IBM Software Group

IBM WebSphere CloudBurst Appliance

Overview

e business on demand.

Updated July 2, 2009

This presentation provides an overview of the new IBM WebSphere® CloudBurst Appliance product.

Agenda

- Background
- Introduction to the CloudBurst appliance and WebSphere Application Server Hypervisor Edition
- Managing WebSphere Application Server life cycle through CloudBurst
- Users and groups
- Other CloudBurst functions
- Data center integration
- Summary and references

The presentation starts with a background of some of the issues facing the enterprise in terms of the increased costs, followed by introduction of CloudBurst and its vision. This is followed by details of how CloudBurst manages the WebSphere Application Server deployment life cycle, followed by other important topics.

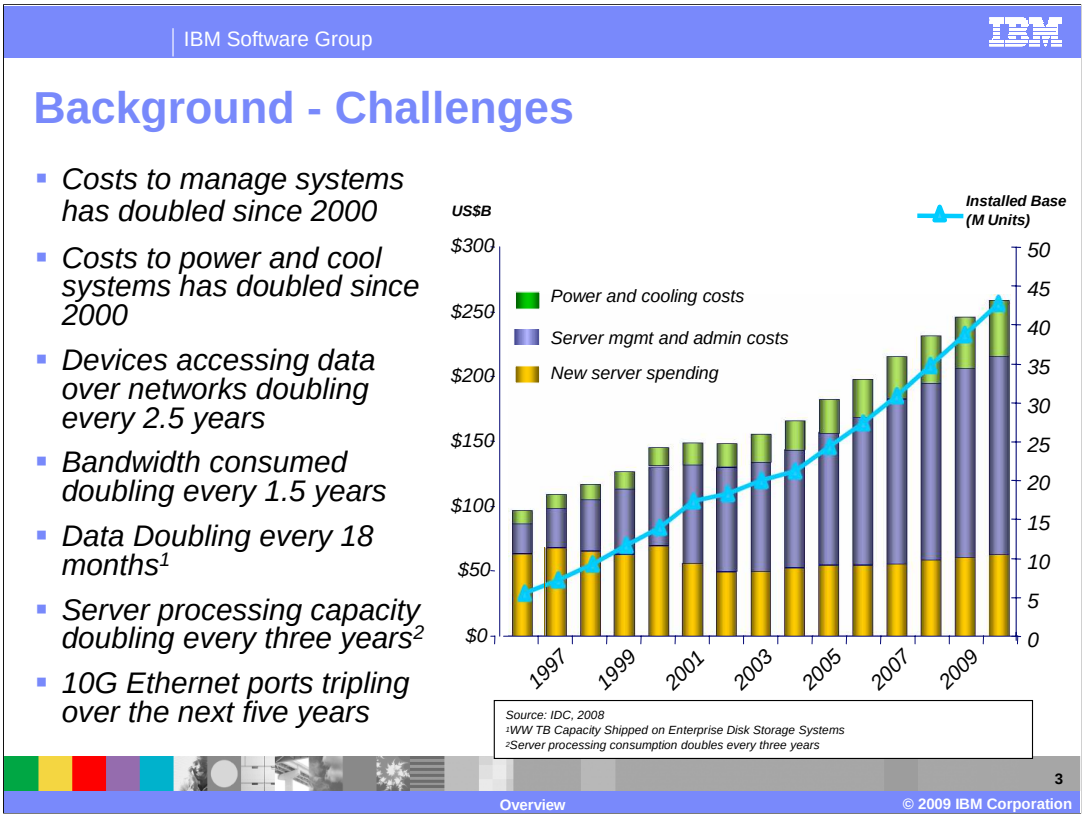

With high operational costs and many under-used servers consuming space and power, enterprise data center staff feels increasing pressure to better use existing resources and reduce costs. The figure on this page shows the increasing costs of power and cooling along with server management and administration costs. These costs are driving an interest in both server consolidation and improved management capabilities to offset the trend. 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.

In addition to consolidation to reduce power and space costs, reducing operational costs is critical. Much of the administrative cost is driven by the number of steps and decisions involved in designing, installing, configuring, and maintaining solutions composed from different software components. For example, administration of a WebSphere solution includes installing, configuring, and applying patches to operating systems, middleware, and application software across multiple servers. A faster, easily repeatable process to get WebSphere Application Server environments up and running for development, test use and for designing, configuring, and applying on-going maintenance to highly available, robust production environments is needed.

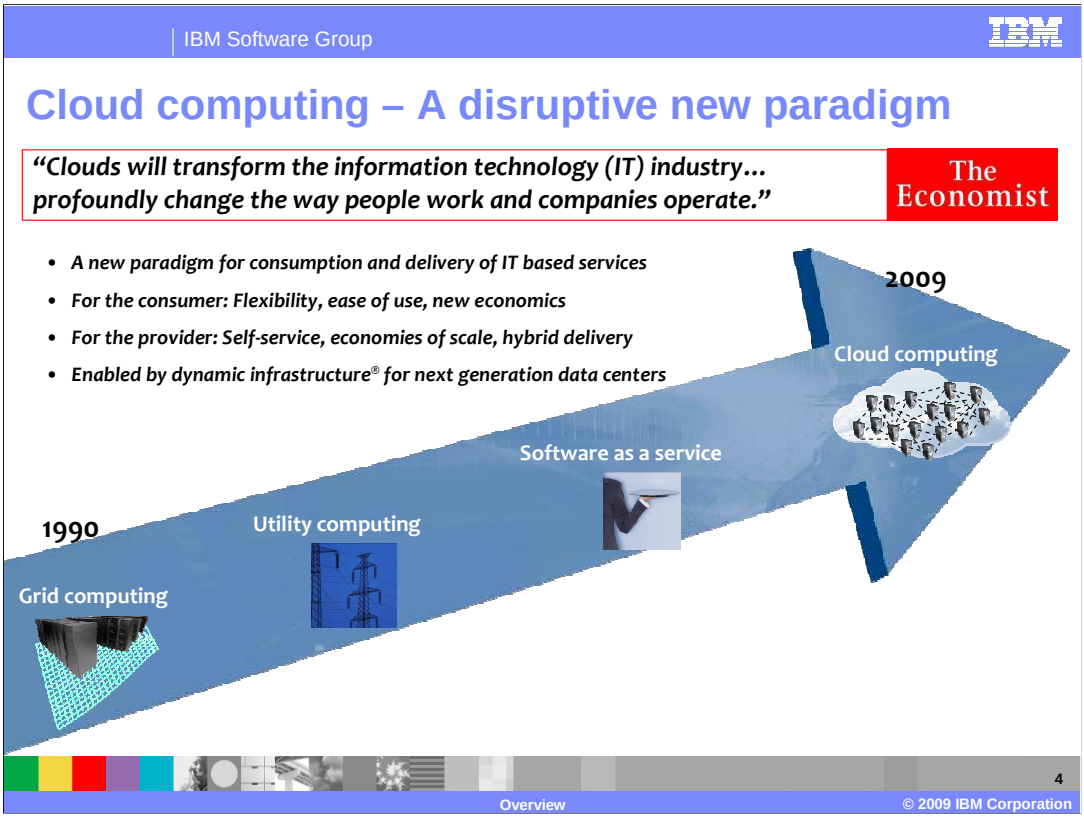

As The Economist reported in October of 2008, cloud computing will transform the IT industry and will impact how businesses operate profoundly.

Cloud computing provides massively scalable computing resources from anywhere. It simplifies service delivery. It enables rapid innovation and creates a dynamic platform for next generation data centers.

Grid computing is solving large problems with parallel computing. Utility computing is offering computing resources as a metered service. Software as a service is network-based subscriptions to applications.

Some say cloud computing is grids or utility computing or software as a service but it is all of those combined. It is anytime, anywhere access to IT resources delivered dynamically as a service.

Section

Introducing IBM WebSphere Cloudburst Appliance and WebSphere Application Server Hypervisor Edition

This section covers the terms and definitions used within the CloudBurst environment.

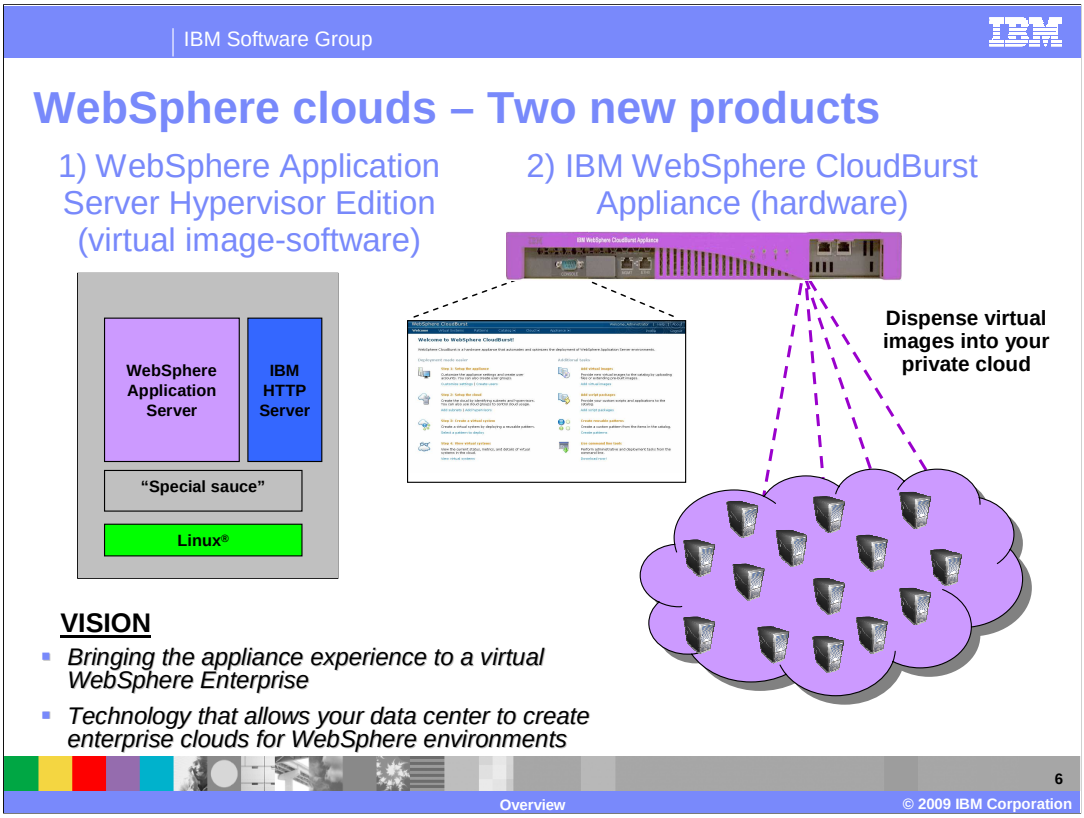

As part of WebSphere's clouds effort, two new products have been introduced.

A new addition to the WebSphere Application Server family is the WebSphere Application Server Hypervisor Edition for V6.1 and V7.0, and second, the IBM WebSphere CloudBurst appliance.

The CloudBurst appliance uses the WebSphere Application Server Hypervisor Edition to create virtual images of WebSphere Application Server, and then dispenses these virtual images into the cloud.

The vision of CloudBurst is to bring an appliance experience to manage the life cycle of WebSphere enterprise virtual environments. It is to provide the technology to allow your data center to create enterprise clouds for WebSphere environments, capitalizing in server virtualization and cloud computing that is gaining popularity in reducing IT costs.

CloudBurst Appliance – Value Proposition

- The business value of WebSphere CloudBurst is to increase efficiency, cost-effectiveness, and usability of WebSphere topologies in a data center by leveraging cloud computing principles
- WebSphere CloudBurst appliance uses key technologies
  - Hardware and virtual appliances
  - Server virtualization capabilities

The emergence of cloud computing within the enterprise aims squarely at reducing IT costs by pooling available resources to reduce the number of physical servers required while also bringing a new level of agility in providing self-service compute resources. At the heart of cloud computing is server virtualization technology, bringing server consolidation, isolation, and rapid provisioning capabilities. These technologies support building next generation data center cloud architectures on many platforms.

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.

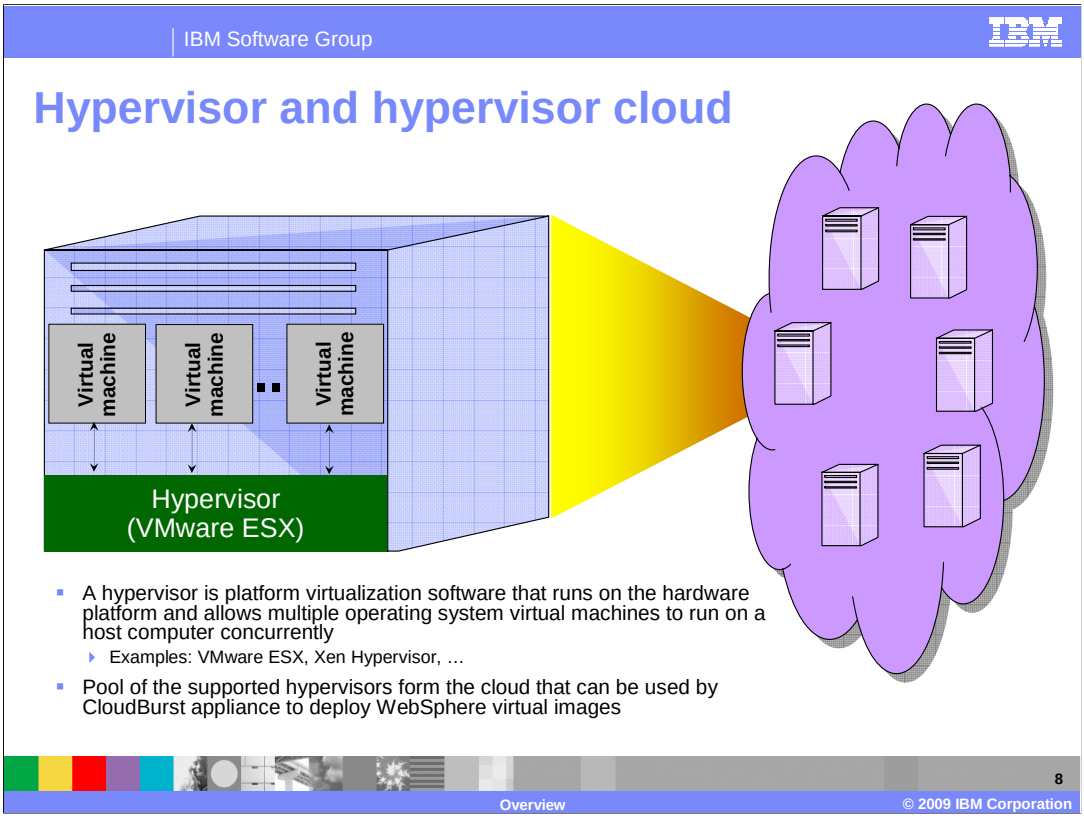

Before diving into the details of CloudBurst appliance, it is necessary to understand what a hypervisor is, and what is meant by the hypervisor cloud.

A hypervisor is a program that allows multiple operating systems – or multiple instances of a single operating system - to run on a host computer concurrently. Examples of hypervisors include VMware ESX and ESXi, Citrix Xen, and z/VM®. The target operating system environment, including any software like WebSphere Application Server that runs on that operating system, is referred to as the virtual machine.

The diagram on the left shows the physical hardware on which the hypervisor is installed, and on which one or more virtual images like VMware can run. A pool of such hypervisors form the cloud - shown on the right - that is then defined within the CloudBurst appliance to be used for deployment of WebSphere Application Server virtual images.

IBM WebSphere Application Server Hypervisor Edition

- WebSphere Application Server Hypervisor Edition is a new addition to WebSphere Application Server family of products optimized to run in virtualized environments
  - Available for V6.1 and V7.0 release
- Contains preinstalled, preconfigured, binary image of the operating system, application server and its profiles, and IBM HTTP Server
- Contains activation code to create virtual machines for different profiles to run on top of the hypervisors
- Image provided using industry standard Open Virtualization format (OVF) 1.0
- Included in CloudBurst appliance, and available as a stand-alone product

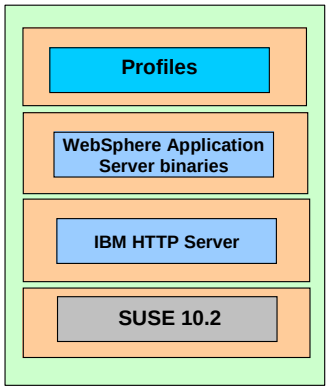

WebSphere Application Server Hypervisor Edition

A new addition to the WebSphere Application Server family of products is the new WebSphere Application Server Hypervisor Edition.

The Hypervisor Edition is optimized for using WebSphere Application Server in virtualized environments on top of the hypervisors. The Hypervisor Edition facilitates more agile usage of WebSphere Application Server by providing a rapid setup or teardown of the Application Server environments.

It contains a preinstalled, preconfigured, OS-included binary image of the application server from which virtual machines can be created and deployed on hypervisors. The Hypervisor Edition is available for WebSphere Application Server V6.1 and V7.0 releases. For each release, the base image contains the Linux operating system, IBM Http Server, WebSphere Application Server binaries, and all the profiles supported for that release. It also contains activation code that is used when creating virtual machines of specific profiles to run on top of the hypervisors.

The IBM provided Hypervisor Edition comes with the SUSE Linux operating system V10.2. Instructions are provided to create your own image using other Linux operating systems, like RedHat. The Hypervisor Edition uses open standard Open Virtualization format, which is an optimized format to store virtual images.

WebSphere Application Server Hypervisor Edition is also available as a stand-alone product.

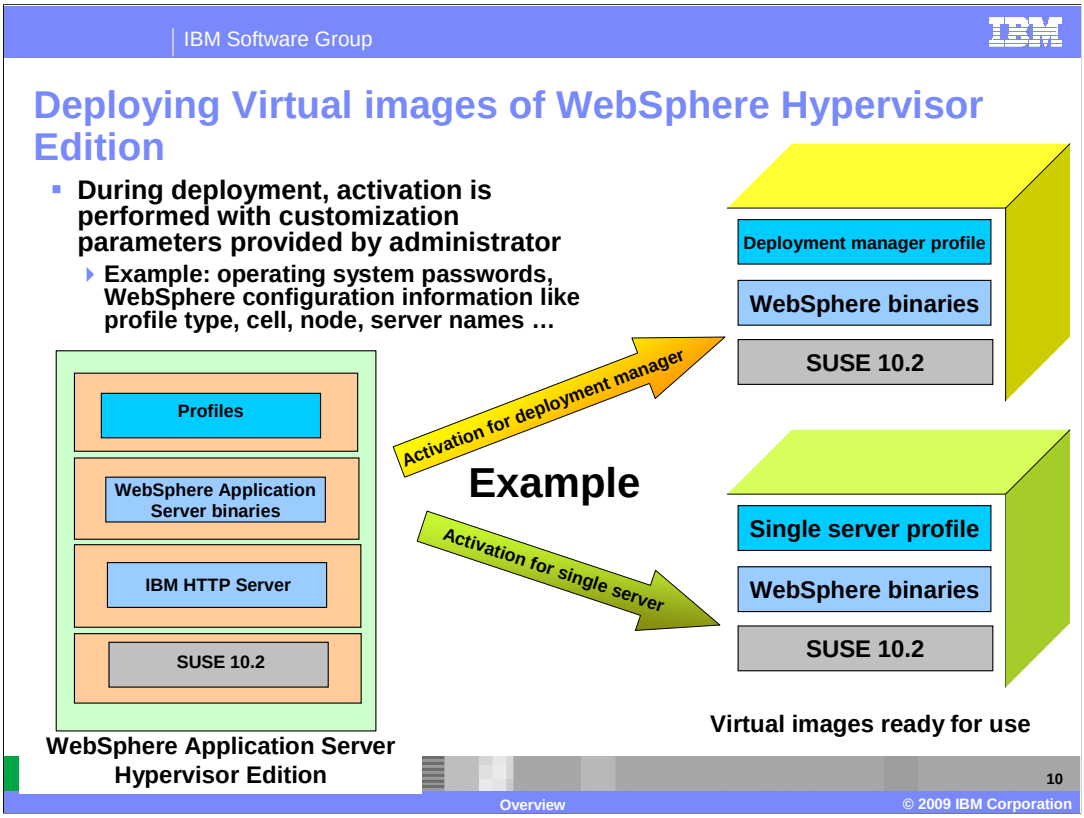

Before the Hypervisor edition can be deployed as a virtual machine, it needs to be activated. The process of activation requires that you provide customization information like operating system root password, WebSphere specific information to create a ready to run virtual image of a specific profile of WebSphere Application Server on the supported hypervisor.

Shown on this page are examples of deploying specific profiles of the WebSphere Hypervisor edition to create virtual images of Deployment manager and Single server. Once the customization parameters are provided, the activation is performed and the images are deployed as virtual machines on the cloud hypervisors.

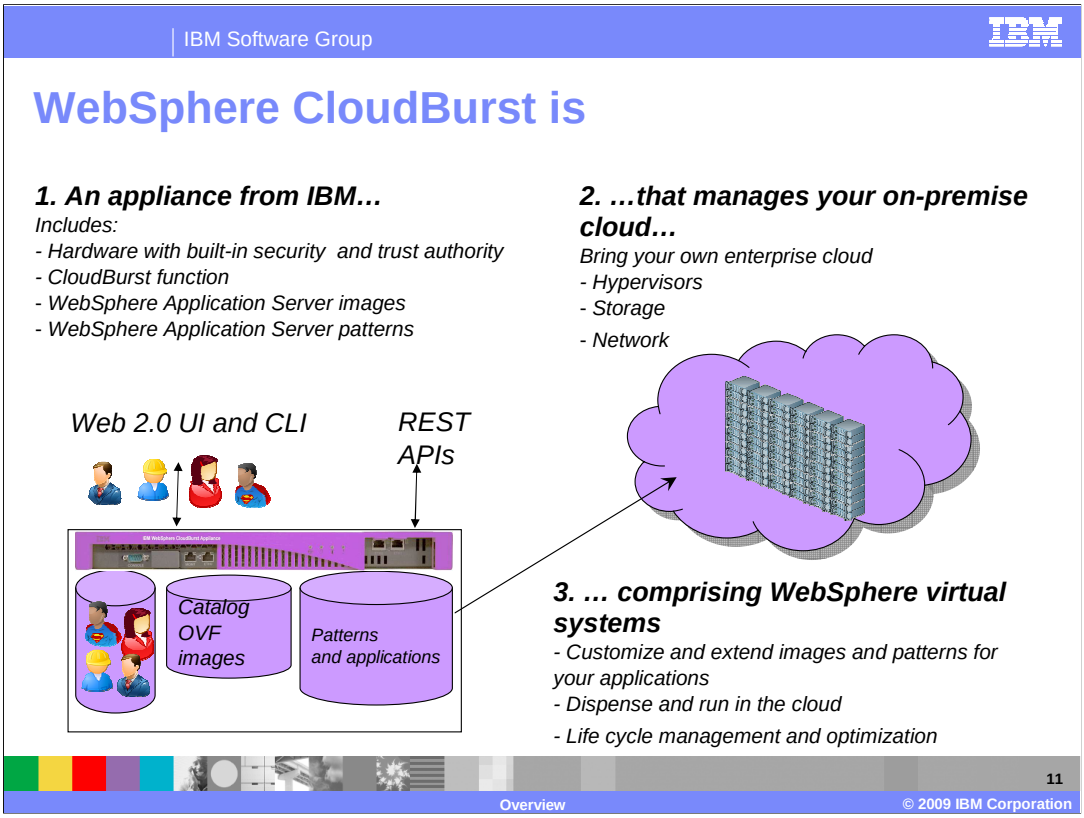

WebSphere Cloudburst is a hardware appliance that provisions hardened WebSphere topologies into a cloud of virtualized hardware. The appliance contains WebSphere Application Server virtual images and patterns.

Your own hardware is used to create the cloud on which WebSphere virtual images are dispensed and managed by the appliance.

The appliance is managed using a browser or using command line tools. External management applications can use the REST APIs exposed by the appliance to manage and view the appliance configuration.

The virtual images include the WebSphere Application Server Hypervisor Edition for V6.1 or V7.0. The Hypervisor Edition images include an operating system and an installed and configured instance of the application server. These images are used by the appliance to create virtual machines and deploy them onto the cloud. These virtual images are stored in the Open Virtualization format (OVF), which is an industry standard specification for packaging and distribution of virtual appliances containing one or more virtual machines.

The appliance can be thought of as a repository of virtual images similar to what CVS is for software programs. Security and trust authority are built in the appliance. Only authorized users and administrators have access to the images. The appliance also contains pre built topology patterns based on years of best practices of using WebSphere Application Server.

Section

WebSphere life cycle through WebSphere CloudBurst: How to deploy WebSphere into the cloud

This section covers the details of each of the WebSphere CloudBurst components, and shows the steps needed to deploy the WebSphere virtual images into the cloud.

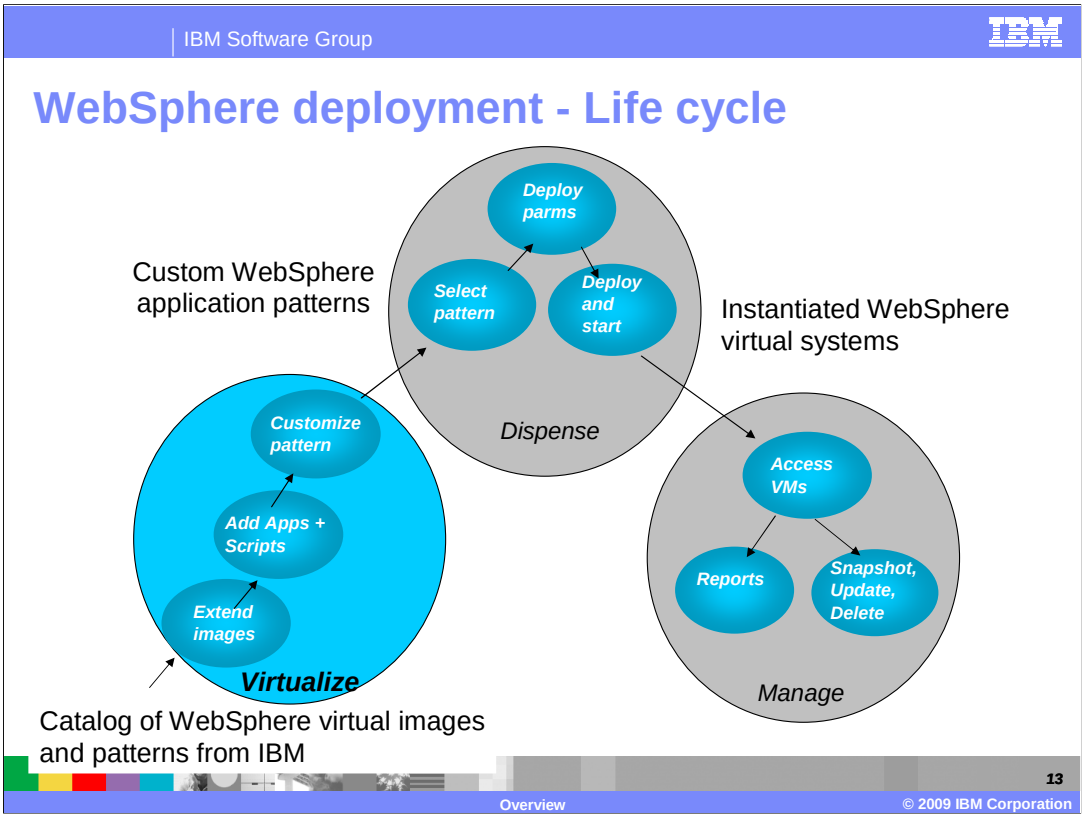

WebSphere CloudBurst manages the entire life cycle of WebSphere Application Server deployment. The three phases are virtualize, dispense and manage the WebSphere Application Server virtual images in the cloud.

Starting from a catalog of WebSphere Application Server Virtual images and pre-configured patterns provided by IBM, CloudBurst administrators can add script packages and customize the patterns to fit specific WebSphere Application Server topologies and configurations.

The custom or pre-built patterns can then be dispensed into a cloud of hypervisors. Here the specific deployment parameters are provided to the pattern, and then the pattern is deployed and started on the cloud.

The last phase of “manage” is where the CloudBurst administrator can access the virtual machines, create, update or delete the virtual machines, and create reports of the utilization of the virtual systems in the cloud.

The next few pages will focus on the “Virtualize” phase of the life cycle.

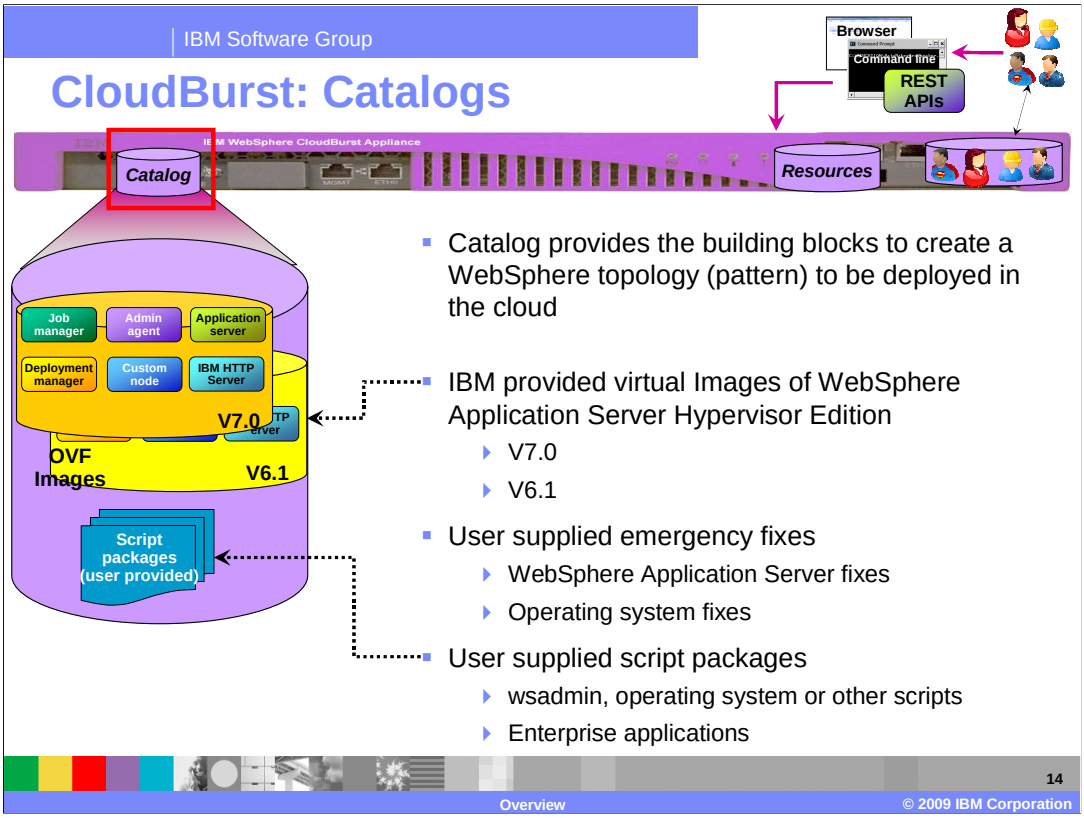

The WebSphere CloudBurst Appliance catalog provides the building blocks to create a WebSphere topology, or pattern, to be deployed in the cloud. All artifacts created by the administrator for deployment start from these building blocks in the catalog.

The catalog contains virtual images of WebSphere Application Server Hypervisor Edition V7.0 and V6.1. Each virtual image contains all profiles specific to that instance of the Application Server. The virtual images follow the OVF standard. Additionally, the virtual images can be extended to include your own software and brought back into the catalog. This is targeted for situations where you want to add some software or files that are to be common among all deployments.

Emergency fixes can be applied directly to a running virtual system. These fixes can be applied to the WebSphere Application Server and the operating system.

To make the deployment easier and more repeatable, the administrator can supply script packages as .zip file or .tar files. These script packages can contain wsadmin scripts or other operating system scripts, along with any enterprise applications or other artifacts. The administrator can then associate the scripts to a given topology for deployment. At deployment time, the script is run in the virtual machine on the hypervisor. This is how the administrator can customize the WebSphere Application Server configuration or operating systems on the virtual machine. Other customization examples include installing an enterprise application or configuring a JDBC™ connection.

Preloaded virtual images

- CloudBurst comes preloaded with three virtual images
  - WebSphere HV 6.1.0.23
  - WebSphere HV (Feature Packs) 6.1.0.23
  - WebSphere HV 7.0.0.3
- Default permissions are assigned to the preloaded images
  - Administrator has read/write authority
  - All other users have read-only authority unless granted additional authority by the administrator

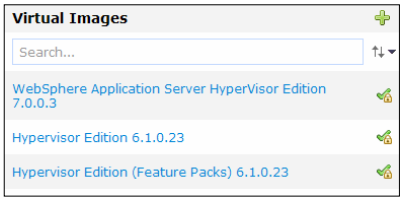

WebSphere CloudBurst comes preloaded with three virtual images. Those virtual images are “WebSphere HV 6.1.0.23”, “WebSphere HV 7.0.0.3” and “WebSphere HV (Feature Packs) 6.1.0.23. By default, the administrator is given read/write permission and all other users are given read-only permission unless the administrator grants additional permission.

Image updates will be provided by IBM. Image updates will include the operating system and WebSphere maintenance streams. The benefit of using the preloaded images is that you do not need to install or configure the operating system or WebSphere Application Server.

Parts of catalog (Individual virtual images)

Catalog parts (individual virtual images)	WebSphere Application Server Hypervisor Edition V7.0	WebSphere Application Server Hypervisor Edition V6.1
Single server	Y	Y
Deployment manager	Y	Y
Custom node	Y	Y
Administrative agent	Y	-
Job manager	Y	-
IBM HTTP Server	Y	Y

The table shows the virtual images included in the each of the releases of the WebSphere Application Server Hypervisor Edition V7.0 and V6.1.

The difference between V6.1 and V7.0 are the new profiles of administrative agent and job manager in V7.0 that do not exist for V6.1.

Script packages

- With the script package, you can:
  - Upload your script package
  - Create environment variables
  - Specify an executable file
  - Supply arguments to the executable file
  - Grant or remove permissions

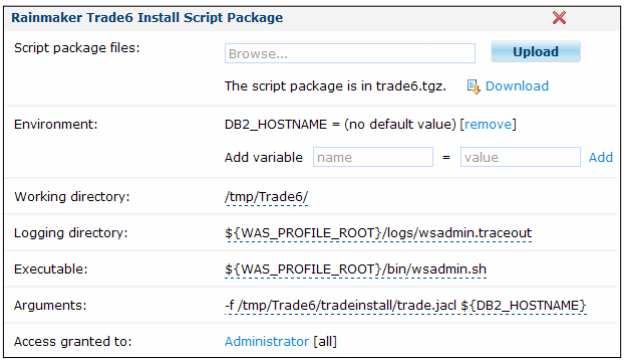

Creating a script package requires you to upload the script package, specify a working directory, an executable file, and a set of arguments.

“Script package files” represents a .zip or .tgz file that contains your scripts and any applications that you want installed. This file is uploaded from your local system to the CloudBurst catalog. It is the copy that resides in the catalog that is used during the pattern deployment process.

During pattern deployment your script package is sent over to the hypervisor along with the virtual image and extracted. The “Working directory” tells CloudBurst where to extract the contents of the script package.

The “Executable” field contains the command that you want CloudBurst to invoke. CloudBurst will invoke this command using root authority.

The “Arguments” field includes any arguments you want passed to the “Executable” when it is called.

The “Logging directory” field lets CloudBurst know where the logs or trace files are to be located after the command runs. CloudBurst uses this information during troubleshooting. You can download the logs or trace files directly from within the CloudBurst administrative console.

There are cases where you do not know the values of something or that it is specific to each deployment in which case you want certain information to be filled out at deployment time. In such cases you can define a variable under “Environment” and not define a value for the variable.

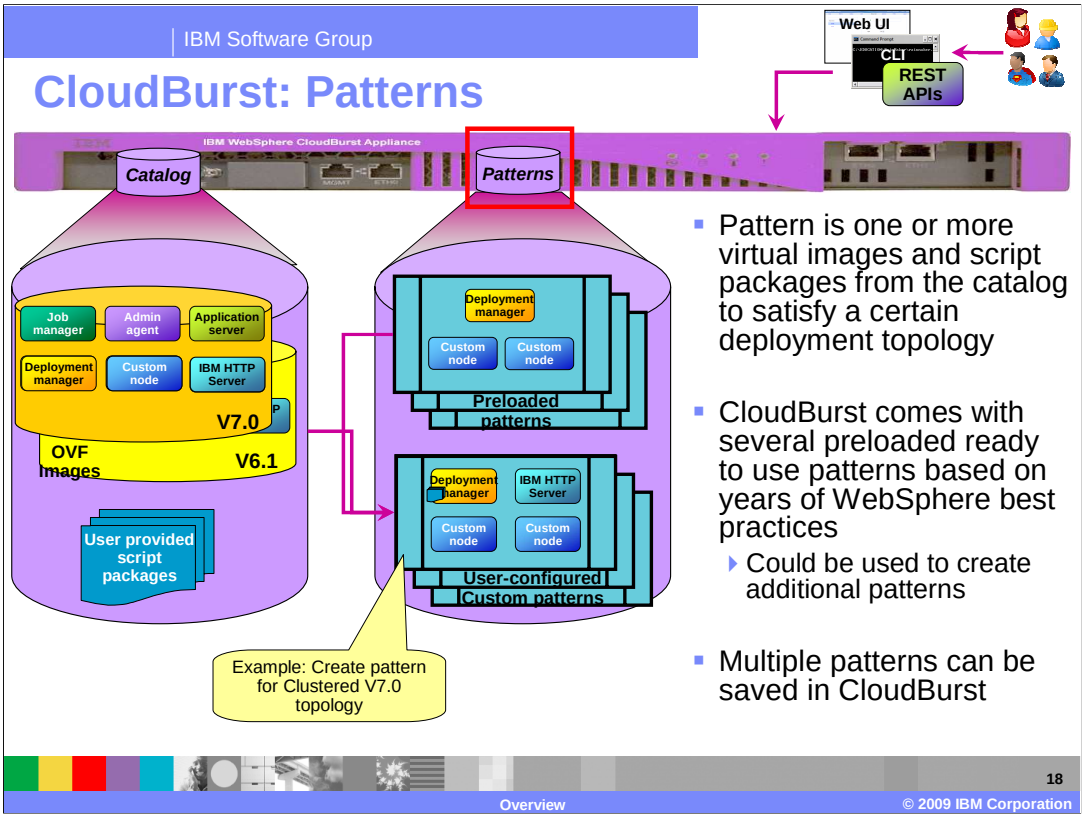

Using the catalog of virtual images of WebSphere Application Server Hypervisor Edition and script packages, you can create patterns or topologies that can then be deployed as virtual machines on the cloud. Patterns are like templates that can be used to deploy on the cloud.

The pattern can be a single server topology where a single server virtual image from the catalog for a given version is used to create a pattern. Or, another example can be a pattern for a Network Deployment cell, where you could select the Deployment manager, custom nodes and IBM HTTP Server and add them to the pattern, which can then be deployed.

[1] The CloudBurst user or administrator starts with an empty pattern or can start with a copy of an existing pattern.

[2] Using the console, they can drag the virtual image parts from the catalog to build the pattern.

Each virtual image in the pattern is deployed as its own virtual machine into the cloud. In the example shown on the slide, four virtual machines are deployed in the cloud for the clustered V7.0 pattern.

[3] CloudBurst contains several pre-configured read-only patterns based on years of WebSphere best practices. These pre-configured patterns can be used as-is for deployment, or used to create copies that can then be customized for a specific topology.

The CloudBurst administrator can create several patterns with different topology and configuration that can then be deployed into the cloud with a push of a button.

Preloaded patterns

- CloudBurst comes preloaded with a set of “best practices” patterns
- The following is a subset of the preloaded patterns
  - WebSphere single server
    - 1 Stand-alone server
    - Total of 1virtual machine
  - WebSphere cluster
    - 1 Deployment manager
    - 2 Custom nodes
    - 1 IBM HTTP server
    - Total of 4 virtual machines
  - WebSphere cluster (large topology)
    - 1 Deployment manager
    - 10 Custom nodes
    - 4 IBM HTTP servers
    - Total of 15 virtual machines
- Patterns can be used as is or used as a starting point when creating your own custom patterns

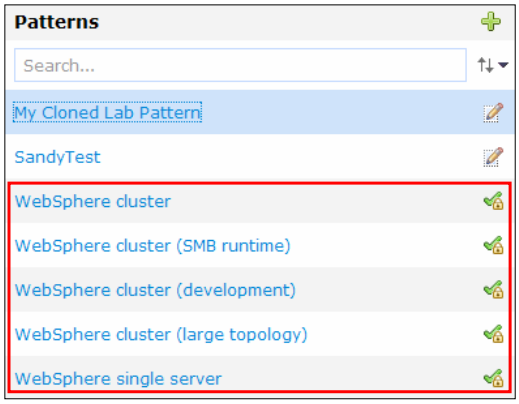

WebSphere CloudBurst comes preloaded with a set of “best practices” patterns, which you should use when creating your own custom patterns. By starting with one of CloudBurst’s existing patterns you are starting with a time tested configuration.

The preloaded patterns range in complexity from a simple single server topology to a highly available cluster topology.

Drag-and-drop pattern assembly

- Assemble patterns by dragging virtual image parts and script packages from the palette on the left and dropping them on to the canvas on the right

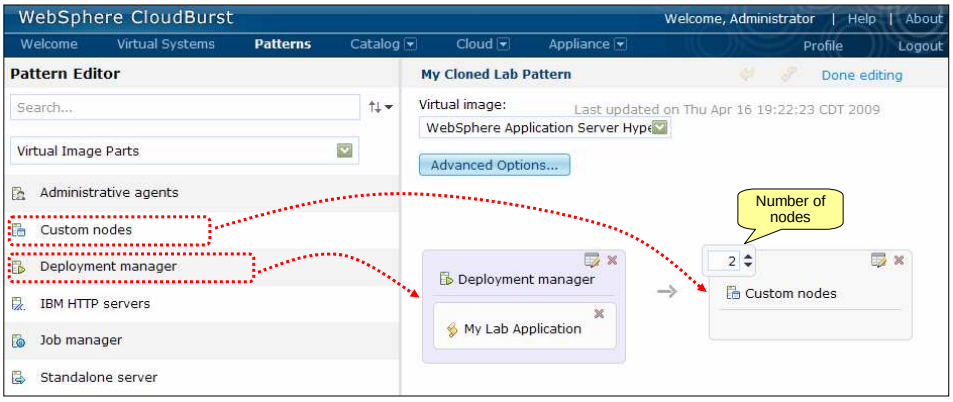

If this is a new pattern you will see a blank canvas on the right and a palette of parts on the left. The palette on the left contains your virtual image parts and script packages which you will use to assemble your pattern. You create a pattern by dragging parts from the palette and dropping them onto the canvas. In order to add multiple parts at the same configuration point, you increment the number located in the upper left corner of the part when available.

The example in this slide shows three nodes - a deployment manager node managing two custom nodes.

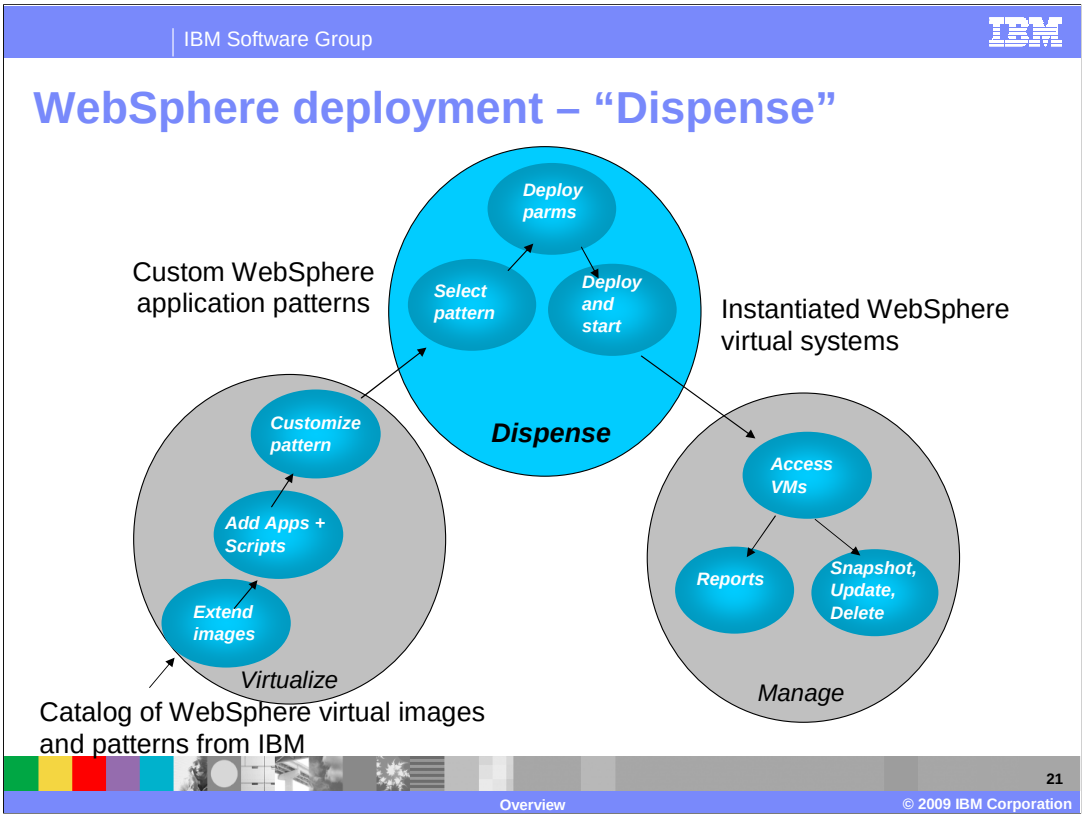

This sub-section focuses on the “Dispense” phase of the life cycle. In this phase, once the pattern is selected and customized, the deployment parameters are entered to create the virtual systems. These virtual systems are then deployed on the hypervisor cloud.

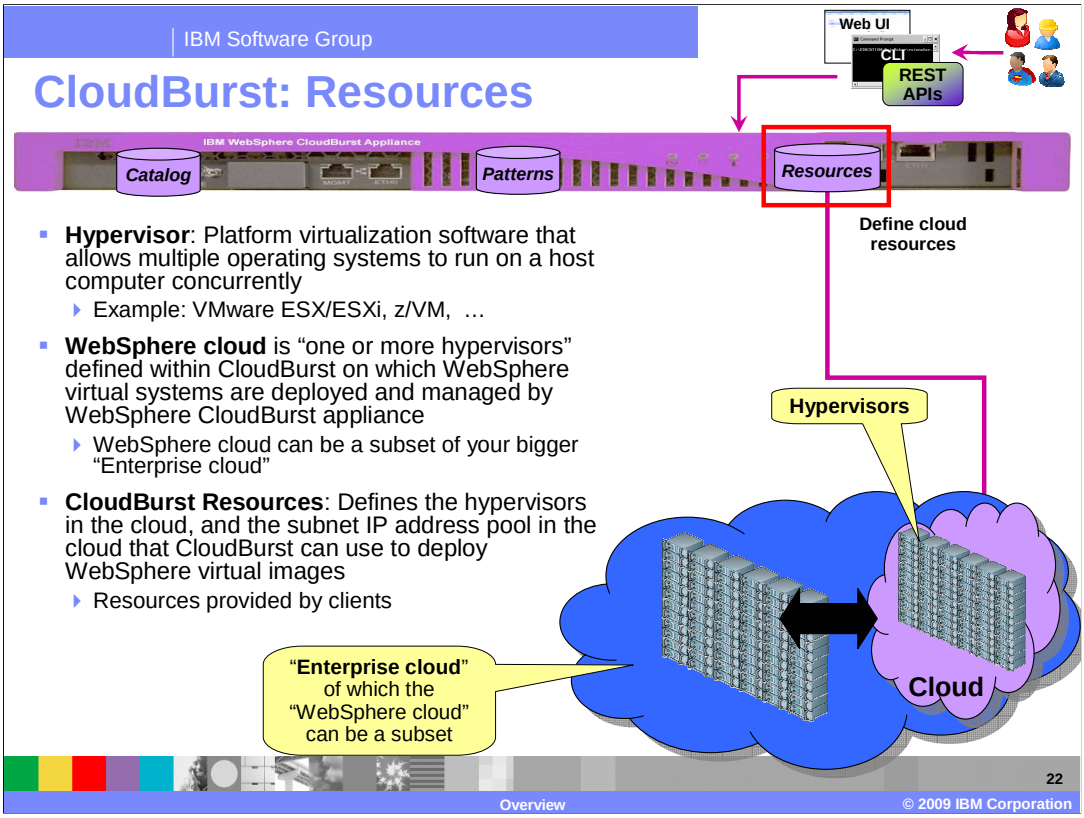

A **hypervisor**, also called *virtual machine monitor* (*VMM*), is software that allows multiple operating systems to run concurrently on a host computer.

These hypervisors form the purple "WebSphere cloud", as shown on the slide, on which WebSphere CloudBurst deploys the virtual machines of WebSphere Hypervisor Edition.

The WebSphere cloud can be a subset of the enterprise cloud, and be managed by an enterprise service management and automation tools within the context of the large enterprise IT resources.

You define the cloud resources within the CloudBurst appliance. The cloud resources are the hypervisors in the cloud, and the subnet IP address pool in the cloud that CloudBurst can use to deploy WebSphere virtual images.

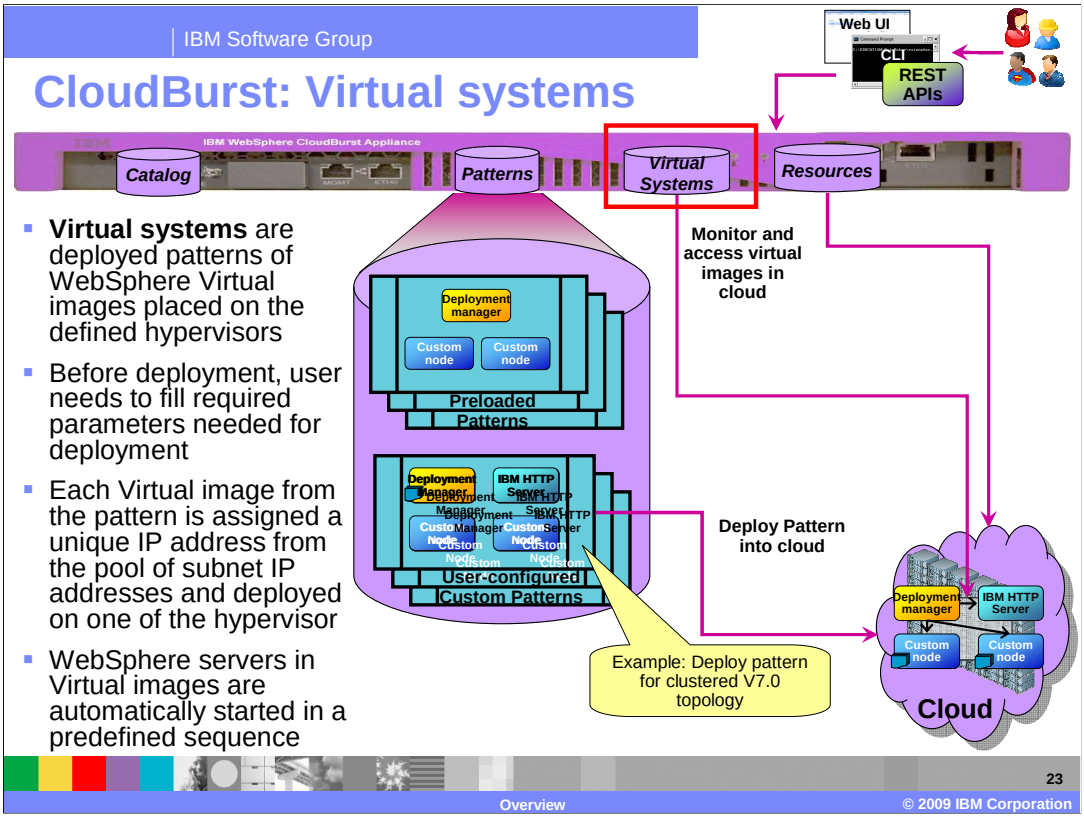

Patterns are the starting point of deployment to the cloud. The deployed patterns are called virtual systems.

Before deployment, you need to fill required parameters needed for deployment, like operating system root password, WebSphere Administrator password, Cell, node server name, and so on. These parameters are used to activate each virtual image in the pattern and create its virtual machine.

Each virtual image is assigned a unique IP address from the pool of subnet IP addresses.

WebSphere servers in virtual images are automatically started in a predefined sequence. Details of the start sequence are provided in the detailed Catalog and Pattern presentation.

[1] Once the virtual system is deployed on the cloud, WebSphere Cloudburst does not play any role in the running of the WebSphere servers in the cloud. Even if the appliance was shut down, the virtual system of WebSphere Application server will continue to run. The Cloudburst Appliance's role is deployment and monitoring, but not affecting or coordinating the actual running of the application servers. The important thing to note is that the Cloudburst appliance does not play any role in the high availability of the deployed WebSphere Application Server in the cloud.

Configure a pattern deployment

- Configuration updates are specific to this deployment (virtual system)
- Root and Virtuser passwords are required for each deployment

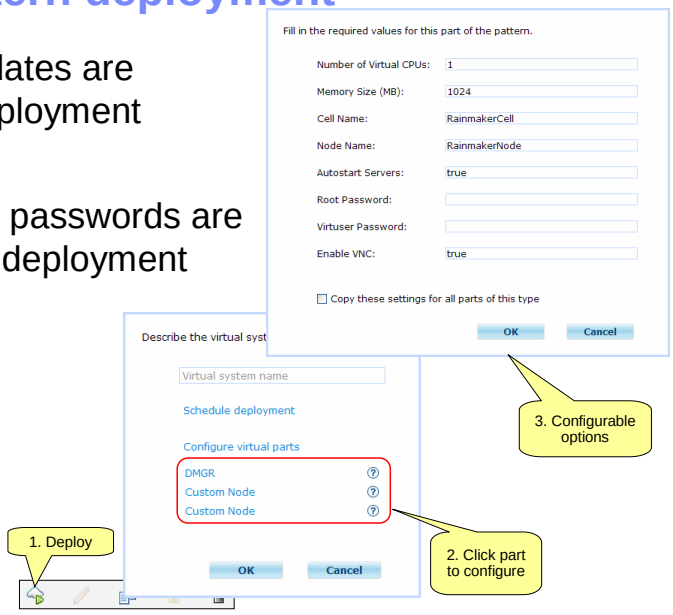

When you deploy a pattern you are given a chance to configure the pattern's parts and script package variables. CloudBurst requires only that the "Root Password" and the "Virtuser Password" be filled in by you; all other fields are given a default value if you do not override them. Script package variable requirements are not dictated by CloudBurst but rather the requirements of the specific script package.

Schedule a pattern deployment

- Deploy immediately
- Deploy at some later date and time
- Run forever or until some later date and time

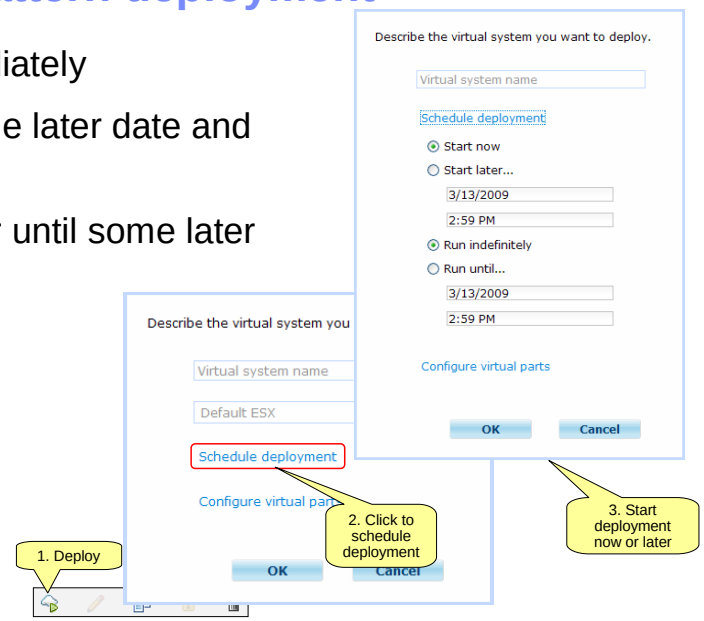

When you deploy a pattern you have the option to deploy immediately or schedule a time for deployment. You also have the option to specify a length for the deployment. You can run the deployment indefinitely or until some future time.

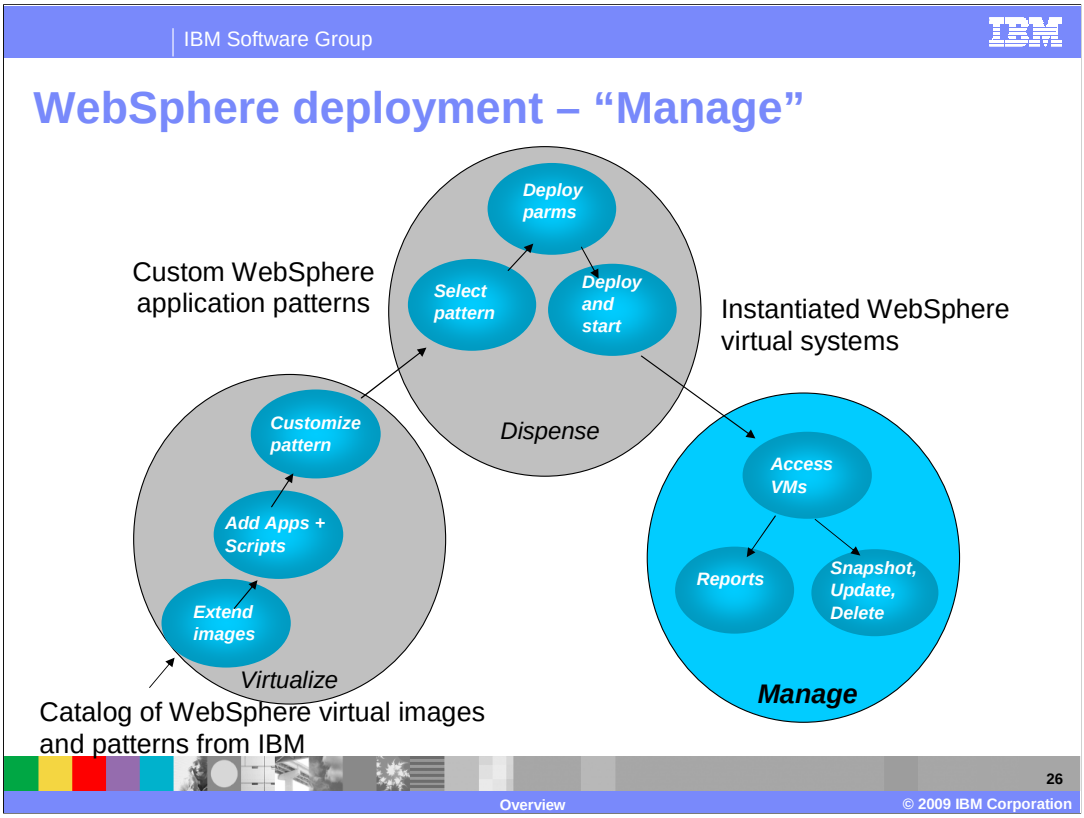

This sub-section focuses on the “Manage” phase of the life cycle. In this phase, the administrator manages the virtual systems, creates reports as needed, is able to update and delete the virtual systems, and takes snap shots of the virtual systems.

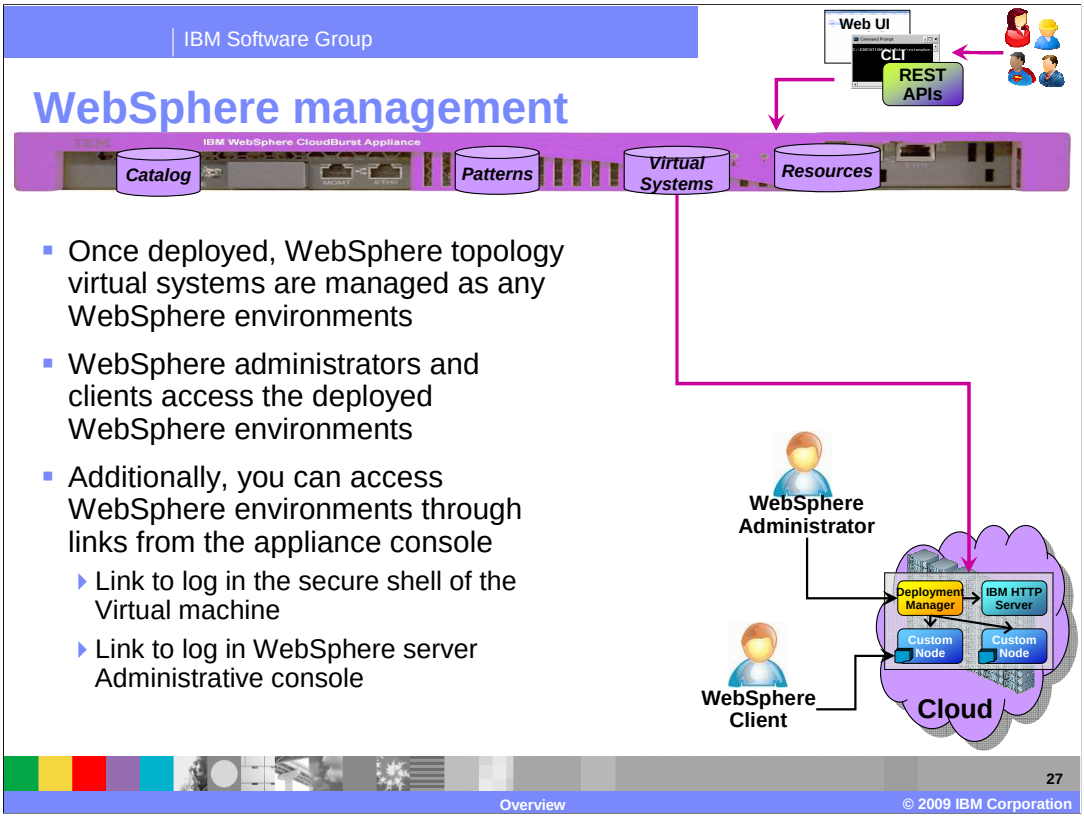

Once the WebSphere virtual systems are deployed, they are managed as any other WebSphere deployments. Each virtual machine has an IP address. You can access the servers knowing the IP address of the virtual machine. You have access to the Administration console or can use wsadmin to connect to the WebSphere server to configure the server or install any applications. Most of the configuration and installing the application is done through the scripts that are added to the pattern before the deployment. However, you can still continue to perform any post deployment configuration changes to the WebSphere servers.

Virtual system administration

- View virtual system and WebSphere Application Server metrics
- Access WebSphere Application Server administrative console
- Access the operating system using SSH or VNC

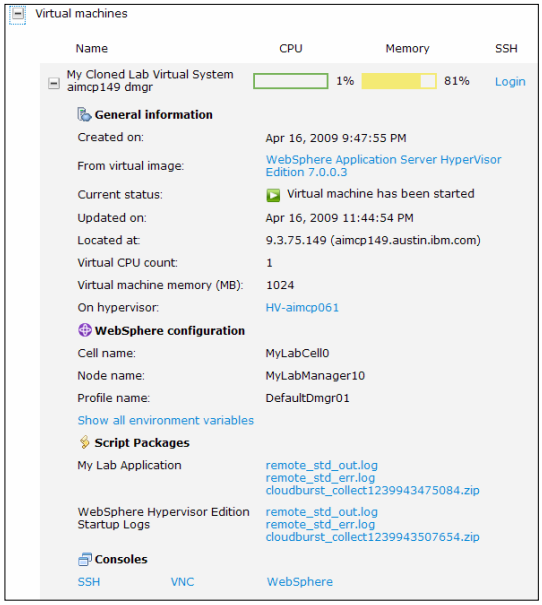

Expanding the “Virtual machines” section will present you with metrics for the virtual machines that make up the virtual system. It contains information such as what hypervisor the virtual machine is installed on, what IP address was assigned, and the cell and node name of the WebSphere Application Server environment.

CloudBurst assumes the responsibility of assigning IP addresses to the virtual machines and placement of those virtual machines in the cloud. Since this is the case, CloudBurst provides you with a WebSphere Application Server administrative console link, an SSH link, and a VNC link into the operating system. These links are found under the “Consoles” section. Once you open the WebSphere Application Server administrative console you can bookmark it for future reference.

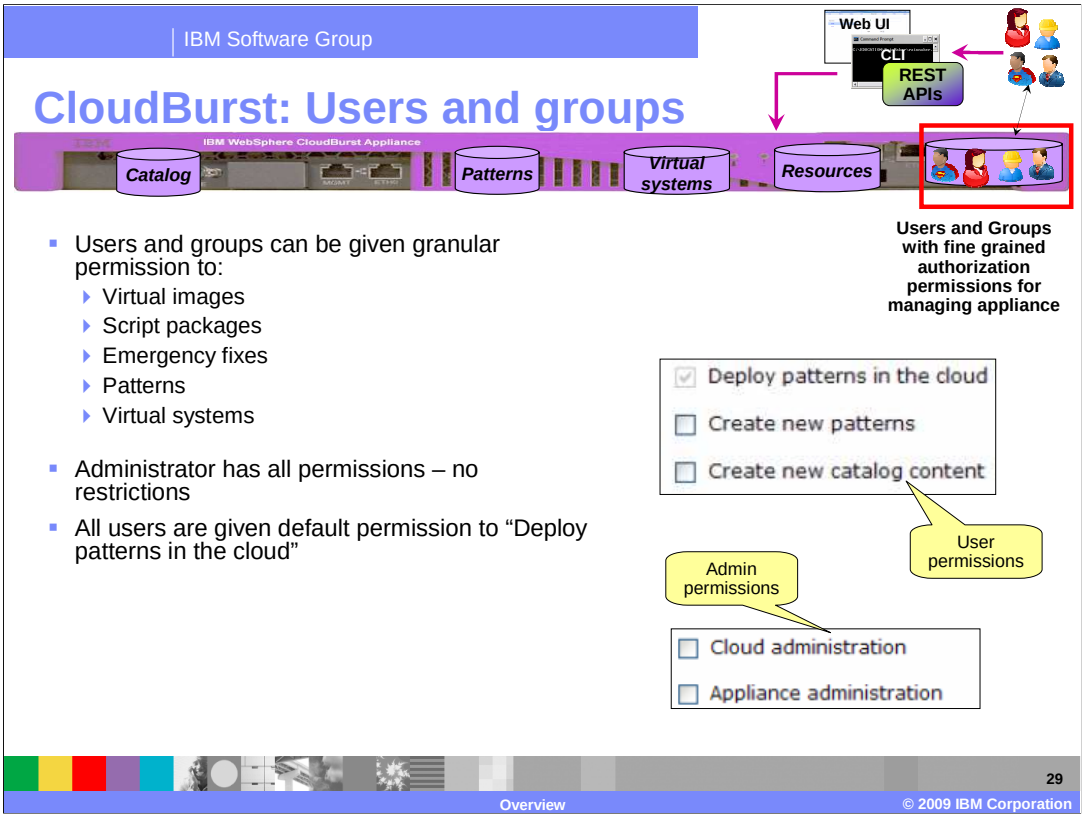

Administrators of Cloudburst can set up fine grain authorization for different administration tasks. These tasks involve managing the cloud resources, managing the appliance, and creating patterns and virtual systems.

Managing the cloud includes managing the cloud resources, catalogs of virtual images of WebSphere Hypervisor Edition, and script packages. Managing the appliance includes setting security for users and groups, maintaining the appliance Ethernet interfaces, and other settings. Creating cloud patterns and virtual systems allows you to create patterns and deploy them on the cloud.

All administrators have permission to perform all administration tasks.

New users by default have only the "Deploy patterns in the cloud" permission. The administrator can give additional permissions to manage the cloud and the appliance.

Virtual images, script packages, emergency fixes, patterns, and virtual system access are controlled based on ownership. Users have permissions to manage only patterns and virtual systems that they created; they do not have permission to manage or view patterns and virtual systems that they do not own. This allows different users to manage their artifacts in isolation from other users.

An owning user can give other users permission to access their patterns and virtual systems, providing the flexibility of multiple users being able to manage a given pattern or virtual system, if needed.

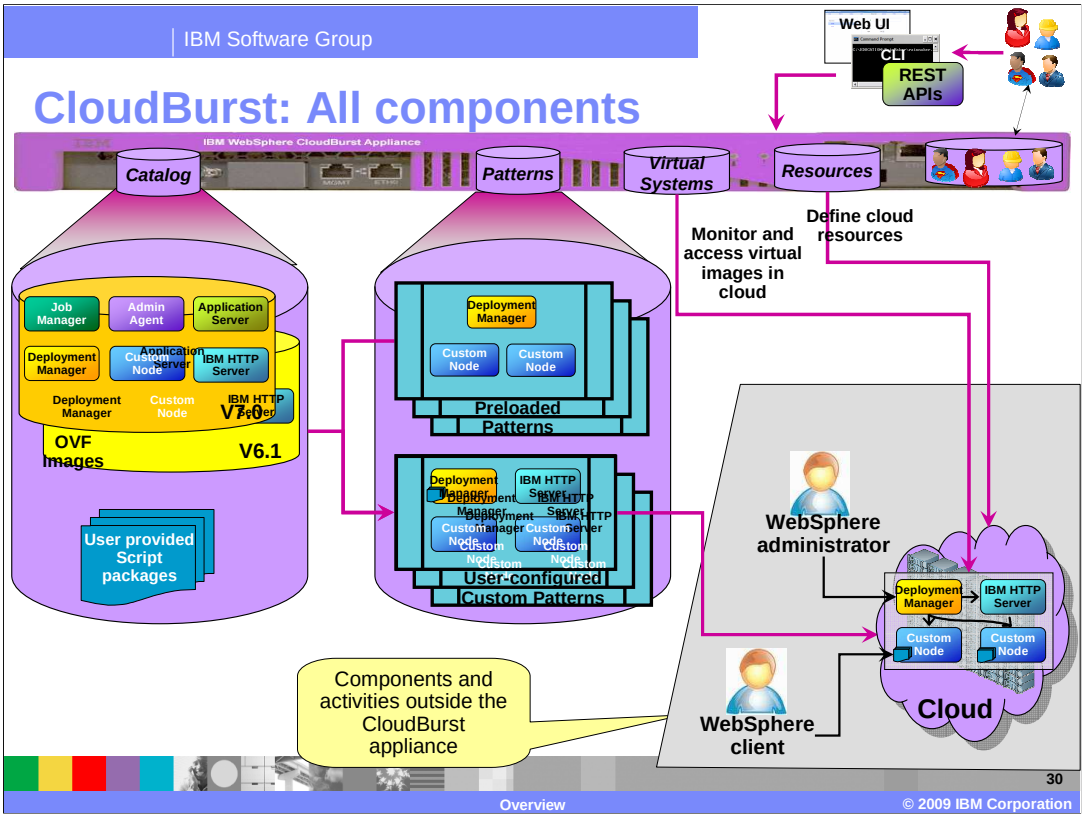

All the major components of the CloudBurst appliance that have been used in this presentation are shown on this slide.

Section

CloudBurst appliance management

This section covers the CloudBurst appliance management.

CloudBurst appliance management

- Administrative console
- Command Line Interface

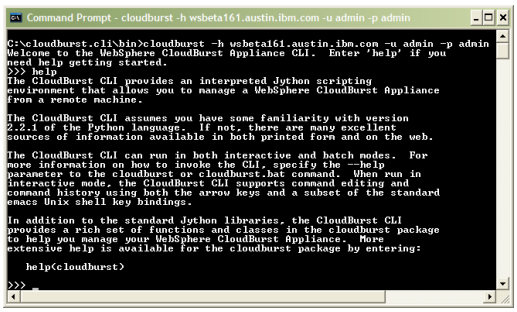

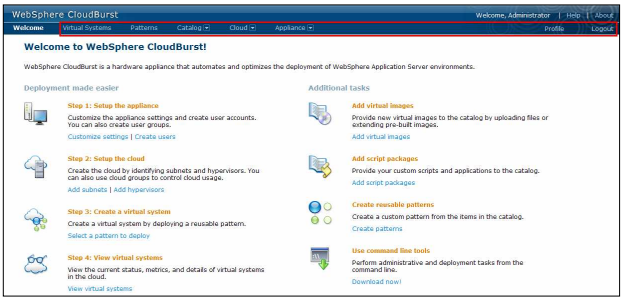

The main Web-based administration console tabs are shown here. These tabs allow you to manage and configure different aspects of the appliance and the environment. Other presentations will go into the details of each aspect.

Other CloudBurst functions

- Securely backup and restore your CloudBurst appliance to a remote location
- Extending existing virtual images to include custom software
- Create and restore snapshots of running virtual systems

The backup and restore functions let you backup a specific state of the appliance and if needed, later, restore the appliance to the saved state. The backup data is saved at a user-specified remote location, and is encrypted so as to allow only the trusted parties to use the saved data during restore.

CloudBurst gives you the capability to extend an existing virtual image, make changes, and then capture those changes back into the catalog.

Creating a snapshot is the process of taking copy of the entire state of a virtual system. Restoring from the snapshot takes you back to the state defined by the snapshot.

Other presentations provide full details on these functions.

WebSphere CloudBurst: Data center integration

- Points of integration into broader data center management solutions:
  - automation / resource provisioning
  - license management
  - monitoring
  - user management
  - metering / billing

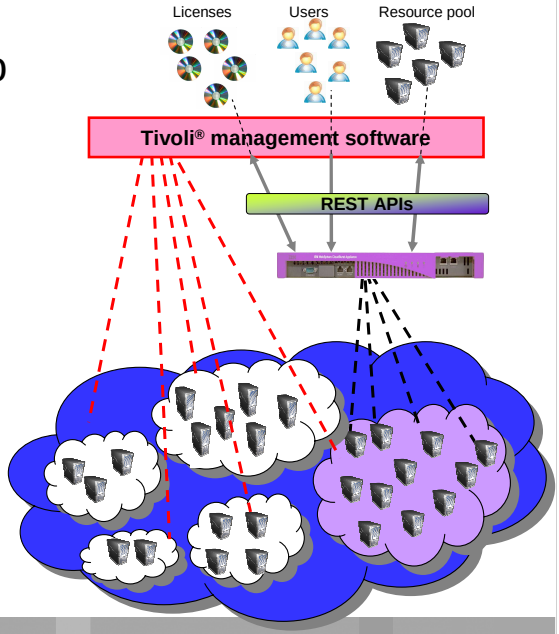

This page shows the role of WebSphere CloudBurst within the enterprise, and being managed by a broader enterprise service management and automation tool like Tivoli management software.

Using CloudBurst's REST based APIs, the enterprise service management tools can manage and configure the entire appliance, similar to what you would do using the browser-based console or command line tool. The enterprise service management tools can define the users, the resources available for the appliance, manage licenses, and essentially drive the appliance to be managed as part of the data center resources.

External enterprise service management tools can dedicate a subset of the resources to the appliance, and on demand, either add or remove resources available to the appliance. This is denoted by the blue "enterprise cloud" and the purple "WebSphere cloud" subset given to the appliance to deploy the WebSphere virtual images.

Section

Summary

This section is the summary.

Summary

- Life cycle support for
  - Virtualize
  - Dispense
  - Manage
- Delivered as a secure appliance
- Integrates into your enterprise

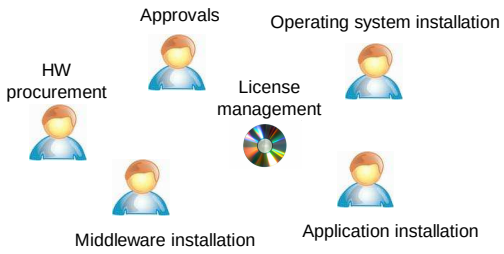

OR

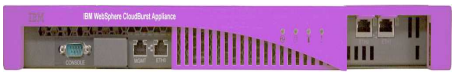

Business value: WebSphere CloudBurst increases efficiency, cost-effectiveness, and usability of WebSphere topologies in a data center by using cloud computing principles

WebSphere CloudBurst lowers cost through higher hardware utilization, less power consumption, and more efficient license management. It provides lower 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.

WebSphere CloudBurst ships with environment “patterns,” which codify 10 years of WebSphere management best practices. WebSphere environments can be customized and captured once, then are able to be dispensed at the push of a button. It increases agility. Rapid setup and teardown of WebSphere Application Server environments means less time spent managing WebSphere environments. And it integrates well into your enterprise.

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.

References

- CloudBurst Information Center
  - http://publib.boulder.ibm.com/infocenter/wscloudb/v1r0/index.jsp

The information center contains much useful information.

Feedback

Your feedback is valuable

You can help improve the quality of IBM Education Assistant content to better meet your needs by providing feedback.

- Did you find this module useful?
- Did it help you solve a problem or answer a question?
- Do you have suggestions for improvements?

Click to send e-mail feedback:

mailto:iea@us.ibm.com?subject=Feedback_about_CloudBurst_Vision_Overview.ppt

This module is also available in PDF format at: ../CloudBurst_Vision_Overview.pdf

You can help improve the quality of IBM Education Assistant content by providing feedback.

Trademarks, copyrights, and disclaimers

IBM, the IBM logo, ibm.com, and the following terms are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both:

Dynamic Infrastructure Tivoli WebSphere z/VM

If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or ™), these symbols indicate U.S. 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 other IBM trademarks is available on the Web at "Copyright and trademark information" at http://www.ibm.com/legal/copytrade.shtml

JDBC, and all Java-based trademarks and logos are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Linux is a registered 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.

Product data has been reviewed for accuracy as of the date of initial publication. Product data is subject to change without notice. This document could include technical inaccuracies or typographical errors. IBM may make improvements or changes in the products or programs described herein at any time without notice. Any statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only. References in this document to IBM products, programs, or services does not imply that IBM intends to make such products, programs or services available in all countries in which IBM operates or does business. Any reference to an IBM Program Product in this document is not intended to state or imply that only that program product may be used. Any functionally equivalent program, that does not infringe IBM's intellectual property rights, may be used instead.

THE INFORMATION PROVIDED IN THIS DOCUMENT IS DISTRIBUTED "AS IS" WITHOUT ANY WARRANTY, EITHER EXPRESS OR IMPLIED. IBM EXPRESSLY DISCLAIMS ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NONINFRINGEMENT. IBM shall have no responsibility to update this information. IBM products are warranted, if at all, according to the terms and conditions of the agreements (for example, IBM Customer Agreement, Statement of Limited Warranty, International Program License Agreement, etc.) under which they are provided. 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 in connection with this publication and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products.

IBM makes no representations or warranties, express or implied, regarding non-IBM products and services.

The provision of the information contained herein is not intended to, and does not, grant any right or license under any IBM patents or copyrights. Inquiries regarding patent or copyright licenses should be made, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. All customer examples described are presented as illustrations of how those customers have used IBM products and the results they may have achieved. The actual throughput or performance that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput or performance improvements equivalent to the ratios stated here.