

IBM WebSphere CloudBurst Appliance V2.0

WebSphere Process Server Hypervisor Edition



This presentation will discuss IBM WebSphere® Process Server Hypervisor Edition and WebSphere CloudBurst™.



Table of contents

- WebSphere Process Server
- Importing the virtual images
- Pattern parts
- Patterns
- Applying fixes

This presentation will first cover a brief overview of WebSphere Process Server, followed by a detailed discussion of WebSphere Process Server Hypervisor Edition.



Section

WebSphere Process Server

3

WebSphere Process Server Hypervisor Edition

© 2010 IBM Corporation

This section will discuss an overview of WebSphere Process Server.

WebSphere Process Server

- IBM WebSphere Process Server is a high-performance business process automation engine to help form processes that meet your business goals
 - Built on top of WebSphere Application Server
 - Supports integration based on SOA and open standards
- Automates business processes that span people, workflows, applications, systems, platforms, and architectures
- Business Process Applications
 - Involve business processes
 - Include tasks requiring human involvement
 - Incorporate business rules
- WS-BPEL specification compliant

IBM WebSphere Process Server is a high performance business process automation engine. It is built on top of IBM WebSphere Application Server and supports business integration based on

Service oriented architecture standards. WebSphere Process Server is an engine that can automate business processes. A business process is a flow of tasks that can involve people, workflows, applications, systems, platforms, and architectures. Various parts of business processes can be automated, while others involve human tasks, such as approving requests for travel as an example. A business process application involves business processes, tasks that can be automated using business rules, and tasks that require human involvement. WebSphere Process Server can run business process applications, and is a WS-BPEL specification compliant business process engine.

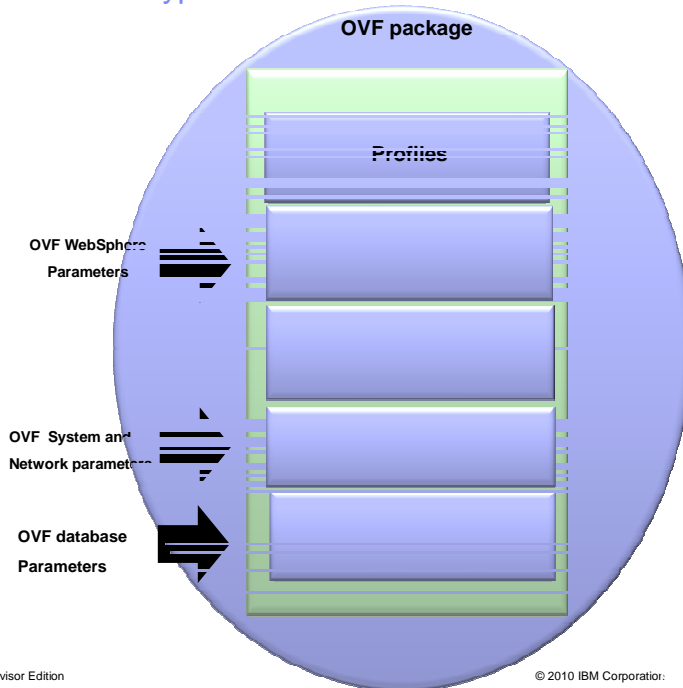
Using WebSphere Process Server and WebSphere CloudBurst

- WebSphere Process Server virtual image available for WebSphere CloudBurst 2.0
- Create WebSphere Process Server pattern
- Edit WebSphere Process Server pattern
- Specify deployment parameters
- Deploy WebSphere Process Server to virtual system
- Access WebSphere Process Server virtual system parts in their virtual machines

When building business process applications for deployment on WebSphere Process Server, tools are required. WebSphere Integration Developer can be used to build and test business process applications. Once a business process application is designed, developed, and tested, it can be deployed on WebSphere Process Server. WebSphere Process Server Hypervisor Edition can be deployed and managed with WebSphere CloudBurst 2.0. You can use WebSphere CloudBurst 2.0 to import the WebSphere Process Server virtual image files into the catalog. You can then create, edit, and work with WebSphere Process Server based patterns, and deploy virtual systems with installed instances of WebSphere Process Server to run business process applications. You manage WebSphere Process Server patterns and deployments in the same manner as you deploy and manage WebSphere Application Server deployments when working with WebSphere CloudBurst. Initially you import the WebSphere Process Server virtual image into the WebSphere CloudBurst catalog. Once that is available, you can create and edit patterns of WebSphere Process Server components using the WebSphere CloudBurst pattern editor. You can import script packages that are used to deploy WebSphere Process Server business process applications into the WebSphere CloudBurst catalog in the same manner as you can for WebSphere Application Server applications. The business process applications can be deployed into a running instance of WebSphere Process Server at the time the process server instance is deployed into the private cloud. WebSphere Process Server pattern parts are deployed into the virtual system into virtual machines on the Hypervisors that are configured for use by WebSphere CloudBurst.

IBM WebSphere Process Server Hypervisor Edition

- Preinstalled, configured, and tuned
- Open Virtualization Format standard packaging
- WebSphere Process Server 6.2 and 7.0
- Two virtual image formats:
 - Stand-alone image
 - VMware ESX, ESXi 3.5.x and above only
 - WebSphere CloudBurst image



6

WebSphere Process Server Hypervisor Edition

© 2010 IBM Corporation.

The WebSphere Process Server Hypervisor Edition comes pre-packaged for use with WebSphere CloudBurst Appliance.

The package includes the operating system and the WebSphere Process Server patterns and pattern parts that are necessary to work with WebSphere Process Deployments.

The virtual images are available as stand-alone Hypervisor edition images that can be used directly with a Hypervisor, and as images that can be imported into the WebSphere CloudBurst catalog.

WebSphere Process Server Hypervisor Edition version 6.2 and version 7.0 are available.

What is included in the WebSphere Process Server Hypervisor Edition

- WebSphere Process Server Hypervisor Edition images
 - WebSphere Process Server Hypervisor Edition V7.0 AIX®(64-bit)
 - WebSphere Process Server Hypervisor Edition V7.0 Novell SUSE Linux® Enterprise Server for x86 (32-bit)
 - WebSphere Process Server Hypervisor Edition V7.0 for Novell SUSE Linux Enterprise Server for System z (64-bit)
 - WebSphere Process Server Hypervisor Edition V6.2 AIX(64-bit)
 - WebSphere Process Server Hypervisor Edition V6.2 Novell SUSE Linux Enterprise Server for x86 (32-bit)
- What the image contains
 - Pre installed guest operating system (SUSE 11 or AIX 6.1)
 - DB2® Enterprise Server Edition version 9.5 for 64-bit or DB2 Workgroup Server Edition version 9.5 for 32-bit
 - Set of patterns that represent common WebSphere Process Server deployment topologies
 - Set of parts that can be used to configure your own custom topology

WebSphere Process Server Hypervisor Edition is available for x86, z/VM and Power platforms.

64-bit versions of the Hypervisor edition are available to be deployed with AIX 6.1 on IBM Power, and SUSE Linux Enterprise Server on System z. 32-bit SUSE Linux Enterprise server editions are available for ESX Hypervisors.

Each of the AIX and SUSE Linux virtual images are available for WebSphere Process Server version 7.0. Virtual images for WebSphere Process Server version 6.2 are currently available for PowerVM and x86 platforms.

The images contain pre-installed guest operating systems, with pattern parts to build a WebSphere Process Server topology.

The pattern parts can be used individually, and can be deployed as part of three pre-defined patterns.

DB2 Enterprise server 9.5 pattern parts are packaged with the AIX and System z virtual images, and DB2 Workgroup server 9.5 pattern parts are packaged with the x86 virtual images.

Importing the virtual images

This section will discuss uploading the WebSphere Process Server Hypervisor Edition virtual images to WebSphere CloudBurst.

Getting started

- Prerequisite requirements of the download target machine
 - Java™ 6 installed
 - 10 gig storage
- Actions:
 - Download the package
 - Unpack the package
 - Invoke image loader (GUI launcher or command line)
 - Supply properties

There is a special image loading tool that is provided with WebSphere Process Server Hypervisor Edition virtual images to help load the virtual image into the WebSphere CloudBurst catalog.

The virtual image is provided in a TAR archive, and that archive includes the image loading tool.

The virtual image archive should be downloaded to a machine on the local network that is hosting the WebSphere CloudBurst appliance.

Scripts are included in the archive to launch the image loading tool. The image loading tool requires Java 6 runtime to be installed on the computer loading the image to the appliance.

When using the image loader tool to load the virtual image files onto the WebSphere CloudBurst appliance, three default patterns are defined as part of the loading process.

Package format

- WebSphere_Process_Server_Hypervisor_Edition_V<version_number>_<guest_operating_system_type>.ova
 - Virtual image that is loaded into the appliance
- loader.bat - command to start the image loading process for Windows®
- loader.sh - command to start the image loading process for UNIX®
- imageLoader folder - Folder that contains the scripts and log files folder

imageLoader	7 items	folder
images	10 items	folder
ibs	7 items	folder
logs	3 items	folder
resources	39 items	folder
scripts	3 items	folder
WebSphere_Process_Server_Hypervisor_Edition_V70_SLES_X86.ovf	30.8 KB	XML document
cb.p12	3.7 KB	PKCS#12 certificate t
WebSphere_Process_Server_Hypervisor_Edition_V70_SLES_X86.ova	6.3 GB	tar archive
loader.bat	993 bytes	plain text document
loader.sh	925 bytes	shell script

This slide shows an overview of the contents of the virtual image archive for WebSphere Process Server Hypervisor Edition.

In the archive is the OVA file to load into the WebSphere CloudBurst catalog. Also included are loader script files to launch the loader tool to import the image into the catalog. Supporting files are included that are used by the image loader tool. Separate script files are included for running in a Windows or Linux environment.

Image loader

- Use the loader command in a command line without GUI
 - `loader.sh -silent -hostname <hostname> -userID <username> -password <password> -language <language> -acceptLicense <true or false>`
 - Virtual image loaded into catalog along with the three predefined patterns
 - Image reference number is used to distinguish between two virtual images
- Use the loader GUI
 - Invoked by using `loader.sh` or `loader.bat` with no arguments
 - Not available on z/VM

```

CMLHV0008I: Welcome to Image Loader, a utility used to load the virtual image on
to the appliance and create predefined patterns for the image.
CMLHV0009I: Image Loader is verifying that you have an active network connection
.
CMLHV0010I: Image Loader established a connection to the following appliance: ai
mcp210.austin.ibm.com
CMLHV0011I: Starting to load the virtual image on to WebSphere CloudBurst Applia
nce.
CMLHV0025I: The WCA version is 2.0.0.0-22581.
CMLHV0026I: Begin to import virtual image WebSphere Process Server 6.2.0.2 for 3
2 bit SUSE 11 168.
CMLHV0015W: The virtual image you are loading already exists on the appliance.
CMLHV0021I: Creating predefined patterns for the virtual image.
CMLHV0023I: Image Loader is creating the following predefined pattern on the app
pliance.   Single Server for WebSphere Process Server V6.2.0.2 on 32 bit SUSE 11
168
CMLHV0023I: Image Loader is creating the following predefined pattern on the app
pliance.   Simulated Scalable Environment for WebSphere Process Server V6.2.0.2 o
n 32 bit SUSE 11 168
CMLHV0023I: Image Loader is creating the following predefined pattern on the app
pliance.   Scalable Environment for WebSphere Process Server V6.2.0.2 on 32 bit S
USE 11 168

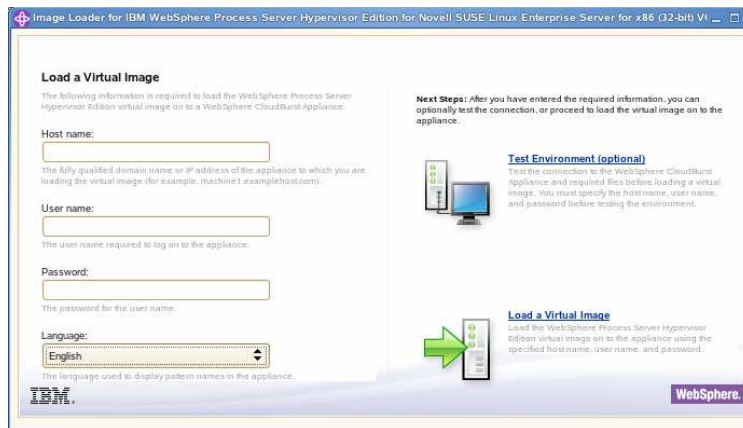
```

The image loading tool can be used either graphically or in a command line format. For z/VM, the command line format must be used.

Running the loader script file without any parameters launches the GUI by default.

Image loader

- Use the loader GUI
 - Invoked using loader.bat/sh without any arguments



When you invoke the loader script without any arguments, the GUI is opened.

The image loader GUI has entry fields for the IP address of the WebSphere CloudBurst appliance, and user and password credentials.

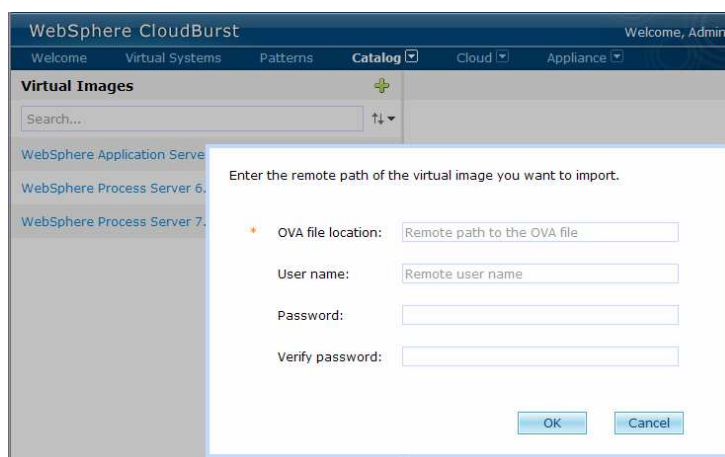
Entering this information is mandatory to get the image uploaded to the WebSphere CloudBurst catalog.

Once you have entered the host, user name, and password information, you can optionally click the test environment link to test the connectivity to your appliance.

When you are ready, you can click the load a virtual image link, and the process of uploading the virtual image to the WebSphere CloudBurst appliance, and creating the three default patterns is started. A progress bar will display the progress of the upload and pattern creation.

Load image using WebSphere CloudBurst console

- Use the WebSphere CloudBurst console
 - Catalog → Virtual Images → Add new Virtual image
 - Provide a path to remote OVA file. You need to configure an http server on the image machine to serve this file
 - Provide username and password of the http server if security is on
 - Note: **Images loaded using the console will not have predefined patterns**



13

WebSphere Process Server Hypervisor Edition


© 2010 IBM Corporation

You can also load the WebSphere Process Server virtual image using the default method in the WebSphere CloudBurst administrative console, or the command line interface.

Loading the virtual image for WebSphere Process Server Hypervisor Edition using this method will not create the three predefined patterns in the pattern editor.

Loader logs

- If using image loader, check for logs under
 - Imageloader/logs/imageloader.log
- You can also check the cloudburst logs by logging into console
 - Appliance → Troubleshooting → Logging



When using the image loader tool to upload the virtual image to the WebSphere CloudBurst catalog, logging information is captured in the imageloader/logs/imageloader.log file at the location where you started the image loader script file.

Additional logs are available from the WebSphere CloudBurst appliance in the troubleshooting page.

Pattern parts

This section will discuss the pattern parts that are available with the WebSphere Process Server Hypervisor Edition virtual images.

WebSphere Process Server Hypervisor Edition - parts

- Database
- Extendable nodes
- Stand-alone server
- Deployment manager
- Custom nodes
- Proxy server
- Basic function nodes
- Basic function control node
- Full function control node

WebSphere Process Server Hypervisor Edition has nine pattern parts.

There are several simpler pattern parts, including a database part, extendable nodes, a stand-alone server, a deployment manager, custom nodes and a proxy server.

There are three pattern parts that are more advanced parts, including a basic function node, a basic function control node, and a full function control node.

There are best practices for working with individual pattern parts when editing a pattern that are discussed later in this presentation.

Parts description

- Stand-alone server
 - Contains WebSphere process server installation with stand-alone server instance and DB2 installed
- Extendable node
 - Contains just the installation of WebSphere Process Server
- Database part
 - Contains the installation of DB2 database.
 - Can be configured for deployment manager or for the clustered topology
- Custom node
 - Provides an unconfigured node for use with a deployment manager or a part with control nodes
- Proxy server
 - Contains a proxy server instance used for load distribution

The stand-alone server part contains an operating system with a pre-installed WebSphere Process Server, with a single server profile created and a DB2 database installed.

The extendable node pattern contains the operating system with only the WebSphere Process Server part installed.

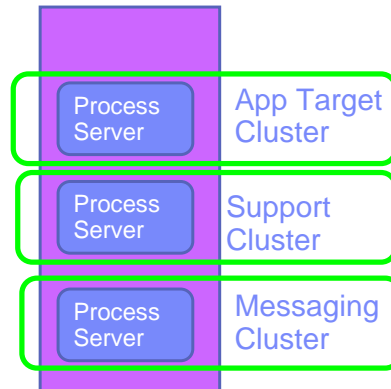
The database pattern part contains only the operating system with a DB2 database instance installed.

The custom node pattern part contains an unconfigured node that can be used in conjunction with a deployment manager, or a pattern part that contains a deployment manager such as one of the control nodes.

The proxy server pattern part contains an operating system deployment with a single proxy server installed to handle load distribution for your deployment.

Parts description - Basic function nodes

- Basic function nodes



The WebSphere Process Server basic function node pattern part is the first part that is more complex than the other simpler pattern parts.


This pattern part contains three server instances deployed in a single virtual machine.

This pattern part cannot be used on its own, it must be deployed as part of a larger pattern that contains a deployment manager in order for it to be of any use.

This pattern part can be used when building a custom pattern with clustering support for WebSphere Process Server deployments.

Parts description - Basic function control nodes

- Basic function control node
 - Remote messaging, remote support Process Server deployment pattern in a single virtual machine.
 - No DB2 server and no proxy server



The second of the more complex pattern parts is the basic function control node.


The word control in the description of this pattern part helps to identify it as a pattern part that contains a deployment manager.

This pattern part is an extension of the basic function node, with a deployment manager as part of the same virtual machine.

This pattern part does not contain a DB2 database or a proxy server as part of its topology.

Parts description - Full function control nodes

- Full function control node
 - Consists of all the Process Server components running in a single virtual machine organized similarly to a production environment
 - Contains deployment manager, a database instance, a proxy server and cluster members



20

WebSphere Process Server Hypervisor Edition

© 2010 IBM Corporation

The final complex pattern part that is provided with WebSphere Process server virtual image is the full function control node.

The use of the word control in the description of this pattern part helps you identify that it is a pattern part that contains a deployment manager.

This pattern part contains a DB2 server instance, a deployment manager, a Proxy Server instance, and a basic control node with three servers.

All of the included components run on a single virtual machine instance.

This pattern part can be used with a basic function node to begin to build a scalable deployment of WebSphere Process Server.

Patterns

This section will discuss the pre-defined patterns that are included with the WebSphere Process Server Hypervisor Edition virtual images.

WebSphere Process Server Hypervisor Edition - Patterns

- WebSphere Process Server comes preconfigured with three patterns
 - Single Server for WebSphere Process Server
 - Simulated Scalable Environment for WebSphere Process Server
 - Scalable Environment for WebSphere Process Server

- Use pattern parts to create custom patterns

WebSphere Process Server Hypervisor Edition has three preconfigured patterns for use with WebSphere CloudBurst.

The patterns are created when the virtual image is added to the WebSphere CloudBurst catalog using the image loading tool.


The three patterns that are defined by the WebSphere Process Server image loading tool include a single server WebSphere Process Server pattern,

a simulated scalable environment pattern that contains a simple gold topology pattern, and a scalable environment pattern that is essentially a scalable

gold topology WebSphere Process Server pattern.

Patterns – Single server


- Uses the stand-alone server part
- Deployed to one virtual machine



Stand-alone profile

Single Java Virtual Machine (JVM)

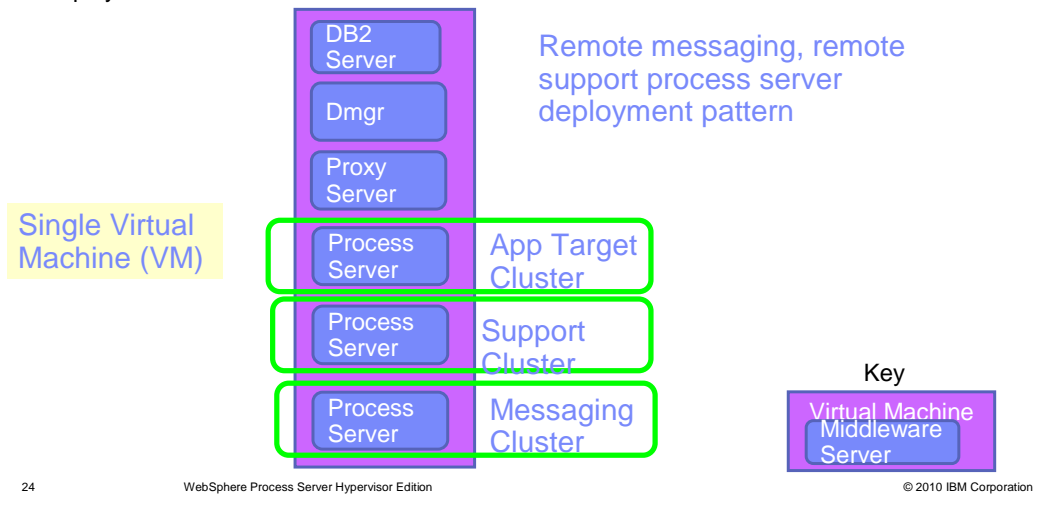
Key



The single server WebSphere Process Server pattern is a single virtual machine that is composed of the operating system, a DB2 server, and a WebSphere Process Server installed with a single stand-alone server instance defined.

Patterns – Simulated scalable environment

- Uses the full function control node part
- Deployed to one virtual machine

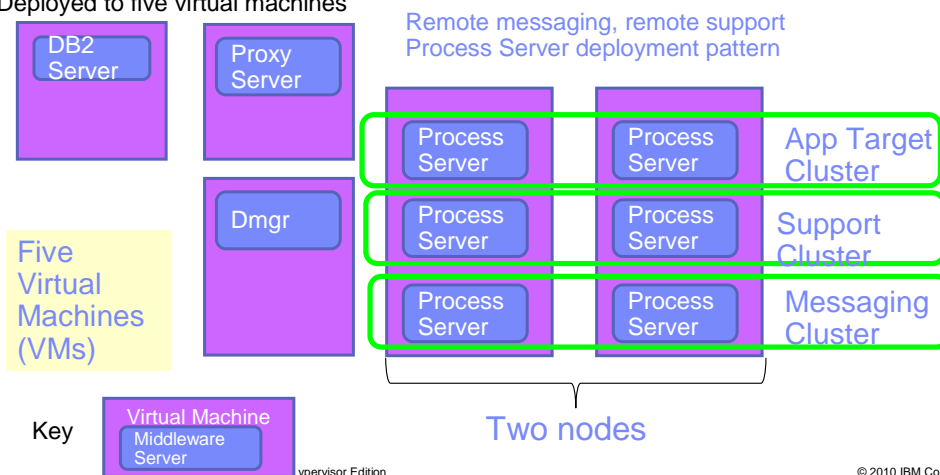


The simulated scalable environment is a pattern that contains the full function control node pattern part. This pattern part includes a DB2 server, a deployment manager, a proxy server, and three WebSphere Process Server clusters. There is a cluster defined that includes a server for user applications, a cluster defined that runs support applications such as CEI and Business Rules Manager and a cluster defined where the messaging infrastructure is configured. All of the components in this pattern are deployed in a single virtual machine, and the pattern is ideal for testing a simple scalable environment.

It is not a suitable production pattern as it does not include failover support in the defined topology.

Patterns – Scalable environment

- Uses the deployment manager, proxy server, database, full function control node and two basic function node parts
- Deployed to five virtual machines



The scalable environment pattern represents a sample production WebSphere Process Server gold topology pattern.

The different pattern parts are broken into several virtual machines that are ideal for a production environment when failover support is important.

The DB2 server, the proxy server, and deployment manager are all deployed in their own virtual machines.

Two more virtual machines are defined with three clusters spanning them. Each cluster has two servers for application target, support, and messaging.

Each of the virtual machines containing the clusters has a single node agent.

Custom patterns – Some guidelines

- Two stand-alone servers in one pattern are two individual entities. No communication occurs between them
- Stand-alone server and deployment manager in one pattern
 - Stand-alone not federated to deployment manager
 - Two individual entities; there is no interaction between these two parts
- Use database part with either basic function control node or a deployment manager
 - Don't add more than database part
- Basic function node and custom node parts should be only used in a pattern that contains a deployment manager or a control node
- Do not include more than one control node in a pattern
- Do not use deployment manager with a control node

There are some guidelines that should be followed when creating customized patterns for WebSphere Process Server deployments.

You should not have more than one stand-alone server in a pattern. If two stand-alone servers are deployed in a pattern, they are deployed as single entities that do not communicate with each other, and provide no advantages to being in a single virtual system.

If a stand-alone server is deployed in the same pattern as a deployment manager, the stand-alone server is not federated to the deployment manager. They are deployed as two separate entities.

The DB2 database server instance should be deployed in a pattern that contains a basic function control node, or a stand-alone deployment manager.

Adding more than one DB2 server instance to your pattern does not provide extra database processing power to your deployment.

The basic function node and custom node parts should be deployed in part of a pattern that contains a deployment manager, such as the basic or full function control nodes, or deployment manager pattern parts. Without a deployment manager, there is nothing to federate the basic function nodes to, and the nodes will just unnecessarily consume resources on your Hypervisor.

More than one deployment manager should not be added to your pattern and does not provide extra computing power to your topology. This includes any pattern parts that contain a deployment manager such as the control nodes, or extra single deployment manager parts.

Applying fixes

This section will discuss applying fixes to WebSphere Process Server deployments with WebSphere CloudBurst.

Applying fixes

- Fix packs, interim fixes, feature packs for WebSphere Process Server
 - Use installation manager
- Fix packs, interim fixes for WebSphere Application Server
 - Use update installer
- Operating system fixes
 - RPM
- GUI selection for applying emergency fixes with WebSphere CloudBurst
 - Installation manager
 - Update installer
 - RPM package manager

The IBM Installation manager and update installer are both installed on the virtual system during deployment of WebSphere Process Server.

You can use the installation manager and update installer WebSphere Process Server has some particular requirements for applying fixes to deployments using the.

I-Fixes and fix packs that contain modifications to the WebSphere Process Server components of a deployed virtual system must be applied using the installation manager.

I-Fixes and fix packs that contain fixes for the underlying WebSphere Application Server components of the WebSphere Process Server deployment are applied using the installation manager.

When applying emergency fixes with WebSphere CloudBurst, additional options are displayed when applying the fixes. You can choose to apply the emergency fix using the update installer or installation manager when applying an emergency fix to a virtual system.

There is also the option to use the RPM Package Manager to install fixes related to the operating system.

Summary

This section will summarize the WebSphere Process Server Hypervisor Edition with WebSphere CloudBurst presentation.

Summary

- WebSphere Process Server Hypervisor Edition images available for WebSphere CloudBurst 2.0
- Load WebSphere Process Server Hypervisor Edition images using the image loader tool
- WebSphere Process Server in CloudBurst
 - Create, clone, and edit patterns
 - Deploy WebSphere Process Server virtual images
 - Deploy WebSphere Process Server applications using script packages

WebSphere Process Server Hypervisor Edition image version 6.2 and 7.0 are available for WebSphere CloudBurst 2.0.

The WebSphere Process Server Hypervisor Edition images can be loaded onto the WebSphere CloudBurst appliance using an image loading tool. Loading the images using the image loader tool generates the default patterns for WebSphere Process Server.

The WebSphere Process Server image provides pattern parts that can be used to create custom patterns. These include WebSphere Process Server instances, DB2 database support, a Proxy server, and a deployment manager. WebSphere Process Server basic function nodes contain three clusters that become targets for user applications, support applications and the messaging infrastructure.



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 WPS.ppt](mailto:iea@us.ibm.com?subject=Feedback%20about%20CB20%20WPS.ppt)

This module is also available in PDF format at: [../CB20_WPS.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, AIX, AIX 6, CloudBurst, DB2, IBM, 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. Java, 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.

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.