IBM Software Group

IBM® WebSphere® Application Server Feature Pack for EJB 3.0

Java™ Persistence API (JPA) code examples

@business on demand.

This presentation will introduce the basic concepts of the Java Persistence API by presenting several code examples.

Agenda

- Entities
- Data retrieval
- Other features

This presentation will begin with examples of how to create and use JPA entities. Data retrieval and queries will be discussed next, followed by some more advanced features of JPA.

Entity definition

EJB 2.1

```
public abstract class CustomerBean implements
    Customer, EntityBean {
 public CustomerBean() { }
 public abstract String getName();
 public abstract void setName(String n);
 public abstract int getAmountSpent();
 public abstract void setAmountSpent(int
    amount);
private EntityContext ctx;

 public String ejbCreateByName(String) throws
    EJBException   { }
 public void setEntityContext (EntityContext
    theCtx) throws EJBException {
  ctx = theCtx; }
 public void unsetEntityContext() throws
    EJBException {
  ctx = null; }
}
```

EJB 3.0

```
@Entity @Table (name="CUSTS")
public class Customer implements Serializable {
 public Customer() { }
 @Id
 public String getName() { return name; }
 public void setName(String n) { name = n; }
 public int getAmountSpent() { return
    amountSpent; }
 public void setAmountSpent(int a) {
    amountSpent = a; }
 private String name;
 private int amountSpent;
}
```

3

This first example compares the definition of an EJB 2.1 entity bean with a JPA Entity in EJB 3.0. The most immediately apparent difference is that unlike an EJB 2.1 entity bean, a JPA entity is a concrete, not abstract, class. You will also notice that the JPA entity does not need to include the boilerplate code, shown in red, that is included at the bottom of the entity bean example. The JPA entity is a plain Java class, with some annotations added to provide metadata and identify the class as an entity. The "@Entity" annotation indicates that the class is a JPA entity. The "@Table" annotation is optional, and specifies the name of the database table that should be used to store instances of this entity. If a table name is not specified, it is implied that the table has the same name as the class. A column in that table will be used for each member variable, and the "@Id" annotation specifies which column should be used as the primary key. Beyond the simplified code in the JPA example, it does not require any additional information to be specified in an XML deployment descriptor.

IBM

Entity creation

EJB 2.1

```
<begin transaction>
Object obj = Context.lookup(
    "java:comp/env/ejb/CustomerHome");
CustomerHome ch = (CustomerHome)
    PortableRemoteObject.narrow(obj,
        CustomerHome.class);

Customer cust1 = ch.create("Joe King");
cust1.setAmountSpent(5);
<end transaction>
```

EJB 3.0

```
<begin transaction>
@PersistenceContext
EntityManager em;

Customer cust1 = new Customer("Joe King");
cust1.setAmountSpent(5);
em.persist(cust1);
<end transaction>
```

4

Entities are created using the "new" command, just like any other Java object. You do not have to deal with the container and perform a lookup, like you are required to do with EJB 2.1. The "@PersistenceContext" annotation is used to inject an instance of an entity manager, which you can then use to interact with the database. In this example, a new customer object is created, and then added to the database using the "persist()" method. For qualities of service, like transaction support or security, you should use the façade model and wrap the entity with a session bean.

Find and update

EJB 2.1

```
<begin transaction>
Object obj = Context.lookup(
    "java:comp/env/ejb/CustomerHome");
CustomerHome ch = (CustomerHome)
    PortableRemoteObject.narrow(obj,
        CustomerHome.class);

Customer cust1 =
    ch.findByPrimaryKey("Joe King");
cust1.setAmountSpent(5);
<end transaction>
```

EJB 3.0

```
<begin transaction>
@PersistenceContext
EntityManager em;

Customer cust1 = em.get(Customer.class,
    "Joe King");
cust1.setAmountSpent(5);

<end transaction>
```

JPA code examples

5

© 2007 IBM Corporation

This example shows a connected update, using the entity manager's "get()" method to find and retrieve the Customer object using its primary key. Again, it is much simpler to do this with the JPA entity than with the EJB 2.1 entity bean. Disconnected updates are also available using the merge() operation to optimistically update entities that have been detached.

Named queries

```
@Entity
@NamedQuery(name="getCustomerByName",
    query="select a from Customer a where a.name = :name")
public class Customer {
...
```

} Define a query

```
@PersistenceContext
EntityManager em;

Query q = em.createNamedQuery("getCustomerByName");
q.setParameter("name", "Joe King");
Customer cust1 = (Customer) q.getSingleResult();
cust1.setAmountSpent(5);
```

} Then use it

In this example, a JPQL (Java Persistence Query Language) query is defined with one variable, "name". In the lower code block, the "name" parameter is set, and then a customer object is retrieved using the getSingleResult() method on the query. The Customer object can then be acted on just like the Customer object that was found in the previous example.

In EJB 2.1, you would use EJB-QL to define the query in the deployment descriptor, and then use the bean's home to call the query that you defined in the deployment descriptor to return a Customer object.

JPQL supports both named parameters and ordered parameters, whereas EJB-QL only supports ordered parameters. In addition to defining the query with an annotation, as shown here, a JPQL query can also be defined in an XML file. Defining queries in XML can be a time saver, so that you do not have to recompile your application to modify a database query.

Native queries

```
orm.xml
…
<named-native-query name="getCustAvg"
    query="select sum(age)/count(*) as avg from CUSTTBL
    where last_purchase_date > :date" />
…
```

Define a query

```
@PersistenceContext
EntityManager em;

Query q = em.createNativeQuery("getCustAvg");
Integer custAvg = (Integer) q.getSingleResult();
```

Then use it

This example shows the definition of a native query in an XML file, and the invocation of that query from Java code. Calling a native query is very similar to using a JPQL query, as was shown previously. Like other query types, native queries can be defined in either Java code or XML files. It may be advantageous to define native queries in XML files so that your application code is not tied to any particular underlying database product. This approach also enables you to fine tune the SQL in your query without having to recompile your Java code.

IBM

Disconnected data

getCustomer()

<begin transaction>
@PersistenceContext
EntityManager em;

Customer cust1 = em.get(Customer.class, "Joe King");
return cust1;
<end transaction>

updateCustomer(Customer c)

<begin transaction>
@PersistenceContext
EntityManager em;

em.merge(c);
<end transaction>

JPA code examples   © 2007 IBM Corporation

Another feature of JPA is that an object can be returned from the database, and then be manipulated outside the scope of a transaction. When the modified object is to be committed to the database, use the entity manager's merge() method to optimistically update the database. The entity, in this case a Customer object, must define a version control field for optimistic updating. The version control field is denoted using the @Version annotation on a variable.

Relationships

- One to one, one to many, and many to many relationships are supported

- Can easily be defined in code
  - Can also be in deployment descriptor, as in EJB 2.1

```
@Entity
public class Order {
    @ManyToOne
    @JoinColumn(name = "ACCOUNT_ACCOUNTID", referencedColumnName =
    "ACCOUNTID")
    private Account acct;

    public Account getAccount(); { return acct; }
    public void setAccount(Account a) { acct = a; }
}
```

JPA supports several types of relationships, including bidirectional ones. Relationships can be one-to-one, one-to-many, or many-to-many. You can use a deployment descriptor to define relationships, as was required with EJB 2.1, or you can use Java annotations, as is shown in this example. This example shows a many-to-one relationship between Order and Account entities. The "@JoinColumn" annotation specifies the mapping between database columns in the different tables. Cascading for related entities can be specified for various operations, such as deleting, and lazy semantics are supported.

Inheritance

```
@Entity
@Table(name="EMP")
@Inheritance(strategy=JOINED)
public abstract class Employee {
    @Id protected Integer empId;
    @Version protected Integer version;
    @ManyToOne protected Address address;
}
```

```
@Entity
@Table(name="FT_EMP")
@DiscriminatorValue("FT")
@PrimaryKeyJoinColumn(name="FT_EMPID")
public class FullTimeEmployee extends Employee {
    // Inherit empId, but mapped in this class to FT_EMP.FT_EMPID
    // Inherit version mapped to EMP.VERSION
    // Inherit address mapped to EMP.ADDRESS foreign key
    // Defaults to FT_EMP.SALARY
    protected Integer salary;
}
```

JPA code examples   © 2007 IBM Corporation

Inheritance is also supported for JPA entities. This example shows a relationship between a base class called "employee" in the first block of code, and the second block of code shows the "FullTimeEmployee" class.

The base class, employee, is mapped to a database table named "EMP" and defines a few basic variables including a primary key and a version identifier. The FullTimeEmployee class extends Employee, adding a salary variable. FullTimeEmployee objects are mapped to the "FT_EMP" table. The empId primary key is inherited from Employee, but for FullTimeEmployees it is stored in a column called "FT_EMPID", as is specified with the "@PrimaryKeyJoinColumn" annotation. As with other JPA features, inheritance can also be specified using XML deployment descriptors.

Section

Summary and references

This section will summarize the presentation.

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

12

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. The code examples in this presentation have shown how to create and use entities, how to query entity data using the entity manager, and how to define relationships and inheritance between entity classes.

Additional resources

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

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

This slide lists some resources that may be helpful for learning more about JPA. The JPA specification is a sub-set of the EJB 3.0 specification.

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_JPAExamples.ppt

This module is also available in PDF format at: ../WASv61_EJB3FP_JPAExamples.pdf

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, 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. Any statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only. References in this document to IBM products, programs, or services does not imply that IBM intends to make such products, programs or services available in all countries in which IBM operates or does business. Any reference to an IBM Program Product in this document is not intended to state or imply that only that program product may be used. Any functionally equivalent program, that does not infringe IBM's intellectual property rights, may be used instead.

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