Relational Population Service (RPS)

Contents

Relational Population Service (RPS)

Introduction

The Relational Population Service (RPS) provides automatic replication of objects from a primary JADE database to one or more relational databases.

RPS enables JADE systems to interoperate with relational databases for reporting, Business Intelligence (BI), Extract, Transform, Load (ETL), and data warehousing purposes.

Microsoft SQL Server (versions 2000, 2005 and 2008) is the first relational database supported by RPS.

RPS enables JADE systems to automatically replicate data into external relational databases in near real-time. Developers can specify the classes, objects, and properties that are to be replicated to the required relational database or databases. This allows an entire JADE database to be replicated, or just the data that is required for external purposes.

RPS is built on Synchronized Database Service (SDS) technology and uses the same mechanisms as SDS for propagating updates to secondary databases. Updates recorded in database journals written by the primary database are mirrored to attached secondary databases and applied.

RPS can populate an entirely independent relational database, or tables within an existing relational database where that database can be updated from multiple sources (that is, via JADE's RPS and also from non-JADE systems). In either case, the tables updated by RPS are defined and controlled by JADE.

Additionally, an RPS node operating in **full replica mode** can also be used for disaster recovery, as it can acquire the primary database role in a negotiated or hostile takeover.

This White Paper walks you through the steps required to set up an RPS environment. It uses JADE's Erewhon Investments demonstration application and Microsoft SQL Server Express (freely downloadable) for the relational database.

You need a license for the primary database node and a license for the RPS node. If you do not have these, please e-mail askjade@jadeworld.com for consideration of licenses for evaluation purposes.

Please refer to the *JADE Roadmap* for details about future RPS enhancements:

www.jadeworld.com/downloads/jade/JADE_Roadmap.pdf

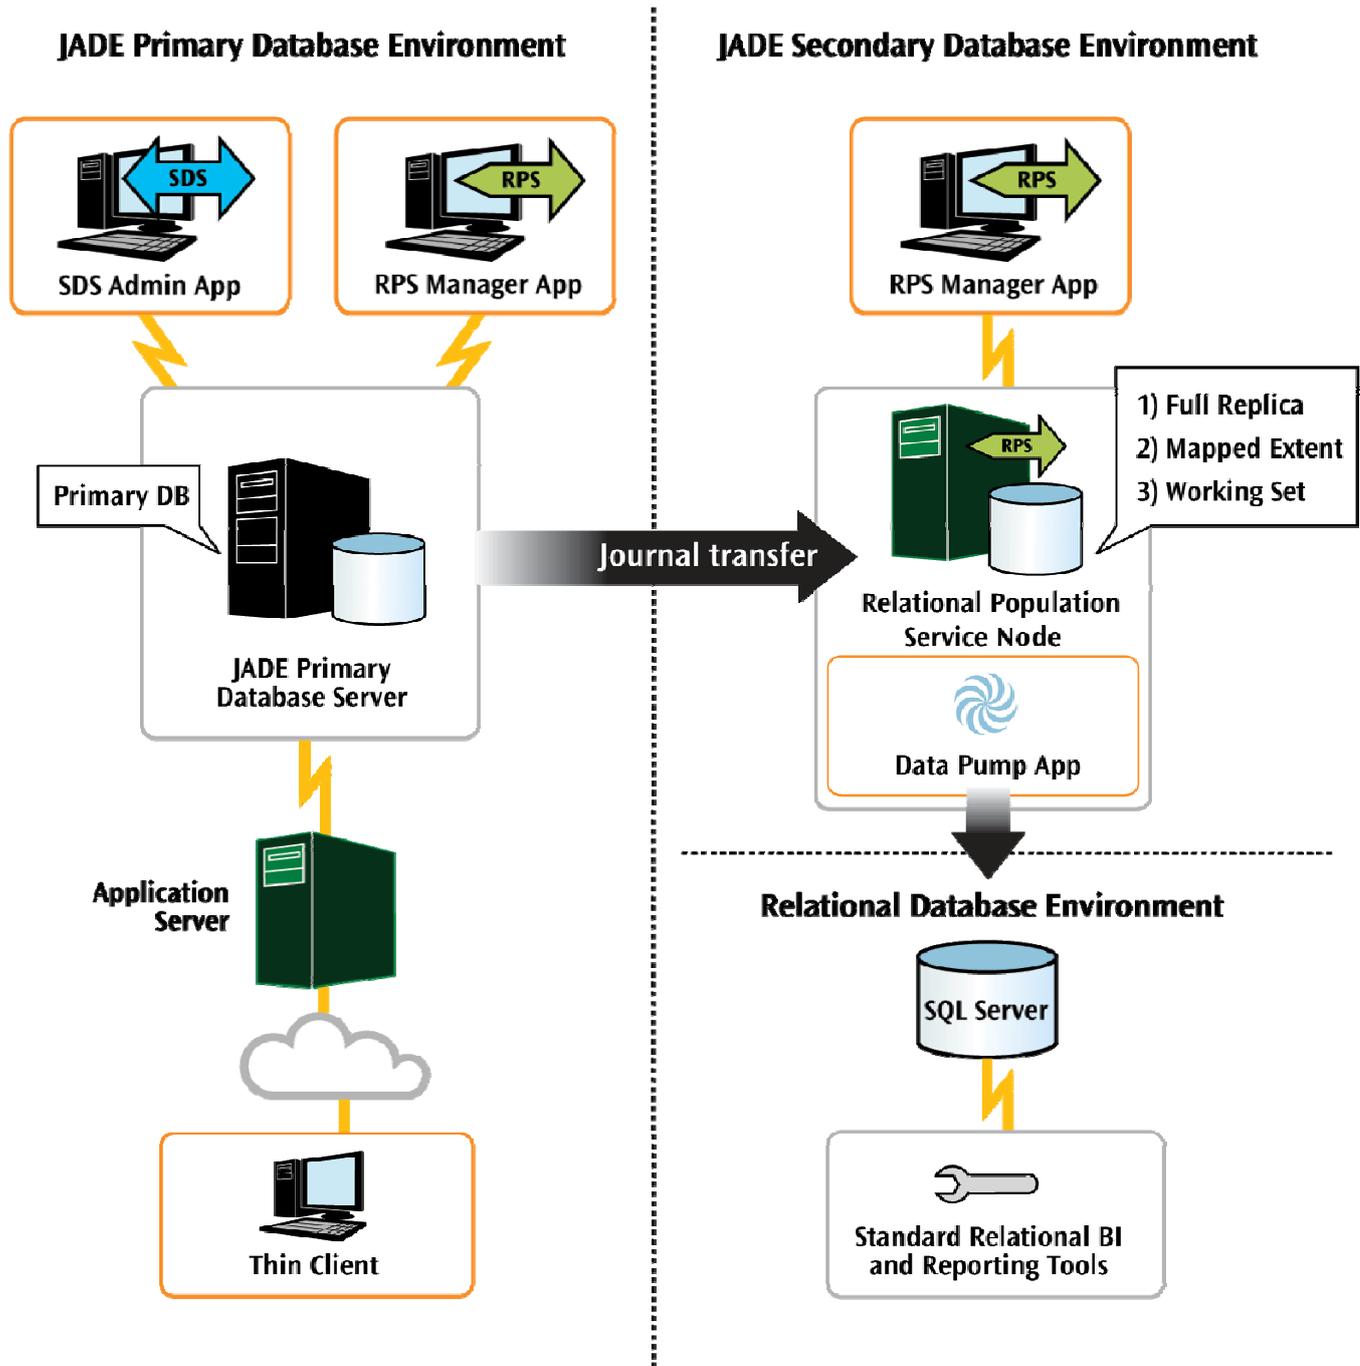Other JADE documentation referred to in this White Paper can be downloaded from:

www.jadeworld.com/jade/doc

Overview

This section provides an overview of the RPS configuration and components.

Configuration

The RPS configuration is shown in the following diagram.

Components

The following components, shown in the previous diagram, are required when using RPS. The setup of these components is described later in the "Setup" section.

JADE Primary Database Server

The JADE primary database server is the database server for the production JADE database.

Relational Population Service (RPS) Node

An RPS node is an SDS secondary database configured with a Relational subrole. This specialised secondary database server can be referred to as an RPS secondary database node or just RPS node, for short. It is convenient to refer to the persistent storage managed by an RPS node as an **RPS database**.

An RPS database can be a full copy or a subset of the primary database, and it always includes a full set of user schema files. Configuring the extent of an RPS database is described later in "Create RPS Database".

The RPS database can be hosted on the same machine or a different machine to the primary database. The examples in this document assume that the RPS database is hosted on the same machine as the primary.

An RPS node participates in a Synchronized Database Environment (SDE). An SDE comprises a primary database server and one or more secondary database servers. A secondary database is either a native secondary or one configured with a relational subrole to make it RPS-capable. This document describes setting up a single RPS node.

Journal Transfer

Objects in the RPS database are kept synchronized with objects in the primary by periodically applying audit records from journals written on the primary and mirrored to all attached secondary databases.

SDS Administration Application

The SDS Administration application provides a facility for administering the primary and one or more secondary database servers in an SDE. For more details, see "Using the SDS Administration Application", in Chapter 6 of the *JADE Administration Guide* for JADE 6.1 or in Chapter 1 of the JADE *Synchronized Database Service (SDS) Administration Guide* for JADE 6.2 and later releases.

RPS Manager Application

Administration of the RPS node is done from the RPS Manager application. This application can be run from either the primary or RPS systems. For more details, see "Managing an RPS Node" in Chapter 9 of the *JADE Administration Guide* for JADE 6.1 or in Chapter 2 of the JADE *Synchronized Database Service (SDS) Administration Guide* for JADE 6.2 and later releases.

Datapump Application

The RPS **Datapump** application executes in the RPS node. It connects to the relational database (RDB) and propagates updates that were made on the primary as described by the RPS Mapping.

Relational Database Environment

The relational database environment is the target Relational Database Management System (RDBMS) database server. The supported target database servers are Microsoft SQL Server versions 2000, 2005 and 2008. These can be hosted on the same machine or a different machine to the primary and RPS database servers.

The example in this document uses SQL Server Express and assumes that the same machine hosts all database servers.

Enhancements in JADE 6.2 (September 2007)

Customer feedback indicated that developers wanted greater control of how JADE objects were mapped to relational tables and they also wanted to be able to perform more operations on objects before they were replicated from JADE to the relational database.

JADE 6.2 delivered more flexibility and made it easier for developers to map JADE objects to relational tables by:

- Allowing methods and virtual properties to be mapped to table columns. This allows derived values to be replicated and enables developers to write methods that aggregate multiple objects into a single table row.

- Allowing developers to interrogate and change rows prior to them being sent to the relational database. This provides developers with full control of table rows before they are sent from JADE to the relational database.

- Enabling developers to specify (in the primary environment) those transactions that are *not* to be replicated to the relational database. This allows selected operations on the primary system (for example, housekeeping and archiving tasks) to be performed without affecting relational database replicas.

Several operational enhancements were also provided, including an RPS Management API (via both JADE classes and Web services), the ability for developers to specify customer-specific mappings, and support for client-server relational database loads.

Setup

This document uses the Erewhon system as an example to demonstrate the setup and running of an RPS environment.

Before You Start

The following must be setup before proceeding.

1. JADE installed.

    This is referred to as the *primary* database environment and these instructions assume that it is installed in the default location. For example, for JADE 6.2 the default installation folder is **C:\Jade62** and for JADE 6.3 it is **C:\Jade63**.

    > **Note**   Throughout this document, **<JadeVer>** is used to refer to the installation folder name. Whenever you see **<JadeVer>**, replace it with the name of your JADE installation folder.

2. RPS licence key applied to the primary database. This license is transferred to the RPS environment when it is created. If you do not have an RPS license key and require one, please e-mail askjade@jadeworld.com.

3. Shortcuts defined for the primary database environment, including the following:
    □   JADE Database Server
    □   JADE Application Server
    □   JADE Presentation Client
    □   JADE RPS Manager
    □   JADE SDS Administration

    All of these shortcuts should have been created as part of the JADE installation.

4. Erewhon schemas loaded into the primary database. These schemas are available from the JADE installation in **C:\<JadeVer>\examples\erewhon**.

    Sample JADE schemas, including Erewhon, can be downloaded from www.jadeworld.com/jade. Look for the "Example Schemas" link in the side bar on the bottom right of the page.

5. Erewhon data loaded.

    Follow the instructions in the **Erewhon.pdf** document, to initialize Erewhon.

6. Microsoft SQL Server Express installed. This is the relational database that is used. If you already have a full version of SQL Server installed, you can proceed using that version instead of SQL Server Express.

    Microsoft SQL Server Management Studio Express installed. This is used to manage your SQL Server Express database.

    Both products can be downloaded and installed with default settings from http://msdn.microsoft.com/vstudio/express/sql/download.

Defining a Secondary RPS Database

This section describes the steps required to define a secondary RPS database.

Set Up RPS Folders

Set up the folders and binary files to be used by the RPS environment, as follows:

1. Create a folder for the RPS environment.

   □ Create **C:\<JadeVer>RPS**

   This is the folder where the RPS binary and database files will physically reside.

   > **Note**   This is creating a single folder with **RPS** appended to the **<JadeVer>** name. It is *not* creating an **RPS** subfolder within the **<JadeVer>** folder.

2. Copy the binary files to run the RPS environment from the primary to the RPS folder.

   □ Copy **C:\<JadeVer>\bin** to **C:\<JadeVer>RPS\bin**

3. Create a folder for the RPS database files.

   □ Create **C:\<JadeVer>RPS\system**

   This is an empty folder that will contain the database files. The database files are created from the primary, as per the "Create RPS Database" section below.

4. Create a folder for the SQL scripts.

   □ Create **C:\<JadeVer>RPS\RPSSQLScripts**

   This is an empty folder that will contain the SQL Scripts used to define the tables in SQL Server. These scripts will be created later, in the "Extract SQL Scripts" section.

Define RPS Mapping

The RPS Mapping defines the classes and attributes to be used in the RDB and the tables and columns to which these will map. The table and column definitions will later be extracted and loaded into the RDB.

1. Bring up the JADE development environment for the primary database.

   a. Start the **JADE Database Server**

   b. Start the **JADE Application Server**

   c. Start the **JADE Presentation Client**

   The RPS Mapping is defined from the JADE development environment.

2. Select the model schema.

   □ From the Schema Browser, click on **ErewhonInvestmentsModelSchema**

   The RPS Mapping is defined in the model schema because this schema contains all persistent class definitions.

3. Select to add or change the RPS Mapping.

   a. Select **Browse | RPS Mappings**

   b. Right-click on **ErewhonRPS** and select **Change**

ErewhonRPS was already defined in the Erewhon schema that you loaded. This definition can be customized as follows. Alternatively, a new RPS Mapping could be defined.

4. Select the Database Type. This specifies the type of the target RDB (your RDB must be the specified type and version or later) as well as the data type mappings that are used to convert JADE types to RDB types.

    ▫   Select **SQL Server 2005**

    Note   Specify **SQL Server 2005** for this setting even if you are using a later version of SQL Server, because it simplifies the ODBC configuration in later steps.

5. Set the Per Statement Exception Policy.

    ▫   Leave the default **Halt** setting for each of these

    If an exception occurs when creating, updating, or deleting a row in the RDB, the Halt selection will cause the transaction to be aborted and database tracking will cease.

6. Set the Logging Options.

    ▫   Leave the **Default** setting

    Exceptions will be logged to the **jommsg.log** file. If selected, the **RPSTable** option also records exception information in a table in the RDB.

7. Set the OID Mapping Options (where OID stands for *Object Identifier*).

    ▫   Leave the **Default** setting

    In the RDB, OIDs will be mapped to a 16-character fixed-length string in the format *ccccc.nnnnnnnnnn*. If selected, the **Split Into Two Columns** option splits OIDs into two integer values containing the class number and instance identifier.

8. Set the Default Database Name.

    ▫   Enter **Erewhon**

    This is the name of the target RDB. This name will be used when the RDB is set up, detailed in the "Create Relational Database" section.

9. Set the Default Connection String. *Do not* enclose the string in quotes.

    ▫   Enter **DSN=SQLServerODBC; Database=Erewhon**

    This is the connection string for the target RDB, including the ODBC DSN name. The ODBC definition will be set up later, in the "Create a New ODBC Data Source" section below. This is required for the RPS database to connect to the SQL Server database.

10. Set the Default Connection Usercode and Default Connection Password.

    ▫   Leave these values blank

    If specified, this user code and password are used on the RPS node to connect to the RDB. They can also be overridden from the RPS node. They are left blank here, as later steps specify that Windows Authentication will be used to connect to the RDB.

11. Set the **Select Top Most Schema** option.

    ▫   Leave the default **ErewhonInvestmentsModelSchema** selection

    All persistent classes in the **Erewhon** system are defined in the model schema, so this is the highest-level schema required for the mapping.

12. Set the **Show Virtual Properties** and **Show Methods** options. The **Show Methods** option is available from JADE 6.2 only.

    a. Leave the default **Do Not Show Virtual Properties** setting

    b. Leave the default **Do Not Show Methods** setting (from JADE 6.2)

    This means that virtual properties and methods will not be displayed and so cannot be selected for the RPS Mapping.

13. Go to **Select Classes for RPS**, by clicking **Next >**.

14. Select the classes that are to be mapped to tables in the RDB.

    ▫ No action is required, as classes are already selected

    The selected classes are listed in the **Selected Classes** list box. Further classes can be added from the **Available Classes** list box or removed from the **Selected Classes** list box. Only user-defined classes that can have persistent instances are available to be included in the RPS Mapping.

15. Go to **Map Classes to Tables**, by clicking **Next >**.

16. Map the selected classes to the tables required in the RDB.

    ▫ No action is required, as classes are already mapped

    The mapped table name can be changed for each class, by entering a new name in the relevant **Table** column cell. The name must be valid in the RDB, with a maximum length of 80 characters.

    In **Historical Mode**, object updates and deletes are mapped to a new row in the target table. When not in Historical Mode, an object update or delete will update or delete the corresponding row.

    When the **No Delete** option is selected, object deletes are not propagated to the target RDB.

    From JADE 6.2, you can select the **Excluded** option if you want to exclude the table from RPS replication but retain the table in the RDB schema definition.

    From JADE 6.2, you can select the **Output Callback** option if you want output for a table to be passed to your application before being replicated to the RDB. This option should be set only if required, to avoid unnecessary callback processing.

    The existing definition has mapped subclasses to the same table as their superclass. This was done from the **Map SubClasses to this Class Map** right-click menu item.

17. Go to **Select Columns for Tables**, by clicking **Next >**.

18. Select the features that are to be mapped to columns in the RDB.

    ▫ No action is required, as properties are already selected

    The properties for the selected table/class are listed in the **Features** list box (this list box is called **Attributes**, in JADE 6.1). Properties can be added or removed by clicking the relevant check boxes.

19. Go to RPS Column Mappings, by clicking **Next >**.

20. Map the selected attributes to the columns required in the RDB.

    ▫ No action is required, as attributes are already mapped

    The mapped column name, mapping type, ID, and key can be changed for each feature, by entering a new value in the relevant row and column. Select the **View Selected Table** option to see the features for one table at a time. The primary key for all tables will be the OID, which is the default.

21. Go to **Summary and Build**, by clicking **Next >**.

22. Check the details and finalize the RPS Mapping.

    ▫ Check that the details listed are correct and then click **Finish**

    This completes the RPS Mapping changes.

23. Reorganize the database, if required.

    a. Select the schema **ErewhonInvestmentsModelSchema**

    From JADE 6.2, if a reorg is required, the Schema Browser window will display two ErewhonInvestmentsModelSchema items. One of the items will have a highlighted background – this is the item that you should select.

    b. Select menu item **Schema | Reorg | Reorg Schema**

    c. Click **Reorg**

    d. When the reorganization has completed, click **OK**

Create RPS Database

The secondary RPS database files will be produced from the primary database in the folder that was specified in the "Set Up RPS Folders" section.

Note Make sure that you have applied your RPS license key to the primary database before proceeding. You can use **C:\<JadeVer>\bin\jadreg.exe** to do this.

1. Start the RPS Manager application for the primary database.

    a. Start the **JADE Database Server** (if not already running)

    b. Start the **JADE RPS Manager**

    The secondary RPS database will be created from the RPS Manager application.

2. Select the RPS Mapping schema.

    a. Select **ErewhonInvestmentsModelSchema**

    b. Select **ErewhonRPS**

    c. Click **OK**

    This specifies the mapping that was defined earlier. This mapping will be used to create the RPS database.

    Note This screen appears only if this is the first time the RPS Manager is run; otherwise the schema and mapping can be selected in the next step, if required.

3. Select to create a full RPS database.

    ▫ Select **RPS | Create RPS Database | Full**

    The **Full** option will replicate all persistent system and user schema database files from the primary database.

    The **Working Set** option would replicate a persistent copy of user schema files and a working set of persistent objects in temporary storage. This temporary storage is re-established whenever the RPS node is restarted.

    The **Mapped Extent** option would replicate user schema files and all user database files that contain classes included in the RPS Mapping defined in the primary database.

    An RPS Database Creation dialog will be displayed.

4. Enter the RPS Directory.

    ▫ Enter **C:\<JadeVer>RPS\system** or use the browse button to select this directory

    This specifies the folder where the secondary RPS database files will be created. This is the same folder as created in the "Set Up RPS Folders" section.

5. Set the Copy Mode.

▫ Leave the default **Online – Updating Allowed** option selected

This option indicates that updates to the primary database are allowed while the data is copied. For more details, see "Online Full Backup" in Chapter 5 of the JADE *Administration Guide* for JADE 6.1 or in Chapter 1 of the JADE *Database Administration Guide* in JADE 6.2.

The **Quiesced – Read Only** option would place the primary database in quiescent read-only state while the data is copied. For more details, see "Online Quiesced Backup" in Chapter 5 of the JADE *Administration Guide* in JADE 6.1 or in Chapter 1 of the JADE *Database Administration Guide* in JADE 6.2.

6. Set the Overwrite Existing Files option.

▫ Leave this as the default unchecked value

Checking this option would allow any existing files in the specified destination RPS folder to be overwritten when the data is copied from the primary database. There should currently be no files in this folder, so the option is not selected.

7. Start the database creation.

▫ Click **OK**

Progress messages will be displayed and when finished, an RPS Database Create Completed dialog will be shown.

8. Finish the database creation.

▫ Click **OK** on the RPS Database Create Completed dialog

You will be returned to the RPS Manager application window with the following progress messages displayed:

```
> RPS 'Full' Database create started.
> RPS Database create completed.
```

Note   You may see a message saying that files have not been copied. You can safely ignore this.

The secondary RPS database has now been created in **C:\<JadeVer>RPS\system**.

Defining Shortcuts and Setting Up the Initialization File

Shortcuts and initialization (**ini**) files to be used by the primary and secondary databases must be set up to define where these are run. The secondary database is the database for the RPS environment.

Set Specific Initialization File Settings

The same initialization (**ini**) file can be used for both the primary and secondary databases. Our examples assume that the same **ini** file is used.

When a common **ini** file is used, separate sections must be defined to specify different values required for the primary and secondary databases. These sections are distinguished by prefixing the section name with a *unique-id*. The format of this is [*unique-id.section-name*].

The *unique-id* must also be specified in the command line for the JADE program that is to use it. Examples of the required command line parameter are given in the following "Create Shortcuts" section.

If separate **ini** files are used for the primary and secondary databases, the *unique-id* prefix is not required. See the JADE *Initialization File Reference* for further details on individual **ini** file settings.

1. Close the JADE database.

    a. Close the **JADE RPS Manager**

    b. Close the **JADE Presentation Client**

    c. Close the **JADE Application Server**

    d. Close the **JADE Database Server**

    The database is closed so that when it is next run, the following **ini** file changes will take effect.

2. Create a copy of the **ini** file.

    ▫ Copy **C:\<JadeVer>\System\jade.ini** as **C:\<JadeVer>\jadeRPS.ini**

    This creates a copy of the default primary database **ini** file so that changes can be made for use by both the primary and secondary databases.

3. Edit the **ini** file.

    ▫ Open the **C:\<JadeVer>\jadeRPS.ini** file in an editor

    The new **ini** file can now be updated with unique sections used by the secondary database.

4. Specify the server details for the secondary database.

    ▫ Add the following section:

    ```
    [RPSServer.JadeServer]
    NodeName=RPSServer
    NetworkSpecification1=TcpIp,Enabled,6007
    ```

    This defines the parameters used to run the secondary JADE server node. **RPSServer** has been used as the *unique-id* to define this as applying to the secondary database only. The *unique-id* must match that specified in the RPS shortcuts defined later in this document.

5. Specify the client details for the secondary database.

   □ Add the following section:

   ```
   [RPSServer.JadeClient]
   ServerNodeSpecifications=TcpIp,LocalHost,6007
   ```

   This defines the parameter used to run the secondary JADE client node. **RPSServer** has been used as the *unique-id* to define this as applying to the secondary database only. The *unique-id* must match that specified in the RPS shortcuts defined later in this document.

6. Specify SDS (that is, secondary) details for the primary database.

   □ Add the following section:

   ```
   [Primary.SyncDbService]
   DatabaseRole=PrimaryRole
   MyName=Primary
   ```

   This section defines the SDS parameters used by the primary database server. An RPS node is a special kind of SDS secondary server and therefore the same SDS **ini** file options are used. Note that the **MyName** parameter must match the extension used in the [ConnectionParams] section defined later in this document.

7. Specify the SDS parameters for the secondary RPS database.

   □ Add the following section:

   ```
   [RPSServer.SyncDbService]
   MyName=RPSServer
   DatabaseRole=SecondaryRole
   PrimaryServerName=Primary
   DatabaseSubrole=RelationalRole
   ```

   This section defines the parameters used for the SDS functions in the secondary RPS database. Note that the **PrimaryServerName** parameter must match the extension used in the [ConnectionParams] section defined later in this document.

8. Specify the connection details for the primary database.

   □ Add the following section:

   ```
   [ConnectionParams.Primary]
   NetworkSpecification1 = TcpIp,Enabled,60001
   ServerNodeSpecifications = TcpIp,127.0.0.1,60001
   RestrictedHostAccess=false
   ```

   These parameters are used to establish connections between the primary and RPS nodes. **Primary** specified after **ConnectionParams** in the section heading must match the **MyName** parameter in the [SyncDbService] section of step 6 above.

9. Update primary database shortcuts to specify the new **ini** file.

   □ Add **ini=C:\<JadeVer>\jadeRPS.ini** to the command line for the following shortcuts (if "ini=" is present already then replace it):

   – JADE Database Server

   – JADE Application Server

   – JADE Presentation Client

   – JADE RPS Manager

   – JADE SDS Administration

When next run, these programs will use the newly created **ini** file instead of the default **ini** file in **C:\<JadeVer>\System\Jade.ini**. Therefore all programs run from the primary and secondary database will use the same **ini** file.

10. Update the primary JADE database server to use the primary **ini** file sections.

   □ Add **name=primary** to the command line for the **JADE Database Server** shortcut

When next run, the primary database server will use the [Primary.SyncDbService] **ini** file section that was defined above, instead of the default [SyncDbService] section.

Create Shortcuts

It is assumed that the shortcuts to run the JADE Database Server, JADE Application Server, JADE Presentation Client, JADE RPS Manager, and JADE SDS Administration are already defined as specified in the "Before You Start" section, earlier in this document.

Define new shortcuts for the following:

- Primary database Erewhon **Administration** application.

   a. Enter the Target value of **C:\<JadeVer>\bin\jade.exe appServer=localHost appServerPort=6006 ini=C:\<JadeVer>\jadeRPS.ini schema=ErewhonInvestmentsViewSchema app=Administration**

   b. Enter the Start In value of **C:\<JadeVer>\bin**

   c. Enter the Name value of **Erewhon Admin**

   This shortcut will run the Erewhon Administration application against the primary database as a presentation client. Note that the **appServerPort** number must be the same as that specified on the JADE application server shortcut. If the JADE application server is on another machine, **appServer=** must specify the relevant DNS name.

- Secondary RPS database server.

   a. Enter the Target value of **C:\<JadeVer>RPS\bin\jadrap.exe path=C:\<JadeVer>RPS\System ini=C:\<JadeVer>\jadeRPS.ini name=RPSServer**

   b. Enter the Start In value of **C:\<JadeVer>RPS\bin**

   c. Enter the Name value of **RPS – JADE Database Server**

   This shortcut will run the JADE database server for the secondary RPS database. Note the **name=** parameter specified in the target. This value is the same as the *unique-id* specified in the **ini** file and it indicates the **ini** file section to be used.

- Secondary database RPS Manager application.

   a. Enter the Target value of **C:\<JadeVer>RPS\bin\jade.exe path=C:\<JadeVer>RPS\System server=multiUser ini=C:\<JadeVer>\jadeRPS.ini schema=JadeMonitorSchema app=RPSManager name=RPSServer**

   b. Enter the Start In value of **C:\<JadeVer>RPS\bin**

   c. Enter the Name value of **RPS – RPS Manager**

This shortcut will run the RPS Manager application against the secondary RPS database as a standard client. Note that the **name=** parameter is also specified in the target for this application, which must be the same as the *unique-id* specified in the **ini** file.

Defining a Relational Database

The RDB used in this document is SQL Server 2005 Express or later release. An ODBC data source must be defined for this and the initial tables, columns, and data must be extracted and loaded from the primary database.

The following instructions demonstrate how to do this.

Create a New ODBC Data Source

This defines a DSN for the SQL Server Express ODBC driver. The RPS secondary JADE database will use this to connect to the SQL Server database.

Note   JADE is supported running on an x64 version of Windows in a 32-bit mode. When creating a DSN on an x64 platform to use with a 32-bit edition of JADE, make sure you use the 32-bit version of the Data Source configuration tool (**c:\windows\syswow64\odbcad32.exe**) to make the DSN accessible to JADE.

1. Start the ODBC Data Source Administrator.

    a. Bring up the **Windows Control Panel**

    b. Select **Administrative Tools**

    c. Select **Data Sources (ODBC)**

    d. Select the **System DSN** sheet

    This opens the dialog to allow a new ODBC data source to be added to your machine.

2. Specify the data source.

    a. Click **Add...**

    b. Select **SQL Native Client**

    c. Click **Finish**

    The data source will be SQL Server. This is the only RDB currently supported by RPS.

3. Specify SQL Server details.

    a. Enter Name of **SQLServerODBC**

    b. Enter Description of **SQL Server**

    c. Enter Server of *machine-name***\SQLExpress**

    d. Click **Next >**

    Note   The name must match the name entered in the **Default Connection String** in the "Define RPS Mapping" section. If SQL Server is running locally, the server name must be specified *without* leading backslashes (\\). Include the backslashes if a DNS name is used.

4. Specify the login details.

   a. Leave the default value of **With Integrated Windows authentication**

   b. Leave the default setting of **Connect to SQL Server to obtain default settings for the additional configuration options**

   c. Click **Next >**

   This defines the login details used when connecting to SQL Server. The Windows login id will be used. When **Next >** is clicked, an attempt is made to connect to SQL Server. If this is unsuccessful, an error will be displayed and the next dialog will not be shown.

   If an error occurs, check that the server name has been entered correctly in the first configuration dialog and that SQL Server is installed and running. If you are still encountering problems, run SQL Server Management Studio. Note the **Server Name** on the logon window, as this should be the same as the server name entered in step 3. Check that you can log on to Management Studio using Windows Authentication. If this logon fails, this problem must be resolved before you can connect using the ODBC data source.

5. Finish defining the data source.

   a. Click **Next >**

   b. Click **Finish**

   This will accept the default values on the last two configuration windows and complete the ODBC setup. An ODBC Microsoft SQL Server Setup dialog will be displayed.

6. Test the Data Source.

   □ Click **Test Data Source**

   This will attempt to connect to SQL Server using the specified configuration. If this fails, check that the server name has been entered correctly in the first configuration dialog and that SQL Server is installed and running. If the test is successful, an SQL Server ODBC Data Source Test dialog will be displayed.

7. Exit from the ODBC Data Source Administrator.

   a. Click **OK** on the SQL Server ODBC Data Source Test dialog

   b. Click **OK** on the ODBC Microsoft SQL Setup dialog

   c. Click **OK** on the ODBC Data Source Administrator dialog with the **SQLServerODBC** data source selected

   d. Close the Administrative Tools window, if required

   SQL Server is now defined as an ODBC data source.

Create Relational Database

From SQL Server, a new RDB is to be created.

1. Start the SQL Server Manager.

   □ From the Start menu, run **SQL Server Management Studio**

   This application is used to manage the SQL Server database.

2. Select to add a new database.

    a. Select **Databases** in the Object Explorer

    b. Right-click and select **New Database**

    The New Database dialog will be shown, to enable the definition of a new database.

3. Add new database.

    a. Enter a database name of **Erewhon**

    b. Click **OK**

    The new database should appear in the Object Explorer under Databases. The database name *must* be the same as that specified in the **Connection String** of the "Define RPS Mapping" section. When adding this database, modify any details you require for your use or change its properties after the database is added.

Extract and Execute SQL Scripts

The following steps create SQL scripts to create the tables and columns required in SQL Server. These tables and columns are those that were defined earlier in the RPS Mapping. Files are also created containing the initial data to be loaded into these tables.

1. Start the servers for the primary and RPS database.

    a. Start the **JADE Database Server**

    b. Start the **RPS – JADE Database Server**

2. Start the RPS Manager application for the RPS database.

    ▫ Start the **RPS – RPS Manager**

    This application is used to create the SQL scripts.

3. Select to generate creation scripts.

    ▫ Select menu item **Extracts | Generate table creation scripts**

    The Generate Creation Scripts dialog will be shown. This is to create a script to load the table definitions into SQL Server.

4. Check the mapping details are correct.

    a. Schema = **ErewhonInvestmentsModelSchema**

    b. Mapping = **ErewhonRPS**

    c. DB Name = **Erewhon**

    These are the details of the JADE schema and the RPS Mapping to be used. If different mapping details are required, you can change them by clicking the **Change** button.

5. Specify tables to be extracted.

    ▫ Select **All**

    This will extract all tables that are specified in the RPS Mapping. Selecting **Selected** would enable a subset of these tables to be used.

6. Enter the script file name.

   ▫ Enter **C:\<JadeVer>RPS\RPSSQLScripts\tableDefn_ErewhonRPS.sql**

   This is the file that will be used to define the tables in SQL Server. The folder is the same as that defined in the "Set Up RPS Folders" section.

7. Create the script.

   ▫ Click **OK**

   The following message will be displayed in the RPS Manager's main window:

   ```
   > Scripts for database creation stored in
         C:\<JadeVer>RPS\RPSSQLScripts\tableDefn_ErewhonRPS.sql
   ```

   The **tableDefn_ErewhonRPS.sql** script contains definitions of the tables to be loaded into SQL Server.

   If you are using **JADE 6.3 or later**, disable **sqlcmd** for script execution. The use of sqlcmd is highly recommended for production systems, but disabling it for this example simplifies the configuration.

   a. Select menu item **File | RPS Node Configuration**

   b. Uncheck **Use Sqlcmd** in the Data Pump Restart Options

   c. Enter **C:\<JadeVer>RPS\RPSSQLScripts** as the **Alter Script Path**

   d. Enter **C:\<JadeVer>RPS\RPSSQLScripts** as the **RDB Path**

   e. Click **OK**

   The configuration options will be applied and the following message will be displayed:

   ```
   Configuration successful
   ```

8. Select to execute the table creation script.

   a. Select menu item **File | Execute SQL Script**

   b. When prompted, select the table creation script file:

   ```
   C:\<JadeVer>RPS\RPSSQLScripts\tableDefn_ErewhonRPS.sql
   ```

   c. Click **Open**

   The script will be executed to create the tables in the relational database. You can use your database's management tool to check that the tables have been created.

9. Select to generate and load the data creation scripts.

   ▫ Select menu item **Extracts | Extract And Load**

   The Data Extract dialog will be shown. This is used to create files containing the data to be loaded into SQL Server and a script to load it.

10. Check the mapping details are correct.

    a. Schema = **ErewhonInvestmentsModelSchema**

    b. Mapping = **ErewhonRPS**

    c. DB Name = **Erewhon**

    These are the details of the JADE schema and the RPS Mapping to be used. If different details are required, you can change them by clicking the **Change** button.

11. Specify the tables to be extracted.

   ▫ Select **All**

   This will extract data for all tables that are specified in the RPS Mapping. Selecting **Selected** would enable data from a subset of these tables to be extracted. This should match the selection from the Generate Creation Scripts dialog.

12. Enter the script file path and data files path.

   ▫ Enter **C:\<JadeVer>RPS\RPSSQLScripts** for both paths

   The script file to be used for loading the data will be created in the specified Script File Path folder. The data files will be created in the specified Data Files Path folder.

13. Enter the RDB path.

   a. Ensure that **Server (BULK INSERT)** is selected (from JADE 6.2 only)

   b. Enter an RDB path of **C:\<JadeVer>RPS\RPSSQLScripts**

   This is the folder where the data files are located, from the point of view of the RDB machine. As the RDB and the JADE RPS node are on the same machine, the same folder is used.

   The RDB Path must be visible to SQL Server. If SQL Server is on a different machine to the JADE RPS node, specify an RDB path relative to the RDB machine. If the location is not visible, the resulting data files will have to be manually moved to the RDB path before loading the SQL Scripts.

   Alternatively, the **Client (bcp)** option can be used from JADE 6.2.

14. Create and run the scripts.

   ▫ Click **OK**

   Details of the data extracted will be displayed in the RPS Manager's main window.

   In the **Data Files Path** folder, a ***.fmt** and ***.tbl** file will be created for each table and these will contain the data to be loaded for each table.

   In the **Script File Path** folder, the **BulkInsert.sql** script will be created. This is the script that will be used to load the table data.

   The scripts are run to load the data into the relational database.

15. Use the SQL Server Manager to check that the data has been loaded.

   a. From the Start menu, run **SQL Server Management Studio Express**

   b. Select the **Erewhon** database

   c. Select menu item **File | New | Query with Current Connection**

   d. Enter the following query:

   ```
   Select * from AddressableEntity
   ```

   This query will output all of the data in the **AddressableEntity** table.

16. Run the new query.

   ▫ Press **F5**

   All data from the **AddressableEntity** table will be displayed.

Operation

This section describes starting RPS, testing RPS, and schema changes.

Starting RPS

All components should now be installed and defined. Start the following (if they are not already running) to see RPS in action.

Run the Primary JADE Database Components

1. Start the primary database server.

    □   Start the **JADE Database Server** (if not already running)

    The primary database server is started to enable the connection to the RPS database and so that changes can be made and propagated to the SQL Server database.

2. Start the SDS Administration application.

    □   Start the **JADE SDS Administration** application

    The SDS Administration application will be used to view the connection between the primary and RPS servers and to help diagnose any problems that might occur. Note that the primary database should currently be listed in this window.

Run the JADE RPS Database Components

1. Start the RPS database server.

    □   Start the **RPS – JADE Database Server** (if not already running)

    The RPS database server is started so that changes made to the primary database can be processed and propagated to the SQL Server database.

    The RPS database will now be shown in the JADE SDS Administration window and will be highlighted with an orange background. This indicates that database tracking is halted. Tracking is halted because we have not yet started the data pump.

2. Start the RPS Manager application.

    □   Start the **RPS - RPS Manager** (if not already running)

    The RPS Manager is used to start the data pump later in following section.

Enable the RPS Data Pump

The **Datapump** application must be running to update the SQL Server database with changes that have been applied to the RPS database.

Perform the following from the RPS - RPS Manager application.

1. Select to configure the RPS node.

    □   Select the menu item **File | RPS Node Configuration**

    The Configure RPS Node dialog will be displayed.

2. Check the configuration details.

    a.   Schema = **ErewhonInvestmentsModelSchema**

b. Mapping = **ErewhonRPS**

c. Jade Database Mode = **Full**

d. Relational Database Name = **Erewhon**

e. Connection String = **DSN=SQLServerODBC; Database=Erewhon**

f. Alter Script Path = **C:\<JadeVer>RPS\RPSSQLScripts**

These details should be pre-filled and match those details entered during the "Define RPS Mapping" section.

3. Unset **Ask for RDB Login** (prior to JADE 6.3).

▢ Uncheck this option

This option controls whether the Relational Database Connection dialog is displayed when you attempt to connect to SQL Server. As **Windows Authentication** was specified for SQL Server login (in the "Create a New ODBC Data Source" section), the RDB Login is not required.

4. Set the **Auto Start Datapump** option.

▢ Enable **Auto Start** by checking this box

The data pump will be started manually in the following steps. However, it is generally a good idea to enable **Auto Start**. With this option selected, the data pump will be started automatically when the RPS node is started. It also allows operations such as database reorganizations to be performed without user intervention. Note that when **Auto Start** is enabled, the **Ask for RDB Login** option (prior to JADE 6.3) must be unchecked.

5. Accept the RPS configuration details.

▢ Click **OK**

From JADE 6.2 only, if you are asked to enter an **RDB Path Name** or an **RDB Server Name**:

a. Select the **Server** option for the **Load Execution Location**

b. Enter **C:\<JadeVer>RPS\RPSSQLScripts** as the RDB **Path Name**

c. Click **OK**

The configuration options will be applied and the following message will be displayed:

```
Configuration successful
```

6. Start the data pump.

▢ Select the menu item **RPS | Start Datapump...**

The data pump will be started. Progress messages will be displayed in the RPS Manager's main window.

7. Check that the data pump is running.

a. Go to the **SDS Administration** application

b. Check that the **RPSServer** database is highlighted in green

This indicates that the data pump is running and that the secondary RPS database is synchronized. To determine what all the background colors mean, select **Options | Show Legend**.

If you have any problems starting the data pump, please refer to the "Troubleshooting" section, later in this document.

Testing RPS

RPS should now be up and running, and any changes made to the primary database should be sent (via journals) to the RPS database and then (via the data pump) to the relational database. Test that this is functioning as expected.

Look at Data in the Relational Database

1. Go to the RDB.

   a. Go to **SQL Server Management Studio Express**

   b. Select the **Erewhon** database

   We will now look at data in this database.

2. Add a query.

   a. Select menu item **File |New | Query with Current Connection**

   b. Enter the following query:

   ```
   SELECT shortDescription, fullDescription
   FROM  SaleItem
   WHERE (shortDescription = 'Lamp')
   ```

   This query will display details of a single object that we will later update from the primary database.

3. Run the query.

   □ Press **F5**

   This will return the following details:

   ```
   Lamp, Purchase this rare stained glass piece from the early 14th
   century.
   ```

   As we will re-run this query later to show changes, keep this window open.

Make a Change to the Primary Database

1. Start the user application in the primary database.

   a. Start the **JADE Application Server**

   b. Start the **Erewhon Administration**

   c. Log on as **Erewhon Investments Inc**

   Data changes will be made from the Erewhon Administration application. The application server is required as Erewhon Administration is run as a presentation client.

2. Select an object to change.

   a. Select menu item **View | Sale Items (by Category)**

   b. Expand the **Antiques | Items For Sale** branch

   c. Select the **ANQU2001 Lamp: $250.00 – Unsold** item

    d.   Click **Edit**

Details of the object will be displayed.

3.   Change some details on the object.

    a.   Add some extra text to the **Full Description**

    b.   Click **OK**

The object will be updated with the new description. This change made on the primary database should have been replicated automatically to the RDB.

Check the Change is Replicated to the Relational Database

1.   Go to the RDB.

    ▫   Go to **SQL Server Management Studio Express**

We will now look at the changed data in this database.

2.   Query the data that has changed

    ▫   Re-run the previous query:

```
SELECT shortDescription, fullDescription
FROM  SaleItem
WHERE (shortDescription = 'Lamp')
```

The query result should include the changes that you made to the Lamp description in the primary database (above). This shows that the change made in the primary database has been replicated to the SQL Server database.

Schema Changes

Structural changes made to the primary database may require changes to the table definitions in the RDB. Some changes will be automatically sent and applied to the RDB, but others require manual table redefinition.

For more details, see "RPS Considerations" in the *Schema Evolution and Instantiation* White Paper.

Appendix

This appendix contains useful references, troubleshooting, and tips.

Useful References

For further information, please refer to the following documents:

- JADE *Administration Guide*, Chapter 9, "Relational Population Service (RPS)" for JADE 6.1 or Chapter 2 of the JADE *Synchronized Database Service (SDS) Administration Guide* for JADE 6.2

- JADE *User's Guide,* Chapter 16, "Using the Relational Population Service (RPS) Wizard" for JADE 6.1 or JADE *Development Environment User's Guide,* Chapter 16, "Using the Relational Population Service (RPS) Wizard" for JADE 6.2

- JADE *Monitor Utility User's Guide*, "Using the SDS Administration Application" for JADE 6.1 or Chapter 1 of the JADE *Synchronized Database Service (SDS) Administration Guide* for JADE 6.2

- JADE *Monitor Utility User's Guide*, "Using the RPS Manager Application" for JADE 6.1 or Chapter 2 of the JADE *Synchronized Database Service (SDS) Administration Guide* for JADE 6.2

- White Paper, *Schema Evolution and Instantiation,* "RPS Considerations"

JADE documentation can be downloaded from www.jadeworld.com/jade/doc.

White Papers can be downloaded from www.jadeworld.com/jade/whitepapers.

Troubleshooting

This section is intended to address common problems that might be encountered when setting up and running with RPS. Follow these steps when problems are encountered.

- Check that all components are running.
  - □ JADE Primary Server
  - □ JADE RPS Server
  - □ RPS Data Pump application
  - □ SQL Server
- Run the SDS Administration application and verify that the primary and secondary servers (that is, databases) are connected.
- Check the logs for errors.

jommsg.log

This log contains information about what is occurring on each database, including successful operations as well as errors. It is usually the first log that you should check when problems are encountered.

SDS and RPS Work Logs

The primary SDS work log is named **sds_<primary server name>_<yyyymmdd>** and contains detailed diagnostic and error information logged by SDS for the primary database.

The RPS work log is named **rps_<rps server name>_<yyyymmdd>** and contains detailed diagnostics and error information logged by the SDS or RPS functions executing in the RPS node.

Application Logs

An application log will be created if errors occur for any of the applications. These will contain details of any exceptions that occur in the relevant applications.

Common Errors

3201 – SDS not initialized

> This error normally means that your primary database is not configured as a primary or that the SDS service failed to initialize. Check the **jommsg.log** for the primary. Also check that the [SyncDbService] and [ConnectionParams.Primary] sections are included in your **ini** file.

RPS instance ID mismatch with Target Relational DB

> The instance ID held by the RPS database does not match the instance id of the target RDB. This id is generated when you create or recreate an RPS database. If the RPS database is recreated, you should also recreate the RDB.

Tips

This section provides the following tips about using RPS.

- Keep your first RPS configuration simple by running all JADE and SQL Server databases on the same machine.

- Make your RPS node easily identifiable from your primary JADE database. Install the RPS database within a separately named directory structure. For example, locate your primary environment in **C:\Jade63** and your RPS environment in **C:\Jade63RPS**.

- Keep your first test mapping small and simple.

- When creating an initial RPS database, various flavours of RPS are available (working set, mapped extent, and full replica). While all modes will work, it is easiest to start with the full replica mode using a small primary database.

- To evaluate RPS with SQL Server, you can download SQL Server Express and SQL Server Management Studio Express free-of-charge from http://msdn.microsoft.com/vstudio/express/sql/download/. Even with certain size and connection limitations, SQL Server Express will provide a solid, full-featured platform for evaluation of RPS without cost or evaluation time limits.

- If you want to recreate your relational database completely, the simplest approach is to drop and recreate the database, and then reload a full table definition script before bulk-loading the data. Dropping and creating a database can be scripted and you can use the Management Studio (or equivalent) to generate the required scripts.

- JADE is supported running on an x64 version of Windows in a 32-bit mode. When creating a DSN on an x64 platform to use with a 32-bit edition of JADE, make sure you use the 32-bit version of the Data Source configuration tool (**c:\windows\syswow64\odbcad32.exe**) to make the DSN accessible to JADE.

- If a **Working Set** RPS database is being used, the RPS node **ini** file setting for **OpenAllFilesOnDbOpen** in the [PersistentDB] section should be set to **False**. This is because not all database files are created in a **Working Set** RPS database.

Glossary

Database Role
In a Synchronized Database Environment, the terms *primary* and *secondary* signify the intended function or role of each copy of the database. The valid roles are **Primary Database Role** and **Secondary Database Role**.

Full Replica Mode
The mode of operation whereby an RPS node retains a full copy of the primary database, regardless of the mapped extent. In this mode, tables in the relational database can be repopulated without impacting the primary database. An RPS node running in full replica mode can also be used for disaster recovery, as it can acquire the primary database role in a takeover situation.

JADE System Model
A single JADE specification or object model, potentially spanning multiple JADE schemas.

Mapped Extent
The set of objects in the production database that are included by all the ClassTable maps in a database map, and are not excluded by filter conditions.

Mapped Extent Mode
An intermediate mode between **Working Set** and **Full Replica** whereby an RPS node retains a copy of the mapped extent, even after user data objects have been replicated to the relational database. In this mode, tables in the relational database can be repopulated without impacting the primary database.

Nodes and Processes
A primary system is comprised of primary nodes. A secondary system is comprised of secondary nodes. Secondary processes execute in a secondary system. In addition, special processes called *replay processes* are created in secondary servers to reflect and represent the original updating processes on the primary database. These replay processes are associated with a single secondary node called a *replay node.*

Primary Database
The database with the **Primary** role is the only database in a Synchronized Database Environment (SDE) where user applications can update persistent objects.

Primary Database Server
The primary database server is the database server node hosting the JADE database with the primary role; that is, this server hosts the primary database.

Primary System
The primary system is the term used to describe the JADE system comprising nodes attached to the primary database server, including the primary database server node itself. User applications executing in the primary system can create and update persistent objects, which reside in the primary database.

Production Database
The source JADE database being replicated (also known as the primary database).

[Relational] AttributeColumn Map
Part of a ClassTable map. This defines how JADE attributes map to relational attributes (table columns).

[Relational] ClassTable Map
Defines how instances of a JADE class are mapped to rows in one or more relational tables.

[Relational] Database Map
Defines a mapping between a JADE system model and a single relational database.

Relational Population Service (RPS)
The service that provides automatic replication of objects from a primary JADE database to one or more relational databases.

RDB
Relational Database

RDBMS
Relational Database Management System

RPS Data Store
A native JADE database maintained by an RPS server node, either as temporary storage or to persist a mapped extent.

RPS Node
The JADE database server node that directly connects to a target RDBMS in order to replicate JADE objects to a relational database.

Secondary Database Role
Secondary databases are kept synchronized with a primary database by replaying journals automatically shipped from the primary system. User applications cannot update persistent objects (apart from environmental objects) in a secondary database. A secondary database can assume the role of the primary if the primary database becomes unavailable.

Secondary Database Server
A secondary database server is a database server node that hosts a database with a secondary database role.

Synchronized Database Environment (SDE)
A Synchronized Database Environment is the collective term that describes an environment comprising of one or more connected JADE systems that provide or subscribe to a Synchronized Database Service. An SDE consists of a primary database (where updating occurs) and one or more secondary databases, which are transactionally consistent copies of the primary database. Includes both SDS and RPS nodes.

Synchronized Database Service (SDS)
The service that enables you to keep one or more secondary databases synchronized with a primary database.

Target Database
An instance of a relational database connected to an RPS node.

Working Set Mode
The mode of operation whereby an RPS node retains user data objects only until they have been replicated to the relational database. This mode minimizes the RPS node's disk storage requirements. However, repopulating one or more tables in the relational database will require the RPS node to retrieve the necessary objects again from the primary database.