

IBM WebSphere CloudBurst Appliance

Patterns



This presentation will cover WebSphere® CloudBurst™ patterns.

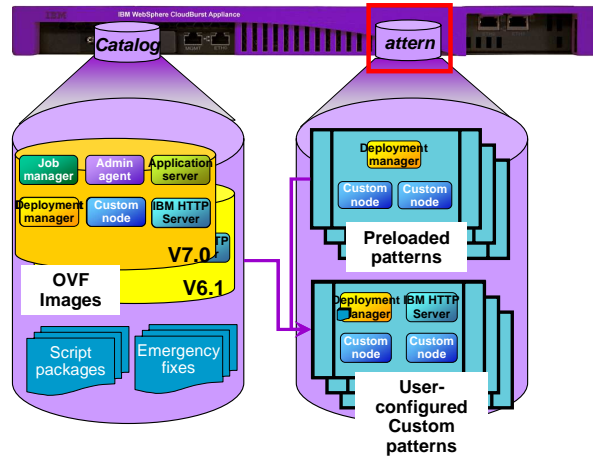
Table of contents

- Pattern defined
- Pattern creation
- Pattern deployment
- Preloaded patterns

This presentation will cover the steps involved in the creation and deployment of patterns. It will also cover a set of “best practices” patterns that come preloaded on the appliance ready for immediate use.

Pattern defined

- Pattern is made up of virtual image parts and script packages
- Industry “best practices” patterns come preloaded on the appliance
- Use as-is or customize to suit your needs



Patterns are a combination of virtual image parts and script packages combined together to create a WebSphere Application Server topology. A pattern is a set of instructions to WebSphere CloudBurst. It tells WebSphere CloudBurst the structure of the WebSphere Application Server environment and how to configure it.

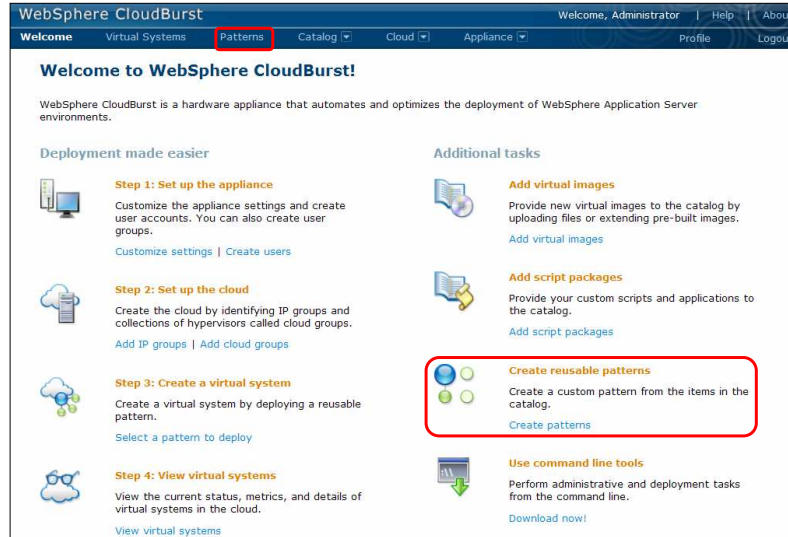
Patterns come preloaded and ready for immediate deployment. If one of the preloaded patterns does not fit your scenario you can take one of the preloaded patterns clone it and configure the clone to fit your scenario.

Pattern creation

This section will cover the steps involved in creating a pattern.

Creating a pattern

- Two ways to get to the pattern creation page



5


Patterns

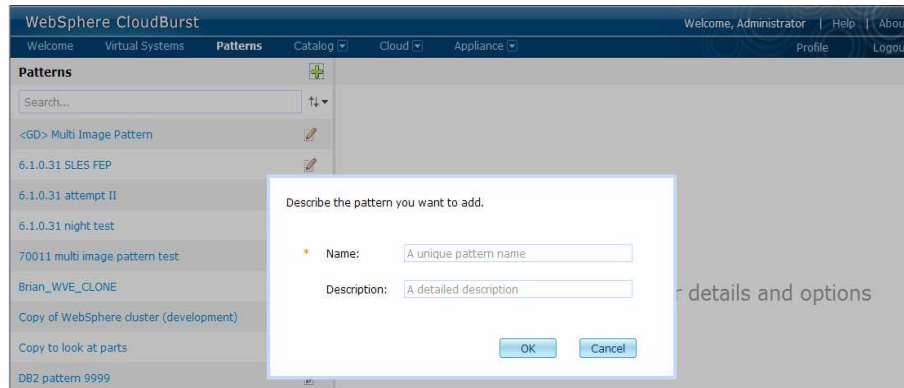
© 2010 IBM Corporation

A pattern is a logical grouping of virtual image parts forming a WebSphere Application Server topology and associated script packages to further define the environment. You deploy a pattern as a single unit into the cloud. The virtual images and script packages will get deployed as part of the pattern deployment.

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)

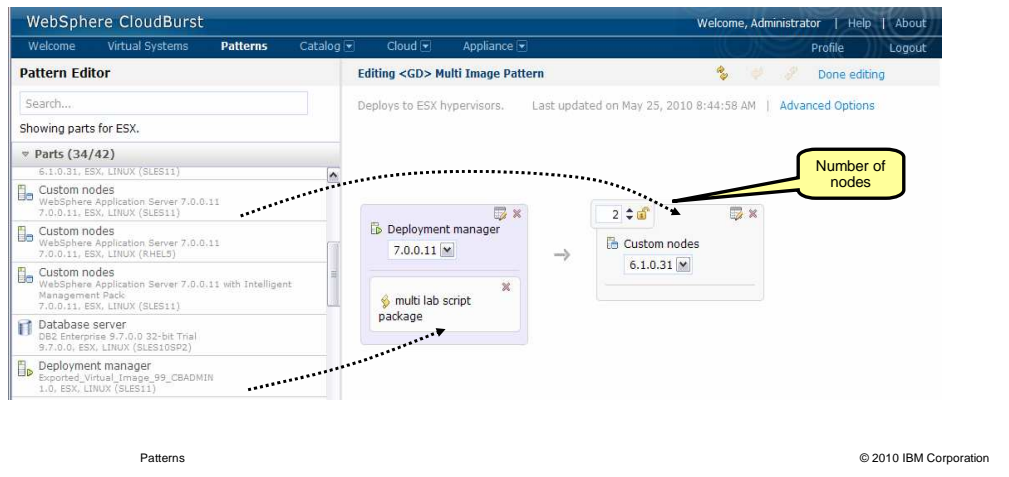
- From the Patterns page, click the green plus icon to create a new pattern 
- Two fields are required to start the process of creating a pattern



A unique name is all that is required to create a pattern. It is good practice to specify a description for your pattern to avoid any possible confusion later.

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 are presented with a blank canvas on the right and a palette of virtual image 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. There is also an option to choose which virtual image will make up this pattern. The virtual image will dictate the available virtual image parts that you can use to make up your pattern.

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

Customizing pattern parts

- Navigate to: **Patterns > your_pattern > Edit > edit_configurable_part**
- Customize each virtual image part to suit your environment
- Two attributes are required to be filled in now or during pattern deployment
 - “root” password
 - “virtuser” password
- Customize script package environment variables
- Customize now or during deployment

The screenshot displays two configuration windows. The top window, titled 'Properties for part Standalone server', contains the following fields:

- Name: StandalonePart
- Virtual CPUs: 1
- Memory size (MB): 2048
- Cell name: CloudBurstCell
- Node name: CloudBurstNode
- Feature packs: none, cea, sca, sdo, xml
- Password (root):

 The bottom window, titled 'Parameters for script createDatabase9999', contains:

- DATABASE_NAME:

 Yellow arrows point from the list items in the text to the 'Standalone server' and 'Script package' components in the interface.

8

Patterns

© 2010 IBM Corporation

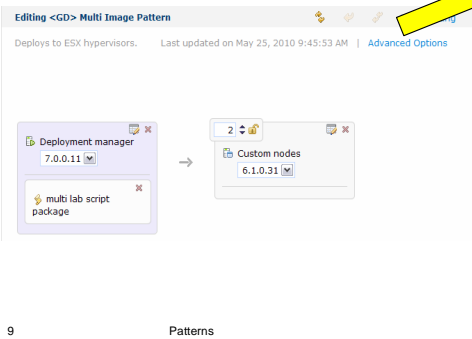
Once you have assembled your high level topology you can further configure your environment by clicking on the edit icon of each configurable virtual image part and script package to bring up a list of configurable options.

You can configure your environment now or during pattern deployment. If you configure it now then the configuration is available each time you deploy the pattern. It essentially becomes part of the template. You also have the option of locking down your configuration options. If for example, your company has a policy where all WebSphere Application Server cells are required to have a cell name “CompanyCell”, you enter “CompanyCell” into the “Cell name” field and click the lock icon located to the right of the field. The locked field is removed from the list of configurable options during pattern deployment enforcing your cell name.

There are two required fields that you must enter values for either now or during pattern deployment. It is preferable that you fill these in at pattern deployment time to ensure that they are unique to the deployment. Those fields are the “root” and “virtuser” password fields. All other fields are given a default value if you do not specify one. It is best practice to configure a unique set of passwords for each deployment, so leaving the password fields blank during pattern creation is preferred. “Root” user is configured as the operating systems root user. “Virtuser” user is an additional non-administrative user which you use to log into the WebSphere Application Server administrative console.

Advanced customization

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



The following options are available for this pattern. The recommended options are common choices for topology patterns like this one.

Define clusters

Once enabled, the number of clusters and application servers can be configured on the deployment manager part.

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 (legacy feature)

- MQ link configuration
- MQ server configuration

Enable session persistence

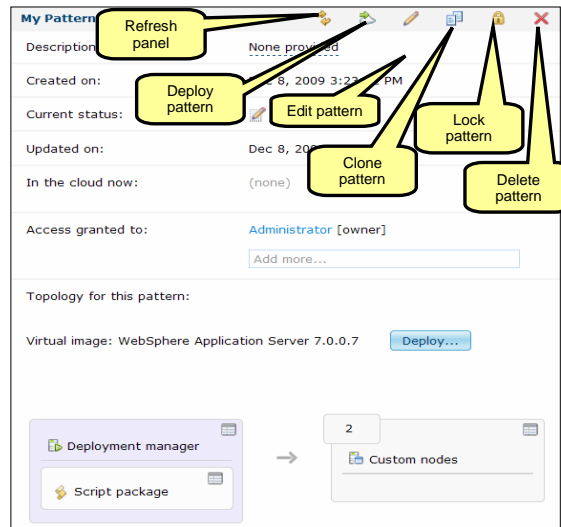
- Memory-memory implemented session persistence
- Database implemented session persistence

OK Cancel

WebSphere CloudBurst provides a set of predefined configuration options of some commonly and typically more complex settings. These can be found under the “Advance Options” button. There are options for configuring clusters, messaging and global security. You can accomplish the same configuration by creating your own script packages. In fact, the WebSphere CloudBurst is using the script package feature to accomplish this.

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
 - See who has permission to view and manipulate this pattern



10

Patterns

© 2010 IBM Corporation

The option ribbon at the top of your pattern provides options to deploy, edit, clone, lock and delete. The “In the cloud now” attribute shows 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 you have created, tested and verified your pattern you should lock down the pattern to prevent any further changes. You can lock a pattern by clicking the lock icon. Locking a pattern is a permanent operation; once locked it cannot be unlocked. All is not lost if at a later time you decide that changes are required. You need to clone the pattern, make changes to the clone, test and verify. Once you are confident in your new pattern you can delete the original pattern.

Cloning a pattern makes an exact copy of the original. Typical use of the clone operation is to take one of the predefined “best practices” patterns, clone the pattern, and use this as a basis for your own custom patterns.

Pattern deployment

This section will cover pattern deployment.

Pattern deployment steps

- Pattern deployment steps
 1. Specify a unique name for the virtual system to be created
 2. Choose an existing cloud group to deploy your pattern into
 3. Schedule a deployment time
 4. Configure virtual image parts and script packages

Deploying a pattern requires four steps. Step one; provide a unique name for the virtual system. Step two; choose an existing cloud group to deploy your pattern into. Step three; schedule a deployment time. And finally, step four; configure any remaining configurable options for your virtual image parts and script packages.

Schedule a pattern deployment

- Deploy immediately
- Deploy at some later date and time
- Run forever or until some later date and time
- Default option is to start now and run indefinitely

1. Deploy

2. Click to schedule deployment

3. Start deployment now or later

13

Patterns

© 2010 IBM Corporation

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 later time. The default option is to start the deployment immediately and run indefinitely.

Configure a pattern deployment

- Configuration updates are specific to this deployment (virtual system)
- Values for “root” and “virtuser” passwords are required by the user

The screenshot shows two overlapping dialog boxes. The background dialog, titled 'Describe the virtual system you want to deploy.', has a 'Virtual system name' field, a 'vCenter Cloud Group' dropdown, and a 'Schedule deployment' button. Below these are expandable sections for 'Configure virtual parts', including 'Deployment manager' and 'Custom nodes'. A callout '1. Deploy' points to the top of this dialog. A callout '2. Click each part to configure' points to the 'Deployment manager' and 'Custom nodes' sections. The foreground dialog, titled 'Fill in the required values for this part of the pattern.', is for 'DMGRPart' and contains fields for 'Virtual CPUs' (set to 1), 'Memory size (MB)' (2048), 'Cell name' (CloudBurstCell), 'Node name' (CloudBurstNode), 'Feature packs' (with 'none' selected), and 'Password (root)'. A callout '3. Configurable options' points to the 'Feature packs' section. 'OK' and 'Cancel' buttons are at the bottom of the foreground dialog.

14


Patterns

© 2010 IBM Corporation

When you deploy a pattern you are given a chance to configure the pattern’s virtual image parts and script package variables. CloudBurst requires only the “Root Password” and the “Virtuser Password” fields 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 WebSphere CloudBurst but rather the requirements of the specific script package.

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



15

Patterns

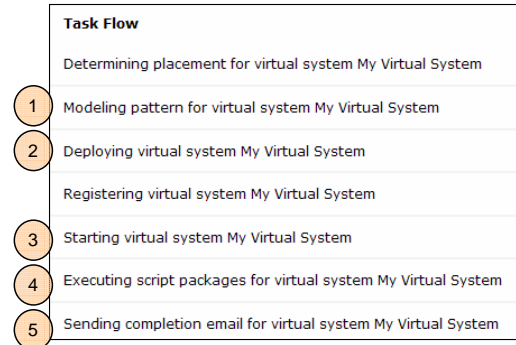
© 2010 IBM Corporation

After you schedule a deployment WebSphere CloudBurst will transfer the virtual images and script packages 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 is started and configured. Only after the entire environment is configured and operational will the script packages be executed.

Pattern deployment – Task queue

Pattern deployment steps broken down

1. WebSphere CloudBurst determines placement of virtual system
2. Virtual system is deployed into the cloud.
3. Virtual machines are configured (OS configured and started, WebSphere Application Server configured and started)
4. Script packages are transferred to target virtual machines and run
5. Email sent out to virtual system owner with success or failure of virtual system deployment



This slide breaks down each step of the pattern deployment process. The screen capture was taken from the “Task Queue”. There are seven steps in total and five are discussed further.

Step one, WebSphere CloudBurst determines placement of the virtual machines that make up the virtual system.

Step two, virtual images are transferred to the destination hypervisors determined in step one. The transfer of virtual images happens once for each hypervisor and virtual image type. The virtual images are cached on the hypervisors, so subsequent deployments do not require the transfer of the virtual image, saving valuable time and resources. The preloaded virtual images average about 6GB in size. The amount of time this step takes is dependent on your network bandwidth and the target hypervisor.

Step three, operating system and WebSphere Application Server are configured and started. Any federation that needs to take place will take place during this step. Once this step completes your WebSphere Application Server environment is fully functional. The amount of time this step takes is dependent on your target hypervisor.

Step four, script packages are transferred using SSH, extracted and started using root authority.

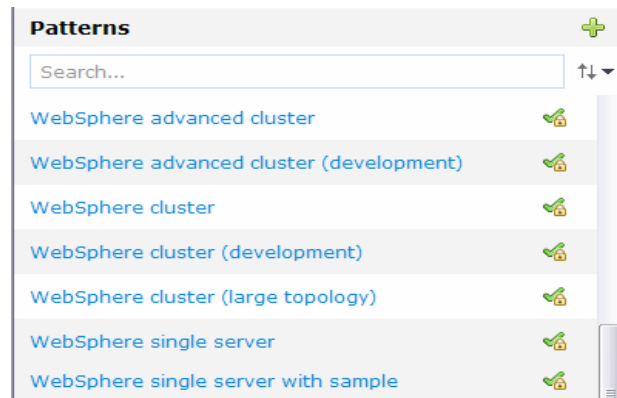
Finally, step five, an email notification is sent out to the owner of the virtual system with a status message.

Preloaded patterns

This section will provide information on the preloaded patterns that come shipped with the appliance.

Preloaded patterns (1 of 4)

- Appliance comes preloaded with a set of “best practices” patterns
- Used as is or used as a starting point when creating your own custom patterns



18

Patterns

© 2010 IBM Corporation

WebSphere CloudBurst comes preloaded with a set of “best practices” patterns. The screen capture on this slide shows a subset of the preloaded patterns. When creating your own custom pattern you should use one of the preloaded “best practices” patterns as a starting point. Again, you can do this by making a clone of the pattern. By starting with one of WebSphere 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.


All of the preloaded patterns are based on the default WebSphere Application Server virtual image. If you want to use another virtual image as your source you will need to clone the pattern, change the virtual image and deploy the new pattern.

Preloaded patterns (2 of 4)

- WebSphere single server / WebSphere single server with sample
 - 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




This slide and the next two slides take a deeper look into the makeup of the preloaded patterns. Described on this slide are the “WebSphere single server (with and without sample)” pattern and the “WebSphere cluster (development)” pattern. Besides the makeup of the pattern the thing to take away from this and the corresponding slides, is that each virtual image part requires its own virtual machine. This has ramifications in the number of IPs required and the amount of memory required to host the virtual machines.

Preloaded patterns (3 of 4)

- 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




Described on this slide are the “WebSphere cluster” pattern and the “WebSphere cluster (large topology)” pattern.

Preloaded patterns (4 of 4)

- WebSphere advanced cluster
 - 1 deployment manager
 - 4 custom nodes
 - 2 On demand routers
 - 2 IBM HTTP server
 - Total of 9 virtual machines



- WebSphere advanced cluster (development)
 - 1 deployment manager
 - 2 custom nodes
 - 1 On demand router
 - Total of 4 virtual machines



Described on this slide are the “WebSphere advanced cluster” pattern and the “WebSphere advanced cluster (development)” pattern.

Summary

- Pattern creation
 - Drag-and-drop assembly
- Pattern deployment
 - 20-30 minute deployment times of a complete WebSphere Application Server environment
 - Only administrative password required for deployment
- Appliances comes preloaded with industry “best practice” patterns

Creating a WebSphere Application Server topology is no longer an extensive drawn out manual process. With WebSphere CloudBurst, comes the Pattern Editor, which allows you to graphically drag and drop your WebSphere Application Server parts forming a WebSphere Application Server pattern which can be used for a later deployment. Unless required by your specific scenario a pattern requires only that you specify an administrative password for deployment.

Finally, besides making it extremely easy to create your own patterns WebSphere CloudBurst takes this one step further by providing you with a set of industry “best practice” patterns that you can use immediately as-is.



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 email feedback:

mailto:iea@us.ibm.com?subject=Feedback_about_CB20_Patterns.ppt

This module is also available in PDF format at: [./CB20_Patterns.pdf](http://CB20_Patterns.pdf)

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



Trademarks, disclaimer, and copyright information

IBM, the IBM logo, ibm.com, CloudBurst, and WebSphere are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of other IBM trademarks is available on the Web at "[Copyright and trademark information](http://www.ibm.com/legal/copytrade.shtml)" at <http://www.ibm.com/legal/copytrade.shtml>

THE INFORMATION CONTAINED IN THIS PRESENTATION IS PROVIDED FOR INFORMATIONAL PURPOSES ONLY. WHILE EFFORTS WERE MADE TO VERIFY THE COMPLETENESS AND ACCURACY OF THE INFORMATION CONTAINED IN THIS PRESENTATION, IT IS PROVIDED "AS IS" WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. IN ADDITION, THIS INFORMATION IS BASED ON IBM'S CURRENT PRODUCT PLANS AND STRATEGY, WHICH ARE SUBJECT TO CHANGE BY IBM WITHOUT NOTICE. IBM SHALL NOT BE RESPONSIBLE FOR ANY DAMAGES ARISING OUT OF THE USE OF, OR OTHERWISE RELATED TO, THIS PRESENTATION OR ANY OTHER DOCUMENTATION. NOTHING CONTAINED IN THIS PRESENTATION IS INTENDED TO, NOR SHALL HAVE THE EFFECT OF, CREATING ANY WARRANTIES OR REPRESENTATIONS FROM IBM (OR ITS SUPPLIERS OR LICENSORS), OR ALTERING THE TERMS AND CONDITIONS OF ANY AGREEMENT OR LICENSE GOVERNING THE USE OF IBM PRODUCTS OR SOFTWARE.

© Copyright International Business Machines Corporation 2010. All rights reserved.