5 hours.

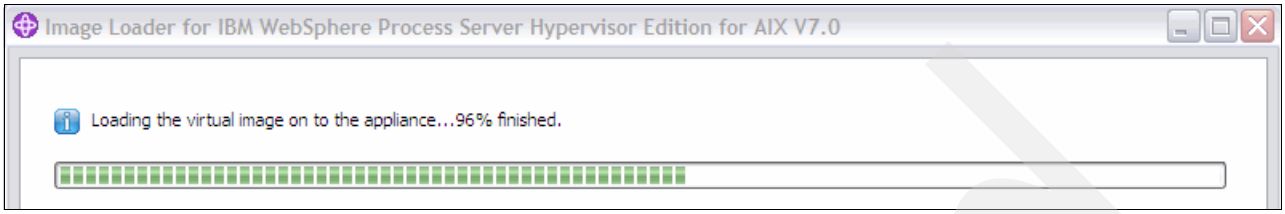

Figure 6-104 The load virtual image progress

6. When the image loader is complete, log in to the WebSphere CloudBurst Appliance web console and select **Catalog** → **Virtual Images** from the menu.
7. Click **WebSphere Process Server 7.0.0.2 for AIX** from the list of virtual images (Figure 6-105).

Figure 6-105 WebSphere Process Server Virtual Image

8. Click the **accept** hyperlink next to the accept category from the virtual image properties window (Figure 6-106).

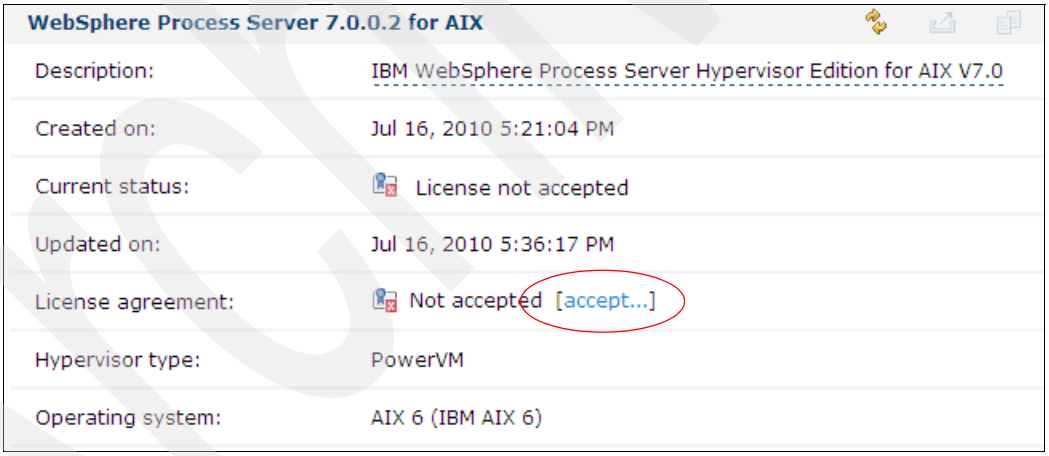

Figure 6-106 Accept the licenses

9. Accept all of the license agreements. Patterns using the WebSphere Process Server virtual image cannot be deployed until all of the licence agreements have been accepted (Figure 6-107).

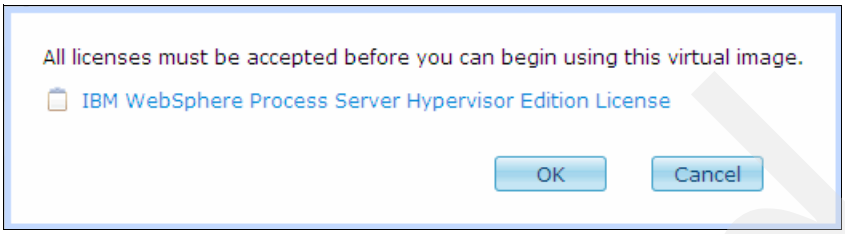

Figure 6-107 Accept the WebSphere Process Server license

10. Click the **Patterns** menu item. Confirm that the predefined WebSphere Process Server patterns have been installed with the product (Figure 6-108). The pattern names are:
    - Scalable Environment for Process Server v7.0.0.2 on AIX
    - Simulated Scalable Environment for Process Server v7.0.0.2 on AIX
    - Single Server for Process Server v7.0.0.2 on AIX

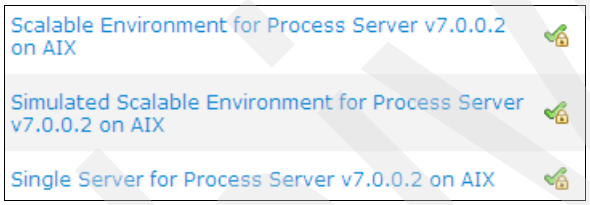

Figure 6-108 Predefined WebSphere Process Server patterns

Patterns using the WebSphere Process Server can now be deployed.

7

Working with patterns

This chapter provides an overview of WebSphere CloudBurst patterns and IBM Hypervisor Edition virtual images.

In this chapter, the following topics are discussed.

7.1 Definition of pattern and components

This section defines patterns and their components.

Patterns are a combination of virtual image parts and scripts packages. A pattern is a complete representation of an application environment. It represents both the node types in the environment and the configuration of the environment. WebSphere CloudBurst provides the capability to deploy these patterns and turn them into a running application environment of patterns (Figure 7-1).

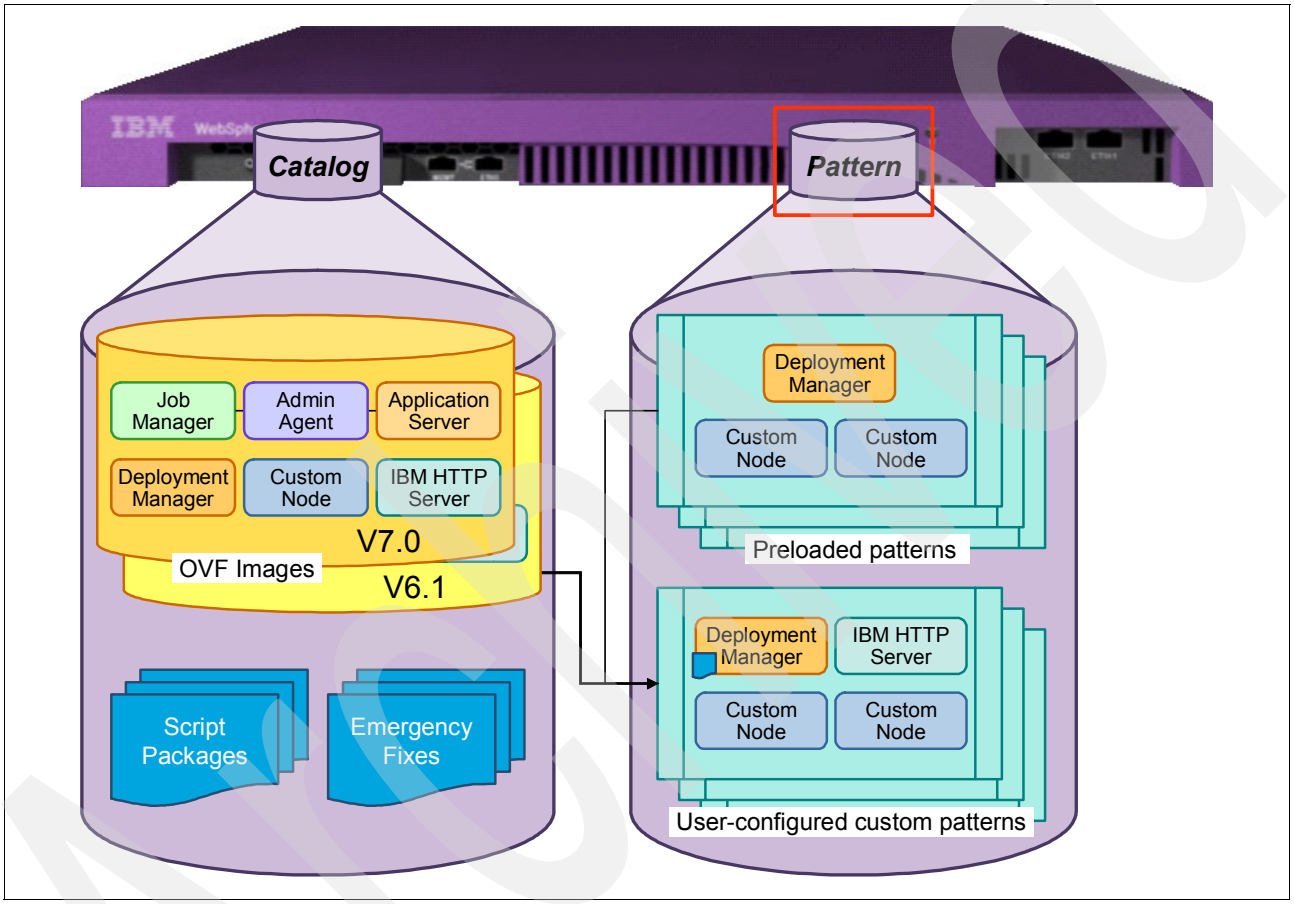

Figure 7-1 Patterns

Pattern components

The pattern components are:

- Virtual image parts

  Virtual image parts are used to create a pattern. Virtual image parts represent a type of node in your application environment. It could be a profile, IBM HTTP Server instance, or database server.

- Script packages

  A script package is a file, such as a script or archive, that represents middleware customization.

  Section 7.3, “Introduction of script packages” on page 169, provides a detailed description of scripts packages.

7.2 Overview IBM Hypervisor Edition virtual images

This section describes IBM Hypervisor Edition virtual images. IBM Hypervisor Edition virtual image is a system platform that contains the operating system. IBM Hypervisor Edition virtual image runs on a hypervisor, which is the virtualization technology that manages the state of the virtual machines. IBM Hypervisor Edition images provide a complete middleware software stack in a single virtual image. The images contain an operating system, selected IBM middleware software, and an activation engine that configures the software at startup.

WebSphere CloudBurst 2.0 enables you to easily manage the IBM middleware products in a private cloud environment across a pool of virtualized servers, providing significant value to your data center. The WebSphere CloudBurst Appliance comes pre-loaded with the IBM Hypervisor Edition virtual images of WebSphere Application Server, WebSphere Application Server Feature Pack, and WebSphere Application Server Intelligent Management Pack.

Catalog

IBM Hypervisor Edition virtual images reside in the WebSphere CloudBurst Appliance catalog.

7.2.1 IBM WebSphere Application Server Hypervisor Edition for PowerVM

The IBM Hypervisor Edition contains an operating system, WebSphere Application Server binaries, IBM HTTP Server, and all the profiles supported for that release (Figure 7-2).

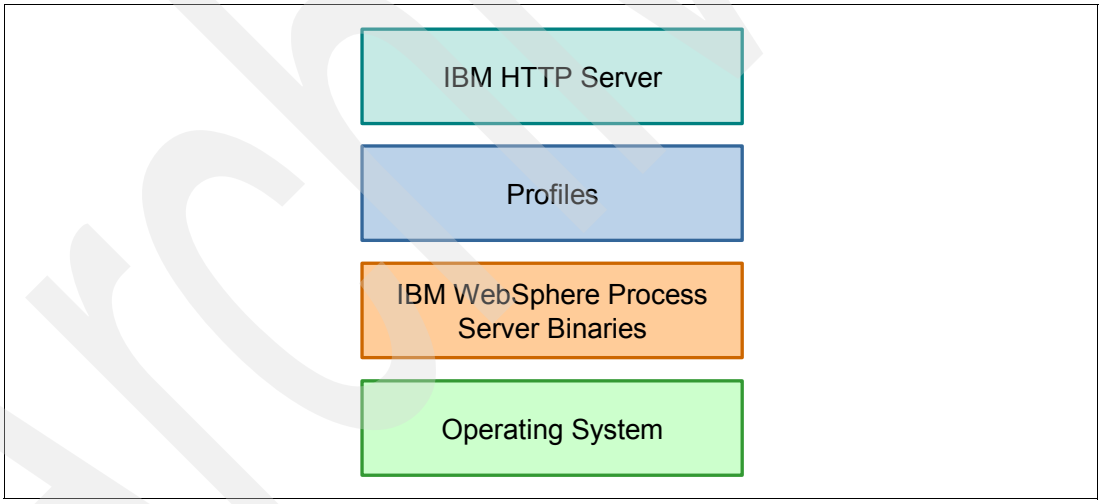

Figure 7-2 WebSphere Application Server Hypervisor Edition virtual image

There are three versions of the IBM WebSphere Application Server Hypervisor Edition virtual image, which we discuss in the following sections.

IBM WebSphere Application Server Hypervisor Edition 6.1 for PowerVM

This PowerVM image supports the AIX 6.1.3 operating system. This image includes parts that represent the following application server topology components (Figure 7-3):

- Stand-alone server
- IBM HTTP server
- Deployment Manager
- Custom Node

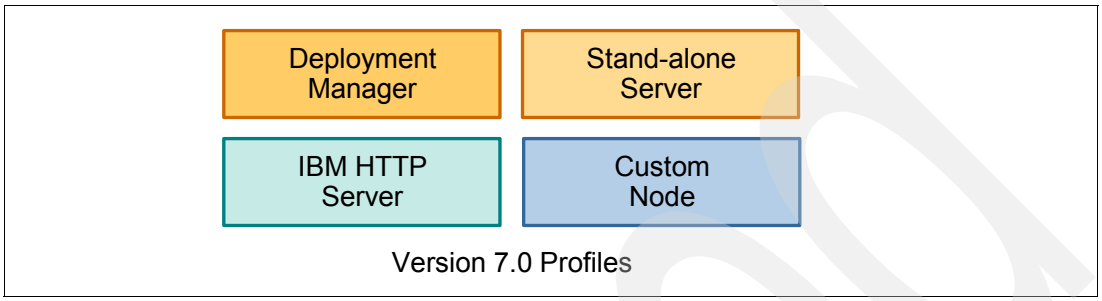

Figure 7-3 IBM WebSphere Application Server Hypervisor Edition 7.0 profiles

Generally, virtual image parts correspond to desired node or profile types such as deployment manager, custom node, and IBM HTTP server. A pattern is made up of a combination of these virtual image parts.

IBM WebSphere Application Server Hypervisor Edition 7.0 for PowerVM

This section describes the IBM WebSphere Application Server Hypervisor Edition 7.0 for PowerVM.

IBM WebSphere Application Server Hypdervisor Edition 7.0 for PowerVM support's AIX 6.1.3 operating system. Figure 7-4 shows parts that represent the application server topology components the image supports below. These parts, with the exception of the on demand router, are not particular to the WebSphere Intelligent Management version of the 7.0 image. The on demand router can only be used when the WebSphere Intelligent Management Pack is enabled.

- Administrative agent
- Custom nodes
- Deployment manager
- IBM HTTP server
- Job manager
- Stand-alone server
- On demand router

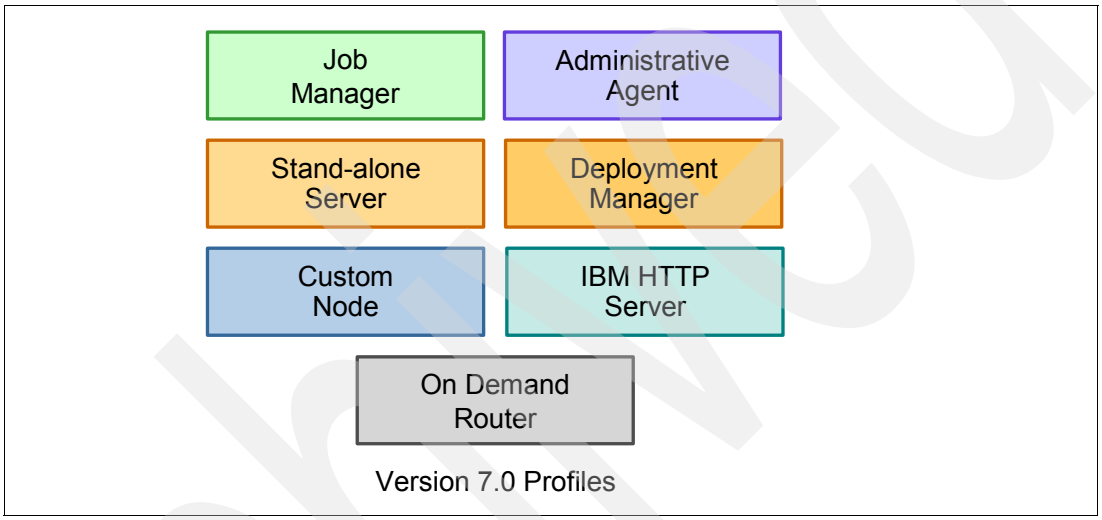

Figure 7-4 IBM WAS Hypervisor Edition 7.0 Intelligent Management Pack Profiles for PowerVM

7.2.2 WebSphere Intelligent Management Pack

This section describes the WebSphere Intelligent Management Pack and WebSphere Virtual Enterprise overview and key benefits.

WebSphere Intelligent Management Pack provides dynamic runtime capabilities similar to those present in WebSphere Virtual Enterprise.

> **Important:** The Intelligent Management Pack features are similar to the WebSphere Virtual Enterprise, but they are not the same.

WebSphere Virtual Enterprise

WebSphere Virtual Enterprise provides the capability of application virtualization. Using application infrastructure virtualization, you can separate applications from the physical infrastructure on which they are hosted. Workloads can be dynamically placed and migrated across a pool of application server resources, which allows the infrastructure to dynamically adapt and respond to business needs. Requests are prioritized and intelligently routed to respond to the most critical applications and users.

Key resources delivered by WebSphere Virtual Enterprise for PowerVM are:

- Logical partitions with fixed CPU resources.
- Logical partitions with fixed fractional allocation of CPU resources.
- Dynamic logical partitioning with whole allocation. This does not support shared allocation.

On demand router

The on demand router is an intelligent HTTP and Session Initiation Protocol (SIP) proxy server in WebSphere Virtual Enterprise. You can configure the on demand router to determine how it handles failure scenarios and how it tunes certain work requests. The on demand router is the point of entry into a WebSphere Virtual Enterprise environment and is a gateway through which HTTP requests and SIP messages flow to back-end application servers.

Key features of WebSphere Intelligent Management Pack for PowerVM

Key features are:

- Improved application performance and delivery response times to meet service level agreements
- Increased application availability and minimized administration costs
- Interruption-free maintenance upgrades
- Advanced options in WebSphere Intelligent Management Pack

Health management

Health management allows you to take a policy driver approach to monitoring your environment and take corrective action when certain predefined criteria are met.

Health management standard policies are:

- Monitor when excessive memory is being consumed.
- Monitor when a memory leak has been detected.
- Monitor when a server reaches a certain age and recycle the server automatically.

On demand router dependent policies apart from standard policies are:

- Monitor incoming requests and take corrective action if a certain predefined threshold is met.
- Monitor the number of time-out requests and take corrective action if the response times exceed a predefined threshold.

Dynamic clustering

Dynamic clustering is an application deployment target that can expand and contract depending on workload. Dynamic clusters operating at the application layer virtualization level means that dynamic clustering takes care of the resources inside the cell.

Key points of dynamic clustering are:

- Dynamic clusters grow and shrink depending on the workload demand.
- Dynamic clusters work closely with the on demand router to ensure the even distribution of workload amongst the cluster members.

Overload protection

Overload protection monitors the usage of memory and CPU. It regulates the rate at which the on demand router forwards traffic to the application server tier to prevent memory and

processor overload. These options are functions for which the on demand router is a key component, and thus depend on an enhanced pattern to be applied.

WebSphere Intelligent Management Pack enablement

The Intelligent Management Pack enablement feature is disabled by default. You need to enable the features of WebSphere Intelligent Management Pack. Follow the steps below to enable the features of Intelligent Management Pack capabilities:

1. Navigate to **Catalog** → **Virtual Images**.
2. Locate the WebSphere Application Server V7.0.0.11 virtual image.
3. Clone the WebSphere Application Server V7.0.0.11.
4. After the clone is completed, select **Enable** from the drop-down menu (Figure 7-5).

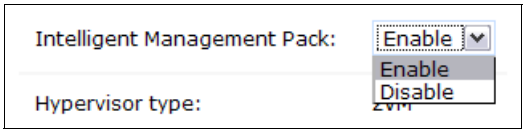

Figure 7-5 WebSphere Intelligent Management Pack Enablement

Intelligent Management Pack capabilities are enabled now on this virtual image.

We need to have and add the product ID for the virtual image. Figure 7-6 shows the product ID.

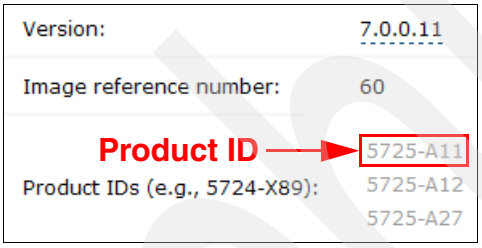

Figure 7-6 Product ID information

We need to configure the available licenses based on the product IDs:

1. Click **Catalog** → **Product Licenses** (7.2.3, “WebSphere Application Server Hypervisor Edition Feature Packs for PowerVM” on page 164).
2. IBM products are located in the same window. You will see the list of IBM products residing on the catalog that have been registered with a product ID.
3. Enablement of the Intelligent Management Pack features adds another product to this list labeled “IBM WebSphere Application Server Hypervisor Edition Intelligent Management Pack.”
4. Enter the number of process value units (PVUs) in the Licenses owned (PVU’s) field.

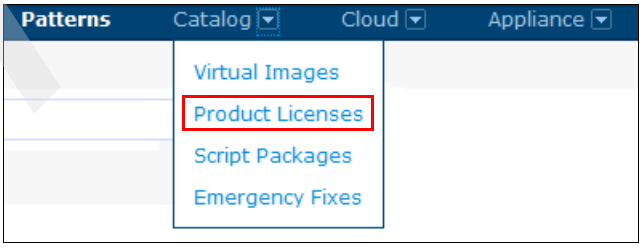

Figure 7-7 Product licenses

7.2.3 WebSphere Application Server Hypervisor Edition Feature Packs for PowerVM

IBM WebSphere Application Server Hypervisor Edition is a self-contained virtual machine image that contains a guest operating system and WebSphere Application Server. IBM WebSphere Application Server is the premier Java Platform, Enterprise Edition (Java EE) and web services-based application server. WebSphere Application Server is built on open standards and helps you deploy and manage applications ranging from simple websites to powerful on demand solutions.

The IBM WebSphere Application Server Hypervisor Edition 6.1 feature packs contain:

- EJB 3.0 Feature Pack
- Web Services Feature Pack
- Web 2.0

The IBM WebSphere Application Server Hypervisor Edition 7 feature packs (except EJB 3.0 and Web Services FPs) contain:

- Feature Pack for Service Component Architecture (SCA)
- Feature Pack for Web 2.0
- Feature Pack for Communication and Enabled Applications
- Feature Pack for XML

7.2.4 WebSphere Process Server Hypervisor Edition for PowerVM

IBM WebSphere Process Server Hypervisor Edition is a self-contained virtual machine image that contains a guest operating system and WebSphere Process Server. Use it to easily create, deploy, and manage WebSphere Process Server environments and applications using hypervisors. Figure 7-8 shows the WebSphere Process Server Hypervisor Edition.

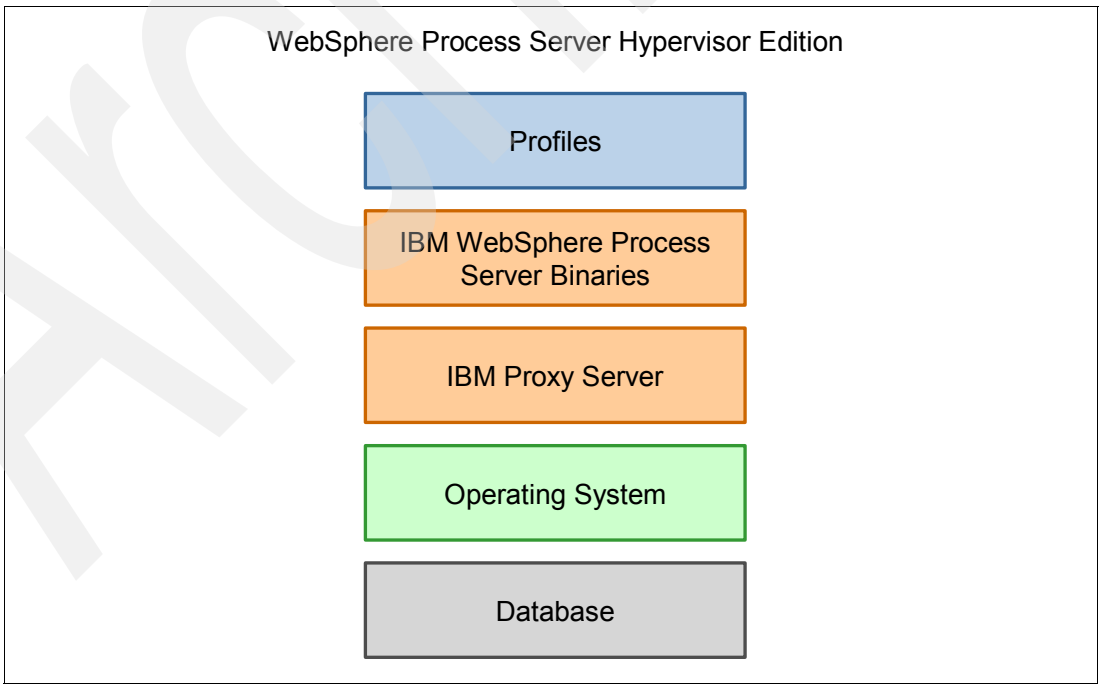

Figure 7-8 WebSphere Process Server Hypervisor Edition

WebSphere Process Server Hypervisor Edition does not come with the appliance. You can download the package from Passport Advantage. There are two images available to download:

- WebSphere Process Server Hypervisor Edition 6.2
- WebSphere Process Server Hypervisor Edition 7.0

The Websphere Hypervisor Edition package contains a virtual image that can be loaded in the WebSphere CloudBurst Appliance image catalog and the image loader utility used to load the virtual image. There are nine parts, four images, and three patterns for each image. The following are the types of virtual parts.

- Database

  This part contains the installation of DB2® database and can be configured for the deployment manager or for the cluster topology.

- Extendable nodes

  This part contains the installation of WebSphere Process Server.

- Stand-alone Server

  This part contains DB2 installed and WebSphere Process Server installation with a stand-alone server instance.

- Deployment Manager

  This part contains a configuration of application servers and nodes in a cell. In this part the deployment management profile is configured to use a DB2 database.

- Custom nodes

  This part provides the unconfigured node for use with a deployment manager or a part with control nodes.

- Proxy Server

  This part contains the proxy server instance used for load distribution.

- Basic function nodes

  This pattern part contains three application server clusters:

  - Messaging cluster
  - Support cluster
  - Application cluster

- Full function control node

  This pattern part contains:

  - Deployment manager
  - DB2 Server
  - Proxy server
  - Two custom nodes with three application server clusters
    - Messaging
    - Support
    - Application

- Basic function control node

  This is a full function control node without a DB2 server or proxy server.

WebSphere Process Server pattern

There are several preconfigured patterns for use with WebSphere CloudBurst. These patterns are ready to use and are added automatically to the appliance after we load the image using image the loader tool. The preconfigured patterns are:

- Single server

  This pattern is a single virtual machine that is composed of the operating system, a DB2 server, and a WebSphere Process server installed with a single stand-alone server instance defined. This pattern tests and verifies application functionality in an environment where instances are created and destroyed frequently (Figure 7-9).

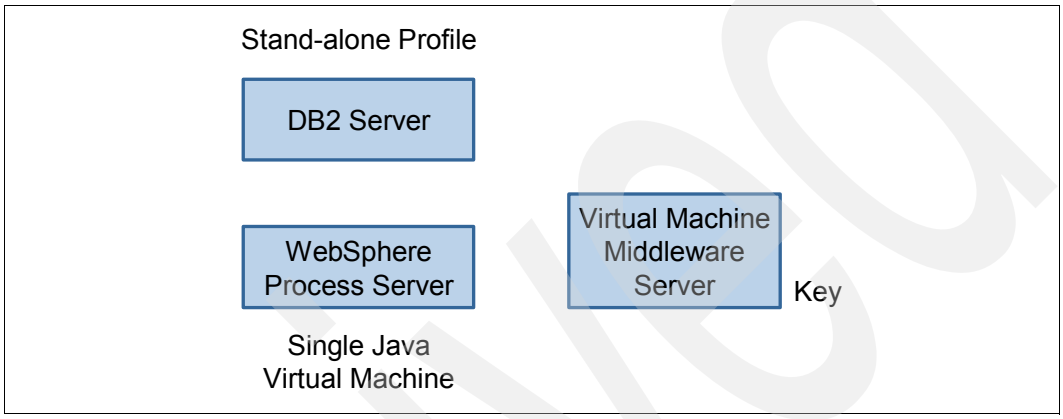

Figure 7-9 WebSphere Process Server Hypervisor Edition - Single server

- Scalable environment

  The scalable environment pattern represents a sample production WebSphere Process Server gold topology pattern. The pattern parts are broken into several virtual machines that are ideal for a production environment deployment manager when failover support is important. The DB2 server, the proxy server, and the deployment server are all deployed in their own virtual machines. Two more virtual machines are defined with three clusters spaning them. Each cluster has two servers for the application target, support, and messaging. Each virtual machine containing the clusters has a single node agent (Figure 7-10).

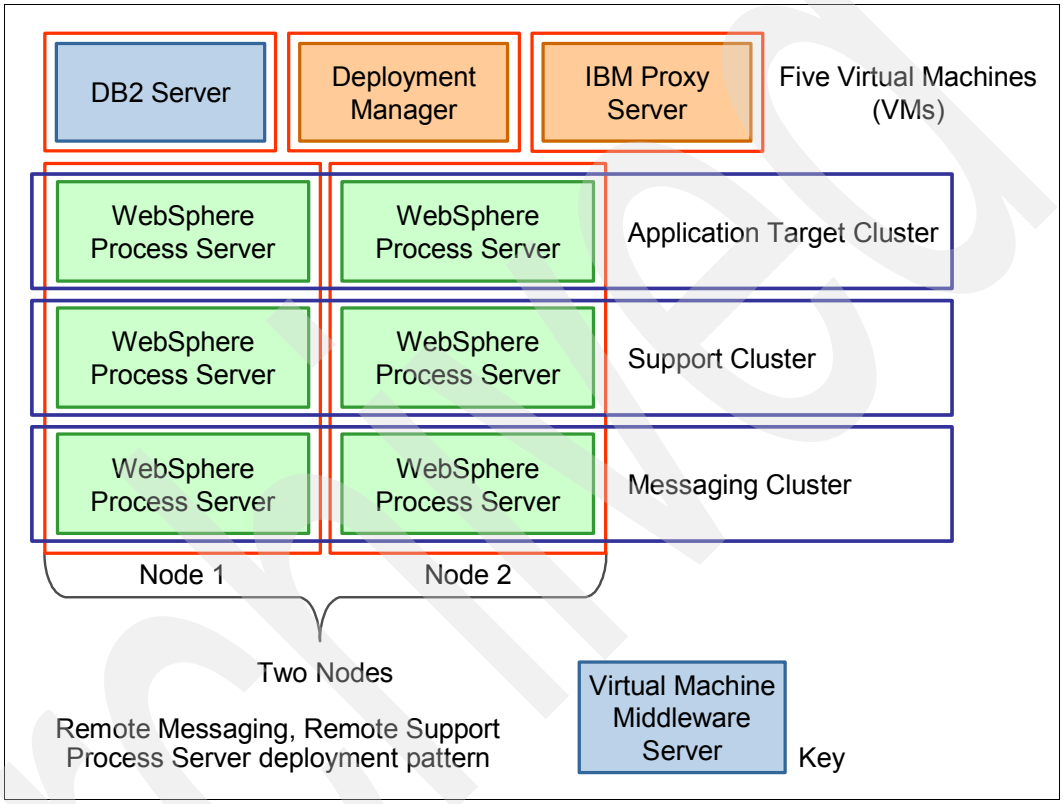

Figure 7-10 WebSphere Process Server Hypervisor Edition - Scalable environment

- Simulated scalable environment

  This is a pattern that contains the full function control node pattern part. This pattern part includes a DB2 server, a deployment manager, a proxy server, and three WebSphere Process Server clusters. There is a cluster defined for the application target, for support, and for messaging (Figure 7-11).

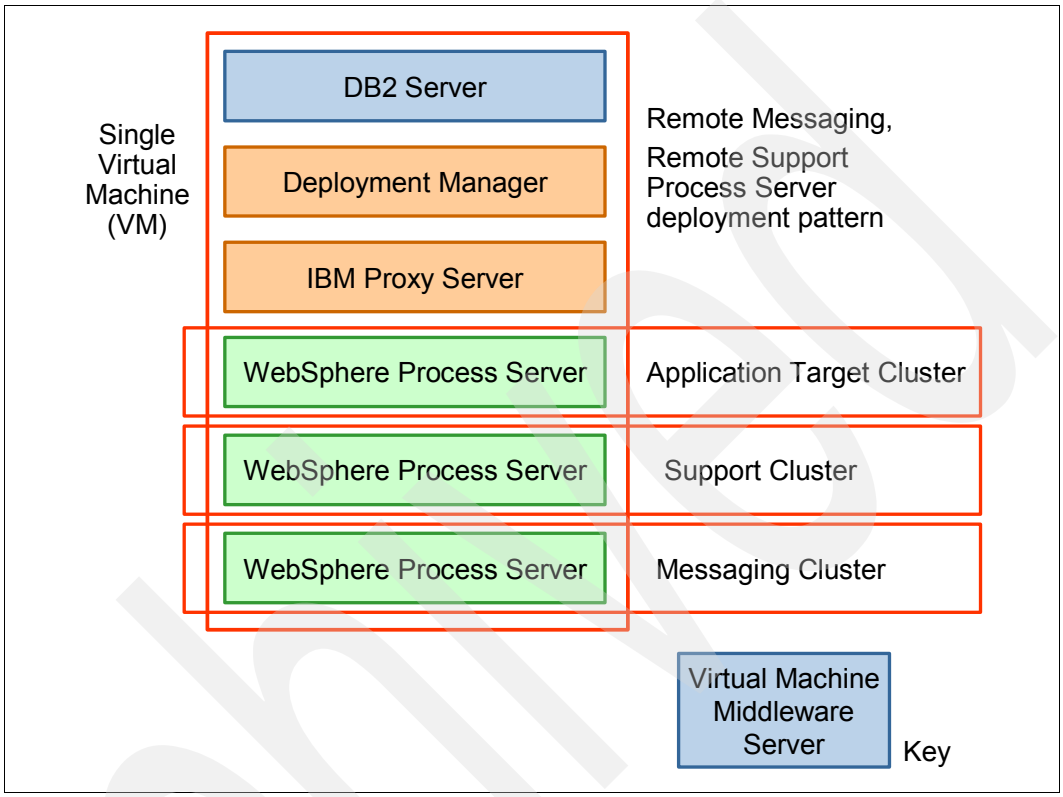

Figure 7-11 WebSphere Process Server Hypervisor Edition - Simulated scalable environment

The following steps are required before loading the WebSphere Process Server Hypervisor Edition image:

1. An IBM Password Advantage account so that you can download the image.
2. An image size is 6062 MB. You need to have 20% extra free space to uncompress and extract the download package.
3. Java Runtime Environment (JRE) Version 6 SR3 or later installed.
4. 10 MB of free space in the home directory of the user who will run image loader.
5. Verify that you can establish an active connection to WebSphere CloudBurst Appliance and that you have one of the following permissions necessary to load virtual images to the catalog:
   - Create new catalog content.
   - Cloud administration.

> **Note:** Image loader is supported on Windows and Linux platforms only.

The procedure is:

1. Go to the directory where you extracted the image and run `loader.bat` from Windows or `loader.sh` from Linux.
2. In the Host Name field enter the fully qualified domain name or IP address of the appliance where you want to load the virtual image and patterns.
3. In the User name and Password fields, enter the user name and password required to log on to the appliance.
4. Optional: In the Language menu, select the language that you want to use for pattern names. The default value is determined by the locale of the machine used to run image loader.
5. Optional: Click **Test Environment** to determine whether image loader can successfully access the appliance.

   Note: If image loader indicates an error with the connection, verify that your host name, user name, and password are correct, and then test the connection again.Click **Load Image** to add the WebSphere Process Server Hypervisor Edition virtual image to the WebSphere CloudBurst Appliance image catalog and to add the predefined patterns to the appliance.

6. When the loading process is complete, close image loader.
7. Log in to WebSphere CloudBurst Appliance, navigate to **Catalog** → **Virtual images**, and accept all licenses for the image that you loaded.

7.3 Introduction of script packages

A script package is designed to customize the parts within a pattern. The intention of using script packages is to further enable you to customize your WebSphere Application Server environment beyond the customization provisions that are standard with WebSphere CloudBurst. Script packages allow you to make late binding (during deployment) customizations and to customize the middleware layer of your pattern.

A script package is associated with a virtual image part. Upon deployment, the virtual image part becomes a virtual machine, and the script package associated with the virtual image part gets executed on that virtual machine. Associating a script package with a virtual image part is as easy as dragging and dropping the script package onto the virtual image part.

Script package scenario

This section describes steps for creating and deploying the script package into the cloud.

After you create file containing your executable and supporting articrafts, you can create the script package in WebSphere CloudBurst:

1. Log in to the WebSphere CloudBurst administrative console.
2. Navigate to **Catalog** → **Script Packages** to create script package (Figure 7-12).

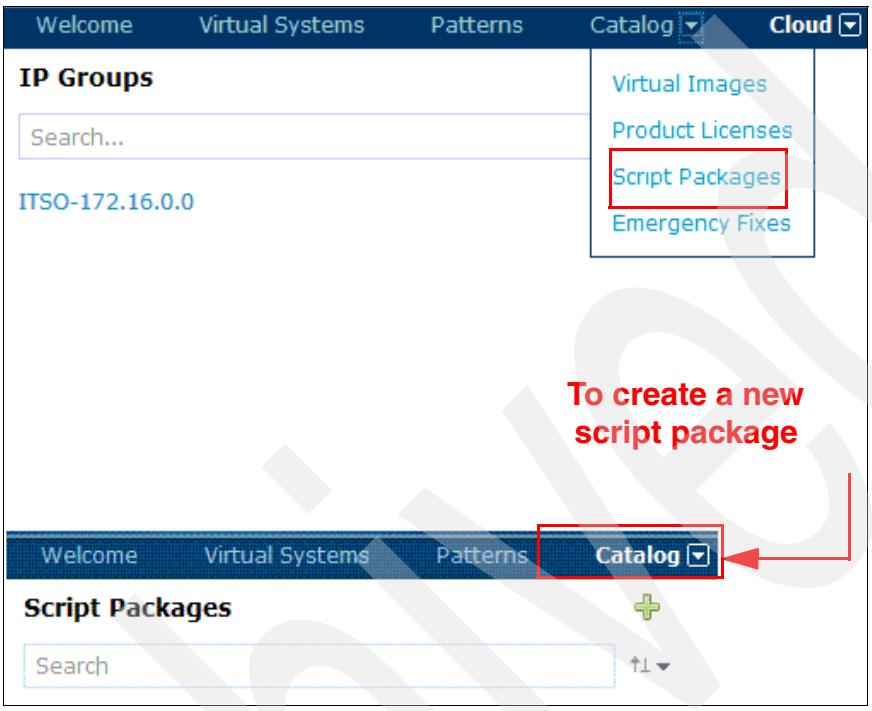

Figure 7-12 Create a script package

3. Upload your archive file to WebSphere CloudBurst. Click **Browse** (Figure 7-13), choose your archive from the local file system, and then click **Upload**. The amount of time this operation takes depends on the size of your archive and the speed of your network connection. When this step completes, you will receive a successful confirmation message below the Upload button.

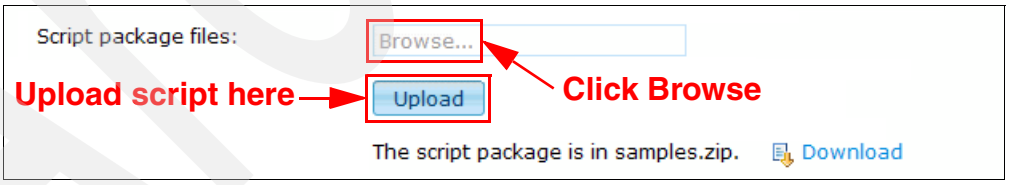

Figure 7-13 Create script package

4. Define the executable that you want WebSphere CloudBurst to invoke (Figure 7-14).

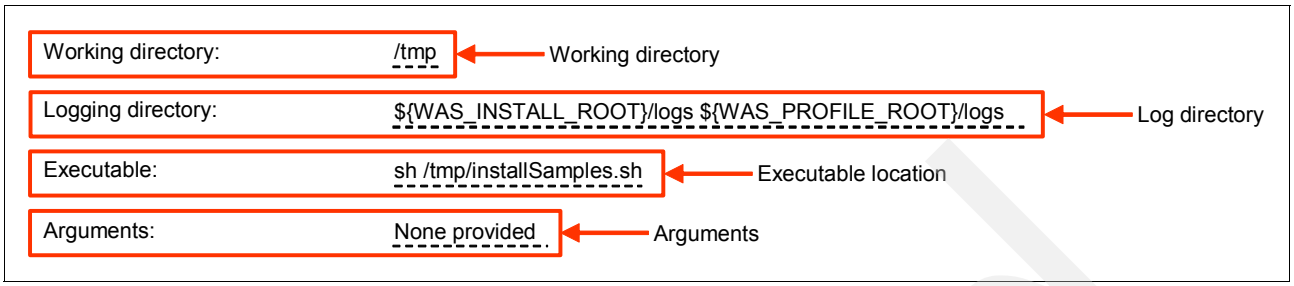

Figure 7-14 Script Package options

5. Define your arguments (Figure 7-14). Arguments defined here are passed to the executable when it is run.
6. Define your working directory (Figure 7-14). The value of the working directory field indicates the directory into which the contents of the archive in your script package will be extracted.
7. The Logging directory field (Figure 7-14) tells WebSphere CloudBurst which directory is used to store logs produced by the execution of the script package.
8. User-defined variables are defined in the Environment field (Figure 7-15). These variables are specific to the script package in which they are defined. After a variable has been defined in the Environment field, it can be used in the scripts defined within the archive or in the definition of the script package itself. You can enter a key and default value or just a key. If you enter the key only, you can override the value at deployment time. In Figure 7-15, define a DB2_HOSTNAME variable with no associated value, so the value will need to be defined before deployment.

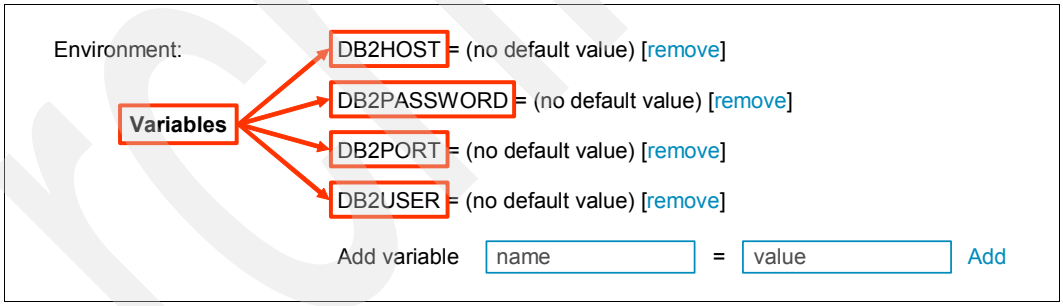

Figure 7-15 Script Package variables

Table 7-1 Subset of predefined WebSphere Application Server environment variables

Variable	Description
CELL_NAME	WebSphere cell name
NODE_NAME	WebSphere node name
WAS_INSTALL_ROOT	WebSphere Application Server installation root
WAS_PROFILE_ROOT	WebSphere Application Server profile root
PROFILE_NAME	WebSphere Application Server profile short name
HOSTNAME	Host name of the virtual machine hosting the WebSphere Application Server instance

Variable	Description
WAS_USERNAME	WebSphere Application Server user name
WAS_PASSWORD	WebSphere Application Server password

7.4 Patterns

This section discusses the pre-loaded patterns that come shipped with the WebSphere CloudBurst Appliance. WebSphere CloudBurst Appliance ships with industry best practices patterns. Figure 7-16 shows a set of pre-loaded patterns that come with the 2.0 appliance. When creating your own custom pattern you can use one of the pre-loaded patterns as a starting point by making a clone of the pattern. By starting with one of the WebSphere CloudBurst existing patterns you are starting with a time-tested approach. The pre-loaded patterns range in complexity from a simple server topology to a highly available cluster topology.

All of the pre-loaded patterns are based on the default WebSphere Application Server virtual image. If you want to use another virtual image as your source you need to clone the pattern, change the virtual image, and deploy the new pattern.

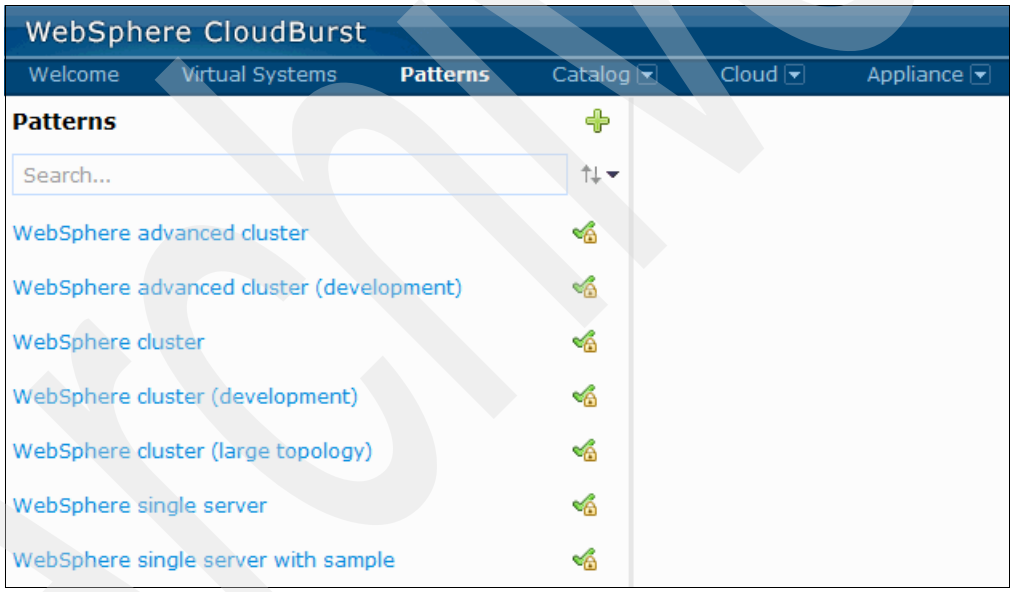

Figure 7-16 Pattern pre-loaded with a 2.0 appliance

There are types of pre-loaded patterns available with IBM WebSphere CloudBurst, as discussed in the following section.

7.4.1 Single server patterns

The single server patterns that come predefined with WebSphere CloudBurst are:

- WebSphere single server

  Single server is a WebSphere Application Server topology or part of a WebSphere Application Server network deployment topology. The single node can be used for a development environment or as part of a multiple-node, production environment in which the application configuration is manually duplicated. Figure 7-17 shows a WebSphere Single Server with one stand-alone server for a total of one virtual machine.

  – One stand-alone server
  – Total of one virtual machine

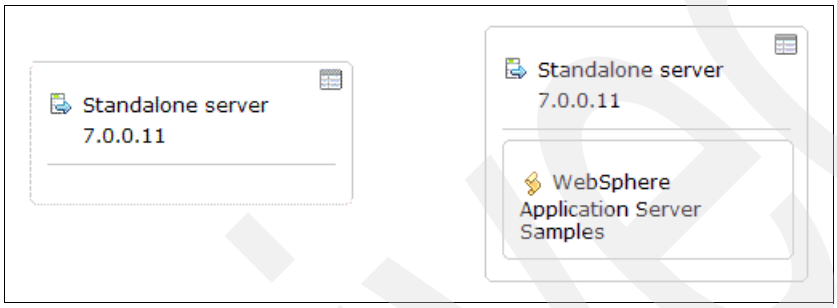

Figure 7-17 WebSphere single server and WebSphere single server with sample

- WebSphere single server with sample

  WebSphere single server with sample is a WebSphere Application Server topology or a part of a WebSphere Application Server Network deployment topology. The single node can be used for a development environment or as part of a multiple-node, production environment in which the application configuration is manually duplicated. This includes the plants by WebSphere example script packages (Figure 7-17).

7.4.2 Cluster patterns

The cluster patterns that come pre-loaded with WebSphere CloudBurst Appliance are:

- WebSphere cluster

  Cluster is a WebSphere Application Server Network Deployment topology for larger scale development or production environments. The IBM HTTP Server resides in a dedicated virtual machine (in an unmanaged node) (Figure 7-18). A WebSphere cluster pattern consists of:

  – One deployment manager
  – Two custom nodes
  – One IBM HTTP Server
  – Total of four virtual machines

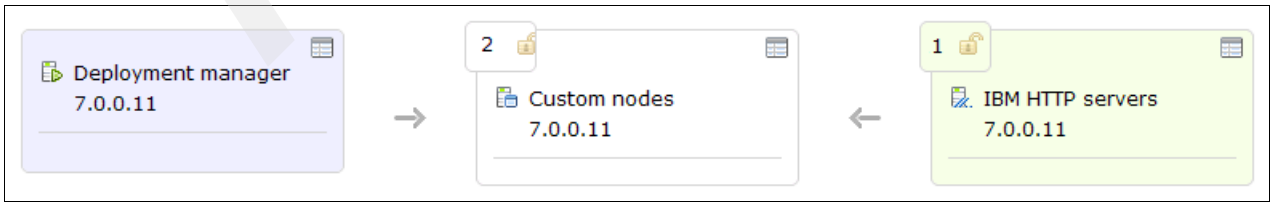

Figure 7-18 WebSphere cluster

- WebSphere advanced cluster

  Advanced cluster is built from a version of the WebSphere Application Server Hypdervisor Edition in which the Intelligent Management Pack is enabled. This pattern is used for large-scale development or production environments (Figure 7-19). A Websphere advanced cluster pattern consists of:

  – One deployment manager
  – Four custom nodes
  – Two on demand routers (See “On demand router” on page 162.)
  – Two IBM HTTP Servers
  – Total of nine virtual machines

  Note: The on demand router concept is discussed in 7.2.2, “WebSphere Intelligent Management Pack” on page 161.

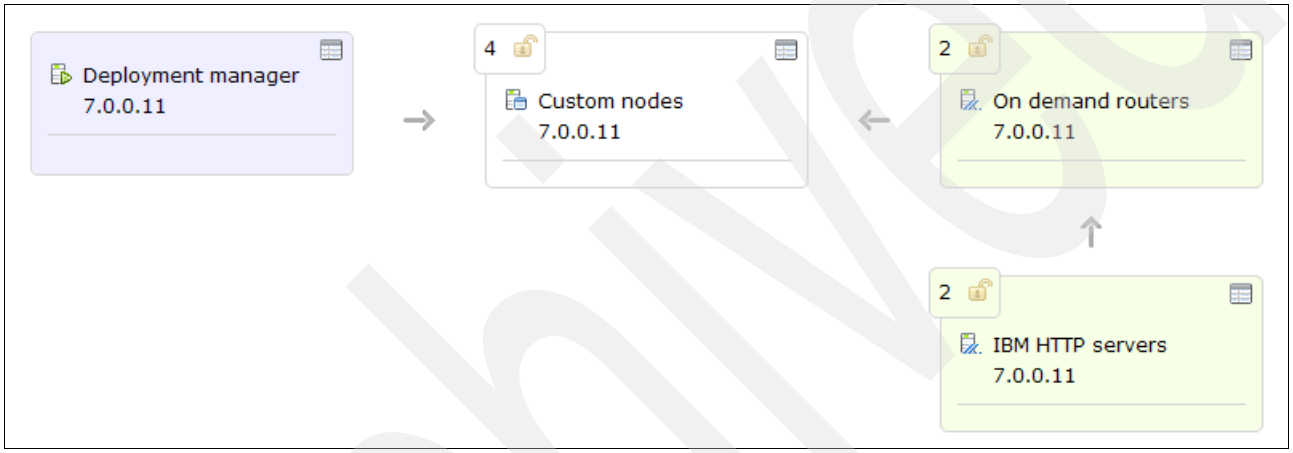

Figure 7-19 WebSphere advanced cluster

- WebSphere cluster (development)

  Cluster (development) is a WebSphere Application Server Network Deployment topology for small-scale development environments. The deployment manager and IBM HTTP Server (in an unmanaged node) are located on the same virtual machine, reducing the number of virtual machines deployed (Figure 7-20). Websphere cluster development pattern consists of:

  – One deployment manager
  – Two custom nodes
  – Total of three virtual machines

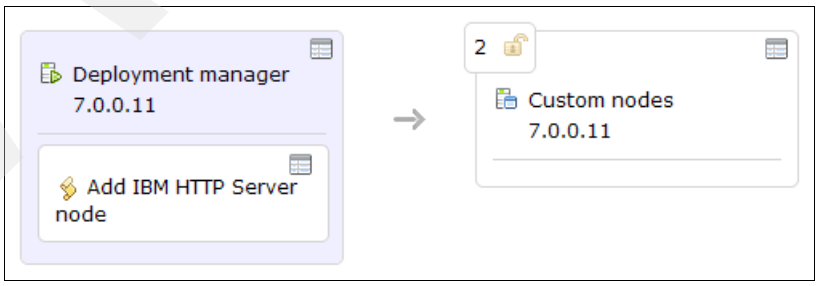

Figure 7-20 WebSphere cluster (development)

- WebSphere advanced cluster (development)

  Advanced cluster (development) is built from a version of the WebSphere Application Server Hypdervisor Edition in which the Intelligent Management Pack is enabled. This pattern is targeted for the development environment, but can be used for production (Figure 7-21). WebSphere advanced cluster development pattern consists of:

  – One deployment manager
  – Two custom nodes
  – One on demand router (“On demand router” on page 162)
  – Total of four virtual machines

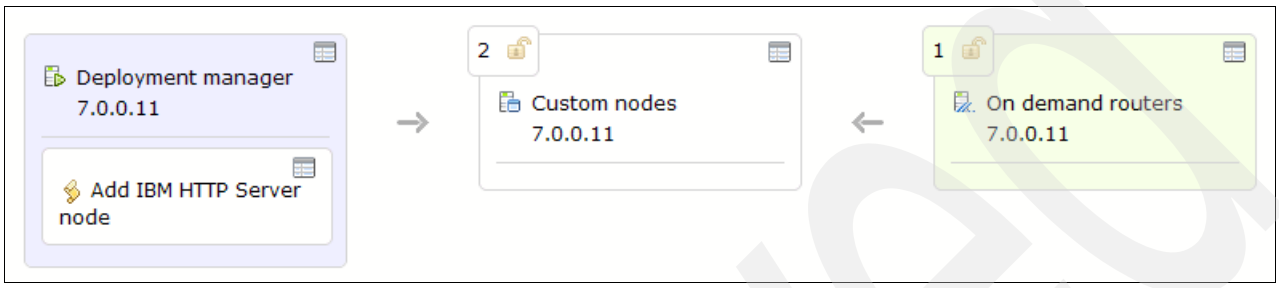

Figure 7-21 WebSphere advanced cluster (development)

- WebSphere cluster (large topology)

  Cluster (large topology) is a WebSphere Application Server Network Deployment topology for large-scale environments. The topology consists of a development manager, ten custom nodes with the same properties, and an IBM HTTP Server (Figure 7-22). WebSphere cluster large topology pattern consists of:

  – One deployment manager
  – Ten custom nodes
  – Four IBM HTTP Server
  – Total of 15 virtual machines

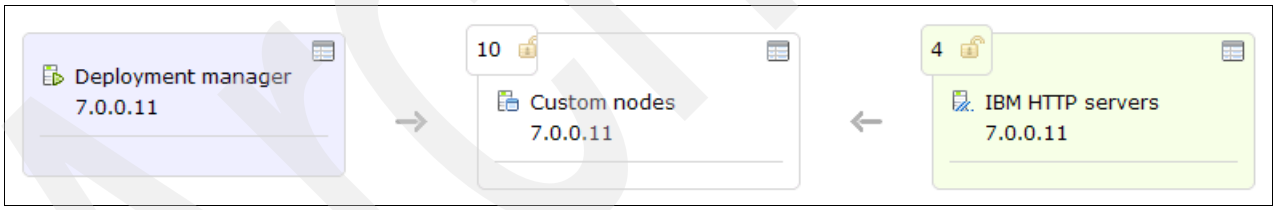

Figure 7-22 WebSphere cluster (large topology)

Attention: The WebSphere advanced cluster pattern and the WebSphere advanced cluster (development pattern are available with the WebSphere Application Server 7.0.0.11 with the Intelligent Management Pack virtual image. Depending on the version of your WebSphere CloudBurst Appliance, you have the following options:

- If your appliance came with IBM WebSphere CloudBurst Appliance Version 2.0 firmware pre-installed, it has this virtual image.
- If you upgrade your appliance to Version 2.0 firmware, you must import the WebSphere Application Server hypervisor Edition 7.0.0.11 virtual image. You need to recreate the patterns from scratch. You must also enable the Intelligent Management Pack function.

You can work with patterns in the following ways:

- Use predefined patterns as they are.
- Copy predefined patterns and customize them.
- Create new patterns.
- Edit any patterns that are read-only.

7.4.3 Customizing patterns

This section covers customization of patterns, which is customizing patterns using existing images. Images are installed and configured once, then copied and reused over and over again. The content of a script package is not defined or restricted by WebSphere CloudBurst. Script packages can be used to further configure your environment.

Pattern creation

This section describes pattern creation. You can follow the steps to create a pattern from a clone.

Cloning pattern

Cloning a pattern creates a copy of the source pattern. A clone of a pattern is a copy that can be edited.

Steps to create pattern using clone method

In this example we use the WebSphere cluster pattern and a virtual image of WebSphere Application Server 7.0.0.11 (PowerVM) with Intelligent Management Pack 6.1.1.2.

1. Navigate to **Patterns**.
2. Select **WebSphere Cluster**.
3. Click **Clone** (Figure 7-23).

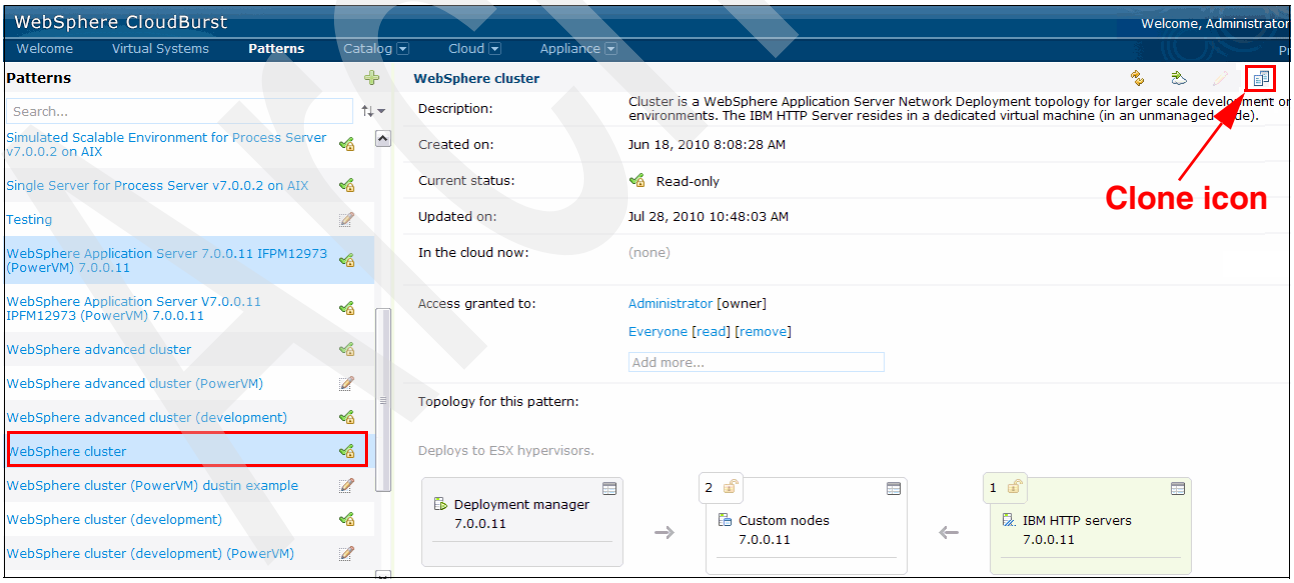

Figure 7-23 Pattern clone

4. On the Pattern page, enter the name of the pattern and a description (Figure 7-24).

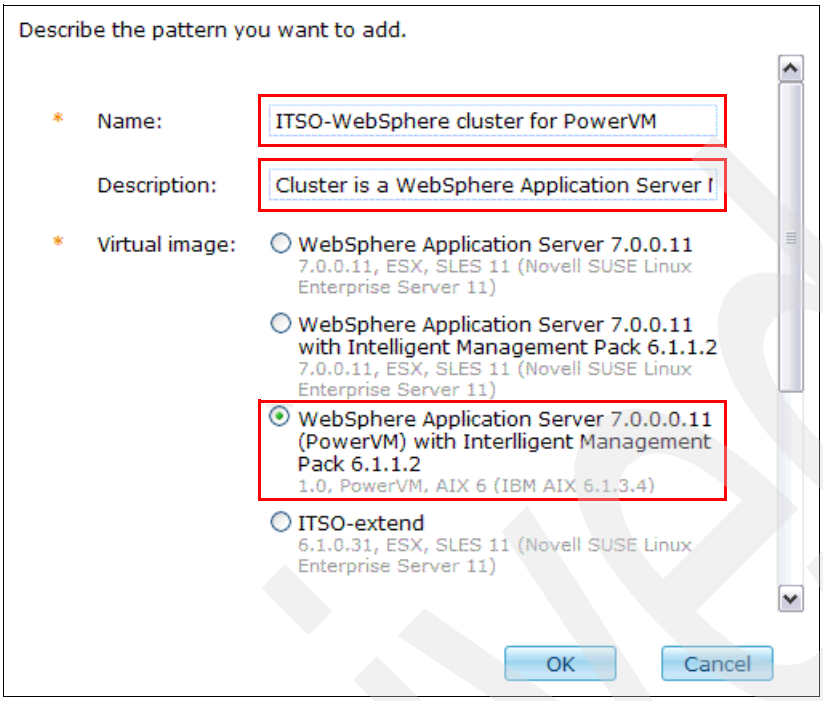

Figure 7-24 Pattern clone - Name, description, and image selection

5. Select **Image** (Figure 7-24).
6. Click **OK** to save your changes. When the information is processed, you are returned to the Patterns panel, and the pattern that you created is added to the list in the left panel. It is selected so that the information about it displays in the right panel.
7. To change the virtual image and the pattern topology, click **Edit** on the top of the right panel of the Patterns view.

8. Edit the parts of canvas as follows:
   - Add extra custom nodes: Drag and drop in to the canvas shown in Figure 7-25.
   - Add custom script package: Drag the script package into the part shown in Figure 7-25.

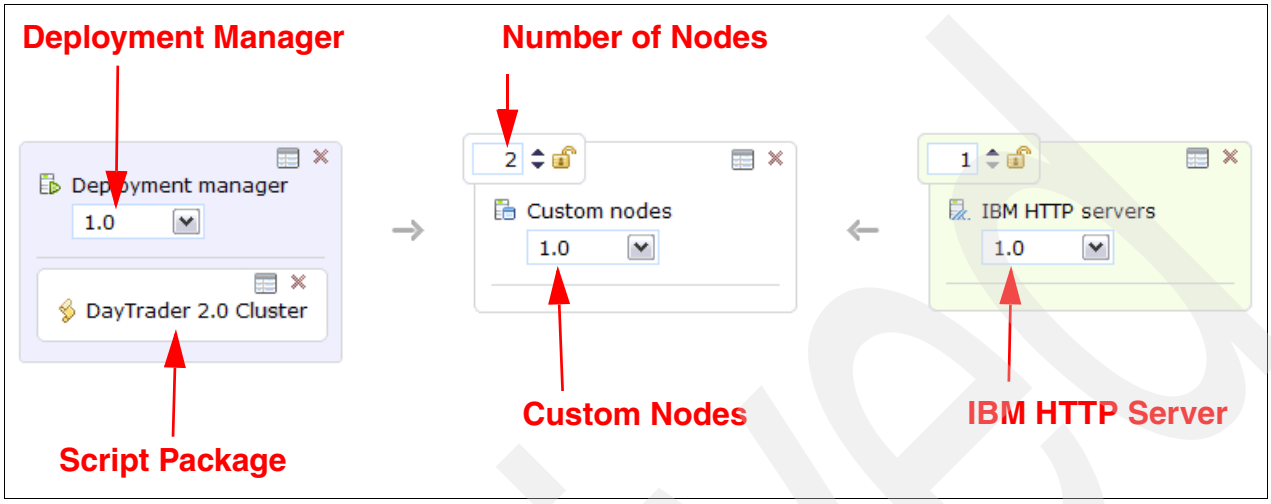

Figure 7-25 Pattern canvas with parts

> **Note:** Uniquely name the pattern and give the specific description of the pattern to avoid confusion later.

9. Click **Advanced options** on the top of the right side of the Patterns view.
10. In the Advanced options page, check **Define Clusters** and uncheck Enable Messaging, and then click **OK** (Figure 7-26).

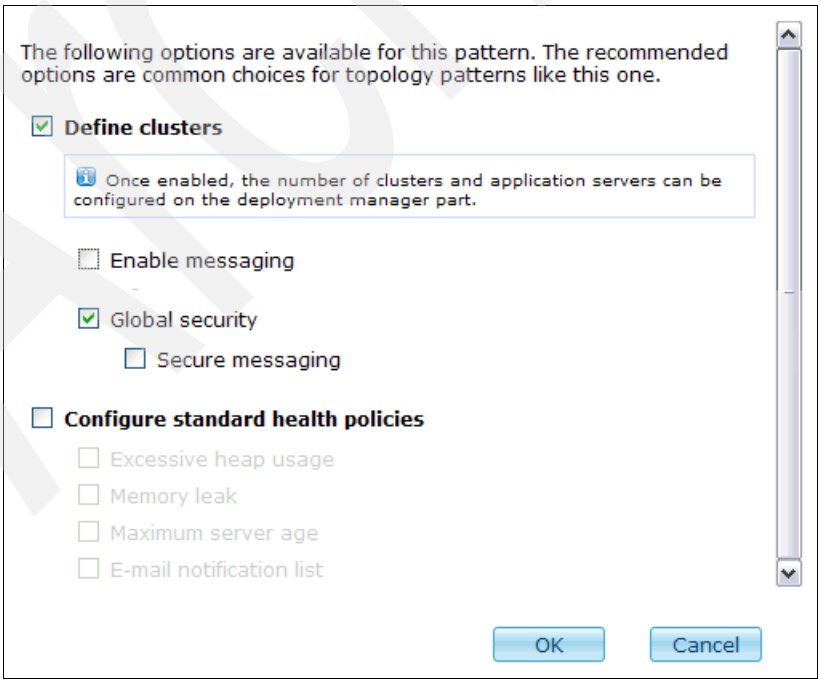

Figure 7-26 Pattern WebSphere Cluster Advanced options window

11. Click **Deployment Manager** to unlock and change the parameters (Figure 7-27 and Figure 7-28).

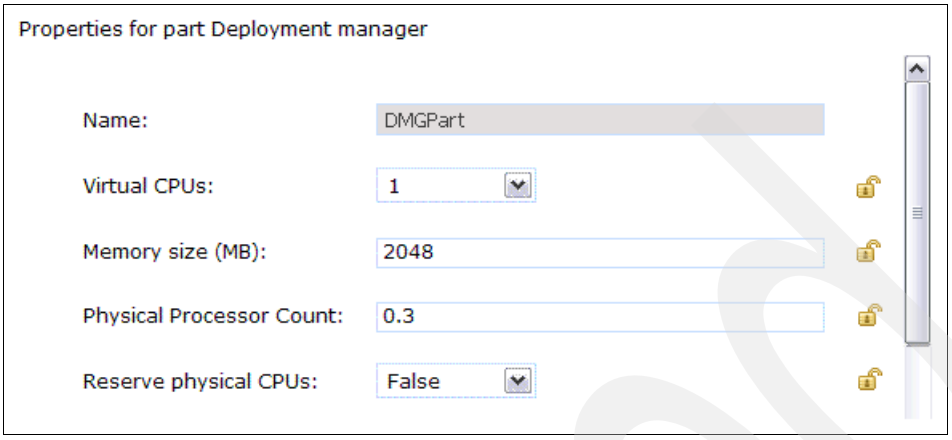

Figure 7-27 Properties of Deployment Manager - Panel 1

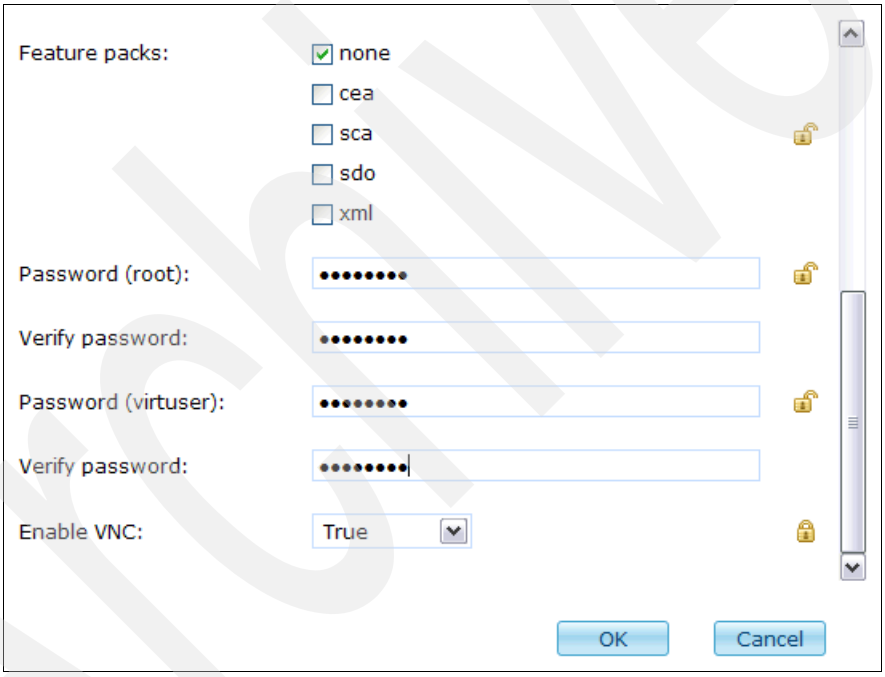

Figure 7-28 Properties of Deployment Manager - Panel 2

> **Note:** virtuser is a operating system and application server used to log in to the WebSphere Application Server Admin Console.

12. You can change the parameters as per your requirement (that is, virtual CPUs, memory size, physical processor count, cell name, node name, password for root, password for virtuser) and click **OK**.
13. Change the properties of IBM HTTP Server.
14. Change the properties of custom nodes.

15. Select **Properties of a Cluster configuration**. Figure 7-19 on page 174 shows the properties that can be modified. You can see the properties in Figure 7-29. You can define the number of clusters and number of nodes per server, and then click **OK**.

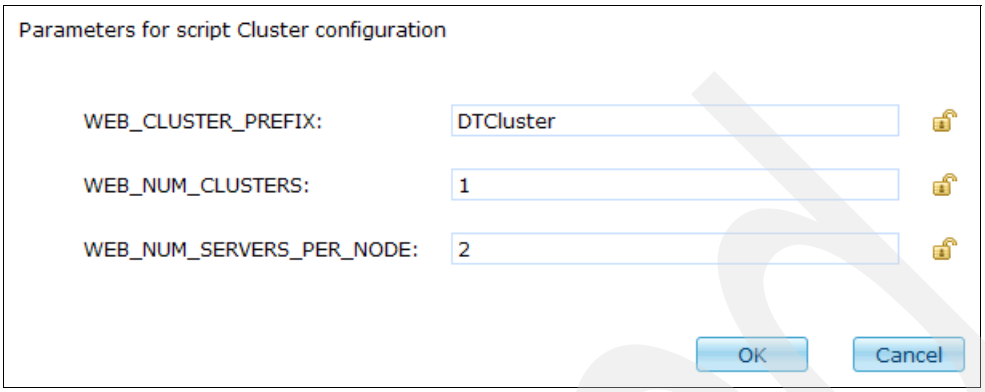

Figure 7-29 Cluster configuration parameters

16. Select **Properties of a script package**. Figure 7-30 shows the properties that can be modified. Here we define DB2HOST, DB2PASSWORD, DB2PORT, and DB2USER, which are part of script package.

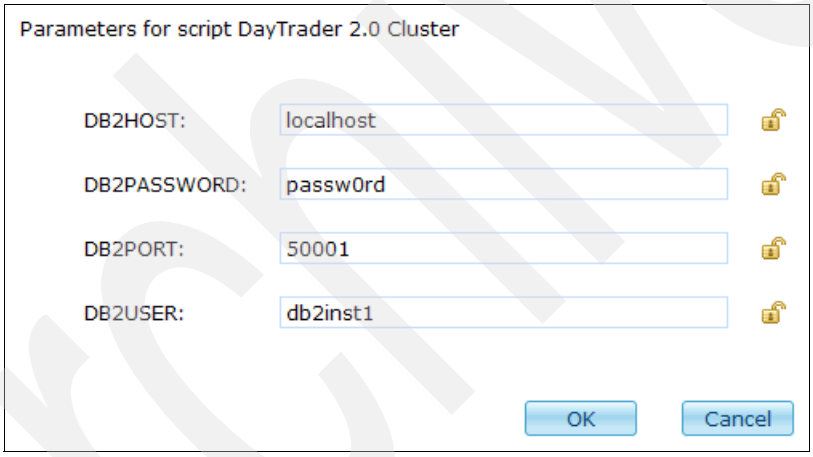

Figure 7-30 Script Package properties parameters

17. Click **Done Editing**.

18. Click **Deploy in the Cloud**. This takes you to the next window, where you can enter the name of the virtual system, schedule the deployment (Figure 7-31), and change the configuration of virtual parts. You can choose the **start now** or the **start later** (specify the date and time of deployment) radio button, and then click **OK**.

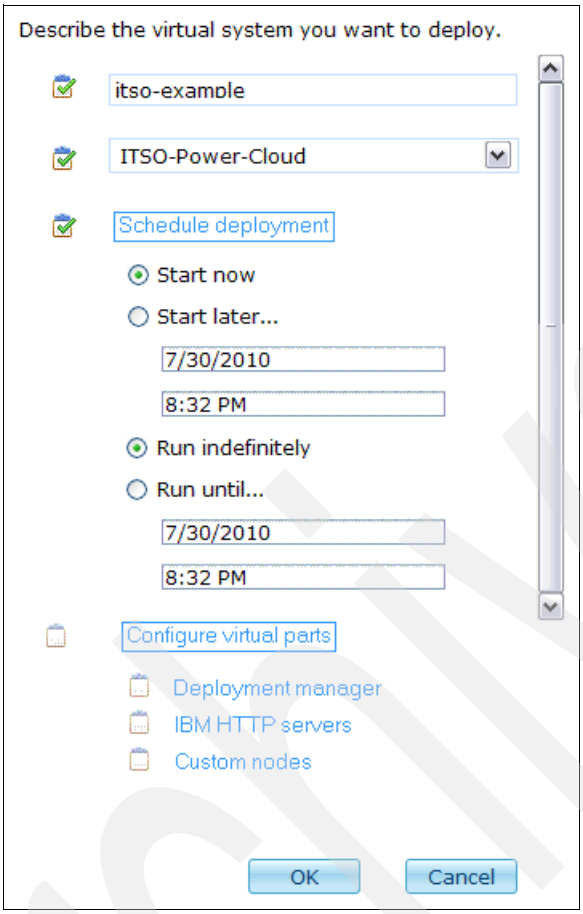

Figure 7-31 Pattern deployment panel

19. Next you are taken to the deployment process. The first step is transferring virtual images to hypervisors (Figure 7-32).

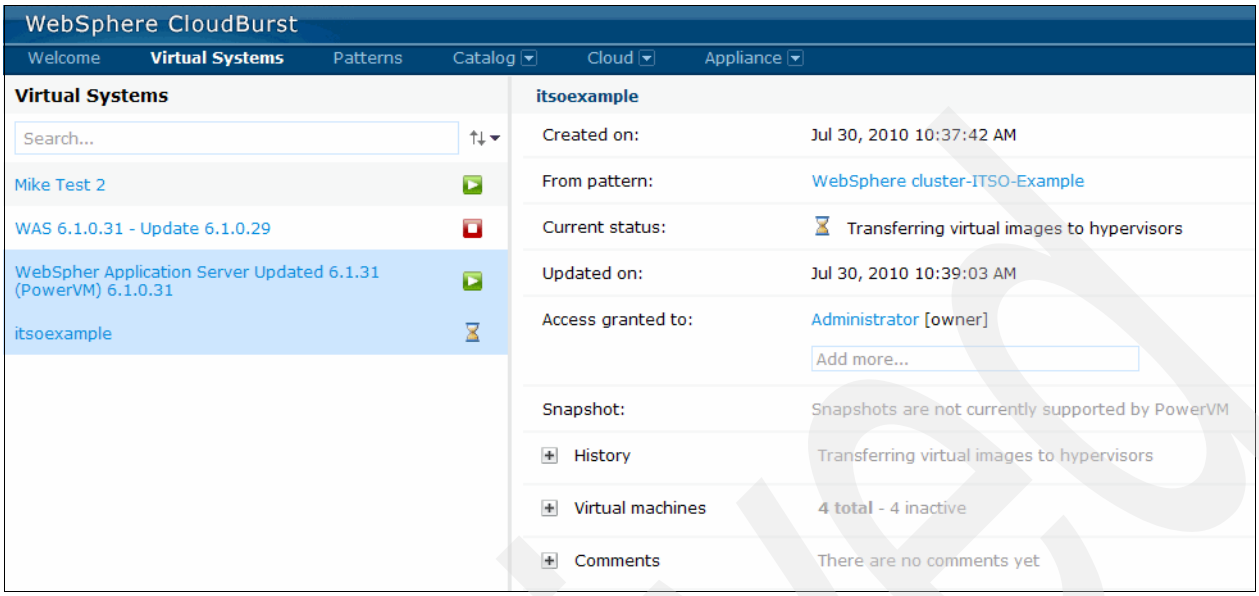

Figure 7-32 Pattern deployment - Transferring image to hypervisor

20. Select **Network Installation Management (NIM) Base operating System Installation** started (Figure 7-33).

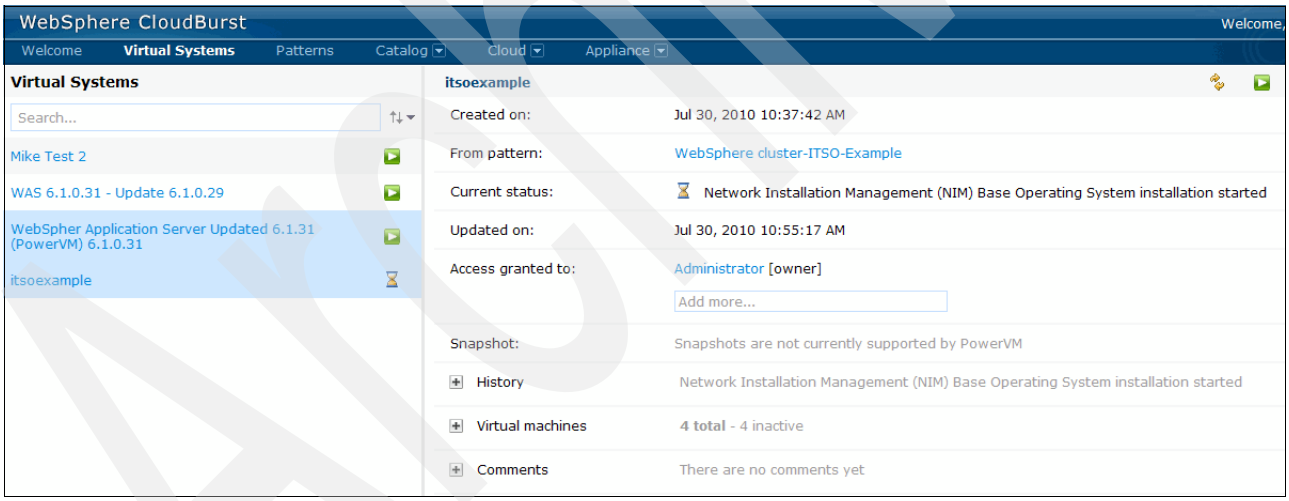

Figure 7-33 Pattern deployment - Network Installation Management (NIM) Base Operating System installation

In the next window, virtual systems has been deployed and is ready to use (Figure 7-34).

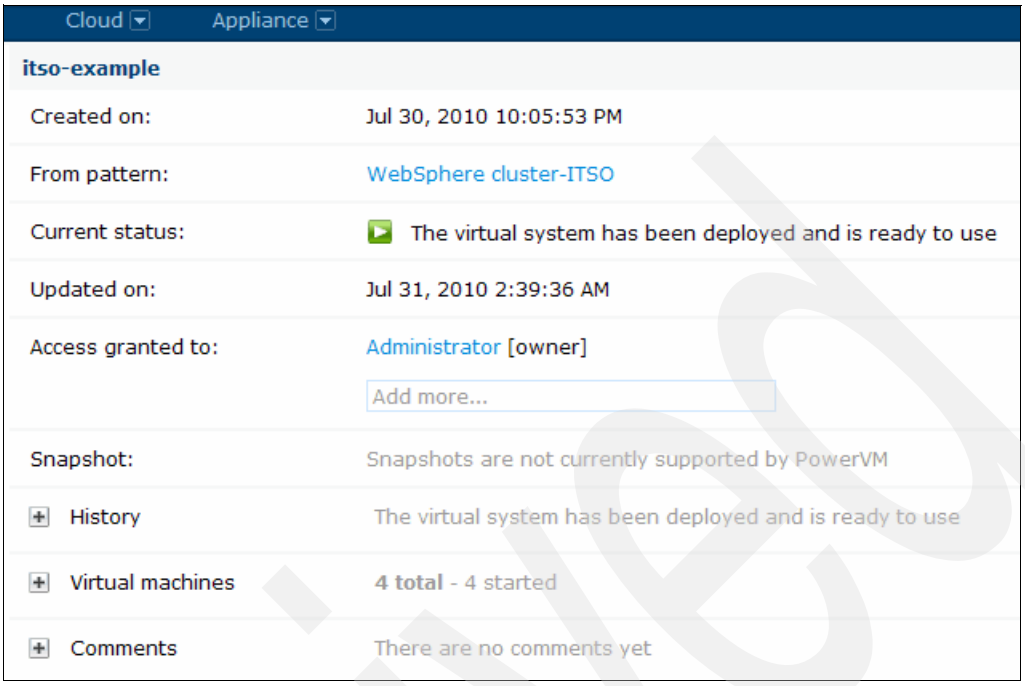

Figure 7-34 Virtual systems deployed

In the next window, you will see that four machines were created (Figure 7-35).

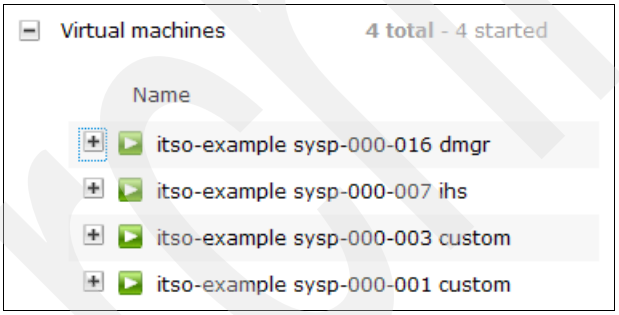

Figure 7-35 Virtual systems - Node information

21. Make sure that virtual machines are started, and then log in to the WebSphere Admin Console by clicking the **WebSphere link** in the Deployement Manager virtual machine section (you can collapse the virtual machines section, and then the deployment managers virtual machine section) (Figure 7-36). You are taken to a page that asks you to log in. The user name is virtuser and the password is the virtuser password that you supplied during deployment. If you cannot get to this link, try using the IP address for the virtual machine instead of the host name. If you still cannot get to it, you can have private IP addresses and you will need to ask the system administrator who manages the device about getting access to them or about getting public IP addresses.

Figure 7-36 WebSphere Admin Console link

22. After you are in the WebSphere Admin Console, expand the **Servers** section and then the **Clusters** section. Click **WebSphere application server clusters**. Select the check box and click **Start**. After few minutes everything starts.

23. Now you can access the Day Trader splash window. Go to the Virtual machine section on the WebSphere CloudBurst Admin Console and expand one of the custom nodes. Note the IP address. Use the URL format `http://<Virtual Machine IP-Address>:9080/daytrader` to access the Day Trader page (Figure 7-37).

Figure 7-37 Day Trader page

Abbreviations and acronyms

CAS	Common Agent Services
HEA	Host Ethernet Adapter
HMC	Hardware Management Console
IBM	International Business Machines Corporation
ILMT	IBM License Metric Tool
ITSO	International Technical Support Organization
IVM	Integrated Virtualization Manager
JDBC	Java Database Connectivity
JRE	Java Runtime Environment
LDAP	Lightweight Directory Access Protocol
LPAR	Logical Partitions
MES	Miscellaneous Equipment Specification
NIM	Network Installation Manager
NTP	Network Time Protocol
OVF	Open Virtualization Format
PVU	Processor Value Unit
RHEL5	RedHat Enterprise Linux 5
RMC	Remote Monitoring and Control
SAN	Storage Area Network
SEA	Shared Ethernet Adapter
SLES11	SUSE Linux Enterprise Server 11
SMTP	Simple Mail Transfer Protocol
SNMP	Simple Network Management Protocol
SSH	Secure Shell
VIOS	Virtual I/O Server

Related publications

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

Online resources

These websites are also relevant as further information sources:

- For information about OpenSSL and OpenSSH versions, refer to IBM Systems Director Server on AIX in the IBM Systems Director 6.1.x Information Center

  http://publib.boulder.ibm.com/infocenter/director/v6r1x/index.jsp?topic=/director.install_6.1/fqm0_t_preparing_to_install_ibm_director_on_aix.html
- Open IBM Systems Director Agents web page

  http://www.ibm.com/systems/software/director/downloads/agents.html
- IBM Systems Director Plug-ins

  http://www.ibm.com/systems/software/director/downloads/plugins.html
- IBM Systems Director 6.1.x Information Center

  http://publib.boulder.ibm.com/infocenter/director/v6r1x/index.jsp
- IBM Fix Central

  http://www.ibm.com/support/fixcentral
- Information Center WebSphere CloudBurst

  ttp://publib.boulder.ibm.com/infocenter/wscloudb/v2r0/topic/com.ibm.websphere.cloudburst.doc/pc/pct_package_ef.html

How to get IBM Redbooks publications

You can search for, view, or download Redbooks publications, Redpapers publications, Technotes, draft publications and Additional materials, as well as order hardcopy Redbooks publications, at this website:

ibm.com/redbooks

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services

WebSphere Cloudburst Appliance and PowerVM

Discover PowerVM WebSphere Cloudburst Environment

Install and administer WebSphere Cloudburst Appliance

Learn by example and practical scenaios

This IBM Redbooks publication discusses the concepts and implementation of PowerVM and the WebSphere CloudBurst appliance. This book is aimed at administrators and developers who have little knowledge of PowerVM, but in-depth knowledge of WebSphere software.

Cloud computing is the pooling of computing resources to provide a single source of computing power to multiple users. A cloud manager provides a self-service portal that maintains permissions and information about cloud objects such as virtual images, patterns, and resources. The WebSphere CloudBurst Appliance represents a cloud manager. It is a secure hardware appliance that optimizes the configuration, deployment, and management of WebSphere Application Server environments in a cloud. It can also be used by service providers providing hosted public clouds and software-as-a-service environments to simplify and standardize repeated deployments of their software.

The book includes an introduction to cloud computing and how the WebSphere CloudBurst appliance fits into business today. It presents the features and benefits of using the WebSphere CloudBurst Appliance and the advantages of using PowerVM. It provides the steps required to implement WebSphere CloudBurst appliance with.

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

SG24-7806-00 ISBN 0738434825