


IBM Software Group

# WebSphere® Process Server V6.0.1


## *Clustering – Administration and Configuration*



@business on demand.

© 2006 IBM Corporation  
Updated May 15, 2006

This presentation covers clustering administration and configuration for WebSphere Process Server version 6.0.1.

IBM Software Group 

## Goals

- Present the administrative tasks involved with setting up a clustered topology with WebSphere Process Server V6.0.1
  - ▶ The 4 cluster deployment pattern presented in the previous section will be used as the vehicle for the discussion.
- Identify the elements that need to be installed and configured.
  - ▶ Discuss the options available.
- Prepare for executing the tutorial which will provide the details for each step.
- Will not cover the installation and configuration of the Business Rules, Selectors or Relationships.
- Will not cover the IHS Web Server and client routing considerations.

WPS V6.0.1 – Clustering Administration 2  
© 2006 IBM Corporation



The previous unit on clustering *theory and concepts*, presented a deployment pattern for separating the *application* components from *Message Engine* components and separating the *infrastructure* components from the *application* components, which led to a 4 cluster topology.

The goal of this presentation is to describe the process of configuring this topology, including all of the basic tasks and issues that might not be apparent to you as a WebSphere administrator.

This discussion will identify the essential elements that must be installed and configured.


Installation and configuration of Business Rules, Selectors or Relationships will not be covered, as there are no special steps for these components with respect to the clustering configuration.

Installation and configuration of the IBM HTTP Server Web Server and client routing will also not be covered in this presentation.



## Agenda

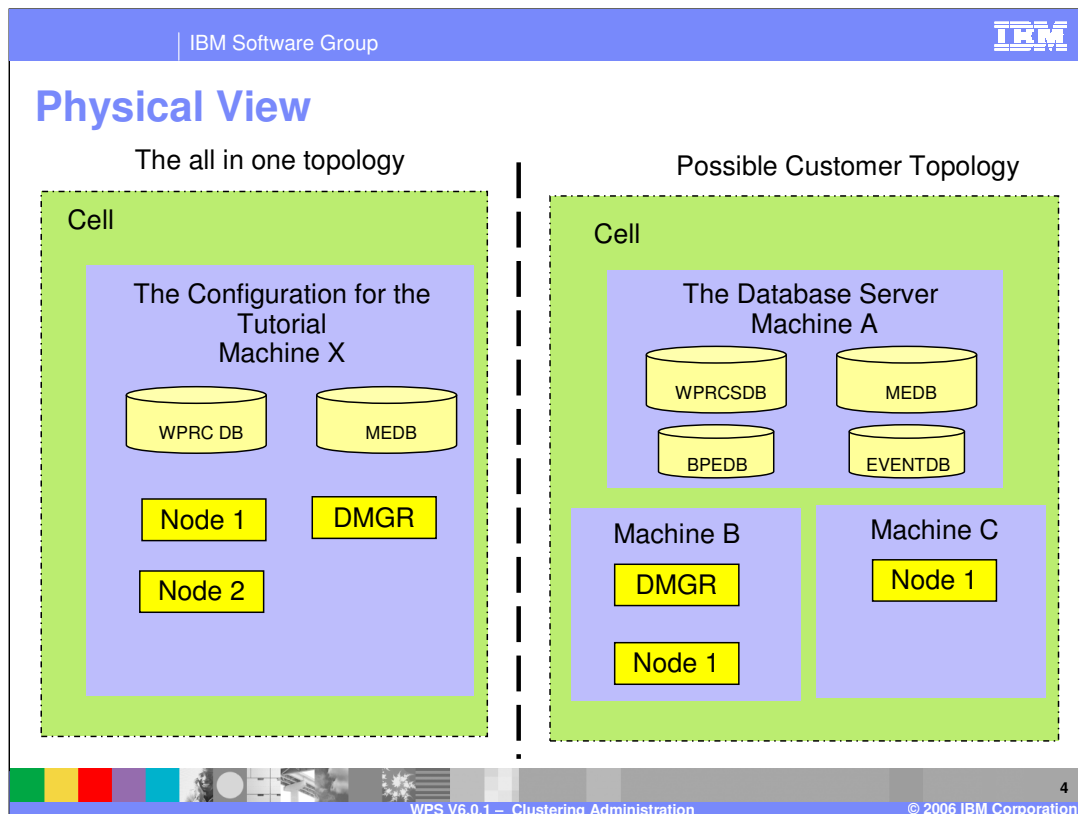
- The Topology
  - ▶ Physical View
  - ▶ Logical View
- The Server Clusters
  - ▶ Four basic kinds
- Overview and Discussion of the Steps Involved for configuring the 4-cluster topology.
- Creating and Configuring Server Clusters
  - ▶ Advanced Configuration Wizard



WPS V6.0.1 – Clustering Administration 3  
© 2006 IBM Corporation

The agenda is to cover the physical and logical view of the topology, followed by a description of the four basic kinds of server clusters.

An overview and description of the steps required to configure the 4 cluster topology and create and configure server clusters will also be provided.



For the purpose of simplicity and maintenance, the tutorial introduced here will only use two databases, as shown on the right.

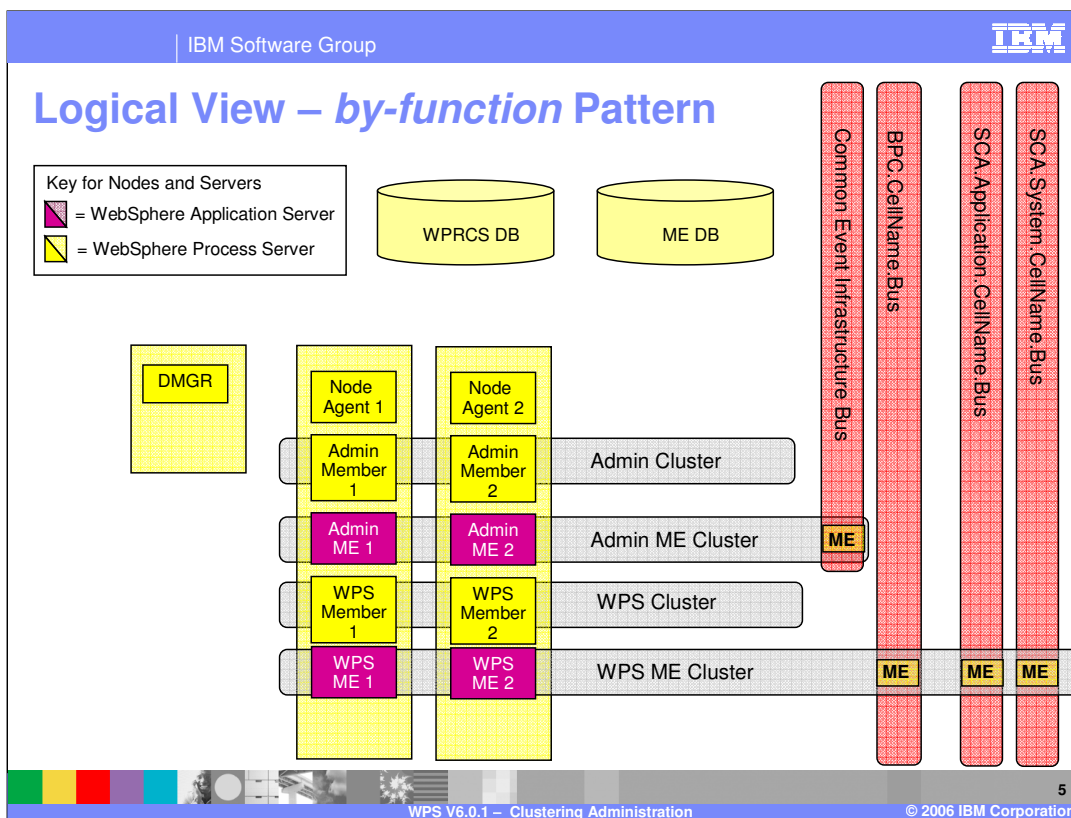
In a more realistic scenario the databases will be hosted on a remote database server with several databases defined, each based on their function and relationship to the applications and components they support, as shown in the cell on the right.

When designing the database topology, the function and scope of each database must be carefully considered and understood.

With the Business Process Choreographer, there must be a separate database ( *BPEDB* ) for each application server or application server cluster hosting the business process applications.

The Common Event Infrastructure, on the other hand, is a special purpose global service, that warrants a separate database. The default name used for this database is the *EVENTDB*.

It is expected that in order to meet scalability and high availability requirements that production environments will have topologies utilizing a database server with multiple databases and application servers distributed over several nodes.



Shown here is the logical view of the *by-function* pattern, with the cluster boundaries determined by function.


This is the topological deployment pattern that was presented in the *Theory and Concepts* discussion.

The next few slides will provide a more detailed discussion of this pattern.

Notice that the Admin and WPS clusters are not directly associated with a message engine and that they both use *WebSphere Process Server*, server types, the yellow servers.


The 2 message engine clusters, host the message engines created by the association with the service integration bus. The application servers in the message engine clusters are *WebSphere Application Server*, server types, the purple servers.

Each message engine will be active in only one application server at time, using the *1 of N* HA policy.

IBM Software Group


## Admin Server Clusters

- AdminMECluster
  - ▶ Administration Message Engine Cluster
  - ▶ Will contain the Message Engine used by the *AdminCluster*
  - ▶ This cluster will be configured so that the Message Engine will be active in one server in this cluster, as determined by the WebSphere High Availability Manager.
  - ▶ active/standby
- AdminCluster
  - ▶ Administration Cluster
  - ▶ Used to host the Common Event Infrastructure and other components used for monitoring, logging, and management of SCA components that are global to the system.
  - ▶ This cluster will be configured so that the Message Engines will be active in all application servers in the *AdminMECluster*.
  - ▶ active/active
- The Administrative cluster pair must be configured first.


WPS V6.0.1 – Clustering Administration
© 2006 IBM Corporation


Take a moment and consider the administration cluster pair.

The AdminMECluster contains the Message Engine used by the Common Event Infrastructure, using the active/standby HA policy.

The AdminCluster hosts the administrative applications and components, including the Common Event Infrastructure server, Common Event Infrastructure Message Driven Beans and Business Rules Manager.

There are dependency relationships between the various server clusters which must be considered. These relationships determine the order in which each cluster is installed, configured and started. The AdminMECluster must be configured before the AdminCluster, and the WPSMECluster must be configured before the WPSCluster.

Since the Common Event Infrastructure is used by the applications in the WPSCluster, the *administrative cluster pair* must be installed and configured before the *application cluster pair*.

IBM Software Group


## Application Server Clusters

- **WPSMECluster**
  - ▶ WPS Message Engine Cluster
  - ▶ Will contain the Message Engine used by the WPSCluster, which hosts the user applications
  - ▶ This cluster will be configured so that the Message Engine will be active in one server in this cluster, as determined by the WebSphere HA Manager.
  - ▶ active/standby
- **WPSCluster**
  - ▶ Used to host the user applications.
    - BPEL business processes
    - Applications that use the Service Component Architecture.
  - ▶ This cluster will be configured so that the Message Engines will be active all the application servers in the *WPSMECluster*.
  - ▶ active/active

WPS V6.0.1 – Clustering Administration
© 2006 IBM Corporation
7


Next, consider the application cluster pair.

The WPSMECluster contains the Message Engines used by the Service Component Architecture and the Business Process Choreographer components, using the active/standby HA policy.

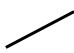



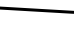
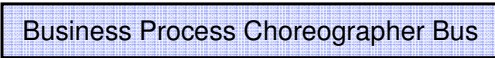
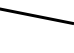
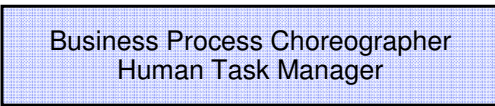
The WPSCluster hosts the BPEL applications and applications using the Service Component Architecture components.



As previously mentioned, the AdminMECluster and the AdminCluster must be configured and started before the WebSphere Process Server application and Message Engine clusters.

Since the WPSCluster depends on the WPSMECluster, the WPSMECluster must be started first.

a4 | IBM Software Group 

## Steps For Setting up the *by-function* Pattern

- Install the Product
- Create the WPRCSDB and MEDB databases
- Setup Databases
- Create DMGR
- Start DMGR
- Create Custom Nodes
- Create AdminMECluster  
- Create AdminCluster  
  - ▶ Configure the server
  - ▶ Grow cluster
- Create WPSMECluster  
- Create WPSCluster  
  - ▶ Configure the server
  - ▶ Grow cluster

  WPS V6.0.1 – Clustering Administration 8  
© 2006 IBM Corporation



The overall process for setting up a clustered environment in WebSphere Process Server is outlined here. The steps are very similar to setting up a typical WebSphere Application Server Network Deployment V6 cell.

First install the product, then create and configure the databases, define the cell by creating the deployment manager node, create custom nodes and add them to the cell.

With the *by-function* deployment pattern, the function of each application or process server cluster must be customized to meet specific needs.


1. The AdminMECluster requires the Service Integration Bus for the Common Event Infrastructure, which entails some manual steps.
2. The AdminCluster hosts the Common Event Infrastructure components, therefore they must be installed in this cluster.
3. The WPSMECluster hosts the Message Engine for the Business Process Choreographer applications, meaning the Service Integration Bus for the Business Process Choreographer must be installed and configured here. There are Business Process Choreographer specific entities that must be defined and associated with the Service Integration Bus manually.
4. The WPSCluster hosts the Business Process Choreographer applications, so the Business Process Choreographer and the Human Task Manager containers must be installed and configured in this cluster.



## Install WebSphere Process Server V6.0.1

- **WebSphere Process Server V6.0.1**
  - ▶ A complete refresh of the WebSphere Process Server product
  - ▶ The prerequisite WebSphere Application Server ND V6.0.2.3 is also installed automatically.
- **Options to install support for deprecated features**
  - ▶ Business Rule Beans
  - ▶ Extended Messaging
  - ▶ Samples
- **Do not run the Profile Wizard directly from the Installation Wizard when prompted at the end of the installation process.**
  - ▶ The Profile Creation Wizard for the Deployment Manger node requires a database and the chances are that this database has not been created yet.



WPS V6.0.1 – Clustering Administration © 2006 IBM Corporation 9

The first step is to install the WebSphere Process Server product.


The latest recommended release is WebSphere Process Server V6.0.1, with a base WebSphere Application Server ND of V6.0.2.3

Check the support site frequently for new updates.

Note: if upgrading from WebSphere Process Server V6.0, there is no automated migration utility. There are a few cases where some applications might need to be re-deployed and components with Visual Snippets must be re-authored. The relevant issues regarding migration and updates are documented in the release notes associated with a given release or update.



WebSphere Business Integration Server Foundation V5.1 applications that utilize Business Rule Beans or Extended Messaging can be run unchanged in WebSphere Process Server by installing the optional libraries. The BPEL applications on the other hand, will need to be migrated using the migration utility provided by the WebSphere Process Server runtime.

At the end of the installation process, there is an opportunity to run the Profile Creation Wizard. Because the Profile Creation Wizard requires a database for the Deployment Manager node, its recommended that this operation be deferred until after the database has been created.

IBM Software Group


## Create the WPRCSDB and MEDB Databases

- Databases and their respective tables can be created in several ways.
  - ▶ There are scripts which are provided as part of the WebSphere Process Server installation and...
  - ▶ There will be some cases where the Wizards in the *Administration Console* will provide the mechanism for populating the databases and tables.
  - ▶ Remote databases must be created ahead of time.
- Databases
  - ▶ WPRCSDB
    - Use the database tools provided by your database provider to create an *empty* database.
    - The tables will be created later by the Administration Console Wizards and Process Choreography scripts.
  - ▶ MEDB
    - Use the database tools provided by your database provider to create an *empty* database.
    - The Administration Console Wizards will create the tables later when configuring the Messaging Engines.
- Deployment Manager Profile Wizard requires a database.
  - ▶ It will provide the opportunity to create a new "*local*" database or use an existing database.
    - Remote databases must already exist and be cataloged.



WPS V6.0.1 – Clustering Administration
© 2006 IBM Corporation
10

The next step is to create the required databases.

At this point it is only necessary to have the databases created and cataloged. The table will be create in subsequent steps using a combination of Administrative Console Wizards and scripts provided with WebSphere Process Server runtime.


### WPRCSDB

Many of the WebSphere Process Server components require the use of a database. The tables for these components can be placed in separate databases or DB2® tablespaces as defined by the application requirements and the database administrator. For the purpose of this discussion and the tutorial, a single database called WPRCSDB will be used to simplify administration and facilitate the discussion.

### MEDB

The MEDB database is used to persist data used by the Messaging Engines. One datastore and schema per Messaging Engine is required.

If you use the Profile Creation Wizard to create the WPRCSDB, it will create a *local* database.

IBM Software Group


## Setup Databases

- Create Tables for the *Business Process Choreography* components
  - ▶ Use the appropriate creation scripts provided in the `WPS_Home\ProcessChoreography` directory
  - ▶ Default database name is BPEDB
- Create the tables for the *Enterprise Service Bus* message logging.
  - ▶ The script is located in `WPS_Home\util\EsbLoggerMediation\databaseType\table.ddl`
- The tables for the *Common Event Infrastructure* components will be created during the configuration of the AdminCluster.
  - ▶ Default database name is EVENTDB

WPS V6.0.1 – Clustering Administration
© 2006 IBM Corporation


The WebSphere Process Server runtime installation provides sample scripts to create and configure the many databases that are required. Because of the custom nature of the topology being developed, the scripts must be modified before being used.

The tutorial describes the details of creating the tables, and generating and running the scripts. Here the focus is on *what* must happen at this point in the setup process, which is to create the databases.

Before creating the WebSphere Process Server node profiles in the next step, create the tables for the Business Process Choreographer and the Enterprise Service Bus.

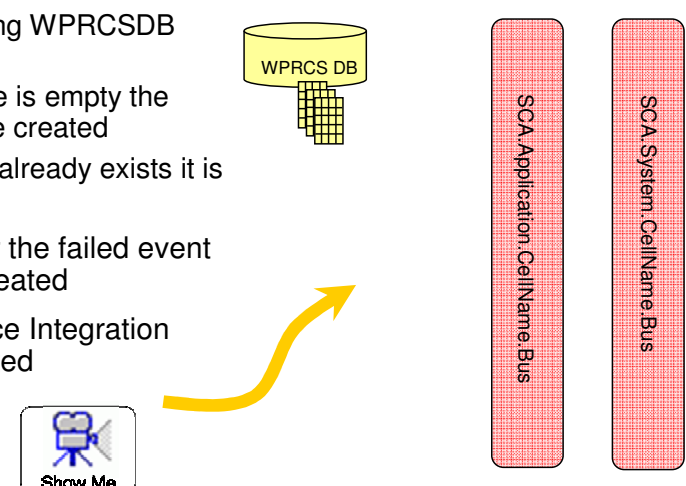
The database tables for the Common Event Infrastructure will be created later, when creating the AdminCluster, and will be placed in the *WPRCSDB* database.

When configuring a production environment, with a separate Common Event Infrastructure database, the default name used is *EVENTDB*.

a3 | IBM Software Group 


## Create Deployment Manager Profile

- Requires a database for 'general' WebSphere Process Server needs.
  - ▶ Use the existing WPRCSDB database.
  - ▶ If the database is empty the schema will be created
  - ▶ If the schema already exists it is left intact.
- Some tables for the failed event manager are created
- The SCA Service Integration Buses are created
  - ▶ Application
  - ▶ System



**Service Integration Buses**

SCA.Application:CellName:Bus      SCA.System:CellName:Bus

 Show Me

WPS V6.0.1 – Clustering Administration 12 © 2006 IBM Corporation


After the databases are established, its time to create the WebSphere Process Server Cell. The first step in creating a cell is to create the deployment manager node by running the Profile Creation Wizard and selecting the Deployment Manager node as the type of node to create.

The Profile Creation Wizard detects the presence or absence of the schema tables in the database and behaves in the appropriate manner, meaning it will use the existing schema and tables if they are present and create them if they are not.


The Profile Creation Wizard also creates the mandatory Service Component Architecture Application and System buses.

Creation at this point really means to define it in the profile. When the Deployment Manager is started with the new profile, the entities are physically created.

To see the Profile Creation Wizard in action, select the “Show Me” link.

IBM Software Group 

## Create Custom Nodes




- Nothing special here.
  - ▶ Standard V6 configuration
- Define nodes based on the topology
  - ▶ Logical and Physical
- Sets up some environmental variables.
- Keep the name of the Profile and the node as short as possible if configuring for a Windows™ platform
  - ▶ Path\Name length limitation.
  - ▶ The cell, node and profile names get used for creating deployed artifacts

WPS V6.0.1 – Clustering Administration 13  
© 2006 IBM Corporation

Once the cell has been created, the application server nodes can be create and added to the cell using the Profile Creation Wizard, specifying, *custom node*, as the node type.







At this point, the biggest decision is whether or not to have more than one node on a machine. This will depend on the physical attributes of the system and your server management strategy. There is nothing specific with respect to WebSphere Process Server for this configuration step.

To see the Profile Creation Wizard with the custom profile option, select the “Show Me” link.

a4 | IBM Software Group 

## Creating and Configuring Clusters

- The strategy is to *create* and *configure* a server cluster with one cluster member and then '*grow*' the cluster by adding additional servers.
- When creating a server cluster, there are initially 3 default templates to choose from.
  - ▶ default = WebSphere Application Server (Can be used with the Message Engine Clusters)
  - ▶ defaultESBServer = Enterprise Service Bus server
  - ▶ defaultProcessServer = WebSphere Process Server (Can be used with the Application Clusters)
- Configuring the cluster will depend on the kind of cluster
  - ▶ Messaging Engine Cluster ( a messaging server )
  - ▶ Application Cluster ( a messaging client )
    - Kind of application
      - WebSphere Process Server (Business Process, Relationships, Business State Machines, ...)
      - Basic WebSphere Application

      14  
WPS V6.0.1 – Clustering Administration © 2006 IBM Corporation



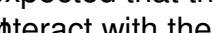

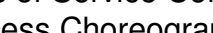
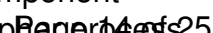
Having established the physical nodes which will host the servers and server clusters, the server clusters can be created.

When creating server clusters the initial server is created and configured per the application requirements and this initial server is then used as a model to create or clone additional servers in the cluster. The strategy employed here is to create and configure each server cluster with one server and then grow the clusters as needed based on the requirements and capacity of the system and hardware.

When creating a server cluster, you can select from one of three default server templates or use an existing server as a template and that template will be applied to all the servers in the cluster. It is important to understand that all three types of servers can exist in a WebSphere Process Server Cell, allowing you to select the right server for the task at hand. For instance, this means that for the Message Engine functionality, you can create a basic WebSphere Application Server without the overhead associated with a WebSphere Process Server. This is shown in the Logical view of the *by-function* Pattern.

Configuring a given cluster will depend on the functionality hosted by that cluster.

A Business Process is just one of several possible kinds of WebSphere Process Server applications. There can also be WebSphere Process Servers that host Relationships, Business Rules, Selectors, Business State Machines and basic Service Component Architecture components. Its expected that these other kinds of Service Component

Architecture components will interact with the Business Process Choreography flows in some manner.       Page 14 of 25

a2 | IBM Software Group | IBM

## Concept: Message Engine, Service Integration Bus, Server

- The server or server cluster hosts a Message Engine.
- The Message Engine is associated with an Service Integration Bus.
- The Message Engine is created when a server or server cluster is added as a member to the Service Integration Bus.
  - ▶ The association can be verified after the fact by inspecting the Service Integration Bus configuration.
- The Message Engine has a unique schema and datastore in a database
  - ▶ The schema name must be set after the Message Engine has been created.
- The Message Engine may have many JSM Destinations.

WPS V6.0.1 – Clustering Administration | © 2006 IBM Corporation | 15

The most important aspect of configuring a cluster in a WebSphere Process Server environment is the messaging engine.

### Attributes of the Message Engine:

- The Message Engine, which is created when the server or server cluster is added to the Service Integration Bus as a bus member, is the link between the Service Integration Bus and the application server or server cluster. To verify the association after the server or cluster is added, you can inspect the Service Integration Bus configuration.
- The datastore has a 1:1 relationship to the Message Engine and is a set of tables qualified by a unique schema.
- When adding the server or server cluster to the Service Integration Bus there is no option to set the schema name. Therefore, it must be set to a unique and meaningful name in a subsequent step.
- Each Message Engine associated with a Service Integration Bus can have many JMS Destinations.
- In the topology shown here, the active/standby pattern, the messaging engine will be active in one of the servers in the application server cluster as determined by the WebSphere High Availability and Workload Managers.

IBM Software Group

## Advanced Configuration Wizard

- A Wizard provided by the Administration Console as an “*additional property*”
- Used to configure WebSphere Process Server and Service Component Architecture features on the server or server cluster
- **Run once only per server or server cluster**

Runtime Configuration Local Topology

**General Properties**

\* Cluster name  
WPS\_Cluster

Bounding node group name  
DefaultNodeGroup

Prefer local

Enable high availability for persistent services

Apply OK Reset Cancel

**Additional Properties**

- [Cluster members](#)
- [Backup cluster](#)
- [Advanced Configuration](#)
- [Business process container](#)
- [Human task container](#)
- [Endpoint Listeners](#)

WPS V6.0.1 - Clustering Administration © 2006 IBM Corporation

When creating servers and server clusters in WebSphere Process Server there are certain process server configuration tasks, such as installing the Business Rules Manager and configuring the service integration bus, that must be considered. What to configure or how, will depend on the functionality to be provided by the server or server cluster.

The **Advanced Configuration Wizard** is a tool provided by the WebSphere Process Server adminconsole to assist the administrator in performing these tasks. It is provided as an *Additional Property* of the server or server cluster.




IBM Software Group

## Advanced Configuration Wizard: Message

- If the Advanced Configuration Wizard has already been run for a given server or server cluster, the following message will be displayed.

Messages

 This cluster or server has already been setup to run Process Server applications and the helper cannot be executed again.

- The Advanced Configuration Wizard cannot be used to determine the state to of the configuration at a later time.

WPS V6.0.1 - Clustering Administration © 2006 IBM Corporation

**Advanced Configuration** is a Wizard, which means that it is run once and the information provided on the panel is used to configure the WebSphere Process Server and Service Component Architecture components on the selected server or server cluster. Because it is not an edit panel, it can only be run once and cannot be used to verify the configuration.

To verify the configuration requires knowledge of what the wizard is doing with respect to the System and Application Service Integration Buses, Message Engines and datastores, along with an understanding of the client/server nature of the messaging infrastructure, whether remote or local.

The following discussion on Advanced Configuration will reinforce these features and concepts.

There are three primary task that the Advanced Configuration Wizard assists with.

1. Installing the Business Rules Manager. If the Business Rules are to be used the Business Rules Manager must be installed some where, on one of the servers or server clusters in the cell.
2. Specifying the JNDI name of the emitter factory to be used by the Common Event Infrastructure. The Common Event Infrastructure is a global service that will be used by components in the WebSphere Process Server servers and server clusters. The JNDI name is how the applications that need to emit events to the infrastructure locate the appropriate emitter to use.
3. Configure the Service Component Architecture components.

The Service Component Architecture options are mutually exclusive, so if there are no Service Component Architecture applications to be hosted in this server or cluster, there is no need to configure the rest.

If there will be Service Component Architecture applications hosted by this server or server cluster and the Message Engine is in a remote destination on another server or server cluster, then pick the server or server cluster that will be hosting the Message Engine. This means that the messaging server cluster must be defined first.

If the Message Engine will be hosted in the server or server cluster currently being configured, then there are additional fields available for defining attributes of the

IBM Software Group

## Advanced Configuration: SCA Applications

- The “Hosting SCA applications” option will determine whether or not to configure the Message Engine on the SCA.application and the SCA.system busses.
  - ▶ With the *by-function deployment* pattern
    - The *AdminCluster* does not have any SCA application requirements
      - Hosts the Common Event Infrastructure and Business Rule Manager applications.
    - The *WPSCluster* does, and the Message Engine is remote.
      - Hosts the Business Process Choreography and other WebSphere Process Server applications.
    - The *WPSMECluster* does and the Message Engine is default (local).
      - Hosts the Message Engine used by the WPSCluster.

WPS V6.0.1 – Clustering Administration
© 2006 IBM Corporation
19

The *by-function* deployment pattern provides examples of using the different options of the Advanced Configuration Wizard.

Although the *AdminCluster* is composed of WebSphere Process Server servers it does not have any SCA application requirements. The hosted applications are not using BPEL or any of the other SCA components. It does host the Business Rules Manager, which a WebSphere Process Server application that uses standard web technology.

The *WPSCluster* will be hosting the Service Component Architecture components, but the messaging engine used with the service integration bus will be remote, it will be located in the *WPSMECluster*.

The *WPSMECluster* will not be hosting the Service Component Architecture components directly but it will be hosting the messaging engine used by the components from the *WPSCluster*.

Each cluster will have to be considered carefully and the relationships between the applications, the cluster, the messaging engine and the service integration bus must be understood.

IBM Software Group IBM

## Advanced Configuration : Remote vs. Default

- *Advanced Configuration* defines whether queue destinations exploited by the *Service Component Architecture* runtime are hosted locally or on a remote server or cluster.
- *Remote* is some other server or server cluster.
- The *default* location is the server or server cluster being configured, that is to say, *local*.

The Message Engine for the WPS Cluster is hosted in another server cluster, so it is **remote**.

WPS Cluster

WPS ME Cluster

SCA:Application:CellName:Bus

SCA:System:CellName:Bus

ME ME

WPS V6.0.1 – Clustering Administration 20  
© 2006 IBM Corporation

The notion of Remote vs. Default can be a bit confusing. It's clear what it means to have a remote message engine. The message engine is hosted in another server or server cluster.

Default on the other hand is not as clear. In this case, default means local to the server or server cluster being configured.

In the picture here, the Message Engine for the *WPSCluster* is hosted in another server cluster, called the *WPSMECluster*, which is **remote** with respect to the *WPSCluster*.

The applications in the *WPSCluster* interact with the service integration bus through the JMS destinations that are configured in that cluster to use the Message Engine in the *WPSMECluster*.

IBM Software Group

## Advanced Configuration: Default Destination

- If the server or server cluster being configured is hosting the *Message Engine(s)* locally, that is *default*, then additional configuration is required.
- Each Message Engine requires an unique schema and in a database.
  - ▶ JDBC™ Provider
  - ▶ Unique Schema Names
  - ▶ User Id and Password
  - ▶ Databases
  - ▶ Table creation
    - Yes or no

The persistent data associated with each of the SCA busses, system and application, can go into different databases if necessary. They are configured independently.

WPS V6.0.1 – Clustering Administration
© 2006 IBM Corporation
21

When the Message Engine is hosted by the server or server cluster being configured, there are additional Message Engine configuration tasks that must be completed.

The Message Engine component creates the association between the service integration bus and the server or server cluster. For each Message Engine there must be a unique schema and table in a database. To access the database there must also be a JDBC provider and datasource along with a userid and password.

The purpose of this part of the Wizard is to gather all this information so that the unique tables, qualified by the schema, can be created in the database and associated with the Message Engine, which is associated with the server or server cluster and the service integration buses for the system and application.

Notice that the persistent data for the System and Application Service Integration Buses can be located in different databases, it is not a recommended practice.

IBM Software Group IBM

## Advanced Configuration: Message Engine


- The *schema* will qualify a set of tables ensuring their uniqueness within the database.
- The *schema* name should be such that it connotes the association between the cluster and the Service Integration Bus.
  - ▶ WPSMEClusterAppBus
  - ▶ WPSMEClusterSysBus
- There can be many JMS Destinations for a given Message Engine
- The Message Engine is running in one of several possible servers in the Message Engine server cluster.
  - ▶ This illustrates the association of the queue, server, Service Integration Bus and db table.

22

WPS V6.0.1 – Clustering Administration © 2006 IBM Corporation

If the database administrator allows for the creation of the tables using the Administrative Console and the **Create Tables** option is checked, the Wizard will create the schema qualified tables in the specified database using the JDBC Provider specified.

The key point shown here is the naming convention used to name the schemas. The schema is used to uniquely identify a set of tables and everything will go into the same database. Giving the schema a name that connotes the relationship between the server or server cluster and the service integration bus will aid in debugging and verifying installations at some point in the future.

IBM Software Group 

## Advanced Configuration: Summary

- Advanced Configuration Wizard, is a Wizard
  - ▶ Cannot be used to determine the current state of the configuration
    - Run once only
- Used to assist in the configuration of the *Service Component Architecture* Service Integration Busses only.
  - ▶ Meaning:
    - SCA.Application.*yourCellName*.Bus
    - SCA.System.*yourCellName*.Bus
- User defined *Service Integration Busses* must be configured manually.
  - ▶ Meaning:
    - BPC.*yourCellName*.Bus
    - CEI.*yourCellName*.Bus

WPS V6.0.1 – Clustering Administration 23  
© 2006 IBM Corporation

The Advanced Configuration Wizard can only be used once to configure service integration busses and cannot be used to verify the configuration at a later time.

It is only used to configure the default Service Component Architecture busses which means that the user defined service integration buses such as the Common Event Infrastructure and the Business Process Choreographer buses, must be configured manually.



## Summary

- WebSphere Application Server ND V6 cell basics
  - ▶ Profile Creation Wizard for WebSphere Process Server
    - ProfileCreator\_wbi
- Message Engine is a key differentiator
  - ▶ Separate Clusters
  - ▶ Active/standby
- Customized cluster configurations
  - ▶ Manually intensive steps
- Advanced Configuration Wizard
  - ▶ Aids in WebSphere Process Server specific tasks



WPS V6.0.1 – Clustering Administration © 2006 IBM Corporation 24

Much of the process for setting up a clustered WebSphere Process Server environment is similar to setting up a WebSphere Application Server ND V6 cell and cluster.

Note that there is a special profile creation wizard that is provided with WebSphere Process Server. Using the wrong one will cause confusion and rework.

When configuring server clusters that will be hosting Business Process Choreographer components, the requirements for the message engine will constrain the possibilities and drive the configuration toward separate clusters for the message engines.

Since each cluster is special and unique, the default scripts provided by the WebSphere Process Server installation must be manually edited to conform to the customized topology.

To ease the burden, the Advanced Configuration Wizard is provided as an aid to manage the WebSphere Process Server unique tasks involved in setting up a WebSphere Process Server cell.



# 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.