



Isabelle Mauny Simon Kapadia

Configuring IBM WebSphere Process Server V6.1 with an Oracle Database

This IBM® Redpaper publication explains how to configure IBM WebSphere® Process Server V6.1 to work with an Oracle® database. It discusses only the database-specific configuration and assumes that readers have some experience with WebSphere Process Server as well as an understanding of the concepts that we discuss.

Introduction

IBM WebSphere Process Server requires a Relational Database system (RDBMS) such as IBM Cloudscape®, IBM DB2® Universal Database[™] (DB2 UDB), or Oracle Corporation Oracle Database Server (from this point forward referred to as *Oracle*). The default configuration for WebSphere Process Server uses Cloudscape, which is a full-featured, robust, small-footprint database server that is simple to deploy and available with a no-charge license from IBM. You can also install WebSphere Process Server using an existing RDBMS, such as Oracle.

This paper explains how to configure WebSphere Process Server with a production implementation of Oracle. In a production environment, the person or team who installs WebSphere Process Server often does not have SYSTEM access to the Oracle server. Therefore, we cannot assume that the user who installs WebSphere Process Server has SYSDBA level access to the Oracle Database that is used.

You need to carry out certain steps as the SYS user, but these steps are defined clearly and can be passed on to the Oracle database administrator (DBA) to evaluate and run. This paper explains at each stage the commands and scripts that you need to run as the SYS user and the permissions that need to be granted to the users who are created for WebSphere Process Server to make the process of configuring Oracle as easy as possible to the user.

Table 1 lists the notations when configuring the required Oracle databases:.

Table I Notations	
(ORA) As SYS	Commands or scripts that need to be run as the SYS user. Pass these commands to the DBA who is responsible for configuring the database.
(ORA) As <user></user>	Commands or scripts that need to be run as the Oracle user <i><user></user></i> . Run these commands either on the Oracle database server itself or on a client that has the Oracle client or the Oracle instant client installed.

Table 1 Notations

Scope of this paper

This paper describes how to configure WebSphere Process Server with Oracle. However, it does not explain how to enable WebSphere Security for WebSphere Process Server or discuss how to create clusters.

WebSphere Process Server ships with WebSphere Application Server Network Deployment. For the purposes of this paper, we configured a single, non-clustered WebSphere Process Server instance. For a detailed explanation of clustering WebSphere Process Server, see the developerWorks® article *Building clustered topologies in WebSphere Process Server V6.1*, which is available at:

http://www.ibm.com/developerworks/websphere/library/techarticles/0803_chilanti/080
3_chilanti.html

This paper explains in detail how clustering should take place. You can use the techniques that we present in this paper just as well with Oracle as with DB2.

WebSphere security is required for the secure use of WebSphere Process Server functionality. In addition, the Human Tasks functionality requires WebSphere security to be configured for authentication. Typically, customers configure WebSphere security with a Lightweight Directory Access Protocol (LDAP) server, and various user IDs and passwords are required to access the messaging buses in a secure manner.

This paper does not explain how to enable and configure WebSphere security. For an explanation of how to configure WebSphere Process Server security, see the developerWorks article *WebSphere Process Server security overview*, which is available at:

http://www-128.ibm.com/developerworks/websphere/library/techarticles/0602_khangaon
kar/0602_khangaonkar.html

The topology for this paper

This paper is based on a very simple topology, shown in Figure 1.



Figure 1 Topology diagram

In this topology, there are two systems:

- ▶ The system named *eagle* is running Red Hat Enterprise Linux® 5 (RHEL5).
- ▶ The system named *oracle10g* is running Windows® XP.

For our test environment for this paper, we installed WebSphere Process Server on the Linux system and Oracle 10g (10.2.0.1 version) on the Windows system.

Connectivity choices

Throughout this paper, we use JDBC[™] connectivity to the Oracle database rather than "native" Oracle Call Interface (OCI). Although OCI can provide better performance in certain circumstances and although Oracle recommends OCI, in general customers tend to prefer JDBC as a simple and standardized connection protocol. The JDBC client is simply a set of classes in a JAR file and is, therefore, very simple to install and use. A simple JDBC Uniform Resource Locator (URL) is provided, which contains all of the information that is required to connect.

The OCI, or *thick*, client requires installation of a software component on the database client machines and the correct configuration of Oracle connectivity. All of the techniques and procedures that we discuss in this paper work in a similar way using the OCI *thick* Oracle client.

High-level overview

At a high-level, the following steps are required to configure WebSphere Process Server with Oracle:

- 1. Prepare the environment.
- 2. Create the database users and tables.
- 3. Configure the databases for XA and the Oracle JVM[™].
- 4. Install the WebSphere Process Server software.

- 5. Create a WebSphere Process Server profile.
- 6. Configure the messaging engines data stores.
- 7. Configure the Common Event Infrastructure (CEI) database.
- 8. Configure the Business Process Choreographer container.

Preparing the environment

To prepare the environment, you need to:

1. Download the latest supported version of the Oracle JDBC drivers from the Oracle Web site at:

http://www.oracle.com/technology/software/tech/java/sqlj_jdbc/index.html

 Then, copy these drivers to the WebSphere Process Server system. We chose to copy the drivers to the /opt/ProcServer1/OracleDrivers directory. Oracle 10.2.0.3 was used during the creation of this paper. With this release, you can use the same drivers for Oracle 9iR2(fn) as shown in the following example.¹

[root@eagle OracleDrivers]# pwd /opt/ProcServer61/OracleDrivers [root@eagle OracleDrivers]# ls -1 ... root 1545954 Feb 29 12:15 ojdbc14.jar ... root 1938906 Feb 29 12:16 ojdbc14 g.jar

Process Server components use, as follows:

- 3. Eventually, share a file system across both systems where database configuration scripts are generated (for example /opt/shared) to avoid copying the scripts manually across
- systems. 4. Plan for the minimum space that is required by Oracle table spaces that WebSphere
 - CEI: 300 MB
 - Business Process Choreographer Container: 1.3 GB for non-development/POC environments.
 - Business Process Choreographer Observer: 600 MB for non-development/POC environments.

Creating the database and users

WebSphere Process Server needs a number of database tables. You can configure the database tables using different methods, and each customer has different requirements and standards that can affect this configuration.

One decision that you need to make is whether to configure everything in a single database, using different schemas to separate the tables, or to configure many separate databases. The first option has the advantage of ease of configuration, especially for many independent environments, while the second option allows for individual tuning of the databases. You should make this decision after considering all of the relevant factors at a given site.

For the purposes of this paper, we configured everything in a single Oracle database. This configuration was an initial configuration, that in a real-world environment, we would probably

¹ Oracle 11g is not supported officially at this time of writing (July 08). For more information, see: http://www-306.ibm.com/software/integration/wps/sysreqs/

modify after reviewing the performance of the databases. You need to complete performance testing with a configuration of WebSphere Process Server, load, and hardware that is similar to the actual production environment. After monitoring the performance testing at all levels (Java[™], Operating System, database, any back-end applications, and so forth), you can make a decision on a production configuration. For example, the load testing might show that it is prudent to create a separate system identifier (SID) for the data stores for the messaging engines, a second SID for all of the WebSphere Process Server internal tables, and a third SID for the Business Process Choreographer container itself. The only way to ascertain this information is by running real performance testing.

Creating the Oracle database

As mentioned previously, in our environment for this paper, we created a database using the database configuration assistant and referenced it by the WPSDB SID. You need to make sure that this database has Unicode support to accommodate all languages (which is needed by the Business Process Choreographer).

Configuring the Oracle JVM

You need to enable the Oracle JVM (or JServer in Oracle 9i terms) for the Oracle database. You can use the Oracle Database Configuration Assistant or specific SQL scripts. See the Oracle documentation for details, which is available at:

http://www.oracle.com/technology/documentation/index.html

You can check whether the Oracle JVM/JServer component is installed by invoking the following query from SQL*Plus, as SYS:

- ► (Oracle 9i) Run describe dbms_java.
- Oracle 10g) select version, status from dba_registry where comp_id='JAVAVM'

In addition, you need to grant the correct privilege to run Java to each of your users. If you do not, you will get XA errors. To grant privileges, execute the following command:

grant javauserpriv to <user>;

This is documented in "Creating the database users and schemas" on page 6 for each database user.

Configuring XA

XA is a two-phase commit protocol that is defined by the X/Open DTP group. The database that WebSphere Process Server uses must be XA enabled. To enable XA, the DBA needs to complete these steps:

- Run the <ORACLE_HOME>/javavm/install/initxa.sql script on the WPSDB database. This script configures the database for XA. More information about this requirement is available in the article Configuring and using XA distributed transactions in WebSphere Studio (see "Resources" on page 32).
- 2. Execute the following queries:

```
grant select on pending_trans$ to public;
grant select on dba_2pc_pending to public;
grant select on dba_pending_transactions to public;
grant select on V$XATRANS$ TO PUBLIC;
```

Tip: If the last query fails, the V\$XATRANS\$ view does not exist. You need to prepare the Oracle database for XA transactions recovery by executing the xaview.sql script that is located in the *<ORACLE_HOME>*/RDBMS/ADMIN directory. Use the following command:

sqlplus SYS/xxxxx@WPSDB @xaview.sql

Additionally, you need to grant each user access to dbms_system as illustrated in the next section, "Creating the database users and schemas".

Creating the database users and schemas

Table 2 lists the database schemas that are required for WebSphere Process Server, along with the names that are used for the user IDs (schemas) in the configuration that we present in this paper.

Schema and user ID used in this paper	Usual name or default name	Purpose of this user
WPCDB	WPRCSDB	Main Process Server configuration
CEIDB	CEI	Common Event Infrastructure configuration
SCASYSMSG	First 3 letters of instance SID + SS00	SCA "System" Bus messaging engine data store configuration
SCAAPPMSG	First 3 letters of instance SID + SA00	SCA "Application" Bus messaging engine data store configuration
CEIMSG	First 3 letters of instance SID + CM00	CEI messaging engine data store configuration
ESBLOG	ESBLOG	ESB Messaging Mediation data store configuration
BPCMSG	First 3 letters of instance SID + BM00	Business Process Choreographer Bus data store configuration
BPCDB	First 3 letters of instance SID + BE00	Business Process Choreographer container configuration
BPCOBS ¹	First 3 letters of instance SID + BC00	Business Process Choreographer Observer
¹ This user is necessary only if you intend to install the Business Process Choreographer Observer.		

Table 2 Database tables

Tip: You can change the user IDs according to your naming conventions. We deliberately used non-default user names for this paper to demonstrate that it can be done, but you can also use the default names if you want. You have to adapt the instructions in this paper to reflect any name changes.

By default, in Oracle terms, a *schema* is synonymous with a *user ID*. That is, if you make a JDBC connection and authenticate as a given user, you have access to the schema with the same name as the user.

According to the Oracle 10g documentation, the connect role and resource role are defined as follows:

- ► Connect role: CREATE SESSION
- Resource role: CREATE CLUSTER, CREATE INDEXTYPE, CREATE OPERATOR, CREATE PROCEDURE, CREATE SEQUENCE, CREATE TABLE, CREATE TRIGGER, CREATE TYPE

Common database user (WPCDB)

The WebSphere Process Server configuration database user, WPCDB, has no special requirements other than it must exist. To create it, your DBA needs to run a command equivalent to the following:

(ORA) As SYS

create user WPCDB identified by passwOrd; grant connect, resource, unlimited tablespace to WPCDB; grant execute on dbms_system to WPCDB; grant javauserpriv to WPCDB;

These commands use the system's default table space. Your DBA might want to specify a different table space and this is fine.

Business Process Choreographer Container User (BPCDB)

To create the BPC container database user, the DBA needs to run a command similar to the following:

(ORA) As SYS

create user BPCDB identified by passwOrd; grant connect, resource, create tablespace, drop tablespace, create view, unlimited tablespace to BPCDB; grant execute on dbms_system to BPCDB; grant javauserpriv to BPCDB;

We describe the creation of table spaces and tables in "Creating the BPC table spaces and tables" on page 27.

Business Process Choreographer Observer user (BPCOBS)

The BPCOBS user is optional. You need it only if you intend to install the BPC Observer.

To create the BPC Observer user, you need to execute the following:

(ORA) As SYS

create user BPCOBS identified by passwOrd; grant connect, resource, create tablespace, drop tablespace, create view, unlimited tablespace to BPCOBS; grant execute on dbms_system to BPCOBS; grant javauserpriv to BPCOBS;

We describe the creation of table spaces and tables in "Creating the BPC table spaces and tables" on page 27.

Common Event Infrastructure user

The Common Event Infrastructure (CEI) scripts that are used to create the entire infrastructure, including users, roles, tables, and views are generated as part of the application server profile creation step. You can also execute these scripts as part of the profile creation process, but this method requires that you provide a user ID with SYSDBA privileges at installation time, which few DBAs will allow you to do.

We describe the CEI setup in details in "Configuring the CEI database" on page 23.

ESB Message Logger user (ESBLOG)

The Message Logger ESB primitive uses the ESBLOG.MSGLOG table to log messages.

Tip: By default, the SQL configuration script for the MSGLOG table uses the ESBLOG schema. (You cannot change this from the installation wizard. Nevertheless, you can configure a different user manually. If you want to use a different user or schema, you must change the instructions in this section accordingly, and then follow the instructions in "Changing the default ESBLOG schema" on page 19).

To create the ESB message logger user, you need to execute the following commands:

(ORA) As SYS

create user ESBLOG identified by passwOrd; grant connect, resource, unlimited tablespace to ESBLOG; grant execute on dbms_system to ESBLOG; grant javauserpriv to ESBLOG;

Messaging engines data stores

Each messaging engine uses the same table names (SIB000, SIB001, and so forth). Therefore, you need to use a specific schema (that is, user) for each messaging engine, as follows.

(ORA) As SYS

create user SCASYSMSG identified by passwOrd; grant connect, resource, unlimited tablespace to SCASYSMSG; grant execute on dbms_system to SCASYSMSG; grant javauserpriv to SCASYSMSG;

create user SCAAPPMSG identified by passwOrd; grant connect, resource, unlimited tablespace to SCAAPPMSG; grant execute on dbms_system to SCAAPPMSG; grant javauserpriv to SCAAPPMSG;

create user CEIMSG identified by passwOrd; grant connect, resource, unlimited tablespace to CEIMSG; grant execute on dbms_system to CEIMSG; grant javauserpriv to CEIMSG;

create user BPCMSG identified by passwOrd; grant connect, resource, unlimited tablespace to BPCMSG; grant execute on dbms_system to BPCMSG; grant javauserpriv to BPCMSG;

Installing WebSphere Process Server

For the purposes of this exercise, WebSphere Process Server V6.1 is installed on the *eagle* system (see Figure 1 on page 3). Version 6.1 is provided as a complete "umbrella installation" image, and it uses the WebSphere Install Factory to install the prerequisite version of WebSphere Application Server Network Deployment, V6.1.0.13.

To start the installation using the graphical installer:

1. Run ./launchpad.sh from the WebSphere Process Server installation CD. In this example, we choose to use a custom installation path of /opt/ProcServer61, rather than the default path of /opt/ibm/WebSphere/ProcServer (primarily for convenience), shown in Figure 2.



Figure 2 Choosing a custom install path

2. When prompted for the type of installation that you want to pursue, select **None** as shown in Figure 3. You then use the profile manager to create a specific profile type.



Figure 3 Do not choose to run the profile creation wizard

3. You receive a warning similar to the one shown in Figure 4 that prompts you to confirm your choice because WebSphere Process Server is non-functional without a profile. Confirm by clicking **Yes**.



Figure 4 Proceed without creating a profile

Fix pack installation

At the time of writing, WebSphere Process Server Fix Pack 1 is available, and we highly recommend that you install it. You can download the fix pack from the following support page:

▶ WebSphere Process Server Universal Test Client V6.1.0 Fix Pack 1 (6.1.0.1)

http://www-1.ibm.com/support/docview.wss?rs=2307&context=SSQH9M&dc=D400&uid=swg 24018729&loc=en US&cs=UTF-8&lang=en&rss=ct2307websphere

To install fixes, you must have installed the IBM Update Installer, which is available on the WebSphere Process Server installation CD. You can also download the installer from the support site (see "Resources" on page 32).

Silent installation

You can also install WebSphere Process Server using a silent installation mode, by providing a response file that contains the required information to determine the components to install and the locations in which to install them. This method is documented in the WebSphere Process Server Information Center (see "Resources" on page 32).

Creating a profile (graphical version)

You need to create a WebSphere Process Server profile by starting the profile management tool with the following command:

/ibm/ProcServer/bin/ProfileManagement/pmt.sh

This command opens a GUI that lets you define the profile that you want to create.

Basic configuration

To configure the profile:

1. Select to create a WebSphere Process Server profile (as shown in Figure 5), and click **Next**.

Select the type of environment to create. Environments:	
Cell (deployment manager and a federated application server)	
Deployment manager	
Application server	
Custom profile	
WebSphere Enterprise Service Bus	
WebSphere Process Server	

Figure 5 Select a WebSphere Process Server profile

2. Select Stand-alone process server profile (Figure 6), and click Next.

elect a profile type for the WebSphere Process Server environment. rofile Types:
Deployment manager profile
Stand-alone process server profile
Custom profile

Figure 6 Select stand-alone profile

3. Select **Advanced profile creation** (as shown in Figure 7) to allow you to choose all database options (otherwise, the installation is done with Derby).

Choose the profile creation process that meets your needs. Pick the Typical option to allow the Profile Ma configuration values to the profile. Pick the Advanced option to specify your own configuration values for Jypical profile creation Create a process server profile that uses default configuration settings. The Profile Management to node, and cell. The tool also assigns unique port values. The administrative console and the defau optionally select whether to enable administrative security. The tool might create a system service on the operating system of your machine and the privileges assigned to your user account. All req Embedded. Create a process server profile using default configuration settings, or specify your own configurat port values. You can optionally choose whether to deploy the administrative console, Sample appl definition. You might have the option to run the process server as a system service depending on t and the privileges assigned to your own configuration values for database and Common database. You can optionally choose to configure the sample Business Pro Rules Manager.

Figure 7 Choose to do an advanced installation

4. Choose the applications that you want to deploy into the newly created server. We recommend that you install the administrative console.



Figure 8 Installing default applications

5. Provide a name and a directory for the profile, as shown in Figure 9. We renamed the profile to *WPS61Oracl*. You can also choose to profile this server for development, which requires less powerful hardware.

pecify a profile name and directory path to contain the files for the run-time environment, such as commar lick Browse to select a different directory.
rofile name:
WPS61Oracl
Profile <u>d</u> irectory:
/opt/ProcServer61/profiles/WPS61Oracl
Create the server using the development template.
Select this option to create a server using configuration settings optimized for development. The devel and allows the server to run on less powerful hardware. Do not use this option for production servers.
☐ <u>M</u> ake this profile the default.
Each installation of WebSphere Process Server always has one default profile. Commands that run with default profile. Select this option to make this profile the new default.

Figure 9 Choose profile options

6. Select a node and host name. The node name is the name by which this node referred in the administrative console. The host name is the name of the host on which the node runs. See Figure 10.

Specify a node name, a host name, and a cell name for this profile.	
N <u>o</u> de name:	
WPS61EagleNode1	
<u>H</u> ost name:	
eagle.mad.es.ibm.com	
C <u>e</u> ll name:	
WPS61EagleCell1	
-	

Figure 10 Select a node and host name

7. Define a default user name and password for this server, as shown Figure 11. You can change the security setup later to use a specific LDAP server for example.

User name:	unty	
wpsadmin		
Password:		
Confirm password:		

Figure 11 Enabling security

8. Define the various port values Figure 12. For the purposes of this paper, we used the default port numbers.

The values in the follo of WebSphere Applica	wing fields define the ports for the proc tion Server, WebSphere Process Server,	ess server and do not con WebSphere Enterprise Se
avoid run-time port co	nflicts, verify that each port value is uni	ique.
	Recommended Port Valdes	
<u>A</u> dministrative conso	le port (Default 9060):	9061 🖨
Administrative <u>c</u> onsole secure port (Default 9043):		9044
HTTP transport port (Default 9080):		9081
H <u>T</u> TPS transport port (Default 9443):	9444
Bootstrap port (Defau	lt 2809):	2810 🚔

Figure 12 Port numbers for the application server

The next screens let you decide whether you want to install this server as a service and define an HTTP server.

CEI and common database configuration

You then reach the CEI configuration. Follow these steps:

1. You need to choose the database type that CEI will use. You need to specify Oracle as the database product, a folder where the CEI database configuration scripts are generated, as well as the database name, in our case *WPSDB*, as shown in Figure 13 on page 14.

The CEI database configuration scripts can be executed as part of the installation process. However, this methods work only if you supply a user with administrative privileges (that is, the ability to create users and roles). Most DBAs are not willing to give such access to users, especially when installing production systems. Therefore, for our example, we chose to delay the execution of the scripts, which we will do later manually.

Figure 13 CEI Database configuration – Part 1

2. Next, you provide the Oracle user a password, as well as the JDBC drivers path, as shown in Figure 14. Although in this example, we plan to run the scripts manually, this information needs to be correct because it is used to create the JDBC data sources to connect to the CEI database.

Additional information is required to complete configuration for the Oracle database.	
User name to authenticate with the database:	
CEIDB	
Password for database authentication:	
•••••	
Confirm password:	
•••••	
Location (directory) of JDBC driver classpath files:	
/opt/ProcServer61/OracleDrivers	

Figure 14 CEI Database configuration – Part 2

3. Now, supply the SID of the instance (WPSDB) as well as the Oracle server location and TNS listener port, as shown in Figure 15. You can use values that are not real for the Oracle admin user and password, because these values are not used to execute the database scripts.

atabase server <u>h</u> ost name (for example IP address):
oracle10g
erver por <u>t</u> :
1521
vent service instance name:
WPSDB
<u>d</u> min user name:
bogus
ass <u>w</u> ord:
onfir <u>m</u> password:
•••••

Figure 15 CEI System user details

- 4. On the next panel, do not configure a sample Business Process Choreographer. We configure the Business Process Choreographer container later, using a wizard in the WebSphere Administrative Console, as the default configuration uses a Cloudscape database.
- 5. Next, decide whether you need the Business Rules manager and then provide details about the WebSphere Process Server common database configuration. Select to use an existing database, choose Oracle 10g for the database product, and supply the SID as the database name.

6. Similarly to CEI, you can change the destination directory for the database scripts and choose to delay their execution, which we chose to do as shown in Figure 16. You can also let the profile creation wizard configure the common database for you. In that case, make sure the WPSDB user ID and password are correctly set. Otherwise, the profile creation will fail while validating the database connectivity.

Various components use WebSphere Process Server common database. Choose a database type and enter
Choose a database product:
Oracle10g
Override the destination directory for generated scripts. Database script output
/opt/shared/WPSDB
 Create a new local database. Use an existing database. Database <u>n</u>ame:
WPSDB
 Delay execution of database scripts for new or existing database. Use a file store for Messaging Engines (MEs). Use this database for Messaging Engines (MEs).

Figure 16 Common database configuration

7. Provide the user ID and password that will be used for the common WPS database. Also, supply the location of the JDBC driver files and the host name of the database as shown in Figure 17.

Additional information is required to complete configuration for the Oracle10g database
Addend mennader brequied to complete comparater for the ordering database.
User name to authenticate with the database:
WPSCDB
Password for database authentication:
Con <u>f</u> irm password:
Location (directory) of JDBC driver classpath files:
/opt/ProcServer61/OracleDrivers
JDBC driver type:
O oci
O Thin
Database server <u>h</u> ost name (for example IP address):
oracle10g
Server port:
1521

Figure 17 Common database details

This concludes the information that you need to supply to the profile creation wizard. This entire process should take around 10 to 15 minutes. Do not start the application server yet, because you have to complete the database configuration first.

What you have created thus far

So far, you have created the following artifacts:

- 1. The definition of an application server that can host business processes and mediations.
- 2. Definitions of JDBC data sources that will allow this server to connect to its main database as well as the CEI event database. Those data sources are usable as is.
- Definitions of JDBC data sources for the messaging engine data stores, which you need to modify to take into account the users that you have created.

You also generated scripts in /opt/shared for CEI and, eventually, the common WebSphere Process Server database.

Common database setup

If you chose not to delay database execution scripts during the profile setup, skip to the next step, which is described in the next section "Starting the application server". Otherwise, you need to execute the scripts in the following order:

```
(ORA) As WPCDB
```

```
sqlplusWPCDB/passwOrd@WPSDB@createTable_AppScheduler.sqlsqlplusWPCDB/passwOrd@WPSDB@createTable_CommonDB.sqlsqlplusWPCDB/passwOrd@WPSDB@createTable_customization.sqlsqlplusWPCDB/passwOrd@WPSDB@createTable_lockmanager.sqlsqlplusWPCDB/passwOrd@WPSDB@createTable_mediation.sqlsqlplusWPCDB/passwOrd@WPSDB@createTable_mediation.sqlsqlplusWPCDB/passwOrd@WPSDB@createTable_Recovery.sqlsqlplusWPCDB/passwOrd@WPSDB@createTable_RelationshipMetadataTable.sqlsqlplusWPCDB/passwOrd@WPSDB@insertTable_CommonDB.sql
```

Starting the application server

Because you have now created the common database infrastructure, you can start the server using the following command:

<WPS_InstallRoot>/profiles/WPS610racl/bin/startServer.sh server1

You will see errors that are related to the messaging engines. You will correct these errors soon.

Configuring the Mediation Message Logger primitive

Message Logger primitives that are deployed in WebSphere ESB (which is part of WebSphere Process Server) by default use a table called ESBLOG.MSGLOG. The ESBLOG schema is hardcoded in the table creation script. You need to create this table and configure the data source which accesses it.

Database setup

You can choose to use the default ESBLOG schema name or change the default name.

Important: If you intend to create a trigger for this table, see: http://www-1.ibm.com/support/docview.wss?rs=2346&uid=swg1IC53621

Using the ESBLOG default schema name

You need to execute the createTable_EsbLoggerMediation.sql script as the ESBLOG user: (ORA) As ESBLOG

sqlplus ESBLOG/passwOrd@WPSDB @createTable_EsbLoggerMediation.sql

Changing the default ESBLOG schema

If you want to use a different SCHEMA name, such as MSGLOGGER, then you need to:

- Edit the createTable_EsbLoggerMediation.sql script and replace ESBLOG with MSGLOGGER.
- 2. Run the following script:

(ORA) As MSGLOGGER

sqlplus MSGLOGGER/passwOrd@WPSDB @createTable_EsbLoggerMediation.sql

- 3. Create a WebSphere variable called ESB_MESSAGE_LOGGER_QUALIFIER, and set it to the value of the schema you are using, here MSGLOGGER. Follow these steps:
 - a. Open the administration console.
 - b. Go to **Environment** \rightarrow **Variables** in the navigation menu.
 - c. Select a scope, and click New.
 - d. Create a variable as shown in Figure 18, and click **OK**.

Name			
ESB_MESSAGE_L	OGGER_QUAL		
/alue			
MSGLOGGER			
Description		1.00	
Schema for Medi	ation Logging Prir	mitive.	

Figure 18 ESB_MESSAGE_LOGGER-QUALIFIER variable

e. When the configuration completes successfully, save the configuration.

Changing the data source configuration

The default configuration uses the WPCDB, because you used this user ID for the common database setup. You, therefore, need to change the J2C Alias for the data source called *ESBLoggerMediationDatasource* as follows:

- 1. Start the administrative console and log in.
- 2. Go to **Resources** \rightarrow **JDBC** \rightarrow **Data sources** in the navigation menu.
- 3. Click the ESBLoggerMediation data source.
- 4. Click JAAS-J2C Authentication Aliases on the right.

5. Click **New**, and create an alias similar to Figure 19 (using the schema name that you decided to use, ESBLOG or other). Click **OK**.

Data sources > ESBLoggerMediationDataSource > JAAS - J2C authentication data > New
Specifies a list of user identities and passwords for Java(TM) 2 connector security to use.
Configuration
General Properties
* Alias
ESBLOG_Alias
* User ID
ESBLOG
* Password
••••••
Description
Userid for MSGLOG table

Figure 19 J2C authentication alias

- Click ESBLoggerMediationDataSource in the navigation menu to go back to the data source editor.
- 7. Select the alias that you just created as the Component-managed authentication alias, and click **OK**.
- 8. When the configuration completes successfully, save your configuration.

Configuring the messaging engines data store

Because you selected to use the common database for messaging engines, each messaging engine is configured to target Oracle. However, all data sources are using the WPCDB user ID, which is not what you want. You want each messaging engine to use its own user and schema. Therefore, for each messaging engine, you need to:

- 1. Update the messaging engine data store definition to use the right schema.
- 2. Update the J2C alias used to connect to the database to use the right user.

You also need to create the supporting tables for the data stores.

Configuring the data source and J2C alias

You have defined three service integration buses so far, and each has its own messaging engine, backed by a message store. We start with the CEI bus:

- 1. Open the administration console and log on.
- 2. Go to Service Integration \rightarrow Buses in the navigation menu.
- Click CommonEventInfrastructure_Bus. Then, click Messaging engines (under Topology).

 Select <NodeName>.server1-CommonEventInfrastructure_Bus. Click Message store. You see something similar to Figure 20.

General Properties	
UUID	
2B6D89017B14CB5B	
* Data source JNDI name	
jdbc/com.ibm.ws.sib/WPS61Node1.server1-CommonEventInfrastructure_Bus	
Schema name	
WPSCM00	
Authentication alias	
CEIME_WPS61Node1.server1_Auth_Alias	
✓ Create tables	

Figure 20 Message store configuration

- Change the schema to reflect the schema you created, in this case CEIMSG. Then, click OK to apply the changes.
- Click J2C Aliases. Select the alias the message engine is using, in this case CEIME_WPS61Node1.server1_Auth_Alias, and change the user ID and password to CEIMSG and its corresponding password. Click OK.
- 7. When the configuration completes successfully, save the configuration.
- 8. Repeat this process for the SCA SYSTEM and SCA APP buses, with respectively the SCASYSMSG and SCAAPPMSG user IDs / schemas.

Creating message stores tables

Tables can be created automatically at application server restart, or you can create them manually.

Creating tables automatically

Each message store has an option to create the tables in the database automatically the first time that the data store is accessed as shown in Figure 20. If you prefer to create the tables yourself, then clear the Create Tables option for each message store, and go to "Creating tables manually". Otherwise, you can skip to "Configuring the CEI database" on page 23.

Creating tables manually

To create the tables for each messaging engine manually, first you have to generate the DDL (SQL) using WebSphere Process Server SI Bus DDL Generator command. You need to adapt the arguments to your target Oracle platform. Use the following command:

/opt/ProcServer/bin/sibDDLGenerator.sh -system oracle -version 10g -platform windows -statementend \; -schema SCAAPPMSG -user scaappmsg > /opt/shared/createSIBSchema SCAAPPMSG.ddl

```
/opt/ProcServer/bin/sibDDLGenerator.sh -system oracle -version 10g -platform
windows -statementend \; -schema SCASYSMSG -user scasysmsg >
/opt/shared/createSIBSchema_SCASYSMSG.ddl
```

```
/opt/ProcServer/bin/sibDDLGenerator.sh -system oracle -version 10g -platform
windows -statementend \; -schema CEIMSG -user ceimsg >
/opt/shared/createSIBSchema_CEIMSG.ddl
```

You can edit the three files to prevent errors during execution. The DDL executes even if you do not edit the files as described here, but we think it is cleaner to proceed this way. Edit the files as follows:

- Starting with WebSphere Application Server V6.1.0.13, the sibDDLGenerator command echoes the command arguments to the console. Therefore, you see them at the beginning of each DDL file. This is a known issue that is documented in the following support technote:
 - PK60247: SIBDDLGENERATOR Command Mistakingly Includes the Arguments that were Passed to the Command

http://www-1.ibm.com/support/docview.wss?rs=180&context=SSEQTP&q1=sibDDLgene
rator&uid=swg1PK60247&loc=en US&cs=utf-8&lang=en

Edit the file to remove these extra lines before submitting it.

2. Each script creates an Oracle user, which you created previously. You can remove the following line that creates the user:

CREATE USER SCASYSMSG IDENTIFIED EXTERNALLY DEFAULT TABLESPACE USERS QUOTA UNLIMITED ON USERS ACCOUNT LOCK;

The scripts are designed to be run by a user with Oracle administrator's rights. For this
paper, we run the scripts as the user who needs to access the tables. Therefore, the user
has all the required permissions. You can remove the following GRANT queries from each
script:

GRANT SELECT,INSERT,UPDATE ON SCASYSMSG.SIBOWNER TO scasysmsg; GRANT SELECT,INSERT ON SCASYSMSG.SIBCLASSMAP TO scasysmsg; GRANT SELECT,INSERT ON SCASYSMSG.SIBLISTING TO scasysmsg; GRANT SELECT,INSERT,DELETE,UPDATE ON SCASYSMSG.SIB000 TO scasysmsg; GRANT SELECT,INSERT,DELETE,UPDATE ON SCASYSMSG.SIB001 TO scasysmsg; GRANT SELECT,INSERT,DELETE,UPDATE ON SCASYSMSG.SIB002 TO scasysmsg; GRANT SELECT,INSERT,DELETE,UPDATE ON SCASYSMSG.SIB002 TO scasysmsg; GRANT SELECT,INSERT,UPDATE,DELETE ON SCASYSMSG.SIBXACTS TO scasysmsg; GRANT SELECT,INSERT,UPDATE,DELETE ON SCASYSMSG.SIBXACTS TO scasysmsg; GRANT SELECT,INSERT,UPDATE ON SCASYSMSG.SIBKEYS TO scasysmsg;

Then, copy the scripts to the Oracle server, and run them against the database as the relevant users:

sqlplus scaappmsg/passwOrd@WPSDB @createSIBSchema_SCAAPPMSG.ddl sqlplus scasysmsg/passwOrd@WPSDB @createSIBSchema_SCASYSMSG.ddl sqlplus ceimsg/passwOrd@WPSDB @createSIBSchema_CEIMSG.ddl

Configuring the CEI database

You use the scripts that are generated during the profile creation to configure the CEI database. These scripts are created under the /opt/shared/CEI_orcl directory, and you need to execute them on the Oracle server system. Table 3 shows the relevant scripts.

Table 3 CEI database scripts

cr_ts.ora	Creates the CEI table spaces for event data
cr_ts_catalog.ora	Creates the CEI table spaces for the event catalog
cr_security.ora	Creates the role and profile
cr_tbl.ora	Creates the event tables
cr_tbl_catalog.ora	Creates the event catalog tables
ins_metadata.ora	Populates static data for the event database
catalogSeed.ora	Populates the base event catalog data
cr_stored_procedure.ora	Creates stored procedures

The Oracle DBA needs to run the scripts as the SYS user. All table spaces names are prefixed by the CEI instance name that you entered at profile creation, in this case WPSDB:

(ORA) As SYS

@cr_ts.ora
@cr_ts_catalog.ora
@cr_security.ora

Tip: The table spaces creation requires about 300 MB of disk space.

The DBA needs to create the user for the CEI database and to configure the role that the last of these scripts created.

(ORA) As SYS

```
CREATE USER CEIDB IDENTIFIED BY passwOrd default tablespace WPSDB_cei_ts_base
temporary tablespace WPSDB_cei_ts_temp quota unlimited on WPSDB_cei_ts_base
profile WPSDB_cei_profile;
grant WPSDB_cei_role, resource, create view to CEIDB;
alter user CEIDB default role WPSDB_cei_role;
grant execute on dbms_system to CEIDB;
grant javauserpriv to CEIDB;
```

Important: The DBA needs to adapt the table spaces names as well as the role used in this command to match the names found in cr_ts.ora and cr_ts_catalog.ora.

Finally, you need to run the remaining scripts as the CEIDB user:

(ORA) As CEIDB

```
sqlplus ceidb/passwOrd@WPSDB @cr_tbl.ora
sqlplus ceidb/passwOrd@WPSDB @cr_tbl_catalog.ora
sqlplus ceidb/passwOrd@WPSDB @ins_metadata.ora
sqlplus ceidb/passwOrd@WPSDB @catalogSeed.ora
sqlplus ceidb/passwOrd@WPSDB @cr stored procedure.ora
```

You now have a working CEI database.

Important: Make sure that the Oracle10g data store helper is used for each CEI-related data source, that is jdbc/event and jdbc/eventcatalog, using the following process.

To make sure that the Oracle10g data store helper is used:

- 1. Open the administrative console and go to **Resources** \rightarrow **JDBC** \rightarrow **Data Sources**.
- 2. For each of the data sources, jdbc/event and jdbc/eventcatalog, set the data store helper class to Oracle10g, as shown in Figure 21.



Figure 21 Data store helper class for Oracle 10G

- 3. Click **OK**.
- 4. When the configuration completes successfully, save the configuration.

Restarting the application server

Because you have made many changes to the server and database configurations, you need to restart the server and check the logs to make sure that no database errors occur.

You can stop the application server using the following command:

<WPS_InstallRoot>/profiles/WPS610racl/bin/stopServer.sh server1

When the server has started, if you have chosen to let the server create the messaging engine data stores, make sure that the tables are created correctly. Also, make sure that the messaging engines start correctly as shown in Figure 22.

Buses :	> CommonEventInfrastructure Bus > Messaging engir	nes	
A messa member	aging engine is a component, running inside a server, that n r. Applications are connected to a messaging engine when a	nanages messaging accessing a service	resources for a bus integration bus.
[⊞] Prefe	rences		
Start Stop mode: Immediate Stop			
Select	Name [‡]	Description 🗘	Status 🗘 👲
Г	WPS61Node1.server1-CommonEventInfrastructure Bus Image: CommonEventInfrastructure Bus Image: Common		
Total 1			

Figure 22 Messaging engine in Started status

Configuring the Business Process Choreographer Container

By default, the Business Process Choreographer (BPC) container and the Human Task Manager are disabled. You need to configure them using the administrative console (or using the command line). The database user or schema is created already. You use it during the BPC configuration. To configure the BPC container:

- 1. Open the administrative console.
- 2. Go to the application servers list from the left navigation menu, and click the server name.
- 3. Under Container Settings on the right- hand side, open the Business Process Choreographer Container Settings, and click **Business Process Choreographer Containers**, as shown in Figure 23.

Application servers > server1 Use this page to configure an application serve applications.	er. An application server is a server that provides services required to run enter
Runtime Configuration	
General Properties	
Name server1	Session management
Node Name	■ SIP Container Settings
WPS61Node1	Web Container Settings
☐ Run in development mode	Business Process Choreographer Container Settings
🔽 Parallel start	Business Process Choreographer Containers

Figure 23 Open the Business Process Choreographer Container section

4. Edit the database configuration to use the user and schema that you defined (that is BPCDB), as shown in Figure 24.

Note: The "Create Tables" option, which similar to the messaging engines, is new in V6.1. This option triggers the creation of necessary tables and an index at the next server restart if those tables do not already exist.

Applicatio	<u>n servers</u> > <u>server1</u> > Business Proc	ess Choreographer Co	ontainers	8			
The Busine can combin	ss Process Choreographer providers su e business process technology with any	upport for business-proce other service offered by	ess applica products	ations. Business p supporting the op	processes can be a pen J2EE architect	utomatic, recoverable ure.	processes, or processes with human inter
To use the additional co	Business Process Choreographer funct onfiguration.	ionality, configure it with	this page.	The business flo	w and the human t	task containers will be	installed and basic configuration will be pe
Configurat	ion						
•	Data Source						
	2						
	Edit Test Connection						
	Database Instance	Schema Name	Create Tables	User Name	Password	Server	Provider
	WPSDB	BPCDB	ন	BPCDB	-	oracle10g	Oracle 10g
					-		

Figure 24 Business Process Choreographer container data source configuration

- 5. There are many configurations for the BPC container, which we do not detail in this paper, but you can do a basic configuration by clearing the e-mail configuration for the Human Task manager and entering the administration password for the administration ID.
- 6. Click OK.

The configuration starts and should complete successfully. Watch for messages such as the messages shown in Figure 25.

2008-03-12 17:53:51	CWSCA3014I: Resources for the SCA application "BPEContainer_WPS61Node1_server1" have been configured successfully.
2008-03-12 17:53:52	SEC J04001: Successfuly updated the application BPEContainer_WPS61Node1_server1 with the appContextIDForSecurity inform
1 2008-03-12 17:53:52	ADMA50111: The cleanup of the temp directory for application BPEContainer_WPS61Node1_server1 is complete.
2008-03-12 17:53:52	ADMA50131: Application BPEContainer_WPS61Node1_server1 installed successfully.
1 2008-03-12 17:53:52	Creating JMS resources for component 'Business Flow Manager'.
2008-03-12 17:53:56	Configuring component 'Human Task Manager' on deployment target 'WPS61Node1/server1'.
1 2008-03-12 17:54:01	ADMA50161: Installation of TaskContainer_WPS61Node1_server1 started.
2008-03-12 17:54:02	ADMA50051: The application TaskContainer_WPS61Node1_server1 is configured in the WebSphere Application Server reposite
1, 2008-03-12 17:54:02	CWSCA30131: Resources for the SCA application "TaskContainer_WPS61Node1_server1" are being configured.
2008-03-12 17:54:02	CWSCA3023I: The EAR file "task_WPS61Node1_server1.ear" is being loaded for the SCA module.
1, 2008-03-12 17:54:02	CWSCA3017I: Installation task "Resource Task for SCA Messaging Binding and EIS Binding" is running.
2008-03-12 17:54:02	CWSCA3017I: Installation task "Resource Task for SCA Messaging Binding and JMS Binding" is running.
1, 2008-03-12 17:54:02	CWSCA3017I: Installation task "SIBus Destination Resource Task for SCA Asynchronous Invocations" is running.
1, 2008-03-12 17:54:04	CWSCA3017I: Installation task "EJB NamespaceBinding Resource Task for SCAImportBinding" is running.
1, 2008-03-12 17:54:04	CWSCA3017I: Installation task "SIBus Destination Resource Task for SCA SOAP/JMS Invocations" is running.
1, 2008-03-12 17:54:04	CWSCA3014I: Resources for the SCA application "TaskContainer_WPS61Node1_server1" have been configured successfully.
1, 2008-03-12 17:54:04	SEC J04001: Successfuly updated the application TaskContainer_WPS61Node1_server1 with the appContextIDForSecurity infor
2008-03-12 17:54:05	ADMA50111: The cleanup of the temp directory for application TaskContainer_WPS61Node1_server1 is complete.
1, 2008-03-12 17:54:05	ADMA50131: Application TaskContainer_WPS61Node1_server1 installed successfully.
2008-03-12 17:54:05	Creating JMS resources for component 'Human Task Manager'.
2008-03-12 17:54:07	The configuration has ended.

Figure 25 BPC container configuration messages

7. When the configuration completes successfully, save the master configuration by clicking **Save the Changes**.

Updating the BPC bus messaging engine settings

When you configure the BPC container, a new bus is created, along with its messaging engine. You need to follow the steps described in "Configuring the messaging engines data store" on page 20 to change the schema name, using the BPCMSG user ID, as shown in Figure 26. You also need to edit the J2C alias called BPCME_00_Auth_alias, using the same user ID.

PCME_00_Auth_Alia:	5 entities and passwords for Java(TM) 2 connector security to use
Configuration	sincles and passwords for Java (199) 2 connector security to use.
General Properties	
* Alias	A.D.
BPCME_00_Auth_	Alias
* User ID	
BPCMSG	
* Password	
* Password	
* Password	

Figure 26 BPC bus messaging engine schema setup

Tables are created automatically at the next server restart. If you want to create the tables manually, follow the instructions in "Creating tables manually" on page 21, with the following **sibDDLGenerator** command:

/opt/ProcServer/bin/sibDDLGenerator.sh -system oracle -version 10g -platform windows -statementend \; -schema BPCMSG -user bpcmsg > /opt/shared/createSIBSchema_BPCMSG.ddl

Creating the BPC table spaces and tables

You now have two options:

1. You can restart the application server immediately. Because you selected the Create tables option when you configured the Business Process Choreographer container, all necessary tables are created under the BPCDB schema.

Important: All tables are created in the default user's table space, which is supported for *development* or proof-of-concepts environments only, *not test, integration or production environments*.

2. You can create the necessary table spaces and tables for the BPC container manually.

Important: If you want to run the Business Process Choreographer container in any mode other than development or POC, you must create tables and table spaces manually.

Creating tables automatically

Restart the server and make sure that the BPC container tables and views are created correctly. (There are 116 tables.) You should see a message in the server SystemOut.log file similar to:

[3/13/08 12:43:27:527 CET] 00000035 CreateSchemaM I CWWBB0625I: Database schema creation started. [3/13/08 12:43:50:731 CET] 00000035 CreateSchemaM I CWWBB0626I: Database schema creation completed successfully.

Creating table spaces and tables manually

Because the Oracle server in our example is remote, you need to copy all the scripts that are needed for the next steps to the Windows system, as follows:

 First, create specific table spaces for the BPC container.You must use a script called createTablespace.sql, which is located is the <<u>WPS_InstallRoot</u>>/dbscripts/ProcessChoreographer/Oracle directory. You need to copy this script to the Oracle server system and run it as the BPCDB user. The script requires that you provide a single parameter, which is the directory into which the table spaces for the user are created.

Tip: The table spaces creation requires about 1.3 GB of disk space.

Use this command:

(ORA) As BPCDB

sqlplus BPCDB/passwOrd@WPSDB @createTablespace.sql f:/Oracle/WPS61TS

This command creates the table spaces that are needed under the f:/Oracle/WPS61TS directory. (This directory must exist.) It creates the files shown in Example 1.

Example 1 Directory of f:\oracle\WPS61TS

Directory of f:\oracle\WPS61TS			
03/13/2008	14:14	104,865,792	AUDITLOG.DBF
03/13/2008	14:14	10,493,952	COMP.DBF
03/13/2008	14:14	262,152,192	INDEXTS.DBF
03/13/2008	14:14	524,296,192	INSTANCE.DBF
03/13/2008	14:14	209,723,392	LOBTS.DBF
03/13/2008	14:14	5,251,072	SCHEDTS.DBF
03/13/2008	14:14	10,493,952	STAFFQRY.DBF
03/13/2008	14:14	104,865,792	TEMPLATE.DBF
03/13/2008	14:14	52,436,992	WORKITEM.DBF
	9 Fi	le(s) 1,284,579	,328 bytes

Now, create the tables. During the BPC configuration, a script called createSchema.sql is generated in the

 WPS_InstallRoot>/profiles/WPSOrcl1/dbscripts/ProcessChoreographer/Oracle/BPCDB directory. Copy this file to the Oracle server, and execute it as the BPCDB user.: (ORA) As BPCDB

sqlplus BPCDB/passwOrd@WPSDB @createSchema.sql

You can now restart the application server and check that the BPC Container starts properly.

Configuring the Business Process Choreographer Observer

You can use Business Process Choreographer Observer to create reports on processes that are complete and to view the status of running processes. Its configuration is done in three steps, which we describe in this section.

Configuring Event Logging

First, make sure that events are being generated by the BPC Container and Human Task Manager by following these steps:

- 1. Log in to the administrative console.
- 2. Navigate to Servers \rightarrow Application Servers, and select the server1 application server.
- 3. Select Business Process Choreographer Containers, as shown in Figure 27.

<u>Application servers</u> > server1	
Use this page to configure an application to run enterprise applications.	server. An application server is a server that provides services required
Runtime Configuration	
General Properties	- Container Settings
server1	Session management
Node Name	 SIP Container Settings
WPS61Test1Node	
🗌 Run in development mode	Business Process Choreographer Container Settings
Parallel start	Business Process Choreographer Containers

Figure 27 Business Process Choreographer container settings

4. Expand the **State Observers** section, and make sure to select all the options, as shown in Figure 28.

tate Observers		
Logging	Business Flow Manager	Human Task Manager
Audit Logging		
Common Event Infrastructure Logging		

Figure 28 State Observers configuration settings

- 5. Click Apply.
- 6. When the configuration completes successfully, save the configuration.

Configuring the Event Observer

To configure the Event Observer:

- 1. Navigate to Servers \rightarrow Application Servers, and select server1.
- 2. Find the Business Integration section, expand the **Business Process Choreographer** section, and click **Business Process Choreographer Event Collector**.
- Complete the database configuration data, using the BPC Observer user that you created (BPCOBS), as shown in Figure 29. Make sure that the Create Tables option is selected if you want to create tables at server restart automatically. Select the server1 server as the Observation Target.

Edit Test Connection	on					
Database Instance	Schema Name	Create Tables	User Name	Password	Server	Provider
WPS	BPCOBS		BPCOBS	•••••	isabelle.mad.	Oracle 1
Observation Target		taiaar				
Managed business process	choreographer con	Lainer				
 Managed business process server=server1,node=WP 	Schoreographer con S61Test1Node	Lainer				

Figure 29 Business Process Choreographer Event Collector configuration

- 4. Click **OK**.
- 5. When the configuration completes successfully, click Save the Changes.

Installing the Observer application

To install the Observer application, you need to:

- 1. Navigate to Servers \rightarrow Application Servers, and select server1.
- 2. Find the Business Integration section, expand the **Business Process Choreographer** section, and click **Business Process Choreographer Observer**.
- 3. Click Add.
- 4. Define a context root for the application (by default it is /bpcobserver).
- 5. Click OK.
- 6. When the configuration completes successfully, click Save the Changes.

Creating the BPC Observer table spaces and tables

You now have two options:

 You can restart the application server immediately. Because you selected the Create Tables option when you configured the Business Process Choreographer observer, all necessary tables are created under the BPCOBS schema.

Important: All tables are created in the default user's table space, which is supported only for *development* or proof-of-concepts environments, not *test*, *integration*, or *production* environments.

2. You can create the necessary table spaces and tables for the BPC observer manually.

Important: If you want to run the Business Process Choreographer observer in any mode other than development or POC, you must follow the manual configuration instructions.

Creating tables automatically

Restart the server and make sure that the BPC observer tables and views are created correctly. (There are 116 tables.). You should see a message in the server SystemOut.log file similar to:

[4/17/08 19:05:14:234 CEST] 00000032 BPCObserverMo I CWWB04010I: The creation of the Business Process Choreographer Observer database schema started

Creating table spaces and tables manually

You can create table spaces, tables, and indexes manually using a script called createSchema_Observer.sql, which is generated under:

<WPS_InstallRoot>/profiles/WPSOrcl1/dbscripts/ProcessChoreographer/Oracle/BPCOBS

You must copy this script to the Oracle server system and run it as the BPCOBS user. It requires that you provide a single parameter, which is the directory into which the table spaces for the user should be created.

Tip: The table spaces creation requires about 600 MB of disk space.

Use this command:

(ORA) As BPCOBS

sqlplus BPCOBS/passw0rd@WPSDB @createSchema_Observer.sql f:/Oracle/WPS61TS

This command creates the table spaces that are needed in the f:/Oracle/WPS61TS directory. (This directory must exist.) The following files are created:

Directory of F:\oracle\WPS61TS 04/18/2008 10:48 262,152,192 OBSVRIDX.DBF 04/18/2008 10:48 209,723,392 OBSVRLOB.DBF 04/18/2008 10:48 104,865,792 OBSVRTS.DBF 3 File(s) 576,741,376 bytes

You can now restart the application server and check that the BPC Observer starts properly.

Using an Oracle cluster (RAC)

By default, all JDBC URLs that are used to connect to the Oracle server are of the following form:

jdbc:oracle:thin:@oracle10g:1521:WPSDB

If your Oracle database is deployed for a cluster of servers, you have to use a URL similar to this one instead:

```
jdbc:oracle:thin:@(DESCRIPTION=(ADDRESS_LIST=
(ADDRESS=(PROTOCOL=TCP)(HOST=oracle10g1)(PORT=1521))
(ADDRESS=(PROTOCOL=TCP)(HOST=oracle10g2)(PORT=1521))
(FAILOVER=on)(LOAD_BALANCE=on))(CONNECT_DATA=(SERVER=DEDICATED)
(SERVICE NAME=WPSDB)))
```

Tip: Use an editor which can match parenthesis to create this string and avoid errors.

Conclusion

In this paper, we explained how to configure all of the databases in our WebSphere Process Server environment to use Oracle in such a way that clarifies which commands need to use the SYS Oracle user and which the user can run for installing WebSphere Process Server.

Resources

For more information about the topics that we discuss in this paper, consult the following resources:

WebSphere Process Server Information Center

http://www-306.ibm.com/software/integration/wps/library/infocenter/

Building clustered topologies in WebSphere Process Server V6.1

http://www.ibm.com/developerworks/websphere/library/techarticles/0803_chilanti/ 0803_chilanti.html

WebSphere Process Server security overview

http://www-128.ibm.com/developerworks/websphere/library/techarticles/0602_khang aonkar/0602_khangaonkar.html

WebSphere Install Factory

http://www-1.ibm.com/support/docview.wss?rs=180&uid=swg24009108

Configuring and using XA distributed transactions in WebSphere Studio

http://www-128.ibm.com/developerworks/websphere/library/techarticles/0407_woolf /0407_woolf.html

WebSphere 6.1 Update Installer

http://www-1.ibm.com/support/docview.wss?rs=180&uid=swg24012718

WebSphere Process Server Fix Pack 1

http://www-1.ibm.com/support/docview.wss?rs=2307&context=SSQH9M&dc=D400&uid=swg 24018729&loc=en_US&cs=UTF-8&lang=en&rss=ct2307websphere

Acknowledgements

The authors want to thank:

- Chidanandan Theyancheri, Vijayavenkatesh Yelanji, Atul R Daherkar, Srikanth Bhattiprolu, and Ekkehard Voesch from the IBM SWG Development Labs for reviewing this paper.
- Frank Gilchrist, Genaro Nieto, Peter Philp, and Muriel Viale from IBM for beta-testing this paper.
- Aaron Miller for his support and help in the writing of this paper.

The team that wrote this paper

This paper was produced by a team of specialists from around the world working at the International Technical Support Organization (ITSO), Raleigh Center.



Isabelle Mauny works for the SWG development labs as a senior architect. Prior to this, she worked several years in the Software Services for WebSphere European team. She has extensive experience of J2EE[™] and BPM and co-wrote several IBM Redbooks® publications about WebSphere Application Server.



Simon Kapadia is the Technical Team Lead for IBM Software Services for WebSphere in the U.K. He works on designing and implementing large distributed computer systems for IBM customers. Simon holds a bachelor's degree in English and Drama, and a Master's in Computing Science. He has owned a computer since six years of age and has turned a life-long hobby into a career. Prior to joining IBM, Simon developed software for digital exchanges at Bell Laboratories, managed networks and Web applications at a Polish ISP, and supported and consulted on DCE, DFS[™], and Encina® for Transarc Corporation.

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

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

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

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 and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

This document REDP-4432-00 was created or updated on July 17, 2008.



Send us your comments in one of the following ways:

- Use the online Contact us review Redbooks form found at: ibm.com/redbooks
- Send your comments in an e-mail to: redbooks@us.ibm.com
- Mail your comments to: IBM Corporation, International Technical Support Organization Dept. HYTD Mail Station P099 2455 South Road Poughkeepsie, NY 12601-5400 U.S.A.



Redbooks (logo) 🔗 🛽

WebSphere®

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. These and other IBM trademarked terms are marked on their first occurrence in this information with the appropriate symbol (® or ™), indicating US registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at http://www.ibm.com/legal/copytrade.shtml

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

Cloudscape®	DFS™
DB2 Universal Database™	Encina®
DB2®	IBM®
developerWorks®	Redbooks®

The following terms are trademarks of other companies:

Oracle, JD Edwards, PeopleSoft, Siebel, and TopLink are registered trademarks of Oracle Corporation and/or its affiliates.

J2EE, Java, JDBC, JVM, and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

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

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.