IBM® WebSphere® Application Server V6.1 Feature Pack for EJB 3.0

Java™ persistence API (JPA) overview

This presentation will cover the Java persistence API support in the WebSphere Application Server V6.1 Feature Pack for EJB 3.0.

Agenda

- JPA introduction

- JPA terms

- Summary and references

This presentation will give a high-level overview of JPA, and explain the core JPA terminology.

Section

JPA introduction

This section will give a high-level overview of JPA.

JPA motivations

- Persistence is a challenge in EJB 2.1
  - Complex programming model
  - Entity beans considered too heavyweight
  - EJB-QL considered too limited

- Other approaches have emerged in response
  - Custom frameworks or plain JDBC
  - Object-relational mapping (ORM) solutions
    - Toplink (now EclipseLink)
    - Hibernate
    - Others

Over the years, persisting data with EJBs has come to be a sore spot for many developers. Entity beans can be complicated and time consuming to write, even with the support of vendor-supplied tools. The heavyweight programming model can be less than ideal for many applications, particularly smaller scale applications. Many Java developers use custom persistence frameworks or plain JDBC to better meet their needs. As a result, several alternatives have emerged, offering lightweight and straightforward solutions for persisting Java objects to a relational database. EclipseLink and Hibernate are two popular solutions. The EJB 3.0 specification introduced JPA in response to these challenges, offering a lightweight, Plain-old Java object (POJO) based persistence framework.

JPA overview

- JPA is a standard persistence and object-relational mapping (ORM) framework for Java
  - ▸ Part of the EJB 3.0 specification (JSR220)
  - ▸ Provided by the javax.persistence package
  - ▸ Enables persisting plain-old Java objects (POJOs) to a relational database

- IBM's implementation of JPA is based on Apache OpenJPA

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. The JPA implementation in WebSphere Application Server is based on the open source Apache OpenJPA project.

JPA features

- Standardized O/R mapping metadata (unlike EJB 2.1)
  - ▶ Metadata is specified with Java annotations or in XML deployment descriptors

- Entities are POJOs
  - ▶ No deployment code or abstract classes
  - ▶ No interfaces required

- An application server is not required
  - ▶ JPA is usable in both Java SE and Java EE environments

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 features (continued)

- Disconnected entities

- Entity life cycle listener and callback support

- Key generation and composite keys

- Unidirectional and bidirectional relationships
  - 1-to-1, 1-to-many, many-to-many

- Several retrieval methods
  - Find by primary key
  - JPQL and native SQL queries

JPA provides many features that you would expect in a persistence framework. You can manipulate entities in a disconnected fashion, and then write updates back to the database when you are done working with them. Entity life cycle listeners and callback methods give you the ability to take action when life cycle events occur, and are optional, just as they are for session beans and message-driven beans in EJB 3.0. There is robust support for key generation, including several different key generation methods, and composite key support. Complex relationship support is also available, enabling unidirectional or bidirectional relationships. 1-to-1, 1-to-many, and many-to-many relationships are supported. Several methods are available for finding entities in a relational database. A standard "find by primary key" method is available, and you can also query the database using either JPQL or native SQL. JPQL is the Java Persistence Query Language, the successor to EJB-QL.

Section

JPA terms

This section will introduce some basic JPA concepts and terms.

Entity

- An *entity* is a plain-old Java object (POJO) that can be managed by a JPA persistence provider
  - ▸ Metadata about the entity is specified using Java annotations or in a deployment descriptor

- Entities are used to persist Java objects to a storage device, such as a relational database

- Entity life cycle is managed by the application
  - ▸ The life cycle of EJB 1 and 2 *entity beans* was managed by the container

9

A JPA Entity is a plain-old Java object that can be persisted to a relational database. The "@Entity" annotation marks a POJO class as an Entity, and enables it to be managed by a persistence provider. It is important to note that JPA Entities are not the same as EJB 1 or 2 "Entity Beans", despite the similar nomenclature. Entity life cycle is managed by an application, and entity beans are managed by the container.

IBM

Persistence unit

- A *persistence unit* defines the scope of entity persistence

- META-INF/persistence.xml defines persistence units, including:
  - ▸ Persistence provider to use
  - ▸ Data sources to use
  - ▸ Object relationship mapping information
  - ▸ Entities accessible for this persistence unit
  - ▸ Locations of the entity classes

- Represents the configuration of an EntityManagerFactory

A persistence unit defines the scope for entity persistence. You define a persistence unit in the persistence.xml file, and specify what entity classes are part of the persistence unit, which persistence provider and data source to use, and other properties about mapping your objects to the database. An EntityManagerFactory uses a persistence unit as its configuration. There is a 1:1 relationship between entityManagerFactories and persistence units.

Persistence context

- A persistence context is the set of entities that are participating in a given unit of work
  - ▸ These entities are managed by an EntityManager

The set of entities that are participating in a given unit of work are known as a "persistence context". Each group of entities in a persistence context is managed by an entity manager.

Entity manager

- An *entity manager* stores and loads entity objects to and from persistent storage

- The EntityManager API enables an application to
  - Persist or remove entities to or from a database
  - Find entities using a primary key or using queries

- Container-managed entity managers
  - Obtained by JNDI lookup or dependency injection
  - More common in Java EE environments

- Application-managed entity managers
  - Obtained from an EntityManagerFactory
  - More common in Java SE or application clients

An entity manager manages the task of persisting and loading objects in a relational database. You can use the EntityManager API to perform JPA tasks such as adding or updating entities to a database, removing entities from a database, or using query methods to find and load particular entities from a database. Entity managers can be either container-managed or application-managed. Most applications in Java EE environments will use container-managed entity managers, because it is the easiest way to use an entity manager. You can get an entity manager either by injection or by JDNI lookup. Since you do not have a container in a Java SE environment, you would use an application-managed entity manager, which you would obtain from an entity manager factory object. The presentation module titled "JPA code examples" will show how to use an entity manager in your application.

IBM

JPQL

- Java Persistence Query Language (JPQL) is an extension of Enterprise Java Beans Query Language (EJB-QL)

- Adds several new features
  - Bulk update and delete
  - *Join* functions
  - *Group by* and *having* functions
  - Dynamic queries and named parameters

Java Persistence Query Language, or JPQL, is the latest revision of EJB-QL, the EJB query language. It introduces several new SQL-like features on top of the capabilities of EJB-QL, including bulk update and delete methods, "join" functions, "group by" and "having" functions, and support for dynamic queries and named query parameters.

Section

Summary and references

This section will summarize the presentation content.

IBM

Summary

- JPA is a standard persistence and object-relational mapping (ORM) framework for Java
  - ▶ Part of the EJB 3.0 specification (JSR220)
  - ▶ Provided by the javax.persistence package
  - ▶ Enables persisting plain-old Java objects (POJOs) to a relational database

- JPA addresses many of the challenges of object persistence in previous versions of the EJB specification

15

JPA overview   © 2007 IBM Corporation

The Java Persistence API is a standard persistence and object-relational mapping framework that is part of the EJB 3.0 specification. It uses plain-old Java objects to represent items in a relational database. It is designed to be a persistence framework for Java EE that improves on prior EJB persistence models by addressing their inherent challenges, by being less complex and lighter-weight.

Additional resources

- EJB 3.0 specification
  - http://java.sun.com/products/ejb/docs.html

- Apache OpenJPA
  - http://openjpa.apache.org/

This slide lists some external resources that you may find useful for learning about JPA. You should also consult the WebSphere Application Server Feature Pack for EJB 3.0 Information Center.

Feedback

Your feedback is valuable

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

- Did you find this module useful?

- Did it help you solve a problem or answer a question?

- Do you have suggestions for improvements?

Click to send e-mail feedback:

mailto:iea@us.ibm.com?subject=Feedback_about_WASv61_EJB3FP_JPAOverview.ppt

This module is also available in PDF format at: ../WASv61_EJB3FP_JPAOverview.pdf

17

JPA overview

© 2007 IBM Corporation

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

Trademarks, copyrights, and disclaimers

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

IBM   WebSphere

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

Product data has been reviewed for accuracy as of the date of initial publication. Product data is subject to change without notice. This document could include technical inaccuracies or typographical errors. IBM may make improvements or changes in the products or programs described herein at any time without notice.

Information is provided "AS IS" without warranty of any kind. THE INFORMATION PROVIDED IN THIS DOCUMENT IS DISTRIBUTED "AS IS" WITHOUT ANY WARRANTY, EITHER EXPRESS OR IMPLIED. IBM EXPRESSLY DISCLAIMS ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NONINFRINGEMENT. IBM shall have no responsibility to update this information. IBM products are warranted, if at all, according to the terms and conditions of the agreements (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.

© Copyright International Business Machines Corporation 2007. All rights reserved.

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