



IBM Software Group

IBM WebSphere® Process Server V6.0.2

Database configuration guide



@business on demand.

© 2007 IBM Corporation
Updated April 10, 2007

This presentation covers database configuration for IBM WebSphere Process Server V6.0.2

Goals

- Understand database configuration during WebSphere Process Server installation
- Understand database usage for WebSphere Process Server components



The goals for this presentation are to understand WebSphere Process Server database configuration and use.

Agenda

- Database configuration overview
- Common database
- Common Event Infrastructure database
- Business Process Choreographer database
- Messaging engine database
- Enterprise Service Bus logger mediation database
- Selector and business rules database

This presentation will begin with an overview, then explain the individual databases used in a WebSphere Process Server environment.

Overview

- WebSphere Process Server V6.0.2 database configuration and its usage
- WebSphere Process Server uses different databases for different components
- SQL Scripts used to create tables

Component	Database (Name is configurable)
Business Process Choreographer	BPEDB
Common Event Infrastructure (CEI)	CEIDB
Relationships	CommonDB (WPRCSDB)
Mediation	CommonDB (WPRCSDB)
Recovery	CommonDB (WPRCSDB)
Application scheduler	CommonDB (WPRCSDB)
LockManager	CommonDB (WPRCSDB)
Selectors and business rules	CommonDB (WPRCSDB)/ RepositoryDB(Cloudscape)
Service Integration (SI) Bus	<User created during Message Engine (ME) configuration>
Enterprise Service Bus (ESB)	EsbLogMedDB

4

WebSphere Process Server uses five different database configurations for different components, as listed on this slide. Based on how your environment is structured, you may be able to combine some of these databases. Scripts to create the databases and associated tables are included with WebSphere Process Server.

The names of databases are configurable; you can change them as needed in your environment. Business Process Choreographer, Common Event Infrastructure, Service Integration Bus and Enterprise Service Bus components makes use of their separate databases by default. Some of these databases can be combined based on your environment. The common database, by default named WPRCSDB, is used for six different components, as shown in this table. Lock Manager is a new component that is added as part of release V6.0.2.

Section

CommonDB (WPRCSDB)



This section covers the database used by several components within WebSphere Process Server.

CommonDB

- The CommonDB database is used by the following components in WebSphere Process Server V6.0.2
 - ▶ Recovery
 - ▶ Relationship
 - ▶ Mediation
 - ▶ AppScheduler
 - ▶ LockManager
 - ▶ Customization (selector and business rule)

As shown previously, the CommonDB database is used by recovery, relationship, mediation, AppScheduler, customization and the new lock manager components.

CommonDB - Supported database types

- MSSQL Server Embedded and Data direct now supported 
- For a list of all the databases supported, refer to information center at https://publib.boulder.ibm.com/infocenter/dmndhelp/v6rxmx/index.jsp?to pic=/com.ibm.wsp.602.ins.doc/doc/cins_db_specs.html
- Cloudscape™ is default for stand-alone profile
- For the Network Deployment environment:
 - ▶ DB2® Universal is the default database type
 - ▶ Cloudscape is not enabled
- Selectors and Business Rules did not support Informix® or MSSQL Server (Embedded and Data Direct) for WebSphere Process Server V6.0 and V6.0.1 but are now supported in V6.0.2 



CommonDB supports most of the major databases, including Cloudscape, DB2, Informix, MSSQL, and Oracle. Cloudscape is the default database type for stand-alone profile. If your environment is Network Deployment, then DB2 is the default selection and Cloudscape is not supported.

Selectors and business rules did not support Informix or MSSQL Server for WebSphere Process Server V6.0 and V6.0.1. They are now supported in release 6.0.2 for the Selectors and Business Rules components.

CommonDB (cont.)

- Database instances
 - ▶ One instance per cell
- Configuration actions during profile creation
 - ▶ The commonDB database can be created as part of the profile creation using the profile action scripts.
 - ▶ Delay commonDB database creation to after profile creation

There is one database instance per cell.

The CommonDB database can be created as part of the profile creation process. The functionality to delay database creation has been added for CommonDB in version 6.0.2. If you select this option in the profile creation wizard, scripts will be generated and exported. These scripts can be run after profile creation to create the databases.

You will see where these scripts are exported later in this presentation. Configuration for the business process choreographer is done after profile creation, so there is no option to export the scripts.

The common event infrastructure database configuration can also be done after the profile creation.

If you decide to export the scripts instead of running them, incorrect values for the database configuration will not cause the profile actions to fail, and the errors will be visible when the exported scripts are run.

CommonDB

- After WebSphere Process Server installation, profile action scripts are in:

- ▶ $\{\text{WPS_HOME}\}/\text{profileTemplates}/\{\text{profileType}\}/\text{actions}/\text{scripts}$ for V6.0.1
- ▶ $\{\text{WPS_HOME}\}/\text{dbscripts}/\text{CommonDB}/\{\text{dbType}\}$ for V6.0.2
- ▶ $\{\text{WPS_HOME}\}/\text{profiles}/\{\text{profileName}\}/\text{dbscripts}/\text{CommonDB}/\{\text{dbType}\}/\{\text{dbName}\}$ for V6.0.2 delayed database creation



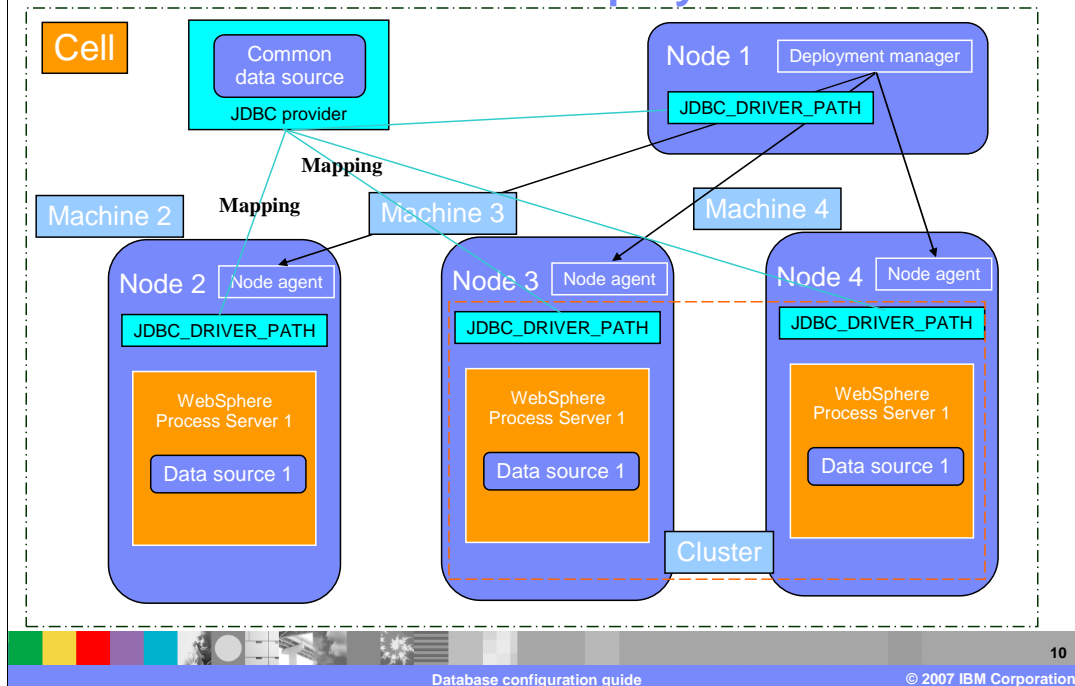
- The scripts:

- ▶ Create new database (valid only for a local database) if asked, depending on user's choice
- ▶ Create JDBC provider and variable "JDBC_DRIVER_PATH"
 - Node level for stand-alone and managed profiles
 - Cell level for deployment manager
- ▶ Create data source upon the JDBC provider.
- ▶ Data source JNDI name: jdbc/WPSDB
- ▶ Each component will run CommonDB ant script "createTable" target to create tables

This slide describes the actions that are performed as part of the profile creation. Scripts are located in the WebSphere Process Server installation folder, as listed here.

For V6.0.2, the JDBC provider name changes to the one defined by the WebSphere Application Server base template located in the jdbc-resource-provider-templates.xml file. If a JDBC provider matching the template already exists at the scope required, it will be used to create the data sources instead of creating a new JDBC provider. The JDBC provider is created at the node level while configuring a stand-alone profile and at the cell level in Network Deployment. The JDBC_DRIVER_PATH variable is created at node level for stand-alone or managed profiles, and at the cell level for Network Deployment environments.

CommonDB in Network Deployment



This is a graphical representation of CommonDB configuration in Network Deployment. You can see that all the nodes make use of the single JDBC provider to access the CommonDB database. The `JDBC_DRIVER_PATH` variable refers to the location of the JDBC drivers on each machine that a node resides on.

Common event infrastructure database (CEIDB)

This section covers the common event infrastructure database.

Common event infrastructure (CEIDB)

- Only created for a Stand-alone profile.
- CEI database is internal and all interactions with it must be performed using the support Common Event Infrastructure (CEI) APIs.
- DB2 Universal OS/390® V7.1 and DB2 Universal OS/390 V8.1 configurable from profile creation wizard
- For list of all supported databases, see

http://publib.boulder.ibm.com/infocenter/dmndhelp/v6rxmx/topic/com.ibm.wsp.602.ins.doc/doc/cins_cei_db_specs.html



The common event infrastructure service is used to process events in WebSphere applications and processes. The common event infrastructure database is created for a stand-alone profile. All interactions should be performed using the supported APIs.

For WebSphere Process Server versions before 6.0.2, DB2 Universal databases for OS/390 were not configurable using the profile creation wizard. With V6.0.2 you can configure both version 7.1 and version 8.1 from the wizard. For list of all the supported databases, follow the information center link provided here.

CEIDB (cont.)



The common event infrastructure database instance is available for a stand-alone profile, therefore a single instance is created. Configuration of the CEI database is done by the profile template ANT script called `configCei.ant`. The ANT script invokes the CEI API with a response file containing all of the necessary properties for the required configuration. As CEI processes this response file, it creates the database or scripts and they are placed under specific database-named folders in the `profilePath`.

Refer to the link provided to get more information on how to configure the CEI database.

Business process choreographer database (BPEDB)

This section covers the business process choreographer database.

Business process choreographer database (BPEDB)

- BPEDB is only used by the business process choreographer component
- BPEDB - Supported database types
 - ▶ Cloudscape is used for the BPC sample configuration.
 - ▶ The network JDBC driver for Cloudscape is not supported, because it has no XA support.
 - ▶ For list of all supported databases, refer to http://publib.boulder.ibm.com/infocenter/dmndhelp/v6rxmx/topic/com.ibm.wsp.602.ins.doc/doc/cins_bpc_db_specs.html

The business process choreographer component provides support for business process applications within IBM WebSphere Application Server. Business processes can be automatic, recoverable processes, or processes that require human interaction. With the process choreographer, you can combine business process technology with any other service offered by products supporting the open J2EE architecture. Business process choreographer supports DB2, Informix, Cloudscape, MS SQL server and Oracle. For a list of all the supported database types and versions, refer to the information center link provided in the slide.

BPEDB - Database instances

- Each deployment target (server or cluster) can have a business process choreographer configuration.
- Each business process choreographer configuration has its own database.
- Databases cannot be shared between business process choreographer configurations.

Each deployment target can have a business process choreographer configuration. Each business process choreographer configuration has its own database. Databases cannot be shared between business process choreographer configurations.

BPEDB - Configuration

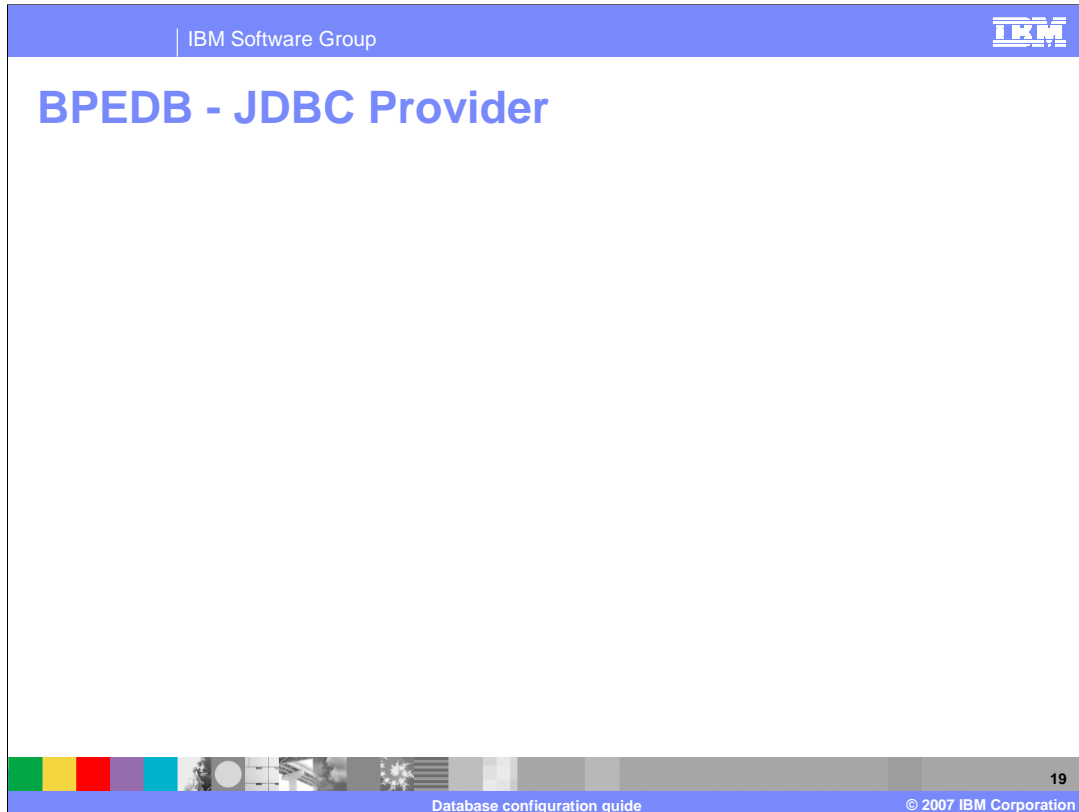
- After profile creation
 - ▶ due to complexity
- “BPC Sample Configuration”
 - ▶ Uses Cloudscape
 - ▶ only for stand-alone configuration
 - Cloudscape not supported in Network Deployment environment
 - ▶ Profile creation and augmentation tool invokes `bpeconfig.jacl`

Configuring business process choreographer can be complex and requires several parameters. Therefore, it is a post-profile-creation step. The "BPC Sample Configuration" is based on Cloudscape and requires a minimal set of parameters. It is not suited for production. The "BPC Sample Configuration" can be selected only when creating a stand-alone profile, because Cloudscape is not supported by business process choreographer in Network Deployment environments. When the "BPC Sample Configuration" option is selected, the profile creation and augmentation tool will invoke the standard business process choreographer configuration script listed here with the appropriate parameters that you have supplied.

BPEDB - SQL scripts



SQL scripts are in the *ProcessChoreographer* directory in the WebSphere Process Server home folder. These scripts are described in detail in the information center. They can be run either manually or be called by `bpeconfig.jacl`, which is what happens during profile creation when the Cloudscape database for the "BPC Sample Configuration" is created.



A JDBC provider is created if necessary, depending on the database type. It is created with the 'server' scope if business process choreographer is configured on a server, and with a 'cell' scope if it is configured on a cluster.

If a suitable JDBC provider already exists at the required scope, it will be re-used. A JDBC provider is suitable if it has the 'name' attribute set to the value that is used by the corresponding WebSphere Application Server base JDBC provider template, and it is for the required database type.

The appropriate variables will be set on any node where there is a server or cluster member where business process choreographer is configured.

Messaging engine database (MEDB)

This section covers the database as used by the message engines for service component architecture. The default database name is SCADB. For the embedded Cloudscape database, each deployment target will have its own database instance.

MEDB - Supported database types

- Messaging engine database (MEDB)
 - ▶ used to store operating information
 - ▶ Store information required for recovery of failed event or failure
- All databases supported by WebSphere Application Server are supported for the messaging database
- For a list of supported database types, refer to <http://www-1.ibm.com/support/docview.wss?rs=180&uid=swg27007251>

The messaging engine database is used to store operating information. Also stored are essential objects that the messaging engine needs for recovery in the event of a failure.

The default database name for the SCA messaging engine is SCADB, for the other messaging engines it is MEDB.

For the embedded Cloudscape database, each messaging engine will have its own database or schema. The default schema name is IBMWSSIB

MEDB - Database instances



You can create a database for each messaging engine or you can use a single database for all the messaging engines. Similarly you can use many database schemas, or a single one. Each messaging engine must have either its own database or its own schema.

One instance of the database is used for the system bus, and other instance of the database is used for the application bus. Another is used to host each messaging engine for the common event infrastructure bus. And yet another is used to host each messaging engine for the business process choreographer bus.

MEDB - Configuration during profile creation



The message engine database is created during installation for a stand-alone profile. In a Network Deployment environment, you can configure the message engines using the wizard available from the administrative console. The message engine database needs to be manually created before you use the wizard to configure message engines.

ESB logger mediation database (ESBDB)

This section covers the ESB logger mediation database.

ESB logger mediation database (ESBDB)

- The EsbLogMedDB is used by the MessageLogger mediation primitive in WebSphere Enterprise Service Bus and in WebSphere Process Server
- Only created for a stand-alone profile
- For a list of all the supported database types, see:
https://publib.boulder.ibm.com/infocenter/dmndhelp/v6rxmx/topic/com.ibm.wsp.602.ins.doc/doc/cins_esb_logger_db_specs.html

The enterprise service bus logger mediation database is used by the message logger mediation primitive in WebSphere Enterprise Service Bus and WebSphere Process Server.

The database created for a stand-alone profile is Cloudscape.

There are .ddl files for other database types to use if you need additional databases either in a stand-alone profile or for a Network Deployment, or clustered, environment.

ESBDB - Actions during profile creation

- For Stand-alone profile
 - ▶ Database is created in <WAS_HOME>/profile/<profile>/databases directory if one does not already exist.
 - ▶ At server scope, script will find or create a suitable database provider and create a data source
- SQL scripts
 - ▶ createMessageLoggerResource.jacl is in <WAS_HOME>/bin and is used during the stand-alone profile creation.
 - ▶ DDL files for all database types are located in the <WAS_HOME>/util/EsbLoggerMediation/<DATABASE_TYPE> directories

For a stand-alone profile, databases are created in the profile's database directory. This action is performed for all WebSphere Process Server and WebSphere Enterprise Service Bus stand-alone profile creations – there is no facility for you to override this data source creation.

Database definition language, or .ddl, files for all database types are located in the directory named for the database in the WebSphere Process Server home directory. If a suitable JDBC provider is configured at the server scope, it is used.

The name of the data source is 'ESB Logger Mediation Data Source'. You can create your own data source and configure a Message Logger mediation to use it.

Section

Selectors and business rules database

This section covers the selectors and business rules database

Selector and business rules database (SBDB)

- Holds selector and business rule artifacts
- Changed through
 - ▶ Administrative console
 - ▶ Business rules manager
- Artifacts in original EAR not updated
- Supports DB2 V7.0, V8.0 for z/OS®, Oracle thin, Oracle OCI, Informix and MSSQL Server
 - ▶ Using Cloudscape requires separate RepositoryDB
 - ▶ For list of all supported databases, see:
https://publib.boulder.ibm.com/infocenter/dmndhelp/v6rxmx/topic/com.ibm.wsps.602.ins.doc/doc/cins_sel_bus_rule_db_specs.html



The selector and business rules components use a database to hold the selector and business rule artifacts that are created in WebSphere Integration Developer and installed on the server. If any changes occur to a selector through the administrative console or to business rules through the business rule manager, the database is updated with the latest information. The original artifacts in the EAR are not synchronized with any updates made once the application is installed.

In WebSphere Process Server V6.0, only Cloudscape and DB2 on the distributed platforms were supported for holding selector and business rule components. In WebSphere Process Server V6.0.1, support was added for DB2 V7.0 and V8.0 for z/OS, Oracle Thin and Oracle OCI databases.

For stand-alone or distributed installations where Cloudscape is selected, the separate RepositoryDB is used.

Version 6.0.2 adds selectors and business rules component support for Informix and MSSQL Server, both Embedded and Data Direct.

SBDB

- Single database instance is available in a stand-alone profile or Network Deployment configuration
- Using separate database instances for different selectors or business rules is not supported
- SQL scripts for table creation at:
 - ▶ Same location as commonDB scripts
- Also in this directory:
 - ▶ configDynamicArtifactRepository.jacl file.
 - This file is used during the stand-alone profile creation when Cloudscape is selected. It configures the appropriate Cloudscape JDBC provider and data source.

Only one database instance is available in a stand-alone profile or Network Deployment configuration. All selectors and business rules for the server in a stand-alone profile or in a cell use the same repository. There is no support for using separate database instances for different selectors or business rules. The repository runtime code uses hard-coded lookups to find the database. In any type of server configuration, the same code is run and therefore the same database is used. Scripts for table creation are available under the scripts folder of WebSphere Process Server home directory. The .jacl script shown here is used during stand-alone profile creation.

Summary

- WebSphere Process Server components use several databases
- Database creation is scripted
- More information at:
https://publib.boulder.ibm.com/infocenter/dmndhelp/v6rxmx/topic/com.ibm.wsps.602.ins.doc/doc/cins_db_specs.html

In summary, the components in WebSphere Process Server use several databases. You can elect to have some of these databases created as part of the installation, or generate the scripts and manually create the databases. This presentation has detailed the different databases and some of the configuration options; for more information, visit the information center

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:

Cloudscape DB2 IBM Informix OS/390 WebSphere z/OS

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 NON-INFRINGEMENT. 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.