

IBM Software Group

WebSphere® Process Server V6.0.1



Clustering – Administration and Configuration Tutorial Overview



@business on demand.


© 2006 IBM Corporation
Updated June 16, 2006

This presentation will provide an overview of the tutorials for configuring a highly available WebSphere Process Server.





Goals

- To present a high level overview of the tutorial for configuring a highly available WebSphere Process Server environment.




WPS V6.0.1 – Clustering Admin Tutorial Overview © 2006 IBM Corporation 2

The process of configuring this topology will touch on all of the basic tasks and illustrate administrative issues that may not be readily apparent.



Agenda


- The Databases
- The Cell
- The AdminMECluster
- The AdminCluster
- The WPSMECluster
- The WPSCluster



WPS V6.0.1 – Clustering Admin Tutorial Overview © 2006 IBM Corporation 3

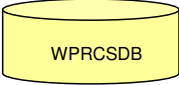
This presentation will briefly discuss the components that need to be installed and configured to create a highly available WebSphere Process Server topology. Note: highly available does not necessarily mean continuously available.

The major components are the databases, the Cell, the Administration cluster pair and the WebSphere Process Server cluster pair.

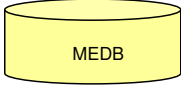
IBM Software Group 

The Databases

- The databases are created before the profiles and nodes.



WPRCSDB



MEDB



- **WPRCSDB**
 - ▶ General database for WebSphere Process Server
 - Event tables for CEI
 - BPC tables for Business Process Choreographer
 - Relationships
 - Business Rules
 - Selectors
 - ESB – MSGLOG table
- **Messaging Engine DB**
 - ▶ Tables for the messaging engines
 - A unique table and schema for each messaging engine defined

WPS V6.0.1 – Clustering Admin Tutorial Overview © 2006 IBM Corporation 4



Note that you should create the database named WPRCSDB before creating the Deployment Manger node.

The decision to use two databases rather than many simplifies the topology for the purpose of the tutorial.

Creation of the databases and tables involves customizing and running scripts that are provided as part of the WebSphere Process Server runtime. Step by step instructions are provided in the tutorial.

Setting up the Cell

- There will be a single Cell
 - ▶ dCell
- Two Systems
 - ▶ At least 2 systems are required for the JVMs and memory requirements.
- Three nodes
 - ▶ Deployment Manager Node
 - ▶ Node 1 (custom)
 - ▶ Node 2 (custom)
- Use the Profile Creation Wizard to create the nodes
 - ▶  Deployment Manager node
 - ▶  Custom nodes
- Configure datasources and WebSphere variables

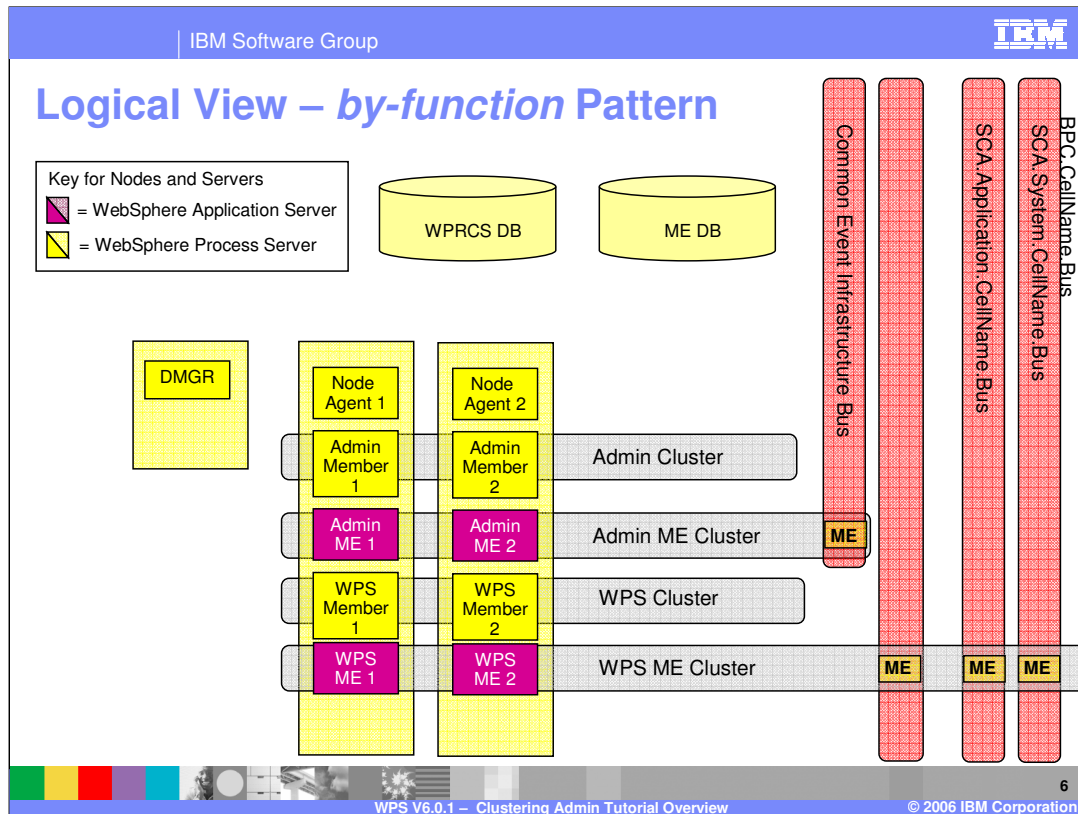
5

WPS V6.0.1 – Clustering Admin Tutorial Overview © 2006 IBM Corporation

The topology being constructed will have eleven JVMs. On a single Windows® XP system, if all the JVMs are started, more than 3GB of ram will be consumed. For this reason it is required or strongly suggested that a minimum of 2 machines be used. The Deployment Manager node can be located on a third system or co-located with a 'custom' node on one of 2 systems.

The Profile Creation Wizard, which is part of the WebSphere Process Server installation, will be used to create the Deployment Manager node and 2 Custom nodes.

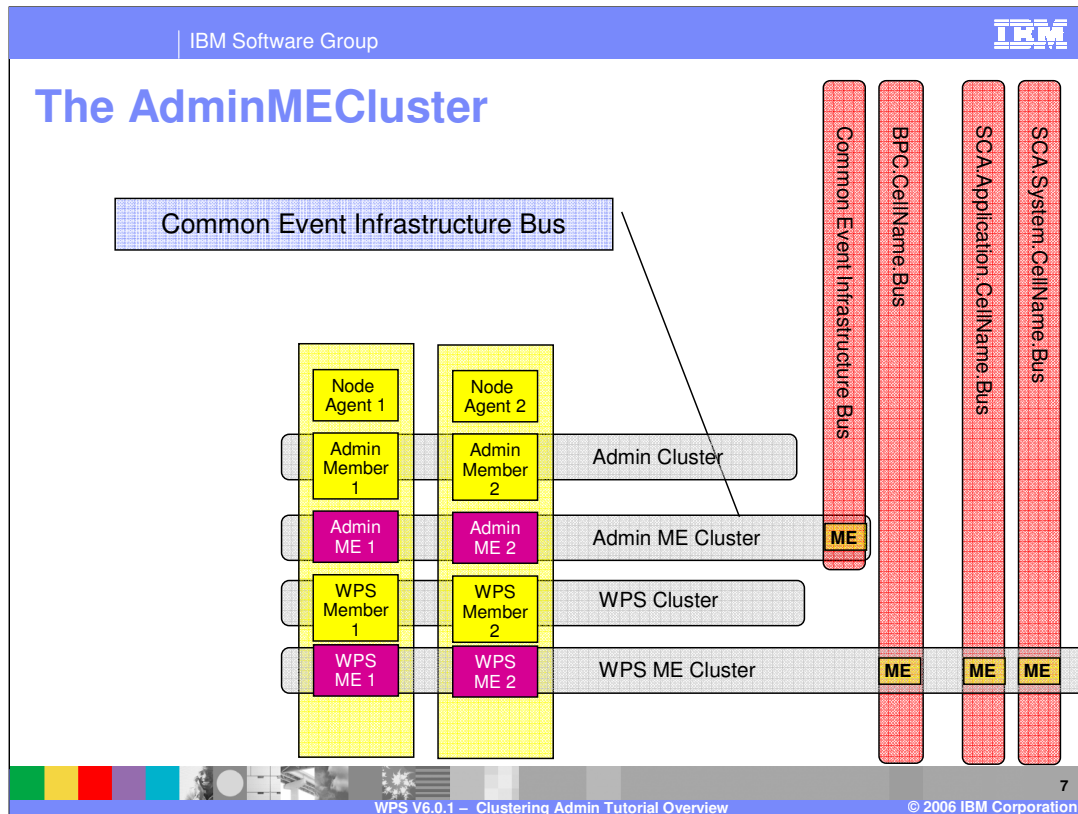
After the creation of the nodes, there will be some global, system level entities that need to be configured before proceeding, such as datasources and WebSphere environment variables.



This is the logical view of the *by-function* pattern, the topological deployment pattern which was presented in the previous *Clustering Concepts* discussion. This is the topology that the tutorial will assist you in configuring.

Each cluster has a specific function within the overall system and therefore the configuration for each must be customized to meet specific functional requirements.

This overview will touch on the key tasks involved in configuring each cluster, the step-by-step details will be in the tutorial.



Focus for a moment on the AdminMECluster.

What are the responsibilities for this server cluster?

This cluster is the 'highly available messaging server' for the Common Event Infrastructure component.

IBM Software Group

Creating and Configuring the AdminMECluster

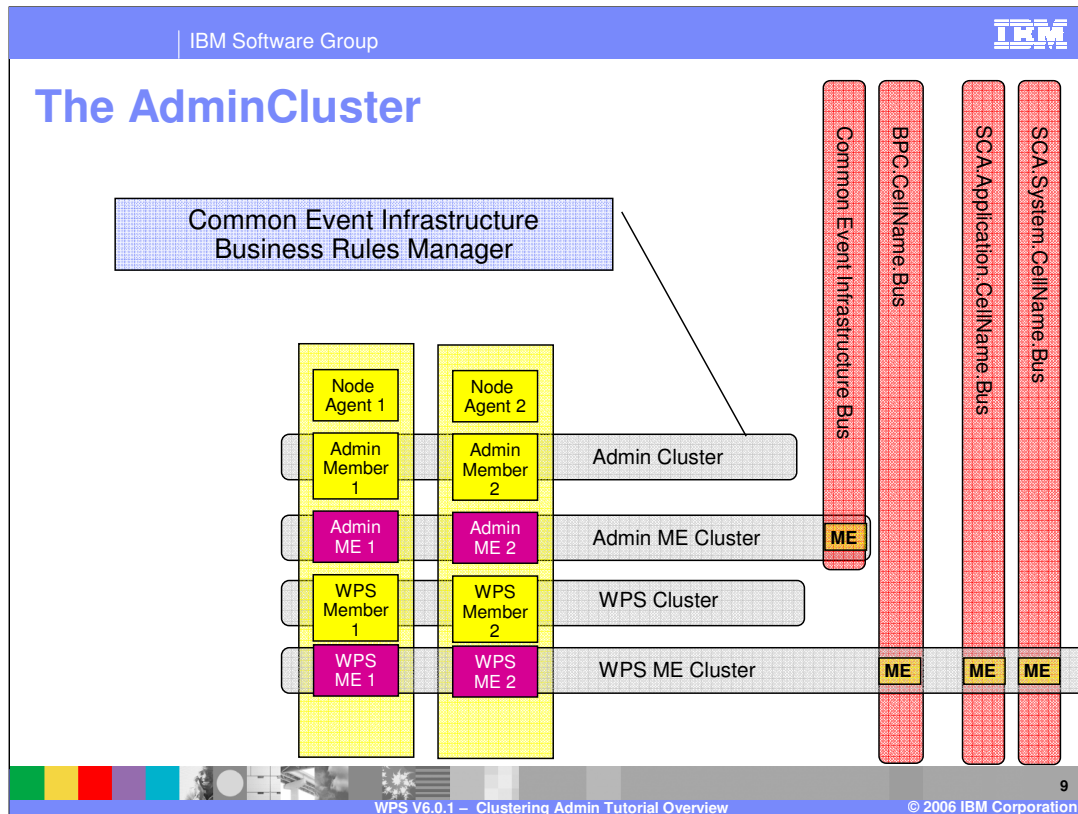
- This must be the first server cluster created.
 - ▶ The application server clusters (WPSCluster and AdminCluster) depend on the Common Event Infrastructure
 - The AdminMECluster being create at this step is the *highly available* 'messaging server' for the Common Event Infrastructure
- Create a new cluster with one cluster member
 - ▶ Server type = default (WebSphere Application Server)
 - ▶ Name = AdminMECluster
- Create and configure the JDBC™ resources required for the Message Engine.
- Create and configure the service integration bus for the Common Event Infrastructure.
 - ▶ CommonEventInfrastructure_Bus
- Add the AdminMECluster to the newly created service integration bus
 - ▶ This creates the Message Engine.
- Configure the newly create Message Engine.
 - ▶ Set the schema name
 - ▶ Set the datastore using the JDBC datasource created above

WPS V6.0.1 – Clustering Admin Tutorial Overview
© 2006 IBM Corporation
8

Shown here are the major steps necessary for creating the Administration Message Engine Cluster.

Notice that the server type is a basic WebSphere Application Server. Since the sole purpose of this cluster is to host the message engines for the Administration Cluster, there will be no WebSphere Process Server application components and the simpler, leaner WebSphere Application Server can be used.

The cluster will be “grown”, or populated with additional servers, after both the AdminMECluster and the AdminCluster have been created and configured. Since all servers in a cluster must be functionally equivalent, the strategy employed here is to create and configure the server completely, then use it as a template for the additional servers.



Next focus on the AdminCluster.


What are the responsibilities for this server cluster?

This cluster will host the applications used by the Common Event Infrastructure and the Business Rules Manager components. It is the 'administrative' cluster, which is also available to host other administrative applications that may be needed for the system.

Since the AdminCluster depends on the AdminMECluster, it must be created and configured after the AdminMECluster.

Notice that there will be no service integration bus or message engine to configure and the server type will be WebSphere Process Server.

Also notice that there will be application components for the Common Event Infrastructure and Business Rules Manager to install.

IBM Software Group 

Creating and configuring the AdminCluster


- Create a new cluster with one cluster member
 - ▶ Nothing special
 - ▶ Server type = defaultProcessServer
 - ▶ Name = AdminCluster
- Create and configure the Common Event Infrastructure components
 - ▶ Tailor response file → generate Scripts → tailor generated scripts →run scripts (for tables and datasources)
 - ▶ Install/deploy the Common Event Infrastructure Server application
 - Using wsadmin, install event-application.ear
 - ▶ Install/deploy the Common Event Infrastructure MDB application
 - Using wsadmin, install event-message.ear

WPS V6.0.1 – Clustering Admin Tutorial Overview 10 © 2006 IBM Corporation

Setting up the AdminCluster entails setting up the Common Event Infrastructure applications that will be running in this cluster.


The Common Event Infrastructure has requirements for a database and datasources, fortunately the tools to create the scripts and the instructions for using them are available as part of the WebSphere Process Server installation.

Because the scenario being developed is not the default, there is some 'post installation' configuration required. These post Common Event Infrastructure configuration steps are covered at a high level here and the details are incorporated into the tutorial.

IBM Software Group 



Creating and configuring the AdminCluster

- Re-configure the *CommonEventInfrastructure_Bus* to utilize the remote Message Engine.
 - ▶ By default the AdminCluster was added to the service integration bus.
 - In the Highly Available Message Engine configuration the AdminCluster will interact with the remote Message Engine in the AdminMECluster, through the JMS destinations.
 - ▶ Remove the AdminCluster from the *CommonEventInfrastructure_Bus*.
 - ▶ Configure the JMS Destinations to be the AdminMECluster
 - ▶ Set the JMS provider to use the CommonEventInfrastructure JMS resources
- Re-configure the Common Event Infrastructure Provider
 - ▶ Set to emit *asynchronous* events.

 WPS V6.0.1 – Clustering Admin Tutorial Overview 11 © 2006 IBM Corporation

The *CommonEventInfrastructure_Bus* is the Service Integration Bus for the Common Event Infrastructure.

The default configuration provided by the Common Event Infrastructure installation scripts, assume that the message engine will be local to the cluster being configured. Recall the discussion previously about remote vs. default (local). With this configuration the message engine will not be local, therefore the AdminCluster must be removed from the Common Event Infrastructure service integration bus and configured for remote access. The Common Event Infrastructure provider must also be set to emit events asynchronously.

Creating and configuring the AdminCluster

- **Advanced Configuration**
 - ▶ Install the Business Rules Manager
 - ▶ Do not configure to 'host SCA applications'
- **Adjust the default transmission profile for the Common Event Infrastructure Provider**
 - ▶ Known defect: 317434
 - ▶ State Observer Plugin for BPEL uses a hardcoded value at initialization time
 - Is dynamically overridden with the configured value at runtime
 - Set the JNDI value in the default *Synchronous Transmission Profile* to be the same as the one used for the AdminCluster.
- **Grow the clusters**
 - ▶ Both, the AdminMECluster and the AdminCluster

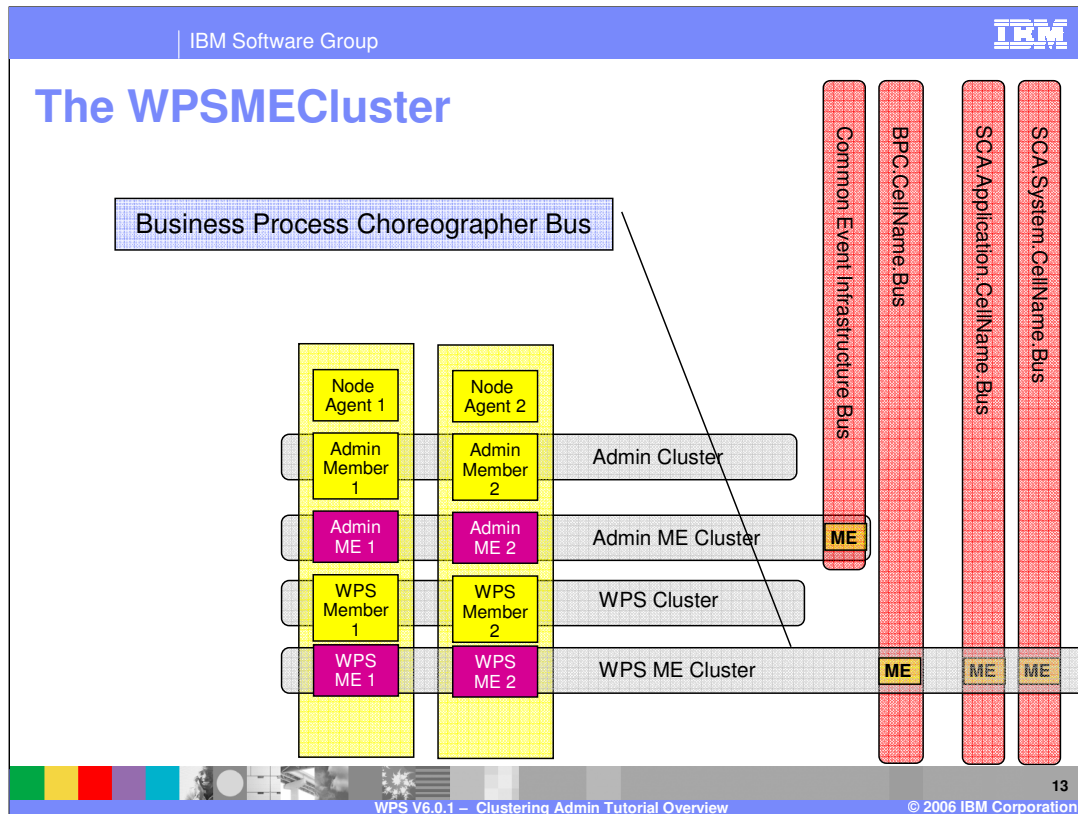
12

WPS V6.0.1 – Clustering Admin Tutorial Overview © 2006 IBM Corporation

The Advanced Configuration Wizard provides the mechanism for configuring WebSphere Process Server features on WebSphere Process Server servers or server clusters.

The adjustment to the default transmission profile is to accommodate a minor bug.

Once the server clusters are configured, additional servers can be added to the other nodes in the cell. The number of servers per node per machine will be limited by the disk and memory resources on the physical box.



Next Focus on the WPSMECluster.

What are the responsibilities for this server cluster?

This cluster is the 'highly available messaging server' for the Business Process Choreographer and Human Task Manager components and the asynchronous Service Component Architecture invocations.

Here there will be the Business Process Choreographer bus and message engines to create and configure. Note that the server type is the default WebSphere Application Server.

IBM Software Group

Creating and configuring the WPSMECluster

- Create a new cluster with one cluster member
 - ▶ Server type = default (WebSphere Application Server)
 - ▶ Name = WPSMECluster
- Configure the *Service Component Architecture* components using the *Advanced Configuration Wizard*.

- ▶ No Business Rules Manager
 - ▶ Will be hosting SCA applications in this cluster
 - ▶ Configure the local Message Engine datastores
- Create and configure the WBI.FailedEvent destination on the SCA.System bus

- ▶ to be used by the applications in the WPSCluster

WPS V6.0.1 – Clustering Admin Tutorial Overview
© 2006 IBM Corporation
14

Each cluster being configured in this topological pattern is slightly different.

The question of 'hosting' Service Component Architecture applications is really asking if there is a need to utilize the SCA Application and System buses. In the configuration being presented, the applications are in the WPSCluster and are using the Business Process Choreographer and Human Task Manager containers, which use messaging and use the SCA buses. Therefore the answer to the question presented by the Advanced Configuration Wizard, is yes, this cluster will be hosting SCA applications.

The applications in the WebSphere Process Server cluster pair and the Administration Cluster pair will both be utilizing the Failed Event Manager Service, therefore the queue destinations need to be defined for both clusters but only one cluster needs to be a member of the SCA.System bus.

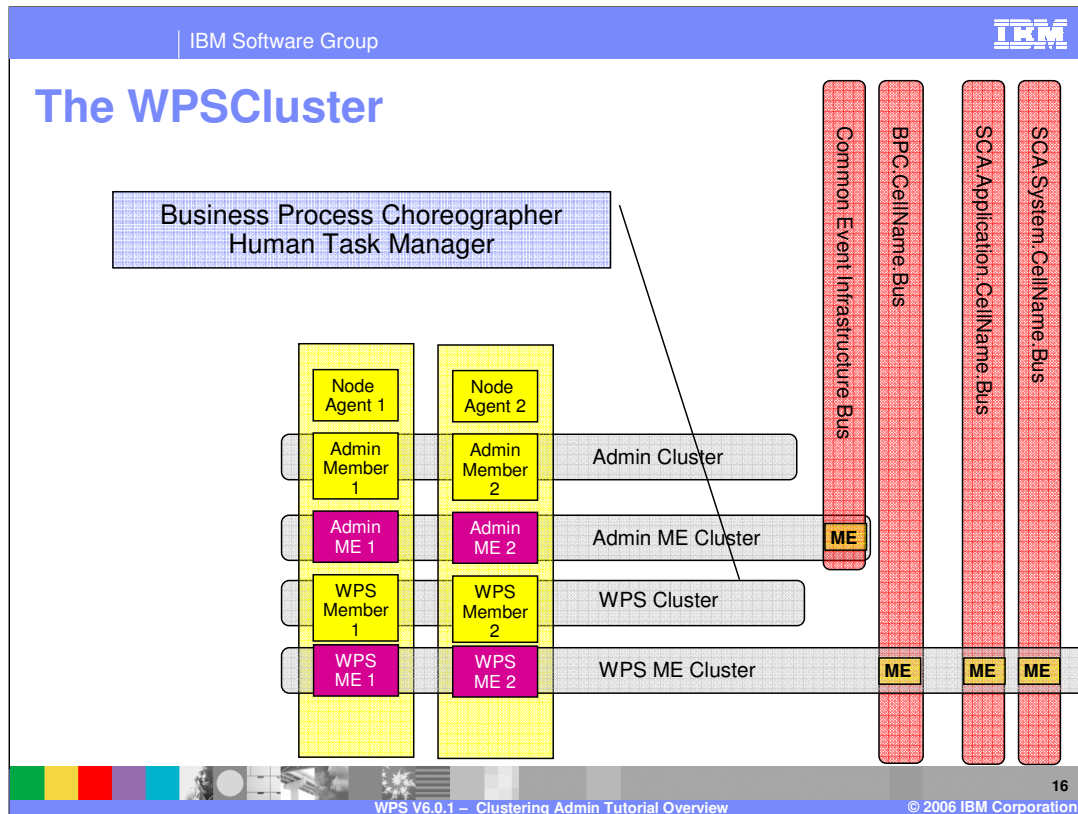
IBM Software Group

Creating and configuring the WPSMECluster

- Create the JDBC resources
 - ▶ JDBC Provider and datasource for the Message Engine database
- Create and configure the service integration bus for the Business Process Choreographer applications
 - ▶ Name = `BPC.yourCellName.Bus`
 - ▶ Configure the Message Engine for this cluster-sib
- Create the JMS destinations to be used by the Business Process Choreographer applications
 - ▶ BPEApiQueue, BPEHIdQueue, BPEIntQueue, BPERetQueue
 - ▶ HTMHIdQueue, HTMIntQueue
- Done with this cluster for the moment.

WPS V6.0.1 – Clustering Admin Tutorial Overview
15 © 2006 IBM Corporation

There are several more steps required for setting up the WebSphere Process Server message engine cluster. When everything is all in the same cluster the Business Process Choreographer setup scripts configure everything. Because in this topology the message engine is separated from the WebSphere Process Server applications, the JMS resources needed by the Business Process Choreographer and Human Task Manager need to be manually configured.



Finally, focus on the WPSCluster.

What are the responsibilities for this server cluster?

This cluster hosts the Business Process Choreographer and Human Task Manager containers and the applications that use them, plus other WebSphere Process Server components such as the Business State Machine, Business Rules, and so on.

There will be no messaging engines or service integration buses to configure.

IBM Software Group

Creating and Configuring the WPSCluster

- Create a new cluster with one cluster member
 - ▶ Nothing special
 - ▶ Server type = defaultProcessServer
 - ▶ Name = WPSCluster
- Configure the *Service Component Architecture* components using the *Advanced Configuration Wizard*.
 - ▶ No Business Rules Manager
 - ▶ Will be hosting SCA applications in this cluster
 - ▶ Configure the remote Message Engine datastores
- Create and configure the JMS resources used by the Business Process Choreography applications.
 - ▶ Connection factories
 - BPECF, BPECFC, HTMCF
 - ▶ Queues
 - BPEApiQueue, BPEHldQueue, BPEIntQueue, BPERetQueue
 - HTMHldQueue, HTMIntQueue
 - ▶ Activation Specification
 - BPEApiActivationSpec, BPEInternalActivationSpec, HTMInternalActivationSpec



Show Me

Show Me

17
WPS V6.0.1 – Clustering Admin Tutorial Overview © 2006 IBM Corporation

Notice the differences with the options for the Advanced Configuration Wizard. This cluster will be using WebSphere Process Server components but the messaging engine is hosted remotely in the WPSMECluster.

Additional work is required to setup the JMS resources used by the Business Process Choreographer and Human Task Manager containers.



Creating and configuring the WPSCluster


- Install the Business Process Choreographer Container
- Install the Human Task Manager Container
- Grow the WPSMECluster
- Grow the WPSCluster

End of Configuration

WPS V6.0.1 – Clustering Admin Tutorial Overview © 2006 IBM Corporation 18

The installation of the Business Process Choreographer and Human Task Manager containers is straightforward and wizards are provided from the WebSphere Process Server Administrative Console. There are scripts which are also available in the WebSphere Process Server installation which can be modified to meet specific requirements for silent installations and configurations.

The configuration of the containers on the WPSCluster completes the basic installation and configuration of the WebSphere Process Server environment. Additional servers can now be added to both clusters to fill out the needs of the system, keeping in mind the limitations of the physical resources available.

IBM Software Group 

Summary

- Messaging Engine clustering is what differentiates the tasks for creating a clustered WebSphere Process Server topology from a WebSphere Application Server ND V6 topology.
 - ▶ remote vs. local
 - ▶ active/standby vs. active/active
- Core functionality requires messaging
 - ▶ Business Process Choreographer
 - ▶ Common Event Infrastructure
 - ▶ Asynchronous SCA invocations
- Advanced Configuration Wizard available for configuring the SCA buses.
- Custom nature of topologies requires thorough knowledge of the messaging, application and SCA relationships.
 - ▶ Default scripts provide guidelines.

WPS V6.0.1 – Clustering Admin Tutorial Overview 19
© 2006 IBM Corporation

When creating a clustered WebSphere Process Server environment that utilizes Business Process Choreographer the key factor to consider is the messaging engine.

Messaging is at the heart of the core functionality provided by WebSphere Process Server. Business Process Choreographer, Common Event Infrastructure and Asynchronous SCA invocations all utilize the default WebSphere messaging.

There are Administrative Console tools, such as the Advanced Configuration wizard, the Business Process Choreographer Container and Human Task Manager Container wizards and the default scripts, which are provided to guide and aid in the installation and configuration process.

Working through the tutorial will provide the additional reinforcement necessary to become comfortable with all the parts and relationships discussed in this module on WebSphere Process Server clustering.

Trademarks, Copyrights, and Disclaimers

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

IBM	CICS	IMS	MQSeries	Tivoli
IBM (logo)	Cloudscape	Informix	OS/390	WebSphere
e (logo) business	DB2	iSeries	OS/400	xSeries
AIX	DB2 Universal Database	Lotus	pSeries	zSeries

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

Microsoft, Windows, Windows NT, and the Windows logo are registered trademarks of Microsoft Corporation in the United States, other countries, or both.

Intel, ActionMedia, LANDesk, MMX, Pentium and ProShare are trademarks of Intel Corporation in the United States, other countries, or both.

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

Linux is a registered trademark of Linus Torvalds.

Other company, product and 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 and/or changes in the product(s) and/or program(s) 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.

Information is provided "AS IS" without warranty of any kind. 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 (e.g., 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.

© Copyright International Business Machines Corporation 2004,2005. All rights reserved.

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.

