

Oracle® WebLogic Server

Programming WebLogic Enterprise JavaBeans, Version 3.0
10g Release 3 (10.3)

August 2008

ORACLE®

Copyright © 2007, 2008, Oracle and/or its affiliates. All rights reserved.

This software and related documentation are provided under a license agreement containing restrictions on use and disclosure and are protected by intellectual property laws. Except as expressly permitted in your license agreement or allowed by law, you may not use, copy, reproduce, translate, broadcast, modify, license, transmit, distribute, exhibit, perform, publish, or display any part, in any form, or by any means. Reverse engineering, disassembly, or decompilation of this software, unless required by law for interoperability, is prohibited.

The information contained herein is subject to change without notice and is not warranted to be error-free. If you find any errors, please report them to us in writing.

If this software or related documentation is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, the following notice is applicable:

U.S. GOVERNMENT RIGHTS Programs, software, databases, and related documentation and technical data delivered to U.S. Government customers are "commercial computer software" or "commercial technical data" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, the use, duplication, disclosure, modification, and adaptation shall be subject to the restrictions and license terms set forth in the applicable Government contract, and, to the extent applicable by the terms of the Government contract, the additional rights set forth in FAR 52.227-19, Commercial Computer Software License (December 2007). Oracle USA, Inc., 500 Oracle Parkway, Redwood City, CA 94065.

This software is developed for general use in a variety of information management applications. It is not developed or intended for use in any inherently dangerous applications, including applications which may create a risk of personal injury. If you use this software in dangerous applications, then you shall be responsible to take all appropriate fail-safe, backup, redundancy, and other measures to ensure the safe use of this software. Oracle Corporation and its affiliates disclaim any liability for any damages caused by use of this software in dangerous applications.

Oracle is a registered trademark of Oracle Corporation and/or its affiliates. Other names may be trademarks of their respective owners.

This software and documentation may provide access to or information on content, products and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services.

Contents

1. Introduction and Roadmap

Document Scope and Audience	1-1
Guide to this Document	1-2
Related Documentation	1-2
Comprehensive Example for the EJB 3.0 Developer	1-3
New and Changed Features in this Release	1-4

2. Understanding Enterprise JavaBeans 3.0

Understanding EJB 3.0: New Features and Changes From EJB 2.X	2-1
Changes in the EJB Programming Model and Requirements Between Versions 2.X and 3.0	2-2
New EJB 3.0 Features.	2-3
WebLogic Server Value-Added EJB 3.0 Features.	2-3
EJB 3.0 Examples	2-4
Programming 3.0 Entities	2-5

3. Simple Enterprise JavaBeans 3.0 Examples

Example of a Simple Stateless EJB	3-1
Example of a Simple Stateful EJB	3-3
Example of an Interceptor Class	3-7
Example of Invoking a 3.0 Entity From A Session Bean	3-8

4. Iterative Development of Enterprise JavaBeans 3.0

Overview of the EJB 3.0 Development Process	4-1
Create a Source Directory	4-3
Program the EJB 3.0 Business Interface	4-4
Business Interface Application Exceptions	4-4
Using Generics in EJBs	4-5
Serializing and Deserializing Business Objects	4-6
Program the Annotated EJB Class	4-6
Optionally Program Interceptors	4-7
Compile Java Source	4-7
Optionally Create and Edit Deployment Descriptors	4-7
Packaging EJBs	4-8
Deploying EJBs	4-9

5. Programming the Annotated EJB 3.0 Class

Overview of Metadata Annotations and EJB 3.0 Bean Files	5-1
Programming the Bean File: Requirements and Changes From 2.X	5-2
Bean Class Requirements and Changes From 2.X	5-2
Bean Class Method Requirements	5-3
Programming the Bean File: Typical Steps	5-3
Specifying the Business and Other Interfaces	5-5
Specifying the Bean Type (Stateless, Stateful, Message-Driven)	5-6
Injecting Resource Dependency into a Variable or Setter Method	5-7
Invoking a 3.0 Entity	5-8
Specifying Interceptors for Business Methods or Life Cycle Callback Events	5-12
Programming Application Exceptions	5-18
Securing Access to the EJB	5-19
Specifying Transaction Management and Attributes	5-22

Complete List of Metadata Annotations By Function	5-22
---	------

6. Using Oracle Kodo with WebLogic Server

Overview of Kodo	6-1
Creating a Kodo Application	6-2
Using Different Kodo Versions	6-2
Configuring Persistence	6-2
Editing the Configuration Property Files	6-3
Using the Configuration Files Together	6-4
Configuring Plug-ins	6-4
Deploying a Kodo Application	6-4
Configuring a Kodo Application	6-5
Using the Administration Console	6-5
Configuring Kodo Without Using the Administration Console	6-5

A. EJB 3.0 Metadata Annotations Reference

Overview of EJB 3.0 Annotations	A-1
Annotations for Stateless, Stateful, and Message-Driven Beans	A-2
javax.ejb.ActivationConfigProperty	A-3
javax.ejb.ApplicationException	A-3
javax.ejb.EJB	A-4
javax.ejb.EJBs	A-5
javax.ejb.Init	A-6
javax.ejb.Local	A-7
javax.ejb.LocalHome	A-8
javax.ejb.MessageDriven	A-9
javax.ejb.PostActivate	A-10
javax.ejb.PrePassivate	A-11

javax.ejb.Remote	A-12
javax.ejb.RemoteHome	A-12
javax.ejb.Remove	A-13
javax.ejb.Stateful	A-14
javax.ejb.Stateless	A-15
javax.ejb.Timeout	A-16
javax.ejb.TransactionAttribute	A-16
javax.ejb.TransactionManagement	A-17
Annotations Used to Configure Interceptors	A-18
javax.interceptor.AroundInvoke	A-18
javax.interceptor.ExcludeClassInterceptors	A-19
javax.interceptor.ExcludeDefaultInterceptors	A-19
javax.interceptor.Interceptors	A-19
Annotations Used to Interact With Entity Beans	A-20
javax.persistence.PersistenceContext	A-20
javax.persistence.PersistenceContexts	A-22
javax.persistence.PersistenceUnit	A-24
javax.persistence.PersistenceUnits	A-25
Standard JDK Annotations Used By EJB 3.0	A-26
javax.annotation.PostConstruct	A-26
javax.annotation.PreDestroy	A-27
javax.annotation.Resource	A-27
javax.annotation.Resources	A-29
Standard Security-Related JDK Annotations Used by EJB 3.0	A-30
javax.annotation.security.DeclareRoles	A-30
javax.annotation.security.DenyAll	A-31
javax.annotation.security.PermitAll	A-31
javax.annotation.security.RolesAllowed	A-31

javax.annotation.security.RunAs	A-32
WebLogic Kodo Annotations	A-32
weblogic.javaee.AllowRemoveDuringTrasaction	A-33
weblogic.javaee.CallByReference	A-33
weblogic.javaee.DisableWarnings	A-34
weblogic.javaee.EJBReference	A-35
weblogic.javaee.Idempotent	A-35
weblogic.javaee.JMSClientID	A-36
weblogic.javaee.JNDIName	A-37
weblogic.javaee.MessageDestinationConfiguration	A-38
weblogic.javaee.TransactionIsolation	A-39
weblogic.javaee.TransactionTimeoutSeconds	A-39

B. Persistence Configuration Schema Reference

persistence-configuration.xml Namespace Declaration and Schema Location.	B-1
persistence-configuration.xml Deployment Descriptor File Structure	B-2
persistence-configuration.xml Deployment Descriptor Elements.	B-6
Example	B-185

Introduction and Roadmap

This section describes the contents and organization of this guide—*Programming WebLogic Enterprise Java Beans, Version 3.0*.

- [“Document Scope and Audience”](#) on page 1-1
- [“Guide to this Document”](#) on page 1-2
- [“Related Documentation”](#) on page 1-2
- [“Comprehensive Example for the EJB 3.0 Developer”](#) on page 1-3
- [“New and Changed Features in this Release”](#) on page 1-4

Document Scope and Audience

This document is a resource for software developers who develop applications that include WebLogic Server Enterprise Java Beans (EJBs), Version 3.0.

The document mostly discusses the new EJB 3.0 programming model, in particular the use of metadata annotations to simplify development. The document briefly discusses the main differences between EJB 3.0 and 2.X for users who are familiar with programming EJB 2.X and want to know why they might want to use the new 3.0 programming model.

This document does not address EJB topics that are the same between versions 2.X and 3.0, such as design considerations, EJB container architecture, deployment descriptor use, and so on. This document also does not address production phase administration, monitoring, or performance

tuning. For links to WebLogic Server documentation and resources for these topics, see [“Related Documentation” on page 1-2](#).

It is assumed that the reader is familiar with Java Platform, Enterprise Edition (Java EE) Version 5 and EJB 2.X concepts.

Guide to this Document

- This chapter, [Chapter 1, “Introduction and Roadmap,”](#) introduces the organization of this guide.
- [Chapter 2, “Understanding Enterprise JavaBeans 3.0,”](#) provides an overview of the new EJB 3.0 features and programming model, as well as a brief description of the differences between EJB 3.0 and 2.X.
- [Chapter 3, “Simple Enterprise JavaBeans 3.0 Examples,”](#) provides simple examples of programming EJBs using the new metadata annotations specified by EJB 3.0.
- [Chapter 4, “Iterative Development of Enterprise JavaBeans 3.0,”](#) describes the EJB implementation process, and provides guidance for how to get an EJB up and running in WebLogic Server.
- [Chapter 5, “Programming the Annotated EJB 3.0 Class,”](#) describes the requirements and typical steps when programming the EJB bean class that contains the metadata annotations.
- [Chapter 6, “Using Oracle Kodo with WebLogic Server,”](#) describes how to use Oracle Kodo to create entity beans. Oracle Kodo is a product that provides the implementation of the Java Persistence API section of the EJB 3.0 specification, as well as other persistence-related technologies such as Java Data Objects (JDO).
- [Appendix A, “EJB 3.0 Metadata Annotations Reference,”](#) provides reference information for the EJB 3.0 metadata annotations, as well as information about standard metadata annotations that are used by EJB.
- [Appendix B, “Persistence Configuration Schema Reference,”](#) provides reference information for the persistence configuration schema.

Related Documentation

This document contains EJB 3.0-specific development information. Additionally, it provides information only for session and message-driven beans. For completed information on general

EJB design and architecture, the EJB 2.X programming model (which is fully supported in EJB 3.0), and programming 3.0 entities, see the following documents:

- [Programming WebLogic Enterprise JavaBeans \(Version 2.X\)](#)
- [Enterprise JavaBeans 3.0 Specification \(JSR-220\)](#)

For comprehensive guidelines for developing, deploying, and monitoring WebLogic Server applications, see the following documents:

- *Developing Applications with WebLogic Server* is a guide to developing WebLogic Server applications.
- *Deploying Applications to WebLogic Server* is the primary source of information about deploying WebLogic Server applications in development and production environments.

Comprehensive Example for the EJB 3.0 Developer

In addition to this document and the basic examples described in [Chapter 3, “Simple Enterprise JavaBeans 3.0 Examples,”](#) Oracle provides a comprehensive example in the WebLogic Server distribution kit. The example illustrates EJB 3.0 in action and provides practical instructions on how to perform key EJB 3.0 development tasks. In particular, the example demonstrates usage of EJB 3.0 with:

- Java Persistence API
- Stateless Session Bean
- Message Driven Bean
- Asynchronous JavaScript based browser application.

The example uses a persistent domain model for entity EJBs.

WebLogic Server optionally installs this comprehensive example in `WL_HOME\samples\server\examples\src\examples\ejb\ejb30`, where `WL_HOME` is the top-level directory of your WebLogic Server installation. On Windows, you can start the examples server, and obtain information about the samples and how to run them from the WebLogic Server Start menu.

Oracle recommends that you run this example before programming your own application that uses EJB 3.0.

New and Changed Features in this Release

For a comprehensive listing of the new WebLogic Server features introduced in this release, see [What's New in WebLogic Server](#) in the *Release Notes*.

For a list of all known and resolved issues in this release, see [Weblogic Server Known and Resolved Issues](#) in *Known and Resolved Issues*.

Understanding Enterprise JavaBeans 3.0

These sections describe the new features and programming model of EJB 3.0.

It is assumed the reader is familiar with Java programming, Java Platform, Enterprise Edition (Java EE) Version 5, and EJB 2.x concepts and features.

- [“Understanding EJB 3.0: New Features and Changes From EJB 2.X” on page 2-1](#)
- [“WebLogic Server Value-Added EJB 3.0 Features” on page 2-3](#)
- [“EJB 3.0 Examples” on page 2-4](#)
- [“Programming 3.0 Entities” on page 2-5](#)

Understanding EJB 3.0: New Features and Changes From EJB 2.X

Enterprise JavaBeans (EJB) is a Java Platform, Enterprise Edition (Java EE) Version 5 technology for the development and deployment of component-based business applications. Although EJB is a powerful and useful technology, the programming model in version 2.X and previous was complex and confusing, requiring the creation of multiple Java files and deployment descriptors for even the simplest of EJB. This complexity hindered the wide adoption of EJBs.

As a consequence, one of the central goals of Version 3.0 of the EJB specification is to make it much easier to program an EJB, in particular by reducing the number of required programming

artifacts and introducing a set of EJB-specific metadata annotations that make programming the bean file easier and more intuitive.

Another goal of the EJB 3.0 specification was to standardize the persistence framework and reduce the complexity of the entity bean programming model and object-relational (O/R) mapping model.

Note: This document does not discuss programming 3.0 entity beans; that information is provided in the [Java Persistence API](#) guide of the [Oracle Kodo documentation](#) set.

The remainder of this section describes, at a high-level, how the programming model and requirements changed in EJB 3.0, as compared to version 2.X, and lists a brief description of the the new features of EJB 3.0.

Changes in the EJB Programming Model and Requirements Between Versions 2.X and 3.0

The following summarizes the changes between EJB 2.X and 3.0:

- You are no longer required to create the EJB deployment descriptor files (such as `ejb-jar.xml`). You can now use metadata annotations in the bean file itself to configure metadata. You are still allowed, however, to use XML deployment descriptors if you want; in the case of conflicts, the deployment descriptor value overrides the annotation value.
- The bean file can be a plain old Java object (or POJO); it is no longer required to implement `javax.ejb.SessionBean` or `javax.ejb.MessageDrivenBean`.
- As a result of not having to implement `javax.ejb.SessionBean` or `javax.ejb.MessageDrivenBean`, the bean file no longer has to implement the lifecycle callback methods, such as `ejbCreate`, `ejbPassivate`, and so on. If, however, you want to implement these callback methods, you can name them anything you want and then annotate them with the appropriate annotation, such as `@javax.ejb.PostActivate`.
- The bean file is required to use a business interface. The bean file can either explicitly implement the business interface or it can specify it using the `@javax.ejb.Remote` or `@javax.ejb.Local` annotations.)
- The business interface is a plain old Java interface (or POJI); it should not extend `javax.ejb.EJBObject` or `javax.ejb.EJBLocalObject`.
- The business interface methods may not throw `java.rmi.RemoteException` unless the business interface extends `java.rmi.Remote`.

Because the EJB 3.0 programming model is so simple, Oracle no longer supports using the EJBGen tags and code-generating tool on EJB 3.0 beans. Rather, you can use this tool *only* on 2.X beans. For information, see [EJBGen Reference](#).

New EJB 3.0 Features

- Bean files can now use metadata annotations to configure metadata, thus eliminating the need for deployment descriptors.
- The only required metadata annotation in your bean file is the one that specifies the type of EJB you are writing (`@javax.ejb.Stateless`, `@javax.ejb.Stateful`, `@javax.ejb.MessageDriven`, or `@javax.persistence.Entity`). The default value for all other annotations reflect typical and standard use of EJBs. This reduces the amount of code in your bean file in the case where you are programming a typical EJB; you only need to use additional annotations if the default values do not suit your needs.
- Bean files supports dependency injection. *Dependency injection* is when the EJB container automatically supplies (or *injects*) a variable or setter method in the bean file with a reference to another EJB or resource or another environment entry in the bean's context.
- Bean files support interceptors, which is a standard way of using aspect-oriented programming with EJB.
- You can configure two types of interceptor methods: those that intercept business methods and those that intercept lifecycle callbacks.
- You can configure multiple interceptor methods that execute in a chain in a particular order.
- You can configure default interceptor methods that execute for all EJBs contained in a JAR file.

WebLogic Server Value-Added EJB 3.0 Features

The following features are not part of the Enterprise JavaBeans 3.0 specification, but rather, are value-added features to further simplify the EJB 3.0 programming model:

- Automatic persistence unit deployment based on variable name.

If you want to query and update an entity in your session bean, you annotate a variable with either the `@javax.persistence.PersistenceContext` or `@javax.persistence.PersistenceUnit` annotation; the variable is then injected with persistence unit information. You can specify the `unitName` attribute to reference a

particular persistence unit in the `persistence.xml` file of the EJB JAR file; in this case, the EJB container automatically deploys the persistence unit and sets its JNDI name to the persistence unit name, as listed in the `persistence.xml` file. You can also simply name the injected variable exactly the same as the persistence unit you want injected into the variable; in this case, you no longer need to specify the `unitName` attribute, but the EJB container still deploys the persistence unit and sets its JNDI name to the persistence unit name.

See [“Invoking a 3.0 Entity” on page 5-8](#) for general information about invoking an entity from a session bean and [“Annotations Used to Interact With Entity Beans” on page A-20](#) for reference information about the annotations.

- Support for vendor-specific subinterfaces when injecting an `EntityManager` from a particular persistence provider.

When you inject a persistence context into a variable using either

```
@javax.persistence.PersistenceContext or
```

```
@javax.persistence.PersistenceUnit,
```

 the standard data type of the injected variable is `EntityManager`. If your persistence provider provides a subinterface of the `EntityManager` (such as `KodoEntityManager` in the case of Oracle Kodo) then as a convenience you can simply set the data type of the injected variable to that subinterface, rather than use the more complicated lookup mechanisms as described in the EJB 3.0 specification. For example:

```
@PersistenceContext private KodoEntityManager em;
```

See [“Invoking a 3.0 Entity” on page 5-8](#) and [KodoEntityManager](#) for general information about `EntityManager` and the Oracle Kodo-provided `KodoEntityManager`.

EJB 3.0 Examples

See [Chapter 3, “Simple Enterprise JavaBeans 3.0 Examples,”](#) for simple examples of stateless and stateful session beans, interceptor classes, and how to invoke an entity. The sections in this guide reference these examples extensively. These examples are meant to simply show how to use the new EJB 3.0 metadata annotations and programming model, rather than how to program the business code of your EJB.

For a more complex example that includes actual business code, see samples directory of the WebLogic Server installation, in particular

```
WL_HOME/samples/server/examples/src/examples/ejb/ejb30, where WL_HOME refers to the installation directory of WebLogic Server, such as /bea/wlserver_10.3.
```


Programming 3.0 Entities

This guide describes how to program 3.0 session and message-driven EJBs, and how to invoke 3.0 entities from a session EJB. It does not describe how to actually program and configure a 3.0 entity. For details instructions on that topic, see [Java Persistence API](#) in the [Oracle Kodo documentation](#).

Simple Enterprise JavaBeans 3.0 Examples

The following sections describe simple Java examples of EJBs that use the new metadata annotation programming model:

- [“Example of a Simple Stateless EJB” on page 3-1](#)
- [“Example of a Simple Stateful EJB” on page 3-3](#)
- [“Example of an Interceptor Class” on page 3-7](#)
- [“Example of Invoking a 3.0 Entity From A Session Bean” on page 3-8](#)

Later procedural sections of this guide that describe how to program an EJB make reference to these examples.

Example of a Simple Stateless EJB

The following code shows a simple business interface for the `ServiceBean` stateless session EJB:

```
package examples;

/**
 * Business interface of the Service stateless session EJB
 */
public interface Service {
    public void sayHelloFromServiceBean();
}
```

```
}
```

The code shows that the `Service` business interface has one method, `sayHelloFromServiceBean()`, that takes no parameters and returns void.

The following code shows the bean file that implements the preceding `Service` interface; the code in bold is described after the example:

```
package examples;

import static javax.ejb.TransactionAttributeType.*;
import javax.ejb.Stateless;
import javax.ejb.TransactionAttribute;
import javax.interceptor.ExcludeDefaultInterceptors;

/**
 * Bean file that implements the Service business interface.
 * Class uses following EJB 3.0 annotations:
 * - @Stateless - specifies that the EJB is of type stateless session
 * - @TransactionAttribute - specifies that the EJB never runs in a
 *     transaction
 * - @ExcludeDefaultInterceptors - specifies any configured default
 *     interceptors should not be invoked for this class
 */

@Stateless
@TransactionAttribute(NEVER)
@ExcludeDefaultInterceptors
public class ServiceBean
    implements Service
{
    public void sayHelloFromServiceBean() {
        System.out.println("Hello From Service Bean!");
    }
}
}
```

The main points to note about the preceding code are:

- Use standard `import` statements to import the metadata annotations you use in the bean file:

```
import static javax.ejb.TransactionAttributeType.*;
import javax.ejb.Stateless;
import javax.ejb.TransactionAttribute;
import javax.interceptor.ExcludeDefaultInterceptors
```

The annotations that apply only to EJB 3.0 are in the `javax.ejb` package. Annotations that can be used by other Java Platform, Enterprise Edition (Java EE) Version 5 components are in more generic packages, such `javax.interceptor` or `javax.annotation`.

- The `ServiceBean` bean file is a plain Java file that implements the `Service` business interface; it is not required to implement any EJB-specific interface. This means that the bean file does not need to implement the lifecycle methods, such as `ejbCreate` and `ejbPassivate`, that were required in the 2.X programming model.
- The class-level `@Stateless` metadata annotation specifies that the EJB is of type stateless session.
- The class-level `@TransactionAttribute(NEVER)` annotation specifies that the EJB never runs inside of a transaction.
- The class-level `@ExcludeDefaultInterceptors` annotation specifies that default interceptors, if any are defined in the `ejb-jar.xml` deployment descriptor file, should never be invoked for any method invocation of this particular EJB.

Example of a Simple Stateful EJB

The following code shows a simple business interface for the `AccountBean` stateful session EJB:

```
package examples;

/**
 * Business interface for the Account stateful session EJB.
 */

public interface Account {

    public void deposit(int amount);
    public void withdraw(int amount);
    public void sayHelloFromAccountBean();

}
```

The code shows that the `Account` business interface has three methods, `deposit`, `withdraw`, and `sayHelloFromAccountBean`.

The following code shows the bean file that implements the preceding Account interface; the code in bold is described after the example:

```
package examples;

import static javax.ejb.TransactionAttributeType.*;

import javax.ejb.Stateful;
import javax.ejb.TransactionAttribute;
import javax.ejb.Remote;
import javax.ejb.EJB;

import javax.annotation.PreDestroy;

import javax.interceptor.Interceptors;
import javax.interceptor.ExcludeClassInterceptors;
import javax.interceptor.InvocationContext;

/**
 * Bean file that implements the Account business interface.
 * Uses the following EJB annotations:
 * - @Stateful: specifies that this is a stateful session EJB
 * - @TransactionAttribute - specifies that this EJB never runs
 *   in a transaction
 * - @Remote - specifies the Remote interface for this EJB
 * - @EJB - specifies a dependency on the ServiceBean stateless
 *   session ejb
 * - @Interceptors - Specifies that the bean file is associated with an
 *   Interceptor class; by default all business methods invoke the
 *   method in the interceptor class annotated with @AroundInvoke.
 * - @ExcludeClassInterceptors - Specifies that the interceptor methods
 *   defined for the bean class should NOT fire for the annotated
 *   method.
 * - @PreDestroy - Specifies lifecycle method that is invoked when the
 *   bean is about to be destroyed by EJB container.
 */

@Stateful
@TransactionAttribute(NEVER)
@Remote({examples.Account.class})
@Interceptors({examples.AuditInterceptor.class})
```

```

public class AccountBean
    implements Account
{
    private int balance = 0;

    @EJB(beanName="ServiceBean")
    private Service service;

    public void deposit(int amount) {
        balance += amount;
        System.out.println("deposited: "+amount+" balance: "+balance);
    }

    public void withdraw(int amount) {
        balance -= amount;
        System.out.println("withdrew: "+amount+" balance: "+balance);
    }

    @ExcludeClassInterceptors
    public void sayHelloFromAccountBean() {
        service.sayHelloFromServiceBean();
    }

    @PreDestroy
    public void preDestroy() {
        System.out.println("Invoking method: preDestroy()");
    }
}

```

The main points to note about the preceding code are:

- Use standard `import` statements to import the metadata annotations you use in the bean file:

```

import static javax.ejb.TransactionAttributeType.*;
import javax.ejb.Stateful;
import javax.ejb.TransactionAttribute;
import javax.ejb.Remote;
import javax.ejb.EJB;

import javax.annotation.PreDestroy;

import javax.interceptor.Interceptors;
import javax.interceptor.ExcludeClassInterceptors;

```

The annotations that apply only to EJB 3.0 are in the `javax.ejb` package. Annotations that can be used by other Java Platform, Enterprise Edition (Java EE) Version 5 components are in more generic packages, such as `javax.interceptor` or `javax.annotation`.

- Import the `InvocationContext` class, used to maintain state between interceptors:

```
import javax.interceptor.InvocationContext;
```

- The `AccountBean` bean file is a plain Java file that implements the `Account` business interface; it is not required to implement any EJB-specific interface. This means that the bean file does not need to implement the lifecycle methods, such as `ejbCreate` and `ejbPassivate`, that were required in the 2.X programming model.
- The class-level `@Stateful` metadata annotation specifies that the EJB is of type `stateful` session.
- The class-level `@TransactionAttribute(NEVER)` annotation specifies that the EJB never runs inside of a transaction.
- The class-level `@Remote` annotation specifies the name of the remote interface of the EJB; in this case it is the same as the business interface, `Account`.
- The class-level `@Interceptors({examples.AuditInterceptor.class})` annotation specifies the interceptor class that is associated with the bean file. This class typically includes a business method interceptor method, as well as lifecycle callback interceptor methods. See [“Example of an Interceptor Class” on page 3-7](#) for details about this class.
- The field-level `@EJB` annotation specifies that the annotated variable, `service`, is injected with the dependent `ServiceBean` stateless session bean context. The data type of the injected field, `Service`, is the business interface of the `ServiceBean` EJB. The following code in the `sayHelloFromAccountBean` method shows how to invoke the `sayHelloFromServiceBean` method of the dependent `ServiceBean`:

```
service.sayHelloFromServiceBean();
```

- The method-level `@ExcludeClassInterceptors` annotation specifies that the `@AroundInvoke` method specified in the associated interceptor class (`AuditInterceptor`) should not be invoked for the `sayHelloFromAccountBean` method.
- The method-level `@PreDestroy` annotation specifies that the EJB container should invoke the `preDestroy` method before the container destroys an instance of the `AccountBean`. This shows how you can specify interceptor methods (for both business methods and lifecycle callbacks) in the bean file itself, in addition to using an associated interceptor class.

Example of an Interceptor Class

The following code shows an example of an interceptor class, specifically the `AuditInterceptor` class that is referenced by the preceding `AccountBean` stateful session bean with the `@Interceptors({examples.AuditInterceptor.class})` annotation; the code in bold is described after the example:

```
package examples;

import javax.interceptor.AroundInvoke;
import javax.interceptor.InvocationContext;

import javax.ejb.PostActivate;
import javax.ejb.PrePassivate;

/**
 * Interceptor class. The interceptor method is annotated with the
 * @AroundInvoke annotation.
 */
public class AuditInterceptor {
    public AuditInterceptor() {}

    @AroundInvoke
    public Object audit(InvocationContext ic) throws Exception {
        System.out.println("Invoking method: "+ic.getMethod());
        return ic.proceed();
    }

    @PostActivate
    public void postActivate(InvocationContext ic) {
        System.out.println("Invoking method: "+ic.getMethod());
    }

    @PrePassivate
    public void prePassivate(InvocationContext ic) {
        System.out.println("Invoking method: "+ic.getMethod());
    }
}
```

The main points to notice about the preceding example are:

- As usual, import the metadata annotations used in the file:

```
import javax.interceptor.AroundInvoke;
import javax.interceptor.InvocationContext;
import javax.ejb.PostActivate;
import javax.ejb.PrePassivate;
```

- The interceptor class is plain Java class.
- The class has an empty constructor:

```
public AuditInterceptor() {}
```

- The method-level `@AroundInvoke` specifies the business method interceptor method. You can use this annotation only once in an interceptor class.
- The method-level `@PostActivate` and `@PrePassivate` annotations specify the methods that the EJB container should call after reactivating and before passivating the bean, respectively.

Note: These lifecycle callback interceptor methods apply only to stateful session beans.

Example of Invoking a 3.0 Entity From A Session Bean

For an example of invoking an entity from a session bean, see the EJB 3.0 example in the distribution kit. After you have installed WebLogic Server, the example is in the following directory:

```
WL_HOME/samples/server/examples/src/examples/ejb/ejb30
```

where *WL_HOME* refers to the directory in which you installed WebLogic Server, such as `/bea/wlserver_10.3`.

Iterative Development of Enterprise JavaBeans 3.0

The sections that follow describe the general EJB 3.0 implementation process, and provide guidance for how to get an EJB 3.0 up and running in WebLogic Server.

For a review of WebLogic Server EJB features, see [“WebLogic Server Value-Added EJB 3.0 Features”](#) on page 2-3.

- [“Overview of the EJB 3.0 Development Process”](#) on page 4-1
- [“Create a Source Directory”](#) on page 4-3
- [“Program the EJB 3.0 Business Interface”](#) on page 4-4
- [“Program the Annotated EJB Class”](#) on page 4-6
- [“Optionally Program Interceptors”](#) on page 4-7
- [“Compile Java Source”](#) on page 4-7
- [“Optionally Create and Edit Deployment Descriptors”](#) on page 4-7
- [“Packaging EJBs”](#) on page 4-8
- [“Deploying EJBs”](#) on page 4-9

Overview of the EJB 3.0 Development Process

This section is a brief overview of the EJB 3.0 development process. It describes the key implementation tasks and associated results.

The following section mostly discusses the EJB 3.0 programming model and points out the differences between the 3.0 and 2.X programming model in only a few places. If you are an experienced EJB 2.X programmer and want the full list of differences between the two models, see [“Understanding EJB 3.0: New Features and Changes From EJB 2.X” on page 2-1](#).

Table 4-1 EJB Development Tasks and Results

#	Step	Description	Result
1	Create a Source Directory	Create the directory structure for your Java source files, and optional deployment descriptors.	A directory structure on your local drive.
2	Program the EJB 3.0 Business Interface	Create the required business interface that describes your EJB.	. java file.
3	Program the Annotated EJB Class	Create the Java file that implements the business interface and includes the EJB 3.0 metadata annotations that describe how your EJB behaves.	. java file.
4	Optionally Program Interceptors	Optionally create the interceptor classes that describe the interceptors that intercept a business method invocation or a life cycle callback event.	. java file.
5	Compile Java Source	Compile source code.	. class file for each class and interface.
6	Optionally Create and Edit Deployment Descriptors	Optionally create the EJB-specific deployment descriptors, although this step is no longer required when using the EJB 3.0 programming model.	<ul style="list-style-type: none"> • ejb-jar.xml, • weblogic-ejb-jar.xml, which contains elements that control WebLogic Server-specific features, and • weblogic-cmp-jar.xml if the bean is a container-managed persistence entity bean.

Table 4-1 EJB Development Tasks and Results

#	Step	Description	Result
7	Packaging EJBs	Package compiled classes and optional deployment descriptors for deployment. If appropriate, you can leave your files unarchived in an exploded directory.	Archive file (either an EJB JAR or Enterprise Application EAR) or equivalent exploded directory.
8	Deploying EJBs	Target the archive or application directory to desired Managed Server, or a WebLogic Server cluster, in accordance with selected staging mode.	The deployment settings for the bean are written to EJBComponent element in <code>config.xml</code> .

Create a Source Directory

Create a source directory where you will assemble the EJB 3.0.

Oracle recommends a *split development directory structure*, which segregates source and output files in parallel directory structures. For instructions on how to set up a split directory structure and package your EJB 3.0 as an enterprise application archive (EAR), see [“Overview of the Split Development Directory Environment”](#) in *Developing Applications with WebLogic Server*.

If you prefer to package and deploy your EJB 3.0 in a JAR file, create a directory for your class files. If you are also creating deployment descriptors (which is optional but supported in the EJB 3.0 programming model) put them in a subdirectory named `META-INF`.

Listing 4-1 Directory Structure for Packaging JAR

```
myEJB/
  META-INF/
    ejb-jar.xml
    weblogic-ejb-jar.xml
    weblogic-cmp-jar.xml
  foo.class
  fooHome.class
  fooBean.class
```

Program the EJB 3.0 Business Interface

The EJB 3.0 business interface is a plain Java interface that describes the full signature of all the business methods of the EJB. For example, assume an `Account` EJB represents a client's checking account; its business interface might include three methods (`withdraw`, `deposit`, and `balance`) that clients can use to manage their bank accounts.

The EJB 3.0 business interface can extend other interfaces. In the case of message-driven beans, the business interface is typically the message-listener interface that is determined by the messaging type used by the bean, such as `javax.jms.MessageListener` in the case of JMS. The interface for a session bean has not such defining type; it can be anything that suits your business needs.

Note: The only requirement for an EJB 3.0 business interface is that it must *not* extend `javax.ejb.EJBObject` or `javax.ejb.EJBLocalObject`, as required in EJB 2.X.

See [“Example of a Simple Stateless EJB” on page 3-1](#) and [“Example of a Simple Stateful EJB” on page 3-3](#) for examples of business interfaces implemented by stateless and stateful session beans.

Business Interface Application Exceptions

When you design the business methods of your EJB, you can define an application exception in the `throws` clause of a method of the EJB's business interface. An *application exception* is an exception that you program in your bean class to alert a client of abnormal application-level conditions. For example, a `withdraw()` method in an `Account` EJB that represents a bank checking account might throw an application exception if the client tries to withdraw more money than is available in their account.

Application exceptions are different from *system exceptions*, which are thrown by the EJB container to alert the client of a system-level exception, such as the unavailability of a database management system. You should not report system-level errors in your application exceptions.

Finally, your business methods should not throw the `java.rmi.RemoteException`, even if the interface is a remote business interface, the bean class is annotated with the `@WebService` JWS annotation, or the method is annotated with `@WebMethod`. The only exception is if the business interface extends `java.rmi.Remote`. If the EJB container encounters problems at the protocol level, the container throws an `EJBException` which wraps the underlying `RemoteException`.

Note: The `@WebService` and `@WebMethod` annotations are in the `javax.jws` package; you use them to specify that your EJB implements a Web Service and that the EJB business will be exposed as public Web Service operations. For details about these annotations and

programming Web Services in general, see [Getting Started With WebLogic Web Services Using JAX-WS](#).

Using Generics in EJBs

The EJB 3.0 programming model supports the use of generics in the business interface at the class level.

Oracle recommends as a best practice that you first define a super-interface that uses the generics, and then have the actual business interface extend this super-interface with a specific data type.

The following example shows how to do this. First, program the super-interface that uses generics:

```
public interface RootI<T> {
    public T getObject();
    public void updateObject(T object);
}
```

Then program the actual business interface to extend `RootI<T>` for a particular data type:

```
@Remote
public interface StatelessI extends RootI<String> { }
```

Finally, program the actual stateless session bean to implement the business interface; use the specified data type, in this case `String`, in the implementation of the methods:

```
@Stateless
public class StatelessSample implements StatelessI {
    public String getObject() {
        return null;
    }
    public void updateObject(String object) {
    }
}
```

If you define the type variables on the business interface or class, they will be erased. In this case, the EJB application can be deployed successfully only when the bean class parameterizes the business interface with upper bounds of the type parameter and no other generic information. For example, in the following example, the upper bound is `Object`:

```
public class StatelessSample implements StatelessI<Object> {
    public Object getObject() {
```

```
        return null;
    }
    public void updateObject(Object object) {
    }
}
```

Serializing and Deserializing Business Objects

Business object serialization and deserialization are supported by the following interfaces, which are implemented by all business objects:

- `weblogic.ejb.spi.BusinessObject`
- `weblogic.ejb.spi.BusinessHandle`

Use the `BusinessObject._WL_getBusinessObjectHandle()` method to get the business handle object and serialize the business handle object.

To deserialize a business object, deserialize the business handle object and use the `BusinessHandle.getBusinessObject()` method to get the business object.

Program the Annotated EJB Class

The EJB bean class is the main EJB programming artifact. It implements the EJB business interface and contains the EJB metadata annotations that specify semantics and requirements to the EJB container, request container services, and provide structural and configuration information to the application deployer or the container runtime.

In the 3.0 programming model, there is only one required annotation: either `@javax.ejb.Stateful`, `@javax.ejb.Stateless`, or `@javax.ejb.MessageDriven` to specify the type of EJB. Although there are many other annotations you can use to further configure your EJB, these annotations have typical default values so that you are not required to explicitly use the annotation in your bean class unless you want it to behave other than in the default manner. This programming model makes it very easy to program an EJB that exhibits typical behavior.

For additional details and examples of programming the bean class, see [Chapter 5, “Programming the Annotated EJB 3.0 Class.”](#)

Optionally Program Interceptors

An interceptor is a method that intercepts the invocation of a business method or a life cycle callback event.

You can define an interceptor method within the actual bean class, or you can program an interceptor class (distinct from the bean class itself) and associate it with the bean class using the `@javax.ejb.Interceptor` annotation.

See “[Specifying Interceptors for Business Methods or Life Cycle Callback Events](#)” on page 5-12 for information on programming the bean class to use interceptors.

Compile Java Source

Once you have written the Java source code for your EJB bean class and optional interceptor class, you must compile it into class files. Typical tools to compile include:

- `weblogic.appc`—The `weblogic.appc` Java class (or its equivalent Ant task `wlappc`) generates and compiles the classes needed to deploy EJBs and JSPs to WebLogic Server. It validates the optional deployment descriptors for compliance with the current specifications at both the individual module level and the application level. The application-level checks include checks between the application-level deployment descriptors and the individual modules as well as validation checks across the modules.

See [Building Modules and Applications Using `wlappc`](#).

- `wlcompile` Ant task—Invokes the `javac` compiler to compile your application's Java components in a split development directory structure.

See [Compiling Applications Using `wlcompile`](#)

- `javac` —The `javac` compiler provided with the Sun Java J2SE SDK provides java compilation capabilities.

See <http://java.sun.com/docs/>.

Optionally Create and Edit Deployment Descriptors

A very important aspect of the new EJB 3.0 programming model is the introduction of metadata annotations. Annotations simplify the EJB development process by allowing a developer to specify within the Java class itself how the bean behaves in the container, requests for dependency injection, and so on. Annotations are an alternative to deployment descriptors that were required by older versions (2.X and earlier) of EJB.

However, EJB 3.0 still fully supports the use of deployment descriptors, even though the standard Java Platform, Enterprise Edition (Java EE) Version 5 ones are not required. For example, you may prefer to use the old 2.X programming model, or might want to allow further customizing of the EJB at a later development or deployment stage; in these cases you can create the standard deployment descriptors in addition to, or instead of, the metadata annotations.

Deployment descriptor elements always override their annotation counterparts. For example, if you specify the `@javax.ejb.TransactionManagement(BEAN)` annotation in your bean class, but then create an `ejb-jar.xml` deployment descriptor for the EJB and set the `<transaction-type>` element to `container`, then the deployment descriptor value takes precedence and the EJB uses container-managed transaction demarcation.

Note: This version of EJB 3.0 also supports all 2.X WebLogic-specific EJB features. However, the features that are configured in the `weblogic-ejb-jar.xml` or `weblogic-cmp-rdbms-jar.xml` deployment descriptor files must continue to be configured that way for this release of EJB 3.0 because currently they do not have any annotation equivalent.

The 2.X version of *Programming WebLogic Enterprise JavaBeans* provides detailed information about creating and editing EJB deployment descriptors, both the Java EE standard and WebLogic-specific ones. In particular, see the following sections:

- [EJB Deployment Descriptors](#) (Overview Information)
- [Editing Deployment Descriptors](#)
- [Deployment Descriptor Schema and Document Type Definitions Reference](#)
- [weblogic-ejb-jar.xml Deployment Descriptor Reference](#)
- [weblogic-cmp-jar.xml Deployment Descriptor Reference](#)

Packaging EJBs

Oracle recommends that you package EJBs as part of an enterprise application. For more information, see [Deploying and Packaging from a Split Development Directory](#) in *Developing Applications with WebLogic Server*.

See [Packaging Considerations for EJBs with Clients in Other Applications](#) in *Programming WebLogic Enterprise JavaBeans* for additional EJB-specific packaging information.

Deploying EJBs

Deploying an EJB enables WebLogic Server to serve the components of an EJB to clients. You can deploy an EJB using one of several procedures, depending on your environment and whether or not your EJB is in production. Deploying an EJB created with the 3.0 programming model is the same as deploying an EJB created with the 2.X programming model.

For general instructions on deploying WebLogic Server applications and modules, including EJBs, see [Deploying Applications to WebLogic Server](#). For EJB-specific deployment issues and procedures, see [Deployment Guidelines For Enterprise Java Beans](#) in the *Programming WebLogic Enterprise JavaBeans* guide, which concentrates on the 2.X programming model.

Iterative Development of Enterprise JavaBeans 3.0

Programming the Annotated EJB 3.0 Class

The sections that follow describe how to program the annotated EJB 3.0 class file:

- [“Overview of Metadata Annotations and EJB 3.0 Bean Files”](#) on page 5-1
- [“Programming the Bean File: Requirements and Changes From 2.X”](#) on page 5-2
- [“Programming the Bean File: Typical Steps”](#) on page 5-3
- [“Complete List of Metadata Annotations By Function”](#) on page 5-22

Overview of Metadata Annotations and EJB 3.0 Bean Files

The new EJB 3.0 programming model uses the [JDK 5.0 metadata annotations](#) feature in which you create an annotated EJB 3.0 bean file and then use the WebLogic compile tool `weblogic.appc` (or its Ant equivalent `wlappc`) to compile the bean file into a Java class file and generate the associated EJB artifacts, such as the required EJB interfaces and deployment descriptors.

The annotated 3.0 bean file is the core of your EJB. It contains the Java code that determines how your EJB behaves. The 3.0 bean file is an ordinary Java class file that implements an EJB business interface that outlines the business methods of your EJB. You then annotate the bean file with JDK 5.0 metadata annotations to specify the shape and characteristics of the EJB, document your EJB, and provide special services such as enhanced business-level security or special business logic during runtime.

See [“Complete List of Metadata Annotations By Function”](#) on page 5-22 for a breakdown of the annotations you can specify in a bean file, by function. These annotations include those described

by the [Enterprise JavaBeans 3.0 specification \(JSR-220\)](#), as well as some described by the [Common Annotations for the Java Platform \(JSR-250\)](#). See [Appendix A, “EJB 3.0 Metadata Annotations Reference,”](#) for reference information about the annotations, listed in alphabetical order.

This topic is part of the iterative development procedure for creating an EJB 3.0, described in [Chapter 4, “Iterative Development of Enterprise JavaBeans 3.0.”](#)

Programming the Bean File: Requirements and Changes From 2.X

The requirements for programming the 3.0 bean class file are essentially the same as the 2.X requirements. This section briefly describes the basic mandatory requirements of the bean class, mostly for overview purposes, as well as changes in requirements between 2.X and 3.0.

See [Programming WebLogic Enterprise JavaBeans](#) for detailed information about the mandatory and optional requirements for programming the bean class.

Bean Class Requirements and Changes From 2.X

The following bullets list the new 3.0 requirements for programming a bean class, as well as the 2.X requirements that no longer apply:

- The class must specify its bean type, typically using one of the following metadata annotations, although you can also override this using a deployment descriptor:
 - `@javax.ejb.Stateless`
 - `@javax.ejb.Stateful`
 - `@javax.ejb.MessageDriven`
 - `@javax.ejb.Entity`

Note: Programming a 3.0 entity is discussed in a separate document. See [Java Persistence API](#) in the [Oracle Kodo documentation](#).

- If the bean is a session bean, the bean class must implement the bean’s business interface(s) or the methods of the bean’s business interface(s), if any.
- Session beans no longer need to implement `javax.ejb.SessionBean`, which means the bean no longer needs to implement the `ejbXXX()` methods, such as `ejbCreate()`, `ejbPassivate()`, and so on.
- Stateful session beans no longer need to implement `java.io.Serializable`.

- Message-driven beans no longer need to implement `javax.ejb.MessageDrivenBean`.

The following requirements are the same as in EJB 2.X and are provided only as a brief overview:

- The class must be defined as `public`, must not be `final`, and must not be `abstract`. The class must be a top level class.
- The class must have a `public` constructor that takes no parameters.
- The class must not define the `finalize()` method.
- If the bean is message-driven, the bean class must implement, directly or indirectly, the message listener interface required by the messaging type that it supports or the methods of the message listener interface. In the case of JMS, this is the `javax.jms.MessageListener` interface.

Bean Class Method Requirements

The method requirements have not changed since EJB 2.X and are provided in this section for a brief overview only.

The requirements for programming the session bean class' methods (that implement the business interface methods) are as follows:

- The method names can be arbitrary.
- The business method must be declared as `public` and must not be `final` or `static`.
- The argument and return value types for a method must be legal types for RMI/IIOP if the method corresponds to a business method on the session bean's remote business interface or remote interface.
- The `throws` clause may define arbitrary application exceptions.

The requirements for programming the message-driven bean class' methods are as follows:

- The methods must implement the listener methods of the message listener interface.
- The methods must be declared as `public` and must not be `final` or `static`.

Programming the Bean File: Typical Steps

The following procedure describes the typical basic steps when programming the 3.0 bean file for an EJB. The steps you follow depends, of course, on what your EJB does.

Refer to [Chapter 3, “Simple Enterprise JavaBeans 3.0 Examples,”](#) for code examples of the topics discussed in the remaining sections.

1. Import the EJB 3.0 and other common annotations that will be used in your bean file. The general EJB annotations are in the `javax.ejb` package, the interceptor annotations are in the `javax.interceptor` package, the annotations to invoke a 3.0 entity are in the `javax.persistence` package, and the common annotations are in the `javax.common` or `javax.common.security` packages. For example:

```
import javax.ejb.Stateless;
import javax.ejb.TransactionAttribute;
import javax.interceptor.ExcludeDefaultInterceptors;
```

2. Specify the business interface that your EJB is going to implement, as well as other standard interfaces. You can either explicitly implement the interface, or use an annotation to specify it.

See [“Specifying the Business and Other Interfaces” on page 5-5.](#)

3. Use the required annotation to specify the type of bean you are programming (session or message-driven).

See [“Specifying the Bean Type \(Stateless, Stateful, Message-Driven\)” on page 5-6.](#)

4. Optionally use dependency injection to use external resources, such as another EJB or other Java Platform, Enterprise Edition (Java EE) Version 5 object.

See [“Injecting Resource Dependency into a Variable or Setter Method” on page 5-7.](#)

5. Optionally create an `EntityManager` object and use the entity annotations to inject entity information.

See [“Invoking a 3.0 Entity” on page 5-8.](#)

6. Optionally program and configure business method or life cycle callback method interceptor method. You can program the interceptor methods in the bean file itself, or in a separate Java file.

See [“Specifying Interceptors for Business Methods or Life Cycle Callback Events” on page 5-12.](#)

7. If your business interface specifies that business methods throw application exceptions, you must program the exception class, the same as in EJB 2.X.

See [“Programming Application Exceptions” on page 5-18](#) for EJB 3.0 specific information.

8. Optionally specify the security roles that are allowed to invoke the EJB methods using the security-related metadata annotations.
See [“Securing Access to the EJB” on page 5-19](#).
9. Optionally change the default transaction configuration in which the EJB runs.
See [“Specifying Transaction Management and Attributes” on page 5-22](#).

Specifying the Business and Other Interfaces

When using the EJB 3.0 programming model to program a bean, you are required to specify a business interface.

There are two ways you can specify the business interface for the EJB bean class:

- By explicitly implementing the business interface, using the `implements` Java keyword.
- By using metadata annotations (such as `javax.ejb.Local` and `javax.ejb.Remote`) to specify the business interface. In this case, the bean class does not need to explicitly implement the business interface.

Typically, if an EJB bean class implements an interface, it is assumed to be the business interface of the EJB. Additionally, the business interface is assumed to be the local interface unless you explicitly denote it as the remote interface, either by using the `javax.ejb.Remote` annotation or updating the appropriate EJB deployment descriptor. You can specify the `javax.ejb.Remote` annotation, as well as the `javax.ejb.Local` annotation, in either the business interface itself, or the bean class that implements the interface.

A bean class can have more than one interface. In this case (excluding the interfaces listed below), you must specify the business interface of the EJB by explicitly using the `javax.ejb.Local` or `javax.ejb.Remote` annotations in either the business interface itself, the bean class that implements the business interface, or the appropriate deployment descriptor.

The following interfaces are excluded when determining whether the bean class has more than one interface:

- `java.io.Serializable`
- `java.io.Externalizable`
- Any of the interfaces defined by the `javax.ejb` package

The following code snippet shows how to specify the business interface of a bean class by explicitly implementing the interface:

```
public class ServiceBean
```

```
implements Service
```

For the full example, see [“Example of a Simple Stateless EJB”](#) on page 3-1

Specifying the Bean Type (Stateless, Stateful, Message-Driven)

There is only one required metadata annotation in a 3.0 bean class: an annotation that specifies the type of bean you are programming. You must specify one, and only one, of the following:

- `@javax.ejb.Stateless`—Specifies that you are programming a stateless session bean.
- `@javax.ejb.Stateful`—Specifies that you are programming a stateful session bean.
- `@javax.ejb.MessageDriven`—Specifies that you are programming a message-driven bean.
- `@javax.ejb.Entity`—Specifies that you are programming an entity bean.

Note: Programming a 3.0 entity is discussed in a separate document. See [Enterprise JavaBeans 3 Persistence](#) in the [Oracle Kodo documentation](#).

Although not required, you can specify attributes of the annotations to further describe the bean type. For example, you can set the following attributes for all bean types:

- `name`—Name of the bean class; the default value is the unqualified bean class name.
- `mappedName`—Product-specific name of the bean.
- `description`—Description of what the bean does.

If you are programming a message-driven bean, then you can specify the following optional attributes:

- `messageListenerInterface`—Specifies the message listener interface, if you haven’t explicitly implemented it or if the bean implements additional interfaces.
- `activationConfig`—Specifies an array of activation configuration name-value pairs that configure the bean in its operational environment.

The following code snippet shows how to specify that a bean is a stateless session bean:

```
@Stateless
public class ServiceBean
    implements Service
```

For the full example, see [“Example of a Simple Stateless EJB”](#) on page 3-1.

Injecting Resource Dependency into a Variable or Setter Method

Dependency injection is when the EJB container automatically supplies (or *injects*) a bean’s variable or setter method with a reference to a resource or another environment entry in the bean’s context. Dependency injection is simply an easier-to-program alternative to using the `javax.ejb.EJBContext` interface or JNDI APIs to look up resources.

You specify dependency injection by annotating a variable or setter method with one of the following annotations, depending on the type of resource you want to inject:

- `@javax.ejb.EJB`—Specifies a dependency on another EJB.
- `@javax.annotation.Resource`—Specifies a dependency on an external resource, such as a JDBC datasource or a JMS destination or connection factory.

Note: This annotation is not specific to EJB; rather, it is part of the common set of metadata annotations used by many different types of Java EE components.

Both annotations have an equivalent grouping annotation to specify a dependency on multiple resources (`@javax.ejb.EJBs` and `@javax.annotation.Resources`).

Although not required, you can specify attributes to these dependency annotations to explicitly describe the dependent resource. The amount of information you need to specify depends upon its usage context and how much information the EJB container can infer from that context. See [“javax.ejb.EJB”](#) on page A-4 and [“javax.annotation.Resource”](#) on page A-27 for detailed information on the attributes and when you should specify them.

The following code snippet shows how to use the `@javax.ejb.EJB` annotation to inject a dependency on an EJB into a variable; only the relevant parts of the bean file are shown:

```
package examples;

import javax.ejb.EJB;

...

@Stateful
public class AccountBean
    implements Account
{
```

```
@EJB(beanName="ServiceBean")
private Service service;

...

public void sayHelloFromAccountBean() {

    service.sayHelloFromServiceBean();

}
```

In the preceding example, the private variable `service` is annotated with the `@javax.ejb.EJB` annotation, which makes reference to the EJB with a bean name of `ServiceBean`. The data type of the service variable is `Service`, which is the business interface implemented by the `ServiceBean` bean class. As soon as the EJB container creates the `AccountBean` EJB, the container injects a reference to `ServiceBean` into the `service` variable; the variable then has direct access to all the business methods of `SessionBean`, as shown in the `sayHelloFromAccountBean` method implementation in which the `sayHelloFromServiceBean` method is invoked.

Invoking a 3.0 Entity

This section describes how to invoke and update a 3.0 entity from within a session bean.

Note: It is assumed in this section that you have already programmed the entity, as well as configured the database resources that support the entity. For details on that topic, see [Java Persistence API](#) in the [Oracle Kodo documentation](#).

An *entity* is a persistent object that represents datastore records; typically an instance of an entity represents a single row of a database table. Entities make it easy to query and update information in a persistent store from within another Java EE component, such as a session bean. A `Person` entity, for example, might include `name`, `address`, and `age` fields, each of which correspond to the columns of a table in a database. Using an `javax.persistence.EntityManager` object to access and manage the entities, you can easily retrieve a `Person` record, based on either their unique id or by using a SQL query, and then change the information and automatically commit the information to the underlying datastore.

The following sections describe the typical programming tasks you perform in your session bean to interact with entities:

- [“Injecting Persistence Context Using Metadata Annotations”](#) on page 5-9
- [“Finding an Entity Using the EntityManager API”](#) on page 5-10

- [“Creating and Updating an Entity Using EntityManager” on page 5-11](#)

Injecting Persistence Context Using Metadata Annotations

In your session bean, use the following metadata annotations inject entity information into a variable:

- `@javax.persistence.PersistenceContext`—Injects a persistence context into a variable of data type `javax.persistence.EntityManager`. A *persistence context* is simply a set of entities such that, for any persistent identity, there is a unique entity instance. The `persistence.xml` file defines and names the persistence contexts available to a session bean.
- `@javax.persistence.PersistenceContexts`—Specifies a set of multiple persistence contexts.
- `@javax.persistence.PersistenceUnit`—Injects a persistence context into a variable of data type `javax.persistence.EntityManagerFactory`.
- `@javax.persistence.PersistenceUnits`—Specifies a set of multiple persistence contexts.

The `@PersistenceContext` and `@PersistenceUnit` annotations perform a similar function: inject persistence context information into a variable; the main difference is the data type of the instance into which you inject the information. If you prefer to have full control over the life cycle of the `EntityManager` in your session bean, then use `@PersistenceUnit` to inject into an `EntityManagerFactory` instance, and then write the code to manually create an `EntityManager` and later destroy when you are done, to release resources. If you prefer that the EJB container manage the life cycle of the `EntityManager`, then use the `@PersistenceContext` annotation to inject directly into an `EntityManager`.

The following example shows how to inject a persistence context into the variable `em` of data type `EntityManager`; relevant code is shown in bold:

```
package examples;

import javax.ejb.Stateless;

import javax.persistence.PersistenceContext;
import javax.persistence.EntityManager;

@Stateless
public class ServiceBean
    implements Service
```

```
{  
    @PersistenceContext private EntityManager em;  
    ...  
}
```

Finding an Entity Using the EntityManager API

Once you have instantiated an `EntityManager` object, you can use its methods to interact with the entities in the persistence context. This section discusses the methods used to identify and manage the life cycle of an entity; see [EntityManager](#) in the [Oracle Kodo documentation](#) for additional uses of the `EntityManager`, such as transaction management, caching, and so on.

Note: For clarity, this section assumes that the entities are configured such that they represent actual rows in a database table.

Use the `EntityManager.find()` method to find a row in a table based on its primary key. The `find` method takes two parameters: the entity class that you are querying, such as `Person.class`, and the primary key value for the particular row you want to retrieve. Once you retrieve the row, you can use standard `getXXX` methods to get particular properties of the entity. The following code snippet shows how to retrieve a `Person` with whose primary key value is 10, and then get their address:

```
public List<Person> findPerson () {  
    Person p = em.find(Person.class, 10);  
    Address a = p.getAddress();  
  
    Query q = em.createQuery("select p from Person p where p.name = :name");  
    q.setParameter("name", "Patrick");  
    List<Person> l = (List<Person>) q.getResultList();  
  
    return l;  
}
```

The preceding example also shows how to use the `EntityManager.createQuery()` method to create a `Query` object that contains a custom SQL query; by contrast, the `EntityManager.find()` method allows you to query using only the table's primary key. In the example, the table is queried for all `Persons` whose first name is `Patrick`; the resulting set of rows populates the `List<Person>` object and is returned to the `findPerson()` invoker.

Creating and Updating an Entity Using EntityManager

To create a new entity instance (and thus add a new row to the database), use the `EntityManager.persist` method, as shown in the following code snippet

```
@TransactionAttribute(REQUIRED)

public Person createNewPerson(String name, int age) {

    Person p = new Person(name, age);

    em.persist(p); // register the new object with the database

    Address a = new Address();

    p.setAddress(a);

    em.persist(a); // depending on how things are configured, this may or
    may not be required

    return p;

}
```

Note: Whenever you create or update an entity, you must be in a transaction, which is why the `@TransactionAttribute` annotation in the preceding example is set to `REQUIRED`.

The preceding example shows how to create a new `Person`, based on parameters passed to the `createNewPerson` method, and then call the `EntityManager.persist` method to automatically add the row to the database table.

The preceding example also shows how to update the newly-created `Person` entity (and thus new table row) with an `Address` by using the `setAddress()` entity method. Depending on the cascade configuration of the `Person` entity, the second `persist()` call may not be necessary; this is because the call to the `setAddress()` method might have automatically triggered an update to the database. For more information about cascading operations, see [Cascade Type](#) in the [Oracle Kodo documentation](#).

If you use the `EntityManager.find()` method to find an entity instance, and then use a `setXXX` method to change a property of the entity, the database is automatically updated and you do not need to explicitly call the `EntityManager.persist()` method, as shown in the following code snippet:

```
@TransactionAttribute(REQUIRED)

public Person changePerson(int id, int newAge) {

    Person p = em.find(Person.class, id);
```

```
p.setAge(newAge);  
return p;  
}
```

In the preceding example, the call to the `Person.setAge()` method automatically triggered an update to the appropriate row in the database table.

Finally, you can use the `EntityManager.merge()` method to quickly and easily update a row in the database table based on an update to an entity made by a client, as shown in the following example:

```
@TransactionAttribute(REQUIRED)  
public Person applyOfflineChanges(Person pDTO) {  
    return em.merge(pDTO);  
}
```

In the example, the `applyOfflineChanges()` method is a business method of the session bean that takes as a parameter a `Person`, which has been previously created by the session bean client. When you pass this `Person` to the `EntityManager.merge()` method, the EJB container automatically finds the existing row in the database table and automatically updates the row with the new data. The `merge()` method then returns a copy of this updated row.

Specifying Interceptors for Business Methods or Life Cycle Callback Events

An interceptor is a method that intercepts a business method invocation or a life cycle callback event. There are two types of interceptors: those that intercept business methods and those that intercept life cycle callback methods.

Interceptors can be specified for session and message-driven beans.

You can program an interceptor method inside the bean class itself, or in a separate interceptor class which you then associate with the bean class with the

`@javax.interceptor.Interceptors` annotation. You can create multiple interceptor methods that execute as a chain in a particular order.

Interceptor instances may hold state. The life cycle of an interceptor instance is the same as that of the bean instance with which it is associated. Interceptors can invoke JNDI, JDBC, JMS, other enterprise beans, and the `EntityManager`. Interceptor methods share the JNDI name space of the

bean for which they are invoked. Programming restrictions that apply to enterprise bean components to apply to interceptors as well.

Interceptors are configured using metadata annotations in the `javax.interceptor` package, as described in later sections.

The following topics discuss how to actually program interceptors for your bean class:

- [Specifying Business or Life Cycle Interceptors: Typical Steps](#)
- [Programming the Interceptor Class](#)
- [Programming Business Method Interceptor Methods](#)
- [Programming Life Cycle Callback Interceptor Methods](#)
- [Specifying Default Interceptor Methods](#)
- [Saving State Across Interceptors With the InvocationContext API](#)

Specifying Business or Life Cycle Interceptors: Typical Steps

The following procedure provides the typical steps to specify and program interceptors for your bean class.

See [“Example of a Simple Stateful EJB” on page 3-3](#) for an example of specifying interceptors and [“Example of an Interceptor Class” on page 3-7](#) for an example of programming an interceptor class.

1. Decide whether interceptor methods are programmed in bean class or in a separate interceptor class.
2. If you decide to program the interceptor methods in a separate interceptor class
 - a. Program the class, as described in [“Programming the Interceptor Class” on page 5-14](#).
 - b. In your bean class, use the `@javax.interceptor.Interceptors` annotation to associate the interceptor class with the bean class. The method in the interceptor class annotated with the `@javax.interceptor.AroundInvoke` annotation then becomes a business method interceptor method of the bean class. Similarly, the methods annotated with the life cycle callback annotations become the life cycle callback interceptor methods of the bean class.

You can specify any number of interceptor classes for a given bean class—the order in which they execute is the order in which they are listed in the annotation. If you specify the interceptor class at the class-level, the interceptor methods apply to all appropriate

bean class methods. If you specify the interceptor class at the method-level, the interceptor methods apply to only the annotated method.

3. In the bean class or interceptor class (wherever you are programming the interceptor methods), program business method interceptor methods, as described in [“Programming Business Method Interceptor Methods”](#) on page 5-14.
4. In the bean class or interceptor class (wherever you are programming the interceptor methods), program life cycle callback interceptor methods, as described in [“Programming Business Method Interceptor Methods”](#) on page 5-14.
5. In the bean class, optionally annotate methods with the `@javax.interceptor.ExcludeClassInterceptors` annotation to exclude any interceptors defined at the class-level.
6. In the bean class, optionally annotate the class or methods with the `@javax.interceptor.ExcludeDefaultInterceptors` annotation to exclude any default interceptors that you might define later. Default interceptors are configured in the `ejb-jar.xml` deployment descriptor, and apply to all EJBs in the JAR file, unless you explicitly use the annotation to exclude them.
7. Optionally specify default interceptors for the entire EJB JAR file, as described in [“Specifying Default Interceptor Methods”](#) on page 5-17.

Programming the Interceptor Class

The interceptor class is a plain Java class that includes the interceptor annotations to specify which methods intercept business methods and which intercept life cycle callback methods.

Interceptor classes support dependency injection, which is performed when the interceptor class instance is created, using the naming context of the associated enterprise bean.

You must include a public no-argument constructor.

You can have any number of methods in the interceptor class, but restrictions apply as to how many methods can be annotated with the interceptor annotations, as described in the following sections.

For an example, see [“Example of an Interceptor Class”](#) on page 3-7.

Programming Business Method Interceptor Methods

You specify business method interceptor methods by annotating them with the `@AroundInvoke` annotation.

An interceptor class or bean class can have only one method annotated with `@AroundInvoke`. To specify that multiple interceptor methods execute for a given business method, you must associate multiple interceptor classes with the bean file, in addition to optionally specifying an interceptor method in the bean file itself. The order in which the interceptor methods execute is the order in which the associated interceptor classes are listed in the `@Interceptor` annotation. Interceptor methods in the bean class itself execute after those defined in the interceptor classes.

You cannot annotate a business method itself with the `@AroundInvoke` annotation.

The signature of an `@AroundInvoke` method must be:

```
Object <METHOD>(InvocationContext) throws Exception
```

The method annotated with the `@AroundInvoke` annotation must always call `InvocationContext.proceed()` or neither the business method will be invoked nor any subsequent `@AroundInvoke` methods. See [“Saving State Across Interceptors With the InvocationContext API” on page 5-17](#) for additional information about the `InvocationContext` API.

Business method interceptor method invocations occur within the same transaction and security context as the business method for which they are invoked. Business method interceptor methods may throw runtime exceptions or application exceptions that are allowed in the `throws` clause of the business method.

For an example, see [“Example of an Interceptor Class” on page 3-7](#).

Programming Life Cycle Callback Interceptor Methods

You specify a method to be a life cycle callback interceptor method so that it can receive notification of life cycle events from the EJB container. Life cycle events include creation, passivation, and destruction of the bean instance.

You can name the life cycle callback interceptor method anything you want; this is different from the EJB 2.X programming model in which you had to name the methods `ejbCreate()`, `ejbPassivate()`, and so on.

You use the following life cycle interceptor annotations to specify that a method is a life cycle callback interceptor method:

- `@javax.ejb.PrePassivate`—Specifies the method that the EJB container notifies when it is about to passivate a stateful session bean.
- `@javax.ejb.PostActivate`—Specifies the method that the EJB container notifies right after it has reactivated a stateful session bean.

- `@javax.annotation.PostConstruct`—Specifies the method that the EJB container notifies before it invokes the first business method and after it has done dependency injection. You typically apply this annotation to the method that performs initialization.

Note: This annotation is in the `javax.annotation` package, rather than `javax.ejb`.

- `@javax.annotation.PreDestroy`—Specifies the method that the EJB container notifies right before it destroys the bean instance. You typically apply this annotation to the method that release resources that the bean class has been holding.

Note: This annotation is in the `javax.annotation` package, rather than `javax.ejb`.

You use the preceding annotations the same way, whether the annotated method is in the bean class or in a separate interceptor class. You can annotate the same method with more than one annotation.

You can also specify any subset or combination of life cycle callback annotations in the bean class or in an associated interceptor class. However, the same callback annotation may not be specified more than once in a given class. If you do specify a callback annotation more than once in a given class, the EJB will not deploy.

To specify that multiple interceptor methods execute for a given life cycle callback event, you must associate multiple interceptor classes with the bean file, in addition to optionally specifying the life cycle callback interceptor method in the bean file itself. The order in which the interceptor methods execute is the order in which the associated classes are listed in the `@Interceptor` annotation. Interceptor methods in the bean class itself execute after those defined in the interceptor classes.

The signature of the annotated methods depends on where the method is defined:

- Life cycle callback methods defined on a bean class have the following signature:

```
void <METHOD>()
```

- Life cycle callback methods defined on an interceptor class have the following signature:

```
void <METHOD>(InvocationContext)
```

See [“Saving State Across Interceptors With the InvocationContext API” on page 5-17](#) for additional information about the `InvocationContext` API.

See [“`javax.ejb.PostActivate`” on page A-10](#), [“`javax.ejb.PrePassivate`” on page A-11](#), [“`javax.annotation.PostConstruct`” on page A-26](#), and [“`javax.annotation.PreDestroy`” on page A-27](#) for additional requirements when programming the life cycle interceptor class.

For an example, see [“Example of an Interceptor Class” on page 3-7](#).

Specifying Default Interceptor Methods

Default interceptor methods apply to *all* components in a particular EJB JAR file or exploded directory, and thus can only be configured in the `ejb-jar.xml` deployment descriptor file and not with metadata annotations, which apply to a particular EJB.

The EJB container invokes default interceptor methods, if any, before *all* other interceptors defined for an EJB (both business and life cycle). If you do not want the EJB container to invoke the default interceptors for a particular EJB, specify the class-level `@javax.interceptor.ExcludeDefaultInterceptors` annotation in the bean file.

In the `ejb-jar.xml` file, use the `<interceptor-binding>` child element of `<assembly-descriptor>` to specify default interceptors. In particular, set the `<ejb-name>` child element to `*`, which means the class applies to all EJBs, and then the `<interceptor-class>` child element to the name of the interceptor class.

The following snippet from an `ejb-jar.xml` file shows how to specify the default interceptor class `org.mycompany.DefaultIC`:

```
<?xml version="1.0" encoding="UTF-8"?>

<ejb-jar version="3.0"
  xmlns="http://java.sun.com/xml/ns/javaee"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="http://java.sun.com/xml/ns/javaee
http://java.sun.com/xml/ns/javaee/ejb-jar_3_0.xsd">
  ...
  <assembly-descriptor>
  ...
    <interceptor-binding>
      <ejb-name>*</ejb-name>
      <interceptor-class>org.mycompany.DefaultIC</interceptor-class>
    </interceptors>
  </assembly-descriptor>
</ejb-jar>
```

Saving State Across Interceptors With the InvocationContext API

Use the `javax.interceptor.InvocationContext` API to pass state information between the interceptors that execute for a given business method or life cycle callback. The EJB Container

passes the same `InvocationContext` instance to each interceptor method, so you can, for example save information when the first business method interceptor method executes, and then retrieve this information for all subsequent interceptor methods that execute for this business method. The `InvocationContext` instance is not shared between business method or life cycle callback invocations.

All interceptor methods must have an `InvocationContext` parameter. You can then use the methods of the `InvocationContext` interface to get and set context information. The `InvocationContext` interface is shown below:

```
public interface InvocationContext {
    public Object getBean();
    public Method getMethod();
    public Object[] getParameters();
    public void setParameters(Object[]);
    public java.util.Map getContextData();
    public Object proceed() throws Exception;
}
```

The `getBean` method returns the bean instance. The `getMethod` method returns the name of the business method for which the interceptor method was invoked; in the case of life cycle callback interceptor methods, `getMethod` returns null.

The `proceed` method causes the invocation of the next interceptor method in the chain, or the business method itself if called from the last `@AroundInvoke` interceptor method.

For an example of using `InvocationContext`, see [“Example of an Interceptor Class” on page 3-7](#).

Programming Application Exceptions

If you specified in the business interface that a method throws an application method, then you must program the exception as a separate class from the bean class.

Use the `@javax.ejb.ApplicationException` annotation to specify that an exception class is an application exception thrown by a business method of the EJB. The EJB container reports the exception directly to the client in the event of the application error.

Use the `rollback` Boolean attribute of the `@ApplicationException` annotation to specify whether the application error causes the current transaction to be rolled back. By default, the current transaction is not rolled back in event of the error.

You can annotate both checked and unchecked exceptions with this annotation.

The following `ProcessingException.java` file shows how to use the `@ApplicationException` annotation to specify that an exception class is an application exception thrown by one of the business methods of the EJB:

```
package examples;

import javax.ejb.ApplicationException;

/**
 * Application exception class thrown when there was a processing error
 * with a business method of the EJB. Annotated with the
 * @ApplicationException annotation.
 */

@ApplicationException()
public class ProcessingException extends Exception {

    /**
     * Catches exceptions without a specified string
     */
    public ProcessingException() {}

    /**
     * Constructs the appropriate exception with the specified string
     *
     * @param message      Exception message
     */
    public ProcessingException(String message) {super(message);}
}
```

Securing Access to the EJB

By default, any user can invoke the public methods of an EJB. If you want to restrict access to the EJB, you can use the following security-related annotations to specify the roles that are allowed to invoke all, or a subset, of the methods:

- `javax.annotation.security.DeclareRoles`—Explicitly lists the security roles that will be used to secure the EJB.
- `javax.annotation.security.RolesAllowed`—Specifies the security roles that are allowed to invoke all the methods of the EJB (when specified at the class-level) or a particular method (when specified at the method-level.)

- `javax.annotation.security.DenyAll`—Specifies that the annotated method can not be invoked by any role.
- `javax.annotation.security.PermitAll`—Specifies that the annotated method can be invoked by all roles.
- `javax.annotation.security.RunAs`—Specifies the role which runs the EJB. By default, the EJB runs as the user who actually invokes it.

The preceding annotations can be used with many Java EE components that allow metadata annotations, not just EJB 3.0.

You create security roles and map users to roles using the WebLogic Server Administration Console to update your security realm. For details, see [Manage Security Roles](#).

The following example shows a simple stateless session EJB that uses all of the security-related annotations; the code in bold is discussed after the example:

```
package examples;

import javax.ejb.Stateless;

import javax.annotation.security.DeclareRoles;
import javax.annotation.security.PermitAll;
import javax.annotation.security.DenyAll;
import javax.annotation.security.RolesAllowed;
import javax.annotation.security.RunAs;

/**
 * Bean file that implements the Service business interface.
 */

@Stateless
@DeclareRoles( { "admin", "hr" } )
@RunAs ( "admin" )

public class ServiceBean
    implements Service
{

    @RolesAllowed ( {"admin", "hr"} )
    public void sayHelloRestricted() {
        System.out.println("Only some roles can invoke this method.");
    }
}
```



```

@DenyAll
public void sayHelloSecret() {
    System.out.println("No one can invoke this method.");
}

@PermitAll
public void sayHelloPublic() {
    System.out.println("Everyone can invoke this method.");
}
}

```

The main points to note about the preceding example are:

- Import the security-related metadata annotations:

```

import javax.annotation.security.DeclareRoles;
import javax.annotation.security.PermitAll;
import javax.annotation.security.DenyAll;
import javax.annotation.security.RolesAllowed;
import javax.annotation.security.RunAs;

```

- The class-level `@DeclareRoles` annotation explicitly specifies that the `admin` and `hr` security roles will later be used to secure some or all of the methods. This annotation is not required; any security role referenced in, for example, the `@RolesReferenced` annotation is implicitly declared. However, explicitly declaring the security roles makes your code easier to read and understand.
- The class-level `@RunAs` annotation specifies that, regardless of the user who actually invokes a particular method of the EJB, the EJB container runs the method as the `admin` role, assuming, of course, that the original user is allowed to invoke the method.
- The `@RolesAllowed` annotation on the `sayHelloRestricted` method specifies that only users mapped to the `admin` and `hr` roles are allowed to invoke the method.
- The `@DenyAll` annotation on the `sayHelloSecret` method specifies that no one is allowed to invoke the method.
- The `@PermitAll` annotation on the `sayHelloPublic` method specifies that all users mapped to any roles are allowed to invoke the method.

Specifying Transaction Management and Attributes

By default, the EJB container invokes a business method within a transaction context. Additionally, the EJB container itself decides whether to commit or rollback a transaction; this is called container-managed transaction demarcation.

You can change this default behavior by using the following annotations in your bean file:

- `javax.ejb.TransactionManagement`—Specifies whether the EJB container or the bean file manages the demarcation of transactions. If you specify that the bean file manages it, then you must program transaction management in your bean file, typically using the Java Transaction API (JTA).
- `javax.ejb.TransactionAttribute`—Specifies whether the EJB container invokes methods within a transaction.

For an example of using the `javax.ejb.TransactionAttribute` annotation, see [“Example of a Simple Stateful EJB” on page 3-3](#).

Complete List of Metadata Annotations By Function

[Appendix A, “EJB 3.0 Metadata Annotations Reference,”](#) provides full reference information about the EJB 3.0 metadata annotations in alphabetical order. The tables in this section group the annotations based on what task they perform.

Annotations to Specify the Bean Type

Table 5-1 Annotations to Specify the Bean Type

Annotation	Description
<code>@javax.ejb.Stateless</code>	Specifies that the bean class is a stateless session bean.
<code>@javax.ejb.Stateful</code>	Specifies that the bean class is a stateful session bean.
<code>@javax.ejb.Init</code>	Specifies the correspondence of a stateful session bean class method with a <code>create<METHOD></code> method for an adapted EJB 2.1 EJBHome and/or EJBLocalHome client view.
<code>@javax.ejb.Remove</code>	Specifies a <code>remove</code> method of a stateful session bean.

Table 5-1 Annotations to Specify the Bean Type

Annotation	Description
<code>@javax.ejb.MessageDriven</code>	Specifies that the bean class is a message-driven bean.
<code>@javax.ejb.ActivationConfigProperty</code>	Specifies properties used to configure a message-driven bean in its operational environment.

Annotations to Specify the Local or Remote Interfaces

Table 5-2 Annotations to Specify the Local or Remote Interfaces

Annotation	Description
<code>@javax.ejb.Local</code>	Specifies a local interface of the bean.
<code>@javax.ejb.Remote</code>	Specifies a remote interface of the bean.

Annotations to Support EJB 2.X Client View

Table 5-3 Annotations to Support EJB 2.X Client View

Annotation	Description
<code>@javax.ejb.LocalHome</code>	Specifies a local home interface of the bean.
<code>@javax.ejb.RemoteHome</code>	Specifies a remote home interface of the bean.

Annotations to Invoke a 3.0 Entity Bean

Table 5-4 Annotations to Invoke a 3.0 Entity Bean

Annotation	Description
<code>@javax.persistence.PersistenceContext</code>	Specifies a dependency on an <code>EntityManager</code> persistence context.
<code>@javax.persistence.PersistenceContexts</code>	Specifies one or more <code>PersistenceContext</code> annotations.
<code>@javax.persistence.PersistenceUnit</code>	Specifies a dependency on an <code>EntityManagerFactory</code> .
<code>@javax.persistence.PersistenceUnits</code>	Specifies one or more <code>PersistenceUnit</code> annotations.

Transaction-Related Annotations

Table 5-5 Transaction-Related Annotations

Annotation	Description
<code>@javax.ejb.TransactionManagement</code>	Specifies the transaction management demarcation type (container- or bean-managed)
<code>@javax.ejb.TransactionAttribute</code>	Specifies whether a business method is invoked within the context of a transaction.

Annotations to Specify Interceptors

Table 5-6 Annotations to Specify Interceptors

Annotation	Description
<code>@javax.interceptor.Interceptors</code>	Specifies the list of interceptor classes associated with a bean class or method.
<code>@javax.interceptor.AroundInvoke</code>	Specifies an interceptor method.

Table 5-6 Annotations to Specify Interceptors

Annotation	Description
<code>@javax.interceptor.ExcludeClassInterceptors</code>	Specifies that, when the annotated method is invoked, the class-level interceptors should <i>not</i> invoke.
<code>@javax.interceptor.ExcludeDefaultInterceptors</code>	Specifies that, when the annotated method is invoked, the default interceptors should <i>not</i> invoke.

Annotations to Specify Life Cycle Callbacks

Table 5-7 Annotations to Specify Life Cycle Callbacks

Annotation	Description
<code>@javax.ejb.PostActivate</code>	Designates a method to receive a callback after a stateful session bean has been activated.
<code>@javax.ejb.PrePassivate</code>	Designates a method to receive a callback before a stateful session bean is passivated.
<code>@javax.annotation.PostConstruct</code>	Specifies the method that needs to be executed after dependency injection is done to perform any initialization.
<code>@javax.annotation.PreDestroy</code>	Specifies a method to be a callback notification to signal that the instance is in the process of being removed by the container

Security-Related Annotations

The following metadata annotations are not specific to EJB 3.0, but rather, are general security-related annotations in the `javax.annotation.security` package.

Table 5-8 Security-Related Annotations

Annotation	Description
<code>@javax.annotation.security.DeclareRoles</code>	Specifies the references to security roles in the bean class.
<code>@javax.annotation.security.RolesAllowed</code>	Specifies the list of security roles that are allowed to invoke the bean's business methods.
<code>@javax.annotation.security.PermitAll</code>	Specifies that all security roles are allowed to invoke the method.
<code>@javax.annotation.security.DenyAll</code>	Specifies that no security roles are allowed to invoke the method.
<code>@javax.annotation.security.RunAs</code>	Specifies the security role which the method is run as.

Context Dependency Annotations

Table 5-9 Context Dependency Annotations

Annotation	Description
<code>@javax.ejb.EJB</code>	Specifies a dependency to an EJB business interface or home interface.
<code>@javax.ejb.EJBs</code>	Specifies one or more <code>@EJB</code> annotations.
<code>@javax.annotation.Resource</code>	Specifies a dependency on an external resource in the bean's environment.
<code>@javax.annotation.Resources</code>	Specifies one or more <code>@Resource</code> annotations.

Timeout and Exceptions Annotations

Table 5-10 Timeout and Exception Annotations

Annotation	Description
<code>@javax.ejb.Timeout</code>	Specifies the timeout method of the bean class.
<code>@javax.ejb.ApplicationException</code>	Specifies that an exception is an application exception and should be reported to the client directly.

Programming the Annotated EJB 3.0 Class

Using Oracle Kodo with WebLogic Server

This chapter provides an overview of developing, deploying, and configuring an Oracle Kodo application using WebLogic Server. The following topics are covered:

- “Overview of Kodo” on page 6-1
- “Creating a Kodo Application” on page 6-2
- “Using Different Kodo Versions” on page 6-2
- “Configuring Persistence” on page 6-2
- “Deploying a Kodo Application” on page 6-4
- “Configuring a Kodo Application” on page 6-5

Overview of Kodo

Oracle Kodo is an implementation of Sun Microsystem’s Java Persistence API (JPA) specification and Java Data Objects (JDO) specification for transparent data objects. Oracle Kodo is available as a stand-alone product and is integrated within WebLogic Server.

This chapter describes how to implement an application using JPA or JDO in WebLogic Server. Within WebLogic Server, the JPA and JDO implementations are part of WebLogic Servers overall Enterprise Java Bean 3.0 persistence implementation.

For general information on creating an application using JPA and JDO, see the *Kodo Developer’s Guide*.

Creating a Kodo Application

The first step in implementing a Oracle Kodo application on WebLogic Server is to write the application code. The following resources provide general information on writing an application that uses Oracle Kodo to manage persistence of your data:

- [Oracle Kodo JPA Tutorials](#)
- [Oracle Kodo JDO Tutorials](#)

Once you are familiar with the steps involved in creating applications using Oracle Kodo and have created your application, the following sections describe how to deploy and configure your application using WebLogic Server.

Using Different Kodo Versions

If you choose to use a different version of Kodo than the one provided by default within WebLogic Server, you must use the `FilteringClassLoader` to exclude the Kodo and OpenJPA libraries from the component classpath.

The following example shows how to exclude these class libraries using `weblogic-application.xml`:

```
<prefer-application-packages>
  <package-name>org.apache.openjpa.*</package-name>
  <package-name>kodo.*</package-name>
</prefer-application-packages>
```

For more information on filtering classloaders, see [“Understanding WebLogic Server Application Classloading”](#) in *Developing Applications With WebLogic Server*.

Configuring Persistence

The following sections describe how to configure persistence.

- [“Editing the Configuration Property Files”](#) on page 6-3
- [“Using the Configuration Files Together”](#) on page 6-4
- [“Configuring Plug-ins”](#) on page 6-4

Editing the Configuration Property Files

Oracle Kodo uses two XML files, listed in the following table, to define configuration properties.

Table 6-1 Persistence Configuration Files

Configuration File	Description
<code>persistence.xml</code>	<p>Kodo configuration parameters defined by the JPA functional specifications. This file is required.</p> <p>The XML schema for structuring this configuration is available at: http://java.sun.com/xml/ns/persistence/persistence_1_0.xsd.</p> <p>For more information, see “Chapter 6. Persistence” in the <i>Kodo Developers Guide</i>.</p>
<code>persistence-configuration.xml</code>	<p>Configuration parameters that are specific to Oracle Kodo. This file is not required when deploying an application. If specified, you must still provide a <code>persistence.xml</code> descriptor.</p> <p>If you do not include <code>persistence-configuration.xml</code> in your deployment, WebLogic Server will create reasonable defaults for each configuration parameter.</p> <p>The XML schema for structuring this configuration is available at: http://www.bea.com/ns/weblogic/persistence-configuration.xsd.</p> <p>Note: The <code>weblogic.jar</code> file must be in the CLASSPATH when using the <code>persistence-configuration.xml</code> file in the Java SE environment.</p>

Edit the contents of the configuration files as required to configure persistence. Persistence units can be packaged as part of a WAR or EJB JAR file, or can be packaged as a JAR file that can be included in a WAR or EAR file. The files should be available as resources in the `META-INF` directory of the root of the persistence unit. For container environments, the root of a persistence unit may be one of the following:

- EJB-JAR file
- `WEB-INF/classes` directory of a WAR file
- JAR file in the `WEB-INF/lib` directory of a WAR file
- JAR file in the root of the EAR

- JAR file in the EAR library directory
- Application client jar file

Using the Configuration Files Together

The following provide considerations for using the configuration files together.

- When using the `persistence-configuration.xml` file, all Kodo-specific properties must be included in the `persistence-configuration.xml` file and not the `persistence.xml` file. In this case, the WebLogic Server Administration Console and WebLogic Scripting Tool (WLST) recognizes the persistence unit as a Kodo persistence unit and provide advanced configuration and tunables support.
- If Kodo-specific properties are included in the `persistence.xml` file, the persistence unit will be treated as a third-party persistence unit by the Administration Console and WLST.
- If the `persistence-configuration.xml` descriptor is available and contains an entry for a given persistence unit, then no Kodo (`kodo`) or OpenJPA (`openjpa`) properties can be specified in the `<properties>` tag of the `persistence.xml` file for that persistence unit.

Configuring Plug-ins

Because Kodo is a highly customizable environment, many configuration properties relate to the creation and configuration of system plugins. Plugin properties have a syntax very similar to that of Java annotations. They allow you to specify both what class to use for the plugin and how to configure the public fields or bean properties of the instantiated plugin instance.

Essentially, plugins are defined using a series of properties using name/value pairs. For example, the following shows how a plugin is defined within `persistence.xml`:

```
<properties>
<property name='myplugin.DataCache'
value='com.bea.MyDataCache(CacheSize=1000, RemoteHost='CacheServer) '>
</properties>
```

Deploying a Kodo Application

Before you deploy a Kodo application, you must perform the following tasks:

- Create a Kodo application, as described in [“Creating a Kodo Application” on page 6-2](#).
- Configure persistence, as described in [“Configuring Persistence” on page 6-2](#).

- Created and archive for your application (.ear or .war).

Once completed, you are ready to deploy your application on WebLogic Server. Once your application is configured, a Kodo application is deployed just like any other application. For complete information on deploying applications, see [Deploying Applications on Oracle WebLogic Server](#).

Configuring a Kodo Application

Note: You cannot create a new persistence unit from the Administration Console. To create a new persistence unit, you must edit `persistence.xml` manually.

Once you have deployed your Kodo application, you can alter the configuration parameters defined in `persistence.xml` and `persistence-configuration.xml`.

The following sections describe how to configure a Kodo application with or without using the Administration Console.

Using the Administration Console

If your deployed application has defined a persistence unit within `persistence.xml`, you can access configuration from within the Administration Console using the following:

1. Select **Deployments**
2. Select the name of the module containing a persistence unit that you want to configure.
3. Select the **Configuration** tab.
4. Select the **Persistence** tab.
5. From the list of persistence units, select the one that you want to configure.

From here, you can access all of the Kodo persistence parameters that can be edited from the Administration Console.

Configuring Kodo Without Using the Administration Console

If you need to alter parameters that are not available using the Administration Console, use one of the following methods:

- Manually edit the `persistence.xml` and `persistence-configuration.xml` files that are archived with the application.

- Use the `SessionHelper` to access and configure the deployment plan. For more information, see “SessionHelper” in “The Tools Package” in [“Understanding the WebLogic Deployment API”](#) in *Programming WebLogic Deployment*.
- Use the WLST `loadApplication()` method to load and update the application deployment plan. For more information, see [“Updating the Deployment Plan”](#) in *WebLogic Scripting Tool*.
- Manually edit your deployment plan. For more information, see “Manually Customizing the Deployment Plan” in [“Exporting an Application for Deployment to New Environments”](#) in *Deploying Applications to WebLogic Server*.

EJB 3.0 Metadata Annotations Reference

The following topics provide reference information about the EJB 3.0 metadata annotations:

- [“Overview of EJB 3.0 Annotations” on page A-1](#)
- [“Annotations for Stateless, Stateful, and Message-Driven Beans” on page A-2](#)
- [“Annotations Used to Configure Interceptors” on page A-18](#)
- [“Annotations Used to Interact With Entity Beans” on page A-20](#)
- [“Standard JDK Annotations Used By EJB 3.0” on page A-26](#)
- [“Standard Security-Related JDK Annotations Used by EJB 3.0” on page A-30](#)
- [“WebLogic Kodo Annotations” on page A-32](#)

Overview of EJB 3.0 Annotations

The new EJB 3.0 programming model uses the [JDK 5.0 metadata annotations](#) feature in which you create an annotated EJB 3.0 bean file and then use the WebLogic compile tool `weblogic.appc` (or its Ant equivalent `wlappc`) to compile the bean file into a Java class file and generate the associated EJB artifacts, such as the required EJB interfaces and deployment descriptors.

The following sections provide reference information for the metadata annotations you can specify in the EJB bean file. Some of the annotations are in the `javax.ejb` package, and are thus specific to EJBs; others are more common and are used by other Java Platform, Enterprise Edition

(Java EE) Version 5 components, and are thus in more generic packages, such as `javax.annotation`.

Annotations for Stateless, Stateful, and Message-Driven Beans

This section provides reference information for the following annotations:

- [“javax.ejb.ActivationConfigProperty” on page A-3](#)
- [“javax.ejb.ApplicationException” on page A-3](#)
- [“javax.ejb.EJB” on page A-4](#)
- [“javax.ejb.EJBs” on page A-5](#)
- [“javax.ejb.Init” on page A-6](#)
- [“javax.ejb.Local” on page A-7](#)
- [“javax.ejb.LocalHome” on page A-8](#)
- [“javax.ejb.MessageDriven” on page A-9](#)
- [“javax.ejb.PostActivate” on page A-10](#)
- [“javax.ejb.PrePassivate” on page A-11](#)
- [“javax.ejb.Remote” on page A-12](#)
- [“javax.ejb.RemoteHome” on page A-12](#)
- [“javax.ejb.Remove” on page A-13](#)
- [“javax.ejb.Stateful” on page A-14](#)
- [“javax.ejb.Stateless” on page A-15](#)
- [“javax.ejb.Timeout” on page A-16](#)
- [“javax.ejb.TransactionAttribute” on page A-16](#)
- [“javax.ejb.TransactionManagement” on page A-17](#)

javax.ejb.ActivationConfigProperty

Description

Target: Any

Specifies properties used to configure a message-driven bean in its operational environment. This may include information about message acknowledgement modes, message selectors, expected destination or endpoint types, and so on.

This annotation is used only as a value to the `activationConfig` attribute of the `@javax.ejb.MessageDriven` annotation.

Attributes

Table A-1 Attributes of the `javax.ejb.ActivationConfigProperty` Annotation

Name	Description	Data Type	Required?
<code>propertyName</code>	Specifies the name of the activation property.	String	Yes
<code>propertyValue</code>	Specifies the value of the activation property.	String	Yes

javax.ejb.ApplicationException

Description

Target: Class

Specifies that an exception is an application exception and that it should be reported to the client application directly, or unwrapped.

This annotation can be applied to both checked and unchecked exceptions.

Attributes

Table A-2 Attributes of the `javax.ejb.ApplicationException` Annotation

Name	Description	Data Type	Required?
rollback	<p>Specifies whether the EJB container should rollback the transaction, if the bean is currently being invoked inside of one, if the exception is thrown.</p> <p>Valid values for this attribute are <code>true</code> and <code>false</code>. Default value is <code>false</code>, or the transaction should <i>not</i> be rolled back.</p>	boolean	No.

`javax.ejb.EJB`

Description

Target: Class, Method, Field

Specifies a dependency or reference to an EJB business or home interface.

You annotate a bean's instance variable with the `@EJB` annotation to specify a dependence on another EJB. WebLogic Server automatically initializes the annotated variable with the reference to the EJB on which it depends; this is also called *dependency injection*. This initialization occurs before any of the bean's business methods are invoked and after the bean's `EJBContext` is set.

You can also annotate a setter method in the bean class; in this case WebLogic Server uses the setter method itself when performing dependency injection. This is an alternative to instance variable dependency injection.

If you apply the annotation to a class, the annotation declares the EJB that the bean will look up at runtime.

Whether using variable or setter method injection, WebLogic Server determines the name of the referenced EJB by either the name or data type of the annotated instance variable or setter method parameter. If there is any ambiguity, you should use the `beanName` or `mappedName` attributes of the `@EJB` annotation to explicitly name the dependent EJB.

Attributes

Table A-3 Attributes of the `javax.ejb.EJB` Annotation

Name	Description	Data Type	Required?
name	<p>Specifies the name by which the referenced EJB is to be looked up in the environment.</p> <p>This name must be unique within the deployment unit, which consists of the class and its superclass.</p>	String	No
beanInterface	<p>Specifies the interface type of the referenced EJB (either a business or home interface).</p> <p>Default value for this attribute is <code>Object.class</code></p>	Class	No
beanName	<p>Specifies the name of the referenced EJB.</p> <p>This attribute corresponds to the name element of the <code>@Stateless</code> or <code>@Stateful</code> annotation in the referenced EJB, which by default is the unqualified name of the referenced bean class.</p> <p>This attribute is most useful when multiple session beans in an EJB JAR file implement the same interface, because the name of each bean must be unique.</p>	String	No
mappedName	<p>Specifies the global JNDI name of the referenced EJB.</p> <p>For example:</p> <pre>mappedName="bank.Account"</pre> <p>specifies that the referenced EJB has a global JNDI name of <code>bank.Account</code> and is deployed in the WebLogic Server JNDI tree.</p> <p>Note: EJBs that use mapped names may not be portable.</p>	String	No
description	Describes the EJB reference.	String	No

`javax.ejb.EJBs`

Description

Target: Class

Specifies an array of `@javax.ejb.EJB` annotations.

Attribute

Table A-4 Attribute of the `javax.ejb.EJBs` Annotation

Name	Description	Data Type	Required?
value	Specifies the array of <code>@javax.ejb.EJB</code> annotations	<code>EJB[]</code>	No

`javax.ejb.Init`

Description

Target: Method

Specifies the correspondence of a method in the bean class with a `createMETHOD` method for an adapted EJB 2.1 `EJBHome` or `EJBLocalHome` client view.

This annotation is used only in conjunction with stateful session beans, or those that have been annotated with the `@javax.ejb.Stateful` class-level annotation,

The return type of a method annotated with the `@javax.ejb.Init` annotation must be `void`, and its parameter types must be exactly the same as those of the referenced `createMETHOD` method or methods.

The `@Init` annotation is required only for stateful session beans that provide a `Remote-Home` or `LocalHome` interface. You must specify the name of the adapted create method of the `Home` or `LocalHome` interface, using the `value` attribute, if there is any ambiguity.

Attributes

Table A-5 Attributes of the `javax.ejb.Init` Annotation

Name	Description	Data Type	Required?
value	<p>Specifies the name of the corresponding <code>createMETHOD</code> method.</p> <p>This attribute is required only when the <code>@Init</code> annotation is used to associate an adapted <code>Home</code> interface of a stateful session bean that has more than one <code>create<METHOD></code> method.</p>	String	No.

`javax.ejb.Local`

Description

Target: Class

Specifies the local interface or interfaces of a session bean. The local interface exposes business logic to local clients—those running in the same application as the EJB. It defines the business methods a local client can call.

You are required to specify this annotation if your bean class implements more than a single interface, not including the following:

- `java.io.Serializable`
- `java.io.Externalizable`
- `javax.ejb.*`

This annotation applies only to stateless or stateful session beans.

Attributes

Table A-6 Attributes of the `javax.ejb.Local` Annotation

Name	Description	Data Type	Required?
value	<p>Specifies the list of local interfaces as an array of classes. You are required to specify this attribute only if your bean class implements more than a single interface, not including the following:</p> <ul style="list-style-type: none"> <code>java.io.Serializable</code> <code>java.io.Externalizable</code> <code>javax.ejb.*</code> 	Class[]	No.

`javax.ejb.LocalHome`

Description

Target: Class

Specifies the local home interface of the bean class.

The local home interface provides methods that local clients—those running in the same application as the EJB—can use to create, remove, and in the case of an entity bean, find instances of the bean. The local home interface also has *home methods*—business logic that is not specific to a particular bean instance.

This attribute applies only to stateless and stateful session beans.

You typically specify this attribute only if you are going to provide an adapted EJB 2.1 component view of the EJB 3.0 bean. You can also use this annotation with bean classes that have been written to the EJB 2.1 APIs.

Attributes

Table A-7 Attributes of the `javax.ejb.LocalHome` Annotation

Name	Description	Data Type	Required?
value	Specifies the local home class.	Class	Yes.

`javax.ejb.MessageDriven`

Description

Target: Class

Specifies that the Enterprise JavaBean is a message-driven bean.

Attributes

Table A-8 Attributes of the `javax.ejb.MessageDriven` Annotation

Name	Description	Data Type	Required?
name	Specifies the name of the message-driven bean. If you do not specify this attribute, the default value is the unqualified name of the bean class.	String	No.
messageListenerInterface	Specifies the message-listener interface of the bean class. You must specify this attribute if the bean class does not explicitly implement the message-listener interface, or if the bean class implements more than one interface other than <code>java.io.Serializable</code> , <code>java.io.Externalizable</code> , or any of the interfaces in the <code>javax.ejb</code> package. The default value for this attribute is <code>Object.class</code> .	Class	No.

Table A-8 Attributes of the `javax.ejb.MessageDriven` Annotation

Name	Description	Data Type	Required?
<code>activationConfig</code>	<p>Specifies the configuration of the message-driven bean in its operational environment. This may include information about message acknowledgement modes, message selectors, expected destination or endpoint types, and so on.</p> <p>You specify activation configuration information using an Array of <code>@javax.ejb.ActivationConfigProperty</code> annotation, specify the property name and value.</p>	Activation ConfigProperty[]	No.
<code>mappedName</code>	<p>Specifies the product-specific name to which the message-driven bean should be mapped.</p> <p>You can also use this attribute to specify the JNDI name of the message destination of this message-driven bean. For example:</p> <pre>mappedName= "my.Queue"</pre> <p>specifies that this message-driven bean is associated with a JMS queue, whose JNDI name is <code>my.Queue</code> and is deployed in the WebLogic Server JNDI tree.</p> <p>Note: If you specify this attribute, the message-driven bean may not be portable.</p>	String	No
<code>description</code>	Specifies a description of the message-driven bean.	String	No

`javax.ejb.PostActivate`

Description

Target: Method

Specifies the life cycle callback method that signals that the EJB container has just reactivated the bean instance.

This annotation applies only to stateful session beans. Because the EJB container automatically maintains the conversational state of a stateful session bean instance when it is passivated, you do not need to specify this annotation for most stateful session beans. You only need to use this annotation, along with its partner `@PrePassivate`, if you want to allow your stateful session

bean to maintain the open resources that need to be closed prior to a bean instance's passivation and then reopened during the bean instance's activation.

Only one method in the bean class can be annotated with this annotation. If you annotate more than one method with this annotations, the EJB will not deploy.

The method annotated with `@PostActivate` must follow these requirements:

- The return type of the method must be `void`.
- The method must not throw a checked exception.
- The method may be `public`, `protected`, `package private` or `private`.
- The method must not be `static`.
- The method must not be `final`.

This annotation does not have any attributes.

javax.ejb.PrePassivate

Description

Target: Method

Specifies the life cycle callback method that signals that the EJB container is about to passivate the bean instance.

This annotation applies only to stateful session beans. Because the EJB container automatically maintains the conversational state of a stateful session bean instance when it is passivated, you do not need to specify this annotation for most stateful session beans. You only need to use this annotation, along with its partner `@PostActivate`, if you want to allow your stateful session bean to maintain the open resources that need to be closed prior to a bean instance's passivation and then reopened during the bean instance's activation.

Only one method in the bean class can be annotated with this annotation. If you annotate more than one method with this annotations, the EJB will not deploy.

The method annotated with `@PrePassivate` must follow these requirements:

- The return type of the method must be `void`.
- The method must not throw a checked exception.
- The method may be `public`, `protected`, `package private` or `private`.

- The method must not be `static`.
- The method must not be `final`.

This annotation does not have any attributes.

javax.ejb.Remote

Description

Target: Class

Specifies the remote interface or interfaces of a session bean. The remote interface exposes business logic to remote clients—clients running in a separate application from the EJB. It defines the business methods a remote client can call.

This annotation applies only to stateless or stateful session beans.

Attributes

Table A-9 Attributes of the javax.ejb.Remote Annotation

Name	Description	Data Type	Required?
value	<p>Specifies the list of remote interfaces as an array of classes.</p> <p>You are required to specify this attribute only if your bean class implements more than a single interface, not including the following:</p> <ul style="list-style-type: none"> • <code>java.io.Serializable</code> • <code>java.io.Externalizable</code> • <code>javax.ejb.*</code> 	Class[]	No.

javax.ejb.RemoteHome

Description

Target: Class

Specifies the remote home interface of the bean class.

The remote home interface provides methods that remote clients—those running in a separate application from the EJB—can use to create, remove, and find instances of the bean.

This attribute applies only to stateless and stateful session beans.

You typically specify this attribute only if you are going to provide an adapted EJB 2.1 component view of the EJB 3.0 bean. You can also use this annotation with bean classes that have been written to the EJB 2.1 APIs.

Attributes

Table A-10 Attributes of the `javax.ejb.RemoveHome` Annotation

Name	Description	Data Type	Required?
value	Specifies the remote home class.	Class	Yes.

`javax.ejb.Remove`

Description

Target: Method

Use the `@javax.ejb.Remove` annotation to denote a remove method of a stateful session bean.

When the method completes, the EJB container will invoke the method annotated with the `@javax.annotation.PreDestroy` annotation, if any, and then destroy the stateful session bean.

Attributes

Table A-11 Attributes of the `javax.ejb.Remove` Annotation

Name	Description	Data Type	Required?
retainIfException	Specifies that the container should not remove the stateful session bean if the annotated method terminates abnormally with an application exception. Valid values are <code>true</code> and <code>false</code> . Default value is <code>false</code> .	boolean	No.

javax.ejb.Stateful

Description

Target: Class

Specifies that the Enterprise JavaBean is a stateful session bean.

Attributes

Table A-12 Attributes of the `javax.ejb.Stateful` Annotation

Name	Description	Data Type	Required?
name	<p>Specifies the name of the stateful session bean.</p> <p>If you do not specify this attribute, the default value is the unqualified name of the bean class.</p>	String	No.
mappedName	<p>Specifies the product-specific name to which the stateful session bean should be mapped.</p> <p>You can also use this attribute to specify the JNDI name of this stateful session bean. WebLogic Server uses the value of the <code>mappedName</code> attribute when creating the bean's global JNDI name. In particular, the JNDI name will be:</p> <p><i>mappedName#name_of_businessInterface</i></p> <p>where <i>name_of_businessInterface</i> is the fully qualified name of the business interface of this session bean.</p> <p>For example, if you specify <code>mappedName="bank"</code> and the fully qualified name of the business interface is <code>com.CheckingAccount</code>, then the JNDI of the business interface is <code>bank#com.CheckingAccount</code>.</p> <p>Note: If you specify this attribute, the stateful session bean may not be portable.</p>	String	No.
description	Describes the stateful session bean.	String	No.

javax.ejb.Stateless

Description

Target: Class

Specifies that the Enterprise JavaBean is a stateless session bean.

Attributes

Table A-13 Attributes of the `javax.ejb.Stateless` Annotation

Name	Description	Data Type	Required?
name	Specifies the name of the stateless session bean. If you do not specify this attribute, the default value is the unqualified name of the bean class.	String	No.
mappedName	Specifies the product-specific name to which the stateless session bean should be mapped. You can also use this attribute to specify the JNDI name of this stateless session bean. WebLogic Server uses the value of the <code>mappedName</code> attribute when creating the bean's global JNDI name. In particular, the JNDI name will be: <i>mappedName#name_of_businessInterface</i> where <i>name_of_businessInterface</i> is the fully qualified name of the business interface of this session bean. For example, if you specify <code>mappedName="bank"</code> and the fully qualified name of the business interface is <code>com.CheckingAccount</code> , then the JNDI of the business interface is <code>bank#com.CheckingAccount</code> . Note: If you specify this attribute, the stateless session bean may not be portable.	String	No.
description	Describes the stateless session bean.	String	No.

javax.ejb.Timeout

Description

Target: Method

Specifies the timeout method of the bean class.

This annotation makes it easy to program an EJB timer service in your bean class. The EJB timer service is an EJB-container provided service that allows you to create timers that schedule callbacks to occur when a timer object expires.

Previous to EJB 3.0, your bean class was required to implement `javax.ejb.TimerObject` if you wanted to program the timer service. Additionally, your bean class had to include a method with the exact name `ejbTimeout`. These requirements are relaxed in Version 3.0 of EJB. You no longer are required to implement the `javax.ejb.TimerObject` interface, and you can name your timeout method anything you want, as long as you annotate it with the `@Timeout` annotation. You can, however, continue to use the pre-3.0 way of programming the timer service if you want.

For details, see [Programming the EJB Timer Service](#).

This annotation does not have any attributes.

javax.ejb.TransactionAttribute

Description

Target: Class, Method

Specifies whether the EJB container invokes an EJB business method within a transaction context.

WARNING: If you specify this annotation, you are also required to use the `@TransactionManagement` annotation to specify container-managed transaction demarcation.

You can specify this annotation on either the bean class, or a particular method of the class that is also a method of the business interface. If specified at the bean class, the annotation applies to all applicable business interface methods of the class. If specified for a particular method, the annotation applies to that method only. If the annotation is specified at both the class and the method level, the method value overrides if the two disagree.

If you do not specify the `@TransactionAttribute` annotation in your bean class, and the bean uses container managed transaction demarcation, the semantics of the `REQUIRED` transaction attribute are assumed.

Attributes

Table A-14 Attributes of the `javax.ejb.TransactionAttribute` Annotation

Name	Description	Data Type	Required?
value	<p>Specifies how the EJB container manages the transaction boundaries when invoking a business method.</p> <p>For details about these values, see the description of the <code>trans-attribute</code> element in the Container-Managed Transactions Elements table.</p> <p>Valid values for this attribute are:</p> <ul style="list-style-type: none"> <code>TransactionAttributeType.MANDATORY</code> <code>TransactionAttributeType.REQUIRED</code> <code>TransactionAttributeType.REQUIRED_NEW</code> <code>TransactionAttributeType.SUPPORTS</code> <code>TransactionAttributeType.NOT_SUPPORTED</code> <code>TransactionAttributeType.NEVER</code> <p>Default value is <code>TransactionAttributeType.REQUIRED</code>.</p>	TransactionAttributeType	No.

`javax.ejb.TransactionManagement`

Description

Target: Class

Specifies the transaction management demarcation type of the session bean or message-driven bean.

A transaction is a unit of work that changes application state—whether on disk, in memory or in a database—that, once started, is completed entirely, or not at all. Transactions can be demarcated—started, and ended with a commit or rollback—by the EJB container, by bean code, or by client code. This annotation specifies whether the EJB container or the user-written bean code manages the demarcation of a transaction.

If you do not specify this annotation in your bean class, it is assumed that the bean has container-managed transaction demarcation.

For additional information about transactions, see [Transaction Design and Management Options](#).

Attributes

Table A-15 Attributes of the `javax.ejb.TransactionManagement` Annotation

Name	Description	Data Type	Required?
value	<p>Specifies the transaction management demarcation type used by the bean class.</p> <p>Valid values for this attribute are:</p> <ul style="list-style-type: none"> <code>TransactionManagementType.CONTAINER</code> <code>TransactionManagementType.BEAN</code> <p>Default value is <code>TransactionManagementType.CONTAINER</code></p>	<code>TransactionManagementType</code>	No.

Annotations Used to Configure Interceptors

This section provides reference information for the following annotations:

- [“`javax.interceptor.AroundInvoke`” on page A-18](#)
- [“`javax.interceptor.ExcludeClassInterceptors`” on page A-19](#)
- [“`javax.interceptor.ExcludeDefaultInterceptors`” on page A-19](#)
- [“`javax.interceptor.Interceptors`” on page A-19](#)

`javax.interceptor.AroundInvoke`

Description

Target: Method

Specifies the business method interceptor for either a bean class or an interceptor class.

You can annotate only *one* method in the bean class or interceptor class with the `@AroundInvoke` annotation; the method cannot be a business method of the bean class.

This annotation does not have any attributes.

javax.interceptor.ExcludeClassInterceptors

Description

Target: Method

Specifies that any class-level interceptors should not be invoked for the annotated method. This does not include default interceptors, whose invocation are excluded only with the `@ExcludeDefaultInterceptors` annotation.

This annotation does not have any attributes.

javax.interceptor.ExcludeDefaultInterceptors

Description

Target: Class, Method

Specifies that any defined default interceptors (which can be specified only in the EJB deployment descriptors, and not with annotations) should not be invoked.

If defined at the class-level, the default interceptors are never invoked for any of the bean's business methods. If defined at the method-level, the default interceptors are never invoked for the particular business method, but they are invoked for all other business methods that do not have the `@ExcludeDefaultInterceptors` annotation.

This annotation does not include any attributes.

javax.interceptor.Interceptors

Description

Target: Class, Method

Specifies the interceptor classes that are associated with the bean class or method. An interceptor class is a class—distinct from the bean class itself—whose methods are invoked in response to business method invocations and/or life cycle events on the bean.

The interceptor class can include both an business interceptor method (annotated with the `@javax.interceptor.AroundInvoke` annotation) and life cycle callback methods (annotated

with the `@javax.annotation.PostConstruct`, `@javax.annotation.PreDestroy`, `@javax.ejb.PostActivate`, and `@javax.ejb.PrePassivate` annotations).

Any number of interceptor classes may be defined for a bean class. If more than one interceptor class is defined, they are invoked in the order they are specified in the annotation.

If the annotation is specified at the class-level, the interceptors apply to all business methods of the EJB. If specified at the method-level, the interceptors apply to just that method. You can specify the same interceptor class to more than one method of the bean class. By default, method-level interceptors are invoked after all applicable interceptors (default interceptors, class-level interceptors, and so on).

Attributes

Table A-16 Attributes of the `javax.interceptor.Interceptors` Annotation

Name	Description	Data Type	Required?
value	Specifies the array of interceptor classes. If there is more than one interceptor class in the array, the order in which they are listed defines the order in which they are invoked.	Class[]	Yes

Annotations Used to Interact With Entity Beans

This section provides reference information about the following annotations:

- [“`javax.persistence.PersistenceContext`” on page A-20](#)
- [“`javax.persistence.PersistenceContexts`” on page A-22](#)
- [“`javax.persistence.PersistenceUnit`” on page A-24](#)
- [“`javax.persistence.PersistenceUnits`” on page A-25](#)

`javax.persistence.PersistenceContext`

Description

Target: Class, Method, Field

Specifies a dependency on a container-managed `EntityManager` persistence context.

You use this annotation to interact with a 3.0 entity bean, typically by performing dependency injection into an `EntityManager` instance.

The `EntityManager` interface defines the methods that are used to interact with the persistence context. A persistence context is a set of entity instances; an entity is a lightweight persistent domain object. The `EntityManager` API is used to create and remove persistent entity instances, to find entities by their primary key, and to query over entities.

Attributes

javax.persistence.PersistenceContexts

Table A-17 Attributes of the javax.persistence.PersistenceContext Annotation

Name	Description	Data Type	Required?
name	<p>Specifies the name by which the <code>EntityManager</code> and its persistence unit are to be known within the context of the session or message-driven bean.</p> <p>You only need to specify this attribute if you use a JNDI lookup to obtain an <code>EntityManager</code>; if you use dependency injection, then you do not need to specify this attribute.</p>	String	No.

Table A-17 Attributes of the `javax.persistence.PersistenceContext` Annotation

Name	Description	Data Type	Required?
unitName	<p>Specifies the name of the persistence unit.</p> <p>If you specify a value for this attribute that is the same as the name of a persistence unit in the <code>persistence.xml</code> file, the EJB container automatically deploys the persistence unit and sets its JNDI name to its persistence unit name. Similarly, if you do not specify this attribute, but the name of the variable into which you are injecting the persistence context information is the same as the name of a persistence unit in the <code>persistence.xml</code> file, then the EJB container again automatically deploys the persistence unit with its JNDI name equal to its unit name.</p> <p>Note: The <code>persistence.xml</code> file is an XML file, located in the <code>META-INF</code> directory of the EJB JAR file, that specifies the database used with the entity beans and specifies the default behavior of the <code>EntityManager</code>.</p> <p>You must specify this attribute if there is more than one persistence unit within the referencing scope.</p>	String	No.
type	<p>Specifies whether the lifetime of the persistence context is scoped to a transaction or whether it extends beyond that of a single transaction.</p> <p>Valid values for this attribute are:</p> <ul style="list-style-type: none"> • <code>PersistenceContextType.TRANSACTION</code> • <code>PersistenceContextType.EXTENDED</code> <p>Default value is <code>PersistenceContextType.TRANSACTION</code>.</p>	PersistenceContextType	No.

Description

Target: Class

Specifies an array of `@javax.persistence.PersistenceContext` annotations.

Attributes

Table A-18 Attributes of the `javax.persistence.PersistenceContexts` Annotation

Name	Description	Data Type	Required?
value	Specifies the array of <code>@javax.persistence.PersistenceContext</code> annotations.	<code>PersistenceContext[]</code>	Yes.

`javax.persistence.PersistenceUnit`

Description

Target: Class, Method, Field

Specifies a dependency on an `EntityManagerFactory` object.

You use this annotation to interact with a 3.0 entity bean, typically by performing dependency injection into an `EntityManagerFactory` instance. You can then use the `EntityManagerFactory` to create one or more `EntityManager` instances. This annotation is similar to the `@PersistenceContext` annotation, except that it gives you more control over the life of the `EntityManager` because you create and destroy it yourself, rather than let the EJB container do it for you.

The `EntityManager` interface defines the methods that are used to interact with the persistence context. A persistence context is a set of entity instances; an entity is a lightweight persistent domain object. The `EntityManager` API is used to create and remove persistent entity instances, to find entities by their primary key, and to query over entities.

Attributes

Table A-19 Attributes of the `javax.persistence.PersistenceUnit` Annotation

Name	Description	Data Type	Required?
name	<p>Specifies the name by which the <code>EntityManagerFactory</code> is to be known within the context of the session or message-driven bean</p> <p>You are not required to specify this attribute if you use dependency injection, only if you also use JNDI to look up information.</p>	String	No
unitName	<p>Refers to the name of the persistence unit as defined in the <code>persistence.xml</code> file. This file is an XML file, located in the <code>META-INF</code> directory of the EJB JAR file, that specifies the database used with the entity beans and specifies the default behavior of the <code>EntityManager</code>.</p> <p>If you set this attribute, the EJB container automatically deploys the referenced persistence unit and sets its JNDI name to its persistence unit name. Similarly, if you do not specify this attribute, but the name of the variable into which you are injecting the persistence context information is the same as the name of a persistence unit in the <code>persistence.xml</code> file, then the EJB container again automatically deploys the persistence unit with its JNDI name equal to its unit name.</p> <p>You are required to specify this attribute only if there is more than one persistence unit in the referencing scope.</p>	String	No

`javax.persistence.PersistenceUnits`

Description

Target: Class

Specifies an array of `@javax.persistence.PersistenceUnit` annotations.

Attributes

Table A-20 Attributes of the `javax.persistence.PersistenceUnits` Annotation

Name	Description	Data Type	Required?
value	Specifies the array of <code>@javax.persistence.PersistenceUnit</code> annotations.	<code>PersistenceUnit[]</code>	Yes

Standard JDK Annotations Used By EJB 3.0

This section provides reference information about the following annotations:

- [“`javax.annotation.PostConstruct`” on page A-26](#)
- [“`javax.annotation.PreDestroy`” on page A-27](#)
- [“`javax.annotation.Resource`” on page A-27](#)
- [“`javax.annotation.Resources`” on page A-29](#)

`javax.annotation.PostConstruct`

Description

Target: Method

Specifies the life cycle callback method that the EJB container should execute before the first business method invocation and after dependency injection is done to perform any initialization.

You may specify a `@PostConstruct` method in any bean class that includes dependency injection.

Only one method in the bean class can be annotated with this annotation. If you annotate more than one method with this annotations, the EJB will not deploy.

The method annotated with `@PostConstruct` must follow these requirements:

- The return type of the method must be `void`.
- The method must not throw a checked exception.

- The method may be `public`, `protected`, `package private` or `private`.
- The method must not be `static`.
- The method must not be `final`.

This annotation does not have any attributes.

javax.annotation.PreDestroy

Description

Target: Method

Specifies the life cycle callback method that signals that the bean class instance is about to be destroyed by the EJB container. You typically apply this annotation to methods that release resources that the bean class has been holding.

Only one method in the bean class can be annotated with this annotation. If you annotate more than one method with this annotations, the EJB will not deploy.

The method annotated with `@PreDestroy` must follow these requirements:

- The return type of the method must be `void`.
- The method must not throw a checked exception.
- The method may be `public`, `protected`, `package private` or `private`.
- The method must not be `static`.
- The method must not be `final`.

This annotation does not have any attributes.

javax.annotation.Resource

Description

Target: Class, Method, Field

Specifies a dependence on an external resource, such as a JDBC data source or a JMS destination or connection factory.

If you specify the annotation on a field or method, the EJB container injects an instance of the requested resource into the bean when the bean is initialized. If you apply the annotation to a class, the annotation declares a resource that the bean will look up at runtime.

Attributes

Table A-21 Attributes of the `javax.annotation.Resource` Annotation

Name	Description	Data Type	Required?
name	<p>Specifies the name of the resource reference.</p> <p>If you apply the <code>@Resource</code> annotation to a field, the default value of the <code>name</code> attribute is the field name, qualified by the class name. If you apply it to a method, the default value is the JavaBeans property name corresponding to the method, qualified by the class name. If you apply the annotation to class, there is no default value and thus you are required to specify the attribute.</p>	String	No
type	<p>Specifies the Java data type of the resource.</p> <p>If you apply the <code>@Resource</code> annotation to a field, the default value of the <code>type</code> attribute is the type of the field. If you apply it to a method, the default is the type of the JavaBeans property. If you apply it to a class, there is no default value and thus you are required to specify this attribute.</p>	Class	No
authenticationType	<p>Specifies the authentication type to use for the resource.</p> <p>You specify this attribute only for resources representing a connection factory of any supported type.</p> <p>Valid values for this attribute are:</p> <ul style="list-style-type: none"> <code>AuthenticationType.CONTAINER</code> <code>AuthenticationType.APPLICATION</code> <p>Default value is <code>AuthenticationType.CONTAINER</code></p>	AuthenticationType	No

Table A-21 Attributes of the `javax.annotation.Resource` Annotation

Name	Description	Data Type	Required?
shareable	<p>Indicates whether a resource can be shared between this EJB and other EJBs.</p> <p>You specify this attribute only for resources representing a connection factory of any supported type or ORB object instances.</p> <p>Valid values for this attribute are <code>true</code> and <code>false</code>. Default value is <code>true</code>.</p>	boolean	No.
mappedName	<p>Specifies the global JNDI name of the dependent resource.</p> <p>For example:</p> <pre>mappedName= "my.Datasource"</pre> <p>specifies that the JNDI name of the dependent resources is <code>my.Datasource</code> and is deployed in the WebLogic Server JNDI tree.</p>	String	No.
description	Specifies a description of the resource.	String	No.

`javax.annotation.Resources`

Description

Target: Class

Specifies an array of `@Resource` annotations.

Attributes

Table A-22 Attributes of the `javax.annotation.Resources` Annotation

Name	Description	Data Type	Required?
value	Specifies the array of <code>@Resource</code> annotations.	<code>Resource[]</code>	Yes.

Standard Security-Related JDK Annotations Used by EJB 3.0

This section provides reference information about the following annotations:

- “[javax.annotation.security.DeclareRoles](#)” on page A-30
- “[javax.annotation.security.DenyAll](#)” on page A-31
- “[javax.annotation.security.PermitAll](#)” on page A-31
- “[javax.annotation.security.RolesAllowed](#)” on page A-31
- “[javax.annotation.security.RunAs](#)” on page A-32

javax.annotation.security.DeclareRoles

Description

Target: Class

Defines the security roles that will be used in the EJB.

You typically use this annotation to define roles that can be tested from within the methods of the annotated class, such as using the `isUserInRole` method. You can also use the annotation to explicitly declare roles that are implicitly declared if you use the `@RolesAllowed` annotation on the class or a method of the class.

You create security roles in WebLogic Server using the Administration Console. For details, see [Manage Security Roles](#).

Attributes

Table A-23 Attributes of the `javax.annotation.security.DeclareRoles` Annotation

Name	Description	Data Type	Required?
value	Specifies an array of security roles that will be used in the bean class.	String[]	Yes.

javax.annotation.security.DenyAll

Description

Target: Method

Specifies that no security role is allowed to access the annotated method, or in other words, the method is excluded from execution in the EJB container.

This annotation does not have any attributes.

javax.annotation.security.PermitAll

Description

Target: Method

Specifies that all security roles currently defined for WebLogic Server are allowed to access the annotated method.

This annotation does not have any attributes.

javax.annotation.security.RolesAllowed

Description

Target: Class, Method

Specifies the list of security roles that are allowed to access methods in the EJB.

If you specify it at the class-level, then it applies to all methods in the bean class. If you specify it at the method-level, then it only applies to that method. If you specify the annotation at both the class- and method-level, the method value overrides the class value.

You create security roles in WebLogic Server using the Administration Console. For details, see [Manage Security Roles](#).

Attributes

Table A-24 Attributes of the `javax.annotation.security.RolesAllowed` Annotation

Name	Description	Data Type	Required?
value	List of security roles that are allowed to access methods of the bean class.	String[]	Yes.

`javax.annotation.security.RunAs`

Description

Target: Class

Specifies the security role which actually executes the EJB in the EJB container.

The security role must exist in the WebLogic Server security realm and map to a user or group. For details, see [Manage Security Roles](#).

Attributes

Table A-25 Attributes of the `javax.annotation.security.RunAs` Annotation

Name	Description	Data Type	Required?
value	Specifies the security role which the EJB should run as.	String	Yes.

WebLogic Kodo Annotations

This section provides reference information for the following WebLogic Kodo Annotations:

- “[weblogic.javaee.AllowRemoveDuringTrasaction](#)” on page A-33
- “[weblogic.javaee.CallByReference](#)” on page A-33
- “[weblogic.javaee.DisableWarnings](#)” on page A-34
- “[weblogic.javaee.EJBReference](#)” on page A-35

- “weblogic.javaee.Idempotent” on page A-35
- “weblogic.javaee.JMSClientID” on page A-36
- “weblogic.javaee.JNDIName” on page A-37
- “weblogic.javaee.MessageDestinationConfiguration” on page A-38
- “weblogic.javaee.TransactionIsolation” on page A-39
- “weblogic.javaee.TransactionTimeoutSeconds” on page A-39

Note: The annotations described in this section are overridden if the comparable configuration is defined in the `weblogic-ejb-jar.xml` deployment descriptor. For more information, see “[weblogic-ejb-jar.xml Deployment Descriptor Reference](#)” in *Programming WebLogic Enterprise JavaBeans*.

weblogic.javaee.AllowRemoveDuringTrasaction

Description

Target: Class (Stateful session EJBs only)

Flag that specifies whether an instance can be removed during a transaction.

Note: This annotation is overridden by the `allow-remove-during-transaction` element in the `weblogic-ejb-jar.xml` deployment descriptor. For more information, see “[weblogic-ejb-jar.xml Deployment Descriptor Reference](#)” in *Programming WebLogic Enterprise JavaBeans*.

weblogic.javaee.CallByReference

Description

Target: Class (Stateful or stateless sessions EJBs only)

Flag that specifies whether parameters are copied—or passed by reference—regardless of whether the EJB is called remotely or from within the same EAR.

Note: Method parameters are *always* passed by value when an EJB is called remotely.

This annotation is overridden by the `enable-call-by-reference` element in the `weblogic-ejb-jar.xml` deployment descriptor. For more information, see

[“weblogic-ejb-jar.xml Deployment Descriptor Reference”](#) in *Programming WebLogic Enterprise JavaBeans*.

weblogic.javaee.DisableWarnings

Description

Target: Class

Specifies that WebLogic Server should disable the warning message whose ID is specified.

Note: This annotation is overridden by the `disable-warning` element in the `weblogic-ejb-jar.xml` deployment descriptor. For more information, see [“weblogic-ejb-jar.xml Deployment Descriptor Reference”](#) in *Programming WebLogic Enterprise JavaBeans*.

Attributes

Table A-26 Attributes of the `weblogic.javaee.DisableWarnings`

Name	Description	Data Type	Required?
WarningCode	<p>Specifies the warning code. Set this element to one of the following four values:</p> <ul style="list-style-type: none"> BEA-010001—Disables this warning message: “EJB class loaded from system classpath during deployment.” BEA-010054—Disables this warning message: “EJB class loaded from system classpath during compilation.” BEA-010200—Disables this warning message: “EJB impl class contains a public static field, method or class.” BEA-010202—Disables this warning message: “Call-by-reference not enabled.” 	String	Yes

weblogic.javaee.EJBReference

Description

Target: Class, Method, Field

Maps EJB reference name to its JNDI name.

Attribute

Table A-27 Attribute of the `weblogic.javaee.EJBReference` Annotation

Name	Description	Data Type	Required?
name	Specifies the name by which the referenced EJB is to be looked up in the environment. This name must be unique within the deployment unit, which consists of the class and its superclass.	String	Yes
jndiName	Specifies the JNDI name of an actual EJB, resource, or reference available in WebLogic Server.	String	Yes

weblogic.javaee.Idempotent

Description

Target: Class

Specifies an EJB that is written in such a way that repeated calls to the same method with the same arguments has exactly the same effect as a single call. This allows the failover handler to retry a failed call without knowing whether the call actually compiled on the failed server. When you enable idempotent for a method, the EJB stub can automatically recover from any failure as long as it can reach another server hosting the EJB.

Note: This annotation is overridden by the `idempotent-method` and `retry-methods-on-rollback` elements in the `weblogic-ejb-jar.xml` deployment descriptor. For more information, see [“weblogic-ejb-jar.xml Deployment Descriptor Reference”](#) in *Programming WebLogic Enterprise JavaBeans*.

Attributes

Table A-28 Attributes of the `weblogic.javaee.Idempotent`

Name	Description	Data Type	Required?
<code>retryOnRollbackCount</code>	Number of times to automatically retry container-managed transactions that have rolled back. This attribute defaults to 0.	int	No

`weblogic.javaee.JMSClientID`

Description

Target: Method

Specifies a client ID for the MDB when it connects to a JMS destination. Required for durable subscriptions to JMS topics.

If you specify the connection factory that the MDB uses in [weblogic.javaee.MessageDestinationConfiguration](#), the client ID can be defined in the `ClientID` element of the associated `JMSConnectionFactory` element in `config.xml`.

If `JMSConnectionFactory` in `config.xml` does not specify a `ClientID`, or if you use the default connection factory, (you do not specify [weblogic.javaee.MessageDestinationConfiguration](#)) the MDB uses the `jms-client-id` value as its client id.

Note: This annotation is overridden by the `jms-client-id` element in the `weblogic-ejb-jar.xml` deployment descriptor. For more information, see [“weblogic-ejb-jar.xml Deployment Descriptor Reference”](#) in *Programming WebLogic Enterprise JavaBeans*.

Attributes

Table A-29 Attributes of the `weblogic.javaee.JMSClientID`

Name	Description	Data Type	Required?
value	Client ID.	String	No
generateUniqueID	Flag that indicates whether or not you want the EJB container to generate a unique client ID for every instance of an MDB. Enabling this flag makes it easier to deploy durable MDBs to multiple server instances in a WebLogic Server cluster.	Class	No

`weblogic.javaee.JNDIName`

Description

Target: Class (Stateful or stateless session EJBs only)

Specifies the JNDI name of an actual EJB, resource, or reference available in WebLogic Server. This annotation is valid on the remote interface and the implementation class, if there is only one remote interface.

Notes: Assigning a JNDI name to a bean is not recommended. Global JNDI names generate heavy multicast traffic during clustered server startup. See “Using EJB Links” in [“Implementing Enterprise Java Beans”](#) in *Programming WebLogic Enterprise JavaBeans* for the better practice.

If you have an EAR library that contains EJBs, you cannot deploy multiple applications that reference the library because attempting to do so will result in a JNDI name conflict. This is because global JNDI name cannot be set for individual EJBs in EAR libraries; it can only be set for an entire library.

This annotation is overridden by the `jndi-name` element in the `weblogic-ejb-jar.xml` deployment descriptor. For more information, see [“weblogic-ejb-jar.xml Deployment Descriptor Reference”](#) in *Programming WebLogic Enterprise JavaBeans*.

Attributes

Table A-30 Attributes of the `weblogic.javaee.JNDIName`

Name	Description	Data Type	Required?
value	JNDI name.	String	Yes.

`weblogic.javaee.MessageDestinationConfiguration`

Description

Target: Class (Message-driven EJBs only)

Specifies the JNDI name of the JMS Connection Factory that a message-driven EJB looks up to create its queues and topics. See [“Configuring MDBs for Destinations”](#) on page 7-18 and [“How to Set connection-factory-jndi-name”](#) on page 7-21.

Note: This annotation is overridden by the `connection-factory-jndi-name` element in the `weblogic-ejb-jar.xml` deployment descriptor. For more information, see [“weblogic-ejb-jar.xml Deployment Descriptor Reference”](#) in *Programming WebLogic Enterprise JavaBeans*.

Attributes

Table A-31 Attributes of the `weblogic.javaee.MessageDestinationConfiguration`

Name	Description	Data Type	Required?
connectionFactoryJNDIName	Connection factory JNDI name. This attribute defaults to an empty string.	String	No
initialContextFactory	WebLogic initial context factory. This attribute defaults to <code>weblogic.jndi.WLInitialContextFactory.class</code> .	Class	No
providerURL	URL of the provider. This attribute defaults to <code>t3://localhost:7001</code> .	String	No

weblogic.javaee.TransactionIsolation

Description

Target: Method

Method-level transaction isolation settings for an EJB.

Note: This annotation is overridden by the `trans-timeout-seconds` element in the `weblogic-ejb-jar.xml` deployment descriptor. For more information, see [“weblogic-ejb-jar.xml Deployment Descriptor Reference”](#) in *Programming WebLogic Enterprise JavaBeans*.

Attributes

Table A-32 Attributes of the `weblogic.javaee.Idempotent`

Name	Description	Data Type	Required?
IsolationLevel	<p>Isolation level. Valid values include:</p> <ul style="list-style-type: none"> <code>READ_COMMITTED</code>—Transaction can view only committed updates from other transactions. <code>READ_UNCOMMITTED</code>—Transactions can view uncommitted updates from other transactions. <code>REPEATABLE_READ</code>—Once the transaction reads a subset of data, repeated reads of the same data return the same values, even if other transactions have subsequently modified the data. <code>SERIALIZABLE</code>—Simultaneously executing this transaction multiple times has the same effect as executing the transaction multiple times in a serial fashion. <p>This attribute defaults to <code>DEFAULT</code>.</p>	int	No

weblogic.javaee.TransactionTimeoutSeconds

Description

Target: Class

Defines the timeout for transactions in seconds.

Attributes

Table A-33 Attributes of the `weblogic.javaee.TransactionTimeoutSeconds`

Name	Description	Data Type	Required?
value	Transaction timeout value in seconds. This attribute defaults to 30 (seconds).	int	No

Persistence Configuration Schema Reference

The following sections describe the namespace, schema location, file structure, and elements in the Kodo-specific deployment descriptor, `persistence-configuration.xml`.

- [“persistence-configuration.xml Namespace Declaration and Schema Location”](#) on page B-1
- [“persistence-configuration.xml Deployment Descriptor File Structure”](#) on page B-2
- [“persistence-configuration.xml Deployment Descriptor Elements”](#) on page B-6

persistence-configuration.xml Namespace Declaration and Schema Location

The correct text for the namespace declaration and schema location for the Kodo `persistence-configuration.xml` file is as follows.

```
<persistence-configuration
xmlns="http://www.bea.com/ns/weblogic/persistence-configuration"
xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xsi:schemaLocation="http://www.bea.com/ns/weblogic/persistence-configuration
http://www.bea.com/ns/weblogic/persistence-configuration/1.0/persistence-configuration.xsd">
...
</persistence-configuration>
```

persistence-configuration.xml Deployment Descriptor File Structure

The `persistence-configuration.xml` deployment descriptor file describes the elements that are unique to Oracle Kodo.

The top-level elements in the Kodo `persistence-configuration.xml` are as follows:

- `persistence-configuration`
 - `persistence-configuration-unit`
 - `aggregate-listeners`
 - `auto-clear`
 - `auto-detaches`
 - `default-broker-factory` | `abstract-store-broker-factory` | `client-broker-factory` | `jdbc-broker-factory` | `custom-broker-factory`
 - `default-broker-impl` | `kodo-broker` | `custom-broker-impl`
 - `default-class-resolver` | `custom-class-resolver`
 - `default-compatibility` | `compatibility` | `default-compatibility`
 - `connection2-driver-name`
 - `connection2-password`
 - `connection2-properties`
 - `connection2-url`
 - `connection2-user-name`
 - `connection-decorators`
 - `connection-driver-name`
 - `connection-factory2-name`
 - `connection-factory2-properties`
 - `connection-factory-mode`
 - `connection-factory-name`
 - `connection-factory-properties`
 - `connection-password`

- `connection-properties`
- `connection-retain-mode`
- `connection-url`
- `connection-user-name`
- `data-caches`
- `default-data-cache-manager` | `kodo-data-cache-manager` |
`data-cache-manager-impl` | `custom-data-cache-manager`
- `data-cache-timeout`
- `access-dictionary` | `db2-dictionary` | `derby-dictionary` |
`empress-dictionary` | `foxpro-dictionary` | `hsql-dictionary` |
`informix-dictionary` | `jdatastore-dictionary` | `mysql-dictionary` |
`oracle-dictionary` | `pointbase-dictionary` | `postgres-dictionary` |
`sql-server-dictionary` | `sybase-dictionary` | `custom-dictionary`
- `default-detach-state` | `detach-options-loaded` |
`detach-options-fetch-groups` | `detach-options-all` |
`custom-detach-state`
- `default-driver-data-source` | `kodo-pooling-data-source` |
`simple-driver-data-source` | `custom-driver-data-source`
- `stack-execution-context-name-provider` |
`transaction-name-execution-context-name-provider` |
`user-object-execution-context-name-provider`
- `none-profiling` | `local-profiling` | `export-profiling` | `gui-profiling`
- `none-jmx` | `local-jmx` | `gui-jmx` | `jmx2-jmx` | `mx4j1-jmx` | `wls81-jmx`
- `dynamic-data-structs`
- `eager-fetch-mode`
- `fetch-batch-size`
- `fetch-direction`
- `fetch-groups`
- `filter-listeners`
- `flush-before-queries`
- `ignore-changes`

Persistence Configuration Schema Reference

- [inverse-manager](#)
- [jdbc-listeners](#)
- [default-lock-manager](#) | [pessimistic-lock-manager](#) | [none-lock-manager](#) | [single-jvm-exclusive-lock-manager](#) | [version-lock-manager](#) | [custom-lock-manager](#)
- [lock-timeout](#)
- [commons-log-factory](#) | [log4j-log-factory](#) | [log-factory-impl](#) | [none-log-factory](#) | [custom-log](#)
- [lrs-size](#)
- [mapping](#)
- [default-mapping-defaults](#) | [deprecated-jdo-mapping-defaults](#) | [mapping-defaults-impl](#) | [persistence-mapping-defaults](#) | [custom-mapping-defaults](#)
- [extension-deprecated-jdo-mapping-factory](#) | [kodo-persistence-mapping-factory](#) | [mapping-file-deprecated-jdo-mapping-factory](#) | [orm-file-jdor-mapping-factory](#) | [table-deprecated-jdo-mapping-factory](#) | [table-jdor-mapping-factory](#) | [custom-mapping-factory](#)
- [default-meta-data-factory](#) | [jdo-meta-data-factory](#) | [deprecated-jdo-meta-data-factory](#) | [kodo-persistence-meta-data-factory](#) | [custom-meta-data-factory](#)
- [default-meta-data-repository](#) | [kodo-mapping-repository](#) | [custom-meta-data-repository](#)
- [multithreaded](#)
- [nontransactional-read](#)
- [nontransactional-write](#)
- [optimistic](#)
- [default-orphaned-key-action](#) | [log-orphaned-key-action](#) | [exception-orphaned-key-action](#) | [none-orphaned-key-action](#) | [custom-orphaned-key-action](#)
- [http-transport](#) | [tcp-transport](#) | [custom-persistence-server](#)
- [default-proxy-manager](#) | [profiling-proxy-manager](#) | [proxy-manger-impl](#) | [custom-proxy-manager](#)

- query-caches
- default-query-compilation-cache | cache-map | concurrent-hash-map | custom-query-compilation-cache
- read-lock-level
- jms-remote-commit-provider | single-jvm-remote-commit-provider | tcp-remote-commit-provider | cluster-remote-commit-provider | custom-remote-commit-provider
- restore-state
- result-set-type
- retain-state
- retry-class-registration
- default-savepoint-manager | in-memory-savepoint-manager | jdbc3-savepoint-manager | oracle-savepoint-manager | custom-savepoint-manager
- schema
- default-schema-factory | dynamic-schema-factory | file-schema-factory | lazy-schema-factory | table-schema-factory | custom-schema-factory
- schemas
- class-table-jdbc-seq | native-jdbc-seq | table-jdbc-seq | time-seeded-seq | value-table-jdbc-seq | custom-seq
- default-sql-factory | kodo-sql-factory | custom-sql-factory
- subclass-fetch-mode
- synchronize-mappings
- transaction-isolation
- transaction-mode
- default-update-manager | constraint-update-manager | batching-operation-order-update-manager | operation-order-update-manager | table-lock-update-manager | custom-update-manager
- write-lock-level
- profiling
 - none-profiling

- `local-profiling`
- `export-profiling`
- `gui-profiling`
- `execution-context-name-provider`
 - `stack-execution-context-name-provider`
 - `transaction-name-execution-context-name-provider`
 - `user-object-execution-context-name-provider`
- `jmx`
 - `none-jmx`
 - `local-jmx`
 - `gui-jmx`
 - `jmx2-jmx`
 - `mx4j1-jmx`
 - `wls81-jmx`

persistence-configuration.xml Deployment Descriptor Elements

The following list of the elements in `persistence-configuration.xml` includes all elements that are supported in this release of Kodo.

abstract-store-broker-factory

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Specifies a plugin string (see “[Plugin Configuration](#)” in *Kodo Developer’s Guide*) describing the `kodo.kernel.BrokerFactory` type to use, in this case abstract store. For more information, see “[kodo.BrokerFactory](#)” in the *Kodo Developer’s Guide*.

Example

```
<abstract-store-broker-factory/>
```

access-dictionary

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Defines configuration values for the Access Dictionary. For a complete description of each of the values that you can specify, see “[Access Dictionary Configuration](#)” in the *Administration Console Online Help*.

Example

The default values for each of the configuration values are shown in the example below.

Persistence Configuration Schema Reference

```
<access-dictionary>
  <char-type-name>CHAR</char-type-name>
  <outer-join-clause>LEFT OUTER JOIN</outer-join-clause>
  <binary-type-name>BINARY</binary-type-name>
  <clob-type-name>CLOB</clob-type-name>
  <supports-locking-with-distinct-clause>>true</supports-locking-with-dis
tinct-clause>
  <simulate-locking>>false</simulate-locking>
  <system-tables>>true</system-tables>
  <concatenate-function>({0}||{1})</concatenate-function>
  <substring-function-name>SUBSTR</substring-function-name>
  <supports-query-timeout>>true</supports-query-timeout>
  <use-set-bytes-for-blobs>>true</use-set-bytes-for-blobs>
  <max-constraint-name-length>18</max-constraint-name-length>
  <search-string-escape>\\</search-string-escape>
  <supports-cascade-update-action>>true</supports-cascade-update-action>
  <string-length-function>CHAR_LENGTH({0})</string-length-function>
  <long-varbinary-type-name>LONGVARBINARY</long-varbinary-type-name>
  <supports-unique-constraints>>true</supports-unique-constraints>
  <supports-restrict-delete-action>>true</supports-restrict-delete-action
>
  <trim-leading-function>LTRIM(LEADING '{1}' from
{0})</trim-leading-function>
  <supports-default-delete-action>>false</supports-default-delete-action>
  <next-sequence-query></next-sequence-query>
  <long-varchar-type-name>LONGVARCHAR</long-varchar-type-name>
  <cross-join-clause>CROSS JOIN</cross-join-clause>
  <max-embedded-clob-size>-1</max-embedded-clob-size>
  <date-type-name>DATE</date-type-name>
  <supports-schema-for-get-tables>>true</supports-schema-for-get-tables>
  <supports-alter-table-with-drop-column>>true</supports-alter-table-with
-drop-column>
  <current-time-function>CURRENT_TIME</current-time-function>
  <requires-condition-for-cross-join>>false</requires-condition-for-cross
-join>
  <ref-type-name>REF</ref-type-name>
  <concatenate-delimiter>'OPENJPATOKEN'</concatenate-delimiter>
  <catalog-separator>.</catalog-separator>
```

```

    <supports-mod-operator>false</supports-mod-operator>
    <schema-case>upper</schema-case>
    <java-object-type-name>JAVA_OBJECT</java-object-type-name>
    <driver-vendor></driver-vendor>
    <supports-locking-with-multiple-tables>true</supports-locking-with-multiple-tables>
    <max-column-name-length>128</max-column-name-length>
    <double-type-name>DOUBLE</double-type-name>
    <use-get-string-for-clobs>false</use-get-string-for-clobs>
    <decimal-type-name>DECIMAL</decimal-type-name>
    <smallint-type-name>SMALLINT</smallint-type-name>
    <date-precision>1000000</date-precision>
    <supports-alter-table-with-add-column>true</supports-alter-table-with-add-column>
    <bit-type-name>BIT</bit-type-name>
    <supports-null-table-for-get-columns>true</supports-null-table-for-get-columns>
    <to-upper-case-function>UPPER( {0} )</to-upper-case-function>
    <supports-select-end-index>false</supprotos-select-end-index>
    <supports-auto-assign>false</supports-auto-assign>
    <store-large-numbers-as-strings>false</store-large-numbers-as-strings>
    <constraint-name-mode>before</constraint-name-mode>
    <allows-alias-in-bulk-clause>true</allows-alias-in-bulk-clause>
    <supports-select-for-update>true</supports-select-for-update>
    <distinct-count-column-separator></distinct-count-column-separator>
    <supports-subselect>true</supports-subselect>
    <time-type-name>TIME</time-type-name>
    <auto-assign-type-name></auto-assign-type-name>
    <use-get-object-for-blobs>false</use-get-object-for-blobs>
    <max-auto-assign-name-length>31</max-auto-assign-name-length>
    <validation-sql></validation-sql>
    <struct-type-name>STRUCT</struct-type-name>
    <varchar-type-name>VARCHAR</varchar-type-name>
    <range-position>0</range-position>
    <supports-restrict-update-action>true</supports-restrict-update-action>
  >
  <auto-assign-clause></auto-assign-clause>
  <supports-multiple-nontransactional-result-sets>true</supports-multipl

```

Persistence Configuration Schema Reference

```
e-nontransactional-result-sets>
  <bit-length-function>(OCTET_LENGTH({0}) * 8)</bit-length-function>
  <create-primary-keys>>true</create-primary-keys>
  <null-type-name>NULL</null-type-name>
  <float-type-name>FLOAT</float-type-name>
  <use-get-bytes-for-blobs>>true</use-get-bytes-for-blobs>
  <table-types>TABLE</table-types>
  <numeric-type-name>NUMERIC</numeric-type-name>
  <table-for-update-clause></table-for-update-clause>
  <integer-type-name>INTEGER</integer-type-name>
  <blob-type-name>BLOB</blob-type-name>
  <for-update-clause>FOR UPDATE</for-update-clause>
  <boolean-type-name>BOOLEAN</boolean-type-name>
  <use-get-best-row-identifier-for-primary-keys>>false</use-get-best-row-
identifier-for-primary-keys>
  <supports-foreign-keys>>true</supports-foreign-keys>
  <drop-table-sql>DROP TABLE {0}</drop-table-sql>
  <use-set-string-for-clobs>>false</use-set-string-for-clobs>
  <supports-locking-with-order-clause>>false</supports-locking-with-order-
clause>
  <platform>Generic</platform>
  <fixed-size-type-names></fixed-size-type-names>
  <store-chars-as-numbers>>true</store-chars-as-numbers>
  <max-indexes-per-table>2147483647</max-indexes-per-table>
  <requires-cast-for-comparisons>>false</requires-cast-for-comparisons>
  <supports-having>>true</supports-having>
  <supports-locking-with-outer-join>>true</supports-locking-with-outer-jo
in>
  <supports-correlated-subselect>>true</supports-correlated-subselect>
  <supports-null-table-for-get-imported-keys>>false</supports-null-table-fo
r-get-imported-keys>
  <bigint-type-name>BIGINT</bigint-type-name>
  <last-generated-key-query></last-generated-key-query>
  <reserved-words></reserved-words>
  <supports-null-update-action>>true</supports-null-update-action>
  <use-schema-name>>true</use-schema-name>
  <supports-deferred-constraints>>true</supports-deferred-constraints>
  <real-type-name>REAL</real-type-name>
```



```

    <requires-alias-for-subselect>>false</requires-alias-for-subselect>
    <supports-null-table-for-get-index-info>>false</supports-null-table-for-
get-index-info>
    <trim-trailing-function>TRIM(TRAILING '{1}' FROM
{0})</trim-trailing-function>
    <supports-locking-with-select-range>>true</supports-locking-with-select-
range>
    <storage-limitations-fatal>>false</storage-limitations-fatal>
    <supports-locking-with-inner-join>>true</supports-locking-with-inner-jo
in>
    <current-timestamp-function>CURRENT_TIMESTAMP</current-timestamp-funct
<cast-function>CAST({0} AS {1})</cast-function>
    <other-type-name>OTHER</other-type-name>
    <max-index-name-length>128</max-index-name-length>
    <distinct-type-name>DISTINCT</distinct-type-name>
    <character-column-size>255</character-column-size>
    <varbinary-type-name>VARBINARY</varbinary-type-name>
    <max-table-name-length>128</max-table-name-length>
    <close-pool-sql></close-pool-sql>
    <current-date-function>CURRENT_DATE</current-date-function>
    <join-syntax>sql92</join-syntax>
    <max-embedded-blob-size>-1</max-embedded-blob-size>
    <trim-both-function>TRIM(BOTH '{1}' FROM {0})</trim-both-function>
    <supports-select-start-index>>false</supports-select-start-index>
    <to-lower-case-function>LOWER({0})</to-lower-case-function>
    <array-type-name>ARRAY</array-type-name>
    <inner-join-clause>INNER JOIN</inner-join-clause>
    <supports-default-update-action>>true</supports-default-update-action>
    <supports-schema-for-get-columns>>true</supports-schema-for-get-columns
>
    <tinyint-type-name>TINYINT</tinyint-type-name>
    <supports-null-table-for-get-primary-keys>>false</supports-null-table-f
or-get-primary-keys>
    <system-schemas></system-schemas>
    <requires-cast-for-math-functions>>false</requires-cast-for-math-functi
ons>
    <supports-null-delete-action>>true</supports-null-delete-action>
    <requires-auto-commit-for-meta-data>>false</requires-auto-commit-for-me

```

```
ta-data>
  <timestamp-type-name>TIMESTAMP</timestamp-type-name>
  <initialization-sql></initialization-sql>
  <supports-cascade-delete-action>true</supports-cascade-delete-action>
  <supports-timestamp-nanos>true</supports-timestamp-nanos>
</access-dictionary>
```

access-unloaded

Range of values:	true false
Default value:	true
Parent elements:	<pre>persistence-configuration persistence-configuration-unit detach-options-all and persistence-configuration persistence-configuration-unit detach-options-fetch-groups and persistence-configuration persistence-configuration-unit detach-options-loaded</pre>

Function

Flag that specifies whether to allow access to unloaded fields of detached objects. Set to `false` to throw an exception whenever an unloaded field is accessed. This option is only available when you use detached state managers.

For more information about the detach state, see [“Defining the Detached Object Graph”](#) in *Kodo Developer’s Guide*.

Example

```
<access-unloaded>true</access-unloaded>
```

action

Range of values:	warn exception
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit inverse-manager

Function

Controls the action taken when the inverse manager detects an inconsistent bidirectional relation. Valid options include:

- warn—Log a warning.
- exception—Throw an exception.

For more information, see [“Managed Inverses”](#) in *Kodo Developer’s Guide*.

Example

```
<action>warn</action>
```

addresses

Range of values:	Valid IP addresses
Default value:	[]
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Specifies a semicolon separated list of IP addresses to which notifications should be sent. For more information, see [“TCP”](#) in the *Kodo Developer’s Guide*.

Example

```
<addresses>[ ]</addresses>
```

advanced-sql

Range of values:	String
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Configures advanced SQL features. For more information, see “[kodo.jdbc.SQLFactory](#)” in the *Kodo Developer’s Guide*.

aggregate-listeners

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Specifies a list of aggregate listeners to make available to all queries. Each listener must implement the `org.apache.openjpa.kernel.exps.AggregateListener` interface. For more information, see “[kodo.AggregateListeners](#)” in *Kodo Developer’s Guide*.

Example

```
<aggregate-listeners>  
  <custom-aggregate-listener>  
    <classname>...</classname>
```

```

    <properties>...</properties>
  </custom-aggregate-listener>
</aggregate-listener>

```

allocate

Range of values:	Integer
Default value:	50
Parent elements:	<p> persistence-configuration persistence-configuration-unit class-table-jdbc-seq and persistence-configuration persistence-configuration-unit native-jdbc-seq and persistence-configuration persistence-configuration-unit table-jdbc-seq and persistence-configuration persistence-configuration-unit value-table-jdbc-seq </p>

Function

Specifies the number of values to allocate on each database trip. Defaults to 50, meaning the class will set aside the next 50 numbers each time it accesses the sequence table. In this case, Kodo only accesses the database to get new sequence numbers once every 50 sequence number requests. For more information, see [“Generators”](#) in the *Kodo Developer’s Guide*.

Example

```
<allocate>50</allocate>
```

assert-allowed-type

Range of values:	true false
Default value:	false
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Flag that specifies whether to immediately throw an exception if you attempt to add an element to a collection or map that is not assignable to the element type declared in metadata. For more information, see “[Custom Proxies](#)” in the *Kodo Developer’s Guide*.

Example

```
<assert-allowed-type>false</assert-allowed-type>
```

auto-clear

Range of values:	<code>persistence-configuration-unit</code> : datastore all <code>kodo-broker</code> : 0 1
Default value:	datastore (0)
Parent elements:	persistence-configuration persistence-configuration-unit and persistence-configuration persistence-configuration-unit kodo-broker

Function

Controls whether an object’s field values are cleared when entering a transaction. Valid values include:

- `datastore (0)`—Clear object field values when entering a datastore transaction. This is the default.
- `all (1)`—Clear object field values when entering any transaction.

Example

```
<auto-clear>datastore</auto-clear>
```

auto-detach

Range of values:	<code>auto-detaches</code> : <code>close</code> <code>commit</code> <code>nontx-read</code> <code>kodo-broker</code> : <code>0</code> <code>1</code> <code>2</code>
Default value:	<code>close (0)</code>
Parent elements:	<code>persistence-configuration</code> <code>persistence-configuration-unit</code> <code>auto-detaches</code> and <code>persistence-configuration</code> <code>persistence-configuration-unit</code> <code>kodo-broker</code> <code>auto-detach</code>

Function

Specifies the event on which the managed instances are automatically detached. When specified as a child of the `auto-detaches` element, valid event types include:

- `close (0)`—Detach all objects when the `PersistenceManager` closes. This is the default.
- `commit (1)`—Detach all objects when a transaction ends.
- `nontx-read (2)`—Reads outside of a transactions will automatically detach instances before returning them.

For more information, see “[Automatic Detachment](#)” in *Kodo Developer’s Guide*.

Example

```
<auto-detach>close</auto-detach>
```

auto-detaches

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Specifies a list of events on which the managed instances are automatically detached. Specifies one or more [auto-detach](#) elements. For more information, see “[Automatic Detachment](#)” in *Kodo Developer’s Guide*.

Example

```
<auto-detaches>  
  <auto-detach>close</auto-detach>  
</auto-detaches>
```


base-name

Range of values:	String
Default value:	n/a
Parent elements:	<p>persistence-configuration</p> <p> profiling</p> <p> export-profiling</p>

Function

Defines the basename of the exported data file to create. See “[Dumping Profiling Data to Disk from a Batch Process](#)” in *Kodo Developers Guide*.

Example

```
<base-name>base</base-name>
```

batching-operation-order-update-manager

Range of values:	n/a
Default value:	n/a
Parent elements:	<p>persistence-configuration</p> <p> persistence-configuration-unit</p>

Function

Defines an update manager capable of statement batching, that ignores foreign keys and table-level locking, but optimizes based on the order in which the application has modified objects. For more information, see “[kodo.jdbc.UpdateManager](#)” and “[Statement Batching](#)” in *Kodo Developer’s Guide*.

Example

```
<batching-operation-order-update-manager>
  <maximize-batch-size>true</maximize-batch-size>
</batching-operation-order-update-manager>
```

buffer-size

Range of values:	Integer
Default value:	10
Parent elements:	persistence-configuration persistence-configuration-unit cluster-remote-commit-provider

Function

Specifies the buffer size of the remote commit provider. For more information, see [“Remote Commit Provider Configuration”](#) in the *Kodo Developer’s Guide*.

Example

```
<buffer-size>10</buffer-size>
```

cache-map

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Defines the cache map used to associate query strings and their parsed form. For more information, see [“Query Compilation Cache”](#) in the *Kodo Developer’s Guide*.

Example

```
<cache-map>
  <cache-size>1000</cache-size>
  <soft-reference-size>-1</soft-reference-size>
</cache-map>
```

cache-size

Range of values:	Integer
Default value:	1000
Parent elements:	<p>persistence-configuration persistence-configuration-unit kodo-concurrent-data-cache</p> <p>and</p> <p>persistence-configuration persistence-configuration-unit lru-data-cache</p> <p>and</p> <p>persistence-configuration persistence-configuration-unit kodo-query-cache</p> <p>and</p> <p>persistence-configuration persistence-configuration-unit lru-query-cache</p>

Function

Set the data or query cache size. For more information, see [“Configuring the Data Cache Size”](#) and [“Query Cache”](#) in *Kodo Developer’s Guide*.

Example

```
<cache-size>1000</cache-size>
```

channel

Range of values:	String
Default value:	<code>openjpa.Runtime</code>
Parent elements:	<code>persistence-configuration</code> <code>persistence-configuration-unit</code>

Function

Specifies the channel to which you want to log. For more information, see [“Orphaned Keys”](#) in the *Kodo Developer’s Guide*.

Example

```
<channel>openjpa.Runtime</channel>
```

class-table-jdbc-seq

Range of values:	n/a
Default value:	n/a
Parent elements:	<code>persistence-configuration</code> <code>persistence-configuration-unit</code>

Function

Specifies a plugin string (see [“Plugin Configuration”](#) in *Kodo Developer’s Guide*) describing the `kodo.kernel.Seq` interface to use to create your own custom generators, in this case `kodo.jdbc.kernel.TableJDBCSeq`.

The `TableJDBCSeq` uses a special single-row table to store a global sequence number. If the table does not already exist, it is created the first time you run the mapping tool on a class that requires it. You can also use the class' `main` method or the `sequencetable` shell/bat script to manipulate the table; see the `TableJDBCSeq.main` method Javadoc for usage details.

For more information, see [“Generators”](#) in the *Kodo Developer’s Guide*.

Example

```
<class-table-jdbc-seq>
  <type>0</type>
  <allocate>50</allocate>
  <table-name>OPENJPA_SEQUENCES_TABLE</table-name>
  <ignore-virtual>false</ignore-virtual>
  <ignore-unmapped>false</ignore-unmapped>
  <table>OPENJPA_SEQUENCES_TABLE</table>
  <primary-key-column>ID</primary-key-column>
  <use-aliases>false</use-aliases>
  <sequence-column>SEQUENCE_VALUE</sequence-column>
  <increment>1</increment>
</class-table-jdbc-seq>
```

classname

Range of values:	Valid classname
Default value:	n/a
Parent elements:	<p> persistence-configuration persistence-configuration-unit aggregate-listener custom-aggregate-listener </p> <p>and</p> <p> persistence-configuration persistence-configuration-unit custom-broker-factory </p> <p>and</p> <p> persistence-configuration persistence-configuration-unit custom-broker-impl </p> <p>and</p> <p> persistence-configuration persistence-configuration-unit custom-class-resolver </p> <p>and</p> <p> persistence-configuration persistence-configuration-unit custom-compatibility </p> <p>and</p> <p> persistence-configuration persistence-configuration-unit custom-connection-decorator </p> <p>and</p> <p> persistence-configuration persistence-configuration-unit custom-data-cache </p> <p>and</p> <p> persistence-configuration persistence-configuration-unit custom-dictionary </p>

Parent elements:	and
(cont)	persistence-configuration persistence-configuration-unit custom-driver-data-source
	and
	persistence-configuration persistence-configuration-unit custom-filter-listener
	and
	persistence-configuration persistence-configuration-unit custom-jdbc-listener
	and
	persistence-configuration persistence-configuration-unit custom-lock-manager
	and
	persistence-configuration persistence-configuration-unit custom-logType
	and
	persistence-configuration persistence-configuration-unit custom-mapping-defaults
	and
	persistence-configuration persistence-configuration-unit custom-meta-data-respository
	and
	persistence-configuration persistence-configuration-unit custom-orphaned-key-action
	and
	persistence-configuration persistence-configuration-unit custom-persistence-server

Parent elements: and
(cont) persistence-configuration
 persistence-configuration-unit
 custom-proxy-manager

and
persistence-configuration
 persistence-configuration-unit
 custom-query-cache

and
persistence-configuration
 persistence-configuration-unit
 custom-query-compilation-cache

and
persistence-configuration
 persistence-configuration-unit
 custom-remote-commit-provider

and
persistence-configuration
 persistence-configuration-unit
 custom-savepoint-manager

and
persistence-configuration
 persistence-configuration-unit
 custom-seq

and
persistence-configuration
 persistence-configuration-unit
 custom-sql-factory

and
persistence-configuration
 persistence-configuration-unit
 custom-update-manager

Function

Name of the class associated with the custom feature.

classpath-scan

Range of values:	true false
Default value:	false
Parent elements:	<p> persistence-configuration persistence-configuration-unit deprecated-jdo-meta-data-factory and persistence-configuration persistence-configuration-unit extension-deprecated-jdo-mapping-factory and persistence-configuration persistence-configuration-unit jdo-meta-data-factory and persistence-configuration persistence-configuration-unit jdor-mapping-factory and persistence-configuration persistence-configuration-unit kodo-persistence-mapping-factory and persistence-configuration persistence-configuration-unit kodo-persistence-meta-data-factory and persistence-configuration persistence-configuration-unit mapping-file-deprecated-jdo-mapping-factory </p>

Parent elements:	and
(cont)	<p>persistence-configuration</p> <p> persistence-configuration-unit</p> <p> orm-file-jdor-mapping-factory</p> <p>and</p> <p>persistence-configuration</p> <p> persistence-configuration-unit</p> <p> table-deprecated-jdo-mapping-factory</p>

Function

Specifies a semicolon-separated list of directories or jar archives listed in your classpath. Each directory and jar archive is scanned for annotated JP entities or JDO metadata files based on your metadata factory. For more information, see [“Metadata Factory”](#) in *Kodo Developer’s Guide*.

Example

```
<classpath-scan>build</classpath-scan>
```

clear-on-close

Range of values:	true false
Default value:	false
Parent elements:	<p>persistence-configuration</p> <p> persistence-configuration-unit</p> <p> tangosol-data-cache</p>

Function

Flag that specifies whether the Tangosol named cache should be completely cleared when the `EntityManagerFactory` or `PersistenceManagerFactory` is closed.

Example

```
<clear-on-close>>false</clear-on-close>
```

client-broker-factory

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

A plugin string (see “[Plugin Configuration](#)” in *Kodo Developer’s Guide*) describing the `kodo.kernel.BrokerFactory` type to use, in this case `remote`. For more information, see “[kodo.BrokerFactory](#)” in the *Kodo Developer’s Guide*.

Example

```
<client-broker-factory/>
```

close-on-managed-commit

Range of values:	true false
Default value:	true
Parent elements:	persistence-configuration persistence-configuration-unit compatibility

Function

Flag that specifies whether the JDO `PersistenceManager` closes after a managed transaction commits, assuming you have invoked the `close` method.

If set to `false`, then the `PersistenceManager` will not close. This means that objects passed to a processing tier in the same JVM will still be usable, as their owning `PersistenceManager` will still be open. This behavior is not in strict compliance with the JDO specification, but is convenient for applications that were coded against Kodo 2.x, which did not close the `PersistenceManager` in these situations.

For more information, see [“Compatibility Configuration”](#) in *Kodo Developer’s Guide*.

Example

```
<close-on-managed-commit>>false</close-on-managed-commit>
```

cluster-remote-commit-provider

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Configures Kodo to share commit notifications among `persistenceManagerFactories` in a cluster. For more information, see [“Remote Commit Provider Configuration”](#) in the *Kodo Developer’s Guide*.

Example

```
<cluster-remote-commit-provider>  
  <name>kodo.RemoteCommitProvider</name>  
  <buffer-size>10</buffer-size>  
  <recover-action>none</recover-action>  
</cluster-remote-commit-provider>
```

commons-log-factory

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Specifies the Apache Jakarta Commons Logging thin library for issuing log messages. The Commons Logging libraries act as a wrapper around a number of popular logging APIs, including the Jakarta Log4J project, and the native `java.util.logging` package in JDK 1.4. If neither of these libraries are available, then logging will fall back to using simple console logging.

For more information, see [“Apache Commons Logging”](#) in *Kodo Developer’s Guide*.

Example

```
<commons-log-factory/>
```

compatibility

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Configure backwards compatibility settings for certain runtime behaviors. For more information, see [“Compatibility Configuration”](#) in *Kodo Developer’s Guide*.

Example

```
<compatibility>
  <copy-object-ids>false</copy-object-ids>
  <close-on-managed-commit>true</close-on-managed-commit>
  <validate-true-checks-store>false</validate-true-checks-store>
  <validate-false-returns-hollow>true</validate-false-returns-hollow>
  <strict-identity-values>false</strict-identity-values>
  <quoted-numbers-in-queries>false</quoted-numbers-in-queries>
</compatibility>
```

concurrent-hash-map

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Defines the concurrent hash map used to associate query strings and their parsed form. For more information, see [“Query Compilation Cache”](#) in the *Kodo Developer’s Guide*.

Example

```
<concurrent-hash-map>
  <max-size>2147483647</max-size>
</concurrent-hash-map>
```

connection-decorators

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

A plugin string (see “[Plugin Configuration](#)” in *Kodo Developer’s Guide*) describing [org.apache.openjpa.lib.jdbc.ConnectionDecorator](#) implementations to install on the connection factory. Specify one or more [custom-connection-decorator](#) elements.

These decorators can wrap connections passed from the underlying `DataSource` to add functionality. Kodo will pass all connections through the list of decorators before using them.

Note: By default Kodo JPA/JDO employs all of the built-in decorators in the `com.solarmetric.jdbc` package already; you do not need to list them here.

For more information, see “[kodo.jdbc.ConnectionDecorators](#)” in *Kodo Developer’s Guide*.

Example

```
<connection-decorators>
  <custom-connection-decorator>
    <classname>...</classname>
    <properties>...</properties>
  </custom-connection-decorator>
</connection-decorators>
```

connection-driver-name

Range of values:	Valid classname
Default value:	n/a
Parent elements:	<p>persistence-configuration persistence-configuration-unit and persistence-configuration persistence-configuration-unit kodo-pooling-data-source and persistence-configuration persistence-configuration-unit simple-driver-data-source</p>

Function

Full class name of either the JDBC `java.sql.Driver` or a `javax.sql.DataSource` implementation to use to connect to the database. For more information, see “[JDBC](#)” in *Kodo Developer’s Guide*.

The following provides sample values:

```
COM.FirstSQL.Dbcp.DbcpDriver
COM.cloudscape.core.JDBCdriver
COM.ibm.db2.jdbc.app.DB2Driver
COM.ibm.db2.jdbc.net.DB2Driver
centura.java.sqlbase.SqlbaseDriver
com.ddtek.jdbc.db2.DB2Driver
com.ddtek.jdbc.oracle.OracleDriver|
com.ddtek.jdbc.sqlserver.SQLServerDriver
com.ddtek.jdbc.sybase.SybaseDriver
com.ibm.as400.access.AS400JDBCdriver
com.imaginary.sql.mssql.MssqlDriver
com.inet.tds.TdsDriver
com.informix.jdbc.IfxDriver
com.internetcds.jdbc.tds.Driver
```



```

com.jnetdirect.jdbc.JSQLDriver
com.mckoi.JDBCdriver
com.microsoft.jdbc.sqlserver.SQLServerDriver
com.mysql.jdbc.DatabaseMetaData
com.mysql.jdbc.Driver
com.pointbase.jdbc.jdbcUniversalDriver
com.sap.dbtech.jdbc.DriverSapDB
com.sybase.jdbc.SybDriver
com.sybase.jdbc2.jdbc.SybDriver
com.thinweb.tds.Driver
in.co.daffodil.db.jdbc.DaffodilDBDriver
interbase.interclient.Driver
intersolv.jdbc.sequelink.SequeLinkDriver
openlink.jdbc2.Driver
oracle.jdbc.driver.OracleDriver
oracle.jdbc.pool.OracleDataSource
org.axiondb.jdbc.AxionDriver
org.enhydra.instantdb.jdbc.idbDriver
org.gjt.mm.mysql.Driver
org.hsql.jdbcDriver
org.hsqldb.jdbcDriver
org.postgresql.Driver
org.sourceforge.jxdmcon.JXDBConDriver|
postgres95.PGDriver
postgresql.Driver
solid.jdbc.SolidDriver
sun.jdbc.odbc.JdbcOdbcDriver
weblogic.jdbc.mssqlserver4.Driver
weblogic.jdbc.pool.Driver

```

Example

```

<connection-driver-name>oracle.jdbc.driver.OracleDriver</connection-driver-
-name>

```

connection-factory-mode

Range of values:	local managed
Default value:	local
Parent elements:	persistence-configuration persistence-configuration-unit

Function

The connection factory mode to use when integrating with the application server's managed transactions. Valid options include:

- `local`—Specifies a standard data source under Kodo control.
- `managed`—Specifies a data source managed by an application server and automatically enlisted in global transactions.

For more information, see [“Managed and XA DataSources”](#) in *Kodo Developer’s Guide*.

Example

```
<connection-factory-mode>local</connection-factory-mode>
```

connection-factory-name

Range of values:	JNDI name
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

The JNDI location of a `javax.sql.DataSource` to use to connect to the database. For more information, see [“JDBC”](#) in *Kodo Developer’s Guide*.

Example

```
<connection-factory-name>java:/OracleSource</connection-factory-name>
```

connection-factory-properties

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

A plugin string (see “[Plugin Configuration](#)” in *Kodo Developer’s Guide*) listing properties for configuration of the `javax.sql.DataSource` in use. For more information, see “[JDBC](#)” in *Kodo Developer’s Guide*.

Example

```
<connection-factory-properties>
  <property>
    <name>QueryTimeout</name>
    <value>5000</value>
    <name>MaxIdle</name>
    <value>1</value>
  </property>
</connection-factory-properties>
```

connection-factory2-name

Range of values:	JNDI name
Default value:	n/a
Parent elements:	<p>persistence-configuration</p> <p>persistence-configuration-unit</p>

Function

The JNDI location of an unmanaged `javax.sql.DataSource` to use to connect to the database. For more information, see [“XA Transactions”](#) in *Kodo Developer’s Guide*.

Example

```
<connection-factory2-name>java://OracleXASource</connection-factory2-name>
```

connection-factory2-properties

Range of values:	n/a
Default value:	n/a
Parent elements:	<p>persistence-configuration</p> <p>persistence-configuration-unit</p>

Function

Properties used to configure the `javax.sql.DataSource` used as the unmanaged `ConnectionFactory`. For more information, see [“Managed and XA DataSources”](#) in *Kodo Developer’s Guide*.

Example

```
<connection-factory2-properties>
  <property>
    <name>MaxActive</name>
```

```

    <value>20</value>
    <name>MaxIdle</name>
    <value>1</value>
  </property>
</connection-factory2-properties>

```

connection-password

Range of values:	Valid password
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Password for the user specified by [connection-user-name](#). For more information, see “[JDBC](#)” in *Kodo Developer’s Guide*.

Example

```
<connection-password>tiger</connection-password>
```

connection-properties

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

A plugin string (see “[Plugin Configuration](#)” in *Kodo Developer’s Guide*) listing properties to configure the driver listed in the [connection-driver-name](#) element. If the given driver class is

a `DataSource`, these properties will be used to configure the bean properties of the `DataSource`. For more information, see [“JDBC”](#) in *Kodo Developer’s Guide*.

Example

```
<connection-properties>
  <property>
    <name>PortNumber</name>
    <value>1521</value>
  </property>
  <property>
    <name>ServerName</name>
    <value>saturn</value>
  </property>
  <property>
    <name>DatabaseName</name>
    <value>solarisid</value>
  </property>
  <property>
    <name>DriverType</name>
    <value>thin</value>
  </property>
</connection-properties>
```

connection-retain-mode

Range of values:	always on-demand persistence-manager transaction
Default value:	on-demand
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Controls how Kodo uses datastore connections. Valid values include:

- `always`—Each `EntityManager` or `PersistenceManager` obtains a single connection and uses it until the `EntityManager` or `PersistenceManager` closes.
- `on-demand`—Connection is obtained only when needed. This option is equivalent to the `transaction` option when datastore transactions are used. For optimistic transactions, though, it means that a connection will be retained only for the duration of the datastore flush and commit process.
- `persistence-manager`—Connection is obtained from the persistence manager.
- `transaction`—Connection is obtained when each transaction begins (optimistic or datastore), and is released when the transaction completes. Non-transactional connections are obtained on-demand.

For more information, see [“Configuring the Use of JDBC Connections”](#) in *Kodo Developer’s Guide*.

Example

```
<connection-retain-mode>on-demand</connection-retain-mode>
```

connection-url

Range of values:	Valid URL
Default value:	n/a
Parent elements:	<p><code>persistence-configuration</code> <code>persistence-configuration-unit</code> and <code>persistence-configuration</code> <code>persistence-configuration-unit</code> <code>kodo-pooling-data-source</code> and <code>persistence-configuration</code> <code>persistence-configuration-unit</code> <code>simple-driver-data-source</code></p>

Function

The JDBC URL for the database. For more information, see [“JDBC”](#) in *Kodo Developer’s Guide*.

Persistence Configuration Schema Reference

The following provides possible values:

```
jdbc:JSQLConnect://<hostname>/database=<database>
jdbc:cloudscape:<database>;create=true
jdbc:daffodilDB_embedded:<database>;create=true
jdbc:datadirect:db2://<hostname>:50000;databaseName=<database>
jdbc:datadirect:oracle://<hostname>:1521;SID=<database>;MaxPooledStatements=0
jdbc:datadirect:sqlserver://<hostname>:1433;SelectMethod=cursor;DatabaseName=<database>
jdbc:datadirect:sybase://<hostname>:5000
jdbc:db2://<hostname>/<database>
jdbc:dbaw://<hostname>:8889/<database>
jdbc:hsqldb:<database>
jdbc:idb:<database>.properties
jdbc:inetdae:<hostname>:1433
jdbc:informix-sqli://<hostname>:1526/<database>;INFORMIXSERVER=<database>
jdbc:interbase://<hostname>//<database>.gdb
jdbc:microsoft:sqlserver://<hostname>:1433;DatabaseName=<database>;SelectMethod=cursor
jdbc:mysql://<hostname>/<database>?autoReconnect=true
jdbc:odbc:<database>
jdbc:openlink://<hostname>/DSN=SQLServerDB/UID=sa/PWD=
jdbc:oracle:thin:@<hostname>:1521:<database>
jdbc:postgresql://<hostname>:5432/<database>
jdbc:postgresql:net//<hostname>/<database>
jdbc:sequelink://<hostname>:4003/[Oracle]
jdbc:sequelink://<hostname>:4004/[Informix];Database=<database>
jdbc:sequelink://<hostname>:4005/[Sybase];Database=<database>
jdbc:sequelink://<hostname>:4006/[SQLServer];Database=<database>
jdbc:sequelink://<hostname>:4011/[ODBC MS Access];Database=<database>
jdbc:solid://<hostname>:<port>/<UID>/<PWD>
jdbc:sybase:Tds:<hostname>:4100/<database>?ServiceName=<database>
jdbc:twtds:sqlserver://<hostname>/<database>
jdbc:weblogic:mssqlserver4:<database>@<hostname>:1433
```


Example

```
<connection-url>
jdbc:oracle:thin:@<hostname>:1521:<database>
</connection-url>
```

connection-user-name

Range of values:	Valid username
Default value:	n/a
Parent elements:	<p>persistence-configuration persistence-configuration-unit and persistence-configuration persistence-configuration-unit kodo-pooling-data-source and persistence-configuration persistence-configuration-unit simple-driver-data-source</p>

Function

Username to use when connecting to the data source specified by [connection-url](#). For more information, see “[JDBC](#)” in *Kodo Developer’s Guide*.

Example

```
<connection-user-name>scott</connection-user-name>
```

connection2-driver-name

Range of values:	Valid classname
Default value:	n/a
Parent elements:	<p>persistence-configuration</p> <p>persistence-configuration-unit</p>

Function

Full class name of either the JDBC `java.sql.Driver` or a `javax.sql.DataSource` implementation to use to connect to the unmanaged database. For more information, see [“Managed and XA DataSources”](#) in *Kodo Developer’s Guide*.

Example

```
<connection2-driver-name>oracle.jdbc.driver.OracleDriver</connection2-driver-name>
```

connection2-password

Range of values:	Valid password
Default value:	n/a
Parent elements:	<p>persistence-configuration</p> <p>persistence-configuration-unit</p>

Function

Password for the unmanaged data source specified by [connection2-url](#).

Example

```
<connection2-password>tiger</connection2-password>
```

connection2-properties

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Properties for the unmanaged data source specified by `connection2-driver-name`. For more information, see “[Managed and XA DataSources](#)” in *Kodo Developer’s Guide*.

Example

```
<connection2-properties>
  <property>
    <name>PortNumber</name>
    <value>1521</value>
  </property>
  <property>
    <name>ServerName</name>
    <value>saturn</value>
  </property>
  <property>
    <name>DatabaseName</name>
    <value>solarisid</value>
  </property>
  <property>
    <name>DriverType</name>
    <value>thin</value>
  </property>
</connection2-properties>
```

connection2-url

Range of values:	Valid URL
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

The JDBC URL for the unmanaged datasource. For more information, see [“Managed and XA DataSources”](#) in *Kodo Developer’s Guide*.

Example

```
<connection2-url>jdbc:oracle:thin:@CROM:1521:KodoDB</connection2-url>
```

connection2-user-name

Range of values:	Valid username
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Username to use when connecting to the data source specified by [connection2-driver-name](#).

Example

```
<connection2-user-name>scott</connection2-user-name>
```

constraint-names

Range of values:	true false
Default value:	false
Parent elements:	persistence-configuration persistence-configuration-unit table-jdor-mapping-factory

Function

Flag that specifies whether to include the names of foreign key and unique constraints in all generated mappings. For more information, see [“Mapping Factory”](#) in *Kodo Developer’s Guide*.

Example

```
<constraint-names>false</constraint-names>
```

constraint-update-manager

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Defines the standard update manager, capable of statement batching and foreign key constraint evaluation. For more information, see [“kodo.jdbc.UpdateManager”](#) and [“Statement Batching”](#) in *Kodo Developer’s Guide*.

Example

```
<constraint-update-manager>
  <maximize-batch-size>true</maximize-batch-size>
</constraint-update-manager>
```

copy-object-ids

Range of values:	true false
Default value:	false
Parent elements:	persistence-configuration persistence-configuration-unit compatibility

Function

Flag that specifies whether to copy `oid` objects before returning them to you. Kodo versions prior to 3.0 always copied `oid` objects before returning them to you. This prevents possible errors resulting from you mutating the `oid` object by reference, but was not very efficient for the majority of users who do not modify `oid` instances. Thus Kodo 3.0 and higher does not copy `oids` by default. Set this property to `true` to force Kodo to copy the `oids`.

For more information, see “[Compatibility Configuration](#)” in *Kodo Developer’s Guide*.

Example

```
<copy-object-ids>false</copy-object-ids>
```

custom-aggregate-listener

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit aggregate-listeners

Function

Defines a custom aggregate listener. Each listener must implement the `org.apache.openjpa.kernel.exps.AggregateListener` interface. For more information, see “[kodo.AggregateListeners](#)” in *Kodo Developer’s Guide*.

Example

```
<aggregate-listeners>
  <custom-aggregate-listener>
    <classname>...</classname>
    <properties>...</properties>
  </custom-aggregate-listener>
</aggregate-listener>
```

custom-broker-factory

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Enables you to specify a custom broker factory. For more information, see “[kodo.BrokerFactory](#)” in *Kodo Developer’s Guide*.

Example

```
<custom-broker-factory>
  <classname>...</classname>
  <properties>...</properties>
</custom-broker-factory>
```

custom-broker-impl

Range of values:	n/a
Default value:	None
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Enables you to define a custom broker implementation. For more information, see [“kodo.BrokerImpl”](#) in *Kodo Developer’s Guide*.

Example

```
<custom-broker-impl>
  <classname>...</classname>
  <properties>...</properties>
</custom-broker-impl>
```

custom-class-resolver

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Enables you to define a custom class resolver for class name resolution. For more information, see [“kodo.ClassResolver”](#) in *Kodo Developer’s Guide*.

Example

```
<custom-class-resolver>
  <classname>...</classname>
  <properties>...</properties>
</custom-class-resolver>
```

custom-compatibility

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Enables you to define your own custom compatibility flag. For more information, see [“Compatibility Configuration”](#) in *Kodo Developer’s Guide*.

Example

```
<custom-compatibility>
  <classname>...</classname>
  <properties>...</properties>
</custom-compatibility>
```

custom-connection-decorator

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit connection-decorators

Function

Enables you to define a custom `org.apache.openjpa.lib.jdbc.ConnectionDecorator` implementation to install on all connection pools. For more information, see “[kodo.jdbc.ConnectionDecorators](#)” in *Kodo Developer’s Guide*.

Example

```
<connection-decorators>
  <custom-connection-decorator>
    <classname>...</classname>
    <properties>...</properties>
  </custom-connection-decorator>
</connection-decorators>
```

custom-data-cache

Range of values:	n/a
Default value:	n/a
Parent elements:	<p>persistence-configuration persistence-configuration-unit data-caches</p>

Function

Enables you to define a custom data cache. For more information, see “[kodo.DataCache](#)” in *Kodo Developer’s Guide*.

Example

```
<data-caches>
  <custom-data-cache>
    <name>...</name>
    <classname>...</classname>
    <properties>...</properties>
  </custom-data-cache>
</data-caches>
```

custom-data-cache-manager

Range of values:	n/a
Default value:	n/a
Parent elements:	<p>persistence-configuration persistence-configuration-unit</p>

Function

Enables you to define a custom data cache manager. For more information, see “[kodo.DataCacheManager](#)” in *Kodo Developer’s Guide*.

Example

```
<custom-data-cache-manager>
  <classname>...</classname>
  <properties>...</properties>
</custom-data-cache-manager>
```

custom-detach-state

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Enables you to define a custom detach state. For more information, see “[kodo.DetachState](#)” in *Kodo Developer’s Guide*.

Example

```
<custom-detach-state>
  <classname>...</classname>
  <properties>...</properties>
</custom-detach-state>
```

custom-dictionary

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Enables you to define a custom dictionary. For more information, see [“Custom Dictionary Configuration”](#) in the *Administration Console Online Help*.

Example

```
<custom-dictionary>
  <name>custom-dictionary</name>
  <classname>..</classname>
  <properties>...</properties>
</custom-dictionary>
```

custom-driver-data-source

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Enables you to define a custom datasource driver. For more information, see [“kodo.jdbc.DriverDataSource”](#) in *Kodo Developer’s Guide*.

Example

```
<custom-driver-data-source>
  <classname>...</classname>
  <properties>...</properties>
</custom-driver-data-source>
```

custom-filter-listener

Range of values:	n/a
-------------------------	-----

Default value:	n/a
-----------------------	-----

Parent elements:	persistence-configuration persistence-configuration-unit filter-listeners
-------------------------	---

Function

Enables you to define a custom filter listener. For more information, see [“kodo.FilterListeners”](#) in *Kodo Developer’s Guide*.

Example

```
<filter-listeners>
  <custom-filter-listener>
    <classname>...</classname>
    <properties>...</properties>
  </custom-filter-listener>
</filter-listeners>
```

custom-jdbc-listener

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit jdbc-listeners

Function

Enables you to define a custom JDBC listener. For more information, see [“kodo.JDBCListeners”](#) in *Kodo Developer’s Guide*.

Example

```
<jdbc-listeners>
  <custom-jdbc-listener>
    <classname>...</classname>
    <properties>...</properties>
  </custom-jdbc-listener>
</jdbc-listeners>
```

custom-lock-manager

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Enables you to define a custom lock manager. For more information, see [“Lock Manager”](#) in *Kodo Developer’s Guide*.

Example

```
<custom-lock-manager>
  <classname>...</classname>
  <properties>...</properties>
</custom-lock-manager>
```

custom-log

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Enables you to define a custom log file. For more information, see [“Custom Log”](#) in *Kodo Developer’s Guide*.

Example

```
<custom-log>
  <classname>...</classname>
  <properties>...</properties>
</custom-log>
```


custom-mapping-defaults

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Enables you to define custom mapping defaults. For more information, see [“Mapping Factory”](#) in *Kodo Developer’s Guide*.

Example

```
<custom-mapping-defaults>
  <classname>...</classname>
  <properties>...</properties>
</custom-mapping-defaults>
```

custom-mapping-factory

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Specifies a custom mapping factory. For more information, see [“Mapping Factory”](#) in *Kodo Developer’s Guide*.

Example

```
<custom-mapping-factory/>
```

custom-meta-data-factory

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Enables you to specify a custom metadata factory. For more information, see [“Metadata Factory”](#) in the *Kodo Developer’s Guide*.

Example

```
<custom-meta-data-factory>
  <classname>...</classname>
  <properties>...</properties>
</custom-meta-data-factory>
```

custom-meta-data-repository

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Enables you to specify a custom metadata repository. For more information, see [“Metadata Factory”](#) in the *Kodo Developer’s Guide*.

Example

```
<custom-meta-data-repository>
  <classname>...</classname>
  <properties>...</properties>
</custom-meta-data-repository>
```

custom-orphaned-key-action

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Enables you to define a custom orphaned key action. For more information, see [“Orphaned Keys”](#) in the *Kodo Developer’s Guide*.

Example

```
<custom-orphaned-key-action>
  <channel>openjpa.Runtime</channel>
  <level>4</level>
  <classname>...</classname>
  <properties>...</properties>
</custom-orphaned-key-action>
```

custom-persistence-server

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Enables you to define a custom persistence server. For more information, see [“Standalone Persistence Server”](#) in the *Kodo Developer’s Guide*.

Example

```
<custom-persistence-server>
  <classname>...</classname>
  <properties>...</properties>
</custom-persistence-server>
```

custom-proxy-manager

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Enables you to define a custom proxy manager. For more information, see [“Custom Proxies”](#) in the *Kodo Developer’s Guide*.

Example

```
<custom-proxy-manager>
  <classname>...</classname>
  <properties>...</properties>
</custom-proxy-manager>
```

custom-query-compilation-cache

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Enables you to define a custom query compilation cache. For more information, see [“Query Compilation Cache”](#) in the *Kodo Developer’s Guide*.

Example

```
<custom-query-compilation-cache>
  <classname>...</classname>
  <properties>...</properties>
</custom-query-compilation-cache>
```

custom-remote-commit-provider

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Enables you to specify a custom remote commit provider. For more information, see “[JMS](#)” in the *Kodo Developer’s Guide*.

Example

```
<custom-remote-commit-provider>
  <name>kodo.RemoteCommitProvider</name>
  <classname>...</classname>
  <properties>...</properties>
</custom-remote-commit-provider>
```

custom-savepoint-manager

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Enables you to specify a custom savepoint manager. For more information, see “[Savepoints](#)” in the *Kodo Developer’s Guide*.

Example

```
<custom-savepoint-manager>
  <classname>...</classname>
  <properties>...</properties>
</custom-savepoint-manager>
```

custom-schema-factory

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Specifies a custom schema factory. For more information, see [“Schema Factory”](#) in the *Kodo Developer’s Guide*.

Example

```
<custom-schema-factory>
  <classname>...</classname>
  <properties>...</properties>
</custom-schema-factory>
```

custom-seq

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Enables you to define a custom generator. For more information, see [“Generators”](#) in the *Kodo Developer’s Guide*.

Example

```
<custom-seq>
  <classname>...</classname>
  <properties>...</properties>
</custom-seq>
```

custom-sql-factory

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Enables you to define a custom SQL factory. For more information, see [“kodo.jdbc.SQLFactory”](#) in the *Kodo Developer’s Guide*.

Example

```
<custom-sql-factory>
  <classname>...</classname>
  <properties>...</properties>
</custom-sql-factory>
```

custom-update-manager

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Enables you to define a custom update manager to use to flush persistent object changes to the datastore. For more information, see “[kodo.jdbc.UpdateManager](#)” in *Kodo Developer’s Guide*.

Example

```
<custom-update-manager>
  <classname>..</classname>
  <properties>...</properties>
</custom-update-manager>
```

data-caches

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Specifies plugins used to cache data loaded from the data store. Must implement org.apache.openjpa.datacache.DataCache. For more information, see “Data Cache” in *Kodo Developer’s Guide*.

Example

```
<data-caches>
  <default-data-cache>...</default-data-cache>
  <kodo-concurrent-data-cache>...</kodo-concurrent-data-cache>
  <gem-fire-data-cache>...</gem-fire-data-cache>
  <lru-data-cache>...</lru-data-cache>
  <tangosol-data-cache>...</tangosol-data-cache>
  <custom-data-cache>...</custom-data-cache>
</data-caches>
```

data-cache-manager-impl

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Enables the `kodo.datacache.DataCacheManager` that manages the system data caches. For more information, see [“kodo.DataCacheManager”](#) in *Kodo Developer’s Guide*.

Example

```
<data-cache-manager-impl/>
```

data-cache-timeout

Range of values:	Integer
Default value:	-1
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Specifies the amount of time in milliseconds that a class’ data is valid. A value of -1 specifies no expiration. See [“Setting the Data Cache Timeout Value”](#) in *Kodo Developer’s Guide*.

Example

```
<data-cache-timeout>-1</data-cache-timeout>
```

db2-dictionary

Range of values:	n/a
Default value:	n/a
Parent elements:	<p>persistence-configuration</p> <p>persistence-configuration-unit</p>

Function

Defines the configuration values for the DB2 Dictionary persistence plugin. For a complete description of each of the values that you can specify, see [“DB2 Dictionary Configuration”](#) in the *Administration Console Online Help*.

Example

The `db2-dictionary` element shares the same subelements as [“access-dictionary”](#) on page B-7.

default-access-type

Range of values:	FIELD PROPERTY
Default value:	FIELD
Parent elements:	<p>persistence-configuration</p> <p>persistence-configuration-unit</p> <p>kodo-persistence-mapping-factory</p> <p>and</p> <p>persistence-configuration</p> <p>persistence-configuration-unit</p> <p>kodo-persistence-meta-data-factory</p>

Function

Specifies the default access type. Valid values include:

- `FIELD`—Injects state directly into your persistent fields, and retrieves changed state from your fields as well.

- **PROPERTY**—Retrieves and loads state through JavaBean "getter" and "setter" methods.

For more information, see [“Field and Property Metadata”](#) in *Kodo Developer’s Guide*.

Example

```
<default-access-type>FIELD</default-access-type>
```

default-broker-factory

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Specifies a plugin string (see [“Plugin Configuration”](#) in *Kodo Developer’s Guide*) describing the `kodo.kernel.BrokerFactory` type to use, in this case the default, which is JDBC. For more information, see [“kodo.BrokerFactory”](#) in the *Kodo Developer’s Guide*.

Example

```
<default-broker-factory/>
```

default-broker-impl

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Specifies a plugin string (see “[Plugin Configuration](#)” in *Kodo Developer’s Guide*) describing the `kodo.kernel.Broker` type to use at runtime, in this case the default. For more information, see “[kodo.BrokerImpl](#)” in *Kodo Developer’s Guide*.

Example

```
<default-broker-impl/>
```

default-class-resolver

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Specifies a plugin string (see “[Plugin Configuration](#)” in *Kodo Developer’s Guide*) describing the `kodo.util.ClassResolver` implementation to use for class name resolution, in this case the default. You may wish to plug in your own resolver if you have special class loading needs. For more information, see “[kodo.ClassResolver](#)” in *Kodo Developer’s Guide*.

Example

```
<default-class-resolver/>
```

default-compatibility

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Sets the default compatibility settings. For more information, see [“Compatibility Configuration”](#) in *Kodo Developer’s Guide*.

Example

```
<default-compatibility/>
```

default-data-cache

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit data-caches

Function

Defines the default data cache. For more information, see [“Data Cache”](#) in *Kodo Developer’s Guide*.

Example

```
<data-caches>
  <default-data-cache>
    <name>kodo.DataCache</name>
```

```

    </default-data-cache>
</data-caches>

```

default-detach-state

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Specifies the default detached state, in this case `loaded`. For more information, see [“Defining the Detached Object Graph”](#) in *Kodo Developer’s Guide*.

Example

```
<default-detach-state/>
```

default-data-cache-manager

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Specifies a plugin string (see [“Plugin Configuration”](#) in *Kodo Developer’s Guide*) describing the `kodo.datacache.DataCacheManager` that manages the system data caches, in this case the default. For more information, see [“Data Cache”](#) in *Kodo Developer’s Guide*.

Example

```
<default-data-cache-manager>...</default-data-cache-manager>
```

default-driver-data-source

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Specifies the default `kodo.jdbc.schema.DriverDataSource` implementation to use to wrap JDBC Driver classes with `javax.sql.DataSource` instances. The provided default implementation, `kodo.jdbc.schema.KodoPoolingDataSource`, will perform connection pooling as described at “[JDBC](#)” in *Kodo Developer’s Guide*.

Example

```
<default-driver-data-source/>
```

default-level

Range of values:	TRACE DEBUG INFO WARN ERROR
Default value:	INFO
Parent elements:	persistence-configuration persistence-configuration-unit log-factory-impl

Function

Specifies the default logging level of unconfigured channels. For more information, see “[Kodo Logging](#)” in *Kodo Developer’s Guide*.

Example

```
<default-level>INFO</default-level>
```

default-lock-manager

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Specifies to use the default lock manager. For more information, see [“Configuring Default Locking”](#) in *Kodo Developer’s Guide*.

Example

```
<default-lock-manager/>
```

default-mapping-defaults

Range of values:	String Standard implementations include: jdo and jpa
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Specifies a plugin string (see [“Plugin Configuration”](#) in *Kodo Developer’s Guide*) describing the `kodo.jdbc.meta.MappingDefaults` to use to define default column names, table names, and constraints for your persistent classes. For more information and a list of standard implementations, see [“Mapping Defaults”](#) in *Kodo Developer’s Guide*.

Example

```
<default-mapping-defaults>jpa</default-mapping-defaults>
```

default-meta-data-factory

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Specifies a plugin string (see [“Plugin Configuration”](#) in *Kodo Developer’s Guide*) describing the `kodo.meta.MetaDataFactory` type to use, in this case the default, which is `kodo.jdo.JDOMetaDataFactory`. For more information, see [“Metadata Factory”](#) in the *Kodo Developer’s Guide*.

Example

```
<default-meta-data-factory/>
```

default-meta-data-repository

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Specifies a plugin string (see [“Plugin Configuration”](#) in *Kodo Developer’s Guide*) describing the `kodo.meta.MetadataRepository` to use to store the metadata for your persistent classes, in

this case the default,. For more information, see [“Metadata Factory”](#) in the *Kodo Developer’s Guide*.

Example

```
<default-meta-data-repository/>
```

default-orphaned-key-action

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Specifies a plugin string (see [“Plugin Configuration”](#) in *Kodo Developer’s Guide*) describing the `kodo.event.OrphanedKeyAction` to invoke when Kodo discovers an orphaned datastore key, in this case the default which is `kodo.event.LogOrphanedKeyAction`. In this case, Kodo logs a message for each orphaned key. For more information, see [“Orphaned Keys”](#) in the *Kodo Developer’s Guide*.

Example

```
<default-orphaned-key-action>
  <channel>openjpa.Runtime</channel>
  <level>4</level>
</default-orphaned-key-action>
```

default-proxy-manager

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Specifies a plugin string (see [“Plugin Configuration”](#) in *Kodo Developer’s Guide*) describing the `kodo.util.ProxyManager` to use for proxying mutable second class objects, in this case the default, which is `kodo.util.ProxyManagerImpl`. For more information, see [“Custom Proxies”](#) in the *Kodo Developer’s Guide*.

Example

```
<default-proxy-manager/>
```

default-query-compilation-cache

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Specifies a plugin string (see [“Plugin Configuration”](#) in *Kodo Developer’s Guide*) describing the Map used to associate query strings and their parsed form, in this case the default. For more information, see [“Query Compilation Cache”](#) in the *Kodo Developer’s Guide*.

Example

```
<default-query-compilation-cache/>
```

default-savepoint-manager

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Specifies a plugin string (see “[Plugin Configuration](#)” in *Kodo Developer’s Guide*) describing the `kodo.kernel.SavepointManager` to use for managing transaction savepoints, in this case the default which is `kodo.kernel.InMemorySavepointManager`. This plugin stores all state, including field values, in memory. Because of this behavior, each set savepoint is designed for small to medium transactional object counts. For more information, see “[Savepoints](#)” in the *Kodo Developer’s Guide*.

Example

```
<default-savepoint-manager/>
```

default-schema-factory

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Specifies a plugin string (see “[Plugin Configuration](#)” in *Kodo Developer’s Guide*) describing the `kodo.jdbc.schema.SchemaFactory` to use to store and retrieve information about the database schema, in this case the default which is `kodo.jdbc.schema.DynamicSchemaFactory`.

The `DynamicSchemaFactory` is the most performant schema factory, because it does not validate mapping information against the database. Instead, it assumes all object-relational

mapping information is correct, and dynamically builds an in-memory representation of the schema from your mapping metadata. When using this factory, it is important that your mapping metadata correctly represent your database's foreign key constraints so that Kodo can order its SQL statements to meet them.

For more information, see [“Schema Factory”](#) in the *Kodo Developer's Guide*.

Example

```
<default-schema-factory/>
```

default-sql-factory

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Specifies a plugin string (see [“Plugin Configuration”](#) in *Kodo Developer's Guide*) describing the `kodo.jdbc.SQLFactory` to use to abstract common SQL constructs, in this case the default. For more information, see [“kodo.jdbc.SQLFactory”](#) in the *Kodo Developer's Guide*.

Example

```
<default-meta-data-factory/>
```

default-update-manager

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Specifies to use the default update manager, `kodo.jdbc.kernel.ConstraintUpdateManager`. For more information, see [“kodo.jdbc.UpdateManager”](#) in *Kodo Developer’s Guide*.

Example

```
<default-update-manager/>
```

deprecated-jdo-mapping-defaults

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Specifies the mapping defaults for JDO 1.0. For more information, see [“Mapping Defaults”](#) in *Kodo Developer’s Guide*.

Example

The `deprecated-jdo-mapping-defaults` element shares the same subelements as [“mapping-defaults-impl”](#) on page B-150.

deprecated-jdo-meta-data-factory

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Specifies a plugin string (see [“Plugin Configuration”](#) in *Kodo Developer’s Guide*) describing the `kodo.meta.MetadataFactory` type to use, in this case `kodo.jdo.DeprecatedJDOMMetadataFactory`. For more information, see [“Metadata Factory”](#) in the *Kodo Developer’s Guide*.

Example

```
<deprecated-jdo-meta-data-factory>
  <use-schema-validation>false</use-schema-validation>
  <urls>t3://localhost:7001/metadata.jar</urls>
  <files>com/file1;com/file2</files>
  <classpath-scan>build</classpath-scan>
  <types>classname1;classname2</types>
  <store-mode>1</store-mode>
  <strict>false</strict>
  <resources>com/aaa/package.jdo;com/bbb/package.jdo</resources>
  <scan-top-down>false</scan-top-down>
</deprecated-jdo-meta-data-factory>
```

derby-dictionary

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Configures the Derby Dictionary. For a complete description of each of the values that you can specify, see [“Derby Dictionary Configuration”](#) in the *Administration Console Online Help*.

Example

The derby-dictionary element shares the same subelements as [“access-dictionary”](#) on [page B-7](#), plus the following subelement (default shown):

```
<derby-dictionary>
...
  <shutdown-on-close>true</shutdown-on-close>
</derby-dictionary>
```

detach-options-all

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Sets the detach state to `all` to detach all fields and relations. Be very careful when using this mode; if you have a highly-connected domain model, you could end up bringing every object in the database into memory. For more information, see [“Defining the Detached Object Graph”](#) in *Kodo Developer’s Guide*.

Example

```
<detach-options-all>
  <detached-state-manager>true</detached-state-manager>
  <detached-state-transient>false</detached-state-transient>
  <access-unloaded>true</access-unloaded>
  <detached-state-field>true</detached-state-field>
</detach-options-all>
```

detach-options-fetch-groups

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Sets the detach state to detach all fields and relations in the default fetch group, and any other fetch groups that you have added to the current fetch configuration. For more information on custom fetch groups, see [“Fetch Groups”](#) in *Kodo Developer’s Guide*. For more information about the detach state, see [“Defining the Detached Object Graph”](#) in *Kodo Developer’s Guide*.

Example

```
<detach-options-fetch-groups>
  <detached-state-manager>true</detached-state-manager>
  <detached-state-transient>false</detached-state-transient>
  <access-unloaded>true</access-unloaded>
  <detached-state-field>true</detached-state-field>
</detach-options-fetch-groups>
```

detach-options-loaded

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Sets the detach state to detach all fields and relations that are already loaded, but do not include unloaded fields in the detached graph. This is the default. For more information about the detach state, see [“Defining the Detached Object Graph”](#) in *Kodo Developer’s Guide*.

Example

```
<detach-options-loaded>
  <detached-state-manager>true</detached-state-manager>
  <detached-state-transient>false</detached-state-transient>
  <access-unloaded>true</access-unloaded>
  <detached-state-field>true</detached-state-field>
</detach-options-loaded>
```

detach-state

Range of values:	1 2 3
Default value:	1
Parent elements:	persistence-configuration persistence-configuration-unit kodo-broker

Function

Configures the detach state. Valid values include:

- 1—(loaded) Detach all fields and relations that are already loaded, but do not include unloaded fields in the detached graph.
- 2—(fetch-groups) Detach all fields and relations in the current fetch configuration.
- 3—(all) Detach all fields and relations. Be very careful when using this mode; if you have a highly-connected domain model, you could end up bringing every object in the database into memory.

Any field that is not included in the set determined by the detach mode is set to its Java default value in the detached instance. For more information, see [“Defining the Detached Object Graph”](#) in *Kodo Developer’s Guide*.

Example

```
<detach-state>1</detach-state>
```

detach-state-field

Range of values:	true false
Default value:	true
Parent elements:	<p>persistence-configuration persistence-configuration-unit detach-options-all</p> <p>and</p> <p>persistence-configuration persistence-configuration-unit detach-options-fetch-groups</p> <p>and</p> <p>persistence-configuration persistence-configuration-unit detach-options-loaded</p>

Function

Flag that specifies whether Kodo should take advantage of a detached state field to make the attach process more efficient. This field is added by the enhancer and is not visible to your application. For more information, see [“Attach Behavior”](#) in *Kodo Developer’s Guide*.

For more information about the detach state, see “[Defining the Detached Object Graph](#)” in *Kodo Developer’s Guide*.

Example

```
<detached-state-field>true</detached-state-field>
```

detached-state-manager

Range of values:	true false
Default value:	true
Parent elements:	<p>persistence-configuration persistence-configuration-unit detach-options-all</p> <p>and</p> <p>persistence-configuration persistence-configuration-unit detach-options-fetch-groups</p> <p>and</p> <p>persistence-configuration persistence-configuration-unit detach-options-loaded</p>

Function

Flag that specifies whether to use a detached state manager. A detached state manager makes attachment much more efficient. Like a detached state field, however, it breaks serialization compatibility with the unenhanced class if it is not transient.

This setting piggybacks on the `detached-state-field` element. If your detached state field is transient, the detached state manager will also be transient. If the detached state field is disabled, the detached state manager will also be disabled. This is typically what you want.

By setting the `detach-state-field` element to `true` (or `transient`) and setting this property to `false`, however, you can use a detached state field without using a detached state manager. This may be useful for debugging or for legacy Kodo users who find differences between Kodo's behavior with a detached state manager and Kodo's older behavior without one.

For more information about the detach state, see [“Defining the Detached Object Graph”](#) in *Kodo Developer’s Guide*.

Example

```
<detached-state-manager>true</detached-state-manager>
```

detached-state-transient

Range of values:	true false
Default value:	true
Parent elements:	<p>persistence-configuration persistence-configuration-unit detach-options-all</p> <p>and</p> <p>persistence-configuration persistence-configuration-unit detach-options-fetch-groups</p> <p>and</p> <p>persistence-configuration persistence-configuration-unit detach-options-loaded</p>

Function

Flag that specifies whether to use a transient detached state field. Setting this to true provides the benefits of a detached state field to local objects that are never serialized, but retains serialization compatibility for client tiers without access to the enhanced versions of your classes.

For more information about the detach state, see [“Defining the Detached Object Graph”](#) in *Kodo Developer’s Guide*.

Example

```
<detached-state-transient>false</detached-state-transient>
```

detached-new

Range of values:	true false
Default value:	true
Parent elements:	persistence-configuration persistence-configuration-unit kodo-broker

Function

Flag that specifies whether to use a non-transient detached state field so that objects crossing serialization barriers can still be attached efficiently. This requires, however, that your client tier have the enhanced versions of your classes and the Kodo libraries.

For more information about the detach state, see [“Defining the Detached Object Graph”](#) in *Kodo Developer’s Guide*.

Example

```
<detached-new>true</detached-new>
```

diagnostic-context

Range of values:	String
Default value:	kodo.ID property, if set
Parent elements:	persistence-configuration persistence-configuration-unit log-factory-impl

Function

String that is prepended to all log messages. The `kodo.ID` property is used by default, if specified. For more information, see [“Kodo Logging”](#) in *Kodo Developer’s Guide*.

Example

```
<diagnostic-context>KodoID</diagnostic-context>
```

dynamic-data-structs

Range of values:	true false
Default value:	false
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Flag that specifies whether to dynamically generate customized structs to hold persistent data. Both the Kodo data cache and the remote framework rely on data structs to cache and transfer persistent state. With dynamic structs, Kodo can customize data storage for each class, eliminating the need to generate primitive wrapper objects. This saves memory and speeds up certain runtime operations.

The cost is a longer warm-up time for the application—generating and loading custom classes into the JVM takes time. Therefore, only set this property to `true` if you have a long-running application where the initial cost of class generation is offset by memory and speed optimization over time.

For more information, see “[kodo.DynamicDataStructs](#)” in *Kodo Developer’s Guide*.

Example

```
<dynamic-data-structs>false</dynamic-data-structs>
```

dynamic-schema-factory

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Specifies a plugin string (see “[Plugin Configuration](#)” in *Kodo Developer’s Guide*) describing the `kodo.jdbc.schema.SchemaFactory` to use to store and retrieve information about the database schema, in this case the default which is `kodo.jdbc.schema.DynamicSchemaFactory`.

The `DynamicSchemaFactory` is the most performant schema factory, because it does not validate mapping information against the database. Instead, it assumes all object-relational mapping information is correct, and dynamically builds an in-memory representation of the schema from your mapping metadata. When using this factory, it is important that your mapping metadata correctly represent your database's foreign key constraints so that Kodo can order its SQL statements to meet them.

For more information, see “[Schema Factory](#)” in the *Kodo Developer’s Guide*.

Example

```
<dynamic-schema-factory/>
```

eager-fetch-mode

Range of values:	join multiple none parallel single
Default value:	parallel
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Configures eager fetching. Eager fetching is the ability to efficiently load subclass data and related objects along with the base instances being queried. For complete details, see [“Eager Fetching”](#) in *Kodo Developer’s Guide*.

You can set this value to one of the following:

- `join`—In this mode, Kodo joins to to-one relations in the configured fetch groups. If Kodo is loading data for a single instance, then Kodo will also join to any collection field in the configured fetch groups. When loading data for multiple instances, though, (such as when executing a Query) Kodo will not join to collections by default. Instead, Kodo defaults to parallel mode for collections, as described below. You can force Kodo to use a join rather than parallel mode for a collection field using the metadata extension described in [“Eager Fetch Mode”](#) in *Kodo Developer’s Guide*.

Under `join` mode, Kodo uses a left outer join (or inner join, if the relations' field metadata declares the relation non-nullable) to select the related data along with the data for the target objects. This process works recursively for to-one joins, so that if `Person` has an `Address`, and `Address` has a `PhoneNumber`, and the fetch groups are configured correctly, Kodo might issue a single select that joins across the tables for all three classes. To-many joins can not recursively spawn other to-many joins, but they can spawn recursive to-one joins.

Under the join subclass fetch mode, subclass data in joined tables is selected by outer joining to all possible subclass tables of the type being queried. Unjoined subclass data is selected with a SQL UNION where possible. As you'll see below, subclass data fetching is configured separately from relation fetching, and can be disabled for specific classes.

- `multiple`—Same as `parallel`.
- `none`—No eager fetching is performed. Related objects are always loaded in an independent select statement. No joined subclass data is loaded unless it is in the table(s)

for the base type being queried. Unjoined subclass data is loaded using separate select statements rather than an `SQL UNION` operation.

- `parallel`—Under this mode, Kodo selects to-one relations and joined collections as outlined in the join mode description above. Unjoined collection fields, however, are eagerly fetched using a separate select statement for each collection, executed in parallel with the select statement for the target objects. The parallel selects use the `WHERE` conditions from the primary select, but add their own joins to reach the related data. Thus, if you perform a query that returns 100 Company objects, where each company has a list of Employee objects and Department objects, Kodo will make three queries. The first will select the company objects, the second will select the employees for those companies, and the third will select the departments for the same companies. Just as for `joins`, this process can be recursively applied to the objects in the relations being eagerly fetched. Continuing our example, if the Employee class had a list of Projects in one of the fetch groups being loaded, Kodo would execute a single additional select in parallel to load the projects of all employees of the matching companies.

Using an additional select to load each collection avoids transferring more data than necessary from the database to the application. If eager joins were used instead of parallel select statements, each collection added to the configured fetch groups would cause the amount of data being transferred to rise dangerously, to the point that you could easily overwhelm the network.

Polymorphic to-one relations to table-per-class mappings use parallel eager fetching because proper joins are impossible. You can force other to-one relations to use parallel rather than join mode eager fetching using the metadata extension described in “[Eager Fetch Mode](#)” in *Kodo Developer’s Guide*.

Setting your subclass fetch mode to parallel affects table-per-class and vertical inheritance hierarchies. Under parallel mode, Kodo issues separate selects for each subclass in a table-per-class inheritance hierarchy, rather than `UNIONing` all subclass tables together as in join mode. This applies to any operation on a table-per-class base class: query, by-id lookup, or relation traversal.

When dealing with a vertically-mapped hierarchy, on the other hand, parallel subclass fetch mode only applies to queries. Rather than outer-joining to subclass tables, Kodo will issue the query separately for each subclass. In all other situations, parallel subclass fetch mode acts just like join mode in regards to vertically-mapped subclasses.

When Kodo knows that it is selecting for a single object only, it never uses parallel mode, because the additional selects can be made lazily just as efficiently. This mode only increases efficiency over join mode when multiple objects with eager relations are being loaded, or when multiple selects might be faster than joining to all possible subclasses.

- `single`—Same as `join`.

Example

```
<eager-fetch-mode>parallel</eager-fetch-mode>
```

empress-dictionary

Range of values:	n/a
Default value:	n/a
Parent elements:	<p>persistence-configuration</p> <p>persistence-configuration-unit</p>

Function

Configures the Empress Dictionary. For a complete description of each of the values that you can specify, see [“Empress Dictionary Configuration”](#) in the *Administration Console Online Help*.

Example

The `empress-dictionary` element shares the same subelements as [“access-dictionary”](#) on [page B-7](#), plus the following subelement (default shown):

```
<empress-dictionary>
...
  <allow-concurrent-read>false</allow-concurrent-read>
</empress-dictionary>
```

EnableLogMBean

Range of values: true | false

Default value:

Parent elements: persistence-configuration
persistence-configuration-unit
local-jmx

Function

Flag that specifies whether the `LogMBean` should be registered. For more information, see [“Optional Parameters in Management Group”](#) in *Kodo Developer’s Guide*.

Example

```
<EnableLogMBean>true</EnableLogMBean>
```

EnableRuntimeMBean

Range of values: true | false

Default value: n/a

Parent elements: persistence-configuration
persistence-configuration-unit
local-jmx

Function

Flag that specifies whether the `RuntimeMBean` should be registered. For more information, see [“Optional Parameters in Management Group”](#) in *Kodo Developer’s Guide*.

Example

```
<EnableRuntimeMBean>true</EnableRuntimeMBean>
```

evict-from-data-cache

Range of values:	true false
Default value:	false
Parent elements:	persistence-configuration persistence-configuration-unit kodo-broker

Function

Flag that specifies whether to evict data from the data cache at specific times. You defined the eviction schedule using the [eviction-schedule](#) element. For more information, see “[Data Cache](#)” in *Kodo Developer’s Guide*.

Example

```
<evict-from-data-cache>true</evict-from-data-cache>
```

eviction-schedule

Range of values:	n/a
-------------------------	-----

Default value:	n/a
-----------------------	-----

Parent elements:	<p>persistence-configuration persistence-configuration-unit data-caches gem-fire-data-cache</p> <p>and</p> <p>persistence-configuration persistence-configuration-unit data-caches kodo-concurrent-data-cache</p> <p>and</p> <p>persistence-configuration persistence-configuration-unit data-caches lru-data-cache</p> <p>and</p> <p>persistence-configuration persistence-configuration-unit data-caches tangosol-data-cache</p>
-------------------------	--

Function

You can clear a data cache at specific times. The `EvictionSchedule` property of Kodo's cache implementation accepts a cron style eviction schedule. The format of this property is a whitespace-separated list of five tokens. You can use the `*` symbol (asterisk) as a wildcard to match all values. The following lists the tokens in the order that they must be specified:

- Minute
- Hour of Day
- Day of Month
- Month

- Day of Week

For more information, see [“Configuring the Eviction Schedule”](#) in *Kodo Developer’s Guide*.

Example

The following setting schedules the default cache to evict values from the cache at 15 and 45 minutes past 3:00 pm on Sunday.

```
<eviction-schedule>15,45 15 * * 1</eviction-schedule>
```

exception-orphaned-key-action

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Specifies a plugin string (see [“Plugin Configuration”](#) in *Kodo Developer’s Guide*) describing the `kodo.event.OrphanedKeyAction` to invoke when Kodo discovers an orphaned datastore key, in this case `kodo.event.ExceptionOrphanedKeyAction`. In this case, Kodo throws an exception for each orphaned key. In JPA, the exception will be `javax.persistence.EntityNotFoundException`. In JDO, the exception type will be `javax.jdo.JDOObjectNotFoundException`. For more information, see [“Orphaned Keys”](#) in the *Kodo Developer’s Guide*.

Example

```
<exception-orphaned-key-action/>
```

exception-reconnect-attempts

Range of values:	Integer
Default value:	0
Parent elements:	<p>persistence-configuration</p> <p>persistence-configuration-unit</p> <p>jms-remote-commit-provider</p>

Function

Specifies the number of time to attempt to reconnect if the JMS system notifies Kodo of a serious connection error. A value of 0 (the default) indicates that Kodo will log the error but otherwise ignore it, hoping the connection is still valid. For more information, see “[JMS](#)” in the *Kodo Developer’s Guide*.

Example

```
<exception-reconnect-attempts>0</exception-reconnect-attempts>
```

execution-context-name-provider

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration

Function

Configure the execution context name provider. For more information, see “[kodo.ExecutionContextNameProvider](#)” in *Kodo Developer’s Guide*.

Example

```
<execution-context-name-provider>
  <stack-execution-context-name-provider>...</stack-execution-context-na
```

```

me-provider>
    <transaction-name-execution-context-name-provider>...</transaction-name-
e-execution-context-name-provider>
    <user-object-execution-context-name-provider>...</user-object-executio
n-context-name-provider>
</execution-context-name-provider>

```

export-profiling

Range of values:	n/a
Default value:	n/a
Parent elements:	<p>persistence-configuration persistence-configuration-unit and persistence-configuration profiling</p>

Function

Configure export profiling data. See [“Dumping Profiling Data to Disk from a Batch Process”](#) in *Kodo Developers Guide*.

Example

```

<profiling>
    <export-profiling>
        <interval-millis>-1</interval-millis>
        <base-name>name</base-name>
    </export-profiling>
</profiling>

```

extension-deprecated-jdo-mapping-factory

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

A plugin string (see [“Plugin Configuration”](#) in *Kodo Developer’s Guide*) describing the `kodo.meta.MetadataFactory` to use to store and retrieve object-relational mapping information for your persistent classes. This setting is valid for JDO 1.0. For more information, see [“Mapping Factory”](#) in *Kodo Developer’s Guide*.

Example

```
<extension-deprecated-jdo-mapping-factory>
  <use-schema-validation>true</use-schema-validation>
  <urls>t3://localhost:7001/metadata.jar</urls>
  <files>com/file1;com/file2</files>
  <classpath-scan>build</classpath-scan>
  <types>classname1;classname2</types>
  <store-mode>1</store-mode>
  <strict>>false</strict>
  <resources>com/aaa/package.jdo;com/bbb/package.jdo</resources>
  <scan-top-down>>false</scan-top-down>
</extension-deprecated-jdo-mapping-factory>
```

fetch-batch-size

Range of values:	Integer value
Default value:	-1
Parent elements:	persistence-configuration persistence-configuration-unit

Function

The number of rows to fetch at one time when scrolling through a result set. The fetch size can also be set at runtime, as described in “[Large Result Sets](#)” in the *Kodo Developer’s Guide*.

Example

```
<fetch-batch-size>-1</fetch-batch-size>
```

fetch-direction

Range of values:	forward reverse unknown
Default value:	forward
Parent elements:	persistence-configuration persistence-configuration-unit

Function

The expected order in which query result lists are accessed. This value affects the type of data structure Kodo uses to store the results, and is passed to the JDBC driver in case it can optimize for certain access patterns. This element accepts the following values, each of which corresponds exactly to the same-named `java.sql.ResultSet` `FETCH` constant:

- `forward`—Query results are listed in order. This is the default.
- `reverse`—Query results are listed in reverse order.
- `unknown`—The order is not specified.

This property can also be varied at runtime, as described in [“Large Result Sets”](#) in the *Kodo Developer’s Guide*.

Example

```
<fetch-direction>forward</fetch-direction>
```

fetch-group

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit fetch-groups

Function

Name of a fetch group. Fetch group names are global, and are expected to be shared among classes. For example, a shopping website may use a detail fetch group in each product class to efficiently load all the data needed to display a product's "detail" page. The website might also define a sparse list fetch group containing only the fields needed to display a table of products, as in a search result.

The following names are reserved for use by Kodo: default, values, all, none, and any name beginning with jdo, ejb, or kodo.

For more information, see [“Custom Fetch Groups”](#) in the *Kodo Developer’s Guide*.

Example

```
<fetch-groups>
  <fetch-group>detail</fetch-group>
</fetch-groups>
```

fetch-groups

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Define one or more custom fetch groups.

By default, Kodo places any field that is eagerly loaded according to the JPA metadata rules into the built-in default fetch group. You may define your own named fetch groups and activate or deactivate them at runtime. Kodo will eagerly-load the fields in all active fetch groups when loading objects from the datastore. For more information, see [“Custom Fetch Groups”](#) in the *Kodo Developer’s Guide*.

Example

```
<fetch-groups>  
  <fetch-group>detail</fetch-group>  
</fetch-groups>
```

field-override

Range of values:	true false
Default value:	true
Parent elements:	<pre>persistence-configuration persistence-configuration-unit kodo-persistence-mapping-factory and persistence-configuration persistence-configuration-unit kodo-persistence-meta-data-factory</pre>

Function

Flag that specifies whether to enable field override.

Example

```
<field-override>true</field-override>
```


file

Range of values:	String
Default value:	See below
Parent elements:	<p> <code>persistence-configuration</code> <code>persistence-configuration-unit</code> <code>file-schema-factory</code> </p> <p>and</p> <p> <code>persistence-configuration</code> <code>persistence-configuration-unit</code> <code>log-factory-impl</code> </p>

Function

If the parent element is `file-schema-factory`, specifies the resource name of the XML schema file. By default, the factory looks for a resource called `package.schema` located in any top-level directory of the CLASSPATH or in the top level of any JAR in your CLASSPATH. For more information, see “[Schema Factory](#)” in the *Kodo Developer’s Guide*.

If the parent element is `log-factory-impl`, specifies the name of the file to which messages are logged. Specify `stdout` or `stderr` to log messages to standard out and standard error, respectively. The default is `stderr`. For more information, see “[Kodo Logging](#)” in *Kodo Developer’s Guide*.

Example

```
<file>stdout</file>
```

file-name

Range of values:	String
Default value:	<code>package.schema</code>
Parent elements:	<p><code>persistence-configuration</code></p> <p><code>persistence-configuration-unit</code></p> <p><code>file-schema-factory</code></p>

Function

Specifies the resource name of the XML schema file. By default, the factory looks for a resource called `package.schema` located in any top-level directory of the `CLASSPATH` or in the top level of any JAR in your `CLASSPATH`. For more information, see “[Schema Factory](#)” in the *Kodo Developer’s Guide*.

Example

```
<file-name>package.schema</file-name>
```

file-schema-factory

Range of values:	n/a
Default value:	n/a
Parent elements:	<p><code>persistence-configuration</code></p> <p><code>persistence-configuration-unit</code></p>

Function

Specifies a plugin string (see “[Plugin Configuration](#)” in *Kodo Developer’s Guide*) describing the `kodo.jdbc.schema.SchemaFactory` to use to store and retrieve information about the database schema, in this case `kodo.jdbc.schema.FileSchemaFactory`.

This factory is a lot like the [table-schema-factory](#), and has the same advantages and disadvantages. Instead of storing its XML schema definition in a database table, though, it stores it in a file.

For more information, see [“Schema Factory”](#) in the *Kodo Developer’s Guide*.

Example

```
<file-schema-factory>  
  <file-name>package.schema</file-name>  
  <file>package.schema</file>  
</file-schema-factory>
```

files

Range of values:	n/a
-------------------------	-----

Default value:	n/a
-----------------------	-----

Parent elements:	<pre>persistence-configuration persistence-configuration-unit deprecated-jdo-meta-data-factory and persistence-configuration persistence-configuration-unit extension-deprecated-jdo-mapping-factory and persistence-configuration persistence-configuration-unit jdo-meta-data-factory and persistence-configuration persistence-configuration-unit jdor-mapping-factory and persistence-configuration persistence-configuration-unit kodo-persistence-mapping-factory and persistence-configuration persistence-configuration-unit kodo-persistence-meta-data-factory</pre>
-------------------------	---

Parent elements:	and
(cont)	<p>persistence-configuration persistence-configuration-unit mapping-file-deprecated-jdo-mapping-factory</p> <p>and</p> <p>persistence-configuration persistence-configuration-unit orm-file-jdor-mapping-factory</p> <p>and</p> <p>persistence-configuration persistence-configuration-unit table-deprecated-jdo-mapping-factory</p>

Function

Specifies a semicolon-separated list of metadata files or JAR archives. Each jar archive will be scanned for annotated JPA entities or JDO metadata files based on your metadata factory. For more information, see “[Metadata Factory](#)” in *Kodo Developer’s Guide*.

Example

```
<files>com/file1;com/file2</files>
```

filter-listeners

Range of values:	n/a
Default value:	n/a
Parent elements:	<p>persistence-configuration persistence-configuration-unit</p>

Function

Defines one or more full plugin strings (see “[Plugin Configuration](#)” in *Kodo Developer’s Guide*) for custom `kodo.FilterListeners` to make available to all queries, in addition to the standard set of listeners. You can also add filter listeners to individual queries, as described in “[Query Language Extensions](#)” in the *Kodo Developer’s Guide*.

For more information, see “[kodo.FilterListeners](#)” in *Kodo Developer’s Guide*.

Example

```
<filter-listeners>
  <custom-filter-listener>detail</custom-filter-listener>
</filter-listeners>
```

foreign-keys

Range of values:	true false
Default value:	false
Parent elements:	persistence-configuration persistence-configuration-unit lazy-schema-factory

Function

Flag that specifies whether to read automatically foreign key information during schema validation. For more information, see “[Schema Factory](#)” in the *Kodo Developer’s Guide*.

Example

```
<foreign-keys>false</foreign-keys>
```

format

Range of values:	String
Default value:	n/a
Parent elements:	<p>persistence-configuration</p> <p>persistence-configuration-unit</p> <p>native-jdbc-seq</p>

Function

Specifies the format. For more information, see [“Generators”](#) in the *Kodo Developer’s Guide*.

foxpro-dictionary

Range of values:	n/a
Default value:	n/a
Parent elements:	<p>persistence-configuration</p> <p>persistence-configuration-unit</p>

Function

Configuration values for the FoxPro Dictionary persistence plugin. For a complete description of each of the values that you can specify, see [“FoxPro Dictionary Configuration”](#) in the *Administration Console Online Help*.

Example

The `foxpro-dictionary` element shares the same subelements as [“access-dictionary”](#) on [page B-7](#).

flush-before-queries

Range of values:	true false with-connection
Default value:	true
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Flag that specifies whether to flush any changes made in the current transaction to the datastore before executing a query. Valid values include:

- `true`—Always flush rather than executing the query in-memory. If the current transaction is optimistic, Kodo will begin a non-locking datastore transaction. This is the default.
- `false`—Never flush before a query.
- `with-connection`—Flush only if the `EntityManager` or `PersistenceManager` has already established a dedicated connection to the datastore, otherwise execute the query in-memory.

This option is useful if you use long-running optimistic transactions and want to ensure that these transactions do not consume database resources until commit. Kodo's behavior with this option is dependent on the transaction status and mode, as well as the configured connection retain mode described earlier in this section.

For more information, see [“Configuring the Use of JDBC Connections”](#) in *Kodo Developer's Guide*.

Example

```
<flush-before-queries>true</flush-before-queries>
```


hsql-dictionary

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Configuration values for the HSQL Dictionary persistence plugin. For a complete description of each of the values that you can specify, see [“HSQL Dictionary Configuration”](#) in the *Administration Console Online Help*.

Example

The `hsql-dictionary` element shares the same subelements as [“access-dictionary”](#) on [page B-7](#).

gem-fire-data-cache

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit data-caches

Function

Integrates with GemStone’s GemFire v5.0.1 caching system. For more information, see [“GemStone GemFire Integration”](#) in *Kodo Developer’s Guide*.

Example

```
<data-caches>
  <gem-fire-data-cache>
```

```
<name>kodo.DataCache</name>
<gem-fire-cache-name>root/kodo-data-cache</gem-fire-cache-name>
<eviction-schedule>15, 45 15 * * 1</eviction-schedule>
</gem-fire-data-cache>
</data-caches>
```

gem-fire-data-cache-name

Range of values:	Valid GemFire region name
Default value:	Data cache: root/kodo-data-cache Query cache: root/kodo-query-cache
Parent elements:	persistence-configuration persistence-configuration-unit data-caches gem-fire-data-cache

Function

GemFire region name. For more information, see [“GemStone GemFire Integration”](#) in *Kodo Developer’s Guide*.

Example

```
<gem-fire-data-cache-name>root/kodo-data-cache</gem-fire-data-cache-name>
```

gui-jmx

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Enable management and invoke the JMX management console in the local JVM. Supports optional parameters in the Management Group, as described in [“Optional Parameters in Management Group”](#) in the *Kodo Developer’s Guide*.

Example

```
<gui-jmx>
  <MBeanServerStrategy>any-create</MBeanServerStrategy>
  <EnableLogMBean>true</EnableLogMBean>
  <EnableRuntimeMBean>true</EnableRuntimeMBean>
</gui-jmx>
```

gui-profiling

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit and persistence-configuration profiling

Function

Turns on the local profiling GUI. See [“Profiling in an Embedded GUI”](#) in the Kodo Developer’s Guide for more information.

Example

```
<profiling>  
  <gui-profiling/>  
</profiling>
```

Host

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit mx4j1-jmx

Function

Hostname on which the RMI registry naming service listens. Defaults to `localhost`. For more information, see [“Optional Parameters in Remote Group”](#) in the *Kodo Developer’s Guide*.

Example

```
<Host>localhost</Host>
```

host

Range of values:	Valid hostname
Default value:	localhost
Parent elements:	persistence-configuration persistence-configuration-unit tcp-transport

Function

Specifies the host name of the server. This setting is used by clients, not by the server. For more information, see [“Standalone Persistence Server”](#) in *Kodo Developer’s Guide*.

Example

```
<host>localhost</host>
```

http-transport

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Specifies the URL of the remote server. For more information, see [“Client Managers”](#) in *Kodo Developer’s Guide*.

Example

```
<http-transport>
  <url>servlet-url</url>
</http-transport>
```

ignore-changes

Range of values:	true false
Default value:	false
Parent elements:	persistence-configuration persistence-configuration-unit kodo-broker

Function

Flag that specifies whether to consider modifications to persistent objects made in the current transaction when evaluating queries. Setting this to `true` allows Kodo to ignore changes and execute the query directly against the datastore. A value of `false` forces Kodo to consider whether the changes in the current transaction affect the query, and if so to either evaluate the query inmemory or flush before running it against the datastore. For more information, see [“kodo.IgnoreChanges”](#) in *Kodo Developer’s Guide*.

Example

```
<ignore-changes>false</ignore-changes>
```

ignore-unmapped

Range of values:	true false
Default value:	false
Parent elements:	persistence-configuration persistence-configuration-unit class-table-jdbc-seq

Function

Flag that specifies whether to ignore unmapped base classes and instead use one row per least-derived mapped class. For more information, see “[Generators](#)” in the *Kodo Developer’s Guide*.

Example

```
<ignore-unmapped>false</ignore-unmapped>
```

ignore-virtual

Range of values:	true false
Default value:	false
Parent elements:	persistence-configuration persistence-configuration-unit class-table-jdbc-seq

Function

Flag that specifies whether to ignore virtual classes. For more information, see “[Generators](#)” in the *Kodo Developer’s Guide*.

Example

```
<ignore-virtual>false</ignore-virtual>
```

in-memory-savepoint-manager

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Specifies a plugin string (see [“Plugin Configuration”](#) in *Kodo Developer’s Guide*) describing the `kodo.kernel.SavepointManager` to use for managing transaction savepoints, in this case `kodo.kernel.InMemorySavepointManager`. This plugin stores all state, including field values, in memory. Because of this behavior, each set savepoint is designed for small to medium transactional object counts. For more information, see [“Savepoints”](#) in the *Kodo Developer’s Guide*.

Example

```
<in-memory-savepoint-manager/>
```


increment

Range of values:	Integer
Default value:	1
Parent elements:	<pre> persistence-configuration persistence-configuration-unit class-table-jdbc-seq and persistence-configuration persistence-configuration-unit native-jdbc-seq and persistence-configuration persistence-configuration-unit table-jdbc-seq and persistence-configuration persistence-configuration-unit time-seeded-seq and persistence-configuration persistence-configuration-unit value-table-jdbc-seq </pre>

Function

Specifies the amount of sequence increments. For more information, see [“Generators”](#) in the *Kodo Developer’s Guide*.

Example

```
<increment>1</increment>
```

indexes

Range of values:	true false
Default value:	false
Parent elements:	persistence-configuration persistence-configuration-unit lazy-schema-factory

Function

Flag that specifies whether to read automatically index information during schema validation. For more information, see [“Schema Factory”](#) in the *Kodo Developer’s Guide*.

Example

```
<indexes>false</indexes>
```

informix-dictionary

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Configuration values for the Informix Dictionary persistence plugin. For a complete description of each of the values that you can specify, see [“Informix Dictionary Configuration”](#) in the *Administration Console Online Help*.

Example

The `informix-dictionary` element shares the same subelements as [“access-dictionary”](#) on [page B-7](#), plus the following subelements (defaults shown):

```

<informix-dictionary>
...
  <lock-mode-enabled>false</lock-mode-enabled>
  <lock-wait-seconds>30</lock-wait-seconds>
  <swap-schema-and-catalog>true</swap-schema-and-catalog>
</informix-dictionary>

```

initial-value

Range of values:	n/a
Default value:	1
Parent elements:	<p>persistence-configuration</p> <p>persistence-configuration-unit</p> <p>native-jdbc-seq</p>

Function

Specifies the initial sequence value. For more information, see [“Generators”](#) in the *Kodo Developer’s Guide*.

Example

```
<initial-value>1</initial-value>
```

interval-millis

Range of values:	n/a
Default value:	n/a
Parent elements:	<p>persistence-configuration persistence-configuration-unit export-profiling</p> <p>and</p> <p>persistence-configuration profiling export-profiling</p>

Function

Number of milliseconds between exports. This value defaults to -1, indicating that there will be a single export upon exit. See [“Dumping Profiling Data to Disk from a Batch Process”](#) in *Kodo Developers Guide*.

Example

```
<interval-millis>-1</interval-millis>
```

inverse-manager

Range of values:	true false
Default value:	false
Parent elements:	<p>persistence-configuration persistence-configuration-unit</p>

Function

A plugin string (see [“Plugin Configuration”](#) in *Kodo Developer’s Guide*) describing a `kodo.kernel.InverseManager` to use for managing bidirectional relations upon a flush. For more information, see [“Managed Inverses”](#) in *Kodo Developer’s Guide*.

Example

```
<inverse-manager>
  <action>warn</action>
  <manage-lrs>false</manage-lrs>
</inverse-manager>
```

jdatastore-dictionary

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Configuration values for the JDataStore Dictionary persistence plugin. For a complete description of each of the values that you can specify, see [“JDataStore Dictionary Configuration”](#) in the *Administration Console Online Help*.

Example

The `jdatastore-dictionary` element shares the same subelements as [“access-dictionary”](#) on [page B-7](#).

jdbc-broker-factory

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

A plugin string (see “[Plugin Configuration](#)” in *Kodo Developer’s Guide*) describing the `kodo.kernel.BrokerFactory` type to use, in this case JDBC. For more information, see “[kodo.BrokerFactory](#)” in the *Kodo Developer’s Guide*.

Example

```
<jdbc-broker-factory/>
```

jdbc-listeners

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Defines one or more full plugin strings (see “[Plugin Configuration](#)” in *Kodo Developer’s Guide*) for custom `kodo.jdbc.kernel.exps.JDBCFilterListeners` to make available to all queries, in addition to the standard set of listeners. You can also add filter listeners to individual queries, as described in “[Query Language Extensions](#)” in the *Kodo Developer’s Guide*.

Example

```
<jdbc-listeners>detail</jdbc-listeners>
```

jdbc3-savepoint-manager

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Specifies a plugin string (see “[Plugin Configuration](#)” in *Kodo Developer’s Guide*) describing the `kodo.kernel.SavepointManager` to use for managing transaction savepoints, in this case the default which is `kodo.jdbc.kernel.JDBCSavepointManager`. This plugin implements savepoints by issuing a flush to the database. For more information, see “[Savepoints](#)” in the *Kodo Developer’s Guide*.

Example

```
<jdbc3-savepoint-manager/>
```

jdo-meta-data-factory

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Specifies a plugin string (see “[Plugin Configuration](#)” in *Kodo Developer’s Guide*) describing the `kodo.meta.MetaDataFactory` type to use, in this case `kodo.jdo.JDOMetaDataFactory`. For more information, see “[Metadata Factory](#)” in the *Kodo Developer’s Guide*.

Example

```
<jdo-meta-data-factory>
  <use-schema-validation>false</use-schema-validation>
  <urls>t3://localhost:7001/metadata.jar</urls>
  <files>com/file1;com/file2</files>
  <classpath-scan>build</classpath-scan>
  <constraint-names>false</constraint-names>
  <types>classname1;classname2</types>
  <store-mode>1</store-mode>
  <strict>false</strict>
  <resources>com/aaa/package.jdo;com/bbb/package.jdo</resources>
  <scan-top-down>false</scan-top-down>
</jdo-meta-data-factory>
```

jms-remote-commit-provider

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Configures the JMS remote commit provider. For more information, see “[JMS](#)” in the *Kodo Developer’s Guide*.

Example

```
<jms-remote-commit-provider>
  <name>kodo.RemoteCommitProvider</name>
  <topic>topic/KodoCommitProviderTopic</topic>
  <exception-reconnect-attempts>0</exception-reconnect-attempts>
  <topic-connection-factory>java://ConnectionFactory</topic-connection-factory>
</jms-remote-commit-provider>
```


jmx

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Enables for the configuration of management capabilities. For more information, see [“Management and Monitoring”](#) in the *Kodo Developer’s Guide*.

Example

```
<jmx>
  <none-jmx/>
  <local-jmx>...</local-jmx>
  <gui-jmx>...</gui-jmx>
  <jmx2-jmx>...</jmx2-jmx>
  <mx4j1-jmx>...</mx4j1-jmx>
  <wls81-jmx>...</wls81-jmx>
</jmx>
```

jmx2-jmx

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Enable remote management via JMX v.1.2 implementations (supporting JSR 160 for remote management). Supports optional parameters in the Management Group and the JSR 160

Group, as described in [“Optional Parameters in Management Group”](#) and [“Optional Parameters in JSR 160 Group”](#) in the *Kodo Developer’s Guide*, respectively.

Example

```
<jmx2-jmx>
  <MBeanServerStrategy>any-create</MBeanServerStrategy>
  <EnableLogMBean>true</EnableLogMBean>
  <EnableRuntimeMBean>true</EnableRuntimeMBean>
  <NamingImpl>mx4j.tools.naming.NamingService</NamingImpl>
  <ServiceURL>service:jmx:rmi://localhost/jndi/jmxservice</ServiceURL>
</jmx2-jmx>
```

JNDIName

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit mx4j1-jmx

Function

JNDI name under which to register the remote JMX adaptor. Defaults to `jrmp`. For more information, see [“Optional Parameters in Remote Group”](#) in the *Kodo Developer’s Guide*.

Example

```
<JNDIName>jrmp</JNDIName>
```

kodo-broker

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Configures the `kodo.BrokerImpl`. For more information, see “[kodo.BrokerImpl](#)” in *Kodo Developer’s Guide*.

Example

```
<kodo-broker>
  <large-transaction>false</large-transaction>
  <auto-clear>datastore</auto-clear>
  <detach-state>1</detach-stage>
  <nontransactional-read>true</nontransactional-read>
  <retain-state>false</retain-state>
  <evict-from-data-cache>false</evict-from-data-cache>
  <detached-new>true</detached-new>
  <optimistic>true</optimistic>
  <nontransactional-write>flase</nontransactional-write>
  <sync-with-managed-transacations>false</sync-with-managed-transactions
>
  <multithreaded>false</multithreaded>
  <populate-data-cache>true</populate-data-cache>
  <ignore-changes>false</ignore-changes>
  <auto-detach>0</auto-detach>
  <restore-state>1</restore-state>
  <order-dirty-objects>false</other-dirty-objects>
</kodo-broker>
```

kodo-concurrent-data-cache

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit data-caches

Function

Enables the basic concurrent data cache. For more information, see [“Data Cache”](#) in *Kodo Developer’s Guide*.

Example

```
<data-caches>  
  <kodo-concurrent-data-cache>  
    <name>default</name>  
    <cache-size>1000</cache-size>  
    <soft-reference-size>0</soft-reference-size>  
    <eviction-schedule>15,45 15 * * 1</eviction-schedule>  
  </kodo-concurrent-data-cache>  
</data-caches>
```

kodo-data-cache-manager

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Enables the `kodo.datacache.DataCacheManager` that manages the system data caches. For more information, see [“kodo.DataCacheManager”](#) in *Kodo Developer’s Guide*.

Example

```
<kodo-data-cache-manager />
```

kodo-mapping-repository

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Specifies a plugin string (see [“Plugin Configuration”](#) in *Kodo Developer’s Guide*) describing the `kodo.meta.MetaDataRepository` to use. For more information, see [“kodo.MetaDataRepository”](#) in the *Kodo Developer’s Guide*.

Example

```
<kodo-mapping-repository>
  <resolve>3</resolve>
  <validate>7</validate>
```

```
<source-mode>7</source-mode>
</kodo-mapping-repository>
```

kodo-persistence-mapping-factory

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

A plugin string (see [“Plugin Configuration”](#) in *Kodo Developer’s Guide*) describing the `kodo.meta.MetadataFactory` to use to store and retrieve object-relational mapping information for your persistent classes. This setting is valid for JPA 1.0. For more information, see [“Mapping Factory”](#) in *Kodo Developer’s Guide*.

Example

```
<kodo-persistence-mapping-factory>
  <urls>t3://localhost:7001/metadata.jar</urls>
  <files>com/file1;com/file2</files>
  <classpath-scan>build</classpath-scan>
  <default-access-type>FIELD</default-access-type>
  <field-override>true</field-override>
  <types>classname1;classname2</types>
  <store-mode>1</store-mode>
  <strict>>false</strict>
  <resources>com/aaa/package.jdo;com/bbb/package.jdo</resources>
</kodo-persistence-mapping-factory>
```

kodo-persistence-meta-data-factory

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Specifies a plugin string (see “[Plugin Configuration](#)” in *Kodo Developer’s Guide*) describing the `kodo.meta.MetaDataFactory` type to use, in this case `org.apache.openjpa.persistence.PersistenceMetaDataFactory`. For more information, see “[Metadata Factory](#)” in the *Kodo Developer’s Guide*.

Example

```
<kodo-persistence-meta-data-factory>
  <urls>t3://localhost:7001/metadata.jar</urls>
  <files>com/file1;com/file2</files>
  <classpath-scan>build</classpath-scan>
  <default-access-type>FIELD</default-access-type>
  <field-override>>true</field-override>
  <types>classname1;classname2</types>
  <store-mode>1</store-mode>
  <strict>>false</strict>
  <resources>com/aaa/package.jdo;com/bbb/package.jdo</resources>
</kodo-persistence-meta-data-factory>
```

kodo-pooling-data-source

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Configures the `kodo.jdbc.schema.KodoPoolingDataSource` which by default performs connection pooling. For more information, see [“JDBC”](#) in *Kodo Developer’s Guide*.

Example

```
<kodo-pooling-data-source>
  <connection-user-name>KodoPool</connection-user-name>
  <login-timeout>10</login-timeout>
  <connection-password>password</connection-password>
  <connection-url>jdbc:hsqldb:db-hypersonic</connection-url>
  <connection-driver-name>org.hsqldb.jdbcDrier</connection-driver-name>
</kodo-pooling-data-source>
```

kodo-sql-factory

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Specifies a plugin string (see [“Plugin Configuration”](#) in *Kodo Developer’s Guide*) describing the `kodo.jdbc.SQLFactory` to use to abstract common SQL constructs. For more information, see [“kodo.jdbc.SQLFactory”](#) in the *Kodo Developer’s Guide*.

Example

```
<kodo-sql-factory>
  <advanced-sql>...</advanced-sql>
</kodo-sql-factory>
```

large-transaction

Range of values:	true false
Default value:	false
Parent elements:	persistence-configuration persistence-configuration-unit kodo-broker

Function

Flag that specifies whether the transaction will generate a large number of objects.

Example

```
<large-transaction>>false</large-transaction>
```

lazy-schema-factory

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Specifies a plugin string (see [“Plugin Configuration”](#) in *Kodo Developer’s Guide*) describing the `kodo.jdbc.schema.SchemaFactory` to use to store and retrieve information about the database schema, in this case `kodo.jdbc.schema.LazySchemaFactory`.

As persistent classes are loaded by the application, Kodo reads their metadata and object-relational mapping information. This factory uses the `java.sql.DatabaseMetaData` interface to reflect on the schema and ensure that it is consistent with the mapping data being read. Because the factory does not reflect on a table definition until that table is mentioned by the mapping information, it is referred to as *lazy*. Use this factory if you want up-front validation that your mapping metadata is consistent with the database during development.

For more information, see “[Schema Factory](#)” in the *Kodo Developer’s Guide*.

Example

```
<lazy-schema-factory>
  <foreign-keys>false</foreign-keys>
  <indexes>false</indexes>
  <primary-keys>false</primary-keys>
</lazy-schema-factory>
```

level

Range of values:	1 2 3 4 5
Default value:	4
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Specifies the level at which to log. Valid values include: 1 (TRACE), 2 (DEBUG), 3 (INFO), 4 (WARN), and 5 (ERROR). For more information, see “[Orphaned Keys](#)” in the *Kodo Developer’s Guide*.

Example

```
<level>4</level>
```

local-jmx

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Enabled local management. This setting is suitable for use in JBoss and other environments where all MBeans should be registered with a JMX MBeanServer for either management via a user defined mechanism, or via a mechanism defined by the MBeanServer. Supports optional parameters in the `Management Group`, as described in “[Optional Parameters in Management Group](#)” in the *Kodo Developer’s Guide*.

Example

```
<local-jmx>
  <MBeanServerStrategy>any-create</MBeanServerStrategy>
  <EnableLogMBean>true</EnableLogMBean>
  <EnableRuntimeMBean>true</EnableRuntimeMBean>
</local-jmx>
```

local-profiling

Range of values:	n/a
Default value:	n/a
Parent elements:	<p>persistence-configuration persistence-configuration-unit and persistence-configuration profiling</p>

Function

Enable profiling without export or GUI. Useful when trying to access the `ProfilingAgent` programmatically. For more information, see [“Profiling”](#) in *Kodo Developer’s Guide*.

Example

```
<profiling>
  <local-profiling/>
</profiling>
```

lock-timeout

Range of values:	Integer
Default value:	-1
Parent elements:	<p>persistence-configuration persistence-configuration-unit</p>

Function

Default amount of time that Kodo waits when trying to obtain a lock. A value of -1 specifies that there is no limit. For more information, see [“Configuring Default Locking”](#) in *Kodo Developer’s Guide*.

Example

```
<lock-timeout>-1</lock-timeout>
```

log-factory-impl

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Configures the default logging system. For more information, see [“Kodo Logging”](#) in *Kodo Developer’s Guide*.

Example

```
<log-factory-impl>
  <diagnostic-context>ID</diagnostic-context>
  <default-level>INFO</default-level>
  <file>stdout</file>
</log-factory-impl>
```

log-orphaned-key-action

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Specifies a plugin string (see “[Plugin Configuration](#)” in *Kodo Developer’s Guide*) describing the `kodo.event.OrphanedKeyAction` to invoke when Kodo discovers an orphaned datastore key, in this case `kodo.event.LogOrphanedKeyAction`. In this case, Kodo logs a message for each orphaned key. For more information, see “[Orphaned Keys](#)” in the *Kodo Developer’s Guide*.

Example

```
<log-orphaned-key-action>
  <channel>openjpa.Runtime</channel>
  <level>4</level>
</log-orphaned-key-action>
```

log4j-log-factory

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Specifies to use Log4J for logging. In a standalone application, Log4J logging levels are controlled by a resource named `log4j.properties`, which should be available as a top-level resource (either at the top level of a jar file, or in the root of one of the `CLASSPATH` directories). When deploying to a web or EJB application server, Log4J configuration is often performed in a

log4j.xml file instead of a properties file. For further details on configuring Log4J, please see the [Log4J Manual](#). For more information, see “[Log4J](#)” in *Kodo Developer’s Guide*.

Example

```
<log4j-log-factory/>
```

login-timeout

Range of values:	Valid classname
Default value:	<code>kodo-pooling-data-source: 10</code> <code>simple-driver-data-source: 0</code>
Parent elements:	<p> <code>persistence-configuration</code> <code>persistence-configuration-unit</code> <code>kodo-pooling-data-source</code> </p> <p>and</p> <p> <code>persistence-configuration</code> <code>persistence-configuration-unit</code> <code>simple-driver-data-source</code> </p>

Function

Specifies the login timeout. For more information, see “[Using the Kodo Data Source](#)” in *Kodo Developer’s Guide*.

Example

```
<login-timeout>0</login-timeout>
```

lrs-size

Range of values:	query last unknown
Default value:	query
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Specifies the size of results sets. This property is only used if you change the fetch batch size from its default of -1 so that Kodo uses the on-demand results loading. Valid values include:

- `query`—The first time you ask for the size of a query result, Kodo will perform a `SELECT COUNT (*)` query to determine the number of expected results. Note that depending on transaction status and settings, this can mean that the reported size is slightly different than the actual number of results available. This is the default.
- `last`—If you have chosen a scrollable result set type, this setting will use the `ResultSet.last` method to move to the last element in the result set and get its index. Unfortunately, some JDBC drivers will bring all results into memory in order to access the last one. Note that if you do not choose a scrollable result set type, then this will behave exactly like `unknown`. The default result set type is `forward-only`, so you must change the result set type in order for this property to have an effect.
- `unknown`—Kodo returns `Integer.MAX_VALUE` as the size for any query result that uses on-demand loading.

For more information, see “[Large Result Sets](#)” in *Kodo Developer’s Guide*.

Example

```
<lrs-size>query</lrs-size>
```


lru-data-cache

Range of values:	n/a
Default value:	n/a
Parent elements:	<p>persistence-configuration persistence-configuration-unit data-caches</p>

Function

Sets the least-recently-used (LRU) data caching option. For more information, see [“Configuring the LRU Caching Option”](#) in *Kodo Developer’s Guide*.

Example

```
<data-caches>
  <lru-data-cache>
    <name>kodo.DataCache</name>
    <cache-size>1000</cache-size>
    <soft-reference-size>0</soft-reference-size>
    <eviction-schedule>15,45 15 * * 1</eviction-schedule>
  </lru-data-cache>
</data-caches>
```

manage-lru

Range of values:	true false
Default value:	false
Parent elements:	persistence-configuration persistence-configuration-unit inverse-manager

Function

Flag that specifies whether large result set fields are included from management. For more information, see [“Managed Inverses”](#) in *Kodo Developer’s Guide*.

Example

```
<manage-lrs>false</manage-lrs>
```

mapping

Range of values:	String
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit and persistence-configuration persistence-configuration-unit orm-file-mapping-factory

Function

When a child element of [persistence-configuration-unit](#), specifies the symbolic name of the object-to-datastore mapping to use. For more information, see [“Mapping Metadata”](#) (JPA) or [“Mapping Metadata Placement”](#) (JDO) in *Kodo Developer’s Guide*.

When a child element of `orm-file-jdor-mapping-factory`, specifies the logical name of these mappings. Mapping files are suffixed with `logicalname.orm`. If not specified, the `kodo.Mapping` configuration property is used. For more information, see [“Mapping Factory”](#) in *Kodo Developer’s Guide*.

Example

```
<mapping>org/mag/Magazine-hsql.orm</mapping>
```

mapping-column

Range of values:	String
Default value:	MAPPING_DEF
Parent elements:	<p> <code>persistence-configuration</code> <code>persistence-configuration-unit</code> <code>table-deprecated-jdo-mapping-factory</code> </p> <p>and</p> <p> <code>persistence-configuration</code> <code>persistence-configuration-unit</code> <code>table-jdor-mapping-factory</code> </p>

Function

Specifies the name of the column that holds the XML mapping. For more information, see [“Mapping Factory”](#) in *Kodo Developer’s Guide*.

Example

```
<mapping-column>MAPPING_DEF</mapping-column>
```

mapping-defaults-impl

Range of values:	n/a
Default value:	n/a
Parent elements:	<p>persistence-configuration</p> <p>persistence-configuration-unit</p>

Function

Configures the mapping defaults if you set `default-mapping-defaults` to `jdo`. For a complete description of each of the values that you can specify, see [“Mapping Defaults”](#) in *Kodo Developer’s Guide*.

Example

```
<mapping-defaults-impl>
  <use-class-criteria>false</use-class-criteria>
  <base-class-strategy>ColumnPerLockGroupVersionStrategy</base-class-strategy>
  <version-strategy>version-number</version-strategy>
  <discriminator-column-name>test</discriminator-column-name>
  <subclass-strategy>vertical</subclass-strategy>
  <index-version>false</index-version>
  <index-logical-foreign-keys>true</index-logical-foreign-keys>
  <null-indicator-column-name>null</null-indicator-column-name>
  <foreign-key-delete-action></foreign-key-delete-action>
  <join-foreign-key-delete-action>1</join-foreign-key-delete-action>
  <discriminator-strategy>final</discriminator-strategy>
  <defer-constraints>false</defer-constraints>
  <field-strategies>none</field-strategies>
  <version-column-name>version</version-column-name>
  <data-store-id-column-name>ID</data-store-id-column-name>
  <index-discriminator>true</index-discriminator>
  <store-enum-ordinal>false</store-enum-ordinal>
  <order-lists>true</order-lists>
  <order-column-name>false</order-column-name>
```

```

    <add-null-indicator>false</add-null-indicator>
    <store-unmapped-object-id-string>false</store-unmapped-object-id-string>
  </mapping-defaults-impl>

```

mapping-file-deprecated-jdo-mapping-factory

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

A plugin string (see [“Plugin Configuration”](#) in *Kodo Developer’s Guide*) describing the `kodo.meta.MetadataFactory` to use to store and retrieve object-relational mapping information for your persistent classes. This setting is valid for JDO 1.0. For more information, see [“Mapping Factory”](#) in *Kodo Developer’s Guide*.

Example

```

<mapping-file-deprecated-jdo-mapping-factory>
  <use-schema-validation>true</use-schema-validation>
  <urls>t3://localhost:7001/metadata.jar</urls>
  <files>com/file1;com/file2</files>
  <classpath-scan>build</classpath-scan>
  <single-file>false</single-file>
  <types>classname1;classname2</types>
  <store-mode>1</store-mode>
  <strict></strict>
  <resources>com/aaa/package.jdo;com/bbb/package.jdo</resources>
  <scan-top-down>false</scan-top-down>
</mapping-file-deprecated-jdo-mapping-factory>

```

max-active

Range of values:	Integer
Default value:	2
Parent elements:	<p>persistence-configuration persistence-configuration-unit tcp-remote-commit-provider</p>

Function

Specifies the maximum allowed number of TCP sockets (channels) to open simultaneously between each peer in the cluster. For more information, see [“TCP”](#) in the *Kodo Developer’s Guide*.

Example

```
<max-active>2</max-active>
```

max-idle

Range of values:	Integer
Default value:	2
Parent elements:	<p>persistence-configuration persistence-configuration-unit tcp-remote-commit-provider</p>

Function

Specifies the maximum allowed number of TCP sockets (channels) to open simultaneously between each peer in the cluster. For more information, see [“TCP”](#) in the *Kodo Developer’s Guide*.

Example

```
<max-idle>2</max-idle>
```

max-size

Range of values:	Integer
Default value:	2147483647
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Defines the maximum size of the concurrent hash map used to associate query strings and their parsed form. For more information, see [“Query Compilation Cache”](#) in the *Kodo Developer’s Guide*.

Example

```
<max-size>2147483647</max-size>
```

maximize-batch-size

Range of values: true | false

Default value: true

Parent elements: persistence-configuration
 persistence-configuration-unit
 constraint-update-manager

and

persistence-configuration
 persistence-configuration-unit
 batching-operation-order-update-manager

and

persistence-configuration
 persistence-configuration-unit
 table-lock-update-manager

Function

Flag that specifies whether to sort statements in order to optimize batch size when statement batching is on. For more information, see “[kodo.jdbc.UpdateManager](#)” in *Kodo Developer’s Guide*.

Example

```
<maximize-batch-size>true</maximize-batch-size>
```


MBeanServerStrategy

Range of values:	<code>any-create create agentID:<agentID></code>
Default value:	<code>any-create</code>
Parent elements:	<code>persistence-configuration</code> <code>persistence-configuration-unit</code> <code>local-jmx</code>

Function

If JMX-based management is enabled, Kodo needs to locate an existing `MBeanServer` or create a new one, based on one of the following options:

- `any-create`—Locate an existing `MBeanServer`. If multiple are found, use the first one. If none are found, create a new server. This is the default.
- `create`—Create a new server. Do not attempt to find an existing `MBeanServer`.
- `agentId:<agentID>`—Locate an existing `MBeanServer` with the given agent id. If not found, create a new server.

For more information, see [“Optional Parameters in Management Group”](#) in *Kodo Developer’s Guide*.

Example

```
<MBeanServerStrategy>any-create</MBeanServerStrategy>
```

multithreaded

Range of values:	true false
Default value:	false
Parent elements:	<p>persistence-configuration persistence-configuration-unit</p> <p>and</p> <p>persistence-configuration persistence-configuration-unit kodo-broker</p>

Function

Flag that specifies whether persistent instances and Kodo components are accessed by multiple threads at once. For more information, see [“kodo.Multithreaded”](#) in *Kodo Developer’s Guide*.

Example

```
<multithreaded>false</multithreaded>
```

mx4j1-jmx

Range of values:	n/a
Default value:	n/a
Parent elements:	<p>persistence-configuration persistence-configuration-unit</p>

Function

Enable remote management via MX4J v.1.x implementations (supporting versions of the JMX specification prior to 1.2). Supports optional parameters in the `Management Group` and the `Remote Group`, as described in [“Optional Parameters in Management Group”](#) and [“Optional Parameters in Remote Group”](#) in the *Kodo Developer’s Guide*, respectively.

Example

```
<mx4j1-jmx>
  <MBeanServerStrategy>any-create</MBeanServerStrategy>
  <EnableLogMBean>true</EnableLogMBean>
  <EnableRuntimeMBean>true</EnableRuntimeMBean>
  <Host>localhost</Host>
  <Port>1099</Port>
  <JNDIName>jrmp</JNDIName>
</mx4j1-jmx>
```

mysql-dictionary

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Configuration values for the MySQL Dictionary persistence plugin. For a complete description of each of the values that you can specify, see [“MySQL Dictionary Configuration”](#) in the *Administration Console Online Help*.

Example

The `mysql-dictionary` element shares the same subelements as [“access-dictionary”](#) on [page B-7](#), plus the following subelements (defaults shown):

```
<mysql-dictionary>
...
  <table-type>innodb</table-type>
  <use-clobs>true</use-clobs>
  <driver-deserializes-blobs>true</driver-deserializes-blobs>
</mysql-dictionary>
```

name

Range of values: n/a

Default value: n/a

Parent elements: persistence-configuration
 persistence-configuration-unit
 data-cache

and

persistence-configuration
 persistence-configuration-unit
 db-dictionary

and

persistence-configuration
 persistence-configuration-unit
 jms-remote-commit-provider

and

persistence-configuration
 persistence-configuration-unit
 property-type

and

persistence-configuration
 persistence-configuration-unit
 query-cache

and

persistence-configuration
 persistence-configuration-unit
 single-jvm-remote-commit-provider

and

persistence-configuration
 persistence-configuration-unit
 tcp-remote-commit-provider

and

persistence-configuration
 persistence-configuration-unit
 cluster-remote-commit-provider

Parent elements (cont): and

persistence-configuration
 persistence-configuration-unit
 custom-remote-commit-provider

Function

Specifies the name of the corresponding element.

name-column

Range of values:	String
Default value:	NAME
Parent elements:	<code>persistence-configuration</code> <code>persistence-configuration-unit</code> <code>table-deprecated-jdo-mapping-factory</code> and <code>persistence-configuration</code> <code>persistence-configuration-unit</code> <code>table-jdor-mapping-factory</code>

Function

Name of the column that holds the mapping name. For class mappings, the mapping name is the class name. For named queries and sequences, it is the sequence or query name. For more information, see [“Mapping Factory”](#) in *Kodo Developer’s Guide*.

Example

```
<name-column>NAME</name-column>
```

NamingImpl

Range of values:	Valid classname
Default value:	n/a
Parent elements:	<p>persistence-configuration</p> <p>persistence-configuration-unit</p> <p>jmx2-jmx</p>

Function

Classname of the naming service implementation to start in order to register the RMI connector with a JNDI name. Defaults to `mx4j.tools.naming.NamingService`, which is appropriate for MX4J v. 2.x. If set to the empty string, no naming service will be started. This setting is appropriate if a naming service is already running, or if a non-RMI connector is used. For more information, see [“Optional Parameters in JSR 160 Group”](#) in *Kodo Developer’s Guide*.

Example

```
<NamingImpl>mx4j.tools.naming.NamingService</NamingImpl>
```

native-jdbc-seq

Range of values:	n/a
Default value:	n/a
Parent elements:	<p>persistence-configuration</p> <p>persistence-configuration-unit</p>

Function

Specifies a plugin string (see [“Plugin Configuration”](#) in *Kodo Developer’s Guide*) describing the `kodo.kernel.Seq` interface to use to create your own custom generators, in this case `kodo.jdbc.kernel.NativeJDBCSeq`.

Many databases have a concept of *native sequences*—a built-in mechanism for obtaining incrementing numbers. For example, in Oracle, you can create a database sequence with a

statement like `CREATE SEQUENCE MYSEQUENCE`. Sequence values can then be atomically obtained and incremented with the statement `SELECT MYSEQUENCE.NEXTVAL FROM DUAL`. Kodo provides support for this common mechanism of sequence generation with the `NativeJDBCSeq`.

For more information, see [“Generators”](#) in the *Kodo Developer’s Guide*.

Example

```
<native-jdbc-seq>
  <type>0</type>
  <allocate>50</allocate>
  <table-name>DUAL</table-name>
  <initial-value>1</initial-value>
  <sequence>OPENJPA_SEQUENCE</sequence>
  <sequence-name>OPENJPA_SEQUENCE</sequence-name>
  <format>format</format>
  <increment>1</increment>
</native-jdbc-seq>
```

none-jmx

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

No management is turned on. This is the default. For more information, see [“Configuration”](#) in *Kodo Developer’s Guide*.

Example

```
<none-jmx/>
```


none-lock-manager

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Specifies the `kodo.kernel.NoneLockManager` which does not perform any locking. For more information, see [“Lock Manager”](#) in *Kodo Developer’s Guide*.

Example

```
<none-lock-manager/>
```

none-log-factory

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Disable logging. For more information, see [“Disable Logging”](#) in *Kodo Developer’s Guide*.

Example

```
<none-log-factory/>
```

none-orphaned-key-action

Range of values:	n/a
Default value:	n/a
Parent elements:	<p>persistence-configuration</p> <p>persistence-configuration-unit</p>

Function

Specifies a plugin string (see “[Plugin Configuration](#)” in *Kodo Developer’s Guide*) describing the `kodo.event.OrphanedKeyAction` to invoke when Kodo discovers an orphaned datastore key, in this case `kodo.event.NoneOrphanedKeyAction`. In this case, Kodo ignores orphaned keys. For more information, see “[Orphaned Keys](#)” in the *Kodo Developer’s Guide*.

Example

```
<none-orphaned-key-action/>
```

none-profiling

Range of values:	n/a
Default value:	n/a
Parent elements:	<p>persistence-configuration</p> <p>persistence-configuration-unit</p> <p>and</p> <p>persistence-configuration</p> <p>profiling</p>

Function

Turns off profiling of your code. For more information, see “[Profiling](#)” in *Kodo Developer’s Guide*.

Example

```
<profiling>
  <none-profiling/>
</profiling>
```

nontransactional-read

Range of values:	true false
Default value:	persistence-configuration-unit : false kodo-broker : true
Parent elements:	persistence-configuration persistence-configuration-unit and persistence-configuration persistence-configuration-unit kodo-broker

Function

Flag that specifies whether the Kodo runtime allows you to read data outside of a transaction. For more information, see “[kodo.NontransacationalRead](#)” in *Kodo Developer’s Guide*.

Example

```
<nontransactional-read>true</nontransactional-read>
```

nontransactional-write

Range of values:	true false
Default value:	false
Parent elements:	<p>persistence-configuration persistence-configuration-unit and persistence-configuration persistence-configuration-unit kodo-broker</p>

Function

Flag that specifies whether the you can modify persistent objects and perform persistence operations outside of a transaction. Changes will take effect in the next transaction. For more information, see “[kodo.NontransacationalWrite](#)” in *Kodo Developer’s Guide*.

Example

```
<nontransactional-write>false</nontransactional-write>
```

num-broadcast-threads

Range of values:	n/a
Default value:	n/a
Parent elements:	<p>persistence-configuration persistence-configuration-unit tcp-remote-commit-provider</p>

Function

Specifies the number of threads to create for the purpose of transmitting events to peers. You should increase this value as the number of concurrent transactions increases. The maximum number of concurrent transactions is a function of the size of the connection pool. See the

MaxActive property of `kodo.ConnectionFactoryProperties` in [“Using the Kodo Data Source”](#) in *Kodo Developer’s Guide*. Setting a value of 0 will result in behavior where the thread invoking `commit` will perform the broadcast directly. For more information, see [“TCP”](#) in the *Kodo Developer’s Guide*.

Example

```
<num-broadcast-threads>2</num-broadcast-threads>
```

operation-order-update-manager

Range of values:	n/a
Default value:	n/a
Parent elements:	<code>persistence-configuration</code> <code>persistence-configuration-unit</code>

Function

Defines an update manager that writes SQL in object-level operation order. For more information, see [“kodo.jdbc.UpdateManager”](#) in *Kodo Developer’s Guide*.

Example

```
<operation-order-update-manager/>
```

optimistic

Range of values:	true false
Default value:	true
Parent elements:	<p>persistence-configuration persistence-configuration-unit</p> <p>and</p> <p>persistence-configuration persistence-configuration-unit kodo-broker</p>

Function

Flag that specifies whether the datastore transactional mode is optimistic. For more information, see [“kodo.Optimistic”](#) in *Kodo Developer’s Guide*.

Example

```
<optimistic>true</optimistic>
```

oracle-dictionary

Range of values:	n/a
Default value:	n/a
Parent elements:	<p>persistence-configuration persistence-configuration-unit</p>

Function

Configuration values for the Oracle Dictionary persistence plugin. For a complete description of each of the values that you can specify, see [“Oracle Dictionary Configuration”](#) in the *Administration Console Online Help*.

Example

The `oracle-dictionary` element shares the same subelements as “[access-dictionary](#)” on [page B-7](#), plus the following subelements (defaults shown):

```
<oracle-dictionary>
...
  <use-triggers-for-auto-assign>false</use-triggers-for-auto-assign>
  <auto-assign-sequence-name>false</auto-assign-sequence-name>
  <use-set-form-of-use-for-unicode>true</use-set-form-of-use-for-unicode>
</oracle-dictionary>
```

oracle-savepoint-manager

Range of values:	n/a
Default value:	n/a
Parent elements:	<code>persistence-configuration</code> <code>persistence-configuration-unit</code>

Function

Specifies a plugin string (see “[Plugin Configuration](#)” in *Kodo Developer’s Guide*) describing the `kodo.kernel.SavepointManager` to use for managing transaction savepoints, in this case the default which is `kodo.jdbc.sql.OracleSavepointManager`. This plugin operates similarly to [jdbc3-savepoint-manager](#), but it uses Oracle-specific calls. This plugin requires using the Oracle JDBC driver and database, versions 9.2 or higher. This plugin implements savepoints by issuing a flush to the database. For more information, see “[Savepoints](#)” in the *Kodo Developer’s Guide*.

Example

```
<oracle-savepoint-manager/>
```

orm-file-jdor-mapping-factory

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

A plugin string (see [“Plugin Configuration”](#) in *Kodo Developer’s Guide*) describing the `kodo.jdo.jdbc.ORMFileJDORMMappingFactory` that stores mapping metadata in `.orm` files. This setting is valid for JDO 1.0. For more information, see [“Mapping Factory”](#) in *Kodo Developer’s Guide*.

Example

```
<orm-file-jdor-mapping-factory>
  <use-schema-validation>true</use-schema-validation>
  <mapping>mapping</mapping>
  <urls>t3://localhost:7001/metadata.jar</urls>
  <files>com/file1;com/file2</files>
  <classpath-scan>build</classpath-scan>
  <types>classname1;classname2</types>
  <store-mode>1</store-mode>
  <strict>>false</strict>
  <resources>com/aaa/package.jdo;com/bbb/package.jdo</resources>
  <scan-top-down>>false</scan-top-down>
</orm-file-jdor-mapping-factory>
```


order-dirty-objects

Range of values:	true false
Default value:	false
Parent elements:	persistence-configuration persistence-configuration-unit kodo-broker

Function

Flag that specifies whether Kodo maintains the order that changes were made to dirty objects when writing the changes back to the database.

Example

```
<order-dirty-objects>false</order-dirty-objects>
```

Password

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit wls81-jmx

Function

Password that Kodo should use to access the WebLogic MBeanServer. This must be set. For more information, see [“Optional Parameters in WebLogic 8.1 Group”](#) in the *Kodo Developer’s Guide*.

Example

```
<Password>admin</Password>
```

persistence-configuration

Requirements:	n/a
Default value:	n/a
Parent elements:	n/a

Function

Root element of the persistence deployment descriptor.

persistence-configuration-unit

Requirements:	n/a
Default value:	n/a
Parent elements:	persistence-configuration

Function

Contains the deployment information for a persistence unit that is available in WebLogic Server.

pessimistic-lock-manager

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Specifies the `kodo.jdbc.kernel.PessimisticLockManager`, which uses `SELECT FOR UPDATE` statements (or the database's equivalent) to lock the database rows corresponding to

locked objects. This lock manager does not distinguish between read locks and write locks; all locks are write locks. For more information, see [“Lock Manager”](#) in *Kodo Developer’s Guide*.

Example

```
<pessimistic-lock-manager>
  <version-check-on-read-lock>false</version-check-on-read-lock>
  <version-update-on-write-lock>false</version-update-on-write-lock>
</pessimistic-lock-manager>
```

persistence-mapping-defaults

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Configures the mapping defaults if you set `default-mapping-defaults` to `jpa`. For a complete description of each of the values that you can specify, see [“Mapping Defaults”](#) in *Kodo Developer’s Guide*.

Example

The `persistence-mapping-defaults` element shares the same subelements as [“mapping-defaults-impl”](#) on page B-150.

pointbase-dictionary

Range of values:	n/a
Default value:	n/a
Parent elements:	<p>persistence-configuration</p> <p>persistence-configuration-unit</p>

Function

Configuration values for the Pointbase Dictionary persistence plugin. For a complete description of each of the values that you can specify, see [“Pointbase Dictionary Configuration”](#) in the *Administration Console Online Help*.

Example

The `pointbase-dictionary` element shares the same subelements as [“access-dictionary”](#) on [page B-7](#).

populate-data-cache

Range of values:	true false
Default value:	true
Parent elements:	<p>persistence-configuration</p> <p>persistence-configuration-unit</p> <p>kodo-broker</p>

Function

Flag that specifies whether to populate the data cache. If your transaction will visit objects that you know are very unlikely to be accessed by other transactions, for example an exhaustive report run only once a month, you can turn off population of the datacache so that the transaction does not fill the entire data cache with objects that will not be accessed again.

Example

```
<populate-data-cache>true</populate-data-cache>
```

Port

Range of values:	n/a
Default value:	n/a
Parent elements:	<p>persistence-configuration</p> <p>persistence-configuration-unit</p> <p>mx4j1-jmx</p>

Function

Port on which the RMI registry naming service listens. Defaults to 1099. For more information, see [“Optional Parameters in Remote Group”](#) in the *Kodo Developer’s Guide*.

Example

```
<Port>1099</Port>
```

port

Range of values:	Valid port
Default value:	See below (depends on parent element)
Parent elements:	<p> persistence-configuration persistence-configuration-unit tcp-transport </p> <p>and</p> <p> persistence-configuration persistence-configuration-unit tcp-remote-commit-provider </p>

Function

Specifies the port on which the server will listen. This setting is used by clients, not by the server.

If the parent element is [tcp-transport](#), the element defaults to 5637. For more information, see “[Standalone Persistence Server](#)” in *Kodo Developer’s Guide*.

If the parent element is [tcp-remote-commit-provider](#), the element defaults to 5636. For more information, see “[TCP](#)” in *Kodo Developer’s Guide*.

Example

```
<port>5637</port>
```

postgres-dictionary

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Configuration values for the Postgres Dictionary persistence plugin. For a complete description of each of the values that you can specify, see [“Postgres Dictionary Configuration”](#) in the *Administration Console Online Help*.

Example

The `postgres-dictionary` element shares the same subelements as [“access-dictionary”](#) on [page B-7](#), plus the following subelements (defaults shown):

```
<postgres-dictionary>
...
  <all-sequences-sql>SELECT NULL AS SEQUENCE_SCHEMA, relname AS
SEQUENCE_NAME FROM pg_class WHERE relkind='S'</all-sequences-sql>
  <named-sequences-from-all-schemas-sql>SELECT NULL AS SEQUENCE_SCHEMA,
relname AS SEQUENCE_NAME FROM pg_class WHERE relkind='S' AND relname =
?</named-sequences-from-all-schemas-sql>
  <all-sequences-from-one-schema-sql>SELECT NULL AS SEQUENCE_SCHEMA,
relname AS SEQUENCE_NAME FROM pg_class, pg_namespace WHERE relkind='S' AND
pg_class.relnamespace = pg_namespace.oid AND nspname =
?</all-sequences-from-one-schema-sql>
  <named-sequences-from-one-schema-sql>SELECT NULL AS SEQUENCE_SCHEMA,
relname AS SEQUENCE_NAME FROM pg_class, pg_namespace WHERE relkind='S' AND
pg_class.relnamespace = ?</named-sequences-from-one-schema-sql>
  <supports-set-fetch-size>true</supports-set-fetch-size>
</postgres-dictionary>
```

primary-key-column

Range of values: String

Default value: ID

Parent elements: persistence-configuration
 persistence-configuration-unit
 class-table-jdbc-seq
and
persistence-configuration
 persistence-configuration-unit
 table-jdbc-seq
and
persistence-configuration
 persistence-configuration-unit
 table-schema-factory
and
persistence-configuration
 persistence-configuration-unit
 value-table-jdbc-seq

Function

Specifies the name of the table's numeric primary key column. For more information, see [“Generators”](#) or [“Schema Factory”](#) in the *Kodo Developer's Guide*.

Example

```
<primary-key-column>ID</primary-key-column>
```


primary-key-value

Range of values:	String
Default value:	DEFAULT
Parent elements:	<p>persistence-configuration</p> <p>persistence-configuration-unit</p> <p>value-table-jdbc-seq</p>

Function

Specifies the primary key that is used by this instance. For more information, see [“Generators”](#) in the *Kodo Developer’s Guide*.

Example

```
<primary-key-value>DEFAULT</primary-key-value>
```

primary-keys

Range of values:	true false
Default value:	false
Parent elements:	<p>persistence-configuration</p> <p>persistence-configuration-unit</p> <p>lazy-schema-factory</p>

Function

Flag that specifies whether to read automatically primary key information during schema validation. For more information, see [“Schema Factory”](#) in the *Kodo Developer’s Guide*.

Example

```
<primary-keys>>false</primary-keys>
```

profiling

Range of values:	n/a
Default value:	n/a
Parent elements:	weblogic-enterprise-bean

Function

Enables you to configure profiling of your application code. For more information, see [“Profiling”](#) in *Kodo Developer’s Guide*.

Example

```
<profiling>
  <export-profiling>
    <interval-millis>-1</interval-millis>
    <base-name>name</base-name>
  </export-profiling>
</profiling>

<profiling>
  <gui-profiling/>
</profiling>

<profiling>
  <local-profiling/>
</profiling>

<profiling>
  <none-profiling/>
</profiling>
```

profiling-proxy-manager

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Configures the profiling proxy manager. For more information, see [“Custom Proxies”](#) in the *Kodo Developer’s Guide*.

Example

```
<profiling-proxy-manager>  
  <assert-allowed-type>false</assert-allowed-type>  
  <track-changes>true</track-changes>  
</profiling-proxy-manager>
```

properties

Range of values:	n/a
-------------------------	-----

Default value:	n/a
-----------------------	-----

Parent elements:	<pre>persistence-configuration persistence-configuration-unit aggregate-listener custom-aggregate-listener and persistence-configuration persistence-configuration-unit custom-broker-factory and persistence-configuration persistence-configuration-unit custom-broker-impl and persistence-configuration persistence-configuration-unit custom-class-resolver and persistence-configuration persistence-configuration-unit custom-compatibility and persistence-configuration persistence-configuration-unit custom-connection-decorator and persistence-configuration persistence-configuration-unit custom-data-cache and persistence-configuration persistence-configuration-unit custom-dictionary</pre>
-------------------------	---

Parent elements:	and
(cont)	persistence-configuration persistence-configuration-unit custom-driver-data-source
	and
	persistence-configuration persistence-configuration-unit custom-filter-listener
	and
	persistence-configuration persistence-configuration-unit custom-jdbc-listener
	and
	persistence-configuration persistence-configuration-unit custom-lock-manager
	and
	persistence-configuration persistence-configuration-unit custom-logType
	and
	persistence-configuration persistence-configuration-unit custom-mapping-defaults
	and
	persistence-configuration persistence-configuration-unit custom-meta-data-respository
	and
	persistence-configuration persistence-configuration-unit custom-orphaned-key-action
	and
	persistence-configuration persistence-configuration-unit custom-persistence-server

Parent elements: and
(cont) persistence-configuration
 persistence-configuration-unit
 custom-proxy-manager

and
persistence-configuration
 persistence-configuration-unit
 custom-query-cache

and
persistence-configuration
 persistence-configuration-unit
 custom-query-compilation

and
persistence-configuration
 persistence-configuration-unit
 custom-remote-commit-provider

and
persistence-configuration
 persistence-configuration-unit
 custom-savepoint-manager

and
persistence-configuration
 persistence-configuration-unit
 custom-seq

and
persistence-configuration
 persistence-configuration-unit
 custom-sql-factory

and
persistence-configuration
 persistence-configuration-unit
 custom-update-manager

Function

Specifies one or more [property](#) elements.

property

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit properties

Function

Enables you to specify the name and value of a property.

Example

```
<property>
  <name>timeout</name>
  <value>1000</value>
</property>
```

proxy-manger-impl

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Configures the default proxy manager. For more information, see [“Custom Proxies”](#) in the *Kodo Developer’s Guide*.

Example

```
<proxy-manager-impl>
  <assert-allowed-type>>false</assert-allowed-type>
```

```

    <track-changes>true</track-changes>
</proxy-manager-impl>

```

query-caches

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

A plugin string (see [“Plugin Configuration”](#) in *Kodo Developer’s Guide*) describing the `kodo.datacache.QueryCache` implementation to use for caching of queries loaded from the data store. For more information, see [“Query Cache”](#) in *Kodo Developer’s Guide*.

Example

```

<query-caches>
  <default-query-cache>...</default-query-cache>
</query-caches>

<query-caches>
  <kodo-concurrent-query-cache>...</kodo-concurrent-query-cache>
</query-caches>

<query-caches>
  <gem-fire-query-cache>...</gem-fire-query-cache>
</query-caches>

<query-caches>
  <lru-query-cache>...</lru-query-cache>
</query-caches>

<query-caches>
  <tangosol-query-cache>...</tangosol-query-cache>
</query-caches>

```



```

<query-caches>
  <disabled-query-cache>...</disabled-query-cache>
</query-caches>

<query-caches>
  <custom-query-cache>...</custom-query-cache>
</query-caches>

```

quoted-numbers-in-queries

Range of values:	true false
Default value:	false
Parent elements:	persistence-configuration persistence-configuration-unit compatibility

Function

Flag that specifies whether to interpret quoted numbers in query strings as numbers, as opposed to strings. Set to `true` to mimic the behavior of Kodo 3.1 and earlier and treat quoted numbers as numbers. For more information, see [“Compatibility Configuration”](#) in *Kodo Developer’s Guide*.

Example

```
<quoted-numbers-in-queries>false</quoted-numbers-in-queries>
```

read-lock-level

Range of values:	<code>none</code> <code>read</code> <code>write</code> numeric values for lock-manager specific lock levels
Default value:	<code>read</code>
Parent elements:	<code>persistence-configuration</code> <code>persistence-configuration-unit</code>

Function

Sets the default read level at which to lock objects retrieved during a non-optimistic transaction. Valid values include:

- `none`—No lock level.
- `read`—Read lock level.
- `write`—Write lock level.
- Numeric values for lock-manager specific lock levels.

Note: For the default JDBC lock manager, the `read` and `write` lock levels are equivalent.

For more information, see “[Object Locking](#)” in the *Kodo Developer’s Guide*.

Example

```
<read-lock-level>read</read-lock-level>
```

recover-action

Range of values:	none clear
Default value:	none
Parent elements:	persistence-configuration persistence-configuration-unit cluster-remote-commit-provider

Function

Specifies the action to take during recovery. Valid values include:

- none—No action is performed.
- clear—Clears the notifications.

For more information, see [“Remote Commit Provider Configuration”](#) in the *Kodo Developer’s Guide*.

Example

```
<recover-action>none</recover-action>
```

recovery-time-millis

Range of values:	Integer
Default value:	15000
Parent elements:	persistence-configuration persistence-configuration-unit tcp-remote-commit-provider

Function

Specifies the amount of time to wait in milliseconds before attempting to reconnect to a peer of the cluster when connectivity to the peer is lost. For more information, see [“TCP”](#) in the *Kodo Developer’s Guide*.

Example

```
<recovery-time-millis>15000</recovery-time-millis>
```

resources

Range of values:	Valid resource paths
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit tangosol-data-cache

Function

Specifies a semicolon-separated list of resource paths to metadata files or jar archives. Each jar archive will be scanned for annotated JPA entities or JDO metadata files based on your metadata factory. For more information, see [“Metadata Factory”](#) in *Kodo Developer’s Guide*.

Example

```
<resources>com/aaa/package.jdo;com/bbb/package.jdo</resources>
```

restore-state

Range of values:	<p><code>persistence-configuration-unit</code>: all immutable none true false</p> <p><code>kodo-broker</code>: 1 2 3 4 5</p>
Default value:	none
Parent elements:	<p>persistence-configuration persistence-configuration-unit and persistence-configuration persistence-configuration-unit kodo-broker</p>

Function

Flag that specifies whether to restore managed fields to their pre-transaction values when a rollback occurs.

Valid values include:

- `all` (1)—Restores all managed values.
- `none` (2)—Do not roll back state. The object becomes hollow and will reload its state the next time it is accessed.
- `immutable` (3)—Restores immutable values (primitives, primitive wrappers, strings) and clears mutable values so that they are reloaded on next access.
- `true` (4)—Same as `all`.
- `false` (5)—Same as `none`.

For more information, see “[Restoring State](#)” in *Kodo Developer’s Guide*.

Example

```
<restore-state>none</restore-state>
```

result-set-type

Range of values:	forward-only scroll-sensitive scroll-insensitive
Default value:	forward-only
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Specifies the JDBC result set type to use when fetching return lists. This setting can also be varied at runtime.

Valid values include:

- `forward-only`—Result set whose cursor may move only forward. The result set is not updatable.
- `scroll-sensitive`—Scrollable result set that is sensitive to changes made by others. The result set is updatable.
- `scroll-insensitive`—Scrollable result set that is not sensitive to changes made by others. The result set is updatable.

For more information, see [“Large Result Sets”](#) in *Kodo Developer’s Guide*.

Example

```
<result-set-type>forward-only</result-set-type>
```

retain-state

Range of values:	true false
Default value:	See below (depends on parent)
Parent elements:	<p>persistence-configuration persistence-configuration-unit and persistence-configuration persistence-configuration-unit kodo-broker</p>

Function

Flag that specifies whether persistent fields retain their values on transaction commit. This element defaults to `true` when a child element of the `persistence-configuration-unit` and `false` when a child element of the `kodo-broker` element.

For more information, see “[kodo.RetainState](#)” in *Kodo Developer’s Guide*.

Example

```
<retain-state>true</retain-state>
```

retry-class-registration

Range of values:	true false
Default value:	false
Parent elements:	<p>persistence-configuration persistence-configuration-unit</p>

Function

Flag that specifies whether to log a warning and defer registration instead of throwing an exception when a persistent class cannot be fully processed. This element should only be used in complex classloader situations where security is preventing Kodo from reading registered

classes. Setting this to `true` unnecessarily may obscure more serious problems. For more information, see “[kodo.RetryClassRegistration](#)” in *Kodo Developer’s Guide*.

Example

```
<retry-class-registration>false</retry-class-registration>
```


scan-top-down

Range of values:	<code>true false</code>
Default value:	<code>false</code>
Parent elements:	<pre> persistence-configuration persistence-configuration-unit deprecated-jdo-meta-data-factory and persistence-configuration persistence-configuration-unit extension-deprecated-jdo-mapping-factory and persistence-configuration persistence-configuration-unit jdo-meta-data-factory and persistence-configuration persistence-configuration-unit jdor-mapping-factory and persistence-configuration persistence-configuration-unit mapping-file-deprecated-jdo-mapping-factory and persistence-configuration persistence-configuration-unit orm-file-jdor-mapping-factory and persistence-configuration persistence-configuration-unit table-deprecated-jdo-mapping-factory </pre>

Function

Flag that specifies whether to search for metadata files top-down in the package tree. Kodo defaults to bottom-up scanning. For example, when scanning metadata for class `C`, Kodo looks first for `C.jdo`, then `package.jdo` in `C`'s package, then `package.jdo` in the parent package, and

so on to the package root. For more information, see [“Metadata Factory”](#) in *Kodo Developer’s Guide*.

Example

```
<scan-top-down>>false</scan-top-down>
```

schema

Range of values:	String
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Specifies the default schema name to prepend to unqualified table names. In addition, specifies the schema in which Kodo creates new tables. For more information, see [“Default Schema”](#) in the *Kodo Developer’s Guide*.

Example

```
<schema>SCHEMA</schema>
```

schema-column

Range of values:	String
Default value:	SCHEMA_DEF
Parent elements:	<p>persistence-configuration</p> <p>persistence-configuration-unit</p> <p>table-schema-factory</p>

Function

Specifies the name of the table's string column for holding the schema definition as an XML string. For more information, see [“Schema Factory”](#) in the *Kodo Developer's Guide*.

Example

```
<schema-column>SCHEMA_DEF</schema-column>
```

schemas

Range of values:	n/a
Default value:	n/a
Parent elements:	<p>persistence-configuration</p> <p>persistence-configuration-unit</p>

Function

Specifies a comma-separated list of the schemas and/or tables used for your persistent data. For more information, see [“Schemas List”](#) in the *Kodo Developer's Guide*.

Example

```
<schemas>BUSOBSJS,GENERAL.ADDRESS,.SYSTEM_INFO </schemas>
```

sequence

Range of values:	n/a
Default value:	OPENJPA_SEQUENCE
Parent elements:	persistence-configuration persistence-configuration-unit native-jdbc-seq

Function

Specifies the name of the database sequence. For more information, see [“Generators”](#) in the *Kodo Developer’s Guide*.

Example

```
<sequence>OPENJPA_SEQUENCE</sequence>
```

sequence-column

Range of values:	String
-------------------------	--------

Default value:	SEQUENCE_VALUE
-----------------------	----------------

Parent elements:	<p>persistence-configuration persistence-configuration-unit class-table-jdbc-seq</p> <p>and</p> <p>persistence-configuration persistence-configuration-unit table-jdbc-seq</p> <p>and</p> <p>persistence-configuration persistence-configuration-unit value-table-jdbc-seq</p>
-------------------------	--

Function

Specifies the name of the column that holds the current sequence value. For more information, see [“Generators”](#) in the *Kodo Developer’s Guide*.

Example

```
<sequence-column>SEQUENCE_VALUE</sequence-column>
```

sequence-name

Range of values:	String
Default value:	OPENJPA_SEQUENCE
Parent elements:	<p>persistence-configuration</p> <p>persistence-configuration-unit</p> <p>native-jdbc-seq</p>

Function

Specifies the name of the database sequence. For more information, see [“Generators”](#) in the *Kodo Developer’s Guide*.

Example

```
<sequence-name>OPENJPA_SEQUENCE</sequence-name>
```

ServiceURL

Range of values:	Valid service URL
Default value:	service:jmx:rmi://localhost/jndi/jmxservice
Parent elements:	<p>persistence-configuration</p> <p>persistence-configuration-unit</p> <p>jmx2-jmx</p>

Function

JMX service URL name under which to register the JMX Connector Server. This value defaults to `service:jmx:rmi://localhost/jndi/jmxservice`, indicating that the RMI connector will be used and registered under the JNDI name `jmxservice`. For more information, see [“Optional Parameters in JSR 160 Group”](#) in *Kodo Developer’s Guide*.

Example

```
<ServiceURL>service:jmx:rmi://localhost/jndi/jmxservice</ServiceURL>
```

simple-driver-data-source

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Configures the Kodo datasource implementation. For more information, see [“Using the Kodo DataSource”](#) in *Kodo Developer’s Guide*.

Example

```
<simple-driver-data-source>
  <connection-user-name>user</connection-user-name>
  <login-timeout>10</login-timeout>
  <connection-password>password</connection-password>
  <connection-url>jdbc:hsqldb:db-hypersonic</connection-url>
  <connection-driver-name>org.hsqldb.jdbcDrier</connection-driver-name>
</simple-driver-data-source>
```

single-file

Range of values:	true false
Default value:	false
Parent elements:	persistence-configuration persistence-configuration-unit mapping-file-deprecated-jdo-mapping-factory

Function

Flag specifying whether or not to use a single file.

Example

```
<single-file>>false</single-file>
```

single-jvm-exclusive-lock-manager

Range of values:	n/a
Default value:	n/a
Parent elements:	<p>persistence-configuration</p> <p>persistence-configuration-unit</p>

Function

Specifies the `kodo.kernel.SingleJVMExclusiveLockManager`. This lock manager uses in-memory mutexes to obtain exclusive locks on object ids. It does not perform any database-level locking. Also, it does not distinguish between read and write locks; all locks are write locks. For more information, see [“Lock Manager”](#) in *Kodo Developer’s Guide*.

Example

```
<single-jvm-exclusive-lock-manager>
  <version-check-on-read-lock>>false</version-check-on-read-lock>
  <version-check-on-write-lock>>false</version-check-on-write-lock>
</single-jvm-exclusive-lock-manager>
```


single-jvm-remote-commit-provider

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Configures Kodo to share commit notifications among `persistenceManagerFactories` in the same JVM. For more information, see [“Remote Commit Provider Configuration”](#) in the *Kodo Developer’s Guide*.

Example

```
<single-jvm-remote-commit-provider>  
  <name>kodo.RemoteCommitProvider</name>  
</single-jvm-remote-commit-provider>
```

soft-reference-size

Range of values:	Integer
Default value:	-1
Parent elements:	<p> persistence-configuration persistence-configuration-unit kodo-concurrent-data-cache and persistence-configuration persistence-configuration-unit kodo-concurrent-query-cache and persistence-configuration persistence-configuration-unit lru-data-cache and persistence-configuration persistence-configuration-unit lru-query-cache and persistence-configuration persistence-configuration-unit query-compilation-cache </p>

Function

Sets the number of soft references that Kodo maintains. When the maximum cache or query size is reached, random entries are moved to a soft reference map so that they are maintained for awhile longer. By default, this value is set to -1 which specifies that an unlimited number of soft references are maintained. Set the property to 0 to disable soft references. For more information, see “[Configuring the Data Cache Size](#)” and “[Query Cache](#)” in *Kodo Developer’s Guide*.

Example

```
<soft-reference-size>-1</soft-reference-size>
```

so-timeout

Range of values:	Integer
Default value:	0
Parent elements:	<p>persistence-configuration</p> <p>persistence-configuration-unit</p> <p>http-transport</p>

Function

Specifies the socket read timeout in milliseconds. For more information, see [“Standalone Persistence Server”](#) in *Kodo Developer’s Guide*.

Example

```
<so-timeout>0</so-timeout>
```

sql-server-dictionary

Range of values:	n/a
Default value:	n/a
Parent elements:	<p>persistence-configuration</p> <p>persistence-configuration-unit</p>

Function

Configuration values for the SQLServer Dictionary persistence plugin. For a complete description of each of the values that you can specify, see [“SQLServer Dictionary Configuration”](#) in the *Administration Console Online Help*.

Example

The `sql-server-dictionary` element shares the same subelements as [“access-dictionary”](#) on [page B-7](#), plus the following subelements (defaults shown):

```
<sql-server-dictionary>
...
  <unique-identifier-as-varbinary>true</unique-identifier-as-varbinary>
</sql-server-dictionary>
```

stack-execution-context-name-provider

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Provider examines the current thread's stack and builds an execution context name based on that information. For more information, see [“Configuring the Execution Context Name Provider”](#) in *Kodo Developer's Guide*.

Example

```
<stack-execution-context-name-provider>
  <style>full</style>
</stack-execution-context-name-provider>
```

store-mode

Range of values:	Integer
-------------------------	---------

Default value:	n/a
-----------------------	-----

Parent elements:	<p>persistence-configuration persistence-configuration-unit deprecated-jdo-meta-data-factory</p> <p>and</p> <p>persistence-configuration persistence-configuration-unit extension-deprecated-jdo-mapping-factory</p> <p>and</p> <p>persistence-configuration persistence-configuration-unit jdo-meta-data-factory</p> <p>and</p> <p>persistence-configuration persistence-configuration-unit jdor-mapping-factory</p> <p>and</p> <p>persistence-configuration persistence-configuration-unit kodo-persistence-mapping-factory</p> <p>and</p> <p>persistence-configuration persistence-configuration-unit kodo-persistence-meta-data-factory</p> <p>and</p> <p>persistence-configuration persistence-configuration-unit mapping-file-deprecated-jdo-mapping-factory</p>
-------------------------	--

Parent elements: and
(cont) persistence-configuration
 persistence-configuration-unit
 orm-file-jdor-mapping-factory

and
persistence-configuration
 persistence-configuration-unit
 table-deprecated-jdo-mapping-factory

and
persistence-configuration
 persistence-configuration-unit
 table-jdor-mapping-factory

Function

Specifies the data store mode.

Example

```
<store-mode>1</store-mode>
```

strict

Range of values:	true false
Default value:	false
Parent elements:	<p>persistence-configuration persistence-configuration-unit deprecated-jdo-meta-data-factory</p> <p>and</p> <p>persistence-configuration persistence-configuration-unit extension-deprecated-jdo-mapping-factory</p> <p>and</p> <p>persistence-configuration persistence-configuration-unit jdo-meta-data-factory</p> <p>and</p> <p>persistence-configuration persistence-configuration-unit jdor-mapping-factory</p> <p>and</p> <p>persistence-configuration persistence-configuration-unit kodo-persistence-mapping-factory</p> <p>and</p> <p>persistence-configuration persistence-configuration-unit kodo-persistence-meta-data-factory</p> <p>and</p> <p>persistence-configuration persistence-configuration-unit mapping-file-deprecated-jdo-mapping-factory</p>

Parent elements:	and
(cont)	persistence-configuration persistence-configuration-unit orm-file-jdor-mapping-factory
	and
	persistence-configuration persistence-configuration-unit table-deprecated-jdo-mapping-factory
	and
	persistence-configuration persistence-configuration-unit table-jdor-mapping-factory

Function

Flag that specifies whether strict mode is enabled.

Example

```
<strict>true</strict>
```

strict-identity-values

Range of values:	true false
Default value:	false
Parent elements:	persistence-configuration persistence-configuration-unit compatibility

Function

Flag that specifies whether to require identity values used for finding application identity instances to be of the exact right type. By default, Kodo allows stringified identity values, and performs conversions between numeric types. For more information, see “[Compatibility Configuration](#)” in *Kodo Developer’s Guide*.

Example

```
<strict-identity-values>false</strict-identity-values>
```

style

Range of values:	line partial full
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit stack-execution-context-name-provider

Function

Provider examines the current thread's stack and builds an execution context name based on that information and the following setting:

- `line`—Name will be the first non-Kodo stack line, including class name, method name, and line number, if available.
- `partial`—Full stack trace minus the Kodo lines will be used.
- `full`—Entire stack trace will be used.

For more information, see [“Configuring the Execution Context Name Provider”](#) in *Kodo Developer’s Guide*.

Example

```
<style>full</style>
```

subclass-fetch-mode

Range of values:	join multiple none parallel single
Default value:	join
Parent elements:	<p>persistence-configuration</p> <p> persistence-configuration-unit</p> <p> stack-execution-context-name-provider</p>

Function

Specifies the method to use to select subclass data when it is in other tables. For a description of the valid values, see under “[eager-fetch-mode](#)” on page B-93. For complete details, see “[Eager Fetching](#)” in *Kodo Developer’s Guide*.

Example

```
<subclass-fetch-mode>join</subclass-fetch-mode>
```

sybase-dictionary

Range of values:	n/a
Default value:	n/a
Parent elements:	<p>persistence-configuration</p> <p> persistence-configuration-unit</p>

Function

Configuration values for the Sybase Dictionary persistence plugin. For a complete description of each of the values that you can specify, see “[Sybase Dictionary Configuration](#)” in the *Administration Console Online Help*.

Example

The `sybase-dictionary` element shares the same subelements as “[access-dictionary](#)” on [page B-7](#), plus the following subelements (defaults shown):

```
<sybase-dictionary>
...
  <create-identity-column>true</create-identity-column>
  <identity-column-name>UNQ_INDEX</identity-column-name>
</sybase-dictionary>
```

sync-with-managed-transactions

Range of values:	true false
Default value:	false
Parent elements:	persistence-configuration persistence-configuration-unit kodo-broker

Function

Flag that specifies whether to sync with managed transactions. For more information, see “[kodo.BrokerImpl](#)” in *Kodo Developer’s Guide*.

Example

```
<sync-with-managed-transactions>false</sync-with-managed-transactions>
```

synchronize-mappings

Range of values:	true false
Default value:	false
Parent elements:	persistence-configuration-unit

Function

Controls whether Kodo attempts to run the mapping tool on all persistent classes to synchronize their mappings and schema at runtime. This element save you the trouble of running the mapping tool manually, and is useful for rapid test and debug cycles. For more information, see [“Runtime Forward Mapping”](#) in Kodo Developer’s Guide.

Example

```
<synchronize-mappings>false</synchronize-mappings>
```

table

Range of values:	String
Default value:	See below
Parent elements:	<p> persistence-configuration persistence-configuration-unit class-table-jdbc-seq and persistence-configuration persistence-configuration-unit table-jdbc-seq and persistence-configuration persistence-configuration-unit table-jdor-mapping-factory and persistence-configuration persistence-configuration-unit table-schema-factory and persistence-configuration persistence-configuration-unit value-table-jdbc-seq </p>

Function

Specifies the name of the table. For more information, see “[Mapping Factory](#)” in *Kodo Developer’s Guide*. The default value depends on the parent element:

- `class-table-jdbc-seq`: OPENJPA_SEQUENCE_TABLE
- `table-jdbc-seq`: OPENJPA_SEQUENCE_TABLE
- `table-jdor-mapping-factory`: KODO_JDO_MAPPINGS
- `table-schema-factory`: OPENJPA_SCHEMA
- `value-table-jdbc-seq`: OPENJPA_SEQUENCES_TABLE

Example

```
<table>KODO_JDO_MAPPING</table>
```

table-deprecated-jdo-mapping-factory

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

A plugin string (see “[Plugin Configuration](#)” in *Kodo Developer’s Guide*) describing the `kodo.jdo.jdbc.TableJDORMappingFactory` that stores mapping metadata as XML strings in a database table. This setting is valid for JDO 1.0. For more information, see “[Mapping Factory](#)” in *Kodo Developer’s Guide*.

Example

```
<table-deprecated-jdo-mapping-factory>
  <table-name>JDO_MAPPING</table-name>
  <urls>t3://localhost:7001/metadata.jar</urls>
  <classpath-scan>build</classpath-scan>
  <types>classname1;classname2</types>
  <mapping-column>MAPPING_DEF</mapping-column>
  <store-mode></store-mode>
  <strict>>false</strict>
  <name-column>NAME</name-column>
  <use-schema-validation>>false</use-schema-validation>
  <single-row>>false</single-row>
  <files>com/file1;com/file2</files>
  <scan-top-down>>false</scan-top-down>
  <resources>com/aaa/package.jdo;com/bbb/package.jdo</resources>
</table-deprecated-jdo-mapping-factory>
```

table-jdbc-seq

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Specifies a plugin string (see [“Plugin Configuration”](#) in *Kodo Developer’s Guide*) describing the `kodo.kernel.Seq` interface to use to create your own custom generators, in this case `kodo.jdbc.kernel.TableJDBCSeq`.

The `TableJDBCSeq` uses a special single-row table to store a global sequence number. If the table does not already exist, it is created the first time you run the mapping tool’s on a class that requires it. You can also use the class’ main method or the `sequencetable` shell/bat script to manipulate the table; see the `TableJDBCSeq.main` method Javadoc for usage details. For more information, see [“Generators”](#) in the *Kodo Developer’s Guide*.

Example

```
<table-jdbc-seq>
  <type>0</type>
  <allocate>50</allocate>
  <table-name>OPENJPA_SEQUENCE_TABLE</table-name>
  <table>OPENJPA_SEQUENCE_TABLE</table>
  <primary-key-column>ID</primary-key-column>
  <sequence-column>SEQUENCE_VALUE</sequence-column>
  <increment>1</increment>
</table-jdbc-seq>
```

table-jdor-mapping-factory

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

A plugin string (see “[Plugin Configuration](#)” in *Kodo Developer’s Guide*) describing the `kodo.jdo.jdbc.TableJDORMappingFactory` that stores mapping metadata as XML strings in a database table. This setting is valid for JDO 2.0. For more information, see “[Mapping Factory](#)” in *Kodo Developer’s Guide*.

Example

```
<table-jdor-mapping-factory>
  <use-schema-validation>false</use-schema-validation>
  <type-column>0</type-column>
  <constraint-names>false</constraint-names>
  <table>KODO_JDO_MAPPINGS</table>
  <types>classname1;classname2</types>
  <store-mode></store-mode>
  <mapping-column>false</mapping-column>
  <strict>false</strict>
  <name-column>NAME</name-column>
</table-jdor-mapping-factory>
```


table-lock-update-manager

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Defines an update manager capable of statement batching, that ignores foreign keys and autoincrement, but optimizes for table level locking. This is useful when using table-level locking and trying to avoid deadlocks. For more information, see “[kodo.jdbc.UpdateManager](#)” and “[Statement Batching](#)” in *Kodo Developer’s Guide*.

Example

```
<table-lock-update-manager>  
  <maximize-batch-size>true</maximize-batch-size>  
</table-lock-update-manager>
```

table-name

Range of values:	String
Default value:	See below
Parent elements:	<p> persistence-configuration persistence-configuration-unit class-table-jdbc-seq </p> <p>and</p> <p> persistence-configuration persistence-configuration-unit native-jdbc-seq </p> <p>and</p> <p> persistence-configuration persistence-configuration-unit table-deprecated-jdo-mapping-factory </p> <p>and</p> <p> persistence-configuration persistence-configuration-unit table-jdbc-seq </p> <p>and</p> <p> persistence-configuration persistence-configuration-unit table-schema-factory </p> <p>and</p> <p> persistence-configuration persistence-configuration-unit value-table-jdbc-seq </p>

Function

Specifies the name of the table.

The default varies based on the parent element, as follows:

- **class-table-jdbc-seq**: OPENJPA_SEQUENCES_TABLE. For more information, see “Generators” in the *Kodo Developer’s Guide*.

- `native-jdbc-seq`: DUAL. For more information, see “[Generators](#)” in the *Kodo Developer’s Guide*.
- `table-deprecated-jdo-mapping-factory`: JDO_MAPPING. For more information, see “[Mapping Factory](#)” in *Kodo Developer’s Guide*.
- `table-jdbc-seq`: OPENJPA_SEQUENCE_TABLE. For more information, see “[Generators](#)” in the *Kodo Developer’s Guide*.
- `table-schema-factory`: OPENJPA_SCHEMA. For more information, see “[Schema Factory](#)” in the *Kodo Developer’s Guide*.
- `value-table-jdbc-seq`: OPENJPA_SEQUENCE_TABLE. For more information, see “[Generators](#)” in the *Kodo Developer’s Guide*.

Example

```
<table-name>JDO_MAPPING</table-name>
```

table-schema-factory

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Specifies a plugin string (see “[Plugin Configuration](#)” in *Kodo Developer’s Guide*) describing the `kodo.jdbc.schema.SchemaFactory` to use to store and retrieve information about the database schema, in this case `kodo.jdbc.schema.TableSchemaFactory`.

This schema factory stores schema information as an XML document in a database table it creates for this purpose. If your JDBC driver doesn’t support the `java.sql.DatabaseMetaData` standard interface, but you still want some schema validation to occur at runtime, you might use this factory. It is not recommended for most users, though, because it is easy for the stored XML schema definition to get out-of-sync with the actual database.

For more information, see “[Schema Factory](#)” in the *Kodo Developer’s Guide*.

Example

```
<table-schema-factory>
  <schema-column>SCHEMA_DEF</schema-column>
  <table-name>OPENJPA_SCHEMA</table-name>
  <table>OPENJPA_SCHEMA</table>
  <primary-key-column>ID</primary-key-column>
</table-schema-factory>
```

tangosol-cache-name

Range of values:	String
Default value:	kodo
Parent elements:	persistence-configuration persistence-configuration-unit tangosol-data-cache

Function

Name of the Tangosol Coherence cache to use. For more information, see [“Tangosol Integration”](#) in *Kodo Developer’s Guide*.

Example

```
<tangosol-cache-name>kodo</tangosol-cache-name>
```

tangosol-cache-type

Range of values:	1 2 3
Default value:	2
Parent elements:	<p>persistence-configuration</p> <p> persistence-configuration-unit</p> <p> tangosol-data-cache</p>

Function

Type of Tangosol Coherence cache to use. Valid values include:

- 1—(named) Cache is looked up via the `com.tangosol.net.CacheFactory.getCache(String)` method. This method looks up the cache by the name as defined in the Coherence configuration.
- 2—(distributed)
- 3—(remote)

For more information, see [“Tangosol Integration”](#) in *Kodo Developer’s Guide*.

Example

```
<tangosol-cache-type>2</tangosol-cache-type>
```

tangosol-data-cache

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit data-caches

Function

Integrates with Tangosol's Coherence caching system. For more information, see [“Tangosol Integration”](#) in *Kodo Developer's Guide*.

Example

```
<data-caches>
  <tangosol-data-cache>
    <name>kodo.dataCache</name>
    <clear-on-close>false</clear-on-close>
    <tangosol-cache-type>2</tangosol-cache-type>
    <tangosol-cache-name>kodo</tangosol-cache-name>
    <eviction-schedule>15,45 15 * * 1</eviction-schedule>
  </tangosol-data-cache>
</data-caches>
```

tcp-remote-commit-provider

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Configures the TCP remote commit provider. For more information, see “TCP” in the *Kodo Developer’s Guide*.

Example

```
<tcp-remote-commit-provider>
  <name>kodo.RemoteCommitProvider</name>
  <max-idle>2</max-idle>
  <num-broadcast-threads>2</num-broadcast-threads>
  <recovery-time-millis>15000</recovery-time-millis>
  <max-active>2</max-active>
  <port>5636</port>
  <addresses>[ ]</addresses>
</tcp-remote-commit-provider>
```

tcp-transport

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Specifies the transport layer for the remote communication. For more information, see [“Standalone Persistence Server”](#) in *Kodo Developer’s Guide*.

Example

```
<tcp-transport>  
  <so-timeout>0</so-timeout>  
  <host>localhost</host>  
  <port>5637</port>  
</tcp-transport>
```

time-seeded-seq

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Specifies a plugin string (see [“Plugin Configuration”](#) in *Kodo Developer’s Guide*) describing the `kodo.kernel.Seq` interface to use to create your own custom generators, in this case `kodo.jdbc.kernel.TimeSeededSeq`.

This type uses an in-memory static counter, initialized to the current time in milliseconds and monotonically incremented for each value requested. It is only suitable for single-JVM environments. For more information, see “[Generators](#)” in the *Kodo Developer’s Guide*.

Example

```
<time-seeded-seq>
  <type>0</type>
  <increment>1</increment>
</time-seeded-seq>
```

topic

Range of values:	Valid topic
Default value:	topic/KodoCommitProviderTopic
Parent elements:	persistence-configuration persistence-configuration-unit jms-remote-commit-provider

Function

Specifies the topic to which the remote commit provider should publish notifications and subscribe for notifications sent from other JVMs. For more information, see “[JMS](#)” in the *Kodo Developer’s Guide*.

Example

```
<topic>topic/KodoCommitProviderTopic</topic>
```

topic-connection-factory

Range of values:	Valid connection factory
Default value:	java:/ConnectionFactory
Parent elements:	<p>persistence-configuration</p> <p>persistence-configuration-unit</p> <p>jms-remote-commit-provider</p>

Function

Specifies the JNDI name of a `javax.jms.TopicConnectionFactory` factory to use for finding topics. This setting may vary depending on the application server in use; consult the application server's documentation for details about the default JNDI name for the topic instance. For WebLogic, the JNDI name for the `TopicConnectionFactory` is `javax.jms.TopicConnectionFactory`. For more information, see [“JMS”](#) in the *Kodo Developer's Guide*.

Example

```
<topic-connection-factory>java:/ConnectionFactory</topic-connection-factory>
```

track-changes

Range of values:	true false
Default value:	true
Parent elements:	<p>persistence-configuration</p> <p>persistence-configuration-unit</p>

Function

Flag that specifies whether to use smart proxies. For more information, see [“Custom Proxies”](#) in the *Kodo Developer's Guide*.

Example

```
<track-changes>true</track-changes>
```

transaction-isolation

Range of values:	default none read-committed read-uncommitted repeatable-read serializable
Default value:	default
Parent elements:	persistence-configuration-unit

Function

Specifies the JDBC transaction isolation level to use.

- `default`—JDBC driver’s default level. This is the default.
- `none`—No transaction isolation.
- `read-committed`—Dirty reads are prevented; non-repeatable reads and phantom reads can occur.
- `read-uncommitted`—Dirty reads, non-reputable reads, and phantom reads can occur.
- `repeatable-read`—Dirty reads and non-repeatable reads are prevented; phantom reads can occur.
- `serializable`—Dirty reads, non-reputable reads, and phantom reads are prevented.

For more information, see [“Setting the Transaction Isolation”](#) in Kodo Developer’s Guide.

Example

```
<transaction-isolation>default</transaction-isolation>
```

transaction-mode

Range of values:	local managed
Default value:	local
Parent elements:	persistence-configuration-unit

Function

Specifies the transaction mode to use. You can override this setting per session.

- `local`—Perform transaction operations locally. This is the default.
- `managed`—Integrate with the application server’s managed global transactions.

For more information, see [“Integrating with the Transaction Manager”](#) in Kodo Developer’s Guide.

Example

```
<transaction-mode>local</transaction-mode>
```

transaction-name-execution-context-name-provider

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Provider returns the current transaction’s name, as defined by WebLogic transaction naming semantics. For more information, see [“Configuring the Execution Context Name Provider”](#) in *Kodo Developer’s Guide*.

Example

```
<transacation-name-execution-context-name-provider/>
```

type

Range of values:	n/a
Default value:	0
Parent elements:	<p> persistence-configuration persistence-configuration-unit class-table-jdbc-seq </p> <p>and</p> <p> persistence-configuration persistence-configuration-unit native-jdbc-seq </p> <p>and</p> <p> persistence-configuration persistence-configuration-unit table-jdbc-seq </p> <p>and</p> <p> persistence-configuration persistence-configuration-unit time-seeded-seq </p> <p>and</p> <p> persistence-configuration persistence-configuration-unit value-table-jdbc-seq </p>

Function

Specifies the type of generator. For more information, see “[Generators](#)” in the *Kodo Developer’s Guide*.

Example

```
<type>0</type>
```

type-column

Range of values:	0 1 2 3
Default value:	0
Parent elements:	persistence-configuration persistence-configuration-unit table-jdor-mapping-factory

Function

Specifies the name of the column that holds the mapping type. Valid types constants include:

- 0—Class mapping.
- 1—Named sequence.
- 2—System-level named query.
- 3—Class-level named query.

For more information, see [“Mapping Factory”](#) in *Kodo Developer’s Guide*.

Example

```
<type-column>0</type-column>
```

types

Range of values:	Valid persistent class names
Default value:	n/a
Parent elements:	<p>persistence-configuration</p> <p>persistence-configuration-unit</p> <p>tangosol-data-cache</p>

Function

Specifies a semicolon-separated list of fully-qualified persistent class names. For more information, see [“Metadata Factory”](#) in *Kodo Developer’s Guide*.

Example

```
<types>classname1; classname2</types>
```

URL

Range of values:	Valid URL
Default value:	t3://localhost:7001
Parent elements:	<p>persistence-configuration</p> <p>persistence-configuration-unit</p> <p>wls81-jmx</p>

Function

URL to which Kodo should connect to access the WebLogic MBeanServer. For more information, see [“Optional Parameters in WebLogic 8.1 Group”](#) in the *Kodo Developer’s Guide*.

Example

```
<URL>t3://localhost:7001</URL>
```

url

Range of values:	Valid URL
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit http-transport

Function

URL to which Kodo should connect to access the remote server. For more information, see [“Client Managers”](#) in *Kodo Developer’s Guide*.

Example

```
<url>t3://localhost:7001/remote-server</url>
```


urls

Range of values:	n/a
-------------------------	-----

Default value:	n/a
-----------------------	-----

Parent elements:	<p>persistence-configuration persistence-configuration-unit deprecated-jdo-meta-data-factory</p> <p>and</p> <p>persistence-configuration persistence-configuration-unit extension-deprecated-jdo-mapping-factory</p> <p>and</p> <p>persistence-configuration persistence-configuration-unit jdo-meta-data-factory</p> <p>and</p> <p>persistence-configuration persistence-configuration-unit jdor-mapping-factory</p> <p>and</p> <p>persistence-configuration persistence-configuration-unit kodo-persistence-mapping-factory</p> <p>and</p> <p>persistence-configuration persistence-configuration-unit kodo-persistence-meta-data-factory</p> <p>and</p> <p>persistence-configuration persistence-configuration-unit mapping-file-deprecated-jdo-mapping-factory</p>
-------------------------	--

Parent elements:	and
(cont)	<p>persistence-configuration persistence-configuration-unit orm-file-jdor-mapping-factory</p> <p>and</p> <p>persistence-configuration persistence-configuration-unit table-deprecated-jdo-mapping-factory</p>

Function

Specifies a semicolon-separated list of URLs of metadata files or jar archives. Each jar archive will be scanned for annotated JPA entities or JDO metadata files based on your metadata factory. For more information, see [“Metadata Factory”](#) in *Kodo Developer’s Guide*.

Example

```
<urls>t3://localhost:7001/metadata.jar</urls>
```

use-aliases

Range of values:	true false
Default value:	false
Parent elements:	<p>persistence-configuration persistence-configuration-unit class-table-jdbc-seq</p>

Function

Flag that specifies whether to use each class’ entity name as the primary key value of each row rather than the full classname. For more information, see [“Generators”](#) in the *Kodo Developer’s Guide*.

Example

```
<use-aliases>>false</use-aliases>
```

use-schema-validation

Range of values:	true false
-------------------------	--------------

Default value:	n/a
-----------------------	-----

Parent elements:	<p>persistence-configuration persistence-configuration-unit deprecated-jdo-meta-data-factory</p> <p>and</p> <p>persistence-configuration persistence-configuration-unit extension-deprecated-jdo-mapping-factory</p> <p>and</p> <p>persistence-configuration persistence-configuration-unit jdo-meta-data-factory</p> <p>and</p> <p>persistence-configuration persistence-configuration-unit jdor-mapping-factory</p> <p>and</p> <p>persistence-configuration persistence-configuration-unit mapping-file-deprecated-jdo-mapping-factory</p> <p>and</p> <p>persistence-configuration persistence-configuration-unit orm-file-jdor-mapping-factory</p> <p>and</p> <p>persistence-configuration persistence-configuration-unit table-deprecated-jdo-mapping-factory</p> <p>and</p> <p>persistence-configuration persistence-configuration-unit table-jdor-mapping-factory</p>
-------------------------	---

Function

Flag that specifies whether to use schema validation.

Example

```
<use-schema-validation>true</use-schema-validation>
```

user-object-execution-context-name-provider

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Provider returns the user object specified in the current EntityManager/PersistenceManager whose key is `com.solarmetric.profile.ExecutionContextNameProvider`. User objects can be set using the `OpenJPAEntityManager.putUserObject(Object, Object)` or `PersistenceManager.putUserObject(Object, Object)`. For more information, see [“Configuring the Execution Context Name Provider”](#) in *Kodo Developer’s Guide*.

Example

```
<user-object-execution-context-name-provider/>
```

UserName

Range of values:	n/a
Default value:	n/a
Parent elements:	<p>persistence-configuration</p> <p>persistence-configuration-unit</p> <p>wls81-jmx</p>

Function

Specifies the username that Kodo should use to access the WebLogic MBeanServer. This must be set. For more information, see [“Optional Parameters in WebLogic 8.1 Group”](#) in the *Kodo Developer’s Guide*.

Example

```
<wls81-jmx>
  <MBeanServerStrategy>any-create</MBeanServerStrategy>
  <EnableLogMBean>true</EnableLogMBean>
  <EnableRuntimeMBean>true</EnableRuntimeMBean>
  <URL>t3://localhost:7001</URL>
  <UserName>admin</UserName>
  <Password>admin</Password>
</wls81-jmx>
```

validate-false-returns-hollow

Range of values:	true false
Default value:	true
Parent elements:	persistence-configuration persistence-configuration-unit compatibility

Function

Flag that specifies whether to return hollow objects (objects for which no state has been loaded) from calls to `PersistenceManager.getObjectById(oid, false)`. This is the default behavior of Kodo 4.0 and above. Previous Kodo versions, however, always loaded the object from the datastore. Set this property to `false` to get the old behavior.

For more information, see [“Compatibility Configuration”](#) in *Kodo Developer’s Guide*.

Example

```
<validate-false-returns-hollow>true</validate-false-returns-hollow>
```

validate-true-checks-store

Range of values:	true false
Default value:	false
Parent elements:	persistence-configuration persistence-configuration-unit compatibility

Function

Flag that specifies whether to check the datastore to be sure the given `oid` exists. Prior to Kodo 4.0, calling `PersistenceManager.getObjectById(oid, true)` always checked the datastore to be sure the given `oid` existed, even when the corresponding object was already cached. Kodo

4.0 and above does not check the datastore when the instance is already cached and loaded. Set this property to `true` to mimic previous behavior.

For more information, see [“Compatibility Configuration”](#) in *Kodo Developer’s Guide*.

Example

```
<validate-true-checks-store>false</validate-true-checks-store>
```

value-table-jdbc-seq

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Specifies a plugin string (see [“Plugin Configuration”](#) in *Kodo Developer’s Guide*) describing the `kodo.kernel.Seq` interface to use to create your own custom generators, in this case

```
kodo.jdbc.kernel.ValueTableJDBCSeq.
```

This `Seq` is like the `table-jdbc-seq`, but has an arbitrary number of rows for sequence values, rather than a fixed pattern of one row per class. For more information, see [“Generators”](#) in the *Kodo Developer’s Guide*.

Example

```
<value-table-jdbc-seq>
  <type>0</type>
  <allocate>50</allocate>
  <table-name>OPENJPA_SEQUENCE_TABLE</table-name>
  <primary-key-value>DEFAULT</primary-key-value>
  <table>OPENJPA_SEQUENCE_TABLE</table>
  <primary-key-column>ID</primary-key-column>
  <sequence-column>SEQUENCE_VALUE</sequence-column>
  <increment>1</increment>
</value-table-jdbc-seq>
```

version-check-on-read-lock

Range of values:	true false
Default value:	false
Parent elements:	persistence-configuration persistence-configuration-unit pessimistic-lock-manager

Function

Specifies whether to perform version checking and incrementing behavior of the pessimistic lock manager on read operations. For more information, see [“Lock Manager”](#) in *Kodo Developer’s Guide*.

Example

```
<version-check-on-read-lock>false</version-check-on-read-lock>
```

version-check-on-write-lock

Range of values:	true false
Default value:	false
Parent elements:	persistence-configuration persistence-configuration-unit pessimistic-lock-manager

Function

Specifies whether to perform version checking and incrementing behavior of the pessimistic lock manager on write operations. For more information, see [“Lock Manager”](#) in *Kodo Developer’s Guide*.

Example

```
<version-check-on-write-lock>false</version-check-on-write-lock>
```


version-lock-manager

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Specifies the `kodo.kernel.VersionLockManager`. This lock manager does not perform any exclusive locking, but instead ensures read consistency by verifying that the version of all read-locked instances is unchanged at the end of the transaction. Furthermore, a write lock will force an increment to the version at the end of the transaction, even if the object is not otherwise modified. This ensures read consistency with non-blocking behavior. For more information, see [“Lock Manager”](#) in *Kodo Developer’s Guide*.

Example

```
<version-lock-manager>
  <version-check-on-read-lock>false</version-check-on-read-lock>
  <version-check-on-write-lock>false</version-check-on-write-lock>
</version-lock-manager>
```

wls81-jmx

Range of values:	n/a
Default value:	n/a
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Enable management and invoke the JMX management console in the local JVM. Supports optional parameters in the Management Group and WebLogic 8.1 Group, as described in

“Optional Parameters in Management Group” and “Optional Parameters in WebLogic 8.1 Group” in the *Kodo Developer’s Guide*.

Example

```
<wls81-jmx>
  <MBeanServerStrategy>any-create</MBeanServerStrategy>
  <EnableLogMBean>true</EnableLogMBean>
  <EnableRuntimeMBean>true</EnableRuntimeMBean>
  <URL>t3://localhost:7001</URL>
  <UserName>admin</UserName>
  <Password>admin</Password>
</wls81-jmx>
```

write-lock-level

Range of values:	none read write numeric values for lock-manager specific lock levels
Default value:	read
Parent elements:	persistence-configuration persistence-configuration-unit

Function

Sets the default write level at which to lock objects retrieved during a non-optimistic transaction. Valid values include:

- none—No lock level.
- read—Read lock level.
- write—Write lock level.
- Numeric values for lock-manager specific lock levels.

Note: For the default JDBC lock manager, the `read` and `write` lock levels are equivalent. For more information, see “Object Locking” in the *Kodo Developer’s Guide*.

Example

```
<write-lock-level>read</write-lock-level>
```

Persistence Configuration Schema Reference