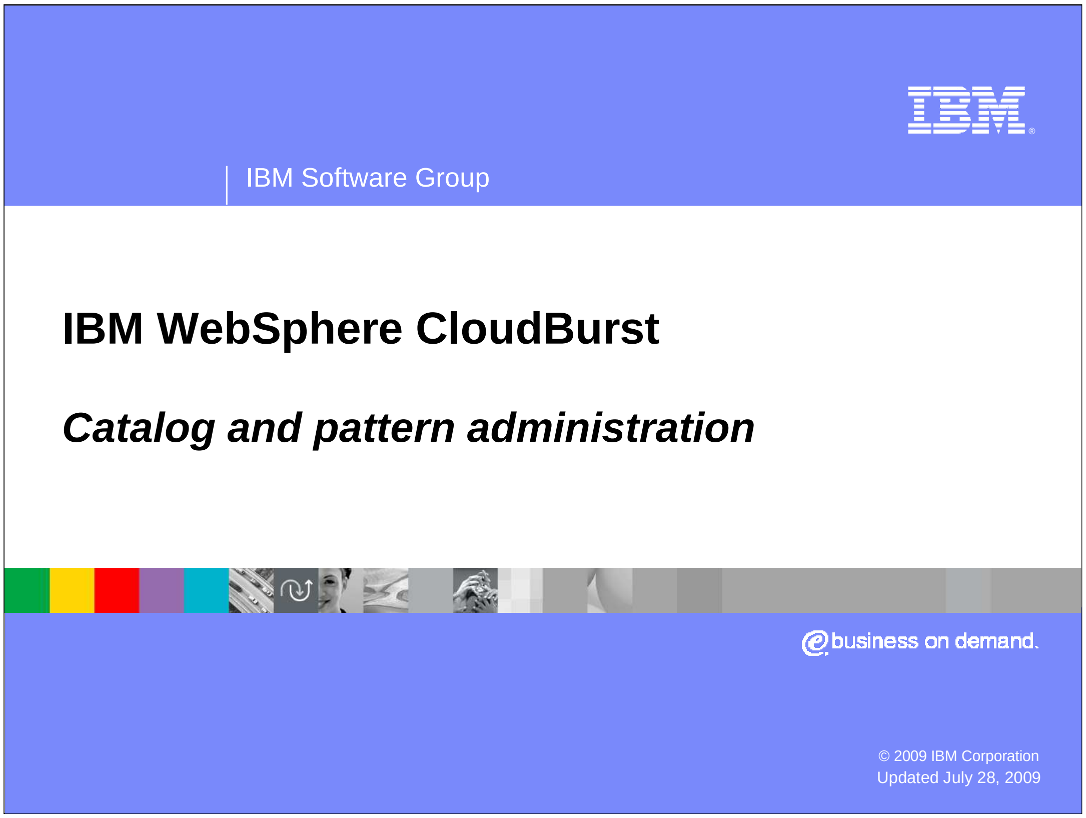

IBM Software Group

IBM WebSphere CloudBurst

Catalog and pattern administration

@business on demand.

Updated July 28, 2009

This presentation covers CloudBurst’s catalog and pattern features.

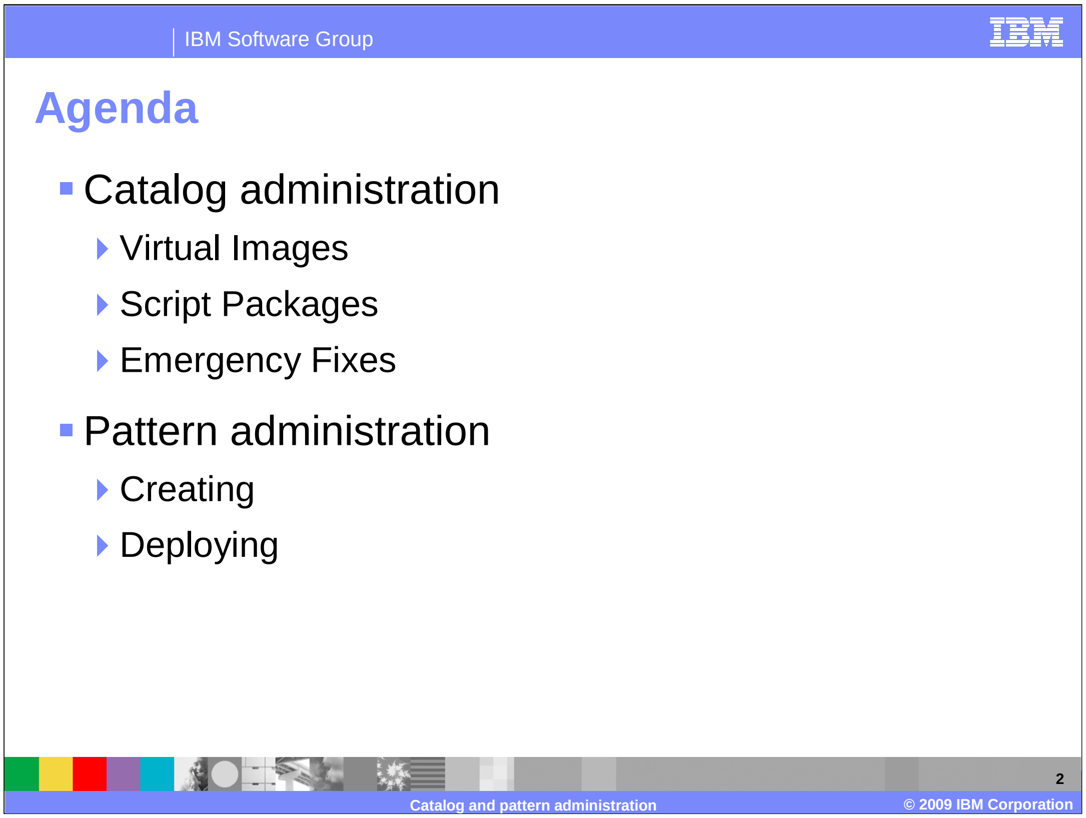

Agenda

- Catalog administration
  - Virtual Images
  - Script Packages
  - Emergency Fixes
- Pattern administration
  - Creating
  - Deploying

This presentation covers the process of creating and working with virtual images and script packages. It also goes over the steps involved in creating and deploying patterns into the CloudBurst cloud.

In this presentation you will cover catalog administration, including virtual images, script packages, and emergency fixes. Once you have covered the components of the catalog you will learn about patterns and how the individual components of the catalog make up a pattern.

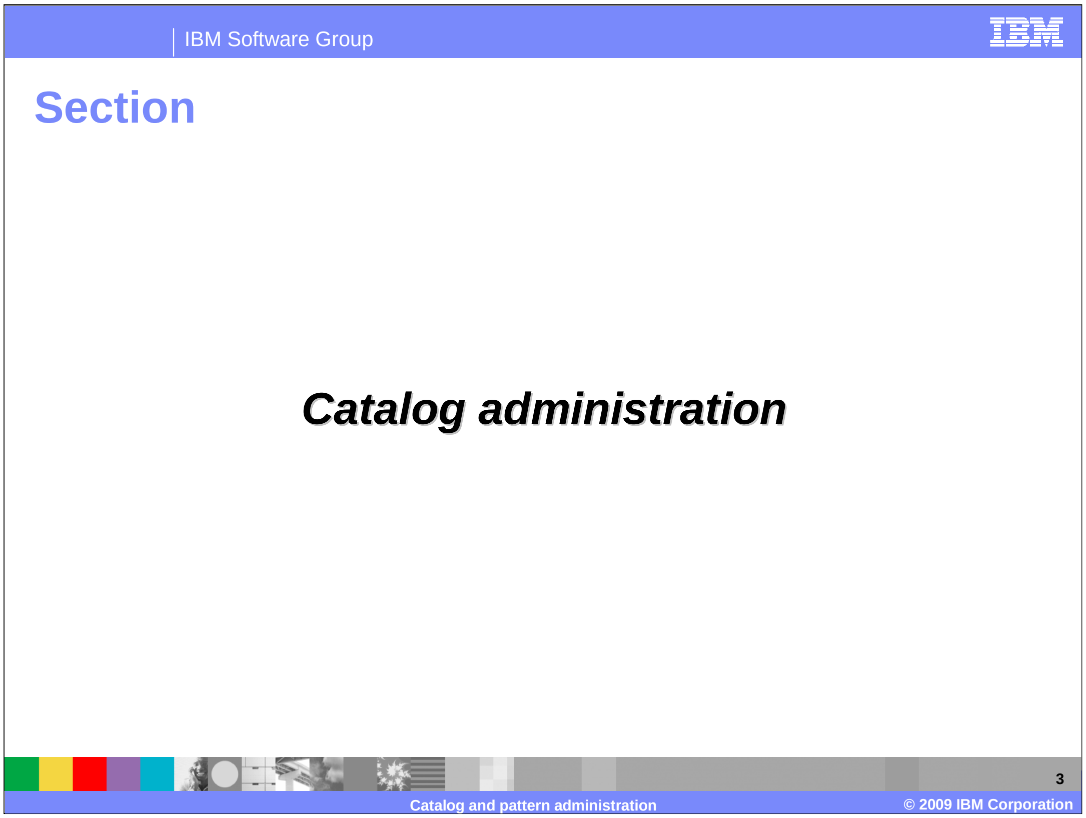

Section

Catalog administration

This section covers catalog administration.

Components administered by “Catalog”

- Virtual images
- Script packages
- Emergency fixes

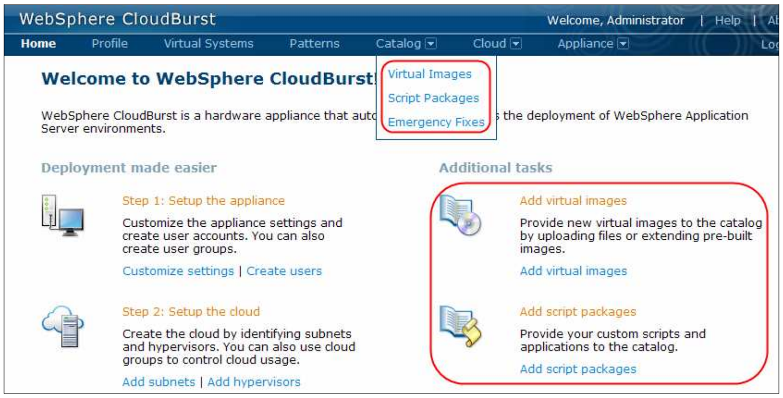

Welcome to WebSphere CloudBurst!

WebSphere CloudBurst is a hardware appliance that auto… the deployment of WebSphere Application Server environments.

Deployment made easier

Step 1: Setup the appliance

Customize the appliance settings and create user accounts. You can also create user groups.

Customize settings | Create users

Step 2: Setup the cloud

Create the cloud by identifying subnets and hypervisors. You can also use cloud groups to control cloud usage.

Add subnets | Add hypervisors

Additional tasks

Add virtual images

Provide new virtual images to the catalog by uploading files or extending pre-built images.

Add virtual images

Add script packages

Provide your custom scripts and applications to the catalog.

Add script packages

There are two ways to create resources in CloudBurst; you can click the links in the main welcome page located under the “Additional tasks” section, or click the “Catalog” drop-down located on the top option banner.

This section covers the creation and setup of virtual images, script packages, and emergency fixes, all of which will reside in CloudBurst’s catalog.

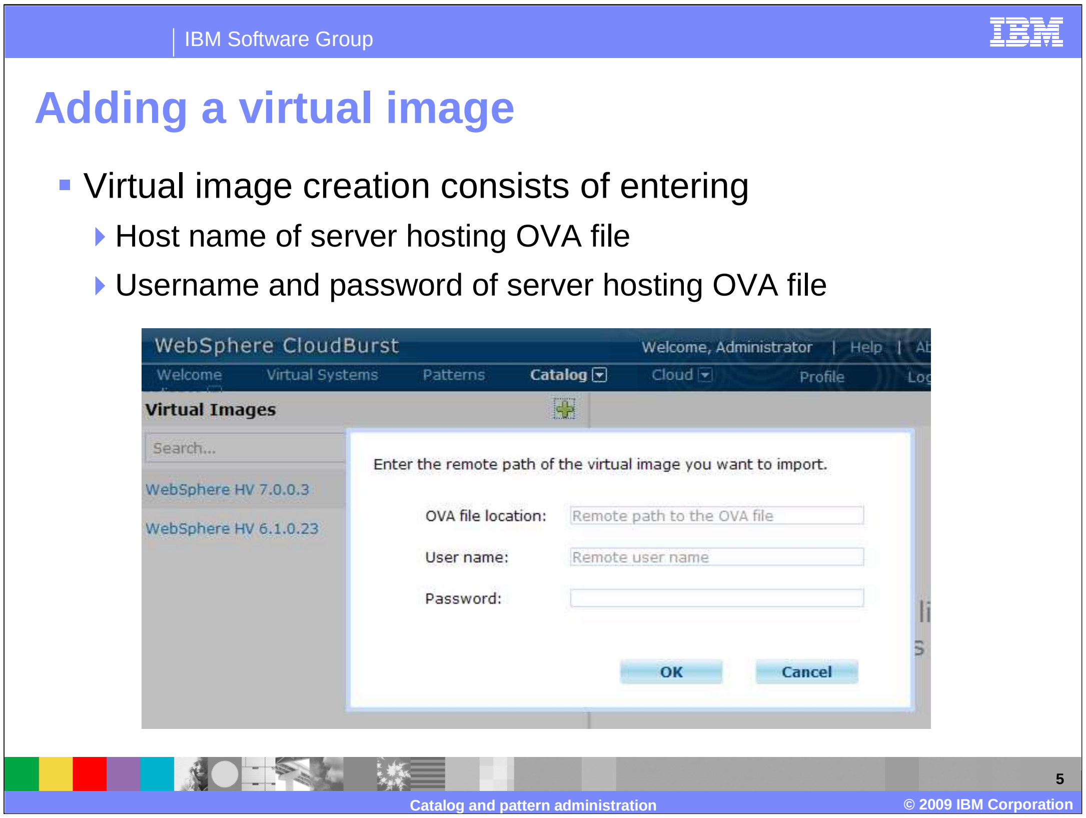

Adding a virtual image

- Virtual image creation consists of entering
  - Host name of server hosting OVA file
  - Username and password of server hosting OVA file

Enter the remote path of the virtual image you want to import.

OVA file location: Remote path to the OVA file

User name: Remote user name

Password:

Adding an existing Open Virtualization Appliance (OVA) to CloudBurst's catalog is as simple as specifying a location, user name, and password of the system hosting the OVA file. Once this step is complete, it will take some time to upload the virtual image from the source location into CloudBurst's catalog.

Currently the "OVA file location" only supports HTTP.

IBM Software Group

Virtual image attributes

- Contains information specific to this virtual image such as
  - Which parts make up with virtual image
  - Which patterns and virtual systems are using parts from this virtual image
  - Who has been granted access to this virtual image
- Extend and capture a virtual image

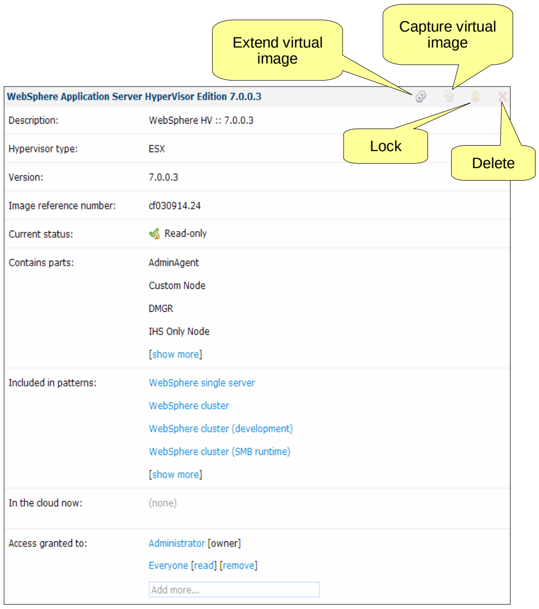

WebSphere Application Server HyperVisor Edition 7.0.0.3

Description:	WebSphere HV :: 7.0.0.3
Hypervisor type:	ESX
Version:	7.0.0.3
Image reference number:	cf030914.24
Current status:	Read-only
Contains parts:	AdminAgent Custom Node DMGR IHS Only Node [show more]
Included in patterns:	WebSphere single server WebSphere cluster WebSphere cluster (development) WebSphere cluster (SMB runtime) [show more]
In the cloud now:	(none)
Access granted to:	Administrator [owner] Everyone [read] [remove] Add more...

6

Catalog and pattern administration

The virtual images attributes page contains information about what parts make up the virtual image, what patterns and virtual systems were built using this virtual image, and who is granted permission to view, modify, and update this virtual image.

“Image reference number” is the version of the virtual image that resides in the CloudBurst catalog. This number is maintained by CloudBurst.

CloudBurst views the virtual image as a series of parts that are defined in the “Contains parts” attribute. These parts can be used to build up a pattern.

“Included in patterns” and “In the cloud now” show you the patterns and virtual systems that are currently using this virtual image. As long as a pattern or virtual system is using this virtual image you will not be allowed to delete this virtual image.

“Access granted to” shows who has authority to view and modify this virtual image. The user who creates the virtual image has exclusive rights to it unless they grant others access to the virtual image.

In the top right corner, you see “Lock” and “Delete” icons. The “Delete” icon lets you delete a virtual image from the catalog. This will only be allowed if there are no patterns or virtual systems currently using it. The “Lock” icon locks the virtual image from any further changes.

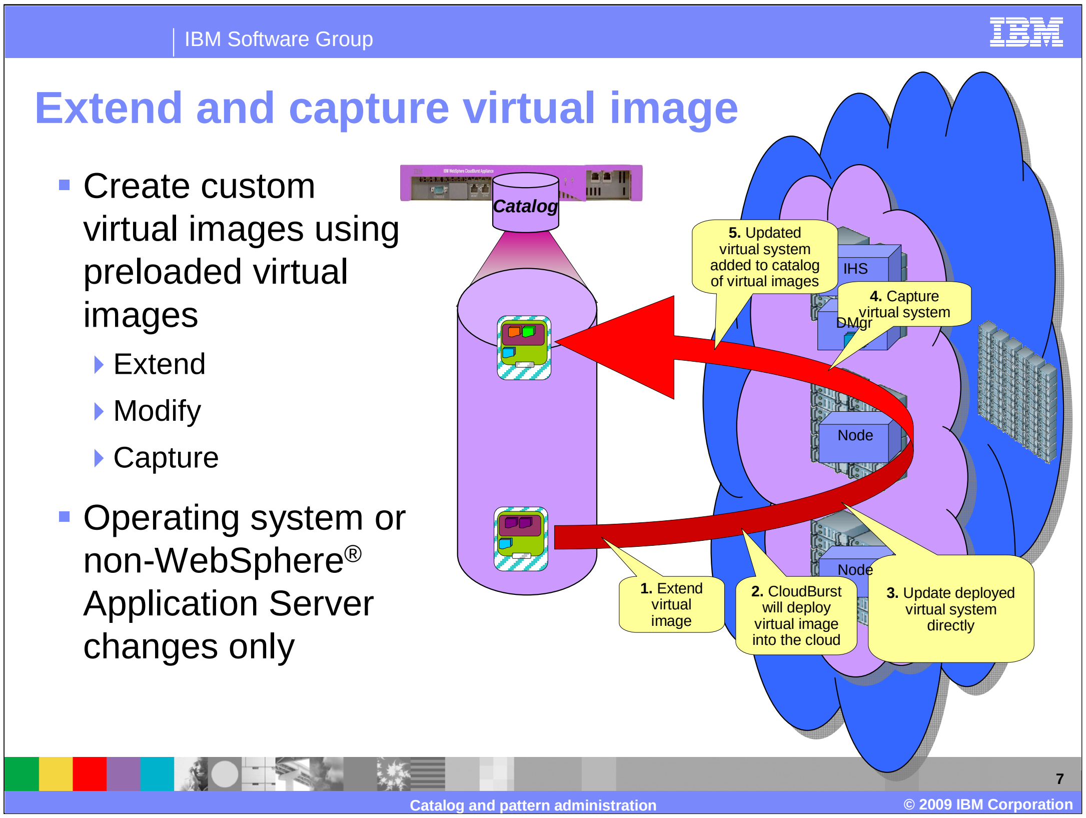

IBM Software Group

Extend and capture virtual image

- Create custom virtual images using preloaded virtual images
  - Extend
  - Modify
  - Capture
- Operating system or non-WebSphere® Application Server changes only

5. Updated virtual system added to catalog of virtual images

4. Capture virtual system

1. Extend virtual image

2. CloudBurst will deploy virtual image into the cloud

3. Update deployed virtual system directly

7

Catalog and pattern administration

© 2009 IBM Corporation

CloudBurst gives you the capability to extend an existing virtual image, make changes, and then capture those changes back into the catalog. You should extend an existing virtual image if you have changes that are not unique to a specific deployment. So for example if your company mandates the use of auditing software this may be a perfect use of the extend and capture capabilities of CloudBurst. The extend and capture feature is meant for operating system and non-WebSphere Application Server profile changes. Any modifications to WebSphere Application Server should be done using script packages.

The first step to create a new virtual image is to choose an existing virtual image as your starting point. Once you have decided on the virtual image, you then click the “Extend” icon near the top of the virtual image attributes page. CloudBurst will then deploy the virtual image into the cloud. You then interact with the deployed virtual image directly by SSH to update the operating system or installing new products.

Once you are sure your environment is completely set up, then click the “Capture” icon located in the top right corner of the virtual image attribute’s page. CloudBurst will capture the changes to the deployed virtual image and store them in the catalog as a new virtual image that you can deploy. Note that if you make any modifications to WebSphere Application Server profiles, those changes will be lost. As part of the capture process the profiles are reset.

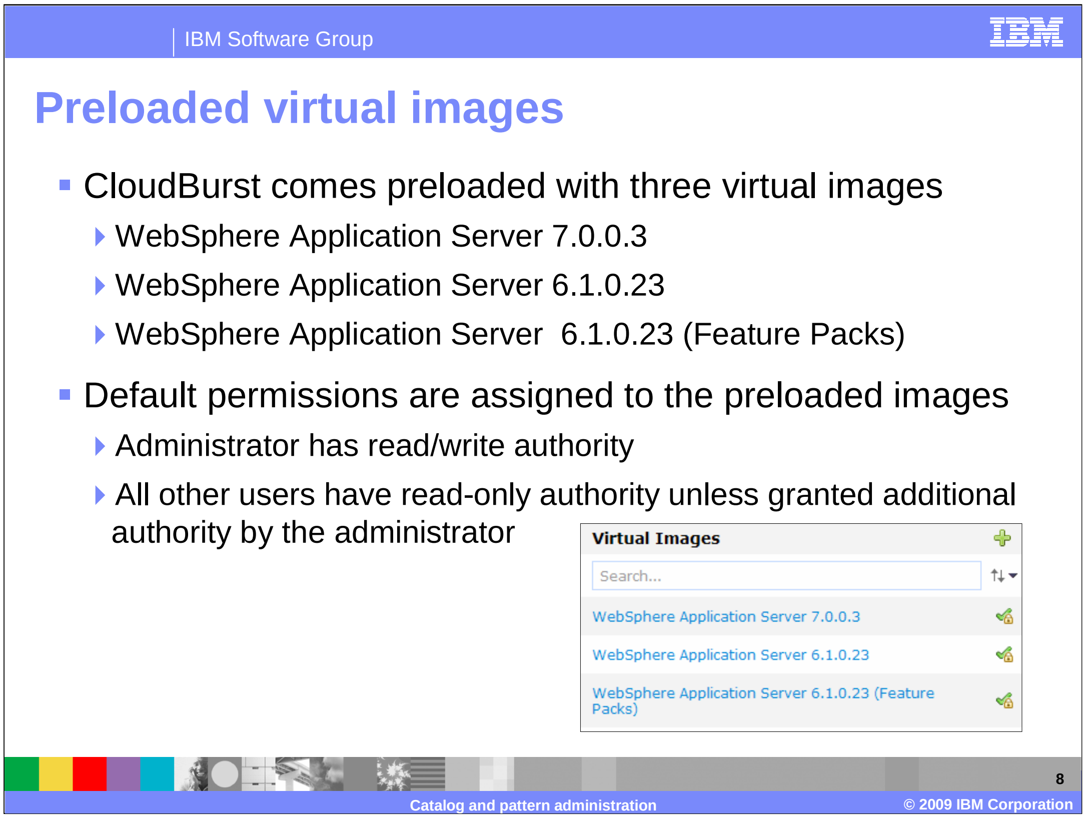

IBM Software Group

Preloaded virtual images

- CloudBurst comes preloaded with three virtual images
  - WebSphere Application Server 7.0.0.3
  - WebSphere Application Server 6.1.0.23
  - WebSphere Application Server 6.1.0.23 (Feature Packs)
- 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

Virtual Images

Search...

WebSphere Application Server 7.0.0.3

WebSphere Application Server 6.1.0.23

WebSphere Application Server 6.1.0.23 (Feature Packs)

8

Catalog and pattern administration

© 2009 IBM Corporation

WebSphere CloudBurst comes preloaded with three versions of WebSphere Application Server Hypervisor Edition. Those versions are “WebSphere Application Server 7.0.0.3”, “WebSphere Application Server 6.1.0.23” and “WebSphere Application Server 6.1.0.23 (Feature Packs)”.

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.

Adding a script package

- Only the name of the script is required to start the process of creating a script package

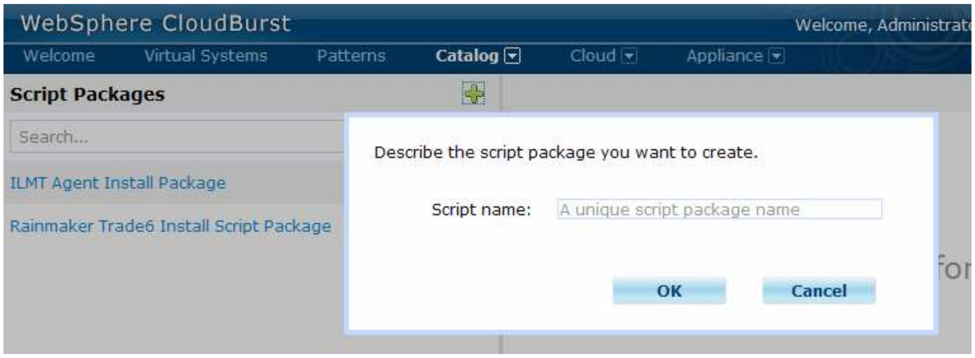

Describe the script package you want to create.

Script name: A unique script package name

Script packages contain scripts and other artifacts that you want run against your pattern after it has been deployed into the cloud. The execution of script packages happens after the environment has been fully initialized, meaning both the operating system and WebSphere Application Server have been started and any configuration changes - such as changes to cell and node names, and federating nodes - have taken place.

Script package attributes

- Script package attributes view
  - Upload your script package
  - Create environment variables
  - Specify an executable
  - Supply arguments to the executable
  - Grant and remove permissions

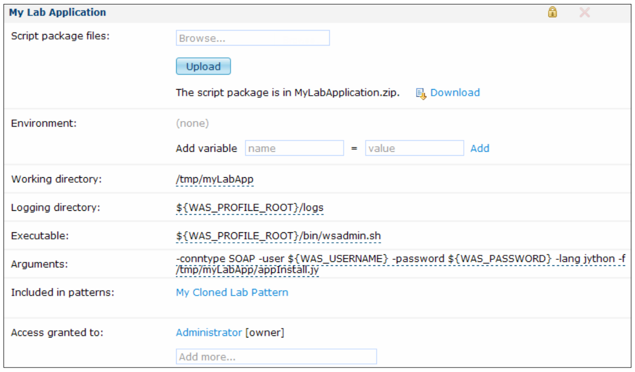

My Lab Application

Script package files:	Browse... Upload The script package is in MyLabApplication.zip. Download
Environment:	(none) Add variable name = value Add
Working directory:	/tmp/myLabApp
Logging directory:	${WAS_PROFILE_ROOT}/logs
Executable:	${WAS_PROFILE_ROOT}/bin/wsadmin.sh
Arguments:	-conntype SOAP -user ${WAS_USERNAME} -password ${WAS_PASSWORD} -lang jython -f /tmp/myLabApp/appInstall.jy
Included in patterns:	My Cloned Lab Pattern
Access granted to:	Administrator [owner] Add more...

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

“Script package files” represent 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 is the command that you want CloudBurst to invoke. CloudBurst will in invoke this command using root authority.

The “Arguments” field are any arguments you want passed to the executable file 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.

The screen capture shown here provides an example. In this illustration, the script package MyLabApplication.zip will be sent to the hypervisor during deployment and extracted into the directory “/tmp/myLabApp”. CloudBurst will then run the “wsadmin.sh” command located under the profile bin directory passing in the appInstall.jy file. Variables WAS_USERNAME and WAS_PASSWORD are CloudBurst provided environment variables.

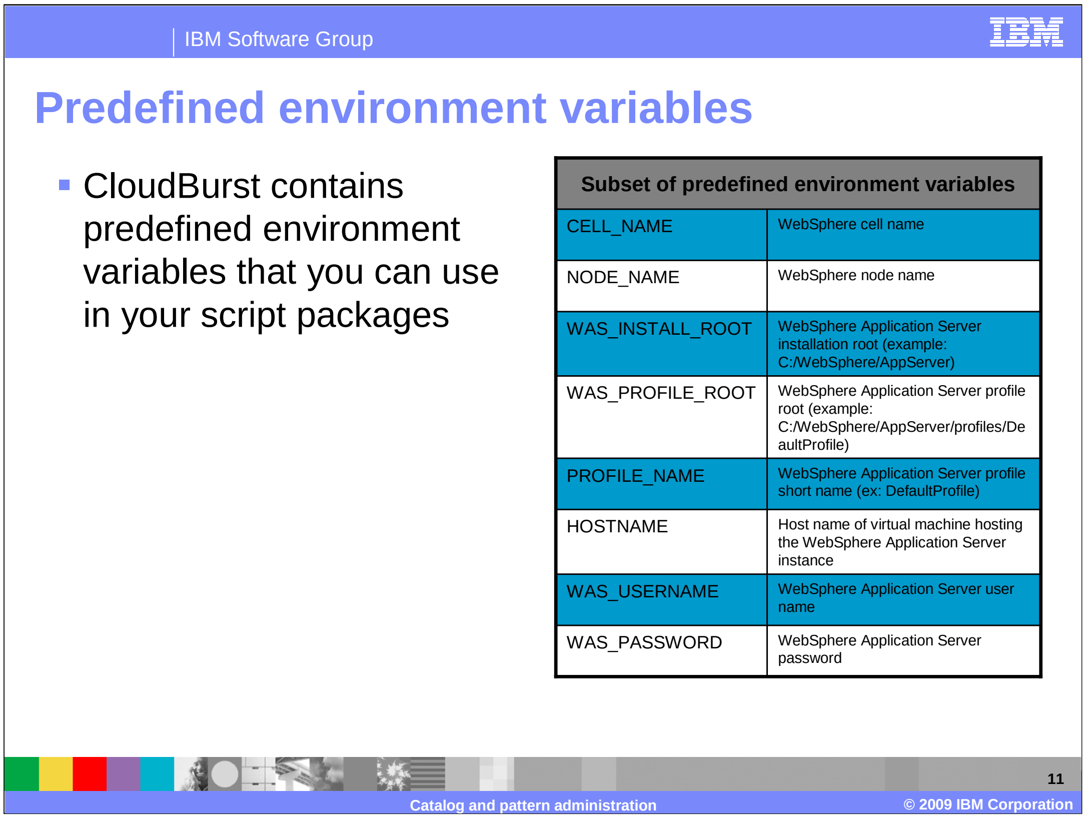

IBM Software Group

Predefined environment variables

- CloudBurst contains predefined environment variables that you can use in your script packages

Subset of predefined environment variables	
CELL_NAME	WebSphere cell name
NODE_NAME	WebSphere node name
WAS_INSTALL_ROOT	WebSphere Application Server installation root (example: C:/WebSphere/AppServer)
WAS_PROFILE_ROOT	WebSphere Application Server profile root (example: C:/WebSphere/AppServer/profiles/DefaultProfile)
PROFILE_NAME	WebSphere Application Server profile short name (ex: DefaultProfile)
HOSTNAME	Host name of virtual machine hosting the WebSphere Application Server instance
WAS_USERNAME	WebSphere Application Server user name
WAS_PASSWORD	WebSphere Application Server password

11

Catalog and pattern administration

© 2009 IBM Corporation

There are many predefined variables that can be used by your scripts. The values are specific to each deployment. Some values do not change such as WAS_INSTALL_ROOT. This slide shows you a subset of the available variables.

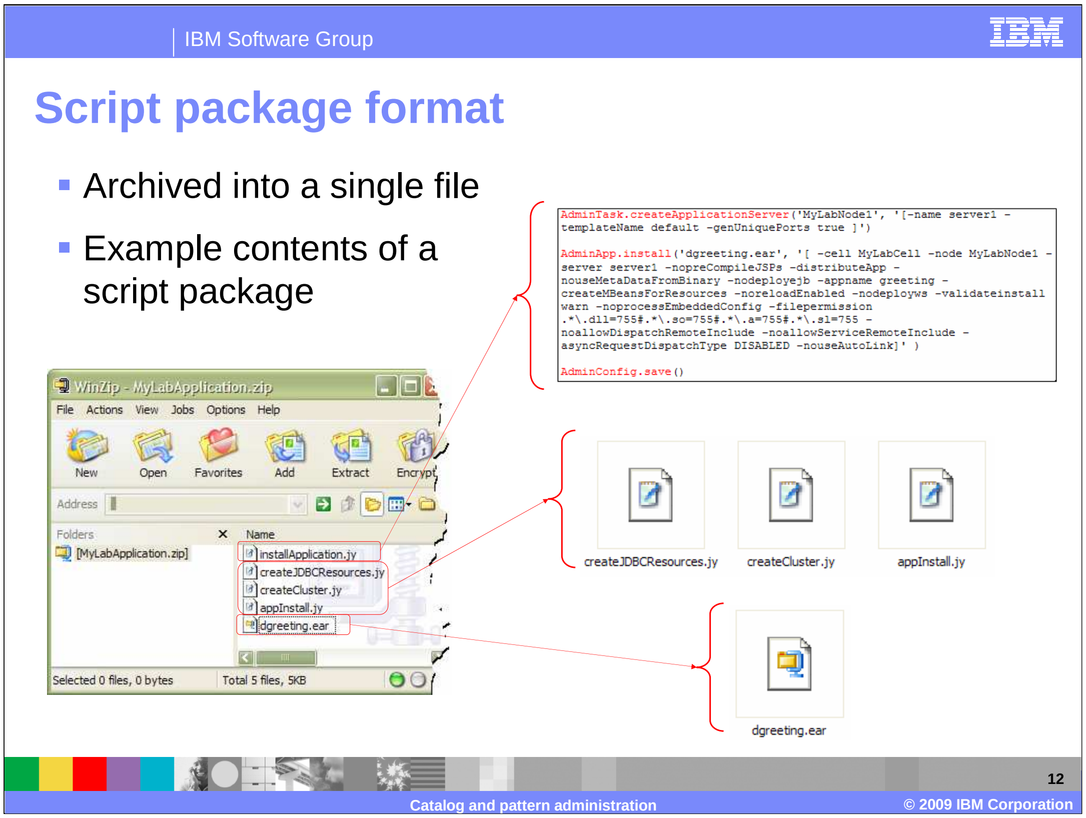

Script package format

- Archived into a single file
- Example contents of a script package

```
AdminTask.createApplicationServer('MyLabNode1', '[-name server1 -
templateName default -genUniquePorts true ]')

AdminApp.install('dgreeting.ear', '[ -cell MyLabCell -node MyLabNode1 -
server server1 -nopreCompileJSPs -distributeApp -
nouseMetaDataFromBinary -nodeployejb -appname greeting -
createMBeansForResources -noreloadEnabled -nodeployws -validateinstall
warn -noprocessEmbeddedConfig -filepermission
.*\.dll=755#.*\.so=755#.*\.a=755#.*\.sl=755 -
noallowDispatchRemoteInclude -noallowServiceRemoteInclude -
asyncRequestDispatchType DISABLED -nouseAutoLink]' )

AdminConfig.save()
```

As has been previously mentioned, script packages can be used for any setup and not just application installation. You are also not locked in to calling wsadmin.sh as your executable file. You can make a call to any executable file residing on the system.

The example script package in this slide contains four scripts and an EAR file. There are scripts to create JDBC™ resources and clusters, and scripts to install the application. There is nothing special about these scripts. The scripts are standard wsadmin Jython scripts. The scripts are extracted into the virtual machine and are run under root authority like any other script. Your scripts do not need to be aware that CloudBurst is involved.

IBM Software Group

Adding an emergency fix

- The name of the emergency fix and a short description are required to start the process of creating an emergency fix

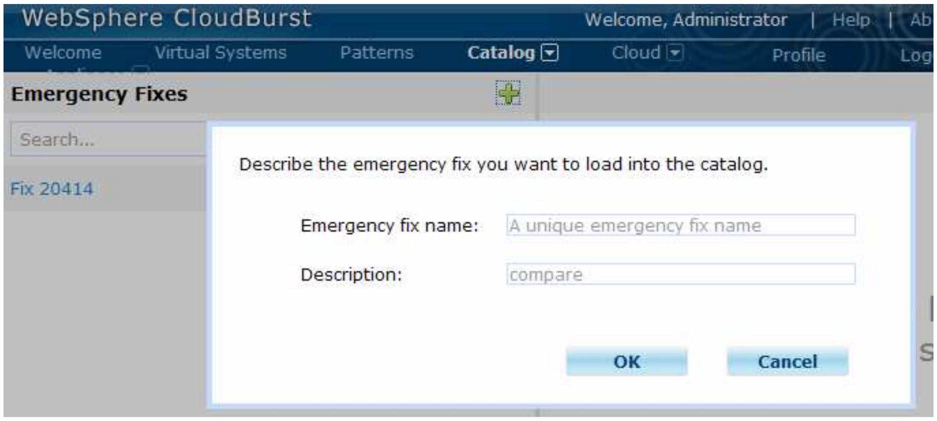

Describe the emergency fix you want to load into the catalog.

Emergency fix name: A unique emergency fix name

Description: compare

13

Catalog and pattern administration

© 2009 IBM Corporation

CloudBurst comes preloaded with images and those images are maintained by IBM, meaning when fixes do become available they will be included in the updated image. However this may not happen as quickly as needed or if you have created your own virtual images you will need another mechanism by which you can fix the system in an emergency.

Emergency fixes allow you to update virtual images that have been deployed into the cloud. Theses updates apply only to the virtual image deployed into the cloud. What this means is that the next time you deploy the image, those fixes will not be present. If you want to update your virtual image you will need to use the extend and capture capabilities. To do this, extend your virtual image, update the deployed virtual image, and then capture the results.

Emergency fix attributes

- Emergency fix attributes view
  - Actual file containing the fix
  - Give other users access to this fix
  - Set the severity level of this fix
  - Set which images this fix can be applied against

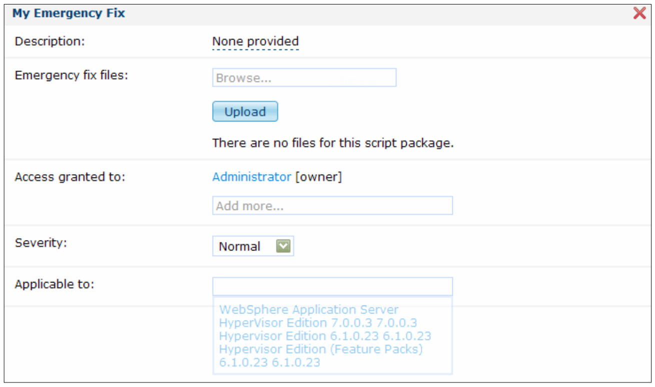

My Emergency Fix

Description:	None provided
Emergency fix files:	Browse... Upload There are no files for this script package.
Access granted to:	Administrator [owner] Add more...
Severity:	Normal
Applicable to:	WebSphere Application Server HyperVisor Edition 7.0.0.3 7.0.0.3 Hypervisor Edition 6.1.0.23 6.1.0.23 Hypervisor Edition (Feature Packs) 6.1.0.23 6.1.0.23

The emergency fix attributes page allows you to upload the fix, grant users access to the fix, set the severity and associate the fix with a virtual image.

Like virtual images and script packages, emergency fixes are uploaded to CloudBurst and reside in the catalog. Before deploying an emergency fix against a virtual system you need to specify which virtual image this fix is to be applied too. So for example if you associate this emergency fix with “WebSphere HV 7.0.0.3” you will only be able to deploy this against the WebSphere Application Server 7.0.0.3 image and no other virtual image.

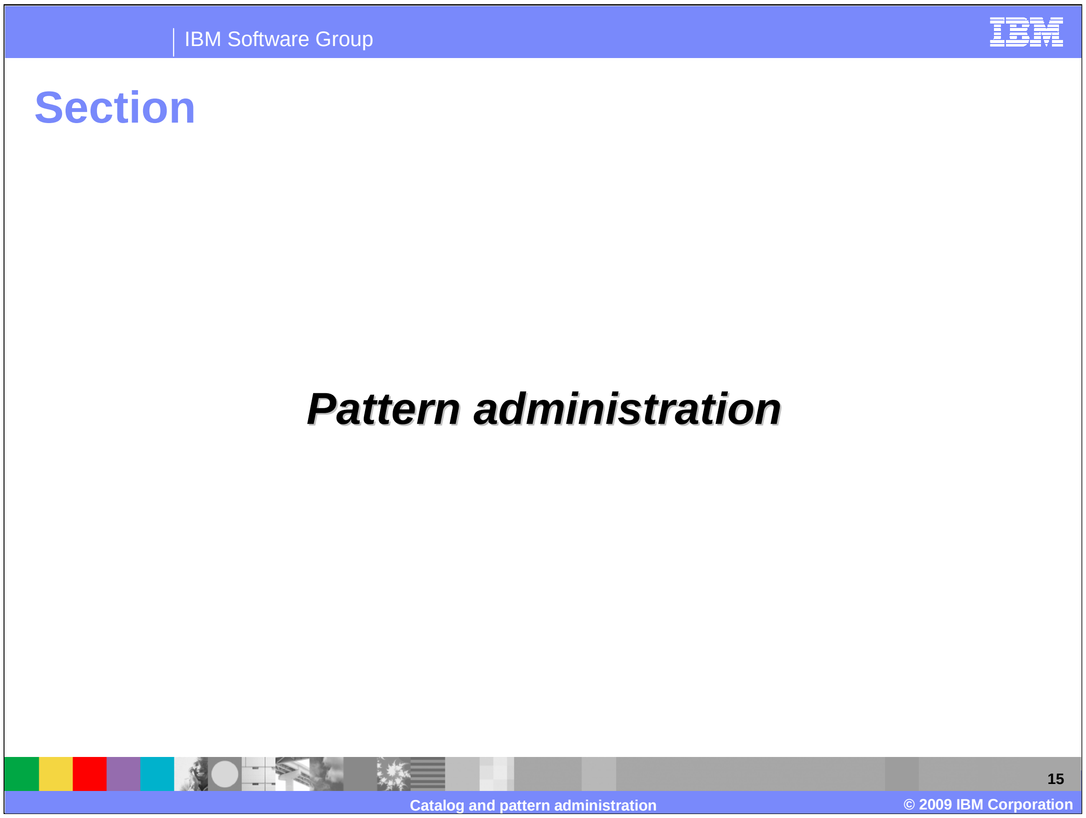

IBM Software Group

Section

Pattern administration

15

Catalog and pattern administration

© 2009 IBM Corporation

This section covers pattern administration.

IBM Software Group

IBM

Creating a pattern

- Two ways to start the process of creating a pattern

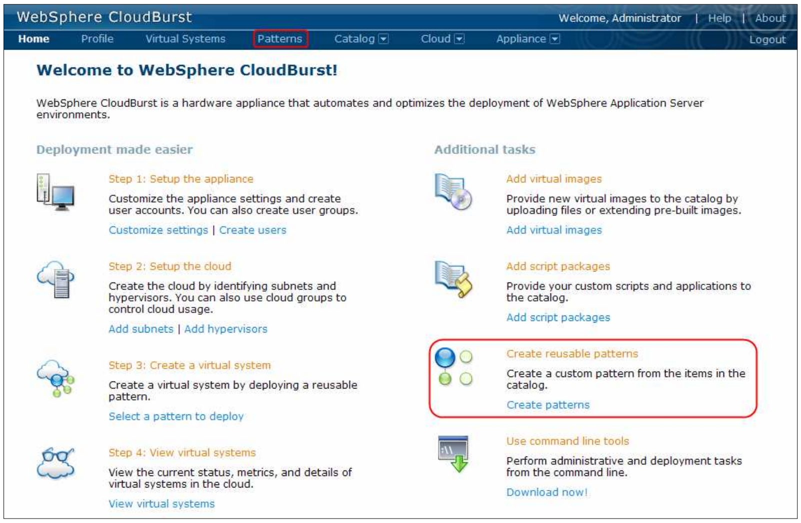

WebSphere CloudBurst

Welcome, Administrator | Help | About

Home | Profile | Virtual Systems | Patterns | Catalog | Cloud | Appliance | Logout

Welcome to WebSphere CloudBurst!

WebSphere CloudBurst is a hardware appliance that automates and optimizes the deployment of WebSphere Application Server environments.

Deployment made easier

Step 1: Setup the appliance

Customize the appliance settings and create user accounts. You can also create user groups.

Customize settings | Create users

Step 2: Setup the cloud

Create the cloud by identifying subnets and hypervisors. You can also use cloud groups to control cloud usage.

Add subnets | Add hypervisors

Step 3: Create a virtual system

Create a virtual system by deploying a reusable pattern.

Select a pattern to deploy

Step 4: View virtual systems

View the current status, metrics, and details of virtual systems in the cloud.

View virtual systems

Additional tasks

Add virtual images

Provide new virtual images to the catalog by uploading files or extending pre-built images.

Add virtual images

Add script packages

Provide your custom scripts and applications to the catalog.

Add script packages

Create reusable patterns

Create a custom pattern from the items in the catalog.

Create patterns

Use command line tools

Perform administrative and deployment tasks from the command line.

Download now!

16

Catalog and pattern administration

© 2009 IBM Corporation

Patterns are essentially a logical grouping of virtual images and script packages into a topology that can then be deployed into the cloud. When you deploy into the cloud you are deploying a pattern and not a virtual image or script package. The virtual images and script packages will get deployed as part of the pattern because they make up the pattern.

There are two ways to create a pattern; either from the main welcome page or from the “Pattern” hyperlink on the top banner.

Creating a pattern (continued)

- Four fields are required to start the process of creating a pattern

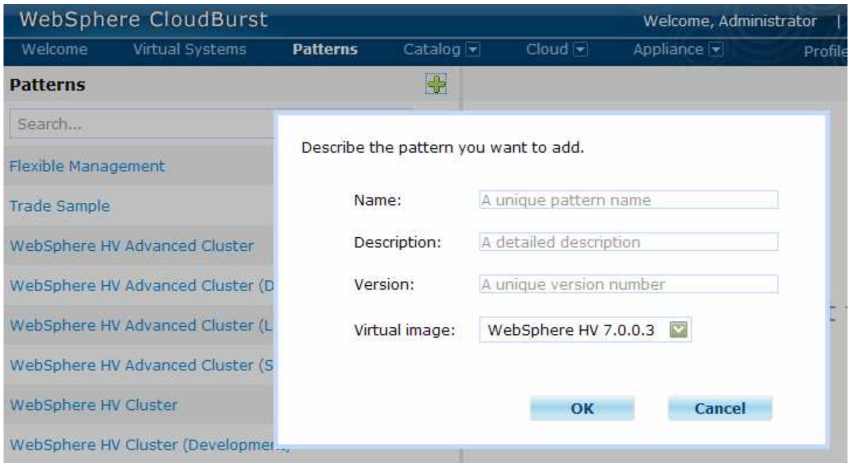

Describe the pattern you want to add.

Name: A unique pattern name

Description: A detailed description

Version: A unique version number

Virtual image: WebSphere HV 7.0.0.3

To start the process of creating a pattern you need to specify a name, description, version and the virtual image that this pattern will be assembled from.

“Name” and “Description” should be something that uniquely identifies the pattern within CloudBurst.

The use of “Version” is not as well defined as “Name” and “Description”. When you make updates to an existing pattern’s parts, you should also update the version to capture this change.

Virtual images are made up of parts. A part corresponds to a specific piece of a virtual image that can be assembled into a topology. For example, the “WebSphere HV 7.0.0.3” contains a deployment manager, custom node, stand-alone node, and HTTP server parts, and together these can be assembled to form a topology. When creating your pattern you have to decide which virtual image you want to use as the source of your parts. What this means is that you cannot mix versions of WebSphere Application Server in a single pattern.

Drag and drop pattern assembly

- Pattern assembly by dragging from the palette over to the canvas
  - Virtual image parts
  - Script packages

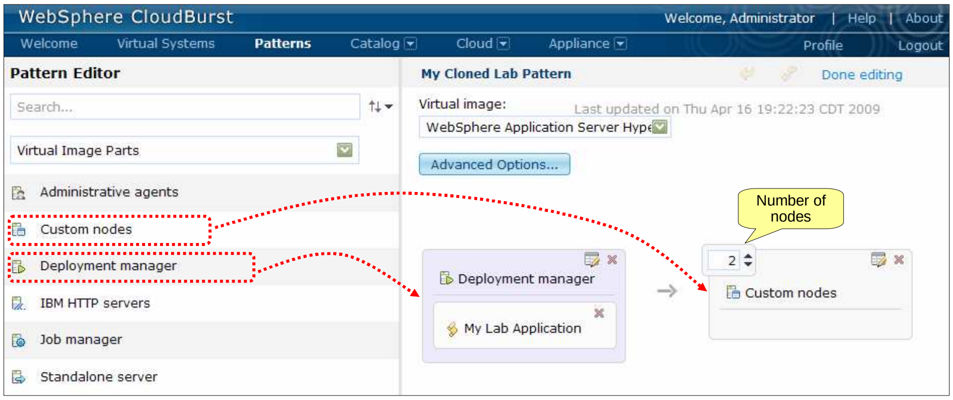

If this is a new pattern you will see a blank canvas on the right and a palette of parts on the left. The palette on the left contains your virtual image parts and script packages which you will use to assemble your pattern. You create a pattern by dragging parts from the palette and dropping them onto 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.

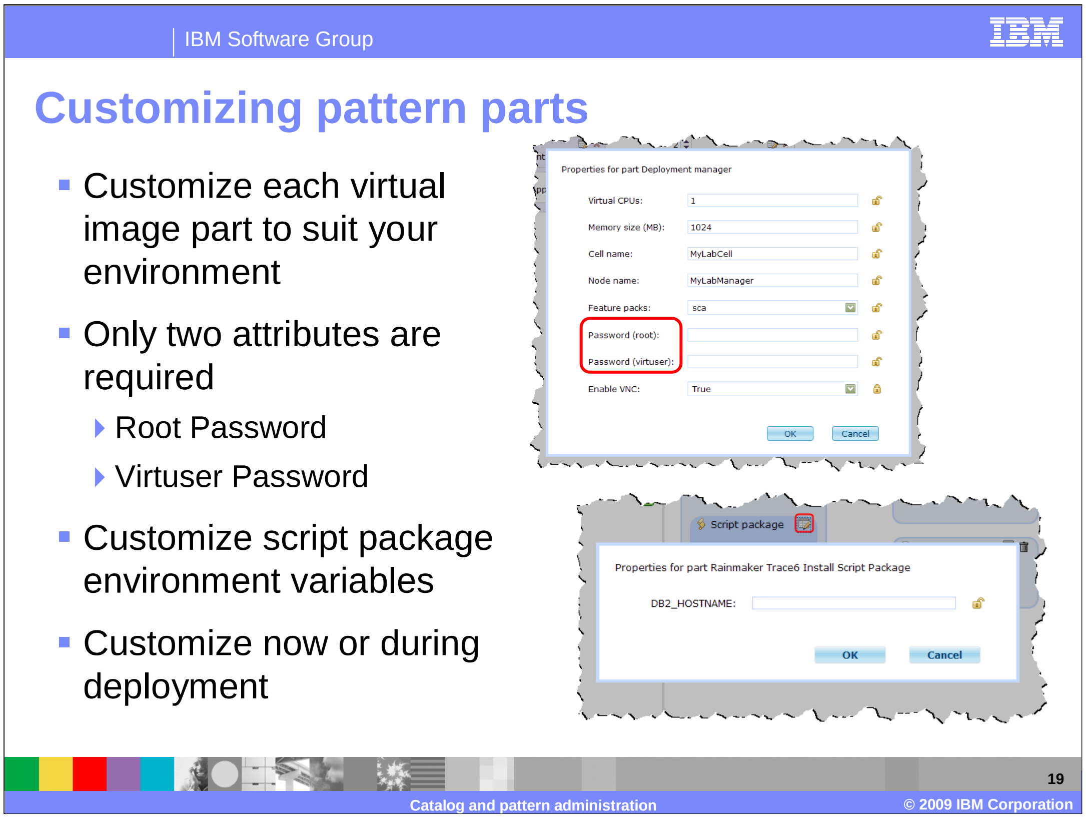

IBM Software Group

Customizing pattern parts

- Customize each virtual image part to suit your environment
- Only two attributes are required
  - Root Password
  - Virtuser Password
- Customize script package environment variables
- Customize now or during deployment

Properties for part Deployment manager

Virtual CPUs: 1
Memory size (MB): 1024
Cell name: MyLabCell
Node name: MyLabManager
Feature packs: sca
Password (root):
Password (virtuser):
Enable VNC: True

Properties for part Rainmaker Trace6 Install Script Package

DB2_HOSTNAME:

19

Catalog and pattern administration

© 2009 IBM Corporation

As you have seen a pattern is made up of a topology parts and associated script packages.

Each part and script package allows for additional configuration. Located in the upper right corner of each part or script package you can click the edit icon to open up the attributes for that specific part or script package. You can choose to configure these additional attributes now during pattern creation or wait until you deploy the pattern. It is important to understand the result or side effect of each choice.

If you choose to configure your parts or script packages during pattern creation then the changes will be part of the pattern's template. Meaning every time you deploy the pattern these changes will be present. You can reconfigure during deployment if needed. If you want to hard code an attribute in the template you configure it during pattern creation and click the lock icon to the right of the field. Locking the field will prevent anyone from changing this value during deployment time.

If you choose to configure your parts or script packages during pattern deployment then the changes will be specific to the deployment. CloudBurst only requires two fields to be filled in and they are "Password (root)" and "Password (virtuser)". All other fields CloudBurst will fill in with default values.

The value you enter for "Root Password" will be used for the operating system password. The value you enter for "Virtuser Password" will be used for the WebSphere Application Server environment.

Advanced customization

- Allows you to configure advanced clustering options
  - Messaging
  - Session persistence
  - Global security

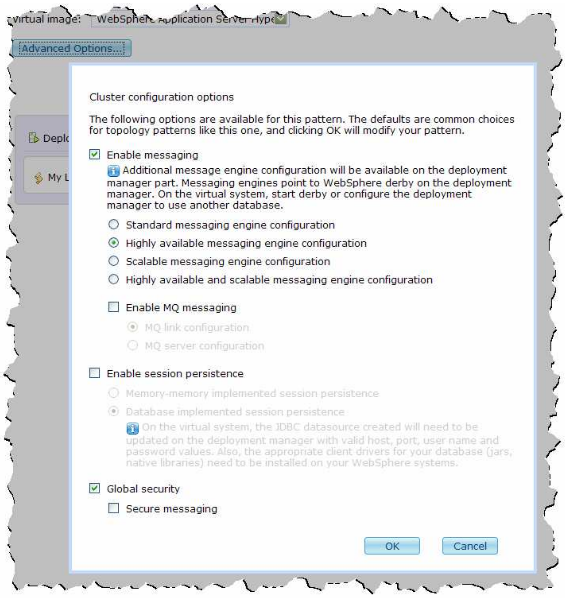

Cluster configuration options

The following options are available for this pattern. The defaults are common choices for topology patterns like this one, and clicking OK will modify your pattern.

☑ Enable messaging

Additional message engine configuration will be available on the deployment manager part. Messaging engines point to WebSphere derby on the deployment manager. On the virtual system, start derby or configure the deployment manager to use another database.

○ Standard messaging engine configuration
● Highly available messaging engine configuration
○ Scalable messaging engine configuration
○ Highly available and scalable messaging engine configuration

□ Enable MQ messaging
● MQ link configuration
○ MQ server configuration

□ Enable session persistence
○ Memory-memory implemented session persistence
● Database implemented session persistence

On the virtual system, the JDBC datasource created will need to be updated on the deployment manager with valid host, port, user name and password values. Also, the appropriate client drivers for your database (jars, native libraries) need to be installed on your WebSphere systems.

☑ Global security
□ Secure messaging

OK Cancel

Advanced options are provided by CloudBurst. These are advance settings such as setting up messaging, enabling security and enabling session persistence. You can apply these to your patterns.

Pattern attributes and options

- From this view you can
  - Deploy, edit, clone, lock and delete your pattern
  - View pattern topology
  - View the virtual systems currently created using this pattern
  - Who has permission to view and manipulate this pattern

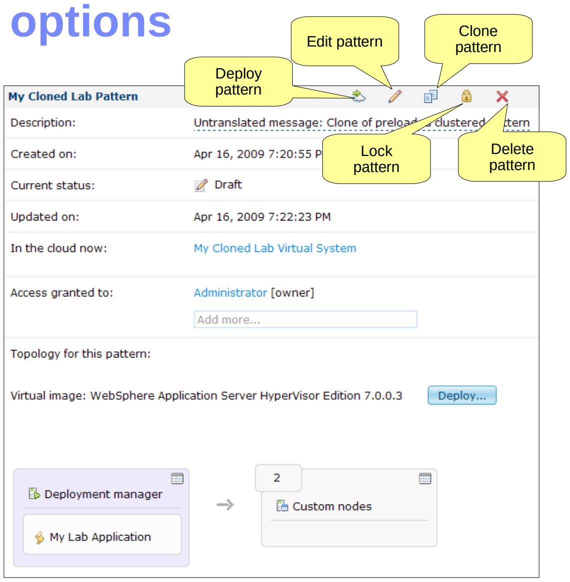

My Cloned Lab Pattern

Description: Untranslated message: Clone of preloa clustered tern

Created on: Apr 16, 2009 7:20:55 P

Current status: Draft

Updated on: Apr 16, 2009 7:22:23 PM

In the cloud now: My Cloned Lab Virtual System

Access granted to: Administrator [owner]

Add more...

Topology for this pattern:

Virtual image: WebSphere Application Server HyperVisor Edition 7.0.0.3 Deploy...

The option ribbon at the top of your pattern provides options to deploy, edit, clone, lock and delete. The “In the cloud now” attribute show which virtual systems were created from this pattern and “Access granted to” attribute shows who has permission to view, modify and update this pattern.

There are two options that warrant a deeper look and those are the “Clone” and “Lock” pattern options.

Once your pattern has been created and tested and you want to prevent any further changes, you can lock the pattern. Locking a pattern prevents any further changes to the pattern. If you decide later that you need to update a locked pattern, you will need to clone the pattern and make updates to the cloned pattern.

The “Clone” option allows you to take an existing pattern and make an exact copy of it. A recommended use of this feature is to clone an existing preloaded “best practices” pattern as a starting point for your custom pattern.

Schedule a pattern deployment

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

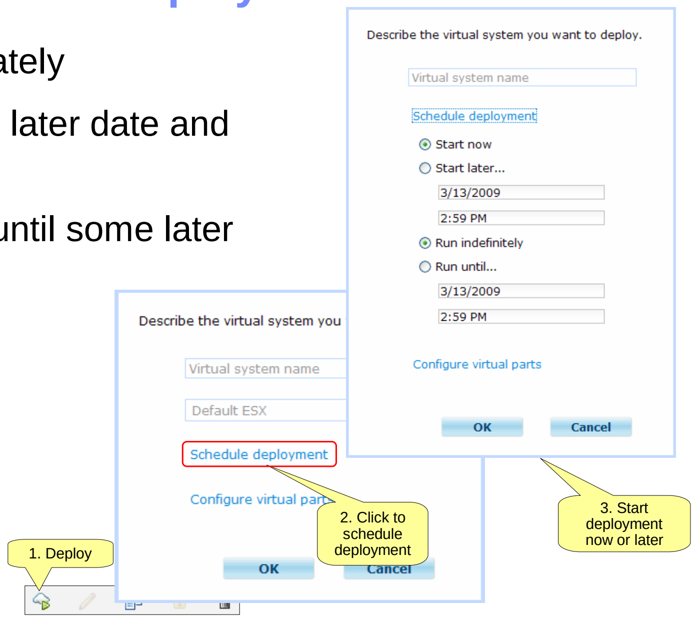

Describe the virtual system you

Virtual system name

Default ESX

Describe the virtual system you want to deploy.

Virtual system name

● Start now
○ Start later...
3/13/2009
2:59 PM
● Run indefinitely
○ Run until...
3/13/2009
2:59 PM

1. Deploy

2. Click to schedule deployment

3. Start deployment now or later

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

Configure a pattern deployment

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

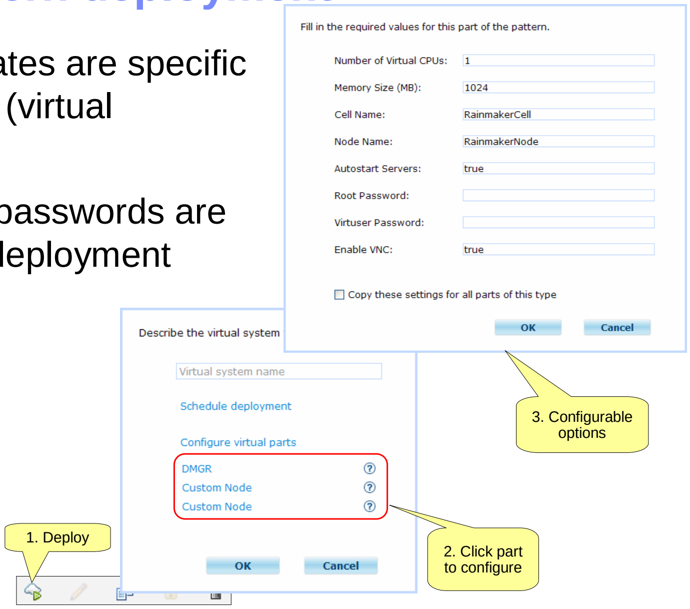

Fill in the required values for this part of the pattern.

Number of Virtual CPUs:	1
Memory Size (MB):	1024
Cell Name:	RainmakerCell
Node Name:	RainmakerNode
Autostart Servers:	true
Root Password:	
Virtuser Password:	
Enable VNC:	true

□ Copy these settings for all parts of this type

When you deploy a pattern you are given a chance to configure the pattern’s parts and script package variables. CloudBurst requires only 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.

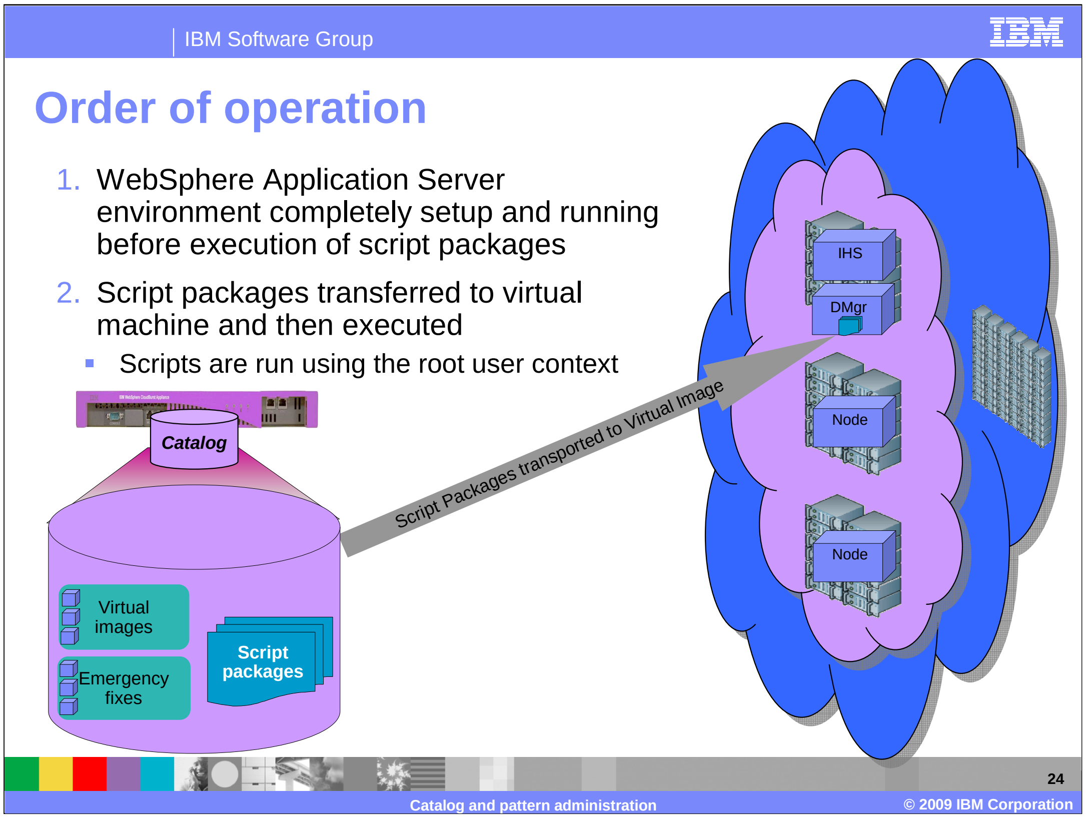

IBM Software Group

Order of operation

1. WebSphere Application Server environment completely setup and running before execution of script packages
2. Script packages transferred to virtual machine and then executed
   - Scripts are run using the root user context

24

Catalog and pattern administration

© 2009 IBM Corporation

After you schedule a deployment CloudBurst will transfer the virtual images and scripts over to the hypervisor. During initial boot of the virtual machine the activation engine will configure the operating system. Once the operating system is set up WebSphere Application Server will be started and configured. Only after the entire environment is configured and operational will the script packages be executed.

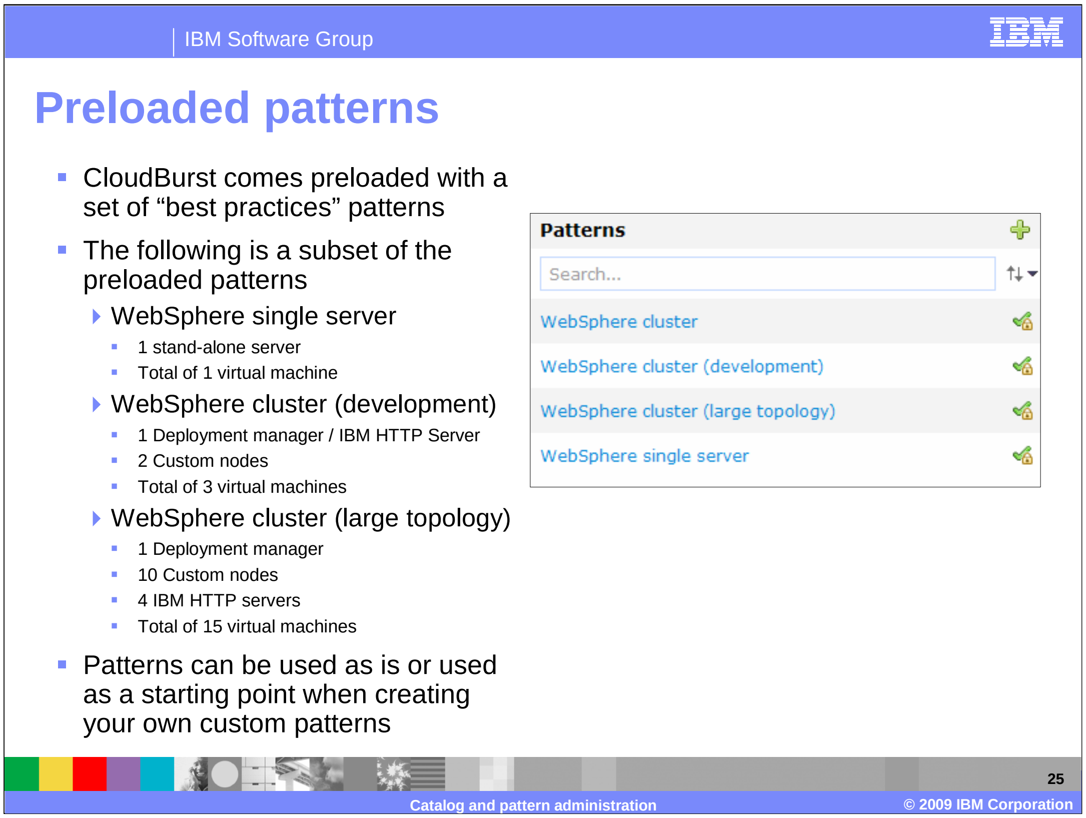

IBM Software Group

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 1 virtual machine
  - WebSphere cluster (development)
    - 1 Deployment manager / IBM HTTP Server
    - 2 Custom nodes
    - Total of 3 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

Patterns

Search...

WebSphere cluster

WebSphere cluster (development)

WebSphere cluster (large topology)

WebSphere single server

25

Catalog and pattern administration

© 2009 IBM Corporation

WebSphere CloudBurst comes preloaded with a set of “best practices” patterns. When creating your own custom pattern you should use one of the preloaded “best practices” patterns. By starting with one of CloudBurst’s existing patterns you are starting with a “best practice, time tested approach”.

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

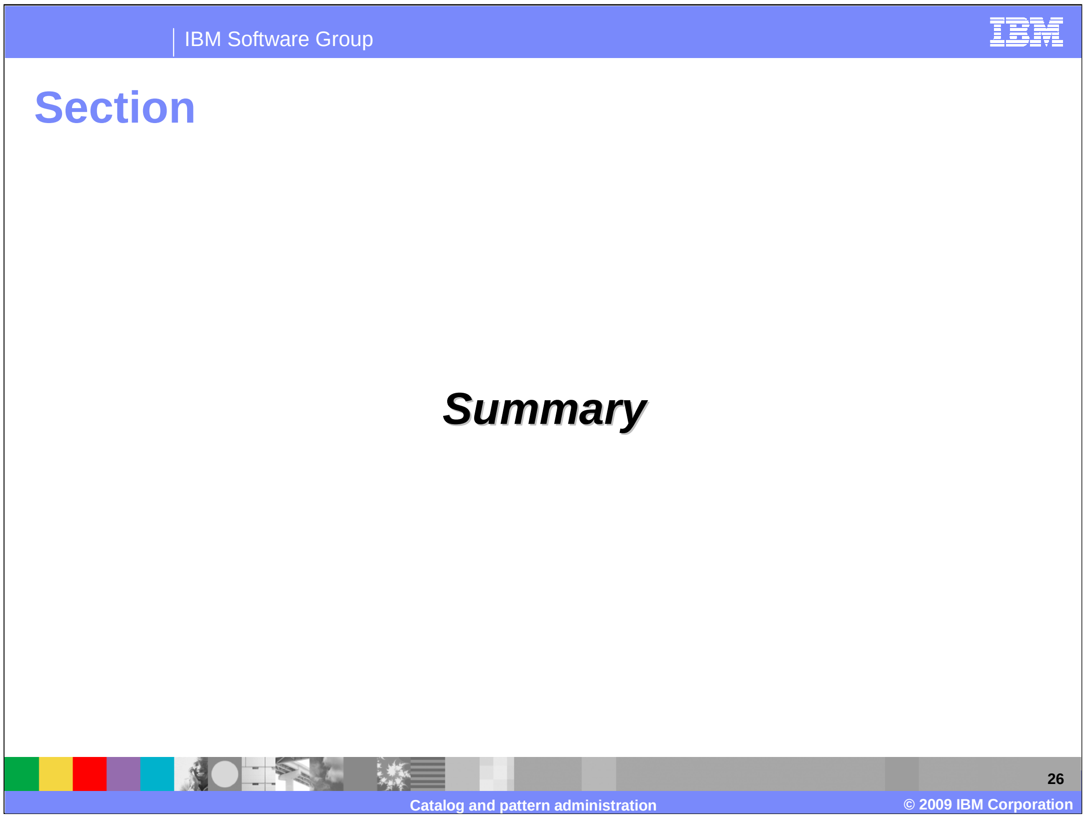

IBM Software Group

Section

Summary

26

Catalog and pattern administration

© 2009 IBM Corporation

And, in summary…

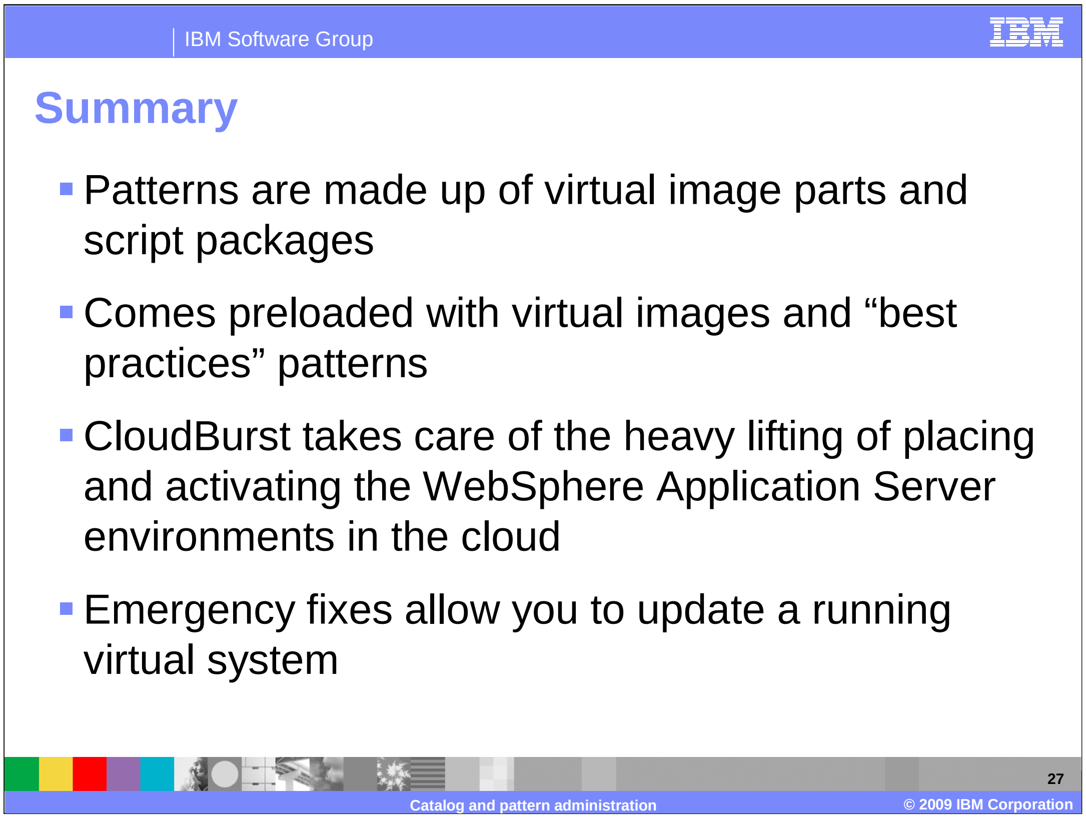

Summary

- Patterns are made up of virtual image parts and script packages
- Comes preloaded with virtual images and “best practices” patterns
- CloudBurst takes care of the heavy lifting of placing and activating the WebSphere Application Server environments in the cloud
- Emergency fixes allow you to update a running virtual system

CloudBurst gives you every tool to create and deploy a WebSphere Application Server topology into the cloud quickly and easily. To make this happen CloudBurst comes preloaded with virtual images and “best practices” patterns ready for deployment. CloudBurst also takes care of the heavy lifting of pattern placement in the cloud and activation of the operating system and WebSphere Application Server components. If a problem happens in one of your virtual systems CloudBurst gives you the option of applying a fix directly to the virtual system in the case of an emergency.

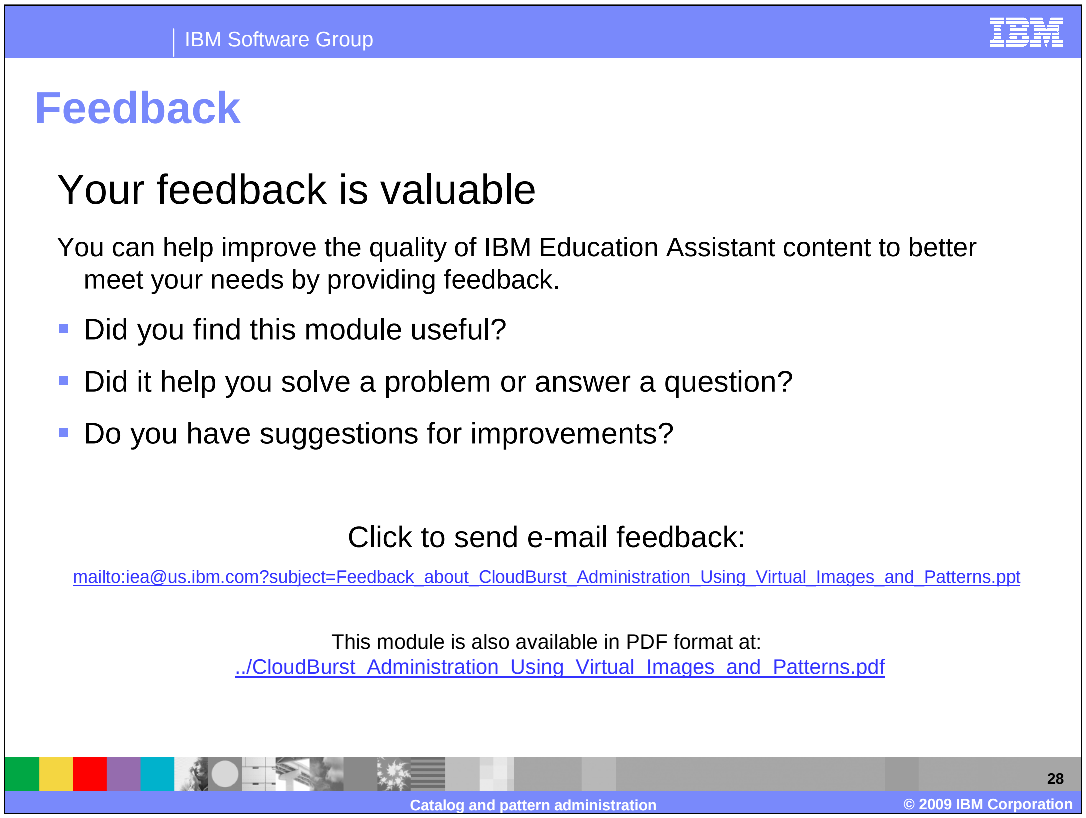

IBM Software Group

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_Administration_Using_Virtual_Images_and_Patterns.ppt

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

28

Catalog and pattern administration

© 2009 IBM Corporation

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

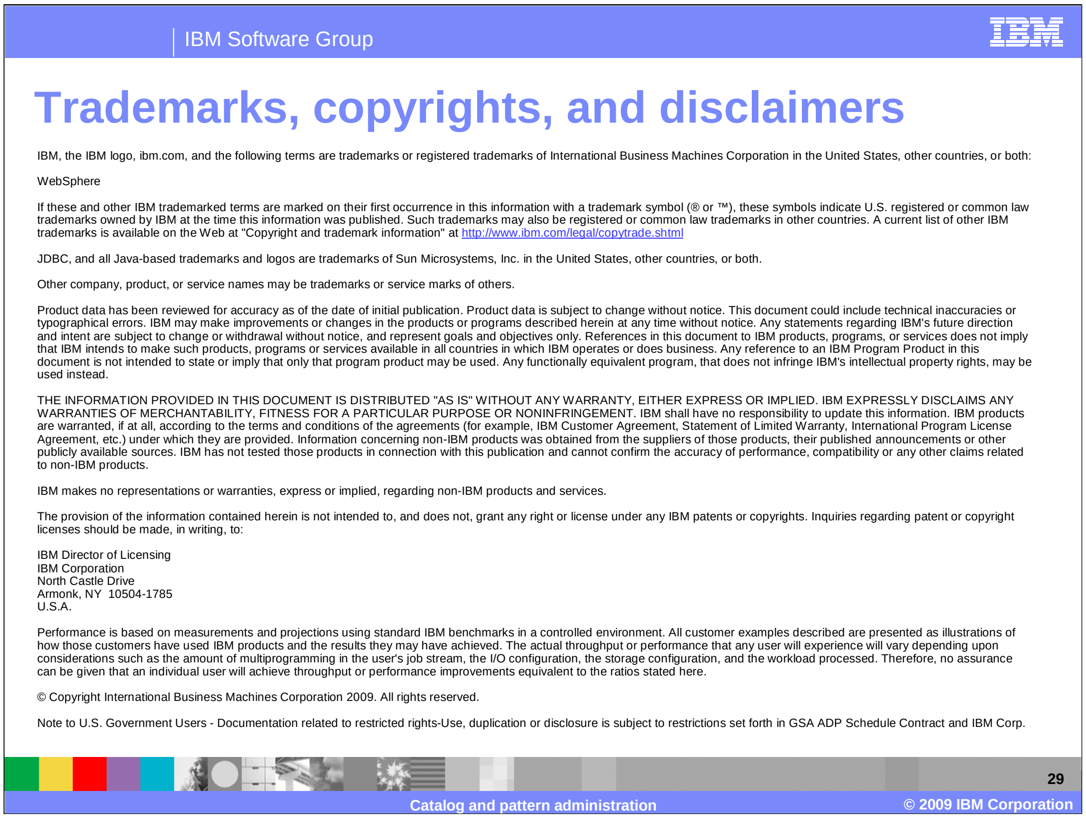

Trademarks, copyrights, and disclaimers

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:

WebSphere

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.

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.

Note to U.S. Government Users - Documentation related to restricted rights-Use, duplication or disclosure is subject to restrictions set forth in GSA ADP Schedule Contract and IBM Corp.