IBM Software Group

IBM® Rational® Application Developer V7.5

Java™ Persistence API tools

@business on demand.

This presentation describes the Java Persistence API (JPA) and the tools available in Rational Application Developer V7.5 to do JPA development.

Agenda

- Overview of the Java Persistence API (JPA)

- JPA development tools
  - ▶ Projects
  - ▶ General development tools
  - ▶ Working with databases

2

The first section of this presentation provides an overview of the Java Persistence API and the JPA tools that are available in the Rational Application Developer workbench. The rest of the presentation covers these development tools in more detail, including how to work with JPA projects, general JPA development tools; like how to create entities from tables and using the new specialized views in the workbench. Finally, how to configure and connect to databases for application testing.

Section

Overview

3

This section includes an overview of the Java Persistence API and JPA development tools in Rational Application Developer.

JPA overview

- JPA is a standard persistence and object-relational mapping (ORM) framework for Java
  - ▶ Part of the EJB 3.0 specification (JSR 220)
  - ▶ Provided by the javax.persistence package
  - ▶ Enables persisting plain-old Java objects (POJOs) to a relational database
    - No deployment code or abstract classes
    - No interfaces required
  - ▶ Metadata is specified with Java annotations or in XML deployment descriptors
- An application server is not required
  - ▶ JPA is usable in both Java SE and Java EE environments

4

The Java Persistence API is a standard framework that provides persistence and object-relational mapping as a part of the EJB 3.0 specification. It is a Java standard that provides many of the benefits of alternative persistence frameworks, with the added benefit of portability to any EJB 3.0 compliant container, without having to install any extra libraries. It allows you to persist plain-old Java objects to a relational database – a much simpler approach than using container-managed persistence (CMP) in EJB 2.1. In JPA, Entities are plain-old Java objects, just like other components in EJB 3.0. An entity typically represents a table in a relational database, and each instance of an entity is a row. Entities are concrete classes, not abstract classes like Entity Beans in EJB 2.1. You do not need to generate deployment code, which speeds deployment, and you do not have to implement any particular interfaces. Another benefit is that all object-relational mapping information is specified in a standard fashion, using Java annotations or XML files. In earlier versions of the EJB specification, there was not a standard way to provide this information, which meant each vendor's implementation was different, and led to a reliance on vendor-specific tools. JPA can also be used without an EJB container, that is, in a Java Standard Edition (SE) environment.
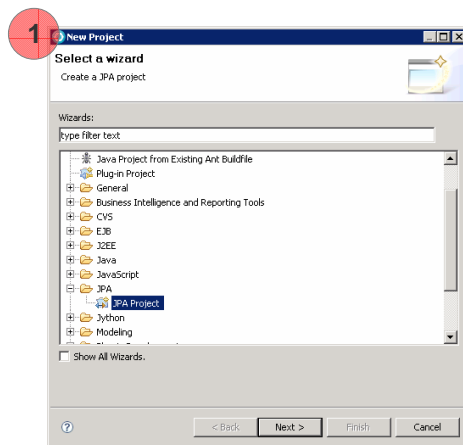
JPA development tools overview

- New JPA projects and JPA views
  - ▸ Add JPA support to existing projects by enabling the JPA facet

- Tools for working with JPA annotations
  - ▸ Java editor provides validation, content assist and quick fixes for JPA annotations
  - ▸ Modify annotation properties in JPA details view or annotations view

- JPA XML editors (persistence.xml, orm.xml)

- Wizards for creating and initializing object-relational mappings
  - ▸ Create tables from entity classes (top-down)
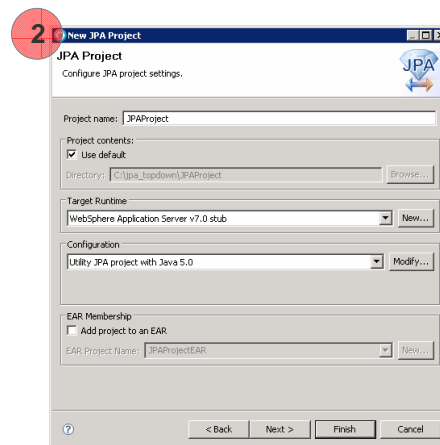 - ▸ Create new entity beans from existing database tables (bottom-up)

With extensible frameworks and tools, you can perform many tasks that are associated with accessing data for a multi-tier enterprise application. You can work with JPA properties in either the JPA Details view or the Annotations view, so that you don't need to keep both views open at once. For clarity, the Annotations view distinguishes between implied and specified annotation attributes. Like in Java EE 5 application development, deployment descriptors in JPA applications are optional – the programming model for the JPA relies on using annotations to specify information about how the Java object will map to a database table. However, deployment descriptors can still be used in a JPA application, and the workbench provides deployment descriptor editors to work with those files. In order to minimize the complexity of mapping between JPA entities and tables, you can use wizards to create and automate initial mappings. You can also receive programming assistance from the tools through dynamic problem identification.

Section

JPA projects

Java Persistence API tools

This section describes how to create and work with JPA projects.

Start the JPA project wizard by selecting **File > New > Project > JPA > JPA Project**

Name the project and choose the **Target Runtime**

While the JPA is a part of the Java EE 5 specification, you can also create stand-alone JPA projects. You can do this either using the JPA project creation wizard or by modifying an existing Java project. In the project creation wizard, you need to choose a project name and specify which target runtime to use for testing.

JPA projects – new project wizard (continued)
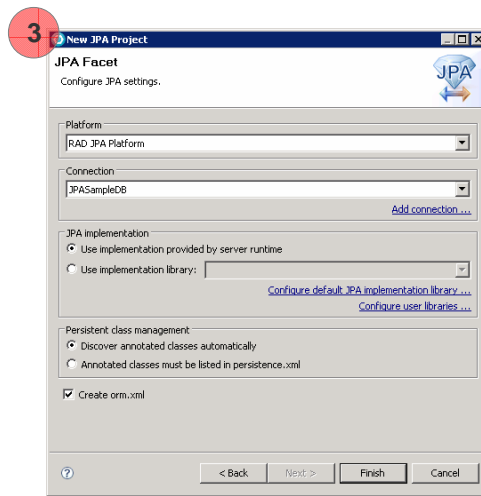
- Configure additional JPA settings on the JPA facet panel
- Add a connection to a database
- Choose which implementation library to use

After choosing the project name and runtime, you can configure additional options associated with the JPA facet. This is the same panel that you see if you go in and edit the project properties later. From here, choose which JPA platform to use. You also have the option of adding a connection to a database. If you do not create a database connection at project creation time, you need to do so later, before you can test your application. You have the option of choosing a specific JPA implementation library to use with your project. If you are using a WebSphere® Application Server V7 runtime and choose to use the implementation provided by the server, then you are using an implementation based on Apache OpenJPA. Deployment descriptors are optional in a JPA project, but you can choose to create the persistence.xml and orm.xml files at project creation time, if needed.

JPA projects – Java project conversion
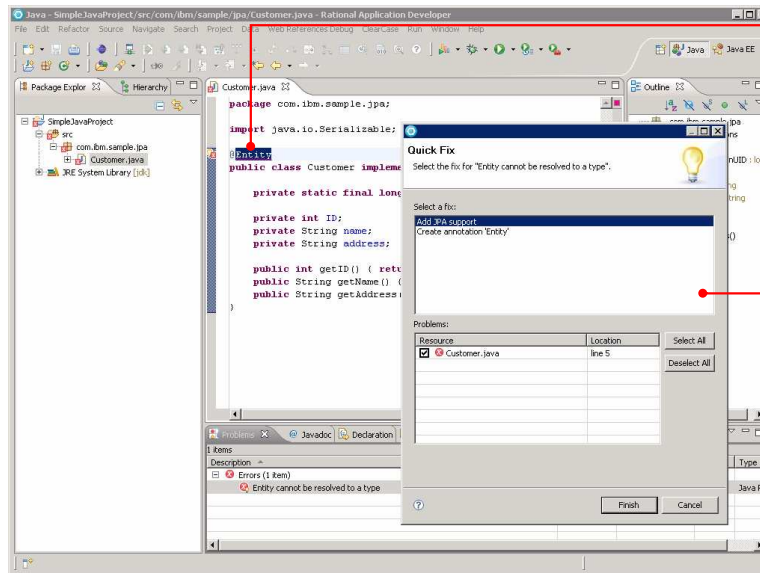
Use JPA in a Java project by adding an appropriate annotation to a class – for example, @Entity

Use the **Quick Fix** menu to select the option to **Add JPA support** to the project

You can also convert a Java project into a JPA project, using Quick Fixes that are available in the Java editor. When you type a JPA annotation into a plain old Java object (POJO), the resulting compilation errors have Quick Fixes for adding the required JPA facet and an appropriate runtime environment (if available) to your project. For example, adding the @Entity annotation to the beginning of a Java class indicates that you want to store that class's fields to a database. The development environment automatically recognizes this as a JPA annotation and includes the appropriate fixes in the Quick Fix menus.

JPA development perspective

Display Java classes, work with deployment descriptors in the main panel

The **Package Explorer** view shows project structure

Work with data connections in the **Data Source Explorer** view

JPA Structure view provides a tree display of the data in a JPA class

Highlight a JPA construct in the JPA Structure view to see configuration details in the **JPA Details** view

Java Persistence API tools

© 2008 IBM Corporation

10

The JPA Development perspective in the Rational Application Developer workbench includes database-related tools and specialized JPA views to assist in JPA application development and testing. The JPA Structure view provides a tree display of the data in a JPA class. Highlight a JPA construct in the JPA structure view to see the configuration details for that item in the JPA Details view. This view provides detailed information on the attributes of a JPA entity, including the name of the table it's being mapped to and all of the attributes of the mapping. Many of the JPA attributes do not need to be specified in your application and have context-appropriate default values; you can see these default values in the JPA Details view. The JPA Development perspective also includes the Data Source Explorer, which you can use to work with data source connections for application testing.
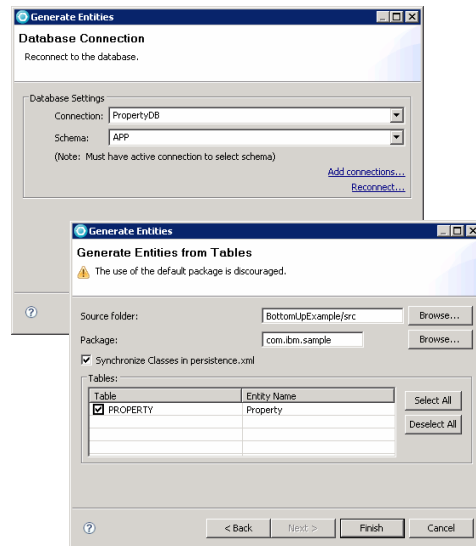
Section

General JPA development tools

11

Java Persistence API tools   © 2008 IBM Corporation

This section describes general JPA development tools available in Rational Application Developer V7.5.

You can generate JPA entity beans from existing database tables; this is called bottom-up mapping. In order to create an entity, you need to configure a database connection for your project and that database needs to contain a table to use to generate the entity. If you have not configured a database connection for your project, you can do so on the Database Connection panel after launching the Generate Entities process for your project. From there, either choose the appropriate existing connection or configure a new connection. Then, select the table to use to create the entity. The development tools will automatically generate a class based on the table structure, and the class will include the appropriate data members and accessor methods. After generating the entity, create a primary key element in the class by using either the Configure JPA Entities tools or by adding the @ID annotation manually to the primary key field in the class.

Tools for annotations

The **Annotations View** provides a graphical interface for working with JPA annotations
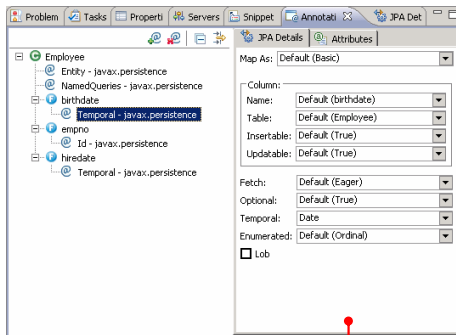
In the **JPA Details** tab, configure how data is stored and fetched in the database
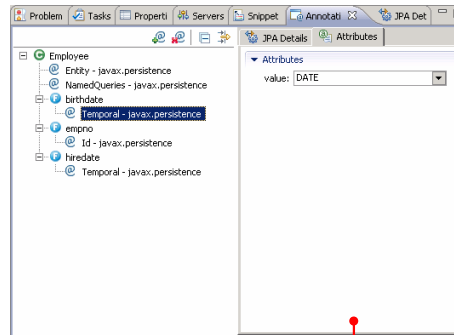
Annotation attributes describe the annotation's properties; configure these in the **Attributes** tab

In the JPA Development perspective, the JPA Details view is integrated with the Annotations view. From this view, you can browse through all of the JPA annotations in your application and configure their attributes. As with other Java EE 5 constructs, in the JPA, annotations take the place of information that was previously provided only in deployment descriptors, and the programming model makes assumptions about what the default attributes should be. For example, if you do not specify a table name to store your JPA entity, the programming model assumes that the class name is the same as the table name. You can see and override these annotation attributes from these views.

JPA details display – mappings

- The **JPA Details** display is available as a separate view and as a part of the **Annotations** view

- Configure how data is persisted
  - ▸ Type of mapping and mapping parameters
    - **Examples:** transient, one-to-many, embedded, many-to-one
    - See documentation for details on available mapping types
  - ▸ Column and table in the database
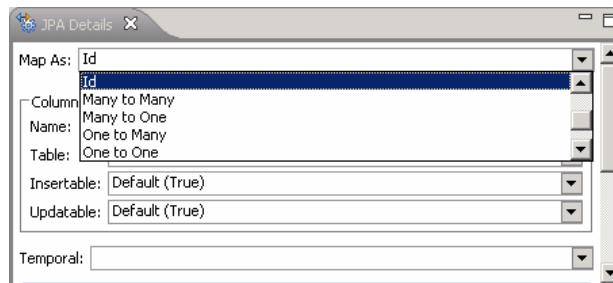
The JPA details view also provides a dropdown menu to configure JPA mapping types and parameters. You can specify the column and table in the database that holds the information and the way that the mapping will happen. For example, you can choose which element is the primary key, if a mapping is one-to-one, one-to-many, or a variety of other mapping types. All of the available types are listed in the mapping menu, and you can read more about them in the product documentation.

Creating a JPA manager
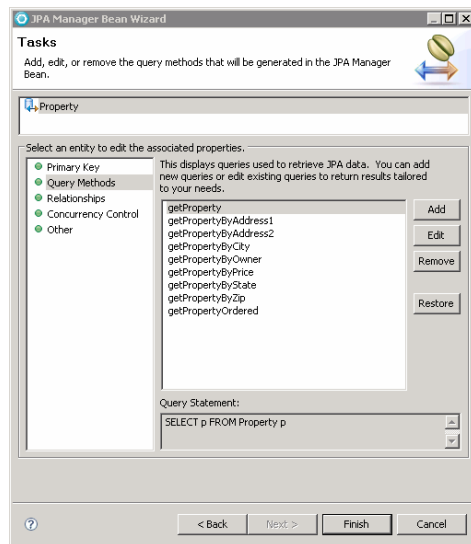
- Right-click a JPA project and select **JPA Tools > Add JPA Manager Beans** to open the wizard
 - ▸ Choose the entities to associate with the manager
  - ▸ Configure tasks for the manager
  - ▸ Configure deployment options

JPA Manager Beans are service beans that act as façades or controllers over a particular JPA entity. They encapsulate and abstract all of the data access code for creating, updating, deleting, and displaying information from your database using JPA entities. JPA Manager Beans are an ideal programming model for use in two-tier Web environments. They fill the role that would normally be filled by a session bean in an EJB environment; all of the business logic related to an entity is performed by the JPA Manager Bean. The use of JPA Manager Beans is not limited to Web applications. They can be used anywhere that you want to take advantage of their data abstraction capabilities, such as an EJB project, a JPA Utility project, or even a plain Java project. Note that if you want to use JPA inside of an EJB project, you might instead want to create an EJB Session Bean to contain all of your JPA logic in order to get the benefits that the EJB container provides. You can automatically generate Java Persistence API (JPA) manager bean classes for JPA entities using the Add JPA Manager Beans wizard. This wizard creates the manager bean and lets you configure which entities to associate with it and what tasks you want the manager to perform.

XML configuration files

- The persistence.xml file defines the database and entity manager options for deploying JPA applications
  - ▸ Packaged in the META-INF directory of your project
  - ▸ Synchronize changes in your Java classes by right-clicking the file and choosing **JPA Tools > Synchronize Classes**
- The orm.xml file stores metadata to describe object-relational mappings
  - ▸ This file is not required – the JPA specification emphasizes using annotations to describe mappings
  - ▸ Mapping information defined in orm.xml overrides default JPA behavior and any mappings defined using annotations
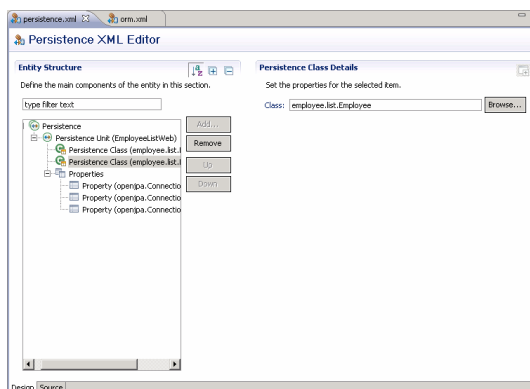
The persistence.xml file describes the details of the persistence units in your JPA project. A persistence unit contains a list of entity beans. You can edit the persistence.xml file with the Persistence XML Editor. The Persistence XML editor simplifies making changes to the persistence.xml file like adding or removing persistence units, or modifying the properties of a persistence unit. After you create persistent classes, you can automatically add them to list of classes in their persistence unit in the persistence.xml file. To automatically add all persistent classes in your JPA project to the persistence.xml file, right-click one of the Java classes and choose the Synchronize Classes option. The persistent classes in your JPA project are automatically discovered and added to the persistence unit in the persistence.xml file.

When you specify object-relational mappings using the JPA tools, the mapping information is contained in the Java class files in the form of Java annotations. However, you can also choose to define the object-relational mappings in XML in a file called orm.xml. Mapping information defined in the orm.xml file automatically overrides both default JPA behavior and any mappings defined using annotations. Therefore, you can, for example, to adapt existing JPA entity beans to a different set of database tables without needing to modify the entity class files. You can use the Object Relational Mapping XML Editor to define object-relational mappings for JPA entity beans in the orm.xml file.

The persistence.xml file and orm.xml file are located in the META-INF directory of the JPA project.
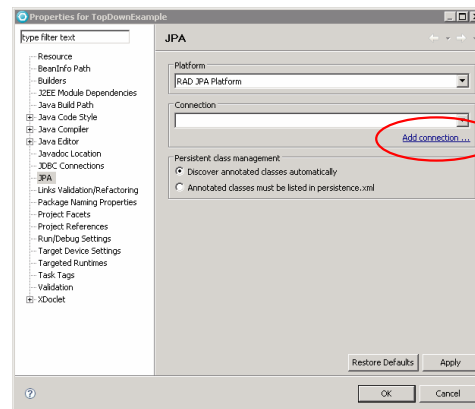
Section

Working with databases

17

This section describes some additional tools for working with databases that can be useful in developing Java Persistence API applications.

Connecting to a database using properties

- You can add a database connection to your JPA project using project properties
  - ▸ Right-click the project and select **Properties > JPA**
  - ▸ Click the text **Add connection…**
  - ▸ Select the connection profile type (for example, Derby)
  - ▸ Choose a name for the connection profile
  - ▸ Choose the database driver to use
  - ▸ Specify other database properties
    - Directory, user name, password

Connection Profile wizard

There are multiple ways to access the database connection wizard from the workbench. One way is to create a database connection while you are creating your JPA project. Another way, described here, is to use the JPA project properties to start the wizard. A connection profile contains the connection property information needed to connect to a database runtime instance. You need to choose the type of database being used (for example, Derby or DB2®), a name for the profile, the database driver you are using, and any other required database parameters to create a connection – like the location of the database and user name and password information if the database is secured.

Using the data source explorer

- Data source explorer is a part of the JPA Development perspective
  - Open the explorer under **Window > Show View > Other > Connectivity > Data Source Explorer**

- Launch the connection profile wizard using **New Connection Profile** icon

- Edit a data table by right-clicking on the table in the **Data Source Explorer** view and choosing **Data > Edit**
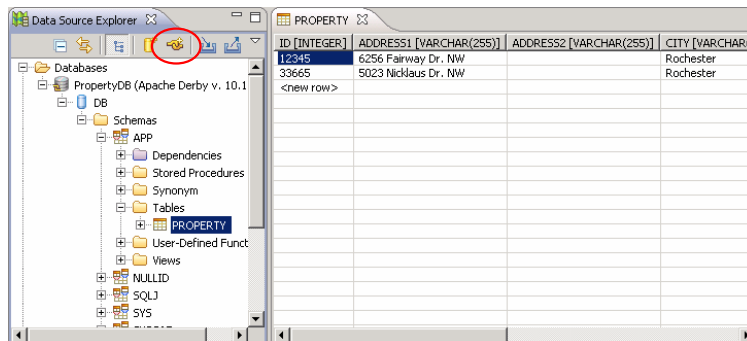
The data source explorer is integrated into the JPA Development perspective, and you can also open it manually under the Window > Show View menu. From the data source explorer, you can create new database connections (using the wizard described on the previous page) and define and edit database tables. To edit a table, right-click it and choose the option to edit the table. The database table entries appear in the table on the right, and you can manually edit them. You can use this editor to populate test data for working with your applications.

Creating a database table from an entity

- Use JPA entities to create database tables top-down
  - In the Package Explorer view, right-click the JPA project and select **JPA Tools > Generate DDL...**
  - The JPA tools create a file named Table.ddl in the META-INF directory of your JPA project
  - Double-click Table.ddl to open it
  - Configure the connection properties using the dropdown menus at the top of the editor
  - Right-click in the editor and select **Execute All**
 - Check the SQL Results view to see if the table generation was successful
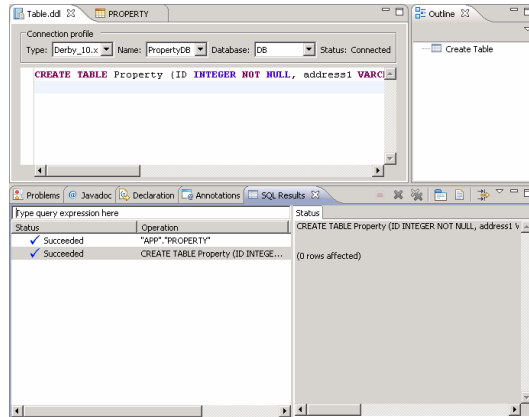
Using the JPA tools, you can generate data definition language files for creating database tables from entity beans that you create. In top-down mapping, you start with entity beans and use them to create your database tables. You start from scratch with the entity definitions and the object-relational mappings, and then you derive database schemas from that data. If you use this approach, you are most likely concerned with creating the architecture of your object model and then writing your entity classes. These entity classes eventually drive the creation of your database model. If you are using a top-down mapping of the object model to the relational model, develop the entity classes, and then use the JPA tools DDL generation capability to create the database tables that are based on the entity classes. The JPA Tools in your project can automatically generate a database definition language, or DDL, file based on one of the entities in your project. Once you have the DDL file, run it from the workbench to create the table.

Defining an application data source

In the Project Explorer, right-click an EAR project and select **Open WebSphere Application Server Deployment**

You can configure other server deployment options further down in the display: J2C options, applications to deploy, shared libraries, and more
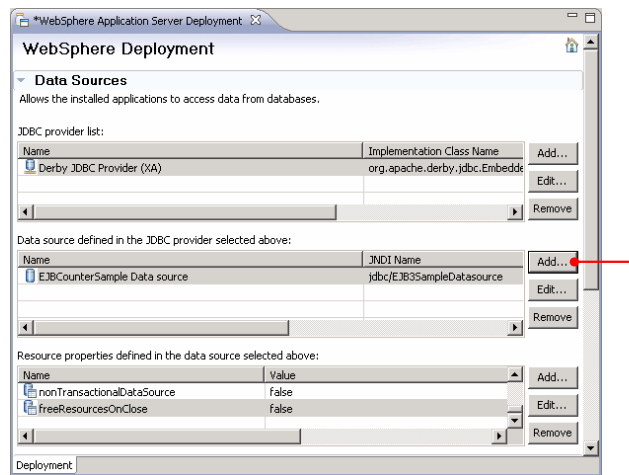
Use wizards to define the data source

IBM Software Group

Java Persistence API tools   © 2008 IBM Corporation

21

If you want to test your application on an application server, like WebSphere Application Server, then you need to define a data source on that server. This assumes that you are working in an EAR project and have a server connection configured for WebSphere Application Server. Right-click the EAR project and choose the Open WebSphere Application Server Deployment option to bring up the WebSphere Deployment editor. From here, you can configure several server deployment options, including data source definitions.

Section

Summary

22

This section contains a summary of the presentation.

Summary

- Rational Application Developer V7.5 provides tools for JPA development
  - ▸ Specialized JPA projects and views
  - ▸ Tools for working with annotations
  - ▸ Top-down and bottom-up mappings for entities and database tables
  - ▸ Wizards for configuring data sources

23

The Java Persistence API is a new programming model for persisting Java objects to database; it is part of the Java EE 5 specification. The Rational Application Developer V7.5 workbench provides tools for doing JPA development, including specialized JPA projects and views, tools for working with annotations, the ability to do top-down and bottom-up mapping of entities and database tables, and wizards for configuring data sources for application testing and deployment.

Feedback

Your feedback is valuable

You can help improve the quality of IBM Education Assistant content to better meet your needs by providing feedback.

- Did you find this module useful?
- Did it help you solve a problem or answer a question?
- Do you have suggestions for improvements?

Click to send e-mail feedback:

mailto:iea@us.ibm.com?subject=Feedback_about_WASv7_RADJPATools.ppt

This module is also available in PDF format at: ../WASv7_RADJPATools.pdf

24

You can help improve the quality of IBM Education Assistant content by providing feedback.

Trademarks, copyrights, and disclaimers

IBM, the IBM logo, ibm.com, and the following terms are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both:

DB2   Rational   WebSphere

If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of other IBM trademarks is available on the Web at "Copyright and trademark information" at http://www.ibm.com/legal/copytrade.shtml

Rational is a trademark of International Business Machines Corporation and Rational Software Corporation in the United States, Other Countries, or both.

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

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

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

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

IBM makes no representations or warranties, express or implied, regarding non-IBM products and services.

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

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

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

© Copyright International Business Machines Corporation 2008. All rights reserved.

Note to U.S. Government Users - Documentation related to restricted rights-Use, duplication or disclosure is subject to restrictions set forth in GSA ADP Schedule Contract and IBM Corp.