



IBM Software Group

WebSphere® Process Server V6.0.2

Clustering – Administration and configuration



@business on demand.

© 2007 IBM Corporation
Updated April 6, 2007

This presentation covers clustering administration and configuration for WebSphere Process Server version 6.0.2.

Goals

- Present the administrative tasks involved with setting up a clustered topology with WebSphere Process Server V6.0.2
 - ▶ The four-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 IBM HTTP Server and client routing considerations

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 four-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

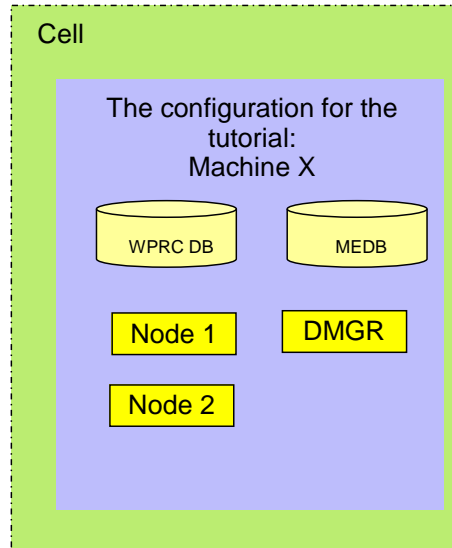
- The topology
 - ▶ Physical view
 - ▶ Logical view
- The server clusters
 - ▶ Four basic kinds
- Overview and discussion of the steps involved for configuring the four-cluster topology.
- Creating and configuring server clusters
 - ▶ Business integration configuration wizard

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.

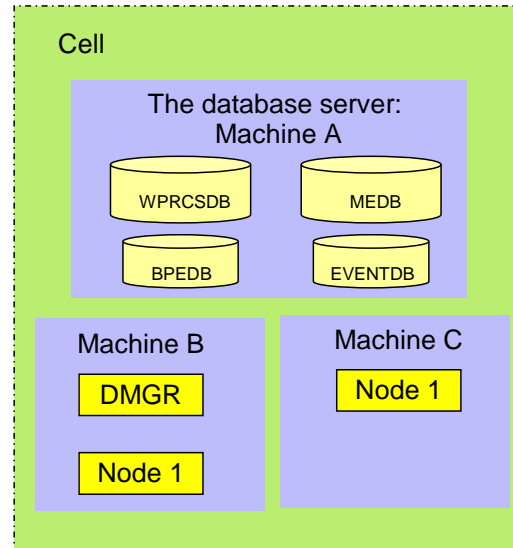
An overview and description of the steps required to configure the four-cluster topology and create and configure server clusters will also be provided.

Physical view

The all in one topology



Possible customer topology



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

In a more realistic scenario the databases will be hosted on a remote database server with several databases defined.

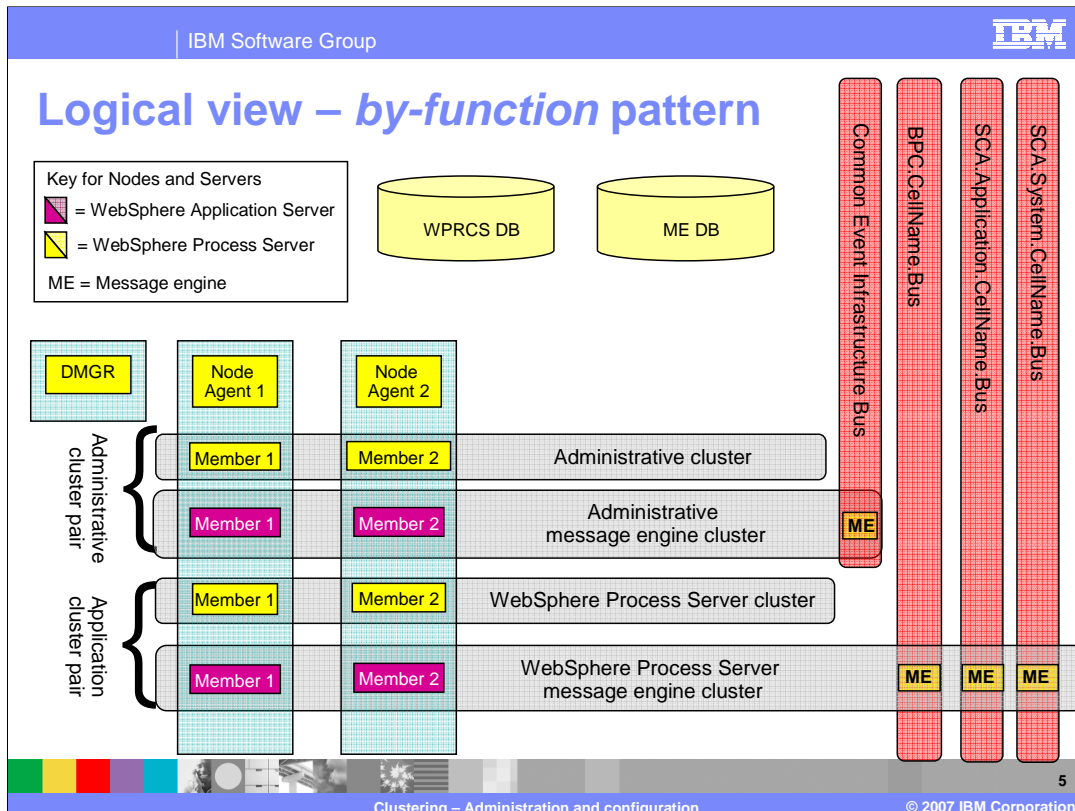
Each database will be based on their function and relationship to the applications and components they support, as shown in the cell on the right.

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 Administrative and WebSphere Process Server 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* High Availability policy.

Administrative 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.

6

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 High Availability 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*.

Application server clusters

- WPSMECluster
 - ▶ WebSphere Process Server message engine cluster
 - ▶ Will contain the message engine used by the WPSCluster, which hosts 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 high availability 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

7

Next, consider the application cluster pair.

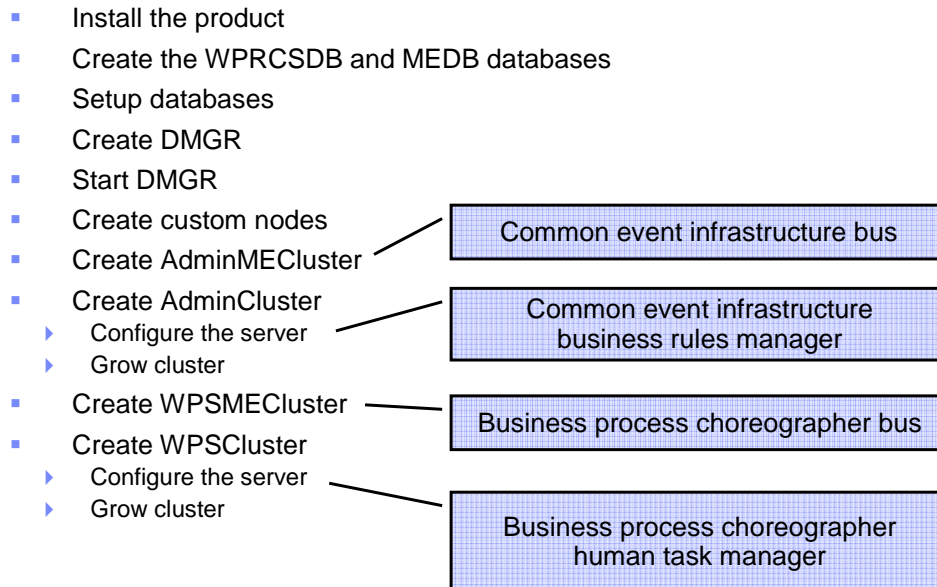
The WPSMECluster contains the Message Engines used by the Service Component Architecture and the Business Process Choreographer components, using the active/standby High Availability policy.

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.

Steps for setting up the *by-function* pattern



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.

Since each cluster is highly customized to meet a special function the default scripts and installation processes are no longer sufficient.

The tutorial that accompanies this discussion will provide the detailed guidance for the required customizations.

The remainder of this discussion will cover each of these steps at a high level.

Install WebSphere Process Server V6.0.2

- WebSphere Process Server V6.0.2
 - ▶ A complete refresh of the WebSphere Process Server product
 - ▶ The prerequisite WebSphere Application Server Network Deployment V6.0.2.17 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 manager node requires a database and the chances are that this database has not been created yet.

The first step is to install the WebSphere Process Server product. This may be a new installation of the WebSphere Process Server V6.0.2 or an upgrade from V6.0.1

The latest recommended release is WebSphere Process Server V6.0.2, with a base WebSphere Application Server Network Deployment of V6.0.2.17
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.

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.

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 created 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 data store and schema per Messaging Engine is required.

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

Setup databases

- Create tables for the *business process choreography* components
 - ▶ Use the appropriate creation scripts provided in the `WPS_Home\dbscripts\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

11

Clustering – Administration and configuration

© 2007 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*.

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 also created



note: in V6.0.1 the SCA.SYSTEM and the SCA.APPLICATION buses were created during profile creation.

This is no longer the case with V6.0.2. They will be automatically created at a later time, when appropriate.

12

Clustering – Administration and configuration

© 2007 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. This means it will use the existing schema and tables if they are present and create them if they are not.

Be aware of new behavior in V6.0.2 in the area of the Service Component Architecture Service Integration bus creation. This behavior has been modified to allow for the federation of a stand-alone server into the deployment manager node.

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.

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 and name length limitation.
 - ▶ The cell, node and profile names get used for creating deployed artifacts

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, pause this presentation and click the “Show Me” icon.

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 three 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

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.

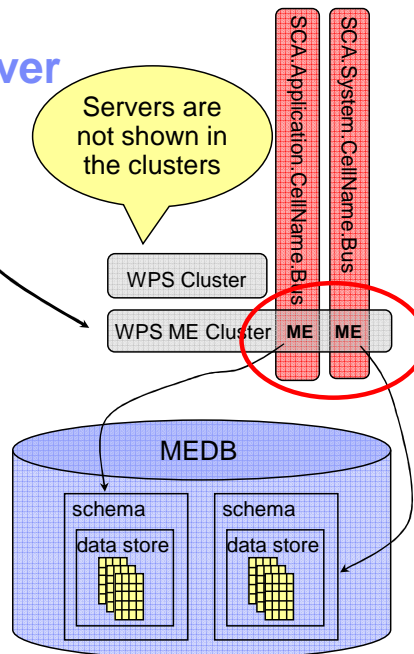
Configuring a given cluster will depend on the functionality hosted by that cluster.

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 Choreographer process flows in some manner.

Subsequent slides will present key features of the Administration Console that are used to install and configure WebSphere Process Server specific services for clustering. For a more detailed walk through, there are several demos and a tutorial that provide a step by step guide for creating and configuring this 4 cluster topology.

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 data store in a database
 - ▶ The schema name must be set after the message engine has been created.
- The message engine may have many JSM destinations.



15

The most important aspect of configuring a cluster in a WebSphere Process Server environment is the messaging engine. The concept is presented at this time because the creation and configuration of the messaging engine is done using the Business Integration Configuration Wizard and is not obvious to the inexperienced administrator. Configuring the messaging engine is part of configuring the SCA Destinations and will be covered in more detail when configuring the WPSMECluster.

Attributes of the Message Engine:

- The Message Engine is created when the server or server cluster is added to the Service Integration Bus as a bus member. The message engine becomes 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 data store 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 messaging engine cluster as determined by the WebSphere High Availability and Workload Managers.

Business integration configuration

- A configuration 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

16

Clustering – Administration and configuration

© 2007 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 **Business Integration 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 configuration.

The Business Integration Configuration Wizard is a replacement for the Advanced Configuration Wizard provided with WebSphere Process Server V6.0.1. It provides the same basic functionality in an easier to use format.

IBM Software Group IBM

New V602

Business integration configuration

- Used to install and configure system services on servers and server clusters.
 - ▶ SCA resources
 - ▶ Business rules manager
 - ▶ Business process container
 - ▶ Human task manager container.
- Can be used to verify current configurations.
- Cannot be used to uninstall currently configured services.

Configure your Network Deployment Environment

Select your clusters and servers to configure

Add your servers and clusters that you would like to configure by clicking on Add. Select in the table the services that you would like to configure for the given target.

If a particular service has already been installed, then this is shown in the table. You may not uninstall an already installed and configured service by unchecking the selection.

If you configure a service on more than one host, then default configurations are applied if no configuration panel is shown or the configuration panel settings apply to the services on all hosts. If you select one application, then you are able to fine tune each aspect of your service configuration during this wizard session. Run the wizard for each service that you need to have different custom configurations.

Select a cluster to configure

Admin

Select a server to configure

Add

Remove

Select	Server/Cluster Name	Node	Setup SCA Destination		Business Rule Manager	Business Processes	Human Tasks
			Local	Remote			
<input type="checkbox"/>	Admin		<input type="radio"/>	<input checked="" type="radio"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

Total 1

Next Cancel

Clustering – Administration and configuration © 2007 IBM Corporation 17

The Business Integration Configuration Wizard is used to install and configure system services on servers and server clusters in WebSphere Process Server 6.0.2.

More than one service can be installed or configured for a given server or server cluster with a single pass through the wizard but the wizard cannot be used to uninstall a service.

If a particular service has already been installed, then this is shown in the table.

It is also possible to install and configure a service on multiple servers or server clusters at the same time by using the drop down list and the 'add' button to extend the list.

If you configure a service on more than one host and there are no configuration panels, then default configurations are applied to all of the target hosts. If configuration options are available, then configuration panels will be displayed and the options will be applied to all of the selected hosts.

If you select one application, then you are able to fine tune each aspect of your service configuration during this wizard session.

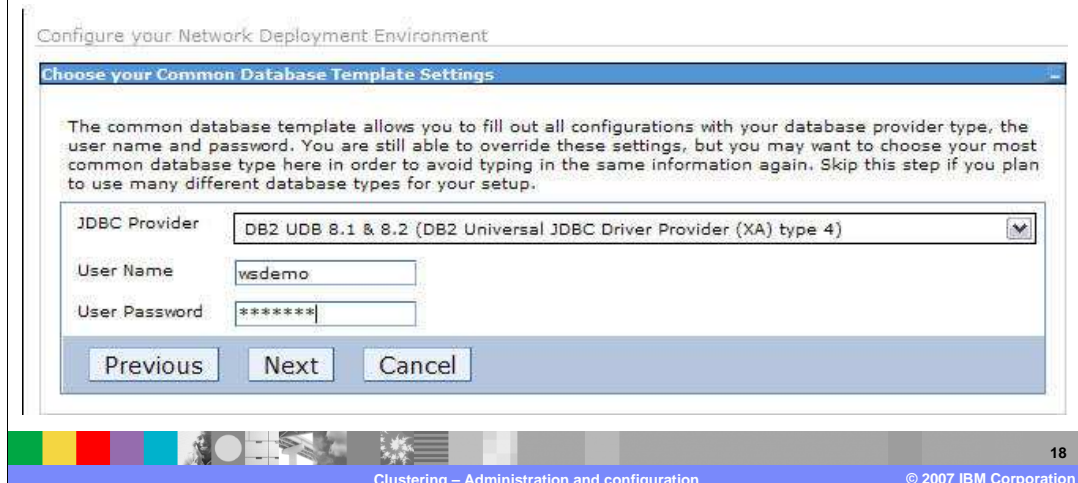
Run the wizard for each service that you need to have different custom configurations.

In this illustration the Administrative Cluster is being configured. The SCA destination is remote and the Business Rules Manager will be installed to run on this cluster. In the topology being presented, the Business Process Choreographer container and the Human Task Manager container will be installed and configured on the WPSCluster, therefore they are not selected here.

The subsequent panels displayed by the wizard, will depend on the options selected on this panel.

The database template

- A database template for the most commonly used database types
- It may be overridden if necessary or skipped entirely



Continuing with the scenario to configure the SCA Destinations and Business Rules Manager on the Administrative Cluster. The first task the administrator must complete, is to define a common database template that will provide the default JDBC Provider and security information for the parts of the wizard that require this information. This reduces the amount of input the administrator has to provide while going through the Wizard.

The information in the template may be overridden at later steps if necessary and if many different database types will be used then this step can be skipped.

IBM Software Group IBM

SCA destinations

- Will the server or server cluster being configured host service component architecture applications?
- If so, then will the queue destinations exploited by the service component architecture runtime be hosted locally or on a remote server or server cluster?

Config your Network Deployment Environment

Configure selected clusters and servers

Step 1: Use Remote Destination Location for Service Component Architecture

Admin

Step 2: Common Event Infrastructure Destination

Step 3: Business Rule Manager

Step 4: Summary

Use Remote Destination Location for Service Component Architecture

Do not host SCA applications

Use a remote destination location

Configure a destination location

Next Cancel

The message engine for the WebSphere Process Server cluster is hosted in another server cluster, so it is **remote**.

WebSphere Process Server cluster

WebSphere Process Server message engine cluster **ME ME**

19

Clustering – Administration and configuration © 2007 IBM Corporation

First, notice that this is the first step of a 4 step process as shown in the left part of the guided task. The current step is highlighted in yellow.

There are three options to consider.

Option 1:

Do not host Service Component Architecture applications in the server or server cluster being configured. This is the simple case. If this is selected then there are no further choices. This is the choice shown in the current step of the guided task and is used when configuring the Administrative Cluster.

<give the student time to absorb the visual >

Option 2:

Host Service Component Architecture applications in this server or server cluster but use a message engine that is on a different server or server cluster. This case is illustrated with the WPS Cluster.

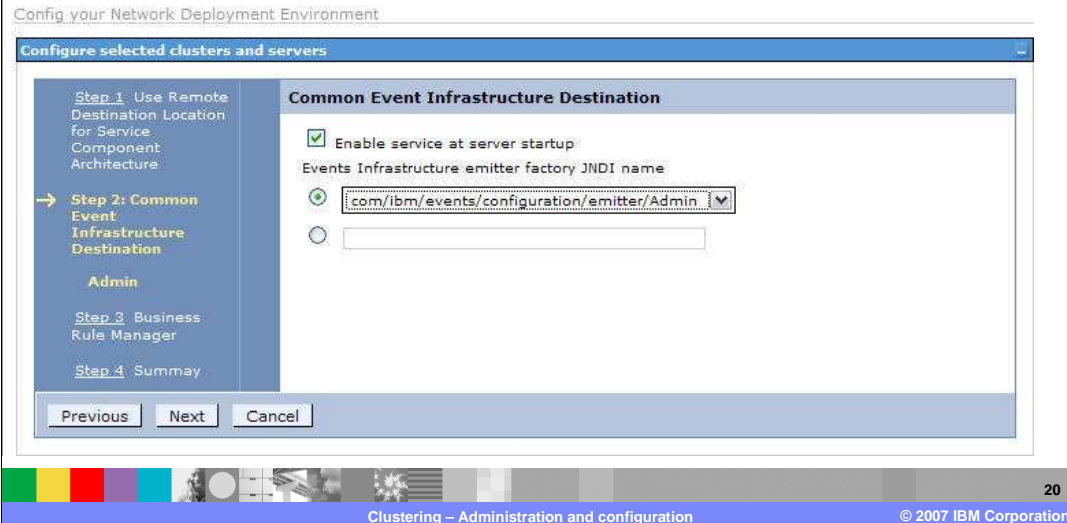
<give the student time to absorb the visual >

Option 3:

Host Service Component Architecture applications in this server or server cluster but use a message engine that is configured locally on the server or server cluster being configured. This case is illustrated with the WPS ME Cluster. When this option is selected there additional configuration information will be presented on the panel.

Common event infrastructure destinations

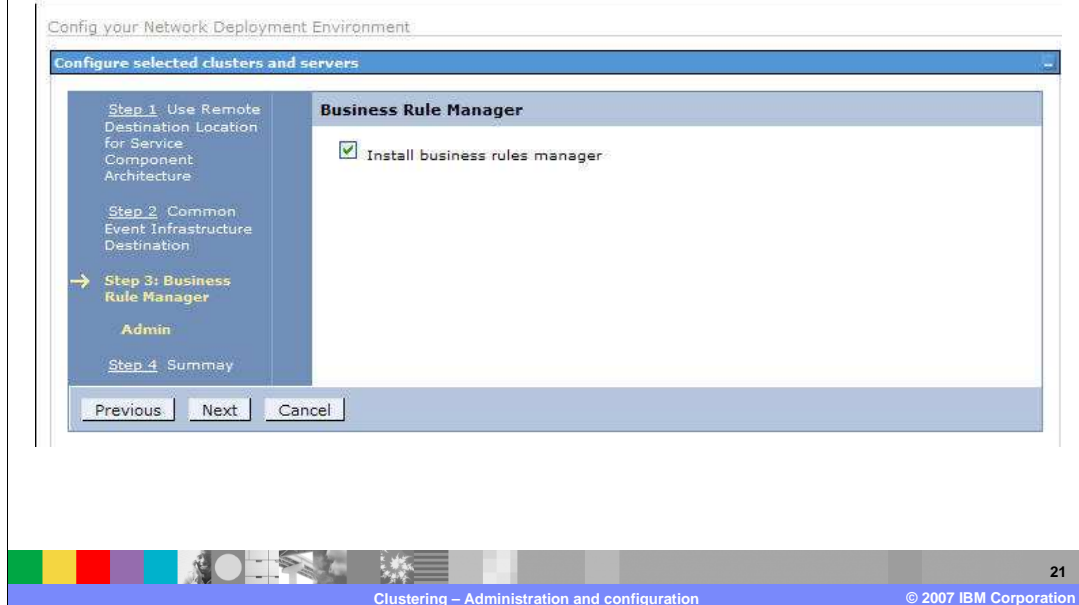
- The common event infrastructure is a global service that will be used by components in the WebSphere Process Server servers and server clusters.
- The JNDI name is how the applications that need to emit events to the infrastructure locate the appropriate emitter to use.



Continuing with the configuration of the Administrative Cluster, step 2 is the specification of the resources used for the Common Event Infrastructure.

Here the system administrator can specify whether the Common Event Infrastructure service is enabled at startup and where the emitter factory can be located in the JNDI namespace.

Business rules manager



Step 3 is where the system administrator indicates that the Business Rules Manager application is to be installed. For the configuration of the Administrative Cluster the Business Rules manager was selected on the initial page of the Business Integration Configuration Wizard, therefore it is automatically selected here, providing a verification step for the administrator.

Summary for configuration of the AdminCluster

Config your Network Deployment Environment

Configure selected clusters and servers

Step 1 Use Remote Destination Location for Service Component Architecture

Step 2 Common Event Infrastructure Destination

Step 3 Business Rule Manager

→ Step 4: Summary

Summary

Remote	
Do not host SCA applications	true
Common Event Infrastructure Destination	
Enable service at server startup	true
Events Infrastructure emitter factory JNDI name	com/ibm/events/configuration/emitter/Admin
Business Rule Manager	
Install business rules manager	true

Previous Finish Cancel

22

Clustering – Administration and configuration © 2007 IBM Corporation

Regardless of the path traversed through the Business Integration Configuration Wizard or the number of steps taken, at the end there will be a final summary and an opportunity to verify and confirm the choices made before committing the configuration changes.

The WPSMECluster configuration

- WPSMECluster
 - ▶ WebSphere Process Server message engine cluster
 - ▶ Will contain the message engine used by the WPSCluster, which hosts the end-user applications

Select	Server/Cluster Name	Node	Setup SCA Destination		Business Rule Manager	Business Processes	Human Tasks
			Local	Remote			
<input checked="" type="checkbox"/>	WPSMECluster		<input checked="" type="radio"/>	<input type="radio"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

Total 1

Configuring the WPSMECluster provides an example of the steps and thought processes involved in configuring the Service Component Architecture destinations locally.

The WPSMECluster is where the message engine used by the WPSCluster resides. This scenario will illustrate how to configure the message engines used by the Service Component Architecture Service Integration Buses.

Local destination configuration

- If the server or server cluster being configured is hosting the *message engines* locally then additional configuration is required
- Each message engine requires a unique schema 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 buses, system and application, can go into different databases if necessary. They are configured independently.

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 data source along with a user ID and password.

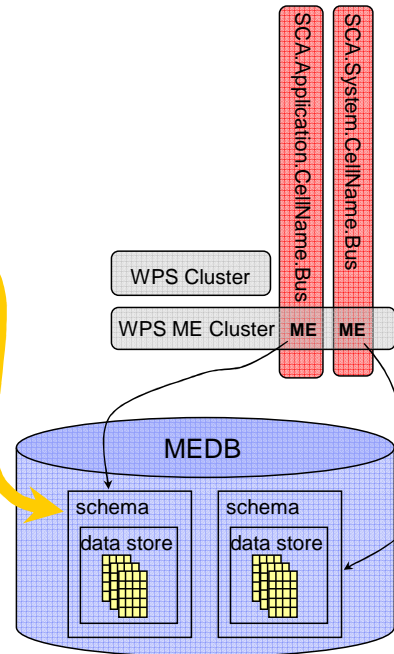
The purpose of this part of the Wizard is to gather all this information so that the unique tables, qualified by the schema, can be created in the database.

The Wizard also creates the association between the Message Engine, the server or server cluster and the service integration buses.

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

Message engine schema

- 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



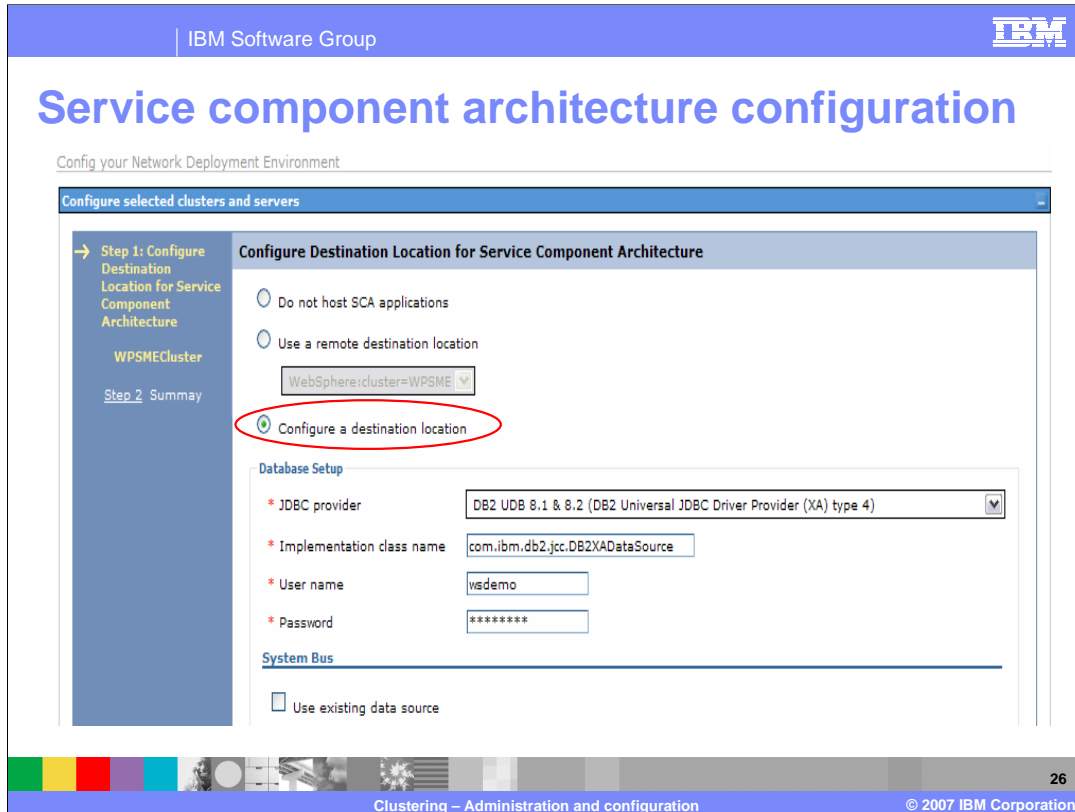
25

All of the tables for the message engine data stores can go into the same database as long as they have a unique schema name.

The key point shown here is the naming convention used to name the schemas. 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.

For example;

WPSMEClusterAppBus indicates that the data store qualified by this schema is for the message engine that is used by the WPSMECluster for the SCA.APPLICATION bus.



Shown here is the top half of the Service Component Architecture Configuration when the option to 'configure a destination location' is selected.

The database information has already been filled in based on the information from the database template which was configured before getting to this step.

Service component architecture configuration

The screenshot displays the configuration interface for the Service Component Architecture. It is divided into two main sections: System Bus and Application Bus. Each section has a 'Use existing data source' checkbox (unchecked) and a dropdown menu set to 'WPSDB'. Below these are 'Schema name' fields and 'Database properties' text areas. In the System Bus section, the schema name is 'WPSMEClusterAppBus' and the database properties include 'databaseName=MEDB', 'driverType=4', 'serverName=localhost', 'portNumber=50000', 'description=SIB DataSource for SCA', 'traceLevel=', 'traceFile=', 'fullyMaterializeLobData=true', 'resultSetHoldability=2', 'currentPackageSet=', and 'readOnly=false'. In the Application Bus section, the schema name is 'WPSMEClusterSysBu' and the database properties are similar but include 'portNumber=50000'. At the bottom of the Application Bus section, the 'Create tables' checkbox is checked. A footer bar at the bottom contains a color bar, the text 'Clustering - Administration and configuration', the page number '27', and the copyright notice '© 2007 IBM Corporation'.

Paging down the configuration panel exposes the sections used for configuring the data source used by the message engines of the Service Component Architecture Application and System buses. As discussed previously, it is very important to specify a schema name that relates the cluster, the message engine and the bus. For example WPSMEClusterAppBus and WPSMEClusterSysBus.

Pay attention to the database properties too, especially the database name.

Notice that this is also the place where automatic creation of the database tables can be specified.

The WPSCluster configuration

- Hosts the WebSphere Process Server applications
 - Business processes, business rules, state machines, ...
- Uses remote service component architecture destinations

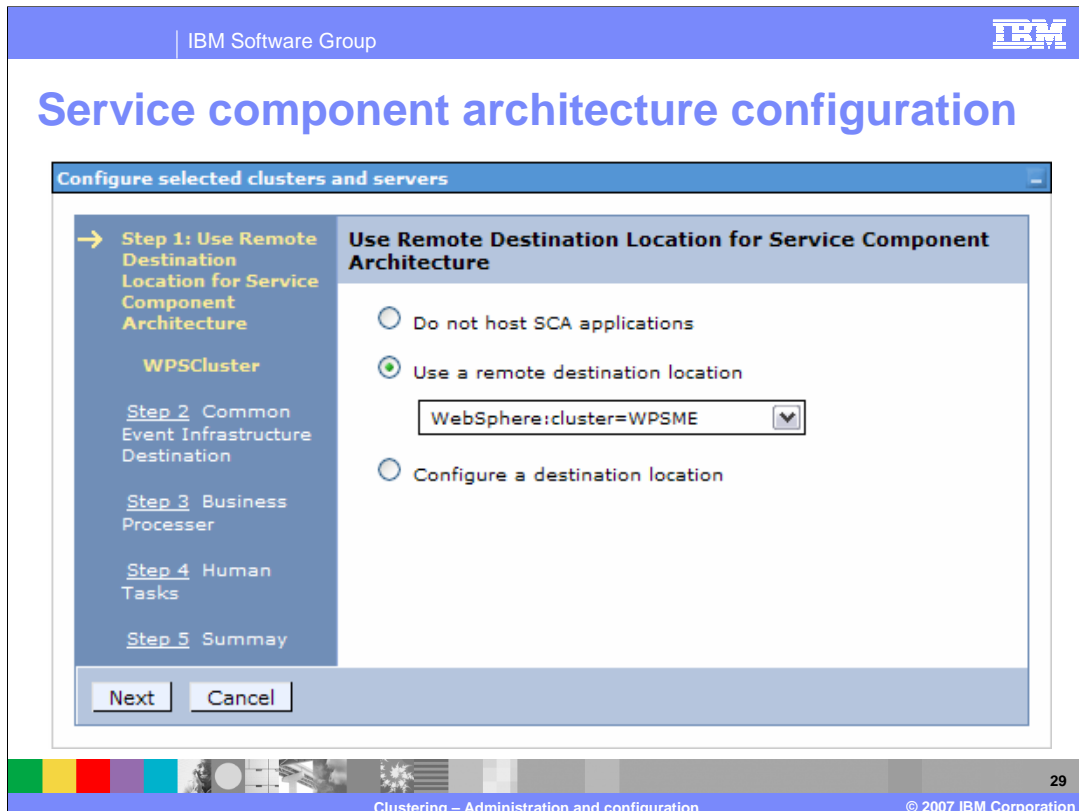
Select	Server/Cluster Name	Node	Setup SCA Destination		Business Rule Manager	Business Processes	Human Tasks
			Local	Remote			
<input checked="" type="checkbox"/>	WPSCluster		<input type="radio"/>	<input checked="" type="radio"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>

Total 1

Next Cancel

Because the WPSCluster will be hosting the Business Processes using Business Process Choreographer and Human Tasks, the Business Process and Human Task containers need to be installed and configured.

Although there may be some applications running in the WPSCluster that use Business Rules, the Business Rules Manager is not installed here because it is a separate application used to manage the business rule entities at runtime.



After filling out the database template, the Business Integration Configuration Wizard begins a 5-step guided activity based on the selections entered on the initial input panel.

Step 1 is to configure the remote destination used by the Service Component Architecture. Since the destination is remote the administrator has to select from a list of available destinations. This means that the server or server cluster hosting the message engine, that is to say the local destination, must be created and configured before configuring the remote destination.

For the 4-cluster topology presented here, this means that the WPSMECluster must be created and configured before the WPSCluster.

Common event infrastructure configuration

- Used by service component architecture applications.

The screenshot shows a configuration window titled "Configure selected clusters and servers". On the left is a navigation pane with five steps: Step 1 (Use Remote Destination Location for Service Component Architecture), Step 2 (Common Event Infrastructure Destination), Step 3 (Business Processer), Step 4 (Human Tasks), and Step 5 (Summay). Step 2 is selected and highlighted. The main area is titled "Common Event Infrastructure Destination" and contains the following options:

- Enable service at server startup
- Events Infrastructure emitter factory JNDI name
 - com/ibm/events/configuration/emitter/Admin
 - [Empty text box]

At the bottom of the window are three buttons: "Previous", "Next", and "Cancel".

30

Clustering – Administration and configuration

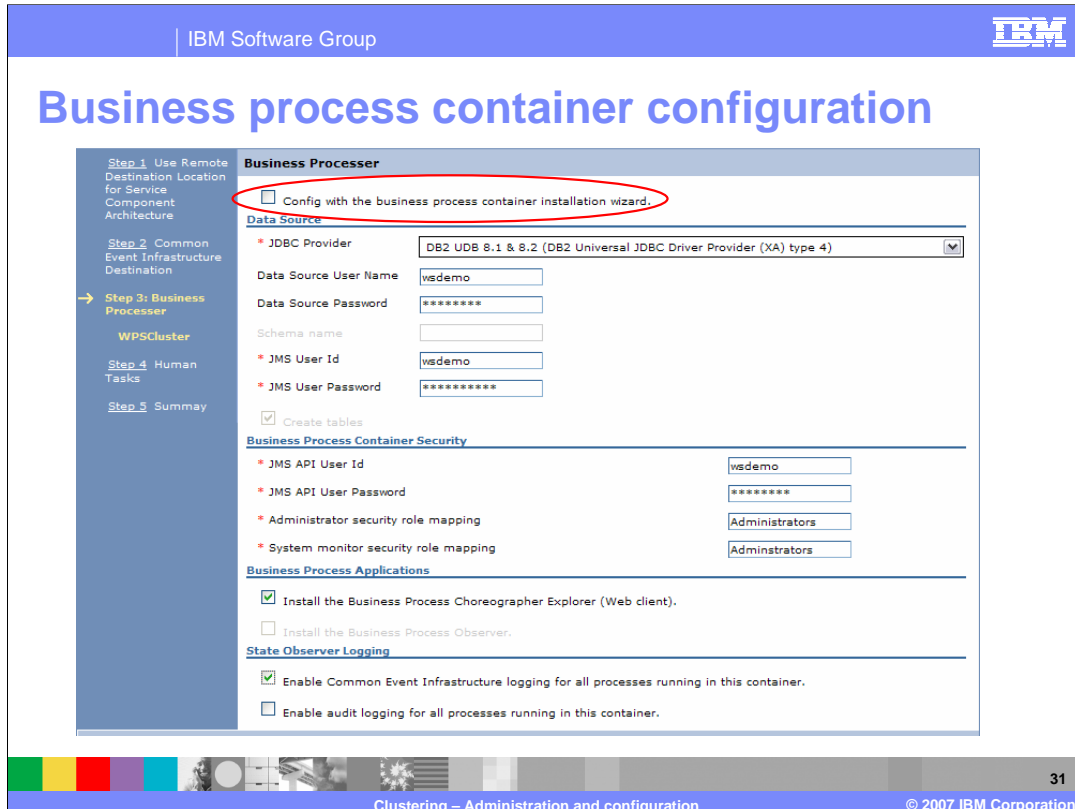
© 2007 IBM Corporation

Step 2 provides the administrator the opportunity to configure the resources need for the Common Event Infrastructure.

This is the client side configuration necessary for the Service Component Architecture applications that emit CEI events.

The applications running in the WPSCluster need to be able to locate the emitter factory so they can generate and send the events to the Common Event Infrastructure server, wherever it happens to reside.

The second radio button can be used to specify the JDI name if the Common Event Infrastructure server has not yet been configured.



With step 3 comes the opportunity to install and configure the Business Process Container. The very first option on the page is a checkbox indicating whether to use the information on this panel for configuring the Business Process Container or to defer the installation until later and use the Process Container Installation Wizard. The Process Container Wizard is invoked from the same place as the Business Integration Configuration Wizard, off the server or server cluster configuration page.

Checking this option will disable all the fields on this panel.

When configuring the Business Process Container from this panel, defaults will be used. If the system administrator requires more control over the specification of the container properties, then the option to use the Business Process Container Wizard should be selected.

Example:

The administrator may need to use a database name other than BPEDB, or corporate guidelines dictate a special context root for the BPC Explorer.

Human task container configuration

Human Tasks

Configure with the human task container installation wizard.

JMS Provider

* Jms User Id: wsdemo

* Jms User Password: *****

Human Task Manager Security

* Escalation user ID: wsdemo

* Escalation password: *****

* Administrator security role mapping: Adminstrators

* System monitor security role mapping: Adminstrators

Mail session for Human Task Manager

Mail session for Human Task Manager

State Observer Logging

Enable Common Event Infrastructure logging for all processes running in this container.

Enable audit logging for all processes running in this container.

Previous Next Cancel

Step 4 is the high level, minimal input, option for configuring the Human Task Container. Here again, the first option is to use the information from this panel or to defer the installation and configuration to the Human Task Container wizard where there will be more options available.

Summary for the WebSphere Process Server cluster configuration

Step 1 Use Remote Destination Location for Service Component Architecture

Step 2 Common Event Infrastructure Destination

Step 3 Business Processer

Step 4 Human Tasks

→ Step 5: Summary

Summary

Remote	
Use a remote destination location	WebSphere:cluster=WPSME
Common Event Infrastructure Destination	
Enable service at server startup	true
Events Infrastructure emitter factory JNDI name	com/ibm/events/configuration/emitter/Admin
Business Processes	
JDBC Provider	DB2 UDB 8.1 & 8.2 (DB2 Universal JDBC Driver Provider (XA) type 4)
Data Source User Name	wsdemo
Data Source Password	*****
JMS User Id	wsdemo
JMS User Password	*****
JMS API User Id	wsdemo
JMS API User Password	*****
Administrator security role mapping	Administrators
System monitor security role mapping	Administrators
Install the Business Process Choreographer Explorer (Web client).	true
Install the Business Process Observer.	false
Enable Common Event Infrastructure logging for all processes running in this container.	true
Enable audit logging for all processes running in this container.	false
Human Tasks	
Escalation user ID	wsdemo
Escalation password	*****
Administrator security role mapping	Administrators
System monitor security role mapping	Administrators
Jms User Id	wsdemo
Jms User Password	*****
Mail session for Human Task Manager	on
Enable Common Event Infrastructure logging for all processes running in this container.	true
Enable audit logging for all processes running in this container.	false

Clustering – Administration and configuration
© 2007 IBM Corporation

The final step is the summary page, which provides one last chance to cancel (the buttons are not shown here) before committing the configuration.

Study it carefully. It is a good practice to take a do a screen capture of this information and save it for the record.

Business integration configuration: Summary

- Can configure multiple server clusters at once if the same options can be used for all the clusters
- Used for specifying the JNDI name for the common event infrastructure emitter factory.
- Used for installing the business rules manager application
- 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

The Business Integration Configuration Wizard is a key tool for setting up sever clusters correctly. Based on initial input form an administrator, it will create a task list and guide the administrator through the necessary steps for the selected options.

In some cases, such as the with the business process and human task containers, it may be necessary to specify non-default options. This is accomplished by deferring the installation and configuration to the Business Process and Human Task container installation wizards.

It is also important to note that the Business Process Integration wizard does not do everything that is need for a complex topology. The administrator must still have a solid knowledge of the Service Component Architecture and the service integration buses and know how to configure them for other applications that use them, such as the Common Event Infrastructure and the Business Process Choreographer.

Summary

- WebSphere Application Server Network Deployment 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
- Business integration configuration wizard
 - ▶ Aids in WebSphere Process Server specific tasks

Much of the process for setting up a clustered WebSphere Process Server environment is similar to setting up a WebSphere Application Server Network Deployment V6 cell and with server clusters.

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 Business Integration Configuration Wizard is provided as an aid to manage the WebSphere Process Server unique tasks involved in setting up a WebSphere Process Server cell.

Feedback

Your feedback is valuable

You can help improve the quality of IBM Education Assistant content to better meet your needs by providing feedback.

- Did you find this module useful?
- Did it help you solve a problem or answer a question?
- Do you have suggestions for improvements?

[Click to send e-mail feedback](#)



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

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 WebSphere

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

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

Product data has been reviewed for accuracy as of the date of initial publication. Product data is subject to change without notice. This document could include technical inaccuracies or typographical errors. IBM may make improvements or changes in the products or programs described herein at any time without notice. Any statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only. References in this document to IBM products, programs, or services does not imply that IBM intends to make such products, programs or services available in all countries in which IBM operates or does business. Any reference to an IBM Program Product in this document is not intended to state or imply that only that program product may be used. Any functionally equivalent program, that does not infringe IBM's intellectual property rights, may be used instead.

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 (for example, IBM Customer Agreement, Statement of Limited Warranty, International Program License Agreement, etc.) under which they are provided. Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products in connection with this publication and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products.

IBM makes no representations or warranties, express or implied, regarding non-IBM products and services.

The provision of the information contained herein is not intended to, and does not, grant any right or license under any IBM patents or copyrights. Inquiries regarding patent or copyright licenses should be made, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. All customer examples described are presented as illustrations of how those customers have used IBM products and the results they may have achieved. The actual throughput or performance that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput or performance improvements equivalent to the ratios stated here.

© Copyright International Business Machines Corporation 2007. 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.