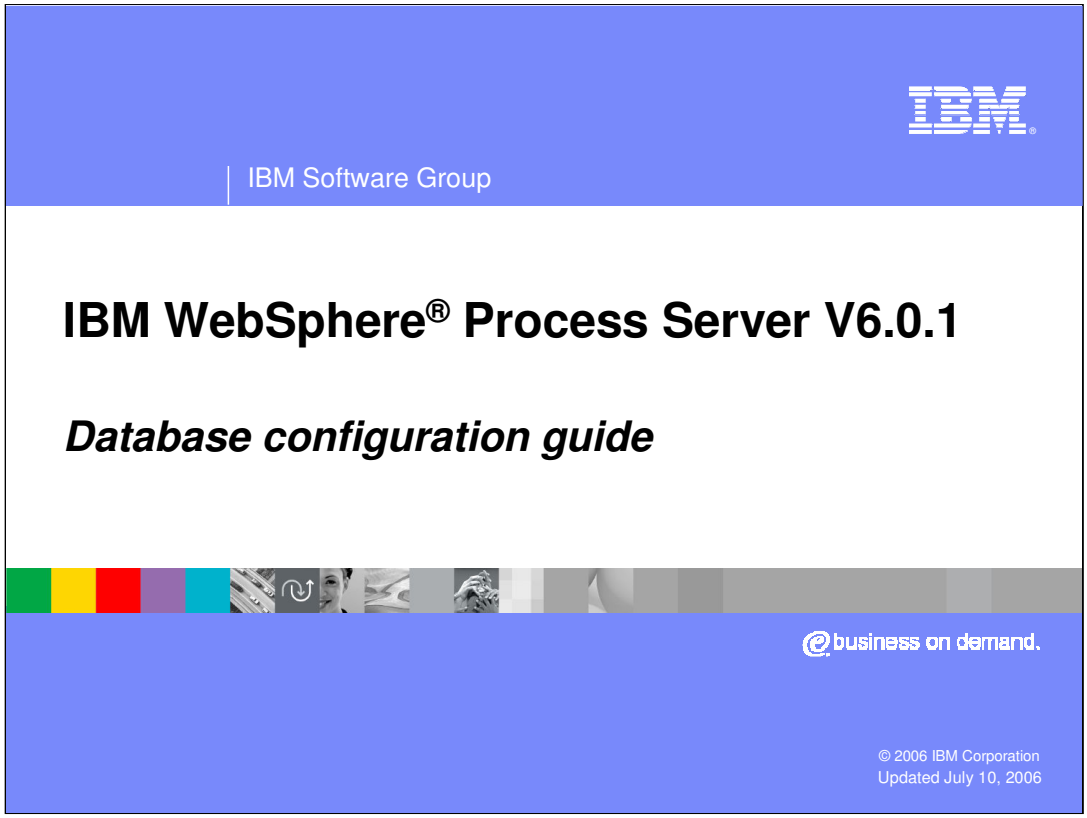

This presentation covers Database Configuration for IBM WebSphere Process Server V6.0.1

Goals

- Understand database configuration during WebSphere Process Server installation
- Understand database usage for WebSphere Process Server components
- Document supported database types
- Restrictions on database types

The goals for this presentation are to understand WebSphere Process Server database configuration, usage, support, and restrictions.

Agenda

- Database configuration overview
- Common Database
- Common Event Infrastructure Database
- Business Process Choreographer Database
- Messaging Engine Database
- ESB Logger Mediation Database
- Selector / 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 V6.0.1 Database configuration and its usage
- WebSphere Process Server uses different databases for different components
- SQL Scripts used to create tables

WebSphere Process Server uses five different database configurations for different components, as listed on the next 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.

Databases used by WebSphere Process Server components

Component	Database (Name is configurable)
Business Process Choreographer	BPEDB
Common Event Infrastructure (CEI)	CEIDB
Relationships	CommonDB (WPRCSDB)
Mediation	CommonDB (WPRCSDB)
Recovery	CommonDB (WPRCSDB)
App Scheduler	CommonDB (WPRCSDB)
Selectors/Business Rules	CommonDB (WPRCSDB)/ RepositoryDB(Cloudscape)
Service Integration (SI) Bus	<User created during ME configuration>
Enterprise Service Bus (ESB)	EsbLogMedDB

Names of Databases are configurable; you can change them as needed in your environment. The common database, by default named WPRCSDB, is used for five different components, as shown in this table.

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 V6.0.1
  - Recovery
  - Relationship
  - Mediation
  - AppScheduler
  - Customization (Selector/Business Rule)

As shown previously, the CommonDB database is used by these components.

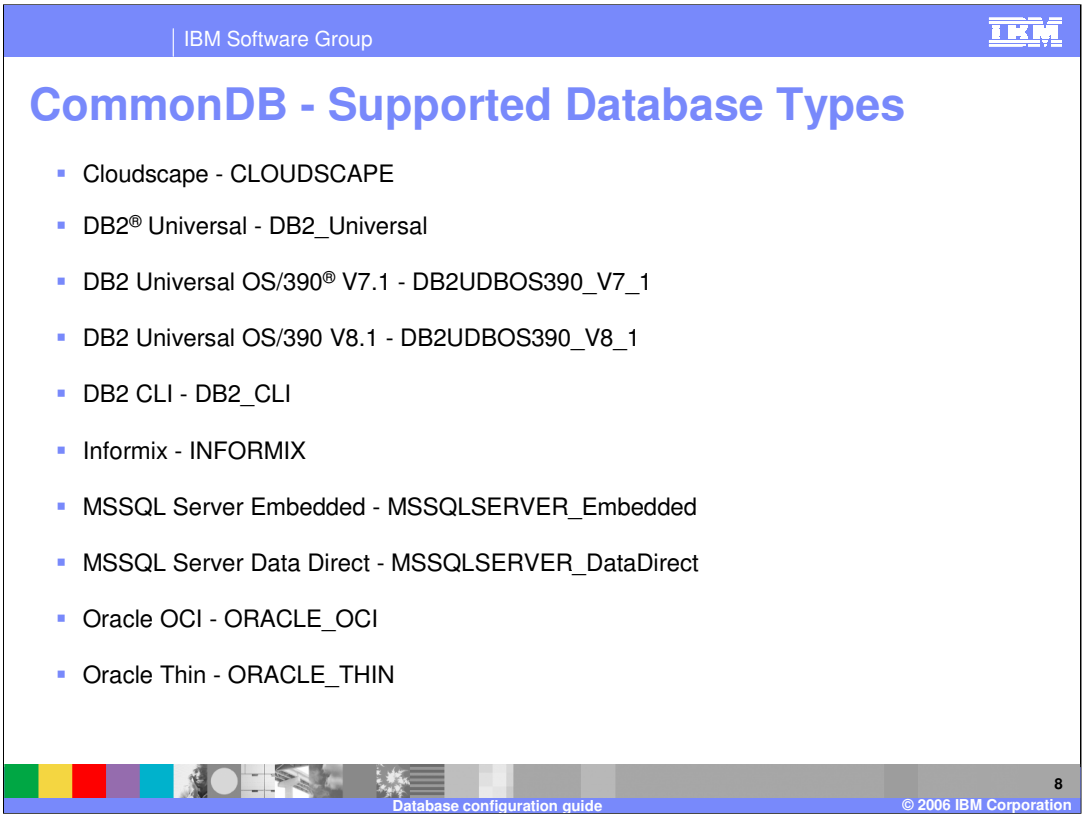

CommonDB supports most of the major databases, including Cloudscape, DB2, Informix, MSSQL, and Oracle.

CommonDB (cont.)

- Cloudscape is default for stand-alone profile
- For the Network Deployment environment:
  - DB2 Universal is the default database type
  - Cloudscape is not enabled
- Selectors/Business Rules do not use the Common database when Cloudscape is selected
- Selectors/Business Rules do not support Informix or MSSQL Server (Embedded and Data Direct) for WebSphere Process Server V6.0 and V6.0.1.x.

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. Note that the Selectors and Business Rules components are not supported on Cloudscape, Informix, or MSSQL.

CommonDB (cont.)

- DB Instances
  - One DB instance per cell
- Configuration actions during profile creation
  - The Common DB database is created as part of the profile creation using the profile action scripts.

There is one DB instance per cell. The CommonDB database is created as part of the profile creation. The next three slides describe the actions that are performed as part of the profile creation.

CommonDB - Stand-alone profile

- After WebSphere Process Server installation, profile action scripts are in:
  - ${WAS_HOME}/profileTemplates/**default**.wbiserver/actions
  - ${WAS_HOME}/profileTemplates/**default**.wbiserver/actions/scripts
- The scripts:
  - Create new database if asked, depending on selection in DB Panel
  - Create JDBC™ provider at node level
    - Variable "JDBC_DRIVER_PATH" is created at node level
  - Create data source upon the JDBC provider
  - Each component will run the 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 WebSphere Process Server installation folder, as listed here. The JDBC provider is created at the node level while configuring a stand-alone profile.

CommonDB - Deployment manager profile

- After WebSphere Process Server installation, profile action scripts are in:
  - ${WAS_HOME}/profileTemplates/**dmgr**.wbiserver/actions
  - ${WAS_HOME}/profileTemplates/**dmgr**.wbiserver/actions/scripts
- The scripts:
  - Create new database if asked
  - Create JDBC provider at cell level
    - Variable "JDBC_DRIVER_PATH" is created at cell level
  - Create data source upon the JDBC provider (cell level).
  - Each component will run the CommonDB ant script "createTable" target to create tables.

The JDBC provider and datasource are created at the cell level while configuring a Deployment manager profile.

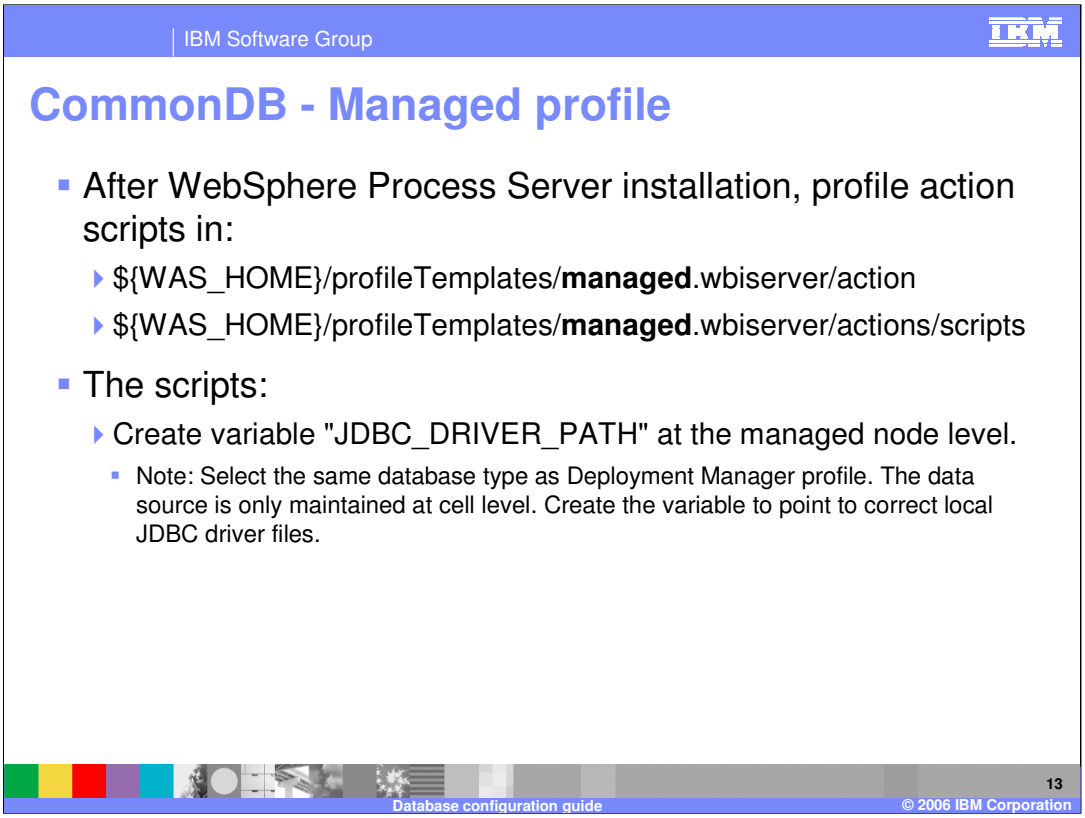

While configuring a managed Profile, *you must select **the same database type** as for the Deployment manager profile.* The data source is maintained at cell level.

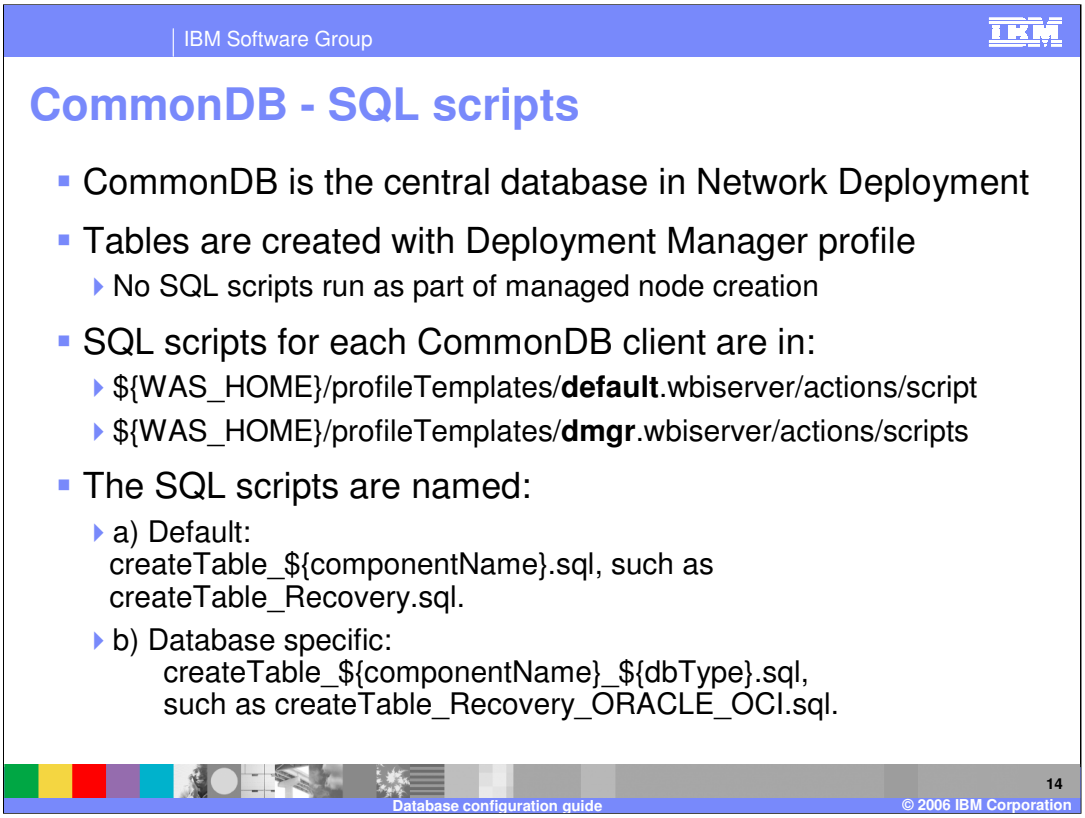

CommonDB is the central database in Network Deployment. Tables are created with Deployment manager profile and so there will not be any SQL scripts run as part of managed node creation. SQL scripts are named to refer to the component and database they are targeted for. If a database-specific script is not found, the default SQL script will be used. Therefore, each component only needs to provide database-specific SQL when it is different from the default one.

CommonDB - JDBC Provider

- New JDBC provider is created depending on database type
  - Stand-alone - server scope
  - Network Deployment - cell scope
  - JDBC Provider refers to variable "JDBC_DRIVER_PATH" to locate local JDBC drivers
    - The variable is specified in cell level and each node level, to point to correct local path.
- JDBC provider: WPSDefaultDatasource_${dbType}
- The data source name is independent of database type
- Data source JNDI name: jdbc/WPSDB

JDBC provider is configured at the server scope for stand-alone profile and Cell level in Network Deployment environment respectively. JDBC Provider refers to variable "JDBC_DRIVER_PATH" to locate local JDBC drivers

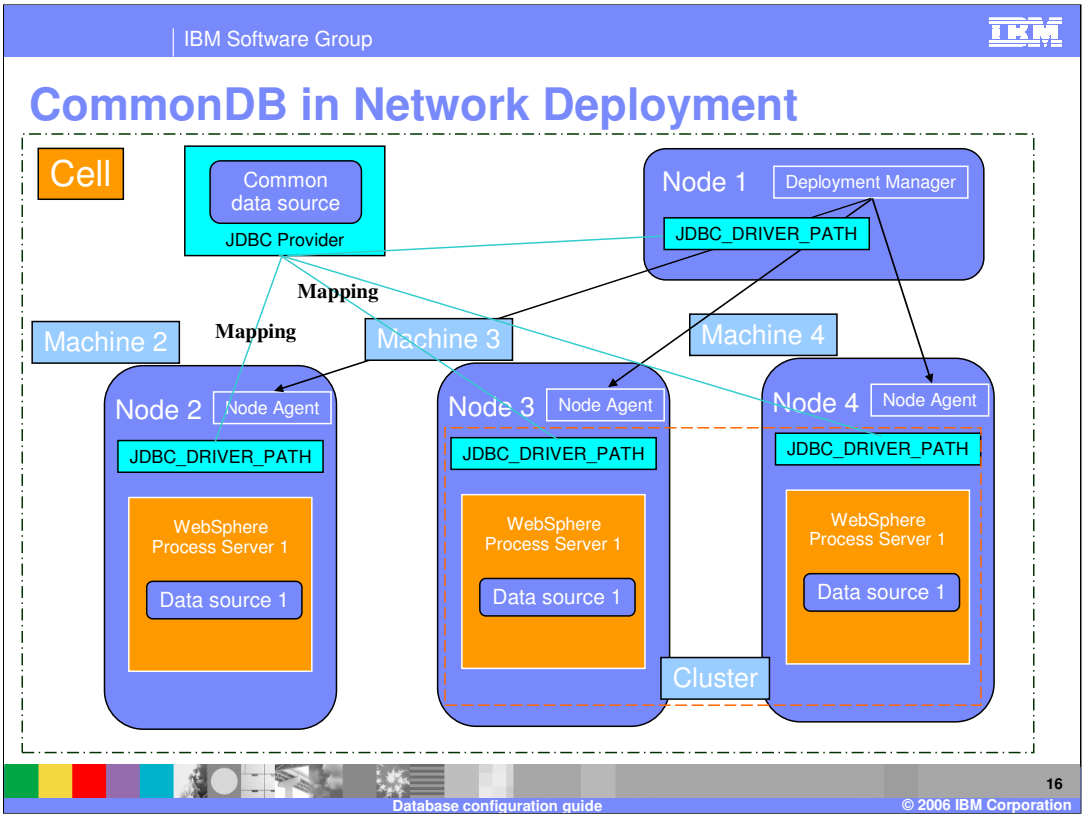

This is a graphical representation of Common database configuration in Network Deployment

CommonDB - Restrictions

- "Create new database" is not available for:
  - DB2 Universal OS/390 V7.1 - DB2UDBOS390_V7_1
  - DB2 Universal OS/390 V8.1 - DB2UDBOS390_V8_1
  - Oracle OCI - ORACLE_OCI
  - Oracle Thin - ORACLE_THIN

Restrictions for CommonDB. The “Create new database” option is not available for DB2 universal OS/390 V7.1 , OS/390 V8.1 , Oracle OCI and Oracle Thin.

CommonDB - Tables

- CommonDB scripts only create static tables during profile creation
- Tables that are created by different components:

Components	Table Names
Recovery	FAILEDEVENTS FAILEDEVENTBOTYPES FAILEDEVENTMESSAGE
Mediation	MEDIATION_TICKETS
Relationship	Dynamic table, be created at runtime.
AppScheduler	WSCH_LMGR WSCH_LMPR WSCH_TASK WSCH_TREG
Customization (Selection/BR)	BYTESTORE BYTESTOREOVERFLOW APPTIMESTAMP

CommonDB script creates many different static tables for different components as listed here.

CommonDB - Exported scripts

- From WebSphere Process Server 6.0.1.2, database scripts are exported to:
  - *${profilePath}/databases/${dbName}*
- The scripts are created for both the options "Create new database" or "Reuse existing database".
- The scripts only contain basic creating database/table/index statements
  - You should leverage database native commands to run the scripts.
  - This is subject to change.

Database scripts for both "Create new database" and "Reuse existing database“ can be found in specific database named folders inside WebSphere Process Server profile path. You should leverage database native commands to run them.

CommonDB – User ID privileges

- The specified user ID should be able to create table spaces, tables, indexes and stored procedures
- For "Create new database" option, the specified user ID should have permissions to create a new database

If you choose to create a new database then your user ID should have sufficient permissions to create a database.

Section

Common Event Infrastructure (CEIDB)

This section covers the Common Event Infrastructure database.

Common Event Infrastructure (CEIDB)

- This database is only created for a Stand-alone profile. The CEI database is internal and all interactions with it must be performed using the support CEI APIs.
- For V6.0.1 see:
  - http://publib.boulder.ibm.com/infocenter/dmndhelp/v6rxmx/topic/com.ibm.wsps.ins.doc/doc/tcei_install_configureDb.html

Common Event Infrastructure database is internal and all interactions should be performed using supported APIs.

CEIDB - Supported Database Types

- WebSphere Process Server V6.0.1
  - Cloudscape V5.1
  - DB2 Universal V8.1
  - DB2 Universal V8.2
  - Oracle V9.1
  - Oracle V10.1
  - DB2 Universal OS/390 V7.1 - not directly supported by profile wizard
  - DB2 Universal OS/390 v8.1 - not directly supported by profile wizard

Common Event Infrastructure supports databases as listed here.

CEIDB (cont.)

- CEIDB - DB Instances
  - This is only for a Stand-alone profile, so there is only one instance.
- Configuration actions during V6.0.1profile creation
  - Configuration of the CEI database is done by the profile template ANT script configCei.ant invoking the CEI API with a response file containing all the necessary properties for the required configuration.

The Common Event Infrastructure database instance is available for a stand-alone profile, therefore a single instance is created. This database is configured by an ANT script called configCei.ant.

CEIDB - SQL scripts

- In V6.0.1, Database configuration is accomplished by creating a response file invoking the CEI API with it as a parameter. As CEI processes this response file and creates the DB (or just the scripts) it will place the scripts in:
  - *profilePath*/default/event/dbscripts/*dbtype*

In Version 6.0.1, Database configuration is accomplished by creating a response file and invoking the CEI API with the name of the response file as a parameter. As CEI processes this response file, it creates the Database or Scripts and they are placed under specific database named folder inside profilePath.

CEIDB - JDBC Provider

- V6.0.1 for DB2ZOS, supported JDBC drivers:
  - On native z/OS®:
    - DB2 Universal JDBC Driver Provider (XA)
  - On distributed systems:
    - DB2 Legacy CLI-based Type 2 JDBC Driver (XA)
    - DB2 Universal JDBC Driver Provider (XA)
    - JDBC_PROVIDER="DB2 Universal JDBC Driver Provider (XA)"

JDBC providers for DB2ZOS are as shown in the list here.

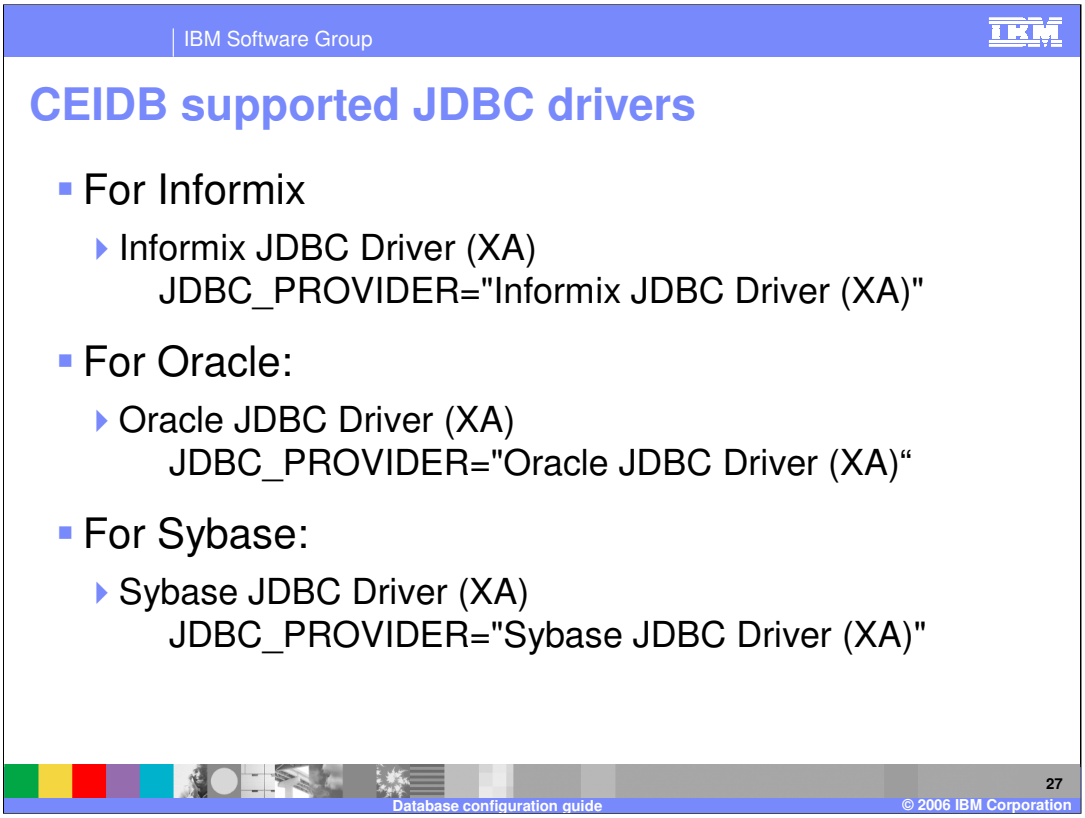

JDBC providers for Informix, Oracle and Sybase are as shown here with the values for the JDBC_PROVIDER variable that are used for each specific database.

CEIDB supported JDBC drivers (cont.)

- For Cloudscape:
  - Cloudscape JDBC Provider (XA)
    JDBC_PROVIDER="Cloudscape JDBC Provider (XA)"
- For DB2:
  - DB2 Universal JDBC Driver Provider (XA)
  - DB2 Legacy CLI-based Type 2 JDBC Driver (XA)
    JDBC_PROVIDER="DB2 Universal JDBC Driver Provider (XA)"
  - jndiName="jdbc/cei“
  - jndiName="jdbc/eventcatalog"

JDBC provider for Cloudscape and DB2 are shown here.

CEIDB supported JDBC drivers (cont.)

- For SQL Server:
  - DataDirect ConnectJDBC type 4 driver for MS SQL Server (XA)
  - WebSphere embedded ConnectJDBC driver for MS SQL Server (XA)
    JDBC_PROVIDER="WebSphere embedded ConnectJDBC driver for MS SQL Server (XA)"

This slide documents the JDBC provider information for MS SQL server.

CEIDB (cont.)

- Restrictions
  - In V6.0.1, DB2 Universal OS/390 V7.1 and V8.1 are not directly supported by profile wizard
- Tables
  - There are many tables for V6.0.x.
  - Look under the generated scripts to see what tables are generated for the given database product

DB2 Universal OS/390 V7.1 and V8.1 are not directly supported by profile wizard. Tables generated can be found In the scripts for specific database type

CEIDB (cont.)

- Exported scripts
  - In profilePath/event/dbscripts/dbtype, there are shell scripts which are to be used to run the SQL scripts generated
- User ID permissions
  - Generally the specified user ID should be able to create tablespaces, tables, indexes and stored procedures.
  - For "Create new database" option, the specified user ID should be able to create a new database.

If you choose to create a new database then your user ID needs to have permissions to create a database, and to create tablespaces, tables, indexes and stored procedures.

Section

Business Process Choreographer (BPEDB)

This section covers the Business Process Choreographer database.

Business Process Choreographer (BPEDB)

- BPEDB is only used by Business Process Choreographer
- BPEDB - Supported Database Types
  - Cloudscape V5.1
    - Cloudscape JDBC Provider (XA)
  - DB2 UDB v8.1 and v8.2
    - DB2 Legacy CLI-based Type 2 JDBC Driver (XA)
    - DB2 Universal JDBC Driver Provider (XA)
  - DB2 z/OS V7 and v8
    - DB2 Legacy CLI-based Type 2 JDBC Driver (XA)
    - DB2 Universal JDBC Driver Provider (XA))
  - Informix v9.4
    - Informix JDBC Driver (XA)

Business Process Choreographer supports DB2, Informix, Cloudscape, MS SQL server and Oracle.

BPEDB (cont.)

- SQL Server 2000
  - WebSphere embedded ConnectJDBC driver for MS SQL Server (XA)
  - DataDirect ConnectJDBC type 4 driver for MS SQL Server (XA)
- Oracle 9i and 10g
  - Oracle JDBC Driver (XA)
- Notes:
  - Cloudscape is used for the BPC Sample Configuration.
  - The network JDBC driver for Cloudscape is not supported, because it has no XA support.

The network JDBC driver for Cloudscape is not supported, because it has no XA support.

BPEDB - DB 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 (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.

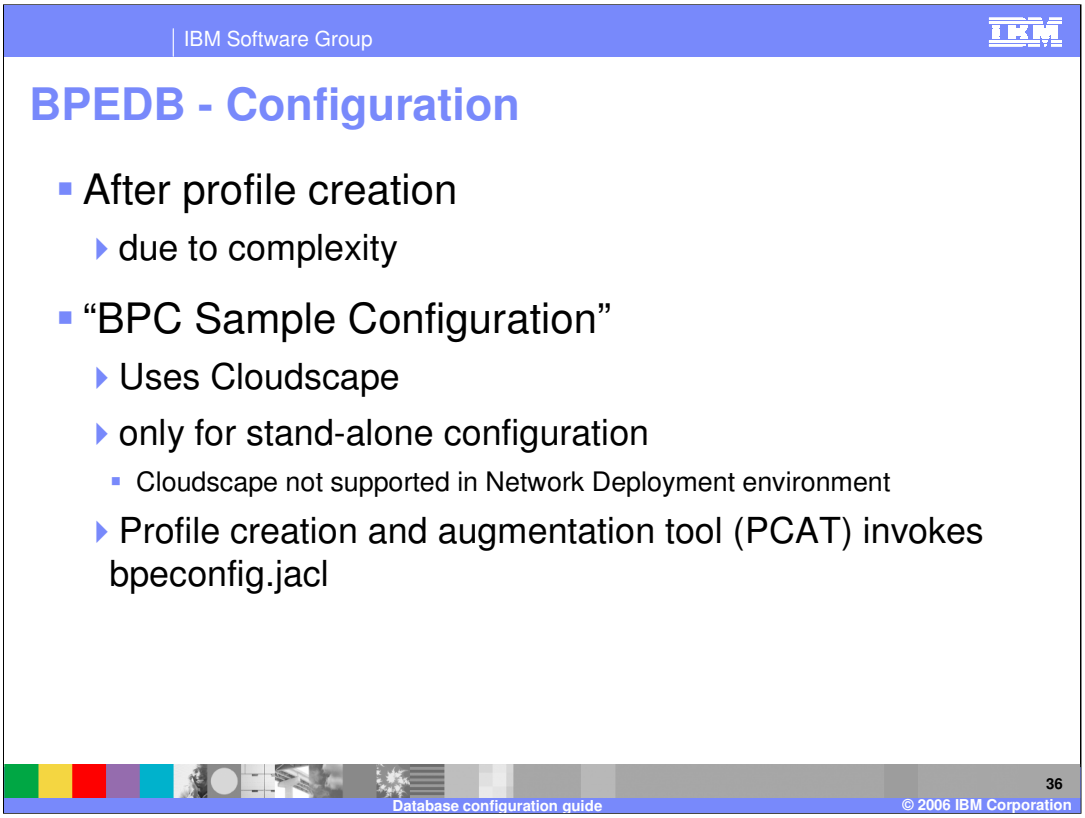

Configuring Business Process Choreographer can be complex and requires a lot of parameters. Therefore, it is a post profile creation step. There is a so called "BPC Sample Configuration" which is based on Cloudscape and requires a minimal set of parameters. It is not suited for production purposes. The "BPC Sample Configuration" can be selected when creating a Stand-alone profile only, 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 (PCAT) will invoke the standard Business Process Choreographer configuration script (bpeconfig.jacl) with the appropriate parameters which you have supplied.

BPEDB - SQL scripts

- All scripts to create and manage the BPEDB are in:
  - *${WAS_HOME}/ProcessChoreographer*
- Scripts described in detail in the Information Center
- SQL Scripts can:
  - create the database (for simple setups)
  - create tablespaces and dbspaces and the schema (for advanced setups)
  - clear the schema
  - drop the schema
  - drop the tablespaces and dbspaces
- Run SQL scripts manually or let bpeconfig.jacl run them
  - this is what happens during profile creation when the Cloudscape database for the "BPC Sample Configuration" is created

SQL scripts are in the *ProcessChoreographer* directory inside 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.

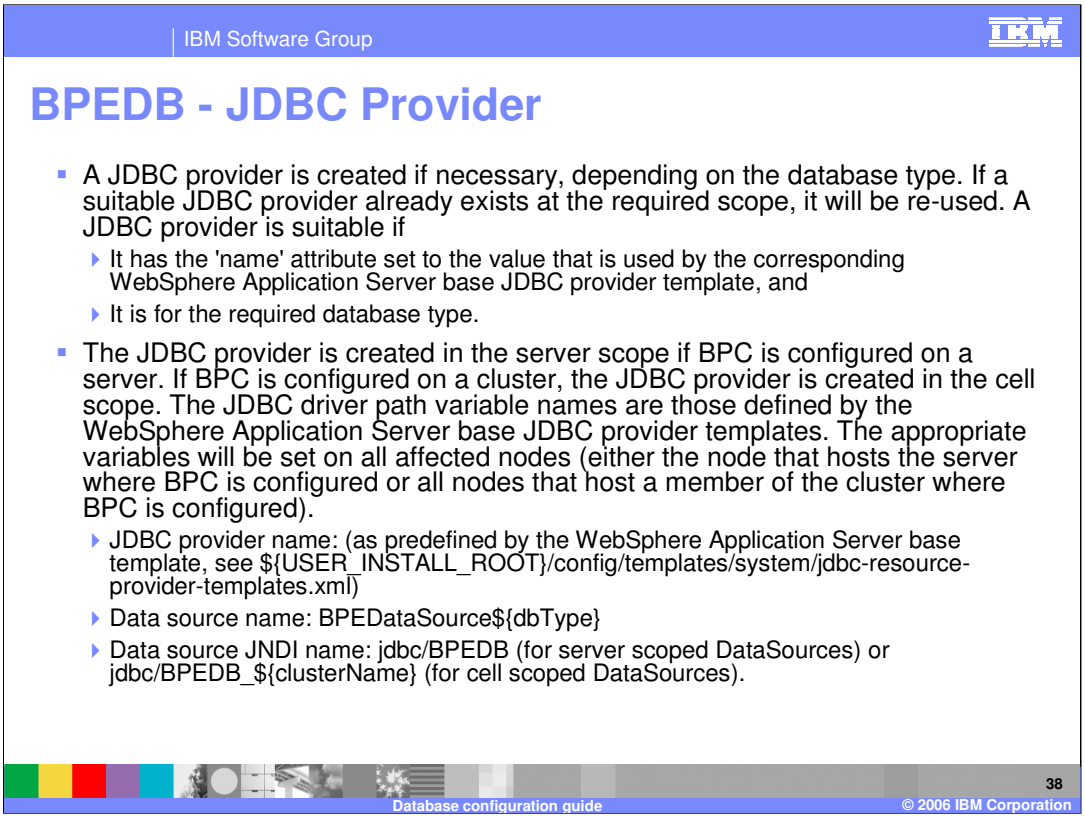

A JDBC provider is created if necessary, depending on the database type. JDBC provider is created in the server scope if Business Process Choreographer is configured on a server and in cell scope if its configured on a cluster.

BPEDB - Restrictions

- Restrictions
  - "Create new database" is not available for:
    - DB2 Universal OS/390 V7.1
    - DB2 Universal OS/390 V8.1
    - Oracle OCI
    - Oracle Thin
  - Cloudscape is not supported in Network Deployment
  - On z/OS, only Cloudscape, DB2 V7, and DB2 V8 are supported.
- BPEDB Tables
  - Too many to list here (about 100).

Create new database is disabled for DB2 Universal OS/390 V7.1 , V8.1, Oracle OCI and Oracle thin.

BPEDB (cont.)

- User ID privileges
  - Generally the specified user ID must have the appropriate authority or privilege to create tablespaces, tables, indexes and stored procedures
  - For the "Create new database" option, the specified user ID must have permission to create a new database
- Exported scripts
  - Business Process Choreographer does not export any scripts. All scripts needed to handle the BPEDB are shipped as part of the product in: ${WAS_HOME}/ProcessChoreographer

If you choose the “Create new database” option, then the specified user ID must have sufficient rights to create database. Necessary scripts can be found in ProcessChoreagrapher directory of WebSphere Process Server home folder.

Section

Messaging Engine Database (MED)

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.

MED - Supported Database Types

- The exact database version will have to be adjusted to the database types supported by WebSphere Application Server V6.0.2.
  - DB2 UDB 8.1 & 8.2 (DB2 Legacy CLI-based Type 2 JDBC Driver (XA))
- DB2 UDB 8.1 & 8.2 (DB2 Universal JDBC Driver Provider (XA))
- DB2 z/OS 7 & 8 (DB2 Legacy CLI-based Type 2 JDBC Driver (XA))
- DB2 z/OS 7 & 8 (DB2 Universal JDBC Driver Provider (XA))
- Informix 9.4 (Informix JDBC Driver (XA))
- Oracle 9i & 10g OCI (Oracle JDBC Driver (XA))
- Oracle 9i & 10g thin (Oracle JDBC Driver (XA))
- SQL Server 2000 (DataDirect ConnectJDBC type 4 driver for MS SQL Server (XA))
- SQL Server 2000 (WebSphere embedded ConnectJDBC driver for MS SQL Server (XA))

Message engine supports DB2, Informix, Oracle, and MS SQL server databases as shown here.

MED - DB Instances

- Each deployment target (server or cluster) can have its own database instance to host the message engine.
- One database is used to host the system bus.
- The other database is used to host the application bus.
- New data sources will be added to each JDBC provider.
- Data source naming:
  - System bus: _<node>.<server>-SCA.SYSTEM.<cell>.Bus
  - Application bus: _<node>.<server>-SCA.APPLICATION.<cell>.Bus
- Cloudscape database naming:
  - System bus:
    - ${USER_INSTALL_ROOT}/databases/com.ibm.ws.sib/(<node>.<server>|<cluster>)-SCA.SYSTEM.<cell>.Bus
  - Application bus:
    - ${USER_INSTALL_ROOT}/databases/com.ibm.ws.sib/(<node>.<server>|<cluster>)-SCA.APPLICATION.<cell>.Bus

One instance of the database is used for system bus and other instance of the database is used for application bus.

MED - Configuration during profile creation

- Stand-alone
  - The message engine database gets created during the installation.
- Network Deployment
  - The message engines need to be configured in the Advanced Configuration panel. The following administrative tasks are called:
    - Remote Destination Location:
      - configSCAForServer (remoteDestLocation is true)
      - configSCAForCluster (remoteDestLocation is true)
    - Local Destination Location:
      - configSCAForServer
      - configSCAForCluster
    - For both tasks, the following parameters are passed: appBusDataSource, systemBusDataSource, meAuthAlias, appBusSchemaName, systemBusSchemaName, createTables

Message engines database needs to be configured in the advanced configuration panel for network deployment where as it is created as part of installation for stand-alone.

MED (cont.)

- SQL Scripts
  - No SQL scripts. The ME DB needs to be created manually before its configured using the 'Advanced Configuration' panel.
- JDBC Provider
- The JDBC Provider is reused when:
  - The JDBC provider implementation class matches the one chosen in the advanced configuration (this usually means that if the same database types are used then the implementation classes usually match)
  - If no matching JDBC provider is found in the resource.xml file, then the jdbc-resource-provider-templates.xml under templates/system (profiles configuration) is searched for a matching JDBC provider. The provider is matched also against the implementation class.

The Message Engine Database should be created manually before it is configured using the 'Advanced Configuration' panel. The JDBC provider implementation class has to match with the one chosen in the advanced configuration.

MED (cont.)

- Restrictions
  - No known restrictions
- Tables
  - SIB000
  - SIB001
  - SIB002
  - SIBCLASSMAP
  - SIBKEYS
  - SIBLISTING
  - SIBOWNER
  - SIBXACTS

Tables created for the Message Engine database are shown here

MED (cont.)

- Exported scripts
  - None
- User ID permissions
  - Specified user ID must have permission to create tablespaces, tables, indexes and stored procedures

The User ID for message engine database must have the appropriate rights for creating tablespace, tables, indexes and stored procedures.

Section

ESB Logger Mediation Database (ESBDB)

This section covers the ESB Logger Mediation Database.

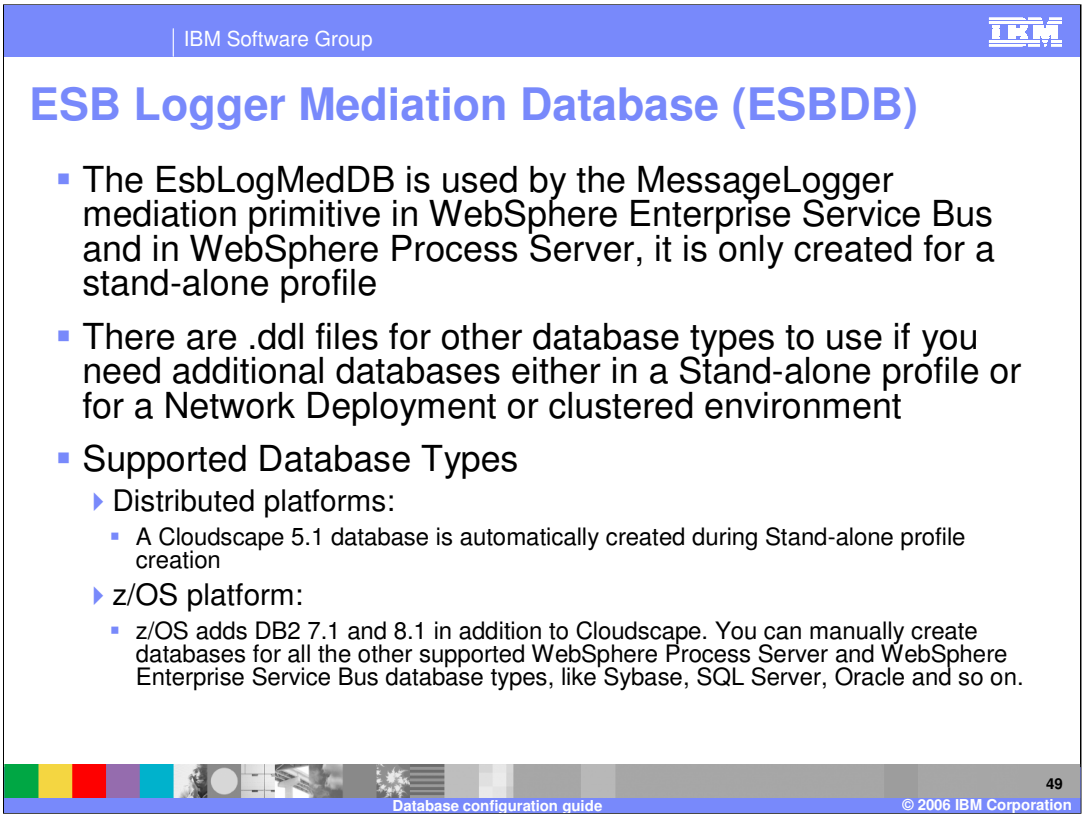

Enterprise service bus LogMedDB is used by the MessageLogger mediation primitive in WebSphere Enterprise Service Bus and WebSphere Process Server. Supported databases include DB2, Sybase, MS SQL server, Oracle and so on.

ESBDB - DB Instances

- Only one instance is automatically created for a Stand-alone profile
  - You can manually create as many other instances as you need
  - if needed, each message logger mediation primitive can be configured to use a different data source and therefore a different database

One instance is auto created for stand-alone profiles and you can manually create additional instances if required. The database is created in the databases directory of the WebSphere Process Server home folder if one does not already exist.

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 datasource

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 datasource creation.

ESBDB (cont.)

- 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
- JDBC Provider
  - If a suitable JDBC provider is obtained at server scope then this is used. The name of the data source is 'ESB Logger Mediation Data Source'. You can create your own data source to use if you want to configure a Message Logger mediation to use a different data source.

DDL files for all database types are located in the database named directory of the WebSphere Process Server home folder. If a suitable JDBC provider is obtained at server scope then this 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.

ESBDB - Restrictions

- For distributed only Cloudscape 5.1 is created in a Stand-alone profile. For z/OS Cloudscape 5.1, DB2 7.1 and DB2 8.1 can be created in a Stand-alone profile. No automatic creation is performed for Network Deployment.
- Currently the MessageLogger mediation runtime code assumes a fixed table schema name for distributed although in z/OS this has been changed to allow different schema names.

For distributed platforms only, Cloudscape 5.1 is created in a stand-alone profile, No auto creation is performed for Network Deployment.

ESBDB (cont.)

- Tables
  - Only one table is created called MSGLOG. The table contains 6 columns.
- Export scripts
  - None
- User ID permissions needed
  - Create tablespaces, tables, indexes and stored procedures

A single table MSGLOG is created. User ID must have permissions to create tablespace, tables, indexes and stored procedures.

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 and Oracle OCI
  - Using Cloudscape requires separate RepositoryDB

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 has been added for DB2 V7.0 for z/OS, DB2 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.

SBDB – DB instance

- 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

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.

SBDB - Actions during profile creation

- For Stand-alone profile:
  - As part of profile creation, if Cloudscape is selected the configDynamicArtifactRepository.ant script is run. This script calls the configDynamicArtifactRepository.jacl command to create the Cloudscape JDBC provider at the Node level for the RepositoryDB.
  - The RepositoryDB Database is created as part of server installation and not as part of profile creation. This is created in <WAS_HOME>/cloudscape/databases/Repository DB. All profiles for the server installation will use the same database. Multiple servers can not be started and use rules as Cloudscape only allows one JVM access to the database at a time.
    For Network Deployment profile/configuration
  - The Common DB for the cell is expected to be setup with the appropriate tables for Selector and Business Rules. All Selectors and Business Rules for the server in a stand-alone profile or Network Deployment profile will use this database.

The RepositoryDB database is created as part of server installation. configDynamicArtifactRepository.jacl is used to create a Cloudscape JDBC provider. Selectors and Business Rules for the server in a stand-alone profile or Network Deployment profile will use a common database .

SBDB - SQL scripts

- SQL scripts for table creation at:
  - <WAS_HOME>profileTemplates\default.wbiserver\actions\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 datasource.

Scripts for table creation are available under the scripts folder of WebSphere Process Server home directory. configDynamicArtifactRepository.jacl is used during stand-alone profile creation.

SBDB - JDBC Provider

- If Cloudscape is selected for the stand-alone profile, the configDynamicArtifactRepository.jacl file is run. It will create a Cloudscape JDBC provider at the Node level. It will also create a datasource for the RepositoryDB with a JNDI name of jdbc/wbiserver/DynamicArtifactRepositoryDataSource. In a Network Deployment environment, the Common DB datasource (jndi name jdbc/WPSDB) is used with its appropriate JDBC provider.
- If both JDBC Providers and datasources are configured in a system, the repository code does a look up to find the provider. It first does a lookup of the JNDI name **jdbc/wbiserver/DynamicArtifactRepositoryDataSource**. If that is not found, it will then do a lookup of the JNDI name jdbc/WPSDB. If neither is found, then a NamingException is thrown. These lookup values are hard-coded in the repository code. Any changes to the JNDI names in the datasource configuration will cause the repository code to fail.

The configDynamicArtifactRepository.jacl script creates a Cloudscape JDBC provider scoped at the node level. This also creates a datasource for RepositoryDB. If both JDBC Providers and datasources are configured in a system, the repository code does a look up to find the provider starting with the JNDI name jdbc/wbiserver/DynamicArtifactRepositoryDataSource .

SBDB - Restrictions

- DB2 V7.0 for z/OS limits primary key 256 bytes.
  - Ensure the Name and Namespace of the Selectors and Business Rules do not exceed 245 bytes.
  - These values can be modified in WebSphere Integration Developer.
  - The Type value is not changeable and is approximately 8 bytes.
  - There are Technotes on this limitation:
    - Technote 1226352 (http://www-1.ibm.com/support/docview.wss?rs=0&q1=1226352&uid=swg21226352&loc=en_US&cs=utf-8&cc=us&lang=en)
    - Technote 1226351 (http://www-1.ibm.com/support/docview.wss?rs=0&q1=1226351&uid=swg21226352&loc=en_US&cs=utf-8&cc=us&lang=en)
- In WebSphere Process Server V6.0 and V6.0.1, Selectors and Business Rules do not support Informix or Microsoft® SQL Server.
- When applications which contain Selectors or Business Rules are uninstalled from the server or cell, the artifacts stored in the CommonDB or RepositoryDB are not removed. These must be removed manually per the instructions in the Information Center (http://publib.boulder.ibm.com/infocenter/dmndhelp/v6rxmx/index.jsp?topic=/com.ibm.wsps.adm.doc/doc/tadm_rembruleselectordatarepository.html)

The primary key for Selectors and Business Rules is composed of the Name, Namespace, and Type. DB2 V7.0 for z/OS has a limitation on the length of the primary key to not exceed 256 bytes. WebSphere Process Server V6.0 and V6.0.1.x, Selectors and Business Rules do not support Informix or Microsoft SQL Server .

SBDB (cont.)

- Tables
  - For Selectors and Business Rules, three tables are used to store the artifacts.
  - ByteStore
  - ByteStoreOverflow
  - AppTimestamp
- Exported Scripts
  - The scripts are exported to the same location as CommonDB (except in the scenario of Cloudscape where no scripts are exported)

Three tables are generated for selectors and business rules as shown.

SBDB – User ID privileges

- Generally the specified user ID must have the appropriate authority or privilege to create tablespaces, tables, indexes and stored procedures.

The specified User ID should have permission to create tablespace, tables, indexes and stored procedures.

Summary

- WebSphere Process Server components use several databases
- Database creation is scripted
- More information at http://publib.boulder.ibm.com/infocenter/dmndhelp/v6rxmx/index.jsp

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

Trademarks, copyrights, and disclaimers

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

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

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

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

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

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

Linux is a registered trademark of Linus Torvalds.

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

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

Information is provided "AS IS" without warranty of any kind. THE INFORMATION PROVIDED IN THIS DOCUMENT IS DISTRIBUTED "AS IS" WITHOUT ANY WARRANTY, EITHER EXPRESS OR IMPLIED. IBM EXPRESSLY DISCLAIMS ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NONINFRINGEMENT. IBM shall have no responsibility to update this information. IBM products are warranted, if at all, according to the terms and conditions of the agreements (e.g., IBM Customer Agreement, Statement of Limited Warranty, International Program License Agreement, etc.) under which they are provided. Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products in connection with this publication and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. IBM makes no representations or warranties, express or implied, regarding non-IBM products and services.

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

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

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