

IBM Operational Decision Manager
Version 8 Release 6

*Configuring Operational Decision
Manager on JBoss 6.1*



Note

Before using this information and the product it supports, read the information in "Notices" on page 55.

This edition applies to version 8, release 6, modification 0 of Operational Decision Manager and to all subsequent releases and modifications until otherwise indicated in new editions.

© **Copyright IBM Corporation 2012, 2014.**

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Chapter 1. Configuring on JBoss 6.1

EAP	1
Configuring Rule Execution Server on JBoss 6.1 EAP	1
Before you start	1
Step 1: Selecting and applying the persistence type	2
Step 2: Restricting database user permissions	3
Step 3: Establishing the database credentials	4
Step 4: Configuring security	4
Step 5: Creating a driver and data source	6
Installing a JDBC driver as the core module	6
Creating the data source	7
Step 6: Deploying the execution unit (XU) RAR file	7
Step 7: Deploying the Rule Execution Server management EAR	9
Step 8: Creating a schema for the Rule Execution Server database	9
Creating a database schema by using the Rule Execution Server console	9
Installation Settings wizard overview	9
Opening the Rule Execution Server console	10
Step 1: Welcome to the Installation Settings wizard	11
Step 2: Choose the database schema	11
Step 3: Review the database schema	12
Step 4: The Installation Settings wizard report	12
Creating a database schema by running SQL scripts	13
Step 9: Deploying the hosted transparent decision service	14
Step 10: Verifying the configuration	14
Optional configuration steps	15
Repackaging the Rule Execution Server archive using Ant	15
Configuring Rule Execution Server for different environments	15
Rule Execution Server deployment for high availability and scalability	16
Cluster configuration and topology	16
Deploying Rule Execution Server in a cluster	17
Configuring the Decision Center consoles on JBoss 6.1 EAP	19
Before you start	19
Step 1: Restricting database user permissions	20
Step 2: Creating a data source and connection pool	21
Step 3: Configuring security on JBoss	22
Defining groups and permissions	22
Declaring custom groups	23
Step 4: Deploying the Decision Center EAR on JBoss	24

Step 5: Verifying the deployment of the Decision Center consoles	25
Step 6: Completing your Decision Center configuration	25
Completing the configuration from the Decision Center Enterprise console	25
Installation Settings wizard overview	25
Step 1: Configure the database	26
Step 2: Set up message files	27
Step 3: Set up groups	28
Step 4: Set the persistence locale	28
Step 5: Set configuration parameters	29
Completing the configuration by using Ant tasks	30
Setting up the Ant tasks environment	30
Creating the database schema	32
Defining and uploading message files	36
Setting the persistence locale	37
Adding or removing configuration parameters	37
Repackaging the Decision Center archive	39
Configuring the search function of the Decision Center Business console	39
Using a remote Solr search engine	39
Setting parameters for the Solr search engine	40
Additional steps to configure testing and simulation	41
Before you start	41
Step 1: Creating Decision Warehouse database resources	42
Step 2: Creating Decision Runner database resources	44
Step 3: Deploying the testing and simulation archive	44
Step 4: Checking the availability of the testing and simulation services	45
Step 5: Using an Ant task to package the SSP archive	46
Distributing rule testing to multiple servers	47
Tuning the Decision Runner web application	48

Chapter 2. Verifying your configuration of Decision Center. 51

Publishing a project to Decision Center	51
Opening the Decision Center consoles	52

Notices	55
Trademarks	57

Index	59
------------------------	-----------

Chapter 1. Configuring on JBoss 6.1 EAP

Configuring Rule Execution Server on JBoss 6.1 EAP

To use Rule Execution Server on a new instance of JBoss, you must establish your database, deploy the provided archives for this server, and perform a number of configuration tasks.

JBoss 6.1 EAP can be used in one of two operating modes: as a stand-alone server or as a managed domain.

As a stand-alone server, the operating mode is similar to the previous versions of JBoss with the option to run the server in a clustered configuration.

In a managed domain, you can run multiple instances of JBoss 6.1 EAP and manage them from a single point of control. A domain consists of a domain controller, one or more host controllers, and zero or more server groups per host.

See About stand-alone servers and managed domains

Decision Server requires a full profile. A sample full profile is provided for both stand-alone and domain modes.

The JBoss Management Console and Management CLI interfaces allow administrators to manage application deployments in a production environment. To test application deployment, developers have the option to use a file system deployment scanner, or to deploy and undeploy using Maven.

Before you start

To configure Rule Execution Server on JBoss you must follow a series of steps

The first two steps concern your choice of persistence and permissions on the database. These steps are meant for a database administrator. Then, steps 3 to 5 depend on whether:

- You want to create an empty database for Rule Execution Server before or during the deployment of the WAR file.
- You want to create or clear the database schema using the Installation Settings wizard in the Rule Execution Server console, or using the SQL scripts directly.



















There are also some “Optional configuration steps” on page 15.

This configuration guide presents generic instructions for all supported databases. To help beginner users, the examples are given for the embedded Derby database.

Note: Make sure to clean up prior installations. Delete the content of the /tmp directory in the folder in which you deploy Rule Execution Server: <JBOSS_HOME>/standalone/deployments. This ensures that JSP pages are recompiled during the installation.

What steps to follow

The following table summarizes the steps that you follow to configure Rule Execution Server on JBoss depending on the persistence type that you choose.

Installation steps	Persistence		
	File	Data source	JDBC
"Step 1: Selecting and applying the persistence type"		Default persistence mode	
"Step 2: Restricting database user permissions" on page 3	Not applicable		
"Step 3: Establishing the database credentials" on page 4	Not applicable		
"Step 5: Creating a driver and data source" on page 6	Not applicable		Not applicable
"Step 4: Configuring security" on page 4			
"Step 6: Deploying the execution unit (XU) RAR file" on page 7			
"Step 7: Deploying the Rule Execution Server management EAR" on page 9			
"Step 8: Creating a schema for the Rule Execution Server database" on page 9	"Creating a database schema by using the Rule Execution Server console" on page 9	Not applicable	
	"Creating a database schema by running SQL scripts" on page 13	Not applicable	
"Step 9: Deploying the hosted transparent decision service" on page 14	Optional	Optional	Optional
"Step 10: Verifying the configuration" on page 14	Recommended	Recommended	Recommended

Related concepts:

Troubleshooting Rule Execution Server on JBoss

Step 1: Selecting and applying the persistence type

You can change the default datasource RuleApp and Java™ XOM persistence settings by running an Ant script that generates a new Rule Execution Server management archive.

Typically, you do this if you are in development mode. This step does not apply to beginners who work with the embedded Derby database.

For you to change the persistence settings, the distribution provides an Ant script in the `<ODM_InstallDir>/executionserver/bin/resetup.xml` file. Use it to create a new instance of the Rule Execution Server management archive and, in the case of a Java EE application server, the execution unit (XU).

Solaris users

If you use file-based persistence on Solaris, your file system must support all characters used in directory and file names that are present in the ruleset path (RuleApp name and ruleset name). Set the **LANG** system property with the encoding that is compatible with your package and rule names, for example `en_US.UTF-8`.

Decision Warehouse

If you select the file persistence type for RuleApps, you cannot use Decision Warehouse.

MySQL persistence

If you choose to use MySQL as a persistence back end, add or set the following properties in the MySQL configuration file: `my.ini` on Windows or `my.cnf` on UNIX operating systems:

```
sql-mode=STRICT_ALL_TABLES
max_allowed_packet=1073741824
```

For more information about these settings, see the MySQL 5.0 reference manual: 5.1.7. Server SQL Modes and 5.1.4 Server System Variables.

Related tasks:

“Repackaging the Rule Execution Server archive using Ant” on page 15

When you need to repackage a Rule Execution Server archive to configure the Rule Execution Server, you can use an Ant task, provided that you have set up the Ant task environment.

Step 2: Restricting database user permissions

If Rule Execution Server data is stored in a database, the database administrator might require that you provide the specific permissions to access the database.

Note: This step applies when database access needs to be restricted. If you manage the database yourself (for example, you use an embedded database for test purposes) or if you do not need further restrictions, skip this step and proceed to the next configuration step.

Connection to the Rule Execution Server database, as established in the data source credentials, and any subsequent requests to the database are handled through a database user. This database user (name and password), for example `resdbUser`, is defined by the database administrator and has no relation to the standard Rule Execution Server groups.

The following table gives the typical list of permissions that the database administrator must define on the Rule Execution Server database, with attention given to the type of operations. Some supported databases do not require all these permissions.

Database permission	Operation	
	Browse and edit rulesets and RuleApps	Create the Rule Execution Server schema
CREATE ANY INDEX	Not required	Required
DROP ANY INDEX	Not required	Required
CREATE ANY SEQUENCE	Not required	Required
DROP ANY SEQUENCE	Not required	Required
SELECT ANY SEQUENCE	Required	Not required
CREATE ANY TABLE	Not required	Required
DROP ANY TABLE	Not required	Required
INSERT ANY TABLE	Required	Not required
SELECT ANY TABLE	Required	Not required
UPDATE ANY TABLE	Required	Not required
DELETE ANY TABLE	Required	Not required
CREATE ANY TRIGGER	Not required	Required
CREATE ANY VIEW	Not required	Required
DROP ANY VIEW	Not required	Required

Step 3: Establishing the database credentials

You must establish the credentials of the database dedicated to Rule Execution Server if you are using database persistence.

These credentials are required to establish the datasource, which you will do in the next step.

Note: Skip to the next step if you are a beginner using the embedded Derby database.

If you set persistence to file, you can skip all the database-related tasks and proceed directly to “Step 7: Deploying the Rule Execution Server management EAR” on page 9.

If a database does not already exist, create one now by following the instructions for that database type.

After the database is created, you can use the Installation Settings wizard in the Rule Execution Server console or you can run the provided SQL scripts to create the database schema so that it includes the tables and views.

Step 4: Configuring security

You control access to Rule Execution Server and enforce security by defining user groups and associated roles.

JBoss management interfaces security

The management interfaces in JBoss 6.1 EAP are secured by default to prevent security violation from remote systems. Local non-HTTP access is protected by a SASL mechanism which consists in a negotiation between the client and the server

when the client connects the first time from the localhost. To access the JBoss administration console, you must register an Administrator user as follows:

1. Open a command prompt in the <JBOSS_HOME>/bin directory.
2. Enter `./add-user[.sh|.bat] <username> <password> <ManagementRealm>`
 For example: `./add-user[.sh|.bat] admin admin?12 ManagementRealm`

Rule Execution Server security

The main groups and their associated default user and password are summarized in the following table.

Group	Use	Default user/password
resAdministrators	Gives a user full administrator rights to: <ul style="list-style-type: none"> • Access and use the Rule Execution Server console to populate the database schema • Deploy, browse, and modify RuleApps • Monitor the decision history, purge and back up the history • Run diagnostics and view server information 	resAdmin/resAdmin+0
resDeployers	Gives a user the rights to: <ul style="list-style-type: none"> • Deploy, browse, and modify RuleApps • Test rulesets 	resDeployer/resDeployer+0
resMonitors	Gives a user the rights to: <ul style="list-style-type: none"> • View RuleApps • Monitor decision history and access Decision Center reports 	resMonitor/resMonitor+0

You can add roles as required by Decision Server by editing the `application-roles.properties` file.

1. In Standalone mode:
`<JBOSS_HOME>/standalone/configuration/application-roles.properties`
2. In Domain mode:
`<JBOSS_HOME>/domain/configuration/application-roles.properties`

Here is an example of the `application-roles.properties` file:

```
resAdmin=resAdministrators,resDeployers,resMonitors
resDeployer=resDeployers,resMonitors
resMonitor=resMonitors
```

You can add users as required by Decision Server:

1. Open a command prompt in the <JBOSS_HOME>/bin directory.
2. Enter the script `./add-user[.sh|.bat] -a <username> <password>`

Here are examples based on the users listed in the table above:

- `./add-user[.sh|.bat] -a -u resAdmin -p resAdmin+0 -ro "resAdministrators,resDeployers,resMonitors"`
- `./add-user[.sh|.bat] -a -u resDeployer -p resDeployer+0 -ro "resDeployers,resMonitors"`
- `./add-user[.sh|.bat] -a -u resMonitor -p resMonitor+0 -ro "resMonitors"`

Step 5: Creating a driver and data source

If you use data source persistence, you must create a driver before creating the data source.

Installing a JDBC driver as the core module

You must install a driver before creating the data source.

Before you begin

Download the JDBC driver from your database vendor.

Procedure

1. Create a file path structure under the `<JBoss_HOME>/modules` directory. For example, for a Derby driver, create the following directory structure:
`<JBoss_HOME>/modules/com/ibm/derby/main.`
2. Copy the JDBC driver JAR file into the main subdirectory.
3. In the main subdirectory, create a `module.xml` file. Example for Derby:

```
<module xmlns="urn:jboss:module:1.1" name="com.ibm.derby">
  <properties>
    <property name="jboss.api" value="unsupported"/>
  </properties>
  <resources>
    <resource-root path="derby.jar"/>
    <!-- Insert resources here -->
  </resources>
  <dependencies>
    <module name="javax.api"/>
    <module name="javax.transaction.api"/>
    <module name="javax.servlet.api" optional="true"/>
  </dependencies>
</module>
```

The module XSD is defined in the `<JBoss_HOME>/docs/schema/module-1_2.xsd` file.

4. Start the server:
 - a. Go to `<JBoss_HOME>/bin`
 - b. In stand-alone mode, run `./standalone[.sh|.bat] -c standalone-full.xml`
 - c. Start the Management CLI: `<JBoss_HOME>/bin/jboss_cli[.sh|.bat]`
 - d. Enter connect.
5. To add the JDBC driver module as a driver, run the following CLI command:

```
/subsystem=datasources/jdbc-driver=DRIVER_NAME:add(driver-name=DRIVER_NAME,driver-module-name=MOD
```

Example:

```
/subsystem=datasources/jdbc-driver=derby:add(driver-name=derby,driver-module-name=com.ibm.derby,d
```

Results

The JDBC driver is now installed and set up as the core module. It can be referenced by a data source.

Creating the data source

You can configure the data source by using the JBoss Management Console or the Management CLI. Here, you use the Management Console.

Procedure

1. Log in to the JBoss Management Console. By default, `http://localhost:9990/console/App.html#server-overview`.
2. Navigate to the **Datasources** panel.
 - a. Choose one of the following options.
 - In Standalone mode:
 - Click the **Profile** tab in the top right of the console.
 - In Domain mode:
 - 1) Click the **Profiles** tab in the top right of the console.
 - 2) Select the appropriate profile in the top left.
 - 3) Expand the **Subsystems** menu in the left.
 - b. Click **Connector > Datasources**
3. Select the **XA Datasources** tab.
4. Click the **Add** button at the top of the **XA Datasources** panel, and then specify the details for the new data source.
 - a. Name:
 - **Name:** RESDatasource
 - **JNDI:** java:/jdbc/resdatasource
 - b. JDBC Driver:
 - **Name:** choose the driver
 - **XA Data Source Class:** org.apache.derby.jdbc.EmbeddedXADataSource
 - c. XA Properties (at least the connection URL):
 - **Name:** CreateDatabase **Value:** Create
 - **Name:** DatabaseName **Value:** /tmp/data/resdb
 - d. Connection Settings:
 - **Username:** the user name for a new connection. Example: RES
 - **Password:** the password for the new connection. Example: RES
 - **Security domain:** the security domain that defines the **javax.security.auth.Subject** parameter used to distinguish connections in the pool. Example: <Empty>
5. Select the newly created data source and then click the **Enable** button.

Example

You can also create the data source by using a script. Here is an example script for Derby:

```
./jboss-cli.sh --connect 'xa-data-source add --name=RESDatasource --jndi-name=java:/jdbc/resdatasource'
./jboss-cli.sh --connect '/subsystem=datasources/xa-data-source=RESDatasource/xa-datasource-properties'
./jboss-cli.sh --connect '/subsystem=datasources/xa-data-source=RESDatasource/xa-datasource-properties'
./jboss-cli.sh --connect 'xa-data-source enable --name=RESDatasource'
```

Step 6: Deploying the execution unit (XU) RAR file

Next, you deploy the resource adapter archive (RAR) for the execution unit (XU).

About this task

The execution unit (XU) is a resource adapter for Java EE application servers. It is supplied as a RAR archive that you must deploy on your application server. Deploy the XU by copying the XU RAR and XML files to the deployment directory of your application server.

In addition to the procedure described here, you also have the option to deploy the XU inside the application (embed the XU into the EAR). There might be some instances where, due to your application constraints, you have to deploy the XU inside the application. It is your decision to choose the appropriate deployment mode of the XU (either embed into the EAR or deploy as a global connector). Refer to the application server documentation for instructions on packaging a connector into an EAR.

Note:

- When the XU is deployed as a global connector, its third-party libraries (such as ASM) might be used by the deployed J2EE applications instead of the libraries deployed in the application server.
- When the XU is deployed as a global connector, use a parent `last` setting for its J2EE application if your J2EE application does not support the version of the libraries from independent software vendors that are distributed with Decision Server. If this is not possible, you might have to embed the XU into the EAR that is executing the rules.
- When using an embedded XU packaging, use a parent `last` setting for the code library if the version of the libraries from independent software vendors that are deployed at the code library level of the application server is not compatible with the XU.

Tip: When the default configuration in shared mode is not appropriate for your use case, you can configure Rule Execution Server so that it is scoped to a single Java™ EE application. See *Configuring Rule Execution Server in scoped mode*.

To deploy the XU RAR, copy the file `jrules-res-xu-JBoss61EAP.rar` from `<ODM_InstallDir>/executionserver/applicationservers/JBoss6EAP` into the `<JBoss_HOME>/standalone/deployments` directory.

When deployed to JBoss 6.1 EAP, the XU uses the JDK logging API.

Procedure

To configure the XU log level with the JBoss Management Console, proceed as follows:

1. Log in to the JBoss Management Console.
2. Navigate to the logging subsystem configuration.
 - Standalone server
 - a. Click **Profile** in the top right of the console.
 - b. Click **Core > Logging** on the left.
 - Managed domain
 - a. Click **Profiles** in the top right of the console.
 - b. Select the profile on the left, and then click **Core > Logging**.
3. Click the **Log Categories** tab.
4. Click the **Add** button and define the following details.

- **Name:** `com.ibm.rules.res.execution`
 - **Log Level:** specifies which message levels are logged by this logger. Message levels lower than this value are discarded.
 - **User Parent Handler:** specifies whether or not this logger sends its output to the parent logger.
5. Save.
 6. Restart the server.

Step 7: Deploying the Rule Execution Server management EAR

To deploy the Rule Execution Server management EAR archive, you must copy it to the deployment directory of your application server.

Procedure

Copy the file `jrules-res-management-JBoss6EAP.ear` from `<ODM_InstallDir>/executionserver/applicationservers/JBoss6EAP/` to the `<JBoss_HOME>/standalone/deployments` directory.

Step 8: Creating a schema for the Rule Execution Server database

After you have created a dedicated Rule Execution Server database, you can create the schema for the database by running SQL scripts from either the Rule Execution Server console or the SQL tool of your database.

Creating a database schema by using the Rule Execution Server console

To create a schema for the Rule Execution Server database, you can use the Installation Settings wizard if you work on Windows and other supported distributed platforms.

Installation Settings wizard overview:

On Windows and distributed platforms only, you can use the Installation Settings wizard of the Rule Execution Server console to choose a database type and create a schema that contains the necessary tables and views.

You can use the Installation Settings wizard to configure Rule Execution Server with database persistence.

Note: You must have created the data source connection before you use the Installation Settings wizard.

The Installation Settings wizard creates all the required tables for Rule Execution Server and Decision Warehouse.

If you are using file persistence or have an existing database schema, the Installation Settings wizard does not open when you sign in to the Rule Execution Server console. If you want to modify the database schema after having created the database tables already, you must run the SQL scripts in the database client.

The combination of persistence settings for RuleApps and managed Java XOMs affects the way in which you use the Installation Settings wizard.

- If the RuleApp persistence and the Java XOM persistence are both set to file, no wizard is presented.
- If the RuleApp persistence and the Java XOM persistence are both set to datasource or jdbc, the **RuleApp persistence details** part of the Installation Settings wizard opens for you to create the schema for RuleApps and the Decision Warehouse trace when you sign in to the Rule Execution Server console. After you have completed this step, the **Java XOMs persistence details** part of the wizard opens for you to configure the database for Java XOM persistence.
- If the RuleApp persistence is set to file and the Java XOM is set to datasource or jdbc, you see only **Java XOMs persistence details**, and you cannot use the Decision Warehouse.
- If the RuleApp persistence is set to datasource and the Java XOM is set to file or is not defined, only **RuleApp persistence details** opens for you to create the schema for RuleApps when you sign in to the Rule Execution Server console. The wizard does not show **Java XOMs persistence details**.

The following table summarizes the cases.

Persistence		RuleApps	
		file	datasource/JDBC
Java XOMs	file	No Installation Settings wizard	RuleApps persistence details only
	datasource	Java XOMs persistence details only No Decision Warehouse	Complete Installation Settings wizard

Opening the Rule Execution Server console:

To open the Rule Execution Server console and the Installation Settings wizard, you must sign in with administrator rights.

Procedure

1. Start the application server by double-clicking run.bat in <JBoss_HOME>/bin
2. Open the Rule Execution Server console in a web browser by typing res at the root URL on the host machine:
http://localhost:8080/res
 - If your browser is not running on the same host as the application server, replace localhost with the address of the machine.
 - If the web application is mapped to a host which is defined on a port that is different from the default of the server, change the port number to the host port number.
3. Sign in to the Rule Execution Server console as the administrator. For example:

User ID

resAdmin+0

Password

resAdmin+0

Results

If you use database persistence and the database schema is empty, the Installation Settings wizard opens and you can use it to complete the installation.

Step 1: Welcome to the Installation Settings wizard:

If you open the Rule Execution Server console with `datasource` or `jdbc` as the persistence setting and an empty database schema, the Installation Settings wizard opens.

The wizard can display the following parts:

- **RuleApp persistence details:** This part opens if you set `datasource` persistence for RuleApps, regardless of the persistence type for managed Java XOMs.
- **Java XOMs persistence details:** This part opens after **RuleApp persistence details** if you have set `datasource` or `jdbc` persistence for both RuleApps and Java XOMs. In this case, you go through the same steps twice.

The wizard starts with **Java XOMs persistence details** if you have set the persistence type to `file` for RuleApps, and to `datasource` or `jdbc` for managed Java XOMs.

Both parts of the wizard are similar, and you use them in the same way:

1.

The Welcome page provides the following information:

- Persistence details about the type of database used. This includes information about the driver and JDBC URL.
- A brief description of the purpose of the Installation Settings wizard.
- A diagnostic report that provides information about why the persistence check failed (because you have not created the database schema).

2. Click **Next** to proceed.

Step 2: Choose the database schema:

Select a schema for your database. The wizard includes settings for different databases, or you can select a customized SQL script.

About this task

You select an available Rule Execution Server database schema or you upload a custom schema.

Procedure

1. In the **Database schema selected** field, select an available type of database schema.

If you select a `db2` or `db2_os390` schema, an extra field opens so that you can enter the name of the buffer pool, which is used to create the Decision Warehouse tablespace. This buffer pool must have a page size of 32K. Check the DB2[®] documentation for information about how to create a 32K buffer pool.

Note:

The scripts for creating the Decision Warehouse database on DB2 are written for databases that use automatic storage. When you use the Installation Settings

wizard, you create both the Rule Execution Server and the Decision Warehouse database, so your database must use automatic storage.

If you have not configured your DB2 database to use automatic storage, you cannot use the Rule Execution Server console to create the Rule Execution Server tables. In this case see “Creating a database schema by running SQL scripts” on page 13.

2. Optional: Select **custom** if you want to use a customized SQL script, and then click **Browse** to select the location of the custom script.
3. Click **Next** to review the database schema.

Step 3: Review the database schema:

After you have selected a database schema, you confirm the creation of a schema for Rule Execution Server.

About this task

You can also use SQL drop statements that flush data from an existing table, and view the SQL statements. Ensure that you have a backup of your database resources.

Procedure

1. Select from the following options:

Create SQL schema “resdbUser”

Select this option to run the SQL statement for the schema type selected in the previous step.

Keep drop SQL statements

Select this option to flush data from an existing Rule Execution Server database.

Show SQL statements

Click this option to display the SQL statements.

2. Click **Execute** to start the options that you have selected.

Step 4: The Installation Settings wizard report:

After you have selected and confirmed the schema, the Installation Settings wizard reports the status of the schema creation.

Procedure

1. Click **Show execution details** to view the list of executed SQL statements.
2. Click **Finish** to open the Explorer tab in the Rule Execution Server console.

What to do next

If you have just worked in **RuleApps persistence details** and the persistence setting for managed Java XOMs is `datasource` or `jdbc`, the **Java XOMs persistence details** part of the wizard opens for you to repeat this procedure.

Creating a database schema by running SQL scripts

After you have created an empty database, you can create the Rule Execution Server database schema by running SQL scripts if you do not want to use the Installation Settings wizard.

The SQL scripts can be found in this directory: `<ODM_InstallDir>/executionserver/databases`. This directory also contains a readme file where you can find additional information about the scripts.

The script that creates the database schema is named `repository_<DatabaseName>.sql`.

Note:

If you want to use Decision Warehouse, you can also create the required database table by running the script `trace_<DatabaseName>.sql`. If you are also persisting the Java XOM in a database, you must create these tables by running the `xomrepository_<DatabaseName>.sql` script.

If you use Command Editor to run the scripts, you must log in with the credentials you use for the data source for Rule Execution Server.

Use any tool that can handle SQL to import and run the SQL scripts. The tools provided for each database include:

Database	Database tool
IBM® DB2	DB2 command line processor
Derby	ij command line processor
MySQL	mysql command line processor
Oracle	sqlplus command line processor
Postgre SQL	Postgre SQL command line tool
SQL Server	Query Tool
Sybase	isql command line processor

To access the database, the database user must have a user ID and a password. The database user must also have:

- complete privileges on the tables and view of the schema (create, insert, delete)
- *create index* privileges

On Oracle, the database user must also have *create trigger* and *create sequence* privileges.

When using an Oracle database, run all the scripts in the SQL Plus client.

When using DB2, the scripts that create the Rule Execution Server database tables are written for databases that use automatic storage. The following conditions apply:

- BP32K is the buffer pool that is expected in SYSCAT.BUFFERPOOLS. If BP32K is not there, you can use the existing buffer pool or create a new buffer pool named BP32K. Use the following command to query SYSCAT.BUFFERPOOLS for the existing buffer pool:

```
Select * from SYSCAT.BUFFERPOOLS
```

Otherwise, use the following command to create a buffer pool named BP32K:

```
CREATE BUFFERPOOL BP32K SIZE 2000 PAGESIZE 32K
```

- You must update the trace_db2.sql script and select the custom option in the Installation Settings wizard to run it. Modify the following line in the script to specify storage for the tablespace:

```
CREATE TABLESPACE RESDWTS PAGESIZE 32K BUFFERPOOL BP32K;
```

The following line gives an example of the tablespace specification in the script:

```
CREATE TABLESPACE RESDWTS PAGESIZE 32K MANAGED BY Database  
USING [ FILE 'C:\DB2\Container.file' 640 ] BUFFERPOOL BP32K;
```

- You might have to further modify the script based on your database settings.

Step 9: Deploying the hosted transparent decision service

You can deploy the hosted transparent decision service on your application server.

Procedure

1. Copy the file jrules-res-htds-JBOSS61EAP.ear from <ODM_InstallDir>/executionserver/applicationserver/JBoss6EAP to the <JBoss_HOME>/standalone/deployments directory.
2. Set the **ruleset.xmlDocumentDriverPool.maxSize** ruleset property to the appropriate value.
See Setting the ruleset.xmlDocumentDriverPool.maxSize property.

What to do next

For more information about checking that the hosted transparent decision service has been deployed successfully, refer to the Rule Execution Server console online help.

Step 10: Verifying the configuration

You can verify that Rule Execution Server has been successfully configured by running Rule Execution Server diagnostics.

About this task

If the diagnostics are performed before any XUs have been started, the test is passed and a message displays to verify that no XU(s) have been initialized.

Note: To let a scalable number of users access resources through the Java components, JCA assigns the task of implementing connection pooling to application server vendors.

Procedure

1. Open the Rule Execution Server console by typing res at the root URL on the host machine:
`http://localhost:8080/res`
If your browser is not running on the same host as the application server, replace localhost with the address of the machine. If the web application is mapped to a host that is defined on a different port to 8080, change the port number from 8080 to the host port number.
2. Sign in to the Rule Execution Server console.
3. Click the **Diagnostics** tab.
4. Click **Run Diagnostics**.

Optional configuration steps

You can enhance your configuration with additional options, such as multiserver configuration, or deploying onto a server cluster.

Repackaging the Rule Execution Server archive using Ant

When you need to repackage a Rule Execution Server archive to configure the Rule Execution Server, you can use an Ant task, provided that you have set up the Ant task environment.

Before you begin

Before you run the `res-setup` Ant task, you must set up the Ant task environment correctly. For more information, see [Setting up your environment to automate processes](#).

About this task

By default, persistence is set to `datasource`. To change the persistence type, you can use an Ant task. An Ant script is provided with the distribution for this purpose. The script creates new archives that use a specific persistence mode.

The following procedure repackages the archives to change the persistence mode to `file`.

Procedure

To repackage an archive file using Ant:

Write the code that creates a new XU (execution unit) RAR file and a new management EAR file that set file persistence:

```
ant -Dxu.in=../applicationserver/JBoss6EAP/jrules-res-xu-JBoss61EAP.rar
-Dxu.out=myxu.rar
-Dconsole.ear.in=../applicationserver/JBoss6EAP/jrules-res-management-JBoss61EAP.ear
-Dconsole.ear.out=mymanagement.ear
-Dpersistence.type=file -f ressetup.xml setup
```

Configuring Rule Execution Server for different environments

It is very likely that the development of your BRMS requires more than a single deployment of Rule Execution Server.

The development lifecycle of a business rule application is similar to any other software development process including stages for implementation, testing, deployment and maintenance. At the very least, you are likely to need an environment for your development team, one for your QA team, and another one for in-production applications. In cases where you configure Rule Execution Server in a single cell, it is good practice to isolate the rulesets that you use on each server, and ensure that the execution units (XUs) do not interfere with each other.

Use the following instructions to set up your different environments in a single cell.

To set up a data source for each environment:

Use a unique JNDI name. For example:

- `jdbc/resdatasourceEnv1`
- `jdbc/resdatasourceEnv2`

•

To configure a XU for each environment and define a J2C connection factory:

1. Open the `<ODM_InstallDir>executionserver/applicationservers/JBoss6EAP/jrules-res-xu-JBOSS61EAP.ear` archive with a decompression tool.
2. In the `jrules-res-xu-JBOSS61EAP.ear /Meta-Inf/ironjacamar.xml` file, make the following highlighted change: `<connection-definition class-name="ilog.rules.res.xu.spi.IlrManagedXUConnectionFactory" enabled="true" jndi-name="java:/eis/XUConnectionFactoryEnv1" pool-name="eis/XUConnectionFactory" use-ccm="true" use-java-context="true">`. Execution components that invoke this XU must be modified and use this JNDI name instead of the default.
3. Save the file.
4. Repeat the same process for XUs in other environments.

To deploy the Rule Execution Server console for each environment:

1. Modify the deployment descriptor of the Rule Execution Server console EAR. In the `web.xml`, uncomment the `JMX_XU_QUERY_PART` parameter and specify `xuName=xuEnv1`.
2. Deploy the Rule Execution Server console EAR on the server.
3. In the resource reference settings in the application server:
 - a. Set the JNDI name for the data source as: `jdbc/resdatasourceEnv1`.
 - b. Set the JNDI name for the XU as: `eis/XUConnectionFactoryEnv1`.
4. Repeat the process to deploy the Rule Execution Server console for the other environments.

If you deploy to a cluster, synchronize your changes across the cluster after you complete the configuration.

Invoke the XU instances to register the XU with the Rule Execution Server console.

Rule Execution Server deployment for high availability and scalability

Use Rule Execution Server in a cluster or other multi-server environments.

Cluster configuration and topology:

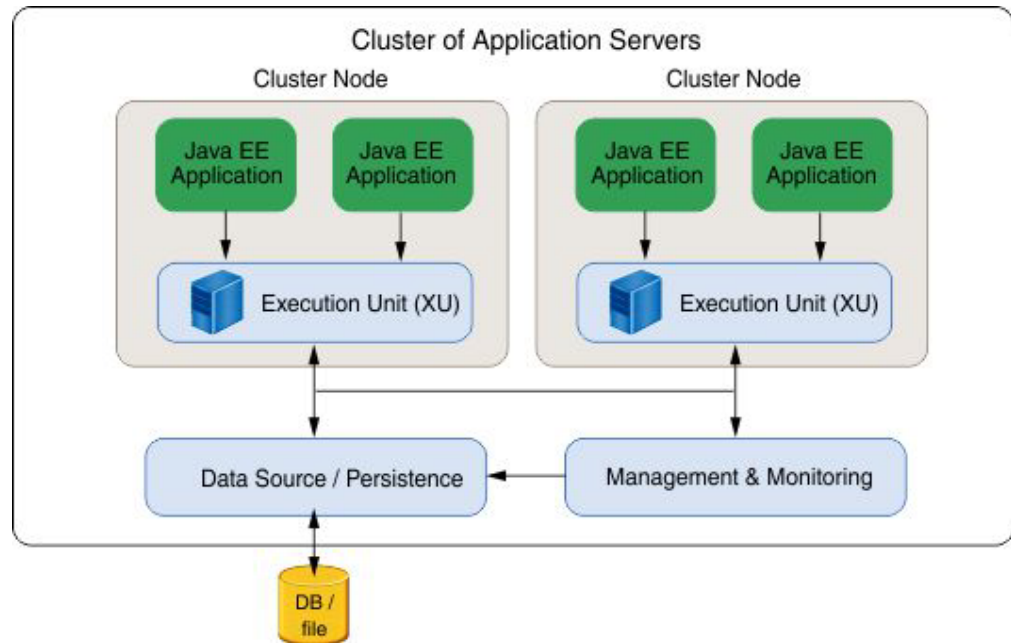
When you deploy Rule Execution Server onto server clusters, each node hosts one execution unit (XU), which is used only by a local rule session. The cluster topology significantly affects the notification mechanism.

Within the Java EE framework, clusters provide mission-critical services to ensure minimal downtime and maximum scalability. A cluster is a group of application servers that transparently run your Java EE application as if it were a single entity.

Cluster implementations on Java EE application servers come with their own set of terminology. Each of the following terms are important to understand how your cluster performs:

- Cluster and component failover services
- HTTP session failover
- Single points of failure in a cluster topology

On a cluster configuration, deploy an execution unit (XU) on each node. There is one XU for each node of a cluster. Use the administration console of application servers to handle cluster deployment. A XU instance can be used only by a local (same node) rule session. The rule session and the XU communicate through direct Java method calls, so the XU does not require serialization.



A cluster that uses Rule Execution Server involves a collaboration between the Rule Execution Server MBeans. The topology of the cluster has significant influence on the management of the notification mechanism when a resource is changed.

The management model is likely to use several times a basic scenario of a distributed notification mechanism within a cluster to interact with the various execution unit (XU) instances. A XU message-driven rule bean (MBean) is deployed with the XU to collaborate with the Rule Execution Server JMX infrastructure.

The following sequence applies:

1. A management client sets a resource on a ruleset MBean.
2. The ruleset makes a query to the MBean server to retrieve all the XU instances in the cluster. This operation requires a specific implementation for each application server.
3. A notification is sent to each instance.

Deploying Rule Execution Server in a cluster:

When you deploy the different Rule Execution Server components (console and runtime) on a JBoss application server, it is important to activate a single instance of the console in a JBoss cluster.

About this task

There is no synchronization between deployed Rule Execution Server consoles.

The Rule Execution Server console automatically discovers all the XUs deployed in the cluster before any ruleset has been executed.

The notification mechanism works by transmitting messages to all the XU instances on cluster members. The management model finds the list of members in a cluster with the TCP/IP Notification protocol, then uses the TCP/IP remote connection on each cluster member to notify all XU instances.

Procedure

To deploy the Rule Execution Server in a cluster.

1. Configure the Rule Execution Server console for TCP/IP management. See https://infocenter.ilogfr-bso.fr.ibm.com/infocenter/dmanager/trunk/topic/com.ibm.odm.dserver.rules.res.managing/topics/tsk_res_rescons_config_tcpip.html
2. Configure the execution units (XU) to connect to a TCP/IP management server. See https://infocenter.ilogfr-bso.fr.ibm.com/infocenter/dmanager/trunk/topic/com.ibm.odm.dserver.rules.res.managing/topics/tsk_res_config_xu_tcpip.html
The **Hostname** parameter must refer to the host name of the server where the console is deployed.
3. Deploy the Rule Execution Server console modified in step 1 and the XU in a stand-alone instance. See steps 4, 6, 7
4. Deploy the XU with the TCP/IP notification activated in a clustered JBoss domain. Choose the targeted domain and deploy the XU with the TCP/IP activated.
5. Verify the configuration.
 - a. To open the Rule Execution Server console, type `res` at the root URL on the host machine: `http://localhost:8080/res`. If your browser is not running on the same host as the application server, replace `localhost` with the address of the machine. If the web application is mapped to a host that is defined on a different port to 8080, change the port number from 8080 to the host port number.
 - b. Sign in to the Rule Execution Server console.
 - c. Click the **Diagnostics** tab.
 - d. Click **Run Diagnostics**.
 - e. Go to the **Server Info** tab. You should see the XUs attached to the Rule Execution Server console.

Example

You can deploy the XU and Rule Execution Server console configured with TCP/IP notification activated in the Group Member named Others. Proceed as follows:

1. Go to `<JBoss_HOME>/bin`.
2. Enter `./domain.[sh|bat]`
3. Open the JBoss Management Console.
4. Start the **Others Domain Group**.
5. Click the **Profiles** tab.
6. Configure the Rule Execution Server data source for the full and full-ha profiles.
7. Click the **Runtimes** tab.
8. Click **Management Deployments**.

9. Deploy the XU with TCP/IP activated.
 - a. Click **Add**.
 - b. Browse to the XU .rar file that has TCP/IP activated.
 - c. Click **Next** and then **Save**.
10. Deploy the Rule Execution Server console with TCP/IP activated. .
 - a. Click **Add** and then **Add** again.
 - b. Browse to the Rule Execution Server console EAR file with TCP/IP activated.
11. Assign the jrules-res-xu-JBoss61EAP.rar XU to the **Other Server Group**.
12. Assign the jrules-res-xu-JBoss61EAP.rar XU to the **Main Server Group**.
13. Assign the jrules-res-management-JBoss61EAP.ear XU to the **Other Server Group**.
14. Verify the configuration.

Configuring the Decision Center consoles on JBoss 6.1 EAP

To use the Decision Center Enterprise console on a new instance of the JBoss application server, you must deploy the provided archive for this server and perform a number of configuration tasks.

Before you start

Review the steps that you must perform to complete the configuration.

These instructions are intended for users familiar with JBoss and the database they are using. Examples using the Derby embedded database are included throughout to help beginners.

Refer to “What steps to follow” to understand the steps that you must follow to configure Decision Center on your application server.

You can complete the configuration of your application server with the Installation Settings wizard in the Decision Center console or with Ant scripts.


If you want to configure testing and simulation, see “Additional steps to configure testing and simulation” on page 41.






After finishing the configuration, you can use Decision Center (see Opening the Decision Center consoles). There is no rule project the first time you open the console. You must publish a project from Rule Designer (see Publishing a project to Decision Center).

Note: If you have rule projects created with a previous product version, refer to the Migrating topics on how to upgrade the Decision Center database schema.

What steps to follow

The steps that you follow to configure Decision Center on JBoss 5 and 6.1 are summarized in the following table:

Step	Required
“Step 1: Restricting database user permissions” on page 20	

Step		Required
"Step 2: Creating a data source and connection pool" on page 21		
"Step 3: Configuring security on JBoss" on page 22	"Defining groups and permissions" on page 22	
	"Declaring custom groups" on page 23	Optional
"Step 4: Deploying the Decision Center EAR on JBoss" on page 24		
"Step 5: Verifying the deployment of the Decision Center consoles" on page 25		Recommended
"Step 6: Completing your Decision Center configuration" on page 25	"Completing the configuration from the Decision Center Enterprise console" on page 25	
	"Completing the configuration by using Ant tasks" on page 30	Use one of the two methods to complete the installation. Required for distributed platforms only.
Verifying your configuration of Decision Center		

Step 1: Restricting database user permissions

Decision Center data is stored in a database. The database administrator might require that you provide the specific permissions that you need when accessing this database.

Note: This step applies when database access needs to be restricted. If you manage the database yourself (for example, you use an embedded database for test purposes) or if you do not need further restrictions, skip this step and proceed to the next configuration step.

Connection to the Decision Center database, established in the data source credentials, and any subsequent requests to the database are handled through a database user. This database user (name and password), for example `rtsdbUser`, is defined by the database administrator and has no relation to the standard Decision Center groups.

The following table gives the permissions that the database administrator must define on the Decision Center database, with attention given to the type of operations that you want to perform. Some supported databases do not require all the above permissions.

Database permission	Operation			
	Use Decision Center	Create the database schema by using the Decision Center console or Ant tasks	Modify the database schema by using the Decision Center console or Ant tasks	Migrate the database schema
ALTER TABLE	Not required	Not required	Required	Required
CREATE INDEX	Not required	Required	Required	Required
CREATE ROLE	Not required	Not required	Not required	Required
CREATE SEQUENCE	Not required	Required	Required	Required
CREATE TABLE	Not required	Required	Required	Required
CREATE VIEW	Not required	Required	Required	Required
DROP INDEX	Not required	Not required	Required	Required
DROP SEQUENCE	Not required	Not required	Required	Required
DROP TABLE	Not required	Not required	Required	Required
DROP VIEW	Not required	Not required	Required	Required
INSERT TABLE	Required	Required	Required	Required
SELECT SEQUENCE	Required	Required	Required	Required
SELECT TABLE	Required	Required	Required	Required
UPDATE TABLE	Required	Required	Required	Required

Step 2: Creating a data source and connection pool

You declare a data source to store the data that is used by Decision Center.

Before you begin

You must have a database to create a data source.

Note: Beginners who use the Derby embedded database can create the database as part of this procedure.

For a list of supported databases, see the IBM Support site.

Procedure

1. Register the JDBC driver as a core module.
 - a. Under the `<JBoss_HOME>/modules` directory, create a file path structure. For example, for a Derby driver, create the following structure:
`<JBoss_HOME>/modules/org/apache/derby/main`
 - b. Copy the JDBC driver `.jar` file into the main subdirectory.
 - c. In the main subdirectory, create a `module.xml` file. Example for Derby:

```
<module xmlns="urn:jboss:module:1.0" name="org.apache.derby">
  <resources>
    <resource-root path="derby.jar"/>
    <!-- Insert resources here -->
  </resources>
```

```

    <dependencies>
      <module name="javax.api"/>
    </dependencies>
  </module>

```

- d. Launch the JBoss server and run the `jboss-cli` command to register the new module in the server. The command is available in `<JBOSS_HOME>/bin`. It must be executed with the name of the new module as follows:


```
./jboss-cli[.sh|.bat] --connect "/subsystem=datasources/jdbc-driver=derby:add(driver-name=derby,driver-module-name=org.apache.derby)"
```
2. Create the data source.
 - a. Log in to the JBoss Management Console.
 - b. Click the **Profile** tab in the top right of the console.
 - c. From the tree on the left, go to **Subsystems > Connector > Datasources**
 - d. Click the **Datasources** panel.
 - e. Click **Add**, and then enter a name and a JNDI name (`java:/jdbc/i1ogDataSource`) as the data source attributes.
 - f. Select the JDBC driver to be used for the new data source. The JDBC driver registered as a core module in step 1 should appear here.
3. Specify the connection settings.

Connection URL

This setting defines the JDBC connection URL. For example:
`jdbc:derby:c:/rtsdb;create=true`

User name and Password

The tags `username="rtsdbUser"` and `password="rtsdbUser"` are used to access the database.

Tip: If the transaction isolation level is not set to `READ-COMMITTED`, you might encounter database access problems. Configure data source isolation to provide a better concurrency experience.

4. Enable the new data source. After the data source is created, it appears in the Available Datasources table. By default, the new data source is not enabled.
 - a. Select the new data source in the table and then click **Enable**.
 - b. Confirm the operation. The data source is listed as **Enabled**.

Step 3: Configuring security on JBoss

You control access to Decision Center and enforce security by defining user groups and associated roles.

Defining groups and permissions

You control access to Decision Center and enforce security by defining user groups and associated roles.

About this task

Any user of Decision Center must belong to at least one of these mandatory groups:

- `rtsAdministrator`
- `rtsConfigManager`
- `rtsInstaller`

- rtsUser

Adherence to these groups determines what parts of Decision Center a user can access. You must create all these groups. For testing purposes, also create a default user and password for each of these groups. In addition, if you want to follow the Decision Center permissions tutorial in your own installation, you must create the two Validator and Eligibility custom groups.

The following table summarizes the mandatory and custom groups and their associated role, default user, and password.

Group	Use	Default User - Password
rtsAdministrator	Mandatory, gives the user administrator access.	rtsAdmin - rtsAdmin+0
rtsConfigManager	Mandatory, gives the user configuration manager access.	rtsConfig - rtsConfig+0
rtsUser	Mandatory, gives a user standard access.	rtsUser1 - rtsUser1+0
rtsInstaller	Mandatory, gives the user access to the Installation Settings wizard.	rtsAdmin - rtsAdmin+0
Validator	Optional custom group, used in the Decision Center permissions tutorial.	Val - ValVal+0
Eligibility	Optional custom group, used in the Decision Center permissions tutorial.	Eli - EliEli+0

Note: For more information on the Decision Center groups, see Groups.

You declare groups and users by using the JBoss **add-user** command.

The following procedure is a suggested configuration with a default user for each of the basic groups.

Procedure

In the `<JBOSS_HOME>/bin` directory, open a command prompt and run the following **add-user** commands:

```
./add-user[.sh|.bat] -a -u rtsAdmin -p rtsAdmin+0 -ro "rtsAdministrator,rtsInstaller,rtsUser"
./add-user[.sh|.bat] -a -u rtsConfig -p rtsConfig+0 -ro "rtsConfigManager,rtsUser"
./add-user[.sh|.bat] -a -u rtsUser1 -p rtsUser1+0 -ro "rtsUser"
./add-user[.sh|.bat] -a -u Val -p ValVal+0 -ro "Validator,rtsUser"
./add-user[.sh|.bat] -a -u Eli -p EliEli+0 -ro "Eligibility,rtsUser"
```

What to do next

After you deploy the archive in the next step, you must upload all groups to the database from the Installation Settings wizard of the Decision Center console. For more information, see “Step 3: Set up groups” on page 28.

Declaring custom groups

If you create custom groups, you must declare them before you deploy the EAR file.

About this task

The Decision Center EAR file references the basic groups: `rtsUser`, `rtsConfigManager`, `rtsAdministrator`, and `rtsInstaller`.

You must add any custom groups that you declared in “Defining groups and permissions” on page 22. Also, add the Validator and Eligibility groups that you created for the Decision Center tutorials.

Tip:

- The EAR file is compressed. You must open it to extract the files that must be changed, and then replace the files in the EAR file. You can use Ant commands to repackage the EAR file, as explained in “Repackaging the Decision Center archive” on page 39.
- Back up the EAR file before you modify it.
- To use the Decision Center permissions mechanism, you must upload groups to the database. For more information, see “Completing the configuration from the Decision Center Enterprise console” on page 25.

Procedure

1. Add each custom group as a role in the appropriate `application.xml` file.
`<ODM_InstallDir>/teamserver/applicationservers/JBoss6EAP/jrules-teamserver-JBOSS61EAP.ear/META-INF/` For example:
...

```
<security-role>  
  <role-name>my_custom_group</role-name>  
</security-role>
```


...
 2. If you have already deployed the archive, upload the custom groups to the database from the Installation Settings wizard of the Decision Center console. For more information, see “Step 3: Set up groups” on page 28.

What to do next

Otherwise, proceed to “Step 4: Deploying the Decision Center EAR on JBoss” and upload the groups when you complete the configuration.

Step 4: Deploying the Decision Center EAR on JBoss

To deploy the Decision Center EAR on JBoss, you must copy the EAR archive to the deployment directory of your application server.

Procedure

Copy the file `<ODM_InstallDir>/teamserver/applicationservers/JBoss6EAP/jrules-teamserver-JBOSS61EAP.ear` to the deployments folder:
`<JBOSS_HOME>/standalone/deployments`

Important:

Deploying the Decision Center EAR sets the persistence locale. After you save a rule to the database, you cannot change the persistence locale. If you want to install Decision Center in a language other than English, take note of the instructions provided in “Step 6: Completing your Decision Center configuration” on page 25.

Step 5: Verifying the deployment of the Decision Center consoles

After you have finished configuring Decision Center for your application server, verify that you have deployed the archive successfully.

About this task

You start your application server and then use your web browser to open the Decision Center Enterprise console.

Procedure

1. Make sure that your application server is running.
2. Start a new browser instance and enter the default URL to access Decision Center in a web browser.
 - The default URL to access Decision Center is `http://localhost:<PORT_NUMBER>/teamserver`. Set `<PORT_NUMBER>` to the port number of your web application. The default value of the port number is 8080. If your browser is not running on the same host as the application server, replace `localhost` with the address of the host.
 - Starting from WebSphere® Application Server V8.0, each console has a different default URL.
 - For the Enterprise Console: `http://localhost:<PORT_NUMBER>/teamserver`
 - For the Business Console: `http://localhost:<PORT_NUMBER>/decisioncenter`
 - If your browser is not running on the same host as the application server, replace `localhost` with the address of the host.
 - If your web application is mapped to a host on a port that is different from the port number shown, change the port number to your host port number.The Decision Center log in page opens in your browser.
3. Sign in with `rtsAdministrator` rights to start testing.

For example, `rtsAdmin` and `rtsAdmin+0`.

Step 6: Completing your Decision Center configuration

After you have created a data source, defined security settings, and deployed the Decision Center EAR file, you complete the configuration either from the Decision Center console or by running Ant tasks. If your database is DB2 for z/OS®, generate the schema on z/OS for reasons of performance. You can, however, use the wizard to import extension files.

Completing the configuration from the Decision Center Enterprise console

After you have deployed the Decision Center archive to your application server, you can work from the Decision Center Enterprise console to complete or modify the configuration.

Installation Settings wizard overview:

You use the Installation Settings wizard in the Decision Center console to create or modify the database schema, set up message files or groups, or change the persistence locale or configuration parameters.

The Installation Settings wizard opens automatically when you start the Decision Center console to complete an installation.

You can also open the Installation Settings wizard by clicking **Configure** > **Installation Settings Wizard** in the Decision Center console after you have completed your initial installation. If you open Decision Center after you have followed the steps to install the module, only the **Install** tab is available. For more information, see *Opening the Decision Center consoles*.

Note: To access the Installation Settings wizard, you must have both administrator privileges and the `rtsInstaller` role when you sign in.

You use the Installation Settings wizard for the following actions.

Table 1. Actions in the Installation Settings wizard

Action	Description
Configure the database	Mandatory when you complete the configuration with a database on a distributed platform. For more information, see Step 1: Configure the database.
Set up message files	Mandatory during the installation only if you have some custom rule model extension files. For more information, see Step 2: Set up message files.
Set up groups	You must set up the same groups that are declared in the application server if you want to use the Decision Center security and permissions mechanisms. For more information, see Step 3: Set up groups.
Change the persistence locale	Mandatory if the persistence locale is different from the locale <code>en_US</code> . For more information, see Step 4: Set the persistence locale.
Change configuration parameters	Optional. You change some configuration parameters when you customize Decision Center. For more information, see Step 5: Set configuration parameters.

After you have completed the installation, Decision Center is ready to use but does not contain rule projects. If you open Decision Center and no rule projects are available, a message in the **Configure** tab informs you that no project has been found and that you should either publish a rule project by using Rule Designer or contact the administrator.

If you see this message, you must publish a rule project from Rule Designer. For more information, see *Publishing a project to Decision Center*.

More information about using the Installation Settings wizard is available from the Decision Center console online help. To access the online help, click **Help** in the top banner after you have signed in to Decision Center.

Note: If you have rule projects that were created with a previous product version, refer to the *Migrating* topics on how to upgrade the Decision Center database schema.

Step 1: Configure the database:

You use the Installation Settings wizard to configure the database.

About this task

You store the extensions to the Decision Center rule model in two XML files:

- Model description: This file usually has the file name extension `.brmx`.
- Initialization of enumerations and hierarchies: This file usually has file name extension `.brdx`.

For more information about defining common model extensions, see the customization topics.

Procedure

1. When the Installation Settings wizard opens in Decision Center, click **Next**.
2. Select one of the extension files.
 - **Default extensions** (already selected)
 - **Custom extensions (brmx/brdx)**, or
 - **Custom extensions (Zip)**
3. Click **Generate SQL** to generate the script that creates the database tables, which are based on the contents of your rule model files.
4. After the script is generated, select the **Execute the SQL script** check box, and then click **Next**.

Step 2: Set up message files:

Message files contain the display text that is associated with the extensions to the rule model contained in the `.brmx` and `.brdx` files.

About this task

You can find the default message file in `<ODM_InstallDir>/teamserver/bin/defaultextensionmessages_<LOCALE>.properties`.

If you use the default rule model when you create your database, the default message file is automatically sent to the database. To upload your own message files, use the Installation Settings wizard as explained below.

You must have a message file for each locale that you use. Message files are identified by their locale. The contents of the message files must respect the ISO-LATIN-1 standard.

Procedure

To declare a message file in the Installation Settings wizard:

1. Click **New**.
2. Enter a locale.
3. Browse to the location of the message file for this locale.
4. Click **Apply**.

Results

If Decision Center supports this locale, the Installation Settings wizard assigns a locale code so that you can identify it.

Example

For example:

```
status=Status
effectiveDate=Effective Date
expirationDate=Expiration Date
new=New
defined=Defined
```

Step 3: Set up groups:

In addition to creating groups in your application server when you set up security access, you must use the Setup Groups page in the Installation Settings wizard to upload groups to the database.

Before you begin

Before you set up groups in the Enterprise console, make sure to add all the groups that you want to see in the available list when you enforce project security or set permissions in Decision Center. For more information, refer to the topics on Groups and Permissions in the Decision Center help.

About this task

You set up groups only if you want to use the Decision Center project access and permission mechanisms.

Tip: In Decision Center, the groups are the roles in the application server, **not** the groups defined in the user registry. Decision Center uses the group information to verify whether a user belongs to a role in the application server.

You do not have to upload the `rtsAdministrator` or `rtsInstaller` group. The administrator group has access to everything, and an installer user must belong to another group.

You use the Setup Groups page in the Installation Settings wizard to upload the default groups for `rtsUser` and `rtsConfigManager`, and any custom groups, such as `Validator` and `Eligibility` if you want to follow the permissions tutorial.

Procedure

To set up groups:

1. Click **New**.
2. Type the group name.
3. Click **Apply**.
4. Repeat steps 1 to 3 for each group.
5. When you have added all the groups, proceed in one of the following ways:
 - Click **Next** if you want to set a different persistence locale, or configuration parameters.
 - Click **Finish** if you do not want to change these settings.

Step 4: Set the persistence locale:

The persistence locale determines the language in which you store rules in the Decision Center database.

About this task

You set the locale when you deploy the Decision Center EAR or WAR file to your application server. As a consequence, you store the rules in the database in the locale of the Decision Center application.

Changing the persistence locale does not change the language in which Decision Center displays rules. Changing it in Decision Center is necessary only to match the locale of Rule Designer when you synchronize your rule projects, and to access the tutorials in your locale.

Important: You must not change the persistence locale after you have saved a rule to the database.

Procedure

To set the persistence locale:

1. Enter a locale in the **Locale** field.
2. Click **Apply**.
3. Proceed as follows:
 - Click **Next** if you want to set the configuration parameters.
 - Click **Finish** if you do not want to change these settings.

Step 5: Set configuration parameters:

Many tasks that are related to customizing Decision Center require that you add or remove configuration parameters.

About this task

Decision Center uses the following configuration parameters to generate complete URLs in permalinks:

- **teamserver.server.port**: the port number
- **teamserver.server.isSecure**: true if the connection is secure
- **teamserver.server.hostname**: the name of the host.

To create, modify, or delete configuration parameters, you use the Set configuration parameters page in the Installation Manager wizard. You generate these parameters when you sign in to the Decision Center console for the first time after you have configured the database. You can use the Installation Settings wizard to change these parameters at any time.

The following table gives a description of the main configuration parameters available in `teamserver.war/WEB-INF/lib/teamserver-model-XXX.jar/ilog/rules/teamserver/preferences.properties`.

Note:

The parameters in the table include the **teamserver** prefix, which is not in the `preferences.properties` file. You must include the prefix when you set configuration parameters in the Installation Settings wizard.

Parameter	Used to
<code>teamserver.<extractorValidator>.class</code>	Specify a class of ruleset extractor validators to use for the <code>extractorValidator</code> name. The class must implement the <code>ILrExtractorValidator</code> interface. After you define this class, specify this name as the extractor validator to use when defining a ruleset extractor.
<code>teamserver.build.path</code>	Define the location of the IRL cache in the file system. Compute the path as follows: <ul style="list-style-type: none"> • Use this property with the name of the user who started the server as the root for the cache (<code><build.path>_<username></code>). • If this property is not defined, use the system property <code>java.io.tmpdir</code> and add <code>rtscache</code>. For example, <code><temp dir>/rtscache_<username></code>. • If the system property is not defined, use the server directory and add <code>rtscache</code>. For example, <code><server dir>/rtscache_<username></code>.
<code>teamserver.br1.verbalizers</code>	Specify the list of locales for which a BAL verbalizer is defined.
<code>teamserver.br1.verbalizer.<locale></code>	Specify the verbalizer class for the locale. The class must implement the <code>ilog.rules.vocabulary.verbalization.ILrVerbalizer</code> interface.

Procedure

1. Create a parameter or change an existing one.
 - To create a parameter, click **New**.
 - To change a parameter, select the check box next to the parameter and then click **Modify** to change the parameter or click **Delete** to remove it.
2. Click **Apply** to implement your changes.
3. Proceed as follows:
 - Click **Previous** if you want to make changes to previous settings.
 - Click **Finish**. The Installation log opens with a summary of the operations that you performed in the Installation Settings wizard.
4. Click **OK** to finish.

What to do next

You now have to sign in to the Decision Center console. Continue with the section Publishing a project to Decision Center.

Completing the configuration by using Ant tasks

Ant tasks provide an alternative method for completing or modifying the configuration. These tasks perform the same configuration steps as the Installation Settings wizard in the Decision Center console.

Setting up the Ant tasks environment:

To run Decision Center Ant tasks, you must first set up the required environment variables.

To run Ant tasks, you must have version 1.7.1 (or later) of Ant set up on your system. If Ant is not installed or your version is older than version 1.7.1, you must set up your environment to use the correct version of Ant. You can download Ant from the Apache web site, or you can use the Ant 1.7.1 distribution packaged at `<ODM_InstallDir>/shared/tools/ant`, where `<ODM_InstallDir>` is your Operational Decision Manager installation directory.

To test your current version of Ant:

Type the following command in a Windows command prompt or UNIX shell: `ant -version`

To set up your environment to use Ant:

1. Set the `ANT_HOME` environment variable to `<ODM_InstallDir>/shared/tools/ant`.
2. Set the `JAVA_HOME` environment variable to the path to your JDK installation (1.6).
3. Add the directory `<ODM_InstallDir>/shared/tools/ant/bin` to your `PATH` environment variable.

The Decision Center Ant tasks are defined in `<ODM_InstallDir>/teamserver/bin/build.xml` and executed by commands of the form:

```
ant <taskName> <parameters list>
```

Note: To execute these Ant tasks, you must use the same Java virtual machine version and vendor as the one that is used by the application server.

Ant task parameters start with `-D`. Use them to set values such as the following ones:

- `-Dserver.url=<server url>`: Specifies the URL of the target application server.
- `-DdatasourceName=<data source name>`: Specifies the JNDI name of the data source to use for the task. The default value is `jdbc/ilogDataSource`.

```
ant execute-schema -Dserver.url=<protocol://host:port>/teamserver/  
-DdatasourceName=jdbc/ilogDataSource -Dfile=my_sql_file.sql
```

The `<protocol://host:port>` URL is defined in the file `<ODM_InstallDir>/teamserver/bin/teamserver-anttasks.properties`. If your browser is not running on the same host as the application server, replace `localhost` with the address of the machine. If your web application is mapped to a host on a port that is different from the port number shown, change the port number to your host port number.

The file `<ODM_InstallDir>/teamserver/bin/teamserver-anttasks.properties` defines the value of some common parameters and others that depend on the application server used. You do not have to include these parameters in your Ant task command if they are properly defined in this file. The content of the `teamserver-anttasks.properties` file is as follows:

```
# Default properties  
# -----  
rtsAdmin.login=rtsAdmin  
rtsAdmin.password=rtsAdmin  
  
protocol=http  
server.host=localhost  
server.port=8080  
server.url=${protocol}://${server.host}:${server.port}/teamserver
```

```
datasourceName=jdbc/ilogDataSource
outputFile=output.sql

languagePackPath = .
languagePackOutputPath = ./generated

persistenceLocale =
selector =
branch =
override = false
```

Also, take note of any special instructions in this file concerning your application server.

The **appserver.name** property configures the class path for the Ant tasks. If you have to add specific drivers to your class path, you can add them to `<ODM_InstallDir>/teamserver/lib/classpath-teamserver.xml`.

Communication between the Ant tasks and Decision Center supports the HTTP or HTTPS communication protocols. For more information, see Ant task communication protocol.

Creating the database schema:

You can create the database schema in a single operation by using the **set-extensions** Ant task, or choose to create it step by step.

Creating schemas with the set-extensions Ant task:

For convenience, you can create the database schema by using the **set-extensions** Ant task.

About this task

Extensions to the Decision Center rule model are stored in two XML files.

- Model description: This file usually has the file name extension `.brmx`.
- Initialization of enumerations and hierarchies: This file usually has file name extension `.brdx`.

You can use Ant tasks to load the rule model from the two XML files and build the SQL script that is necessary to get the proper database schema.

Note: To run these Ant tasks, you must use the same Java virtual machine version and vendor as the one used by the application server.

Alternatively, you can create the database schema step by step, which is useful if you want to look at the generated SQL schema. See “Creating the schema using a step-by-step sequence” on page 33.

Procedure

Run the **set-extensions** Ant task.

This Ant target runs **gen-create-schema + execute-schema + upload-extensions + upload-roles**. Set the following parameters:

```
-Dserver.url=<server url>
-DdatasourceName=<data source name>
```

-DextensionModel=<model file>

The model description file, with the .brmx extension

-DextensionData=<data file>

The model data description, with .brdx extension

[-DdbName=<database schema name>]

You can use this optional parameter to specify the database schema name. Otherwise, Decision Center uses the database user name as the schema name. However, some databases allow some users to access several schemas, and the default schema does not always reflect the user name.

[-Droles=<role list>]

You can use this optional parameter to upload the list of roles to Decision Center. This list is specified as "role1 role2". For example:

```
ant upload-roles -Droles="rtsUser rtsConfigManager Eligibility Validator"
```

Creating the schema using a step-by-step sequence:

If you want to look at the generated SQL schema, you can create it step-by-step.

Creating the database schema script:

You can create the database schema script by using the **gen-create-schema** Ant task.

Procedure

To create the SQL script that is necessary to create or update the database schema, run the **gen-create-schema** Ant task with these parameters:

-Dserver.url=<server url>

-DdatasourceName=<data source name>

-DextensionModel=<model file>

The model description (.brmx extension).

-DextensionData=<data file>

The model data description (.brdx extension).

[-DdbName=<database schema name>]

You can use this optional parameter to specify the database schema name in which the Decision Center tables are stored. If you do not specify the parameter, Decision Center uses the database user name as the schema name. However, some databases allow some users to access several schemas, and the default schema is not always named the same as the user.

[-DoutputFile=<SQL file>]

The name of the file that stores the generated SQL script. If this parameter is not specified, the task creates a file named output.sql in the directory that is defined as basedir in the build.xml file.

```
ant gen-create-schema -DextensionModel=my_model_file.brmx -DextensionData=my_data_file.brdx -Dout
```

Results

1. The task connects to the specified data source from the application server.
2. The task checks whether this data source points to an existing Decision Center database.
 - If a database does not exist, the task builds the SQL script to create a fresh database schema to store the model.

- If a database does exist, the task builds the SQL script that is necessary to update the existing database schema.

Executing the database schema script:

You execute the database schema script.

Procedure

To execute the SQL script that you created, run the **execute-schema** Ant task with these parameters:

-Dserver.url=<server url>

-DdatasourceName=<data source name>

[-Dfile=<SQL file>]

The name of the file to execute, which corresponds to the script that you created. If you do not specify this parameter, the task attempts to execute a file named `output.sql` in the directory that is defined as `basedir` in the `build.xml` file.

```
ant execute-schema -Dfile=my_sql_file.sql
```

Uploading the database schema extension:

You upload the database schema extension.

Procedure

To store the rule model description in the database schema, run the **upload-extensions** Ant task with these parameters:

-Dserver.url=<server url>

-DdatasourceName=<data source name>

-DextensionModel=<model file>

The model description (`.brmx` extension).

-DextensionData=<data file>

The model data description (`.brdx` extension).

```
ant upload-extensions -DextensionModel=my_model_file.brmx -DextensionData=my_data_file.brdx
```

Results

The description is stored in the database so that Decision Center applications can load it when they start. It is also used by **gen-create-schema** to get the current model description to run a `diff` with the new schema.

In a cluster, you must restart the servers and close all current sessions.

Uploading a list of roles or groups to the database:

In addition to creating groups in your application server when you set up security access, you must upload groups to the database.

Before you begin

1. Add all the groups that you want to see in the available list when you enforce project security or set permissions in Decision Center.

2. Create the default groups for `rtsUser` and `rtsConfigManager`, and upload your custom groups.

About this task

You must upload roles or groups only if you want to use the Decision Center project access and permissions mechanisms. For more information, see the topics on Groups and Permissions in the Decision Center online help.

You do not have to upload the `rtsAdministrator` group or the `rtsInstaller` group. The Administrator group has access to everything and an Installer user must belong to another group.

Note: To run these Ant tasks, you must use the same Java virtual machine version and vendor as the one used by the application server.

Procedure

To store in the database the list of roles or groups to be used by the application, run the **upload-roles** Ant task with the following parameters:

-Dserver.url=<server url>

-DdataSourceName=<data source name>

-Droles=<role list>

<role list> is the list of roles or groups to upload to Decision Center, specified as "group1 group2".

```
ant upload-roles -Droles="rtsUser rtsConfigManager Eligibility Validator"
```

Removing a database schema:

You can create an SQL script to remove (drop) a database schema by using the **gen-drop-schema** Ant task.

About this task

To remove a database schema, you proceed in two steps:

1. Create the SQL script that is necessary to remove the database schema.
2. Execute the SQL script that you created.

Procedure

1. To create the SQL script to delete a database schema, run the **gen-drop-schema** Ant task with the following parameters:

-Dserver.url=<server url>

-DdataSourceName=<data source name>

-DextensionModel=<model file>

The description of the database schema to remove.

[-DdbSchemaName=<database schema name>]

You can use an optional parameter to specify the database schema name. If you do not specify this parameter, Decision Center uses the database user name as the schema name. However, in some databases, users can access several schemas and the default schema is not always named as the user.

[-DoutputFile=<SQL file>]

The name of the file that stores the generated SQL script. If you do not specify this parameter, the task creates a file named `output.sql` in the directory that is defined as `basedir` in the `build.xml` file.

```
ant gen-drop-schema -DextensionModel=my_model_file.brmx -DoutputFile=my_sql_file.sql
```

2. To execute the SQL script that you created, run the **execute-schema** Ant task with these parameters:

-Dserver.url=<server url>

-DdatasourceName=<data source name>

[-Dfile=<SQL file>]

The name of the file to execute, which corresponds to the script that you created. If you do not specify this parameter, the task attempts to execute a file named `output.sql` in the directory that is defined as `basedir` in the `build.xml` file.

```
ant execute-schema -Dfile=my_sql_file.sql
```

The task connects to the specified data source from the application server. It reads the model description that is passed in the parameters, and generates the SQL script to remove the existing schema. Because many database tables are linked through foreign keys, these tables must be removed in a specific order and the script generation handles these constraints.

Example

Here is the complete code sample:

```
ant gen-drop-schema -DextensionModel=my_model_file.brmx -DoutputFile=my_sql_file.sql
ant execute-schema -Dfile=my_sql_file.sql
```

Defining and uploading message files:

You can define and upload message files to Decision Center by using the **upload-messages** Ant task.

Message files contain the display text that is associated with the extensions to the rule model that is contained in the `.brmx` and `.brdx` files. For example:

```
status=Status
effectiveDate=Effective Date
expirationDate=Expiration Date
new=New
defined=Defined
```

The default messages file is provided in: `<ODM_InstallDir>/teamserver/bin/defaulttextextensionmessages_<LOCALE>.properties`

Note: The contents of the messages files must conform to the ISO-LATIN-1 standard.

You must have a messages file for each locale that you use. Upload the messages file to Decision Center by running the **upload-messages** Ant task with these parameters:

- **-Dserver.url=<server url>**
- **-DdatasourceName=<data source name>**
- **-Dlocale=<locale>**
- **-DmessageFile=<message file>**


```
ant upload-messages -Dlocale=en_US -DmessageFile=mymessages.properties
```

Setting the persistence locale:

The persistence locale is used to determine the language in which rules are stored in the Decision Center database.

About this task

The persistence locale is set when you deploy the Decision Center archive to your application server, which means that the rules in the database are stored in the locale of the Decision Center application.

Changing the persistence locale does not change the language in which rules display in Decision Center. Changing the persistence locale in Decision Center is necessary only to match the locale of Rule Designer when you synchronize your rule projects, and to access the tutorials in your locale.

Important: You must not change the persistence locale after you have saved a rule to the database.

Procedure

To set the persistence locale by running an Ant task:

1. Open the `<ODM_InstallDir>/teamserver/bin/teamserver-anttasks.properties` file.

This file defines the value of some common parameters.

2. Add your locale to the **persistenceLocale** property and save the `teamserver-anttasks.properties` file.

For example: `persistenceLocale = fr_FR`

3. Run the Ant task in this form: `ant taskName parameters_list`

Alternatively, you can add the parameter to the command line. For example:

```
ant taskName -DpersistenceLocale=fr_FR
```

Adding or removing configuration parameters:

For many tasks that are related to customizing Decision Center, you must add or remove configuration parameters.

The following configuration parameters, used to generate complete URLs in permalinks, are generated the first time you sign in to Decision Center after the database is configured. You can use the Installation Settings wizard to set these parameters beforehand or change them afterward:

- **teamserver.server.port:** The port number
- **teamserver.server.isSecure:** true if the connection is secure.
- **teamserver.server.hostname:** The name of the host

The following table gives a description of the main configuration parameters available in `teamserver.war/WEB-INF/lib/teamserver-model-XXX.jar/ilog/rules/teamserver/preferences.properties`.

Parameter	Use
<code>teamsserver.<extractorValidator>.class</code>	Specify a ruleset extractor validator class to use for the <code>extractorValidator</code> name. The class must implement the <code>ILrExtractorValidator</code> interface. After this class is defined, specify this name as the extractor validator to use when defining a ruleset extractor.
<code>teamsserver.build.path</code>	Define where the cache of the IRL is located on the file system. The path is computed as follows: <ol style="list-style-type: none"> 1. First, use this property with the name of the user who started the server as the root for the cache: <code><build.path>_<username></code>. 2. If the path is not defined, use the system property <code>java.io.tmpdir</code> and add <code>rtscache</code>. For example, <code><temp_dir>/rtscache_<username></code>. 3. If the system property is not defined, use the server directory and add <code>rtscache</code>. For example, <code><server_dir>/rtscache_<username></code>.
<code>teamsserver.br1.verbalizers</code>	Specify the list of locales for which a BAL verbalizer is defined.
<code>teamsserver.br1.verbalizer.<locale></code>	Specify the verbalizer class for the specified locale. The class must implement the <code>ILrVerbalizer</code> interface.

You can use the following Ant tasks to add or remove configuration parameters:

set-config-param

Sets a configuration parameter for a specified user. If the user is not specified, it sets a *global parameter*.

Parameters:

- `-Dserver.url=<server url>`
- `-DdatasourceName=<data source name>`
- `[-Duser=<username>]`
- `-Dkey=<parameter key>`
- `-Dvalue=<parameter value>`

For example:

```
ant set-config-param -Dkey=locale -Dvalue=en_US
```

remove-config-param

Drops the configuration parameter for a specified user. If the user is not specified, it drops the global configuration parameter.

Parameters:

- `-Dserver.url=<server url>`
- `-DdatasourceName=<data source name>`
- `[-Duser=<username>]`
- `-Dkey=<parameter key>`

print-config-param

Prints the global parameters or specified user parameters if the *username* value is specified. If no key is specified, all keys are printed.

Parameters:

- **-Dserver.url=<server url>**
- **-DdataSourceName=<data source name>**
- **[-Duser=<username>]**
- **-Dkey=<parameter key>**

Repackaging the Decision Center archive:

You can repackage the Decision Center archive by using an Ant task.

When you add new .jar files to the Decision Center archive, you must repackage the archive by running the **repackage-ear** or **repackage-war** Ant task. This task does not use the **server.url** and **dataSourceName** parameters.

If you have customized Decision Center, you must package the custom .jar files before you use the Ant task to repackage the Decision Center.

The **repackage-ear** or **repackage-war** Ant task takes the following parameters:

-DtargetEar=<target ear>

-DtargetWar=<target war> for servers that require WAR files.

-DsourceEar=<source ear>

-DsourceWar=<source war> for servers that require WAR files.

-DdescriptorsDir=<descriptors directory>

A directory that is copied into the META-INF directory of the target EAR (not mandatory).

-DadditionalJars=<"myjar1.jar,myjar2.jar, myjarn.jar">

Additional .jar files to store in the lib directory of the target archive (not mandatory).

-DtmpDir=<directory>

A directory that you can specify to store temporary files (not mandatory).

-DwebResourcesDir=<web resources directory>

A directory that is copied into the WAR library (not mandatory).

-Dconsole=both|enterprise|business

Specifies whether to repackage the Business or the Enterprise WAR files. The default is both.

Configuring the search function of the Decision Center Business console

The search function in the Decision Center Business console is based on the Solr search engine. You must configure the engine to provide this functionality.

Using a remote Solr search engine

You can set up the search function in Decision Center Business console to work with a remote instance of the Apache Solr search engine.

About this task

The search function in Decision Center Business console uses an embedded instance of the Apache Solr search engine. Alternatively, you can have the search function work with a remote instance of the search engine, which you can run on another computer, or the same computer but in a dedicated web application.

Procedure

To configure the search function to run with a remote instance of the Solr search engine:

1. Install the Apache Solr search engine on another computer, or as part of a dedicated web application on your computer.
For information on installing the Solr server, visit the Apache Solr website.
2. Locate the `decisioncenter-solr-home.zip` file in the Decision Center `teamserver` folder on your computer.
The Decision Center installation program placed the folder on your computer.
3. Decompress the file in a directory in the remote instance of the Solr server.
4. Configure the home directory of the remote Solr server to use the location of the decompressed configuration files.
For more information, see the documentation on the Apache Solr website.
5. Configure the Decision Center `preferences.properties` file to point to the URL of the Solr server.

Setting parameters for the Solr search engine

The search function in the Decision Center Business console runs on Apache Solr.

By default, the search uses an embedded instance of the engine. Alternatively, you can run the search on a remote instance of the engine on either another computer or the same computer but in a dedicated web application. For more information, see “Using a remote Solr search engine” on page 39.

You can set three parameters for the search engine:

Table 2. Search parameters

Property	Description
SearchProvider	This parameter takes one of the following values: <ul style="list-style-type: none">• <code>SolrEmbedded</code>: Use this value to select the embedded Solr search engine.• <code>SolrRemote</code>: Use this value to select a remote instance of the Solr search engine.
SolrEmbeddedDataDir	Use this optional parameter to direct the index of the embedded version (<code>SolrEmbedded</code>) to a specific directory on the Decision Center.
SolrRemoteUrl	Use this parameter with <code>SolrRemote</code> to provide the URL of the remote Solr search engine.

You set the configuration parameters for the search engine in the `preferences.properties` file for Decision Center. The `preferences.properties` file

can be placed in any package, and the names of the search properties depend on the location of the file, for example with SearchProvider:

- If you place the preferences.properties file in ../rules/decisioncenter/web/search/, the name of the property must be SearchProvider.
- If you place the preferences.properties file in ../rules/decisioncenter/web/, the name of the property must be search.SearchProvider.

See also `IlrPropertyManager`

The following table provides examples for setting the parameters in preferences.properties to work with the Solr search engine:

Table 3. Configuration table for search properties

Solr server	Parameter settings
External server The URL of the external server depends on the installation. For this example, the remote address of the search engine is http://mysearchserver:8983/solr.	Set the preferences.properties file as follows: SearchProvider=SolrRemote SolrRemoteUrl=http://mysearchserver:8983/solr
Embedded server In this example, you store the index in the c:/temp/DC-SearchIndex directory on your computer.	Set the preferences.properties file as follows: SearchProvider=SolrEmbedded SolrEmbeddedDataDir=c:/temp/DC-SearchIndex

If you do not specify a directory for `SolrEmbeddedDataDir`, Decision Center stores the search index in a temporary directory. When the server stops running, it also stops using the temporary directory. When the server restarts, it creates a new temporary directory, and completely re-indexes the repository.

If you specify a directory for `SolrEmbeddedDataDir`, the directory and its content persist across server restarts, and the server does not re-index the repository with each restart.

Additional steps to configure testing and simulation

You can deploy and configure testing and simulation features to complement your Decision Center configuration on JBoss.

Before you start

You start by installing certain applications, and then you follow a sequence of steps to deploy the testing and simulation features.

Make sure that the following software is installed and configured on the same application server:

- Rule Execution Server
- Decision Center

To use the testing and simulation features in the Decision Center consoles, you must deploy the testing and simulation archive, which contains the following WAR files:

- Scenario Service Provider (SSP): Used to run tests and simulations in the Enterprise console and tests in the Business console.
- Decision Runner: Used to run simulations in the Business console.

What steps to follow

The following table shows the configuration steps for testing and simulation:

Step	Mandatory/Optional
“Step 1: Creating Decision Warehouse database resources”	Mandatory. If you used the Rule Execution Server console to create database resources, you do not have to do this step.
“Step 2: Creating Decision Runner database resources” on page 44	Mandatory. To run simulations in the Business console, you must create dedicated tables for the Decision Runner in your database.
Step 3: Deploying the testing and simulation archive for the first time This archive is the default testing and simulation archive packaged with the installer.	Mandatory. This steps allows you to check the availability of the feature. Note: You can also use the default SSP archive if you have an XML XOM. In this case, you do not have to repackage the archive to include the XML XOM.
“Step 4: Checking the availability of the testing and simulation services” on page 45	Optional
“Step 5: Using an Ant task to package the SSP archive” on page 46	Optional. You can use this task as an alternative method for configuring or modifying the testing and simulation archive.

Step 1: Creating Decision Warehouse database resources

You set up the Decision Warehouse to support testing and simulation services.

To use the Decision Warehouse, you must create dedicated tables in your database. You can use SQL scripts to create these tables. The SQL scripts are in `<ODM_InstallDir>/executionserver/databases`. A readme file in this directory provides additional information about the scripts.

The script that creates the Decision Warehouse database schema is named `trace_<database_name>.sql`.

Note:

The Installation Settings wizard in the Rule Execution Server console creates all the necessary tables for Rule Execution Server and the Decision Warehouse. If you are configuring the testing and simulation features and you have already run the Installation Settings wizard to create the tables, you do not have to create database resources manually. However, if you did not use the wizard to create database resources, you must run the script to create the Decision Warehouse database schema. Also, you can use the Decision Warehouse only if the Rule Execution Server persistence is set to `datasource` or `jdbc`.

Only users belonging to resAdministrators or resMonitors groups can see the Decision Warehouse tab in the Rule Execution Server console after the testing and simulation features are configured.

Use any tool that can handle SQL to import and run the SQL scripts. If you use Command Editor to run the scripts, you must log in with the credentials that you use for the data source for Rule Execution Server.

The following table shows the tools for the supported databases:

Database	Database tool
IBM DB2	DB2 command line processor
Derby	ij command line processor
MySQL	mysql command line processor
Oracle	sqlplus command line processor
Postgre SQL	Postgre SQL command line tool
SQL Server	Query Tool
Sybase	isql command line processor

To access the database, the database user must have the following credentials:

- A user ID and a password
- Complete privileges on the tables and view of the schema (create, insert, delete)
- create index privileges
- On Oracle, create trigger and create sequence privileges. If you use an Oracle database, run all the scripts in the SQL Plus client.

Install a database client for the database that you use (refer to the documentation for the database).

The default CLOB size might not be sufficient for the FULL_EXECUTION_TRACE field in the EXECUTION_TRACES table. You might need a size qualifier if SQL raises exceptions with the *<Lob-Value>* reason code.

Using the DB2 database

When you use DB2 (except on z/OS), the scripts that create the Rule Execution Server database tables are written for databases that use automatic storage.

- BP32K is the buffer pool that is expected in SYSCAT.BUFFERPOOLS. If BP32K is not there, you can use the existing buffer pool or create a new buffer pool named BP32K. Use the following command to query SYSCAT.BUFFERPOOLS for the existing buffer pool:

```
Select * from SYSCAT.BUFFERPOOLS
```

Otherwise, use the following command to create a buffer pool named BP32K:

```
CREATE BUFFERPOOL BP32K SIZE 2000 PAGESIZE 32K
```

- You must update the trace_db2.sql script and select the custom option in the Installation Settings wizard to run it. Modify the following line in the script to specify storage for the table space:

```
CREATE TABLESPACE RESDWTS PAGESIZE 32K BUFFERPOOL BP32K;
```

Here is an example of the table space specification in the script:

```
CREATE TABLESPACE RESDWTS PAGESIZE 32K MANAGED BY Database
USING [ FILE 'C:\DB2\Container.file' 640 ] BUFFERPOOL BP32K;
```

- Depending on your database settings, you might have to modify the script further.

Step 2: Creating Decision Runner database resources

To run simulations in the Business console, you must create dedicated tables in your database for the Decision Runner.

About this task

You can create the tables by using the Rule Execution Server console or an SQL script.

Procedure

Select a method to create the tables:

- To create the tables with the Rule Execution Server console:
 1. Open the Rule Execution Server console (see Rule Execution Server console online help).
 2. Run **Diagnostics**.
 3. In the Decision Runner section, follow the link to the installation wizard and use the wizard to create the tables.

Note: The Decision Runner section shows the link to the wizard only if the tables do not already exist.

- To manually create the tables:
 1. In `<ODM_InstallDir>/executionserver/databases`, select the SQL script that matches your database and run it with the appropriate database tool. The script that creates the tables for the Decision Runner is named `decisionrunner_<database_name>.sql`.

Database	Database tool
IBM DB2	DB2 command line processor
Derby	ij command line processor
MySQL	mysql command line processor
Oracle	sqlplus command line processor
Postgre SQL	Postgre SQL command line tool
SQL Server	Query Tool
Sybase	isql command line processor

Step 3: Deploying the testing and simulation archive

After creating the database resources, you can deploy the testing and simulation archive.

About this task

To use the testing and simulation features, you must deploy the default testing and simulation archive.

You must deploy the archive to the same server as the execution unit (XU). Moreover, for the testing and simulation services to work, the SSP and Decision Runner WARs in the archive must be installed on the same server and port as the Rule Execution Server console.

If a Rule Execution Server console instance is not deployed on the same server and port, you must implement the `IlrSSPResolver` interface for SSP. In the resolver, you can use the server name to return a specific server URL.

For the Decision Runner, if a Rule Execution Server console instance is not deployed on the same server and port, you must set the `RES_URL` init parameter on the Decision Runner web application to the Rule Execution Server console that uses the Decision Runner.

Procedure

Copy the JBoss file `<ODM_InstallDir>/executionserver/applicationservers/JBoss6EAP/jrules-ssp-JBOSS61EAP.ear` to `<JBoss_HOME>/server/standalone/deployments`.

You can use this procedure to deploy the default archive that is packaged with the installer or any subsequent deployment of a repackaged archive.

If you have an XML XOM, you can also use the testing and simulation archive to test your rules. Because the ruleset archive inside a RuleApp includes an XML XOM, you do not have to repackage the archive to include the XOM.

Step 4: Checking the availability of the testing and simulation services

You use URLs to check the availability of the Decision Runner and Scenario Service Provider (SSP) applications.

The URLs display different information on each application. Follow these steps to use the URLs:

1. To check the Decision Runner application:
 - a. Enter the URL `http://<host>:<port>/DecisionRunner` in a web browser.
 - b. Log in to the application in one of the Rule Execution Server roles.
The application displays a home page that contains version and patch-level information.
2. To check the SSP application:
 - a. Enter the URL `http://<host>:<port>/testing` in a web browser.
 - b. Log in to the application in one of the Rule Execution Server roles.

The SSP application displays a home page that contains information about the SSP server:

Version

The version of Decision Server used.

Patch level

The patch level of Decision Server used.

License information

The type of license of this version.

RuleSession

The rule session type (POJO or J2SE).

DAO Factory Class

The Data Access Object (DAO) factory class that is used to persist the trace into the Decision Warehouse.

Job store class

The name of the class that is used to persist the SSP job into a cache to free the memory during long computations.

Job pool size

The size of the pool for asynchronous execution.

Started since

The time and date when the SSP started.

Jobs currently running

The About screen provides information about the jobs that are currently running after you run test suites with SSP:

- Job ID: Listed in the table when a user clicks **Run** in Decision Center.
- Created column: Records the date and time when each job is initialized.
- Status column: Shows the number of tested scenarios as compared to the total number of scenarios.
- Start time: Records the time when a resource is allocated for the job.
- Parts column: Records the number of parts in the job:
 - A job that is not run in parallel has one part.
 - A job that is run in parallel has one or more parts.
- End time: Records the time when the execution of the job is complete, that is, all the scenarios in the job have been executed, the tests have been executed (for test suites), and the KPIs have been computed (for simulations).

Note:

The report for the job is automatically downloaded by Decision Center at the end of the execution. If the scenario suite is run in the background, the user downloads the report by viewing the list of scenario suites, and then clicking the report link when it becomes available. After the report is viewed, the job is removed from the table. The job remains in the table until the report is downloaded.

Step 5: Using an Ant task to package the SSP archive

You package the archive for testing and simulation.

About this task

The **ssp-setup** Ant task updates the SSP archive to your specific configuration and XOM (see `ssp-setup`).

Note:

This method works on Windows and other supported distributed platforms only.

Procedure

1. Define the Ant task in your build file by using the `<taskdef>` Ant element in one of the following ways:
 - Define the task at the top level or within a specific target.

```
<taskdef resource="res-tasks.properties"
classpath="${<InstallDir>}/executionserver/lib/jrules-res-setup.jar"/>
```
 - If the JAR file is available in your system, you can write the following code:

```
<taskdef resource="res-tasks.properties"/>
```
2. Use the **ssp-setup** Ant task to update the SSP artifact.
3. Run the Ant task in one of the following ways:
 - From the command line, run Ant in the appropriate directory, followed by the name of the build file if necessary.
 - From Eclipse, right-click the Ant file and click **Run**.

What to do next

You can now deploy the archive.

Distributing rule testing to multiple servers

When you configure the Rule Execution Server environment on a domain with multiple servers, you can define which Rule Execution Server instances are used to execute rule tests.

Before you begin

1. Package all your executable object models (XOMs) into the `ssp.war` archive. For more information, see Making the XOM accessible by repackaging the SSP.
2. Make sure that the archive is deployed on each server along with a XU connection factory and the appropriate data source definition.

About this task

You create a custom resolver to define which Rule Execution Server instances are used to run rule tests.

Procedure

1. Implement the `IlrSSPResolver` interface.

For a simple implementation, you can use the server name to return a specific server URL. For example, you can have two servers that are defined in Decision Center:

- `testingServer1: http://host1:9080/res`
- `testingServer2: http://host1:9080/res`

And your resolver can evaluate the testing URL as follows:

```
if ( serverName.equals("testingServer1") )
{
    return ( new URL( "http://server1:9080/testing" ) );
}
else if ( serverName.equals("testingServer2") )
{
    return ( new URL( "http://server2:9080/testing" ) );
}
else
    return( new URL( "http://host1:9080/testing" ) );
```

2. Add your class to the `teamservice.war` archive.

- Set the `teamserver.defaultSSPResolver.class` property to that class.

Results

When you run a rule test, the execution is directed by the server that you choose.

What to do next

For a better implementation, you can set a dependency on the project that is being tested so that you can distribute the test execution according to that project.

```
ManagerBean bean = ManagerBean.getInstance() ;
IlrSession session = bean.getSession();
IlrRuleProject project = session.getWorkingBaseline().getProject();
String pname = project.getName();
```

Tuning the Decision Runner web application

You improve the performance of simulation runs in the Business console.

Before you begin

The Business console runs simulations on the Decision Runner web application. A simulation can run on one or more threads, and run a ruleset multiple times, requiring an execution unit (XU) (see Execution unit (XU)).

You can change Decision Runner and XU parameters to make simulations more efficient:

- Decision Runner: Change the maximum number of concurrent threads.
- XU: Change the connection pool size and the connection pool wait policy.

About this task

To complete this task, you must first estimate the maximum number of single-threaded simulations (**X**) and multithreaded simulations (**Y**) that might be started in parallel, and the maximum number of threads that might be used for multithreaded simulations (**Z**). Use this formula to determine the maximum number of concurrent threads: $X + (Y \times Z)$.

Procedure

- Change the maximum number of concurrent threads in the Decision Runner on your server:

Application server	Method
<ul style="list-style-type: none"> WebSphere Application Server 8, 8.5, and 8.5.5 WebLogic 11g (10.3.6) and 12c 	<p>Use the work manager that is associated with the Decision Runner web application to set the maximum number of concurrent threads that are used by simulations.</p> <p>Note: To represent accurately the maximum number of concurrent threads that are used by the Decision Runner, the bounded work manager must work with only the Decision Runner. Otherwise, the number of maximum threads must be set much higher than the required amount to accommodate requests from other components.</p>

Application server	Method
<ul style="list-style-type: none"> • Tomcat 7 • JBoss 5.1.2 and 6.1 	Edit the <code>THREADPOOL_MAXIMUM_SIZE</code> parameter in the web application deployment descriptor (<code>web.xml</code>) of the Decision Runner application. The default value is 10.
<ul style="list-style-type: none"> • Liberty profile 8.5.5.3 	You cannot set the maximum number of concurrent threads for simulations directly on the executor service that is associated with the Decision Runner web application. Executor services on Liberty use the Liberty common thread pool. If necessary, you can tune the maximum number of threads directly on the common thread pool. Note that the common thread pool is shared, and its threads are not used by the Decision Runner only.

- Set the connection pool size for the XU. Change the size to be in line with the capacity of the server and greater than the maximum number of concurrent threads.

If you cannot set the connection pool size to be greater than the maximum number of concurrent threads, you must lower your estimate. For best results, dedicate a XU to simulations. If other applications must use the XU, try to use a number of connections greater than the maximum number of concurrent threads.

Note:

For more information about customizing the connection pool of a XU, see Rule session tuning. For WebSphere Application Server, WebLogic Server and JBoss, follow the steps for Java EE. For Tomcat and Liberty, follow the steps for Java SE.

- Set the connection pool wait policy for the XU so that the connection pool refuses new connections immediately when the pool is full. If simulations fail, check your application server logs for the following Decision Runner error:

- `GBRXU200E: The default connection manager pool is full.`

If you get this message, increase the connection pool size to try to solve the problem.

Alternatively, depending on the application server, you can change the connection pool wait policy so that the connection pool is able to wait for a connection to be released to fulfill a connection request when the pool is full. However, if the wait time is set to an indefinite amount of time, it might cause some threads to hang in the system.

Example

In setting up the Decision Runner, you determine that you must be able to run at the same time 10 single-threaded simulations and 4 multithreaded simulations that have a maximum of 8 threads each. The maximum number of concurrent threads comes to $10 + (4 \times 8) = 42$.

You have a server that can handle a connection pool size of 60. The XU is shared, but the other applications take no more than 10 connections at the same time, so we have 50 connections available for simulations.

Using these parameters, you can have up to 8 more single-threaded simulations when compared to the estimated simulation load ($50 - 42 = 8$), or you can have more than 8 threads for some multithreaded simulations, at full performance and without getting failures. If you exceed these parameters, the connection pool wait policy might make new simulations and some existing running simulations fail.

Chapter 2. Verifying your configuration of Decision Center

You can verify that you have correctly configured Decision Center by publishing some projects, opening the consoles, and running the diagnostics.

Publishing a project to Decision Center

After completing the configuration, Decision Center is ready to be used but does not contain any rule projects. You publish rule projects from Rule Designer.

About this task

To publish a rule project to Decision Center, the project must be imported into your workspace.

The procedure uses the Decision Center tutorial projects as an example of how to import and publish a rule project. If you want to carry out the Decision Center tutorials, you have to publish the following projects:

- loanvalidation-rules (with loanvalidation-xom)
- loanvalidation-rules-dependent
- squery-loanvalidation-rules (with squery-loanvalidation-xom)

Procedure

1. To open Rule Designer, click **Start > All Programs > IBM > *package_group* > Rule Designer**.
2. In Rule Designer click **File > Import > General > Existing Projects into Workspace**, and click **Next**.
3. Click **Select root directory**, browse to <InstallDir>/studio/tutorials/shared, and click **OK**.
4. Select the projects and click **Finish**.
5. Right-click the loanvalidation-rules rule project, and click **Decision Center > Connect**.
6. Complete the Decision Center Configuration dialog as follows. The warning message Connection not established displays until you establish the connection.

User name

rtsAdmin

Password

rtsAdmin

URL http://localhost:<port>/teamserver

Data source

Leave this field empty.

Note: If security is enabled, use https://localhost:<PORT_NUMBER>/teamserver

7. Click **Connect**.

The connection is established when the warning message closes and the **Project configuration** area becomes active.

8. In the **Project configuration** area, check that **Create a new project on Decision Center** is selected, and then click **Finish**.
9. The Synchronize Complete - Decision Center Participant dialog opens when the publishing process is complete. Click **OK** to close this dialog.
10. A dialog opens asking you if you want to change to Team Synchronizing perspective. Click **Yes**.
An empty Synchronize view opens, indicating that the projects in Rule Designer and Decision Center are the same. This means that your rules are now published to Decision Center.
11. Repeat for the other rule projects.

What to do next

You can now open the Decision Center Enterprise console and perform diagnostics.

Opening the Decision Center consoles

After you have deployed the Decision Center EAR or WAR to your application server, you can open the Decision Center consoles.

You can open the consoles by using the following URLs in a web browser:

- **Enterprise console:** `http://localhost:<PORT_NUMBER>/teamserver`
- **Business console:** `http://localhost:<PORT_NUMBER>/decisioncenter`

Note: If your browser is not running on the same host as the application server, replace `localhost` with the address of the machine. If your web application is mapped to a host with a port that is different from the default port, use the port number of the host.

By default, the data source is `jdbc/ilogDataSource`. If you want to specify a different data source, you have to pass it as a request parameter in the URL. For example:

```
http://localhost:7001/teamserver?datasource=jdbc/serverextendedbrm.
```

The locale of the sign-in page is English by default. You can specify a locale parameter in the URL that switches the sign-in page to the required locale. For example:

```
http://localhost:<port>/teamserver?locale=es (assuming that your message files are localized).
```

If you sign in with another locale in the URL and want to change the locale afterward, click **Options** in the top banner of the Enterprise console or **Profile** in the Business console. This saves the locale and restores it the next time you sign in.

If you open Decision Center but no database exists, you automatically access the Installation Settings wizard with only the **Install** tab available.

After completing the installation, Decision Center is ready to use but does not contain a rule project. You have to publish a rule project from Rule Designer.

A diagnostics tool, available in the Configure tab of the Enterprise Console, shows a report on the status of your Decision Center configuration.

To learn more about Decision Center, see [Decision Center](#).

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 grant you any license to these patents. You can send license inquiries, in writing, to:

IBM United Kingdom Laboratories,
Mail Point 151,
Hursley Park,
Winchester,
Hampshire,
England SO21 2JN

For license inquiries regarding double-byte character set (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan, Ltd.
19-21, Nihonbashi-Hakozakicho, Chuo-ku
Tokyo 103-8510, Japan

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

websites. The materials at those websites are not part of the materials for this IBM product and use of those websites 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.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM United Kingdom Laboratories,
Mail Point 151,
Hursley Park,
Winchester,
Hampshire,
England SO21 2JN

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

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. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Each copy or any portion of these sample programs or any derivative work, must include a copyright notice as follows:

© (your company name) (year). Portions of this code are derived from IBM Corp. Sample Programs. © Copyright IBM Corp. _enter the year or years_.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at www.ibm.com/legal/copytrade.shtml.

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

Microsoft, Windows, and the Windows logo are trademarks of Microsoft 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.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

Index

A

- Ant tasks
 - configuring Decision Center 30
 - creating the database schema 33
 - defining and uploading message files 36
 - execute-schema 34
 - executing the database schema 34
 - gen-create-schema 33
 - gen-drop-schema 35
 - print-config-param 37
 - remove-config-param 37
 - removing a database schema 35
 - repackage-war 39
 - repackaging the Decision Center archive 39
 - set-config-param 37
 - setting configuration parameters 37
 - setting the persistence locale 37
 - setting up the environment 31
 - ssp-setup 46
 - to repackage an EAR file 15
 - upload-extensions 34
 - upload-messages 36
 - upload-roles 34
 - uploading groups to the database 34
 - uploading the database schema extension 34
- asynchronous execution
 - pool size 45

C

- clusters
 - JBoss 17
- configuration parameters
 - adding or removing 37
 - setting up 29
- configuring Decision Center
 - declaring custom groups 24
 - deploying Decision Validation Services 44
- configuring Decision Validation Services 46
 - checking availability of Decision Validation Services 45
 - creating database resources 42
- configuring Decision Validation Services on JBoss 41
- connection pool
 - creating 21
- custom groups, declaring
 - Decision Center 24

D

- data source
 - creating 21
- data sources
 - creating 6, 7

- database
 - establishing credentials 4
- database resources
 - for Decision Validation Services 42
- database schema
 - creating 33
 - executing 34
 - for Rule Execution Server
 - choosing, using Installation Settings wizard 11
 - reviewing, using Installation Settings wizard 12
 - removing 35
 - uploading extension 34
- databases
 - creating a schema from the console 9
 - permissions
 - Decision Center 20
 - populating with SQL scripts 13
 - user permissions 3
- Decision Center
 - adding groups 22
 - adding users 22
 - configuring for JBoss 19
 - configuring security 22
 - configuring the database 27
 - configuring with Ant tasks 30
 - deploying on JBoss 24
 - Installation Settings wizard 26
 - opening 52
 - repackaging the archive 39
 - setting up configuration parameters 29
 - setting up groups 28
 - setting up message files 27
 - setting up persistence locale 29
 - verifying deployment 25
- Decision Validation Services
 - checking availability 45
 - configuring 41
 - deploying 44
 - distributing tests to multiple servers 47
 - packaging using Ant 46
- Decision Warehouse
 - persistence setting for execution trace 3
- declaring
 - custom groups in Decision Center 24
- deploying
 - Decision Validation Services 44
- deployment descriptor 4
- diagnostics
 - running 14
- dropping a database schema
 - See removing

E

- EAR archives
 - repackaging using Ant 15

- environment variables
 - for Ant 31

G

- groups
 - adding 22
 - adding in JBoss 4
 - setting up 28
 - uploading to the database 34

H

- hosted transparent decision service
 - deploying on JBoss 14

I

- Installation Settings wizard 9, 26
 - choosing the database schema 11
 - Configure Database 27
 - report 12
 - reviewing the database schema 12
 - Set Configuration Parameters 29
 - Set Persistence Locale 29
 - Set up Groups 28
 - Set up Message File 27
- installing Decision Center
 - restricting user permissions 20

J

- JBoss application server
 - adding groups 4
 - adding users 4
 - clusters 17
 - configuring for Decision Center 19
 - configuring for Rule Execution Server 1
 - configuring security 4, 22
 - creating a connection pool 21
 - creating a data source 6, 7, 21
 - creating a database schema 9
 - deploying Rule Execution Server 9
 - deploying the Decision Center EAR 24
 - deploying the XU RAR 8
 - hosted transparent decision service 14
 - verifying deployment of Decision Center 25
- JDBC driver
 - creating 6

M

- message files
 - defining and uploading 36
 - setting up 27

O

- opening
 - Decision Center 52
 - Rule Execution Server console 10

P

- packaging
 - Decision Validation Services 46
 - Rule Execution Server archives using Ant 15
- permissions 22
 - database access 3
 - Decision Center installation 20
- persistence
 - configuring in Rule Execution Server 3
- persistence locale
 - setting 37
 - setting up 29
- publishing rule projects
 - from Decision Center 51

R

- removing
 - database schema 35
- Rule Execution Server
 - configuring
 - using Ant tasks 15
 - configuring for JBoss 1
 - configuring security 4
 - creating a database schema 9
 - deploying on cluster environments 16
 - deploying on JBoss 9
 - deploying the hosted transparent decision service 14
 - deploying the XU RAR 8
 - distributing DVS tests to multiple instances 47
 - populating a database 13
 - running diagnostics 14
 - setting up the persistence mode 3
- Rule Execution Server console
 - opening 10
- rule projects
 - publishing
 - from Decision Center 51

S

- Scenario Service Provider (SSP)
 - ssp-setup Ant task 46
- security settings
 - configuring 4
- SQL scripts
 - to remove a database schema 35
- ssp-setup, Ant task 46

U

- user permissions 3
- users
 - adding 22

- users (*continued*)
 - adding in JBoss 4

W

- WebLogic Server application server
 - verifying deployment of Decision Center 25

X

- XU RAR
 - deploying on JBoss 8



Printed in USA