

Oracle® Fusion Middleware

Programming Enterprise JavaBeans for Oracle WebLogic
Server

12c Release 1 (12.1.1)

E24972-02

January 2012

This document is a resource for software developers who develop applications that include WebLogic Server Enterprise JavaBeans (EJBs) using the Java Platform, Enterprise Edition 6.

Oracle Fusion Middleware Programming Enterprise JavaBeans for Oracle WebLogic Server, 12c Release 1 (12.1.1)

E24972-02

Copyright © 2007, 2012, Oracle and/or its affiliates. All rights reserved.

Primary Author: Jeff Schieli

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 is software or related documentation that 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 America, Inc., 500 Oracle Parkway, Redwood City, CA 94065.

This software or hardware 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 that may create a risk of personal injury. If you use this software or hardware in dangerous applications, then you shall be responsible to take all appropriate fail-safe, backup, redundancy, and other measures to ensure its safe use. Oracle Corporation and its affiliates disclaim any liability for any damages caused by use of this software or hardware in dangerous applications.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group.

This software or hardware 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

Preface	xxxi
Documentation Accessibility	xxxi
Conventions	xxxi
1 Introduction and Roadmap	
1.1 Document Scope and Audience.....	1-1
1.2 Guide to this Document	1-1
1.3 Related Documentation.....	1-2
1.3.1 EJB Documentation in WebLogic Server.....	1-2
1.3.2 Additional EJB Information	1-3
1.4 Comprehensive Examples for the EJB Developer	1-3
1.4.1 New EJB 3.1 Examples	1-3
1.4.2 EJB 3.0 Example.....	1-3
1.5 New and Changed Features in this Release.....	1-4
2 Understanding Enterprise JavaBeans	
2.1 New Features and Changes in EJB	2-1
2.1.1 What Is New and Changed in EJB 3.1	2-1
2.1.2 What Was New and Changed in EJB 3.0.....	2-3
2.2 Understanding EJB Components.....	2-4
2.2.1 Session EJBs Implement Business Logic.....	2-4
2.2.1.1 Stateful Session Beans	2-4
2.2.1.2 Stateless Session Beans	2-4
2.2.1.3 Singleton Session Beans.....	2-5
2.2.2 Message-Driven Beans Implement Loosely Coupled Business Logic	2-5
2.3 EJB Anatomy and Environment	2-5
2.3.1 EJB Components	2-6
2.3.2 The EJB Container.....	2-6
2.3.3 EJB Metadata Annotations	2-6
2.3.4 Optional EJB Deployment Descriptors	2-7
2.4 EJB Clients and Communications.....	2-7
2.4.1 Accessing EJBs.....	2-7
2.4.2 EJB Communications.....	2-8
2.5 Securing EJBs	2-8

3 Simple Enterprise JavaBeans Examples

3.1	Simple Java Examples of 3.x EJBs.....	3-1
3.1.1	Example of a Simple No-interface Stateless EJB	3-1
3.1.2	Example of a Simple Business Interface Stateless EJB.....	3-2
3.1.3	Example of a Simple Stateful EJB	3-3
3.1.4	Example of an Interceptor Class.....	3-5
3.2	Packaged EJB 3.1 Examples in WebLogic Server	3-6
3.2.1	EJB 3.1: Example of a Singleton Session Bean.....	3-6
3.2.2	EJB 3.1: Example of an Asynchronous Method EJB.....	3-6
3.2.3	EJB 3.1: Example of a Calendar-based Timer EJB	3-7
3.2.4	EJB 3.1: Example of Simplified No-interface Programming and Packaging in a WAR File 3-7	
3.2.5	EJB 3.1: Example of Using a Portable Global JNDI Name in an EJB.....	3-7
3.2.6	EJB 3.0: Example of Invoking an Entity From A Session Bean	3-8

4 Iterative Development of Enterprise JavaBeans

4.1	Overview of the EJB Development Process	4-1
4.2	Create a Source Directory	4-2
4.2.1	Directory Structure for Packaging a JAR	4-3
4.2.2	Directory Structure for Packaging a WAR.....	4-3
4.3	Program the Annotated EJB Class.....	4-3
4.4	Program the EJB Interface.....	4-4
4.4.1	Accessing EJBs Using the No-Interface Client View	4-4
4.4.2	Accessing EJBs Using the Business Interface.....	4-4
4.4.2.1	Business Interface Application Exceptions.....	4-5
4.4.2.2	Using Generics in EJBs.....	4-5
4.4.2.3	Serializing and Deserializing Business Objects.....	4-6
4.5	Optionally Program Interceptors.....	4-6
4.6	Optionally Program the EJB Timer Service.....	4-7
4.6.1	Overview of the Timer Service	4-7
4.6.2	Calendar-based EJB Timers.....	4-7
4.6.3	Automatically-created EJB Timers	4-8
4.6.4	Non-persistent Timers	4-8
4.6.5	Clustered Versus Local EJB Timer Services.....	4-9
4.6.5.1	Clustered EJB Timer Services.....	4-9
4.6.5.2	Local EJB Timer Services	4-9
4.6.6	Configuring Clustered EJB Timers.....	4-9
4.6.7	Using Java Programming Interfaces to Program Timer Objects.....	4-10
4.6.7.1	EJB 3.1 Timer-related Programming Interfaces.....	4-10
4.6.7.2	WebLogic Server-specific Timer-related Programming Interfaces	4-11
4.7	Programming Access to EJB Clients.....	4-13
4.7.1	Remote Clients	4-13
4.7.2	Local Clients	4-13
4.7.3	Looking Up EJBs From Clients	4-14
4.7.3.1	Using Dependency Injection.....	4-14
4.7.3.2	Using the JNDI Portable Syntax	4-14
4.7.3.3	Customizing JNDI Names.....	4-15

4.7.4	Configuring EJBs to Send Requests to a URL.....	4-15
4.7.5	Specifying an HTTP Resource by URL.....	4-15
4.7.6	Specifying an HTTP Resource by Its JNDI Name.....	4-16
4.7.7	Accessing HTTP Resources from Bean Code	4-16
4.7.8	Configuring Network Communications for an EJB.....	4-16
4.8	Programming and Configuring Transactions.....	4-17
4.8.1	Programming Container-Managed Transactions	4-17
4.8.2	Configuring Automatic Retry of Container-Managed Transactions	4-18
4.8.3	Programming Bean-Managed Transactions	4-19
4.8.4	Programming Transactions That Are Distributed Across EJBs	4-20
4.8.4.1	Calling multiple EJBs from a client's transaction context.....	4-20
4.8.4.2	Using an EJB "Wrapper" to Encapsulate a Cross-EJB Transaction.....	4-20
4.9	Compile Java Source	4-21
4.10	Optionally Create and Edit Deployment Descriptors	4-21
4.11	Packaging EJBs	4-22
4.11.1	Packaging EJBs in a JAR	4-22
4.11.2	Packaging an EJB In a WAR.....	4-22
4.12	Deploying EJBs.....	4-23

5 Programming the Annotated EJB Class

5.1	Overview of Metadata Annotations and EJB Bean Files.....	5-1
5.2	Programming the Bean File: Requirements and Changes From 2.x.....	5-2
5.2.1	Bean Class Requirements and Changes From 2.x.....	5-2
5.2.2	Bean Class Method Requirements.....	5-3
5.3	Programming the Bean File.....	5-3
5.3.1	Typical Steps When Programming the Bean File.....	5-4
5.3.2	Specifying the Business and Other Interfaces	5-5
5.3.2.1	Specifying the Business Interface	5-5
5.3.2.2	Specifying the No-interface View	5-5
5.3.3	Specifying the Bean Type (Stateless, Singleton, Stateful, or Message-Driven).....	5-6
5.3.4	Injecting Resource Dependency into a Variable or Setter Method.....	5-7
5.3.5	Invoking a 3.0 Entity	5-8
5.3.5.1	Injecting Persistence Context Using Metadata Annotations.....	5-8
5.3.5.2	Finding an Entity Using the EntityManager API.....	5-9
5.3.5.3	Creating and Updating an Entity Using EntityManager.....	5-10
5.3.6	Specifying Interceptors for Business Methods or Life Cycle Callback Events	5-11
5.3.6.1	Specifying Business or Life Cycle Interceptors: Typical Steps.....	5-12
5.3.6.2	Programming the Interceptor Class.....	5-13
5.3.6.3	Programming Business Method Interceptor Methods.....	5-13
5.3.6.4	Programming Asynchronous Business Methods	5-13
5.3.6.5	Programming Life Cycle Callback Interceptor Methods.....	5-14
5.3.6.6	Specifying Default Interceptor Methods.....	5-15
5.3.6.7	Saving State Across Interceptors With the InvocationContext API.....	5-16
5.3.7	Programming Application Exceptions	5-16
5.3.8	Securing Access to the EJB.....	5-17
5.3.9	Specifying Transaction Management and Attributes.....	5-19
5.4	Complete List of Metadata Annotations By Function.....	5-19

5.4.1	Annotations to Specify the Bean Type.....	5-19
5.4.2	Annotations to Specify the Local or Remote Interfaces	5-20
5.4.3	Annotations to Support EJB 2.x Client View	5-20
5.4.4	Annotations to Invoke a 3.0 Entity Bean	5-20
5.4.5	Transaction-Related Annotations.....	5-21
5.4.6	Annotations to Specify Interceptors.....	5-21
5.4.7	Annotations to Specify Life Cycle Callbacks	5-22
5.4.8	Security-Related Annotations	5-22
5.4.9	Context Dependency Annotations	5-23
5.4.10	Timeout and Exceptions Annotations	5-23
5.4.11	Timer and Scheduling Annotations	5-23

6 Deployment Guidelines for Enterprise JavaBeans

6.1	Before You Deploy an EJB	6-1
6.2	Understanding and Performing Deployment Tasks	6-2
6.3	Deployment Guidelines for EJBs	6-2
6.3.1	Deploying Standalone EJBs as Part of an Enterprise Application.....	6-2
6.3.2	Deploying EJBs as Part of an Web Application.....	6-3
6.3.3	Deploying EJBs That Call Each Other in the Same Application.....	6-3
6.3.4	Deploying EJBs That Use Dependency Injection	6-3
6.3.5	Deploying Homogeneously to a Cluster.....	6-3
6.3.6	Deploying Pinned EJBs to a Cluster.....	6-4
6.3.7	Redeploying an EJB	6-4
6.3.8	Using FastSwap Deployment to Minimize Deployment.....	6-4
6.3.9	Understanding Warning Messages	6-5
6.3.10	Disabling EJB Deployment Warning Messages	6-5

7 Using an Embedded EJB Container in Oracle WebLogic Server

7.1	Overview of the Embeddable EJB Container.....	7-1
7.2	EJB 3.1 Lite Functionality Supported in the Embedded EJB Container.....	7-1

8 Configuring the Persistence Provider in Oracle WebLogic Server

8.1	Overview of Oracle TopLink.....	8-1
8.2	Specifying a Persistence Provider.....	8-2
8.2.1	Setting the Default Provider for the Domain.....	8-2
8.2.2	Specifying the Persistence Provider in an Application	8-3
8.3	Using Oracle TopLink in Oracle WebLogic Server.....	8-3
8.4	Using Oracle Kodo in Oracle WebLogic Server	8-3
8.4.1	Using JPA 1.0 APIs	8-4
8.4.2	Using persistence-configuration.xml	8-4
8.4.3	Updating Applications to Overcome Conflicts with JPA 2.0.....	8-4
8.5	Using a Newer Version of OpenJPA in Oracle WebLogic Server.....	8-5

A EJB Metadata Annotations Reference

A.1	Overview of EJB 3.x Annotations	A-1
A.2	Annotations for Stateless, Stateful, and Message-Driven Beans.....	A-1

A.2.1	javax.ejb.AccessTimeout	A-2
A.2.1.1	Description	A-2
A.2.1.2	Attributes	A-3
A.2.2	javax.ejb.ActivationConfigProperty	A-3
A.2.2.1	Description	A-3
A.2.2.2	Attributes	A-4
A.2.3	javax.ejb.AfterBegin.....	A-4
A.2.3.1	Description	A-4
A.2.4	javax.ejb.AfterCompletion.....	A-4
A.2.4.1	Description	A-5
A.2.5	javax.ejb.ApplicationException	A-5
A.2.5.1	Description	A-5
A.2.5.2	Attributes	A-5
A.2.6	javax.ejb.Asynchronous	A-5
A.2.6.1	Description	A-6
A.2.7	javax.ejb.BeforeCompletion.....	A-6
A.2.7.1	Description	A-6
A.2.8	javax.ejb.ConcurrencyManagement.....	A-6
A.2.8.1	Description	A-6
A.2.8.2	Attributes	A-7
A.2.9	javax.ejb.DependsOn.....	A-7
A.2.9.1	Description	A-7
A.2.9.2	Attributes	A-7
A.2.10	javax.ejb.EJB.....	A-7
A.2.10.1	Description	A-8
A.2.10.2	Attributes	A-8
A.2.11	javax.ejb.EJBs	A-9
A.2.11.1	Description	A-9
A.2.11.2	Attribute.....	A-9
A.2.12	javax.ejb.Init.....	A-9
A.2.12.1	Description	A-9
A.2.12.2	Attributes	A-10
A.2.13	javax.ejb.Local	A-10
A.2.13.1	Description	A-10
A.2.13.2	Attributes	A-10
A.2.14	javax.ejb.LocalBean.....	A-10
A.2.14.1	Description	A-11
A.2.15	javax.ejb.LocalHome	A-11
A.2.15.1	Description	A-11
A.2.15.2	Attributes	A-11
A.2.16	javax.ejb.Lock	A-11
A.2.16.1	Description	A-11
A.2.16.2	Attributes	A-11
A.2.17	javax.ejb.MessageDriven	A-12
A.2.17.1	Description	A-12
A.2.17.2	Attributes	A-12
A.2.18	javax.ejb.PostActivate.....	A-13

A.2.18.1	Description	A-13
A.2.19	javax.ejb.PrePassivate.....	A-14
A.2.19.1	Description	A-14
A.2.20	javax.ejb.Remote	A-14
A.2.20.1	Description	A-14
A.2.20.2	Attributes	A-14
A.2.21	javax.ejb.RemoteHome	A-15
A.2.21.1	Description	A-15
A.2.21.2	Attributes	A-15
A.2.22	javax.ejb.Remove	A-15
A.2.22.1	Description	A-15
A.2.22.2	Attributes	A-15
A.2.23	javax.ejb.Schedule.....	A-16
A.2.23.1	Description	A-16
A.2.23.1.1	Calendar-based Schedule Elements.....	A-16
A.2.23.1.2	Forms of Supported Element Values.....	A-17
A.2.23.1.3	Additional Rules for Schedule Specification Elements.....	A-18
A.2.23.2	Attributes	A-18
A.2.24	javax.ejb.Schedules	A-19
A.2.24.1	Description	A-19
A.2.24.2	Attributes	A-19
A.2.25	javax.ejb.Singleton	A-20
A.2.25.1	Description	A-20
A.2.25.2	Attributes	A-20
A.2.26	javax.ejb.Startup.....	A-20
A.2.26.1	Description	A-21
A.2.27	javax.ejb.StatefulTimeout	A-21
A.2.27.1	Description	A-21
A.2.27.2	Attributes	A-21
A.2.28	javax.ejb.Stateless.....	A-21
A.2.28.1	Description	A-21
A.2.28.2	Attributes	A-21
A.2.29	javax.ejb.Timeout.....	A-22
A.2.29.1	Description	A-22
A.2.30	javax.ejb.TransactionAttribute.....	A-23
A.2.30.1	Description	A-23
A.2.30.2	Attributes	A-23
A.2.31	javax.ejb.TransactionManagement.....	A-24
A.2.31.1	Description	A-24
A.2.31.2	Attributes	A-24
A.3	Annotations Used to Configure Interceptors.....	A-24
A.3.1	javax.interceptor.AroundInvoke	A-24
A.3.1.1	Description	A-24
A.3.2	javax.interceptor.ExcludeClassInterceptors.....	A-25
A.3.2.1	Description	A-25
A.3.3	javax.interceptor.ExcludeDefaultInterceptors.....	A-25
A.3.3.1	Description	A-25

A.3.4	javax.interceptor.Interceptors	A-25
A.3.4.1	Description	A-25
A.3.4.2	Attributes	A-26
A.4	Annotations Used to Interact With Entity Beans	A-26
A.4.1	javax.persistence.PersistenceContext.....	A-26
A.4.1.1	Description	A-26
A.4.1.2	Attributes	A-26
A.4.2	javax.persistence.PersistenceContexts	A-27
A.4.2.1	Description	A-27
A.4.2.2	Attributes	A-28
A.4.3	javax.persistence.PersistenceUnit.....	A-28
A.4.3.1	Description	A-28
A.4.3.2	Attributes	A-28
A.4.4	javax.persistence.PersistenceUnits	A-29
A.4.4.1	Description	A-29
A.4.4.2	Attributes	A-29
A.5	Standard JDK Annotations Used By EJB 3.x.....	A-29
A.5.1	javax.annotation.PostConstruct.....	A-29
A.5.1.1	Description	A-30
A.5.2	javax.annotation.PreDestroy	A-30
A.5.2.1	Description	A-30
A.5.3	javax.annotation.Resource.....	A-30
A.5.3.1	Description	A-30
A.5.3.2	Attributes	A-31
A.5.4	javax.annotation.Resources	A-32
A.5.4.1	Description	A-32
A.5.4.2	Attributes	A-32
A.6	Standard Security-Related JDK Annotations Used by EJB 3.x.....	A-32
A.6.1	javax.annotation.security.DeclareRoles.....	A-32
A.6.1.1	Description	A-32
A.6.1.2	Attributes	A-33
A.6.2	javax.annotation.security.DenyAll.....	A-33
A.6.2.1	Description	A-33
A.6.3	javax.annotation.security.PermitAll	A-33
A.6.3.1	Description	A-33
A.6.4	javax.annotation.security.RolesAllowed.....	A-33
A.6.4.1	Description	A-33
A.6.4.2	Attributes	A-34
A.6.5	javax.annotation.security.RunAs.....	A-34
A.6.5.1	Description	A-34
A.6.5.2	Attributes	A-34
A.7	WebLogic Annotations.....	A-34
A.7.1	weblogic.javaee.AllowRemoveDuringTransaction	A-35
A.7.1.1	Description	A-35
A.7.2	weblogic.javaee.CallByReference	A-35
A.7.2.1	Description	A-35
A.7.3	weblogic.javaee.DisableWarnings.....	A-35

A.7.3.1	Description	A-35
A.7.3.2	Attributes	A-36
A.7.4	weblogic.javaee.EJBReference.....	A-36
A.7.4.1	Description	A-36
A.7.4.2	Attribute.....	A-36
A.7.5	weblogic.javaee.Idempotent.....	A-37
A.7.5.1	Description	A-37
A.7.5.2	Attributes	A-37
A.7.6	weblogic.javaee.JMSClientID.....	A-37
A.7.6.1	Description	A-37
A.7.6.2	Attributes	A-38
A.7.7	weblogic.javaee.JNDIName	A-38
A.7.7.1	Description	A-38
A.7.7.2	Attributes	A-38
A.7.8	weblogic.javaee.JNDINames.....	A-38
A.7.8.1	Description	A-38
A.7.8.2	Attributes	A-39
A.7.9	weblogic.javaee.MessageDestinationConfiguration.....	A-39
A.7.9.1	Description	A-39
A.7.9.2	Attributes	A-39
A.7.10	weblogic.javaee.TransactionIsolation	A-39
A.7.10.1	Description	A-39
A.7.10.2	Attributes	A-40
A.7.11	weblogic.javaee.TransactionTimeoutSeconds	A-40
A.7.11.1	Description	A-40
A.7.11.2	Attributes	A-40

B Using Oracle Kodo with Oracle WebLogic Server

B.1	Overview of Oracle Kodo	B-2
B.2	Creating an Oracle Kodo Application	B-2
B.3	Using Different Oracle Kodo Versions	B-2
B.4	Configuring Persistence	B-2
B.4.1	Editing the Configuration Property Files.....	B-3
B.4.2	Using the Configuration Files Together	B-3
B.4.3	Configuring Plug-ins.....	B-4
B.5	Deploying an Oracle Kodo Application	B-4
B.6	Configuring an Oracle Kodo Application	B-4
B.6.1	Using the Administration Console.....	B-5
B.6.2	Configuring Oracle Kodo Without Using the Administration Console.....	B-5

C Oracle Kodo Persistence Configuration Schema Reference

C.1	persistence-configuration.xml Namespace Declaration and Schema Location.....	C-1
C.2	persistence-configuration.xml Deployment Descriptor File Structure	C-2
C.3	persistence-configuration.xml Deployment Descriptor Elements.....	C-5
C.4	abstract-store-broker-factory	C-5
C.4.1	Function.....	C-5
C.4.2	Example.....	C-5

C.5	access-dictionary	C-5
C.5.1	Function.....	C-6
C.5.2	Example.....	C-6
C.6	access-unloaded	C-8
C.6.1	Function.....	C-9
C.6.2	Example.....	C-9
C.7	action.....	C-9
C.7.1	Function.....	C-9
C.7.2	Example.....	C-9
C.8	addresses	C-9
C.8.1	Function.....	C-10
C.8.2	Example.....	C-10
C.9	advanced-sql.....	C-10
C.9.1	Function.....	C-10
C.10	aggregate-listeners	C-10
C.10.1	Function.....	C-10
C.10.2	Example.....	C-10
C.11	allocate	C-11
C.11.1	Function.....	C-11
C.11.2	Example.....	C-11
C.12	assert-allowed-type.....	C-11
C.12.1	Function.....	C-11
C.12.2	Example.....	C-11
C.13	auto-clear	C-12
C.13.1	Function.....	C-12
C.13.2	Example.....	C-12
C.14	auto-detach	C-12
C.14.1	Function.....	C-12
C.14.2	Example.....	C-13
C.15	auto-detaches	C-13
C.15.1	Function.....	C-13
C.15.2	Example.....	C-13
C.16	base-name.....	C-13
C.16.1	Function.....	C-13
C.16.2	Example.....	C-13
C.17	batching-operation-order-update-manager	C-13
C.17.1	Function.....	C-14
C.17.2	Example.....	C-14
C.18	buffer-size.....	C-14
C.18.1	Function.....	C-14
C.18.2	Example.....	C-14
C.19	cache-map	C-14
C.19.1	Function.....	C-14
C.19.2	Example.....	C-15
C.20	cache-size.....	C-15
C.20.1	Function.....	C-15
C.20.2	Example.....	C-15

C.21	channel.....	C-15
C.21.1	Function.....	C-15
C.21.2	Example.....	C-16
C.22	class-table-jdbc-seq	C-16
C.22.1	Function.....	C-16
C.22.2	Example.....	C-16
C.23	classname	C-16
C.23.1	Function.....	C-17
C.24	classpath-scan	C-17
C.24.1	Function.....	C-18
C.24.2	Example.....	C-18
C.25	clear-on-close	C-18
C.25.1	Function.....	C-18
C.25.2	Example.....	C-18
C.26	client-broker-factory	C-18
C.26.1	Function.....	C-19
C.26.2	Example.....	C-19
C.27	close-on-managed-commit	C-19
C.27.1	Function.....	C-19
C.27.2	Example.....	C-19
C.28	cluster-remote-commit-provider	C-19
C.28.1	Function.....	C-19
C.28.2	Example.....	C-20
C.29	commons-log-factory.....	C-20
C.29.1	Function.....	C-20
C.29.2	Example.....	C-20
C.30	compatibility	C-20
C.30.1	Function.....	C-20
C.30.2	Example.....	C-20
C.31	concurrent-hash-map	C-21
C.31.1	Function.....	C-21
C.31.2	Example.....	C-21
C.32	connection-decorators	C-21
C.32.1	Function.....	C-21
C.32.2	Example.....	C-22
C.33	connection-driver-name.....	C-22
C.33.1	Function.....	C-22
C.33.2	Example.....	C-23
C.34	connection-factory-mode	C-23
C.34.1	Function.....	C-23
C.34.2	Example.....	C-23
C.35	connection-factory-name	C-23
C.35.1	Function.....	C-24
C.35.2	Example.....	C-24
C.36	connection-factory-properties.....	C-24
C.36.1	Function.....	C-24
C.36.2	Example.....	C-24

C.37	connection-factory2-name	C-24
C.37.1	Function.....	C-25
C.37.2	Example.....	C-25
C.38	connection-factory2-properties	C-25
C.38.1	Function.....	C-25
C.38.2	Example.....	C-25
C.39	connection-password.....	C-25
C.39.1	Function.....	C-25
C.39.2	Example.....	C-25
C.40	connection-properties.....	C-26
C.40.1	Function.....	C-26
C.40.2	Example.....	C-26
C.41	connection-retain-mode	C-26
C.41.1	Function.....	C-26
C.41.2	Example.....	C-27
C.42	connection-url.....	C-27
C.42.1	Function.....	C-27
C.42.2	Example.....	C-28
C.43	connection-user-name	C-28
C.43.1	Function.....	C-28
C.43.2	Example.....	C-28
C.44	connection2-driver-name.....	C-29
C.44.1	Function.....	C-29
C.44.2	Example.....	C-29
C.45	connection2-password.....	C-29
C.45.1	Function.....	C-29
C.45.2	Example.....	C-29
C.46	connection2-properties.....	C-29
C.46.1	Function.....	C-29
C.46.2	Example.....	C-30
C.47	connection2-url.....	C-30
C.47.1	Function.....	C-30
C.47.2	Example.....	C-30
C.48	connection2-user-name	C-30
C.48.1	Function.....	C-30
C.48.2	Example.....	C-31
C.49	constraint-names	C-31
C.49.1	Function.....	C-31
C.49.2	Example.....	C-31
C.50	constraint-update-manager	C-31
C.50.1	Function.....	C-31
C.50.2	Example.....	C-31
C.51	copy-object-ids.....	C-31
C.51.1	Function.....	C-32
C.51.2	Example.....	C-32
C.52	custom-aggregate-listener	C-32
C.52.1	Function.....	C-32

C.52.2	Example.....	C-32
C.53	custom-broker-factory.....	C-32
C.53.1	Function.....	C-33
C.53.2	Example.....	C-33
C.54	custom-broker-impl.....	C-33
C.54.1	Function.....	C-33
C.54.2	Example.....	C-33
C.55	custom-class-resolver.....	C-33
C.55.1	Function.....	C-33
C.55.2	Example.....	C-33
C.56	custom-compatibility.....	C-34
C.56.1	Function.....	C-34
C.56.2	Example.....	C-34
C.57	custom-connection-decorator.....	C-34
C.57.1	Function.....	C-34
C.57.2	Example.....	C-34
C.58	custom-data-cache.....	C-34
C.58.1	Function.....	C-35
C.58.2	Example.....	C-35
C.59	custom-data-cache-manager.....	C-35
C.59.1	Function.....	C-35
C.59.2	Example.....	C-35
C.60	custom-detach-state.....	C-35
C.60.1	Function.....	C-35
C.60.2	Example.....	C-36
C.61	custom-dictionary.....	C-36
C.61.1	Function.....	C-36
C.61.2	Example.....	C-36
C.62	custom-driver-data-source.....	C-36
C.62.1	Function.....	C-36
C.62.2	Example.....	C-36
C.63	custom-filter-listener.....	C-36
C.63.1	Function.....	C-37
C.63.2	Example.....	C-37
C.64	custom-jdbc-listener.....	C-37
C.64.1	Function.....	C-37
C.64.2	Example.....	C-37
C.65	custom-lock-manager.....	C-37
C.65.1	Function.....	C-38
C.65.2	Example.....	C-38
C.66	custom-log.....	C-38
C.66.1	Function.....	C-38
C.66.2	Example.....	C-38
C.67	custom-mapping-defaults.....	C-38
C.67.1	Function.....	C-38
C.67.2	Example.....	C-38
C.68	custom-mapping-factory.....	C-39

C.68.1	Function.....	C-39
C.68.2	Example.....	C-39
C.69	custom-meta-data-factory.....	C-39
C.69.1	Function.....	C-39
C.69.2	Example.....	C-39
C.70	custom-meta-data-repository.....	C-39
C.70.1	Function.....	C-39
C.70.2	Example.....	C-40
C.71	custom-orphaned-key-action	C-40
C.71.1	Function.....	C-40
C.71.2	Example.....	C-40
C.72	custom-persistence-server	C-40
C.72.1	Function.....	C-40
C.72.2	Example.....	C-40
C.73	custom-proxy-manager.....	C-41
C.73.1	Function.....	C-41
C.73.2	Example.....	C-41
C.74	custom-query-compilation-cache	C-41
C.74.1	Function.....	C-41
C.74.2	Example.....	C-41
C.75	custom-remote-commit-provider	C-41
C.75.1	Function.....	C-42
C.75.2	Example.....	C-42
C.76	custom-savepoint-manager	C-42
C.76.1	Function.....	C-42
C.76.2	Example.....	C-42
C.77	custom-schema-factory	C-42
C.77.1	Function.....	C-42
C.77.2	Example.....	C-42
C.78	custom-seq	C-43
C.78.1	Function.....	C-43
C.78.2	Example.....	C-43
C.79	custom-sql-factory	C-43
C.79.1	Function.....	C-43
C.79.2	Example.....	C-43
C.80	custom-update-manager	C-43
C.80.1	Function.....	C-44
C.80.2	Example.....	C-44
C.81	data-caches.....	C-44
C.81.1	Function.....	C-44
C.81.2	Example.....	C-44
C.82	data-cache-manager-impl.....	C-44
C.82.1	Function.....	C-44
C.82.2	Example.....	C-45
C.83	data-cache-timeout	C-45
C.83.1	Function.....	C-45
C.83.2	Example.....	C-45

C.84	db2-dictionary	C-45
C.84.1	Function.....	C-45
C.84.2	Example.....	C-45
C.85	default-access-type.....	C-45
C.85.1	Function.....	C-46
C.85.2	Example.....	C-46
C.86	default-broker-factory	C-46
C.86.1	Function.....	C-46
C.86.2	Example.....	C-46
C.87	default-broker-impl	C-46
C.87.1	Function.....	C-46
C.87.2	Example.....	C-47
C.88	default-class-resolver.....	C-47
C.88.1	Function.....	C-47
C.88.2	Example.....	C-47
C.89	default-compatibility	C-47
C.89.1	Function.....	C-47
C.89.2	Example.....	C-47
C.90	default-data-cache.....	C-47
C.90.1	Function.....	C-48
C.90.2	Example.....	C-48
C.91	default-detach-state	C-48
C.91.1	Function.....	C-48
C.91.2	Example.....	C-48
C.92	default-data-cache-manager	C-48
C.92.1	Function.....	C-48
C.92.2	Example.....	C-49
C.93	default-driver-data-source.....	C-49
C.93.1	Function.....	C-49
C.93.2	Example.....	C-49
C.94	default-level	C-49
C.94.1	Function.....	C-49
C.94.2	Example.....	C-49
C.95	default-lock-manager	C-49
C.95.1	Function.....	C-50
C.95.2	Example.....	C-50
C.96	default-mapping-defaults	C-50
C.96.1	Function.....	C-50
C.96.2	Example.....	C-50
C.97	default-meta-data-factory	C-50
C.97.1	Function.....	C-50
C.97.2	Example.....	C-50
C.98	default-meta-data-repository	C-51
C.98.1	Function.....	C-51
C.98.2	Example.....	C-51
C.99	default-orphaned-key-action.....	C-51
C.99.1	Function.....	C-51

C.99.2	Example.....	C-51
C.100	default-proxy-manager	C-51
C.100.1	Function.....	C-52
C.100.2	Example.....	C-52
C.101	default-query-compilation-cache.....	C-52
C.101.1	Function.....	C-52
C.101.2	Example.....	C-52
C.102	default-savepoint-manager	C-52
C.102.1	Function.....	C-52
C.102.2	Example.....	C-53
C.103	default-schema-factory.....	C-53
C.103.1	Function.....	C-53
C.103.2	Example.....	C-53
C.104	default-sql-factory.....	C-53
C.104.1	Function.....	C-53
C.104.2	Example.....	C-53
C.105	default-update-manager	C-54
C.105.1	Function.....	C-54
C.105.2	Example.....	C-54
C.106	deprecated-jdo-mapping-defaults.....	C-54
C.106.1	Function.....	C-54
C.106.2	Example.....	C-54
C.107	deprecated-jdo-meta-data-factory	C-54
C.107.1	Function.....	C-54
C.107.2	Example.....	C-55
C.108	derby-dictionary.....	C-55
C.108.1	Function.....	C-55
C.108.2	Example.....	C-55
C.109	detach-options-all	C-55
C.109.1	Function.....	C-56
C.109.2	Example.....	C-56
C.110	detach-options-fetch-groups	C-56
C.110.1	Function.....	C-56
C.110.2	Example.....	C-56
C.111	detach-options-loaded.....	C-56
C.111.1	Function.....	C-57
C.111.2	Example.....	C-57
C.112	detach-state	C-57
C.112.1	Function.....	C-57
C.112.2	Example.....	C-57
C.113	detached-state-field	C-57
C.113.1	Function.....	C-58
C.113.2	Example.....	C-58
C.114	detached-state-manager.....	C-58
C.114.1	Function.....	C-58
C.114.2	Example.....	C-59
C.115	detached-state-transient.....	C-59

C.115.1	Function.....	C-59
C.115.2	Example.....	C-59
C.116	detached-new	C-59
C.116.1	Function.....	C-60
C.116.2	Example.....	C-60
C.117	diagnostic-context.....	C-60
C.117.1	Function.....	C-60
C.117.2	Example.....	C-60
C.118	dynamic-data-structs.....	C-60
C.118.1	Function.....	C-60
C.118.2	Example.....	C-61
C.119	dynamic-schema-factory.....	C-61
C.119.1	Function.....	C-61
C.119.2	Example.....	C-61
C.120	eager-fetch-mode	C-61
C.120.1	Function.....	C-61
C.120.2	Example.....	C-63
C.121	empress-dictionary	C-63
C.121.1	Function.....	C-63
C.121.2	Example.....	C-63
C.122	EnableLogMBean	C-63
C.122.1	Function.....	C-64
C.122.2	Example.....	C-64
C.123	EnableRuntimeMBean.....	C-64
C.123.1	Function.....	C-64
C.123.2	Example.....	C-64
C.124	evict-from-data-cache	C-64
C.124.1	Function.....	C-64
C.124.2	Example.....	C-64
C.125	eviction-schedule	C-64
C.125.1	Function.....	C-65
C.125.2	Example.....	C-65
C.126	exception-orphaned-key-action	C-65
C.126.1	Function.....	C-66
C.126.2	Example.....	C-66
C.127	exception-reconnect-attempts	C-66
C.127.1	Function.....	C-66
C.127.2	Example.....	C-66
C.128	execution-context-name-provider	C-66
C.128.1	Function.....	C-66
C.128.2	Example.....	C-66
C.129	export-profiling	C-67
C.129.1	Function.....	C-67
C.129.2	Example.....	C-67
C.130	extension-deprecated-jdo-mapping-factory.....	C-67
C.130.1	Function.....	C-67
C.130.2	Example.....	C-67

C.131	fetch-batch-size.....	C-68
C.131.1	Function.....	C-68
C.131.2	Example.....	C-68
C.132	fetch-direction.....	C-68
C.132.1	Function.....	C-68
C.132.2	Example.....	C-68
C.133	fetch-group.....	C-69
C.133.1	Function.....	C-69
C.133.2	Example.....	C-69
C.134	fetch-groups.....	C-69
C.134.1	Function.....	C-69
C.134.2	Example.....	C-69
C.135	field-override.....	C-70
C.135.1	Function.....	C-70
C.135.2	Example.....	C-70
C.136	file.....	C-70
C.136.1	Function.....	C-70
C.136.2	Example.....	C-70
C.137	file-name.....	C-71
C.137.1	Function.....	C-71
C.137.2	Example.....	C-71
C.138	file-schema-factory.....	C-71
C.138.1	Function.....	C-71
C.138.2	Example.....	C-71
C.139	files.....	C-72
C.139.1	Function.....	C-72
C.139.2	Example.....	C-72
C.140	filter-listeners.....	C-72
C.140.1	Function.....	C-73
C.140.2	Example.....	C-73
C.141	foreign-keys.....	C-73
C.141.1	Function.....	C-73
C.141.2	Example.....	C-73
C.142	format.....	C-73
C.142.1	Function.....	C-74
C.143	foxpro-dictionary.....	C-74
C.143.1	Function.....	C-74
C.143.2	Example.....	C-74
C.144	flush-before-queries.....	C-74
C.144.1	Function.....	C-74
C.144.2	Example.....	C-75
C.145	hsql-dictionary.....	C-75
C.145.1	Function.....	C-75
C.145.2	Example.....	C-75
C.146	gem-fire-data-cache.....	C-75
C.146.1	Function.....	C-75
C.146.2	Example.....	C-75

C.147	gem-fire-data-cache-name	C-75
C.147.1	Function.....	C-76
C.147.2	Example.....	C-76
C.148	gui-jmx.....	C-76
C.148.1	Function.....	C-76
C.148.2	Example.....	C-76
C.149	gui-profiling.....	C-76
C.149.1	Function.....	C-77
C.149.2	Example.....	C-77
C.150	Host.....	C-77
C.150.1	Function.....	C-77
C.150.2	Example.....	C-77
C.151	host.....	C-77
C.151.1	Function.....	C-77
C.151.2	Example.....	C-77
C.152	http-transport	C-78
C.152.1	Function.....	C-78
C.152.2	Example.....	C-78
C.153	ignore-changes	C-78
C.153.1	Function.....	C-78
C.153.2	Example.....	C-78
C.154	ignore-unmapped	C-78
C.154.1	Function.....	C-79
C.154.2	Example.....	C-79
C.155	ignore-virtual.....	C-79
C.155.1	Function.....	C-79
C.155.2	Example.....	C-79
C.156	in-memory-savepoint-manager	C-79
C.156.1	Function.....	C-79
C.156.2	Example.....	C-79
C.157	increment.....	C-80
C.157.1	Function.....	C-80
C.157.2	Example.....	C-80
C.158	indexes.....	C-80
C.158.1	Function.....	C-80
C.158.2	Example.....	C-80
C.159	informix-dictionary.....	C-81
C.159.1	Function.....	C-81
C.159.2	Example.....	C-81
C.160	initial-value	C-81
C.160.1	Function.....	C-81
C.160.2	Example.....	C-81
C.161	interval-millis.....	C-81
C.161.1	Function.....	C-82
C.161.2	Example.....	C-82
C.162	inverse-manager.....	C-82
C.162.1	Function.....	C-82

C.162.2	Example.....	C-82
C.163	jdatastore-dictionary.....	C-82
C.163.1	Function.....	C-82
C.163.2	Example.....	C-83
C.164	jdbcbroker-factory	C-83
C.164.1	Function.....	C-83
C.164.2	Example.....	C-83
C.165	jdbclisteners.....	C-83
C.165.1	Function.....	C-83
C.165.2	Example.....	C-83
C.166	jdbcsavepoint-manager.....	C-83
C.166.1	Function.....	C-84
C.166.2	Example.....	C-84
C.167	jdo-meta-data-factory	C-84
C.167.1	Function.....	C-84
C.167.2	Example.....	C-84
C.168	jms-remote-commit-provider.....	C-84
C.168.1	Function.....	C-85
C.168.2	Example.....	C-85
C.169	jmx.....	C-85
C.169.1	Function.....	C-85
C.169.2	Example.....	C-85
C.170	jmx2-jmx	C-85
C.170.1	Function.....	C-85
C.170.2	Example.....	C-86
C.171	JNDIName.....	C-86
C.171.1	Function.....	C-86
C.171.2	Example.....	C-86
C.172	kodobroker	C-86
C.172.1	Function.....	C-86
C.172.2	Example.....	C-86
C.173	kodo-concurrent-data-cache.....	C-87
C.173.1	Function.....	C-87
C.173.2	Example.....	C-87
C.174	kodo-data-cache-manager	C-87
C.174.1	Function.....	C-87
C.174.2	Example.....	C-88
C.175	kodo-mapping-repository	C-88
C.175.1	Function.....	C-88
C.175.2	Example.....	C-88
C.176	kodo-persistence-mapping-factory	C-88
C.176.1	Function.....	C-88
C.176.2	Example.....	C-88
C.177	kodo-persistence-meta-data-factory.....	C-89
C.177.1	Function.....	C-89
C.177.2	Example.....	C-89
C.178	kodo-pooling-data-source	C-89

C.178.1	Function.....	C-89
C.178.2	Example.....	C-89
C.179	kodo-sql-factory	C-90
C.179.1	Function.....	C-90
C.179.2	Example.....	C-90
C.180	large-transaction.....	C-90
C.180.1	Function.....	C-90
C.180.2	Example.....	C-90
C.181	lazy-schema-factory	C-90
C.181.1	Function.....	C-91
C.181.2	Example.....	C-91
C.182	level.....	C-91
C.182.1	Function.....	C-91
C.182.2	Example.....	C-91
C.183	local-jmx	C-91
C.183.1	Function.....	C-92
C.183.2	Example.....	C-92
C.184	local-profiling	C-92
C.184.1	Function.....	C-92
C.184.2	Example.....	C-92
C.185	lock-timeout.....	C-92
C.185.1	Function.....	C-93
C.185.2	Example.....	C-93
C.186	log-factory-impl.....	C-93
C.186.1	Function.....	C-93
C.186.2	Example.....	C-93
C.187	log-orphaned-key-action.....	C-93
C.187.1	Function.....	C-93
C.187.2	Example.....	C-93
C.188	log4j-log-factory	C-94
C.188.1	Function.....	C-94
C.188.2	Example.....	C-94
C.189	login-timeout	C-94
C.189.1	Function.....	C-94
C.189.2	Example.....	C-94
C.190	lrs-size	C-95
C.190.1	Function.....	C-95
C.190.2	Example.....	C-95
C.191	lru-data-cache	C-95
C.191.1	Function.....	C-95
C.191.2	Example.....	C-95
C.192	manage-lru	C-96
C.192.1	Function.....	C-96
C.192.2	Example.....	C-96
C.193	mapping	C-96
C.193.1	Function.....	C-96
C.193.2	Example.....	C-97

C.194	mapping-column.....	C-97
C.194.1	Function.....	C-97
C.194.2	Example.....	C-97
C.195	mapping-defaults-impl	C-97
C.195.1	Function.....	C-97
C.195.2	Example.....	C-97
C.196	mapping-file-deprecated-jdo-mapping-factory	C-98
C.196.1	Function.....	C-98
C.196.2	Example.....	C-98
C.197	max-active	C-98
C.197.1	Function.....	C-99
C.197.2	Example.....	C-99
C.198	max-idle.....	C-99
C.198.1	Function.....	C-99
C.198.2	Example.....	C-99
C.199	max-size.....	C-99
C.199.1	Function.....	C-99
C.199.2	Example.....	C-99
C.200	maximize-batch-size	C-100
C.200.1	Function.....	C-100
C.200.2	Example.....	C-100
C.201	MBeanServerStrategy	C-100
C.201.1	Function.....	C-100
C.201.2	Example.....	C-100
C.202	multithreaded	C-101
C.202.1	Function.....	C-101
C.202.2	Example.....	C-101
C.203	mx4j1-jmx	C-101
C.203.1	Function.....	C-101
C.203.2	Example.....	C-101
C.204	mysql-dictionary	C-101
C.204.1	Function.....	C-102
C.204.2	Example.....	C-102
C.205	name.....	C-102
C.205.1	Function.....	C-103
C.206	name-column	C-103
C.206.1	Function.....	C-103
C.206.2	Example.....	C-103
C.207	NamingImpl	C-103
C.207.1	Function.....	C-103
C.207.2	Example.....	C-104
C.208	native-jdbc-seq.....	C-104
C.208.1	Function.....	C-104
C.208.2	Example.....	C-104
C.209	none-jmx.....	C-104
C.209.1	Function.....	C-104
C.209.2	Example.....	C-105

C.210	none-lock-manager	C-105
C.210.1	Function.....	C-105
C.210.2	Example.....	C-105
C.211	none-log-factory	C-105
C.211.1	Function.....	C-105
C.211.2	Example.....	C-105
C.212	none-orphaned-key-action	C-105
C.212.1	Function.....	C-106
C.212.2	Example.....	C-106
C.213	none-profiling.....	C-106
C.213.1	Function.....	C-106
C.213.2	Example.....	C-106
C.214	nontransactional-read.....	C-106
C.214.1	Function.....	C-107
C.214.2	Example.....	C-107
C.215	nontransactional-write	C-107
C.215.1	Function.....	C-107
C.215.2	Example.....	C-107
C.216	num-broadcast-threads	C-107
C.216.1	Function.....	C-107
C.216.2	Example.....	C-108
C.217	operation-order-update-manager.....	C-108
C.217.1	Function.....	C-108
C.217.2	Example.....	C-108
C.218	optimistic.....	C-108
C.218.1	Function.....	C-108
C.218.2	Example.....	C-108
C.219	oracle-dictionary	C-108
C.219.1	Function.....	C-109
C.219.2	Example.....	C-109
C.220	oracle-savepoint-manager	C-109
C.220.1	Function.....	C-109
C.220.2	Example.....	C-109
C.221	orm-file-jdor-mapping-factory.....	C-109
C.221.1	Function.....	C-110
C.221.2	Example.....	C-110
C.222	order-dirty-objects	C-110
C.222.1	Function.....	C-110
C.222.2	Example.....	C-110
C.223	Password	C-110
C.223.1	Function.....	C-111
C.223.2	Example.....	C-111
C.224	persistence-configuration	C-111
C.224.1	Function.....	C-111
C.225	persistence-configuration-unit.....	C-111
C.225.1	Function.....	C-111
C.226	pessimistic-lock-manager	C-111

C.226.1	Function.....	C-111
C.226.2	Example.....	C-111
C.227	persistence-mapping-defaults.....	C-112
C.227.1	Function.....	C-112
C.227.2	Example.....	C-112
C.228	populate-data-cache	C-112
C.228.1	Function.....	C-112
C.228.2	Example.....	C-112
C.229	Port.....	C-112
C.229.1	Function.....	C-113
C.229.2	Example.....	C-113
C.230	port.....	C-113
C.230.1	Function.....	C-113
C.230.2	Example.....	C-113
C.231	postgres-dictionary.....	C-113
C.231.1	Function.....	C-114
C.231.2	Example.....	C-114
C.232	primary-key-column.....	C-114
C.232.1	Function.....	C-114
C.232.2	Example.....	C-115
C.233	primary-key-value	C-115
C.233.1	Function.....	C-115
C.233.2	Example.....	C-115
C.234	primary-keys.....	C-115
C.234.1	Function.....	C-115
C.234.2	Example.....	C-115
C.235	profiling.....	C-115
C.235.1	Function.....	C-116
C.235.2	Example.....	C-116
C.236	profiling-proxy-manager	C-116
C.236.1	Function.....	C-116
C.236.2	Example.....	C-116
C.237	properties	C-116
C.237.1	Function.....	C-118
C.238	property.....	C-118
C.238.1	Function.....	C-118
C.238.2	Example.....	C-118
C.239	proxy-manger-impl.....	C-119
C.239.1	Function.....	C-119
C.239.2	Example.....	C-119
C.240	query-caches	C-119
C.240.1	Function.....	C-119
C.240.2	Example.....	C-119
C.241	quoted-numbers-in-queries.....	C-120
C.241.1	Function.....	C-120
C.241.2	Example.....	C-120
C.242	read-lock-level.....	C-120

C.242.1	Function.....	C-120
C.242.2	Example.....	C-121
C.243	recover-action	C-121
C.243.1	Function.....	C-121
C.243.2	Example.....	C-121
C.244	recovery-time-millis.....	C-121
C.244.1	Function.....	C-121
C.244.2	Example.....	C-122
C.245	resources.....	C-122
C.245.1	Function.....	C-122
C.245.2	Example.....	C-122
C.246	restore-state.....	C-122
C.246.1	Function.....	C-122
C.246.2	Example.....	C-123
C.247	result-set-type.....	C-123
C.247.1	Function.....	C-123
C.247.2	Example.....	C-123
C.248	retain-state.....	C-123
C.248.1	Function.....	C-124
C.248.2	Example.....	C-124
C.249	retry-class-registration	C-124
C.249.1	Function.....	C-124
C.249.2	Example.....	C-124
C.250	scan-top-down.....	C-124
C.250.1	Function.....	C-125
C.250.2	Example.....	C-125
C.251	schema	C-125
C.251.1	Function.....	C-125
C.251.2	Example.....	C-125
C.252	schema-column	C-125
C.252.1	Function.....	C-126
C.252.2	Example.....	C-126
C.253	schemas.....	C-126
C.253.1	Function.....	C-126
C.253.2	Example.....	C-126
C.254	sequence	C-126
C.254.1	Function.....	C-126
C.254.2	Example.....	C-126
C.255	sequence-column.....	C-127
C.255.1	Function.....	C-127
C.255.2	Example.....	C-127
C.256	sequence-name	C-127
C.256.1	Function.....	C-127
C.256.2	Example.....	C-127
C.257	ServiceURL.....	C-127
C.257.1	Function.....	C-128
C.257.2	Example.....	C-128

C.258	simple-driver-data-source	C-128
C.258.1	Function.....	C-128
C.258.2	Example.....	C-128
C.259	single-file	C-128
C.259.1	Function.....	C-128
C.259.2	Example.....	C-129
C.260	single-jvm-exclusive-lock-manager.....	C-129
C.260.1	Function.....	C-129
C.260.2	Example.....	C-129
C.261	single-jvm-remote-commit-provider	C-129
C.261.1	Function.....	C-129
C.261.2	Example.....	C-129
C.262	soft-reference-size	C-130
C.262.1	Function.....	C-130
C.262.2	Example.....	C-130
C.263	so-timeout	C-130
C.263.1	Function.....	C-130
C.263.2	Example.....	C-131
C.264	sql-server-dictionary.....	C-131
C.264.1	Function.....	C-131
C.264.2	Example.....	C-131
C.265	stack-execution-context-name-provider	C-131
C.265.1	Function.....	C-131
C.265.2	Example.....	C-131
C.266	store-mode	C-131
C.266.1	Function.....	C-132
C.266.2	Example.....	C-132
C.267	strict.....	C-132
C.267.1	Function.....	C-133
C.267.2	Example.....	C-133
C.268	strict-identity-values.....	C-133
C.268.1	Function.....	C-134
C.268.2	Example.....	C-134
C.269	style	C-134
C.269.1	Function.....	C-134
C.269.2	Example.....	C-134
C.270	subclass-fetch-mode	C-134
C.270.1	Function.....	C-134
C.270.2	Example.....	C-135
C.271	sybase-dictionary	C-135
C.271.1	Function.....	C-135
C.271.2	Example.....	C-135
C.272	sync-with-managed-transactions.....	C-135
C.272.1	Function.....	C-135
C.272.2	Example.....	C-135
C.273	synchronize-mappings.....	C-135
C.273.1	Function.....	C-136

C.273.2	Example.....	C-136
C.274	table.....	C-136
C.274.1	Function.....	C-136
C.274.2	Example.....	C-137
C.275	table-deprecated-jdo-mapping-factory	C-137
C.275.1	Function.....	C-137
C.275.2	Example.....	C-137
C.276	table-jdbc-seq.....	C-137
C.276.1	Function.....	C-137
C.276.2	Example.....	C-138
C.277	table-jdor-mapping-factory	C-138
C.277.1	Function.....	C-138
C.277.2	Example.....	C-138
C.278	table-lock-update-manager	C-138
C.278.1	Function.....	C-139
C.278.2	Example.....	C-139
C.279	table-name.....	C-139
C.279.1	Function.....	C-139
C.279.2	Example.....	C-140
C.280	table-schema-factory.....	C-140
C.280.1	Function.....	C-140
C.280.2	Example.....	C-140
C.281	tangosol-cache-name	C-140
C.281.1	Function.....	C-141
C.281.2	Example.....	C-141
C.282	tangosol-cache-type	C-141
C.282.1	Function.....	C-141
C.282.2	Example.....	C-141
C.283	tangosol-data-cache	C-141
C.283.1	Function.....	C-142
C.283.2	Example.....	C-142
C.284	tcp-remote-commit-provider.....	C-142
C.284.1	Function.....	C-142
C.284.2	Example.....	C-142
C.285	tcp-transport	C-142
C.285.1	Function.....	C-143
C.285.2	Example.....	C-143
C.286	time-seeded-seq.....	C-143
C.286.1	Function.....	C-143
C.286.2	Example.....	C-143
C.287	topic.....	C-143
C.287.1	Function.....	C-144
C.287.2	Example.....	C-144
C.288	topic-connection-factory	C-144
C.288.1	Function.....	C-144
C.288.2	Example.....	C-144
C.289	track-changes.....	C-144

C.289.1	Function.....	C-144
C.289.2	Example.....	C-144
C.290	transaction-isolation	C-145
C.290.1	Function.....	C-145
C.290.2	Example.....	C-145
C.291	transaction-mode	C-145
C.291.1	Function.....	C-145
C.291.2	Example.....	C-146
C.292	transaction-name-execution-context-name-provider	C-146
C.292.1	Function.....	C-146
C.292.2	Example.....	C-146
C.293	type.....	C-146
C.293.1	Function.....	C-146
C.293.2	Example.....	C-147
C.294	type-column.....	C-147
C.294.1	Function.....	C-147
C.294.2	Example.....	C-147
C.295	types.....	C-147
C.295.1	Function.....	C-147
C.295.2	Example.....	C-147
C.296	URL	C-148
C.296.1	Function.....	C-148
C.296.2	Example.....	C-148
C.297	url	C-148
C.297.1	Function.....	C-148
C.297.2	Example.....	C-148
C.298	urls.....	C-148
C.298.1	Function.....	C-149
C.298.2	Example.....	C-149
C.299	use-aliases	C-149
C.299.1	Function.....	C-149
C.299.2	Example.....	C-150
C.300	use-schema-validation.....	C-150
C.300.1	Function.....	C-150
C.300.2	Example.....	C-150
C.301	user-object-execution-context-name-provider.....	C-150
C.301.1	Function.....	C-151
C.301.2	Example.....	C-151
C.302	UserName	C-151
C.302.1	Function.....	C-151
C.302.2	Example.....	C-151
C.303	validate-false-returns-hollow	C-151
C.303.1	Function.....	C-152
C.303.2	Example.....	C-152
C.304	validate-true-checks-store.....	C-152
C.304.1	Function.....	C-152
C.304.2	Example.....	C-152

C.305	value-table-jdbc-seq.....	C-152
C.305.1	Function.....	C-153
C.305.2	Example.....	C-153
C.306	version-check-on-read-lock.....	C-153
C.306.1	Function.....	C-153
C.306.2	Example.....	C-153
C.307	version-check-on-write-lock.....	C-153
C.307.1	Function.....	C-154
C.307.2	Example.....	C-154
C.308	version-lock-manager.....	C-154
C.308.1	Function.....	C-154
C.308.2	Example.....	C-154
C.309	wls81-jmx.....	C-154
C.309.1	Function.....	C-154
C.309.2	Example.....	C-155
C.310	write-lock-level.....	C-155
C.310.1	Function.....	C-155
C.310.2	Example.....	C-155

Preface

This preface describes the document accessibility features and conventions used in this guide—*Programming Enterprise JavaBeans for Oracle WebLogic Server*.

WLACH0010 and WLACH0020

Documentation Accessibility

For information about Oracle's commitment to accessibility, visit the Oracle Accessibility Program website at

<http://www.oracle.com/pls/topic/lookup?ctx=acc&id=docacc>.

Access to Oracle Support

Oracle customers have access to electronic support through My Oracle Support. For information, visit

<http://www.oracle.com/pls/topic/lookup?ctx=acc&id=info> or visit

<http://www.oracle.com/pls/topic/lookup?ctx=acc&id=trs> if you are hearing impaired.

Conventions

The following text conventions are used in this document:

Convention	Meaning
boldface	Boldface type indicates graphical user interface elements associated with an action, or terms defined in text or the glossary.
<i>italic</i>	Italic type indicates book titles, emphasis, or placeholder variables for which you supply particular values.
monospace	Monospace type indicates commands within a paragraph, URLs, code in examples, text that appears on the screen, or text that you enter.

Introduction and Roadmap

This section describes the contents and organization of this guide—*Programming Enterprise JavaBeans for Oracle WebLogic Server*.

- [Section 1.1, "Document Scope and Audience"](#)
- [Section 1.2, "Guide to this Document"](#)
- [Section 1.3, "Related Documentation"](#)
- [Section 1.4, "Comprehensive Examples for the EJB Developer"](#)
- [Section 1.5, "New and Changed Features in this Release"](#)

1.1 Document Scope and Audience

This document is a resource for software developers who develop applications that include WebLogic Server Enterprise JavaBeans (EJBs).

The document mostly discusses the Java EE 6-based, EJB 3.1 programming model, in particular the use of metadata annotations to simplify development. This document does not address EJB topics that are different between versions 2.x and 3.x, such as design considerations, EJB container architecture, entity beans, 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 [Section 1.3, "Related Documentation."](#)

It is assumed that the reader is familiar with Java Platform, Enterprise Edition (Java EE) Version 6 and basic EJB programming concepts.

For information on programming and packaging 2.1 EJBs, see *Programming Enterprise JavaBeans, Version 2.1, for Oracle WebLogic Server*.

1.2 Guide to this Document

- This chapter, [Chapter 1, "Introduction and Roadmap,"](#) introduces the organization of this guide.
- [Chapter 2, "Understanding Enterprise JavaBeans,"](#) provides an overview of the new EJB 3.1 features, as well as a brief description of the differences between EJB 3.1 and 3.0.
- [Chapter 3, "Simple Enterprise JavaBeans Examples,"](#) provides examples of programming EJBs using the metadata annotations specified by EJB 3.x.

- [Chapter 4, "Iterative Development of Enterprise JavaBeans,"](#) 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 Class,"](#) describes the requirements and typical steps when programming the EJB bean class that contains the metadata annotations.
- [Chapter 6, "Deployment Guidelines for Enterprise JavaBeans,"](#) discusses EJB-specific deployment issues and procedures.
- [Chapter 7, "Using an Embedded EJB Container in Oracle WebLogic Server,"](#) discusses using an embeddable EJB container in Oracle WebLogic Server.
- [Chapter 8, "Configuring the Persistence Provider in Oracle WebLogic Server,"](#) provides an overview of developing an Oracle TopLink application using Oracle WebLogic Server.
- [Appendix A, "EJB 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, "Using Oracle Kodo with Oracle 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).

Note: Oracle Kodo JPA/JDO is deprecated in this release. Customers are encouraged to use Oracle TopLink, which supports JPA 2.0. Kodo supports only JPA 1.0.

- [Appendix C, "Oracle Kodo Persistence Configuration Schema Reference,"](#) provides reference information for the persistence configuration schema.

1.3 Related Documentation

This document contains EJB 3.1-specific development information. Additionally, it provides information only for session and message-driven beans. For complete information on general EJB design and architecture, the EJB 2.x programming model (which is fully supported in EJB 3.1), see the following documents.

1.3.1 EJB Documentation in WebLogic Server

For information about developing and deploying EJBs with WebLogic Server, see:

- "Enterprise Java Beans (EJBs)" in Introduction to Oracle WebLogic Server.
- For instructions on how to organize and build WebLogic Server EJBs in a split directory environment, see *Developing Applications for Oracle WebLogic Server*.
- For information on programming and packaging 2.x EJBs, see *Programming Enterprise JavaBeans, Version 2.1, for Oracle WebLogic Server*.
- *Deploying Applications to Oracle WebLogic Server* is the primary source of information about deploying WebLogic Server applications in development and production environments.

1.3.2 Additional EJB Information

It is assumed the reader is familiar with programming in Java EE 6 and EJB 3.1 concepts and features. To learn more about basic EJB concepts, such as the benefits of enterprise beans, the types of enterprise beans, and their life cycles, then visit the following Web sites:

- Enterprise JavaBeans 3.1 Specification (JSR-318) at <http://jcp.org/en/jsr/summary?id=318>
- The "Enterprise Beans" chapter of the Java EE 6 Tutorial at <http://download.oracle.com/javaee/6/tutorial/doc/bnblr.html>
- Introducing the Java EE 6 Platform: Part 3 (EJB Technology, Even Easier to Use) at <http://www.oracle.com/technetwork/articles/javaee/javaee6overview-part3-139660.html#ejbeasy>

1.4 Comprehensive Examples for the EJB Developer

In addition to this document and the basic examples described in [Chapter 3, "Simple Enterprise JavaBeans Examples,"](#) Oracle provides comprehensive examples in the WebLogic Server distribution kit.

WebLogic Server optionally installs these in `WL_HOME/samples/server/examples/src/examples`, where `WL_HOME` refers to the directory in which you installed WebLogic Server, such as `/Oracle/Middleware/wlserver_12.1`. 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 these examples before programming your own application that uses EJBs.

1.4.1 New EJB 3.1 Examples

Oracle provides Java EE 6 examples that demonstrate new features in EJB 3.1, such as:

- [Section 3.2.1, "EJB 3.1: Example of a Singleton Session Bean"](#)
- [Section 3.2.2, "EJB 3.1: Example of an Asynchronous Method EJB"](#)
- [Section 3.2.3, "EJB 3.1: Example of a Calendar-based Timer EJB"](#)
- [Section 3.2.4, "EJB 3.1: Example of Simplified No-interface Programming and Packaging in a WAR File"](#)
- [Section 3.2.5, "EJB 3.1: Example of Using a Portable Global JNDI Name in an EJB"](#)

For more information, see these examples in the WebLogic Server distribution kit: `WL_HOME/samples/server/examples/src/examples/javaee6/ejb`.

1.4.2 EJB 3.0 Example

There is also an EJB 3.0 persistence service example, [Section 3.2.6, "EJB 3.0: Example of Invoking an Entity From A Session Bean"](#), that includes actual business code and provides practical instructions on how to perform key EJB 3.0 development tasks. In particular, the example demonstrates usage of EJB 3.x 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. For more information, see the example in the WebLogic Server distribution kit: `WL_HOME/samples/server/examples/src/examples/ejb/ejb30`.

1.5 New and Changed Features in this Release

For a comprehensive listing of the new features in EJB 3.1 features introduced in this release of WebLogic Server, see "Enterprise Java Beans (EJBs)" in *What's New in Oracle WebLogic Server*.

Understanding Enterprise JavaBeans

These sections describe the new features and programming model of EJB 3.1 and also provide a basic overview EJB components, anatomy, and features.

- [Section 2.1, "New Features and Changes in EJB"](#)
- [Section 2.2, "Understanding EJB Components"](#)
- [Section 2.3, "EJB Anatomy and Environment"](#)
- [Section 2.4, "EJB Clients and Communications"](#)
- [Section 2.5, "Securing EJBs"](#)

2.1 New Features and Changes in EJB

These sections summarize the changes in the EJB programming model and requirements between EJB 3.0 and 3.1, as well between EJB 2.x and 3.0.

2.1.1 What Is New and Changed in EJB 3.1

The EJB 3.1 specification provides simplified programming and packaging model changes. The mandatory use of Java interfaces from previous versions has been removed, allowing plain old Java objects to be annotated and used as EJB components. The simplification is further enhanced through the ability to place EJB components directly inside of Web applications, removing the need to produce separate archives to store the Web and EJB components and combine them together in an EAR file.

For a comprehensive listing of the new features in EJB 3.1 features introduced in this release of WebLogic Server, see "Enterprise Java Beans (EJBs)" in *What's New in Oracle WebLogic Server*.

- **Singleton Session Bean** – singleton session beans provide a formal programming construct that guarantees a session bean will be instantiated once per application in a particular Java Virtual Machine (JVM), and that it will exist for the life cycle of the application. With singletons, you can easily share state between multiple instances of an enterprise bean component or between multiple enterprise bean components in the application.
- **Simplified No Interface Client View** – The No-interface local client view type simplifies EJB development by providing local session bean access without requiring a separate local business interface, allowing components to have EJB bean class instances directly injected.
- **Packaging and Deploying EJBs Directly in a WAR File** – EJB 3.1 provides the ability to place EJB components directly inside of Web application archive (WAR)

files, removing the need to produce archives to store the Web and EJB components and combine them together in an enterprise application archive (EAR) file.

- **Portable Global JNDI Names** – The Portable Global JNDI naming option in EJB 3.1 provides a number of common, well-known namespaces in which EJB components can be registered and looked up from using the pattern `java:global[/<app-name>]/<module-name>/<bean-name>`. This standardizes how and where EJB components are registered in JNDI, and how they can be looked up and used by applications.
- **Asynchronous Session Bean Invocations** – An EJB 3.1 session bean can expose methods with asynchronous client invocation semantics. Using the `@Asynchronous` annotation in an EJB class or specific method will direct the EJB container to return control immediately to the client when the method is invoked. The method may return a future object to allow the client to check on the status of the method invocation, and retrieve result values that are asynchronously produced.
- **EJB Timer Enhancements** – The EJB 3.1 Timer Service supports calendar-based EJB Timer expressions. The scheduling functionality takes the form of CRON-styled schedule definitions that can be placed on EJB methods, in order to have the methods be automatically invoked according to the defined schedule. EJB 3.1 also supports the automatic creation of a timer based on metadata in the bean class or deployment descriptor, which allows the bean developer to schedule a timer without relying on a bean invocation to programmatically invoke one of the Timer Service timer creation methods. Automatically created timers are created by the container as a result of application deployment.
- **Embeddable EJB Container** – EJB 3.1 supports an embeddable API for executing EJB components within a Java SE environment. Unlike traditional Java EE server-based execution, embeddable usage allows client code and its corresponding enterprise beans to run within the same virtual machine and class loader. This provides better support for testing, offline processing (e.g., batch jobs), and the use of the EJB programming model in desktop applications.
- **JPA 2.0 Support** – Oracle EclipseLink is the default JPA 2.0 persistence provider that is shipped with Oracle WebLogic Server. Oracle Kodo was deprecated in Oracle Fusion Middleware 11gR1 (11.1.1.1.0) (Oracle WebLogic Server 10.3.1). WebLogic Server runs with the JPA 2.0 JAR in the server `classpath`. Although JPA 2.0 is upwardly compatible with JPA 1.0, JPA 2.0 introduced some methods to existing JPA interfaces that conflict with existing signatures in OpenJPA interfaces.

As a result, applications that continue to use Kodo/JPA as the persistence provider with WebLogic Server 12.1.1 must be recompiled. For more information, see "Updating Applications to Overcome Conflicts" in *Programming Enterprise JavaBeans for Oracle WebLogic Server*

- **JPA 2.0 Support Using the Default TopLink Persistence Provider** – Oracle TopLink, a JPA 2.0 persistence provider, is now the default JPA provider, replacing Kodo, which was the default provider in previous releases. Any application that does not specify a JPA provider in `persistence.xml` will now use TopLink by default. Applications can continue to use Kodo (a JPA 1.0 provider) by explicitly specifying Kodo/OpenJPA as their persistence provider in `persistence.xml`. In addition, a Weblogic Server domain can be configured to use Kodo by default, if desired.

For more information, see [Chapter 8, "Configuring the Persistence Provider in Oracle WebLogic Server."](#)

- **Applications That Use Kodo as the Persistence Provider** – This release of WebLogic Server runs with the JPA 2.0 JAR in the server's classpath. Although JPA 2.0 is upwardly compatible with JPA 1.0, JPA 2.0 introduced some methods to existing JPA interfaces that conflict with existing signatures in OpenJPA interfaces. As a result, applications that continue to use Kodo/JPA as the persistence provider with WebLogic Server 12.1.1 must be recompiled. For more information, see [Section 8.4.3, "Updating Applications to Overcome Conflicts with JPA 2.0."](#)

2.1.2 What Was New and Changed in EJB 3.0

The following summarizes the new functionality and simplifications made in EJB 3.0 to the earlier EJB APIs:

- 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 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.
- 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`.
- Session beans may expose client views via business interfaces. Session beans may either explicitly implement the business interface or they 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`.
- 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.

Because the EJB 3.x programming model is so simple, Oracle no longer supports using the EJBGen tags and code-generating tool on EJB 3.x beans. Rather, you can use this tool *only* on 2.x beans. For information, see the "EJBGen Reference" in *Programming Enterprise JavaBeans, Version 2.1, for Oracle WebLogic Server*.

2.2 Understanding EJB Components

Enterprise JavaBeans (EJB) 3.1 technology is the server-side component architecture for Java EE 6. EJB 3.1 technology enables rapid and simplified development of distributed, transactional, secure and portable applications based on Java technology.

2.2.1 Session EJBs Implement Business Logic

Session beans implement business logic. There are three types of session beans: stateful, stateless, and singleton. Stateful and stateless session beans serve one client at a time; whereas, singleton session beans can be invoked concurrently.

For detailed information about the types of session beans and when to use them, see "What Is a Session Bean" in the "Enterprise Beans" chapter of the Java EE 6 Tutorial at <http://download.oracle.com/javase/6/tutorial/doc/gipjg.html>

2.2.1.1 Stateful Session Beans

Stateful session beans maintain state information that reflects the interaction between the bean and a particular client across methods and transactions. A stateful session bean can manage interactions between a client and other enterprise beans, or manage a workflow.

Example: A company Web site that allows employees to view and update personal profile information could use a stateful session bean to call a variety of other beans to provide the services required by a user, after the user clicks "View my Data" on a page:

- Accept the login data from a JSP, and call another EJB whose job it is to validate the login data.
- Send confirmation of authorization to the JSP.
- Call a bean that accesses profile information for the authorized user.

2.2.1.2 Stateless Session Beans

A stateless session bean does not store session or client state information between invocations—the only state it might contain is not specific to a client, for instance, a cached database connection or a reference to another EJB. At most, a stateless session bean may store state for the duration of a method invocation. When a method completes, state information is not retained.

Any instance of a stateless session bean can serve any client—any instance is equivalent. Stateless session beans can provide better performance than stateful session beans, because each stateless session bean instance can support multiple clients, albeit one at a time. The client of a stateless session bean can be a web service endpoint.

Example: An Internet application that allows visitors to click a "Contact Us" link and send an email could use a stateless session bean to generate the email, based on the "to" and "from" information gathered from the user by a JSP.

2.2.1.3 Singleton Session Beans

Singleton session beans provide a formal programming construct that guarantees a session bean will be instantiated once per application in a particular Java Virtual Machine (JVM), and that it will exist for the life cycle of the application. With singletons, you can easily share state between multiple instances of an enterprise bean component or between multiple enterprise bean components in the application.

Singleton session beans offer similar functionality to stateless session beans but differ from them in that there is only one singleton session bean per application, as opposed to a pool of stateless session beans, any of which may respond to a client request. Like stateless session beans, singleton session beans can implement Web service endpoints. Singleton session beans maintain their state between client invocations but are not required to maintain their state across server crashes or shutdowns.

Example: The Apache Web site provides a "Simple Singleton: ComponentRegistry" example that demonstrates how a singleton bean uses Container-Managed Concurrency to utilize the Read (`@Lock(READ)`) functionality, to allow multi-threaded access to the bean, and the Write (`@Lock(WRITE)`) functionality, to enforce single-threaded access to the bean.

2.2.2 Message-Driven Beans Implement Loosely Coupled Business Logic

A message-driven bean implements loosely coupled or asynchronous business logic in which the response to a request need not be immediate. A message-driven bean receives messages from a JMS Queue or Topic, and performs business logic based on the message contents. It is an asynchronous interface between EJBs and JMS.

Throughout its life cycle, an MDB instance can process messages from multiple clients, although not simultaneously. It does not retain state for a specific client. All instances of a message-driven bean are equivalent—the EJB container can assign a message to any MDB instance. The container can pool these instances to allow streams of messages to be processed concurrently.

The EJB container interacts directly with a message-driven bean—creating bean instances and passing JMS messages to those instances as necessary. The container creates bean instances at deployment time, adding and removing instances during operation based on message traffic.

For detailed information, see *Programming Message-Driven Beans for Oracle WebLogic Server*.

Example: In an on-line shopping application, where the process of taking an order from a customer results in a process that issues a purchase order to a supplier, the supplier ordering process could be implemented by a message-driven bean. While taking the customer order always results in placing a supplier order, the steps are loosely coupled because it is not necessary to generate the supplier order before confirming the customer order. It is acceptable or beneficial for customer orders to "stack up" before the associated supplier orders are issued.

2.3 EJB Anatomy and Environment

These sections briefly describe classes required for each bean type, the EJB run-time environment, and the deployment descriptor files that govern a bean's run-time behavior.

- [Section 2.3.1, "EJB Components"](#)
- [Section 2.3.2, "The EJB Container"](#)

- [Section 2.3.4, "Optional EJB Deployment Descriptors"](#)

2.3.1 EJB Components

Every bean type requires a bean class. [Table 2–1](#) defines the supported client views that make up each type of EJB, and defines any additional required classes.

Note: The EJB 2.1 and earlier API required that Local and Remote clients access the stateful or stateless session bean by means of the session bean's local or remote home and the local or remote component interfaces. These interfaces remain available for use with EJB 3.x; however, the EJB 2.1 Remote and Local client view is not supported for singleton session beans.

For more information see "Create EJB Classes and Interfaces" in *Programming Enterprise JavaBeans, Version 2.1, for Oracle WebLogic Server*.

Table 2–1 Supported Client Views in EJB 3.1

Client Views	Session Bean Types	Additional Required Classes
Remote Client	Stateful, Stateless, and Singleton session beans	Remote business interface that defines the bean's business and lifecycle methods.
Local Client	Stateful, Stateless, and Singleton session beans	Local business interface that defines the bean's business and lifecycle methods.
Local No- interface	Stateful, Stateless, and Singleton session beans	Only requires the bean class.
Web Service Clients	Stateless and Singleton session beans	A Web service endpoint that is accessed as a JAX-WS or JAX-RPC service endpoint using the JAX-WS or JAX-RPC client view APIs.

2.3.2 The EJB Container

An EJB container is a run-time container for beans that are deployed to an application server. The container is automatically created when the application server starts up, and serves as an interface between a bean and run-time services such as:

- Life cycle management
- Code generation
- Security
- Transaction management
- Locking and concurrency control

2.3.3 EJB Metadata Annotations

The WebLogic Server EJB 3.1 programming model uses the Java EE 6 metadata annotations feature in which you create an annotated EJB 3.1 bean file, and then compile the class with standard Java compiler, which can then be packaged into a target module for deployment. At runtime, WebLogic Server parses the annotations and applies the required behavioral aspects to the bean file.

For more information, see [Section 5, "Programming the Annotated EJB Class."](#)

2.3.4 Optional EJB Deployment Descriptors

As of EJB 3.0, you are no longer required to create the EJB deployment descriptor files (such as `ejb-jar.xml`). However, you can still use XML deployment descriptors if you want. In the case of conflicts, the deployment descriptor value overrides the annotation value.

If you are continuing to use deployment descriptors in your EJB implementation, refer to "EJB Deployment Descriptors" in *Programming Enterprise JavaBeans, Version 2.1, for Oracle WebLogic Server*.

The WebLogic Server EJB container supports three deployment descriptors:

- `ejb-jar.xml`—The standard Java EE deployment descriptor. The `ejb-jar.xml` may be used to define EJBs and to specify standard configuration settings for the EJBs. An `ejb-jar.xml` can specify multiple beans that will be deployed together.
- `weblogic-ejb-jar.xml`—WebLogic Server-specific deployment descriptor that contains elements related to WebLogic Server features such as clustering, caching, and transactions. This file is required if your beans take advantage of WebLogic Server-specific features. Like `ejb-jar.xml`, `weblogic-ejb-jar.xml` can specify multiple beans that will be deployed together.
- `weblogic-cmp-jar.xml`—WebLogic Server-specific deployment descriptor that contains elements related to container-managed persistence for entity beans. Entity beans that use container-managed persistence must be specified in a `weblogic-cmp-jar.xml` file.

For descriptions of the WebLogic Server EJB deployment descriptors, refer to "Deployment Descriptor Schema and Document Type Definitions Reference" in *Programming Enterprise JavaBeans, Version 2.1, for Oracle WebLogic Server*.

2.4 EJB Clients and Communications

An EJB can be accessed by server-side or client-side objects such as servlets, Java client applications, other EJBs, web services, and non-Java clients. Any client of an EJB, whether in the same or a different application, accesses it in a similar fashion. WebLogic Server automatically creates implementations of an EJB's remote home and remote business interfaces, which can function remotely.

2.4.1 Accessing EJBs

Clients access enterprise beans either through a no-interface view or through a business interface. A no-interface view of an enterprise bean exposes the public methods of the enterprise bean implementation class to clients. Clients using the no-interface view of an enterprise bean may invoke any public methods in the enterprise bean implementation class or any super-classes of the implementation class. A business interface is a standard Java programming language interface that contains the business methods of the enterprise bean.

The client of an enterprise bean obtains a reference to an instance of an enterprise bean through either dependency injection, using Java programming language annotations, or JNDI lookup, using the Java Naming and Directory Interface syntax to find the enterprise bean instance.

Dependency injection is the simplest way of obtaining an enterprise bean reference. Clients that run within a Java EE server-managed environment, JavaServer Faces web

applications, JAX-RS web services, other enterprise beans, or Java EE application clients, support dependency injection using the `javax.ejb.EJB` annotation.

Applications that run outside a Java EE server-managed environment, such as Java SE applications, must perform an explicit lookup. JNDI supports a global syntax for identifying Java EE components to simplify this explicit lookup. For more information see, [Section 4.7.3.2, "Using the JNDI Portable Syntax"](#).

Because of network overhead, it is more efficient to access beans from a client on the same machine than from a remote client, and even more efficient if the client is in the same application.

For information on programming client access to an EJB, see "Accessing Enterprise Bean" in the "Enterprise Beans" chapter of the Java EE 6 Tutorial at <http://download.oracle.com/javaee/6/tutorial/doc/gipsc.html>.

2.4.2 EJB Communications

WebLogic Server EJBs use:

- T3—To communicate with remote objects. T3 is a WebLogic-proprietary remote network protocol that implements the Remote Method Invocation (RMI) protocol.
- RMI—To communicate with remote objects. RMI enables an application to obtain a reference to an object located elsewhere in the network, and to invoke methods on that object as though it were co-located with the client on the same JVM locally in the client's virtual machine. An EJB with a remote interface is an RMI object. For more information on WebLogic RMI, see *Programming RMI for Oracle WebLogic Server*.
- HTTP—An EJB can obtain an HTTP connection to a Web server external to the WebLogic Server environment by using the `java.net.URL` resource connection factory. For more information, see "Configuring EJBs to Send Requests to an URL" in *Programming Enterprise JavaBeans, Version 2.1, for Oracle WebLogic Server*.

You can specify the attributes of the network connection an EJB uses by binding the EJB to a WebLogic Server custom network channel. For more information, see "Configuring Network Resources" in *Configuring Server Environments for Oracle WebLogic Server*.

2.5 Securing EJBs

By default, any user can invoke the public methods of an EJB. Therefore, if you want to restrict access to the EJB, you can use security-related annotations to specify the roles that are allowed to invoke all, or a subset, of the methods, which is explained in [Section 5.3.8, "Securing Access to the EJB."](#)

In addition, 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" in the *Oracle WebLogic Server Administration Console Help*.

For additional information about security and EJBs, see:

- "Security Fundamentals" in *Understanding Security for Oracle WebLogic Server* has introductory information about authentication, authorization and other security topics.
- "Securing Enterprise JavaBeans (EJBs)" in *Programming Security for Oracle WebLogic Server* provides instructions for configuring authentication and authorization for EJBs.

- *Securing Resources Using Roles and Policies for Oracle WebLogic Server* contains instructions for on configuring authentication and authorization for your EJBs using the Administration Console.

Simple Enterprise JavaBeans Examples

The following sections describe Java examples of EJBs that use the version 3.x programming model:

- [Section 3.1, "Simple Java Examples of 3.x EJBs"](#)
- [Section 3.2, "Packaged EJB 3.1 Examples in WebLogic Server"](#)

3.1 Simple Java Examples of 3.x EJBs

The following sections describe simple Java examples of EJBs that use the new metadata annotation programming model. Some procedural sections in this guide that describe how to program an EJB may reference these examples.

- [Section 3.1.1, "Example of a Simple No-interface Stateless EJB"](#)
- [Section 3.1.2, "Example of a Simple Business Interface Stateless EJB"](#)
- [Section 3.1.3, "Example of a Simple Stateful EJB"](#)
- [Section 3.1.4, "Example of an Interceptor Class"](#)

3.1.1 Example of a Simple No-interface Stateless EJB

The EJB 3.1 no-interface local client view type simplifies EJB development by providing local session bean access without requiring a separate local business interface, allowing components to have EJB bean class instances directly injected.

The following code shows a simple no-interface view for the `ServiceBean` stateless session EJB:

```
package examples;
@Stateless
public class ServiceBean {
    public void sayHelloFromServiceBean() {
        System.out.println("Hello From Service Bean!");
    }
}
```

The main points to note about the preceding code are:

- The EJB automatically exposes the no-interface view because no other client views are exposed and its bean class implements clause is empty.
- The `ServiceBean` bean file is a plain Java file; it is not required to implement any EJB-specific interface.

- The class-level `@Stateless` metadata annotation specifies that the EJB is of type stateless session.

3.1.2 Example of a Simple Business Interface 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 javax.ejb.Stateless;
import javax.interceptor.ExcludeDefaultInterceptors;
/**
 * Bean file that implements the Service business interface.
 * Class uses following EJB 3.x annotations:
 * - @Stateless - specifies that the EJB is of type stateless session
 * - @ExcludeDefaultInterceptors - specifies any configured default
 *   interceptors should not be invoked for this class
 */
@Stateless
@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 javax.ejb.Stateless;
import javax.interceptor.ExcludeDefaultInterceptors
```

The annotations that apply only to EJB 3.1 are in the `javax.ejb` package. Annotations that can be used by other Java EE Version 6 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 `@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.

3.1.3 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 javax.ejb.Stateful;
import javax.ejb.Remote;
import javax.ejb.EJB;
import javax.annotation.PreDestroy;
import javax.interceptor.Interceptors;
import javax.interceptor.ExcludeClassInterceptors;
/**
 * Bean file that implements the Account business interface.
 * Uses the following EJB annotations:
 * - @Stateful: specifies that this is a stateful session EJB
 * - @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
@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 javax.ejb.Stateful;
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 3.1 are in the `javax.ejb` package. Annotations that can be used by other Java EE 6 components are in more generic packages, such as `javax.interceptor` or `javax.annotation`.

- 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 `@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 [Section 3.1.4, "Example of an Interceptor Class"](#) 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.

3.1.4 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 a 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.

3.2 Packaged EJB 3.1 Examples in WebLogic Server

The following sections describe the packaged Java EE 6 examples included with Oracle WebLogic Server, which demonstrate new features in EJB 3.1.

- [Section 3.2.1, "EJB 3.1: Example of a Singleton Session Bean"](#)
- [Section 3.2.2, "EJB 3.1: Example of an Asynchronous Method EJB"](#)
- [Section 3.2.3, "EJB 3.1: Example of a Calendar-based Timer EJB"](#)
- [Section 3.2.4, "EJB 3.1: Example of Simplified No-interface Programming and Packaging in a WAR File"](#)
- [Section 3.2.5, "EJB 3.1: Example of Using a Portable Global JNDI Name in an EJB"](#)
- [Section 3.2.6, "EJB 3.0: Example of Invoking an Entity From A Session Bean"](#)

3.2.1 EJB 3.1: Example of a Singleton Session Bean

This example demonstrates the use of the EJB 3.1 singleton session bean, which provides application developers with a formal programming construct that guarantees a session bean will be instantiated once for an application in a particular Java Virtual Machine (JVM). In this example, a `@Singleton` session bean provides a central counter service. The Counter EJB is called from a Java client to demonstrate it is being used, with the count being consistently incremented by "1" as the client is invoked multiple times.

After you have installed WebLogic Server, the example is in the following directory:

```
WL_HOME/samples/server/examples/src/examples/javaee6/ejb/singletonBean
```

`WL_HOME` refers to the directory in which you installed WebLogic Server, such as `/Oracle/Middleware/wlserver_12.1`.

3.2.2 EJB 3.1: Example of an Asynchronous Method EJB

This example demonstrates the use of the EJB 3.1 asynchronous method invocation. Adding the `@Asynchronous` annotation to an EJB class or specific method will direct the EJB container to return control immediately to the client when the method is invoked. The method may return a `Future` object to allow the client to check on the status of the method invocation, and then retrieve result values that are asynchronously produced.

In this example, an `@Stateless` bean is annotated at the class level, with `@Asynchronous` indicating its methods are all asynchronous, with each of the methods simulating a long-running calculation. A servlet is used to call the various

asynchronous methods, keeping track of the invocation and completion times to demonstrate the asynchronous nature of the method calls.

After you have installed WebLogic Server, the example is in the following directory:

```
WL_HOME/samples/server/examples/src/examples/javaee6/ejb/asyncMethodOfEJB
```

WL_HOME refers to the directory in which you installed WebLogic Server, such as /Oracle/Middleware/wlserver_12.1.

3.2.3 EJB 3.1: Example of a Calendar-based Timer EJB

This example demonstrates the enhanced scheduling capabilities of EJB 3.1. This scheduling functionality takes the form of CRON-styled schedule definitions that can be placed on EJB methods, in order for the methods to be automatically invoked according to the defined schedule. This example shows the use of the `@Schedule` annotation defined for a method of a `@Singleton` session bean, which generates and stores the timestamp of when the method was called. A corresponding servlet is provided, into which the `TimerBean` is injected, which retrieves the list of timestamps to display in a browser.

After you have installed WebLogic Server, the example is in the following directory:

```
WL_HOME/samples/server/examples/src/examples/javaee6/ejb/calendarStyledTimer
```

WL_HOME refers to the directory in which you installed WebLogic Server, such as /Oracle/Middleware/wlserver_12.1.

3.2.4 EJB 3.1: Example of Simplified No-interface Programming and Packaging in a WAR File

This example demonstrates the simplified programming and packaging model changes provided in EJB 3.1. Since the mandatory use of Java interfaces from previous versions has been removed in EJB 3.1, plain old Java objects can be annotated and used as EJB components. The simplification is further enhanced by the ability to place EJB components directly inside of Web applications, thereby removing the need to produce archives to store the Web and EJB components and combine them together in an enterprise archive (EAR) file.

In the example, an `@Stateless` annotation is provided on a plain old Java class that exposes it as an EJB session bean. This is then injected into an `@WebServlet` class using an `@EJB` annotation to demonstrate that it is being used as an EJB module. The EJB session bean and servlet classes are then packaged and deployed together in a single WAR file, which demonstrates the simplified packaging and deployment changes available in Java EE 6.

After you have installed WebLogic Server, the example is in the following directory:

```
WL_HOME/samples/server/examples/src/examples/javaee6/ejb/noInterfaceViewInWAR
```

WL_HOME refers to the directory in which you installed WebLogic Server, such as /Oracle/Middleware/wlserver_12.1.

3.2.5 EJB 3.1: Example of Using a Portable Global JNDI Name in an EJB

This example demonstrates the use of the Portable Global JNDI naming option that is available in EJB 3.1. Portable Global JNDI provides a number of common, well-known namespaces in which EJB components can be registered and looked up from using the pattern `java:global[/<app-name>]/<module-name>/<bean-name>`. This

standardizes how and where EJB components are registered in JNDI and how they can be looked up and used by applications. In this example, a servlet is used to look up an EJB session bean using its portable JNDI name `java:module/HelloBean`.

After you have installed WebLogic Server, the example is in the following directory:

`WL_HOME/samples/server/examples/src/examples/javaee6/ejb/portableGlobalJNDIName`

`WL_HOME` refers to the directory in which you installed WebLogic Server, such as `/Oracle/Middleware/wlserver_12.1`.

3.2.6 EJB 3.0: Example of Invoking an 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`

`WL_HOME` refers to the directory in which you installed WebLogic Server, such as `/Oracle/Middleware/wlserver_12.1`.

Iterative Development of Enterprise JavaBeans

The sections that follow describe the general EJB 3.1 implementation process, and provide guidance for how to get an EJB up and running in WebLogic Server.

- [Section 4.1, "Overview of the EJB Development Process"](#)
- [Section 4.2, "Create a Source Directory"](#)
- [Section 4.3, "Program the Annotated EJB Class"](#)
- [Section 4.4, "Program the EJB Interface"](#)
- [Section 4.5, "Optionally Program Interceptors"](#)
- [Section 4.6, "Optionally Program the EJB Timer Service"](#)
- [Section 4.7, "Programming Access to EJB Clients"](#)
- [Section 4.8, "Programming and Configuring Transactions"](#)
- [Section 4.9, "Compile Java Source"](#)
- [Section 4.10, "Optionally Create and Edit Deployment Descriptors"](#)
- [Section 4.11, "Packaging EJBs"](#)
- [Section 4.12, "Deploying EJBs"](#)

4.1 Overview of the EJB Development Process

This section is a brief overview of the EJB 3.1 development process. It describes the key implementation tasks and associated results.

The following section mostly discusses the EJB 3.1 programming model and points out the differences between the 3.1 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 [Section 2.1, "New Features and Changes in EJB."](#)

Table 4–1 EJB Development Tasks and Results

#	Step	Description	Result
1	Section 4.2, "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	Section 4.3, "Program the Annotated EJB Class"	Create the Java file that implements the interface and includes the EJB 3.1 metadata annotations that describe how your EJB behaves.	.java file.
3	Section 4.4, "Program the EJB Interface"	Create no-interface client views or business interfaces that describe your EJB.	.java file for each interface.
4	Section 4.5, "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 for each interceptor class.
5	Section 4.6, "Optionally Program the EJB Timer Service"	Optionally, create timers that schedule callbacks to occur when a timer object expires for timed event.	Either metadata (for automatic timers) and/or bean class changes (for programmatic timers).
6	Section 4.7, "Programming Access to EJB Clients"	Obtain a reference to an EJB through either dependency injection or JNDI lookup.	Metadata (annotations and/or deployment descriptor settings) and/or code changes to the client.
7	Section 4.8, "Programming and Configuring Transactions"	Program container-managed or bean-managed transactions.	Metadata and possibly logic to handle exceptions (retry logic or calls to <code>setRollbackOnly</code>).
8	Section 4.9, "Compile Java Source"	Compile source code.	.class file for each class and interface.
9	Section 4.10, "Optionally Create and Edit Deployment Descriptors"	Optionally create the EJB-specific deployment descriptors, although this step is no longer required when using the EJB 3.1 programming model.	<ul style="list-style-type: none"> ■ <code>ejb-jar.xml</code>, ■ <code>weblogic-<i>ejb-jar.xml</i></code>, which contains elements that control WebLogic Server-specific features.
10	Section 4.11, "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.
11	Section 4.12, "Deploying EJBs"	Target the archive or application directory to desired Managed Server, or a WebLogic Server cluster, in accordance with selected staging mode.	Deployed EJBs are ready to service invocations.

4.2 Create a Source Directory

Create a source directory where you will assemble the EJB 3.1 module.

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 3.1 as an enterprise application archive (EAR), see "Overview of the Split Development Directory Environment" in *Developing Applications for Oracle WebLogic Server*.

4.2.1 Directory Structure for Packaging a JAR

If you prefer to package and deploy your EJBs in a JAR file, create a directory for your class files. If you are also using the EJB deployment descriptor (which is optional but supported in the EJB 3.1 programming model), you can package it as `META-INF/ejb-jar.xml`.

Example 4-1 Directory Structure for Packaging a JAR

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

For more information see, [Section 4.11.1, "Packaging EJBs in a JAR."](#)

4.2.2 Directory Structure for Packaging a WAR

EJBs can also be packaged directly in a web application module (WAR) by putting the EJB classes in a subdirectory named `WEB-INF/classes` or in a JAR file within `WEB-INF/lib` directory. Optionally, if you are also using the EJB deployment descriptor, you can package it as `WEB-INF/ejb-jar.xml`.

Example 4-2 Directory Structure for Packaging a WAR

```
myEJBwar/  
  WEB-INF/  
    ejb-jar.xml  
    weblogic.xml  
    weblogic-ejb-jar.xml  
  /classes  
    foo.class  
    fooServlet.class  
    fooBean.class
```

Note: EJB 2.1 Entity Beans and EJB 1.1 Entity Beans are not supported within WAR files. These component types must only be packaged in a stand-alone `ejb-jar` file or an `ejb-jar` file packaged within an EAR file.

For more information see, [Section 4.11.2, "Packaging an EJB In a WAR."](#)

4.3 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 3.1 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 Class."](#)

4.4 Program the EJB Interface

Clients access enterprise beans either through a no-interface view or through a business interface.

4.4.1 Accessing EJBs Using the No-Interface Client View

The EJB 3.1 No-interface local client view type simplifies EJB development by providing local session bean access without requiring a separate local business interface, allowing components to have EJB bean class instances directly injected.

The no-interface view has the same behavior as the EJB 3.0 local view. For example, it supports features such as pass-by-reference calling semantics and transaction, and security propagation. However, a no-interface view does not require a separate interface. That is, all public methods of the bean class are automatically exposed to the caller. By default, any session bean that has an empty `implements` clause and does not define any other local or remote client views, exposes a no-interface client view.

There is a code example of a no-interface client view in [Section 3.1.1, "Example of a Simple No-interface Stateless EJB."](#) There is also an example of using the no-interface client view bundled in the WebLogic Server distribution kit. See [Section 3.2.4, "EJB 3.1: Example of Simplified No-interface Programming and Packaging in a WAR File."](#)

For more detailed information about the implementing the no-interface client view, see "Accessing Local Enterprise Beans Using the No-Interface View" in the "Enterprise Beans" chapter of the Java EE 6 Tutorial at <http://download.oracle.com/javaee/6/tutorial/doc/gipjf.html>

4.4.2 Accessing EJBs Using the Business Interface

The EJB 3.1 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 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 3.1 business interface is that it must *not* extend `javax.ejb.EJBObject` or `javax.ejb.EJBLocalObject`, as required in EJB 2.x.

See [Section 3.1.2, "Example of a Simple Business Interface Stateless EJB"](#) and [Section 3.1.3, "Example of a Simple Stateful EJB"](#) for examples of business interfaces implemented by stateless and stateful session beans.

For additional details and examples of specifying the business interface, see [Chapter 5.3.2, "Specifying the Business and Other Interfaces."](#)

4.4.2.1 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 JAX-WS Web Services for Oracle WebLogic Server*.

4.4.2.2 Using Generics in EJBs

The EJB 3.1 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) {
    }
}
```

4.4.2.3 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.

4.5 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 [Section 5.3.6, "Specifying Interceptors for Business Methods or Life Cycle Callback Events"](#) for information on programming the bean class to use interceptors.

4.6 Optionally Program the EJB Timer Service

WebLogic Server supports the EJB timer service defined in the EJB 3.1 Specification. As per the EJB 3.1 "specification, the EJB Timer Service is a "container-managed service that allows callbacks to be scheduled for time-based events. The container provides a reliable and transactional notification service for timed events. Timer notifications may be scheduled to occur according to a calendar-based schedule, at a specific time, after a specific elapsed duration, or at specific recurring intervals."

The Timer Service is implemented by the EJB container. An enterprise bean accesses this service by means of dependency injection, through the `EJBContext` interface, or through lookup in the JNDI namespace.

The Timer Service is intended to be used as a coarse-grained timer service. Rather than having a large number of timer objects performing the same task on a unique set of data, Oracle recommends using a small number of timers that perform bulk tasks on the data. For example, assume you have an EJB that represents an employee's expense report. Each expense report must be approved by a manager before it can be processed. You could use one EJB timer to periodically inspect all pending expense reports and send an email to the corresponding manager to remind them to either approve or reject the reports that are waiting for their approval.

4.6.1 Overview of the Timer Service

The timer service provides methods for the programmatic creation and cancellation of timers, as well as for locating the timers that are associated with a bean. Timers can also be created automatically by the container at deployment time based on metadata in the bean class or in the deployment descriptor. Timer objects can be created for stateless session beans, singleton session beans, message-driven beans, and 2.1 entity beans. Timers cannot be created for stateful session beans.

Note: The calendar-based timer, automatically-created timers, and non-persistent timer functionality is not supported for 2.1 Entity beans.

A timer is created to schedule timed callbacks. The bean class of an enterprise bean that uses the timer service must provide one or more timeout callback methods, as follows:

- **Programmatic Timers** – For programmatically-created timers, this method may be a method that is annotated with the `Timeout` annotation, or the bean may implement the `javax.ejb.TimerObject` interface. The `javax.ejb.TimerObject` interface has a single method, the timer callback method `ejbTimeout`.
- **Automatic Timers** – For automatically-created timers, the timeout method may be a method that is annotated with the `Schedule` annotation.
- **2.1 Entity Bean Timers** – A timer that is created for a 2.1 entity bean is associated with the entity bean's identity. The timeout callback method invocation for a timer that is created for a stateless session bean or a message-driven bean may be called on any bean instance in the pooled state.

4.6.2 Calendar-based EJB Timers

The EJB 3.1 Timer Service supports calendar-based EJB Timer expressions. The scheduling functionality takes the form of CRON-styled schedule definitions that can

be placed on EJB methods, in order to have the methods be automatically invoked according to the defined schedule.

Note: Calendar-based timers are not supported for EJB 2.x entity beans.

Calendar-based timers can be created programmatically using the two methods in the `javax.ejb.TimerService` that accept a `javax.ejb.ScheduleExpression` as an argument. The `ScheduleExpression` is constructed and populated prior to creating the `Timer`. For creating automatic calendar-base timers, the `javax.ejb.Schedule` annotation (and its corresponding `ejb-jar.xml` element) contains a number of attributes that allow for calendar-based timer expressions to be configured.

For detailed information about the seven attributes in a calendar-based time expression, see "Section 18.2.1, Calendar Based Timer Expressions" in the Enterprise JavaBeans 3.1 Specification (JSR-318) at <http://jcp.org/en/jsr/summary?id=318>.

4.6.3 Automatically-created EJB Timers

The EJB 3.1 Timer Service supports the automatic creation of a timer based on metadata in the bean class or deployment descriptor. This allows the bean developer to schedule a timer without relying on a bean invocation to programmatically invoke one of the Timer Service timer creation methods. Automatically created timers are created by the container as a result of application deployment.

The `javax.ejb.Schedule` annotation can be used to automatically create a timer with a particular timeout schedule. This annotation is applied to a method of a bean class (or super-class) that should receive the timer callbacks associated with that schedule. Multiple automatic timers can be applied to a single timeout callback method using the `javax.ejb.Schedules` annotation.

When the clustered EJB Timer implementation is used, each `Schedule` annotation corresponds to a single persistent timer, regardless of the number of servers across which the EJB is deployed.

4.6.4 Non-persistent Timers

By default, EJB timers are persistent. A non-persistent timer is a timer whose lifetime is tied to the JVM in which it is created. A non-persistent timer is considered canceled in the event of application shutdown, container crash, or a failure/shutdown of the JVM on which the timer was started.

Note: Non-persistent timers are not supported for EJB 2.x Entity Beans.

Non-persistent timers can be created programmatically or automatically (using `@Schedule` or the deployment descriptor).

Automatic non-persistent timers can be specified by setting the `persistent` attribute of the `@Schedule` annotation to `false`. For automatic non-persistent timers, the container creates a new non-persistent timer during application initialization for each JVM across which the container is distributed.

4.6.5 Clustered Versus Local EJB Timer Services

You can configure two types of EJB timer services: clustered or local.

4.6.5.1 Clustered EJB Timer Services

Clustered EJB timer services provide the following advantages:

- Better visibility.

Timers are accessible from any node in a cluster. For example, the `javax.ejb.TimerService.getTimers()` method returns a complete list of all stateless session or message-driven bean timers in a cluster that were created for the EJB. If you pass the primary key of the entity bean to the `getTimers()` method, a list of timers for that entity bean are returned.

- Automatic load balancing and failover.

Clustered EJB timer services take advantage of the load balancing and failover capabilities of the Job Scheduler.

For information about the configuring a clustered EJB timer service, see [Section 4.6.6, "Configuring Clustered EJB Timers."](#)

4.6.5.2 Local EJB Timer Services

Local EJB timer services execute only on the server on which they are created and are visible only to the beans on that server. With a local EJB timer service, you do not have to configure a cluster, database, JDBC data source, or leasing service, as you do for clustered EJB timer services.

You cannot migrate a local EJB timer object from one server to another; timer objects can only be migrated as part of an entire server. If a server that contains EJB timers goes down for any reason, you must restart the server or migrate the entire server in order for the timers to execute.

Caution: In a clustered environment, the local timer implementation has severe limitations; therefore Oracle recommends not using local timers in a clustered environment. Instead, use the clustered timer implementation in a clustered environment. A deployment-time warning will be thrown when a local timer implementation is configured to be used in a clustered environment.

4.6.6 Configuring Clustered EJB Timers

Note: To review the advantages of using clustered EJB timers, see [Section 4.6.5, "Clustered Versus Local EJB Timer Services."](#)

To configure the clustering of EJB timers, perform the following steps:

1. Ensure that you have configured the following:
 - A clustered domain. For more information, see "Setting up WebLogic Clusters" in *Using Clusters for Oracle WebLogic Server*.
 - Features of the Job Scheduler, including:
 - * HA database, such as Oracle, DB2, Informix, MySQL, Sybase, or MSSQL.

- * JDBC data source that is mapped to the HA database using the `<data-source-for-job-scheduler>` element in the `config.xml` file.
- * Leasing service. By default, database leasing will be used and the JDBC data source defined by the `<data-source-for-job-scheduler>` element in the `config.xml` file will be used.

For more information about configuring the Job Scheduler, see "The Timer and Work Manager API" in *Timer and Work Manager API (Common) Programmer's Guide for Oracle WebLogic Server*.

2. To enable the clustered EJB timer service, set the `timer-implementation` element in the `weblogic-ejb-jar.xml` deployment descriptor to `Clustered`:

```
<timer-implementation>Clustered</timer-implementation>
```

For more information, see the "timer-implementation" element description in the "weblogic-ejb-jar.xml Deployment Descriptor Reference" in *Programming Enterprise JavaBeans, Version 2.1, for Oracle WebLogic Server*.

Please note the following changes in the behavior of the clustered EJB timer service:

- The `weblogic.ejb.WLTimer*` interfaces are not supported with clustered EJB timer services.
- When creating a new clustered EJB timer using the `createTimer()` method, you may notice a delay in timeout execution during the initial setup of the timer.
- The Job Scheduler provides an "at least once" execution guarantee. When a clustered EJB timer expires, the database is not updated until the timer listener callback method completes. If the server were to crash before the database is updated, the timer expiration would be executed twice.
- Timer configuration options related to the actions to take in the event of failure are not valid for the clustered EJB timer service. These configuration options include: retry delay, maximum number of retry attempts, maximum number of time-outs, and time-out failure actions.
- The Job Scheduler queries the database every 30 seconds to identify timers that are due to expire. Execution may be delayed for timers with an interval duration less than 30 seconds.
- Only transactional timers will be retried in the event of failure.
- Fixed rate scheduling of timer execution is not supported.

4.6.7 Using Java Programming Interfaces to Program Timer Objects

This section summarizes the Java programming interfaces defined in the EJB 3.1 Specification that you can use to program timers. For detailed information on these interfaces, refer to the EJB 3.1 Specification. This section also provides details about the WebLogic Server-specific timer-related interfaces.

4.6.7.1 EJB 3.1 Timer-related Programming Interfaces

EJB 3.1 interfaces you can use to program timers are described in the following table.

Table 4–2 EJB 3.1 Timer-related Programming Interfaces

Programming Interface	Description
<code>javax.ejb.ScheduleExpression</code>	Create calendar-based EJB Timer expressions.
<code>javax.ejb.Schedule</code>	Automatically create a timer with a particular timeout schedule. Multiple automatic timers can be applied to a single timeout callback method using the <code>javax.ejb.Schedules</code> annotation.
<code>javax.ejb.TimerObject</code>	Implement for the enterprise bean class of a bean that will be registered with the timer service for timer callbacks. This interface has a single method, <code>ejbTimeout</code> .
<code>EJBContext</code>	Access the timer service using the <code>getTimerService</code> method.
<code>javax.ejb.TimerService</code>	Create new EJB timers or access existing EJB timers for the EJB.
<code>javax.ejb.Timer</code>	Access information about a particular EJB timer.
<code>javax.ejb.TimerHandle</code>	Define a serializable timer handle that can be persisted. Since timers are local objects, a <code>TimerHandle</code> must not be passed through a bean's remote interface or Web service interface.

For more information on EJB 2.1 timer-related programming interfaces, see the EJB 2.1 Specification.

4.6.7.2 WebLogic Server-specific Timer-related Programming Interfaces

WebLogic Server-specific interfaces you can use to program timers include:

- `weblogic.management.runtime.EJBTimerRuntimeMBean`—provides runtime information and administrative functionality for timers from a particular `EJBHome`. The `weblogic.management.runtime.EJBTimerRuntimeMBean` interface is shown in [Example 4–3](#).

Example 4–3 `weblogic.management.runtime.EJBTimerRuntimeMBean` Interface

```
public interface weblogic.management.runtime.EJBTimerRuntimeMBean {
    public int getTimeoutCount(); // get the number of successful timeout
    notifications that have been made
    public int getActiveTimerCount(); // get the number of active timers for this
    EJBHome
    public int getCancelledTimerCount(); // get the number of timers that have been
    cancelled for this EJBHome
    public int getDisabledTimerCount(); // get the number of timers temporarily
    disabled for this EJBHome
    public void activateDisabledTimers(); // activate any temporarily disabled
    timers
}
```

- `weblogic.ejb.WLTimerService` interface—extends the `javax.ejb.TimerService` interface to allow users to specify WebLogic Server-specific configuration information for a timer. The `weblogic.ejb.WLTimerService` interface is shown in [Example 4–4](#); for information on the `javax.ejb.TimerService`, see the EJB 2.1 Specification.

Note: The `weblogic.ejb.WLTimerService` interface is not supported by the clustered EJB timer service, as described in [Section 4.6.6, "Configuring Clustered EJB Timers."](#)

Example 4–4 *weblogic.ejb.WLTimerService Interface*

```
public interface WLTimerService extends TimerService {
    public Timer createTimer(Date initial, long duration, Serializable info,
        WLTimerInfo wlTimerInfo)
        throws IllegalArgumentException, IllegalStateException, EJBException;
    public Timer createTimer(Date expiration, Serializable info,
        WLTimerInfo wlTimerInfo)
        throws IllegalArgumentException, IllegalStateException, EJBException;
    public Timer createTimer(long initial, long duration, Serializable info
        WLTimerInfo wlTimerInfo)
        throws IllegalArgumentException, IllegalStateException, EJBException;
    public Timer createTimer(long duration, Serializable info,
        WLTimerInfo wlTimerInfo)
        throws IllegalArgumentException, IllegalStateException, EJBException;
}
```

- `weblogic.ejb.WLTimerInfo` interface—used in the `weblogic.ejb.WLTimerService` interface to pass WebLogic Server-specific configuration information for a timer. The `weblogic.ejb.WLTimerInfo` method is shown in [Example 4–5](#).

Note: The `weblogic.ejb.WLTimerService` interface is not supported by the clustered EJB timer service, as described in [Section 4.6.6, "Configuring Clustered EJB Timers."](#)

Example 4–5 *weblogic.ejb.WLTimerInfo Interface*

```
public final interface WLTimerInfo {
    public static int REMOVE_TIMER_ACTION = 1;
    public static int DISABLE_TIMER_ACTION = 2;
    public static int SKIP_TIMEOUT_ACTION = 3;
    /**
     * Sets the maximum number of retry attempts that will be
     * performed for this timer. If all retry attempts
     * are unsuccessful, the timeout failure action will
     * be executed.
     */
    public void setMaxRetryAttempts(int retries);
    public int getMaxRetryAttempts();
    /**
     * Sets the number of milliseconds that should elapse
     * before any retry attempts are made.
     */
    public void setRetryDelay(long millis);
    public long getRetryDelay();
    /**
     * Sets the maximum number of timeouts that can occur
     * for this timer. After the specified number of
     * timeouts have occurred successfully, the timer
     * will be removed.
     */
    public void setMaxTimeouts(int max);
}
```

```

public int getMaxTimeouts();
/**
 * Sets the action the container will take when ejbTimeout
 * and all retry attempts fail. The REMOVE_TIMER_ACTION,
 * DISABLE_TIMER_ACTION, and SKIP_TIMEOUT_ACTION fields
 * of this interface define the possible values.
 */
public void setTimeoutFailureAction(int action);
public int getTimeoutFailureAction();
}

```

- `weblogic.ejb.WLTimer` interface—extends the `javax.ejb.Timer` interface to provide additional information about the current state of the timer. The `weblogic.ejb.WLTimer` interface is shown in [Example 4-6](#).

Note: The `weblogic.ejb.WLTimerService` interface is not supported by the clustered EJB timer service, as described in [Section 4.6.6, "Configuring Clustered EJB Timers."](#)

Example 4-6 *weblogic.ejb.WLTimer Interface*

```

public interface WLTimer extends Timer {
    public int getRetryAttemptCount();
    public int getMaximumRetryAttempts();
    public int getCompletedTimeoutCount();
}

```

4.7 Programming Access to EJB Clients

This section provides some guidelines in determining the client view to provide for accessing an enterprise bean.

4.7.1 Remote Clients

As stated in the EJB 3.1 specification, a remote client accesses a session bean through the bean's remote business interface. For a session bean client and component written to the EJB 2.1 and earlier APIs, the remote client accesses the session bean through the session bean's remote home and remote component interfaces.

Note: The EJB 2.1 and earlier API required that a remote client access the stateful or stateless session bean by means of the session bean's remote home and remote component interfaces. These interfaces remain available for use with EJB 3.x, and are described in "Create EJB Classes and Interfaces" in *Programming Enterprise JavaBeans, Version 2.1, for Oracle WebLogic Server*.

The remote client view of an enterprise bean is location independent. A client running in the same JVM as a bean instance uses the same API to access the bean as a client running in a different JVM on the same or different machine.

4.7.2 Local Clients

As stated in the EJB 3.1 specification, a local client accesses a session bean through the bean's local business interface or through a no-interface client view representing all the

public methods of the bean class. For a session bean or entity bean client and component written to the EJB 2.1 and earlier APIs, the local client accesses the enterprise bean through the bean's local home and local component interfaces. The container object that implements a local business interface or the no-interface local view is a local Java object.

A local client is a client that is collocated in the same application with the session bean that provides the local client view and which may be tightly coupled to the bean. A local client of a session bean may be another enterprise bean or a Web component. Access to an enterprise bean through the local client view requires the collocation in the same application of both the local client and the enterprise bean that provides the local client view. The local client view therefore does not provide the location transparency provided by the remote client view.

4.7.3 Looking Up EJBs From Clients

The client of an enterprise bean obtains a reference to an instance of an enterprise bean through either dependency injection, using Java programming language annotations, or JNDI lookup, using the Java Naming and Directory Interface syntax to find the enterprise bean instance.

Note: For instructions on how clients can look up 2.x or earlier enterprise beans using EJB Links, see "Using EJB Links" in *Programming Enterprise JavaBeans, Version 2.1, for Oracle WebLogic Server*.

4.7.3.1 Using Dependency Injection

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.

For detailed information, see [Section 5.3.4, "Injecting Resource Dependency into a Variable or Setter Method."](#)

4.7.3.2 Using the JNDI Portable Syntax

The Portable Global JNDI naming option in EJB 3.1 provides a number of common, well-known namespaces in which EJB components can be registered and looked up from using the patterns listed in this section. This standardizes how and where EJB components are registered in JNDI, and how they can be looked up and used by applications.

Three JNDI namespaces are used for portable JNDI lookups: `java:global`, `java:module`, and `java:app`.

- The `java:global` JNDI namespace is the portable way of finding remote enterprise beans using JNDI lookups. JNDI addresses are of the following form:

```
java:global[/application name]/module name/enterprise bean
name[/interface name]
```

Application name and module name default to the name of the application and module minus the file extension. Application names are required only if the application is packaged within an EAR. The interface name is required only if the enterprise bean implements more than one business interface.

- The `java:module` namespace is used to look up local enterprise beans within the same module. JNDI addresses using the `java:module` namespace are of the following form:

```
java:module/enterprise bean name[/interface name]
```

The interface name is required only if the enterprise bean implements more than one business interface.

- The `java:app` namespace is used to look up local enterprise beans packaged within the same application. That is, the enterprise bean is packaged within an EAR file containing multiple Java EE modules. JNDI addresses using the `java:app` namespace are of the following form:

```
java:app[/module name]/enterprise bean name[/interface name]
```

The module name is optional. The interface name is required only if the enterprise bean implements more than one business interface.

For example, if an enterprise bean, `MyBean`, is packaged within the Web application archive `myApp.war`, the default module name is `myApp`. (In this example, the module name could be explicitly configured in the `web.xml` file.) The portable JNDI name is `java:module/MyBean`. An equivalent JNDI name using the `java:global` namespace is `java:global/myApp/MyBean`.

4.7.3.3 Customizing JNDI Names

Though global JNDI bindings are registered by default, you can also customize the JNDI names of your EJB client view bindings by using the `weblogic.javaee.JNDIName` and `weblogic.javaee.JNDINames` annotations. For more information, see [Section A.7.7, "weblogic.javaee.JNDIName"](#) and [Section A.7.8, "weblogic.javaee.JNDINames."](#)

For EJBs using deployment descriptors, you can specify the custom JNDI bindings in the `weblogic-ejb-jar.xml` deployment descriptor by using the `jndi-binding` element. For more information, see "EJB Deployment Descriptors" in *Programming Enterprise JavaBeans, Version 2.1, for Oracle WebLogic Server*.

4.7.4 Configuring EJBs to Send Requests to a URL

To enable an EJB to open an `URLConnection` to an external HTTP server using the `java.net.URL` resource manager connection factory type, specify the URL, or specify an object bound in the JNDI tree that maps to a URL, using either the `@Resource` annotation in the bean class, or if using deployment descriptors, by using the `resource-ref` element in `ejb-jar.xml` and the `res-ref-name` element in `weblogic-ejb-jar.xml`.

4.7.5 Specifying an HTTP Resource by URL

When using annotations to specify the URL to which an EJB sends requests:

1. Annotate a URL field in your bean class with `@Resource`.

2. Specify the URL value using the look-up element of `@Resource`.

When using deployment descriptors to specify the URL to which an EJB sends requests:

1. In `ejb-jar.xml`, specify the URL in the `<jndi-name>` element of the `resource-ref` element.
2. In `weblogic-ejb-jar.xml`, specify the URL in the `<jndi-name>` element of the `resource-description` element:

```
<resource-description>
  <res-ref-name>url/MyURL</res-ref-name>
  <jndi-name>http://www.rediff.com/</jndi-name>
</resource-description>
```

WebLogic Server creates a URL object with the `jndi-name` provided and binds the object to the `java:comp/env`.

4.7.6 Specifying an HTTP Resource by Its JNDI Name

When using annotations to specify an object that is bound in JNDI and maps to a URL, instead of specifying a URL:

1. Annotate a URL field in your bean class with `@Resource`.
2. Specify the name by which the URL is bound in JNDI using the look-up element of `@Resource`.

When using deployment descriptors to specify an object that is bound in JNDI and maps to a URL, instead of specifying a URL:

1. In `ejb-jar.xml`, specify the name by which the URL is bound in JNDI in the `<jndi-name>` element of the `resource-ref` element.
2. In `weblogic-ejb-jar.xml`, specify the name by which the URL is bound in JNDI in the `<jndi-name>` element of the `resource-description` element:

```
<resource-description>
  <res-ref-name>url/MyURL1</res-ref-name>
  <jndi-name>firstName</jndi-name>
</resource-description>
```

where `firstName` is the object bound to the JNDI tree that maps to the URL. This binding could be done in a startup class. When `jndi-name` is not a valid URL, WebLogic Server treats it as an object that maps to a URL and is already bound in the JNDI tree, and binds a `LinkRef` with that `jndi-name`.

4.7.7 Accessing HTTP Resources from Bean Code

Regardless of how you specified an HTTP resource—by its URL or a JNDI name that maps to the URL—you can access it from EJB code in this way:

```
URL url = (URL) context.lookup("java:comp/env/url/MyURL");
connection = (URLConnection)url.openConnection();
```

4.7.8 Configuring Network Communications for an EJB

You can control the attributes of the network connection an EJB uses for communications by configuring a custom network channel and assigning it to the EJB. For information about WebLogic Server network channels and associated

configuration instructions see "Configure Network Resources" in *Configuring Server Environments for Oracle WebLogic Server*. After you configure a custom channel, assign it to an EJB using the `network-access-point` element in `weblogic-ejb-jar.xml`.

4.8 Programming and Configuring Transactions

The following sections contain guidelines for programming transactions.

4.8.1 Programming Container-Managed Transactions

Container-managed transactions are simpler to program than bean-managed transactions, because they leave the job of demarcation—starting and stopping the transaction—to the EJB container. You configure the desired transaction behaviors using EJB annotations `javax.ejb.TransactionAttribute` or by using EJB deployment descriptors `ejb-jar.xml` and `weblogic-ejb-jar.xml`.

- For more information about using EJB annotations to specify container-managed transactions in a bean file, see [Chapter 5.3.9, "Specifying Transaction Management and Attributes."](#)
- For more information about using EJB deployment descriptors to specify container-managed transactions, see "Container-Managed Transaction Elements" in *Programming Enterprise JavaBeans, Version 2.1, for Oracle WebLogic Server*.

Key programming guidelines for container-managed transactions include:

- Preserve transaction boundaries—Do not invoke methods that interfere with the transaction boundaries set by the container. Do not use:
 - The `commit`, `setAutoCommit`, and `rollback` methods of `java.sql.Connection`
 - The `getUserTransaction` method of `javax.ejb.EJBContext`
 - Any method of `javax.transaction.UserTransaction`
- Roll back transactions explicitly—To cause the container to roll back a container-managed transaction explicitly, invoke the `setRollbackOnly` method of the `EJBContext` interface. (If the bean throws a non-application exception, typically an `EJBException`, the rollback is automatic.)
- Avoid serialization problems—Many data stores provide limited support for detecting serialization problems, even for a single user connection. In such cases, even with `transaction-isolation` in `weblogic-ejb-jar.xml` set to `TransactionSerializable`, exceptions or rollbacks in the EJB client might occur if contention occurs between clients for the same rows. To avoid such exceptions, you can:
 - Include code in your client application to catch SQL exceptions, and resolve them appropriately; for example, by restarting the transaction.
 - For Oracle databases, use the transaction isolation settings described in "isolation-level" in the "weblogic-ejb-jar.xml Deployment Descriptor Reference" appendix in *Programming Enterprise JavaBeans, Version 2.1, for Oracle WebLogic Server*.

4.8.2 Configuring Automatic Retry of Container-Managed Transactions

In Oracle WebLogic Server, you can specify that, if a business method that has started a transaction fails because of a transaction rollback that is not related to a system exception, the EJB container will start a new transaction and retry the failed method up to a specified number of times. If the method fails for the specified number of retry attempts, the EJB container throws an exception.

Note: The EJB container does not retry any transactions that fail because of system exception-based errors.

To configure automatic retry of container-managed transactions:

1. Make sure your bean is a container-managed session or entity bean.

You can configure automatic retry of container-managed transactions for container-managed session and entity beans only. You cannot configure automatic retry of container-managed transactions for message-driven beans because MDBs do not acknowledge receipt of a message they are processing when the transaction that brackets the receipt of the message is rolled back; messages are automatically retried until they are acknowledged. You also cannot configure automatic retry of container-managed transactions for timer beans because, when a timer bean's `ejbTimeout` method starts and is rolled back, the timeout is always retried.

2. Make sure the business methods for which you want to configure automatic retry of transactions are defined in the bean's remote or local interface or as home methods (local home business logic that is not specific to a particular bean instance) in the home interface; the methods must have one of the following container-managed transaction attributes:
 - `RequiresNew`. If a method's transaction attribute (`trans-attribute` element in `ejb-jar.xml`) is `RequiresNew`, a new transaction is always started prior to the invocation of the method and, if configured, automatic retry of transactions occurs if the transaction fails.
 - `Required`. If a method's transaction attribute (`trans-attribute` element in `ejb-jar.xml`) is `Required`, the method is retried with a new transaction only if the failed transaction was begun on behalf of the method.

For more information on:

- Programming interfaces, see [Section 4.7, "Programming Access to EJB Clients."](#)
 - The `trans-attribute` element in `ejb-jar.xml`, see "Container-Managed Transaction Elements" in *Programming Enterprise JavaBeans, Version 2.1, for Oracle WebLogic Server*, which provides detailed information about creating and editing EJB deployment descriptors.
3. Make sure the methods for which you want to enable automatic retry of transactions are safe to be re-invoked. A retry of a failed method must yield results that are identical to the results the previous attempt, had it been successful, would have yielded. In particular:
 - If invoking a method initiates a call chain, it must be safe to re-invoke the entire call chain when the method is retried.
 - All of the method's parameters must be safe for reuse; when a method is retried, it is retried with the same parameters that were used to invoke the failed attempt. In general, parameters that are primitives, immutable objects, or are references to read-only objects are safe for reuse. If a parameter is a

reference to an object that is to be modified by the method, re-invoking the method must not negatively affect the result of the method call.

- If the bean that contains the method that is being retried is a stateful session bean, the bean's conversational state must be safe to re-invoke. Since a stateful session bean's state is not transactional and is not restored during a transaction rollback, in order to use the automatic retry of transactions feature, you must first be sure the bean's state is still valid after a rollback.
4. Specify the methods for which you want the EJB container to automatically retry transactions and the number of retry attempts you want the EJB container to make in the `retry-methods-on-rollback` element in `weblogic-ejb-jar.xml`.

The `retry-count` sub-element to `retry-methods-on-rollback` can also be modified via the Administration Console.

4.8.3 Programming Bean-Managed Transactions

This section contains programming considerations for bean-managed transactions.

- Demarcate transaction boundaries—To define transaction boundaries in EJB or client code, you must obtain a `UserTransaction` object and begin a transaction before you obtain a Java Transaction Service (JTS) or JDBC database connection. To obtain the `UserTransaction` object, use this command:

```
ctx.lookup("javax.transaction.UserTransaction");
```

After obtaining the `UserTransaction` object, specify transaction boundaries with `tx.begin()`, `tx.commit()`, `tx.rollback()`.

If you start a transaction after obtaining a database connection, the connection has no relationship to the new transaction, and there are no semantics to "enlist" the connection in a subsequent transaction context. If a JTS connection is not associated with a transaction context, it operates similarly to a standard JDBC connection that has `autocommit` equal to `true`, and updates are automatically committed to the data store.

Once you create a database connection within a transaction context, that connection is reserved until the transaction commits or rolls back. To optimize performance and throughput, ensure that transactions complete quickly, so that the database connection can be released and made available to other client requests.

Note: You can associate only a single database connection with an active transaction context.

- Setting transaction isolation level—For bean-managed transactions, you define isolation level in the bean code. Allowable isolation levels are defined on the `java.sql.Connection` interface. For information on isolation level behaviors, see "isolation-level" in the "weblogic-ejb-jar.xml Deployment Descriptor Reference" appendix in *Programming Enterprise JavaBeans, Version 2.1, for Oracle WebLogic Server*.

See [Example 4-7](#) for a code sample.

Example 4-7 Setting Transaction Isolation Level in BMT

```
import javax.transaction.Transaction;
import java.sql.Connection
```

```

import weblogic.transaction.TxHelper;
import weblogic.transaction.Transaction;
import weblogic.transaction.TxConstants;
User Transaction tx = (UserTransaction)
ctx.lookup("javax.transaction.UserTransaction");
//Begin user transaction
    tx.begin();
//Set transaction isolation level to TransactionReadCommitted
Transaction tx = TxHelper.getTransaction();
    tx.setProperty (TxConstants.ISOLATION_LEVEL, new Integer
        (Connection.TransactionReadCommitted));
//perform transaction work
    tx.commit();

```

- Avoid restricted methods—Do not invoke the `getRollbackOnly` and `setRollbackOnly` methods of the `EJBContext` interface in bean-managed transactions. These methods should be used only in container-managed transactions. For bean-managed transactions, invoke the `getStatus` and `rollback` methods of the `UserTransaction` interface.
- Use one connection per active transaction context—You can associate only a single database connection with an active transaction context.

4.8.4 Programming Transactions That Are Distributed Across EJBs

This section describes two approaches for distributing a transaction across multiple beans, which may reside on multiple server instances.

4.8.4.1 Calling multiple EJBs from a client's transaction context

The code fragment below is from a client application that obtains a `UserTransaction` object and uses it to begin and commit a transaction. The client invokes two EJBs within the context of the transaction.

```

import javax.transaction.*;
...
u = (UserTransaction) jndiContext.lookup("javax.transaction.UserTransaction");
u.begin();
account1.withdraw(100);
account2.deposit(100);
u.commit();
...

```

The updates performed by the `account1` and `account2` beans occur within the context of a single `UserTransaction`. The EJBs commit or roll back together, as a logical unit, whether the beans reside on the same server instance, different server instances, or a WebLogic Server cluster.

All EJBs called from a single transaction context must support the client transaction—each beans' `trans-attribute` element in `ejb-jar.xml` must be set to `Required`, `Supports`, or `Mandatory`.

4.8.4.2 Using an EJB "Wrapper" to Encapsulate a Cross-EJB Transaction

You can use a "wrapper" EJB that encapsulates a transaction. The client calls the wrapper EJB to perform an action such as a bank transfer, and the wrapper starts a new transaction and invokes one or more EJBs to do the work of the transaction.

The wrapper EJB can explicitly obtain a transaction context before invoking other EJBs, or WebLogic Server can automatically create a new transaction context, if the

wrapper's `trans-attribute` element in `ejb-jar.xml` is set to `Required` or `RequiresNew`.

All EJBs invoked by the wrapper EJB must support the wrapper EJB's transaction context— their `trans-attribute` elements must be set to `Required`, `Supports`, or `Mandatory`.

4.9 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, typically using the standard Java compiler. The resulting class files can then be packaged into a target module for deployment. Typical tools to compile include:

- `javac`—The `javac` compiler provided with the Java SE SDK provides Java compilation capabilities. See <http://www.oracle.com/technetwork/java/javase/documentation/index.html/>.
- `weblogic.appc`—To reduce deployment time, use the `weblogic.appc` Java class (or its equivalent Ant task `wlappc`) to pre-compile a deployable archive file, (WAR, JAR, or EAR). Precompiling with `weblogic.appc` generates certain helper classes and performs validation checks to ensure your application is compliant with the current Java EE specifications. See "Building Modules and Applications Using `wlappc`" in *Developing Applications for Oracle WebLogic Server*.
- `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`" in *Developing Applications for Oracle WebLogic Server*.

4.10 Optionally Create and Edit Deployment Descriptors

An important aspect of the EJB 3.x programming model was 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.x fully supports the use of deployment descriptors, even though the standard Java EE 6 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 3.1 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 3.1 because currently they do not have any annotation equivalent.

The 2.x version of *Programming Enterprise JavaBeans, Version 2.1, for Oracle WebLogic Server* 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)
- "Edit Deployment Descriptors"
- "Deployment Descriptor Schema and Document Type Definitions Reference"
- "weblogic-ejb-jar.xml Deployment Descriptor Reference"
- "weblogic-cmp-jar.xml Deployment Descriptor Reference"

4.11 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 for Oracle WebLogic Server*.

However, EJB 3.1 simplifies packaging by providing the ability to place EJB components directly inside of Web application archive (WAR) files, removing the need to produce separate archives to store the Web and EJB components and combine them together in an enterprise application archive (EAR) file.

4.11.1 Packaging EJBs in a JAR

WebLogic Server supports the use of `ejb-client.jar` files for packaging the EJB classes that a programmatic client in a different application requires to access the EJB.

Specify the name of the client JAR in the `ejb-client-jar` element of the bean's `ejb-jar.xml` file. When you run the `appc` compiler, a JAR file with the classes required to access the EJB is generated.

Make the client JAR available to the remote client. For Web applications, put the `ejb-client.jar` in the `/lib` directory. For non-Web clients, include `ejb-client.jar` in the client's classpath.

Note: WebLogic Server classloading behavior varies, depending on whether the client is stand-alone. Stand-alone clients with access to the `ejb-client.jar` can load the necessary classes over the network. However, for security reasons, programmatic clients running in a server instance cannot load classes over the network.

4.11.2 Packaging an EJB In a WAR

EJB 3.1 has removed the restriction that enterprise bean classes must be packaged in an `ejb-jar` file. Therefore, EJB classes can be packaged directly inside a Web application archive (WAR) using the same packaging guidelines that apply to Web application

classes. Simply put your EJB classes in the `WEB-INF/classes` directory or in a JAR file within `WEB-INF/lib` directory. Optionally, if you are also using the EJB deployment descriptor, you can package it as `WEB-INF/ejb-jar.xml`. When you run the appc compiler, a WAR file with the classes required to access the EJB components is generated.

In a WAR file there is a single component naming environment shared between all the components (web, enterprise bean, etc.) defined by the module. Each enterprise bean defined by the WAR file shares this single component environment namespace with all other enterprise beans defined by the WAR file and with all other web components defined by the WAR file.

Enterprise beans (and any related classes) packaged in a WAR file have the same class loading requirements as other non-enterprise bean classes packaged in a WAR file. This means, for example, that a servlet packaged within a WAR file is guaranteed to have visibility to an enterprise bean component packaged within the same WAR file, and vice versa.

Caution: EJB 2.1 Entity Beans and EJB 1.1 Entity Beans are not supported within WAR files. These component types must only be packaged in a stand-alone `ejb-jar` file or an `ejb-jar` file packaged within an EAR file. Applications that violate this restriction will fail to deploy.

There is an example of using the simplified WAR packaging method bundled in the WebLogic Server distribution kit. See [Section 3.2.4, "EJB 3.1: Example of Simplified No-interface Programming and Packaging in a WAR File."](#)

4.12 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.

For general instructions on deploying WebLogic Server applications and modules, including EJBs, see *Deploying Applications to Oracle WebLogic Server*. For EJB-specific deployment issues and procedures, see [Section 6.3.1, "Deploying Standalone EJBs as Part of an Enterprise Application."](#) and [Section 6.3.2, "Deploying EJBs as Part of a Web Application."](#)

For more information about deploying an EJB created with the 2.x programming model, see "Deployment Guidelines For Enterprise JavaBeans" in the *Programming Enterprise JavaBeans, Version 2.1, for Oracle WebLogic Server* guide, which concentrates on the 2.x programming model.

Programming the Annotated EJB Class

The sections that follow describe how to program the annotated EJB 3.1 class file:

- [Section 5.1, "Overview of Metadata Annotations and EJB Bean Files"](#)
- [Section 5.2, "Programming the Bean File: Requirements and Changes From 2.x"](#)
- [Section 5.3, "Programming the Bean File"](#)
- [Section 5.4, "Complete List of Metadata Annotations By Function"](#)

5.1 Overview of Metadata Annotations and EJB Bean Files

The WebLogic Server EJB 3.1 programming model uses the Java EE 6 metadata annotations feature in which you create an annotated EJB 3.1 bean file, compile the class with the standard Java compiler, and the resulting class can then be packaged into a target module for deployment. At runtime, WebLogic Server parses the annotations and applies the required behavioral aspects to the bean file.

Tip: To reduce deployment time, you can also use the WebLogic compile tool `weblogic.appc` (or its Ant equivalent `wlappc`) to pre-compile a deployable archive file, (WAR, JAR, or EAR). Precompiling with `weblogic.appc` generates certain helper classes and performs validation checks to ensure your application is compliant.

The annotated 3.1 bean file is the core of your EJB. It contains the Java code that determines how your EJB behaves. The 3.1 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 6 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 [Section 5.4, "Complete List of Metadata Annotations By Function"](#) for a breakdown of the annotations you can specify in a bean file, by function. These annotations include those described by the following specifications:

- Enterprise JavaBeans 3.1 Specification (JSR-318) at <http://jcp.org/en/jsr/summary?id=318>
- Common Annotations for the Java Platform (JSR-250) at <http://www.jcp.org/en/jsr/detail?id=250>

See [Appendix A, "EJB 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 3.1, described in [Chapter 4, "Iterative Development of Enterprise JavaBeans."](#)

For more information on general EJB design and architecture, see the following documents:

- Enterprise JavaBeans web site at <http://www.oracle.com/technetwork/java/ejb-141389.html>
- Introducing the Java EE 6 Platform: Part 3 (EJB Technology, Even Easier to Use) at <http://www.oracle.com/technetwork/articles/javaee/javaee6overview-part3-139660.html#ejbeasy>
- The Java EE 6 Tutorial at <http://download.oracle.com/javaee/6/tutorial/doc/bnblr.html>

5.2 Programming the Bean File: Requirements and Changes From 2.x

The requirements for programming the 3.1 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 3.1.

See *Programming Enterprise JavaBeans, Version 2.1, for Oracle WebLogic Server* for detailed information about the mandatory and optional requirements for programming the bean class.

5.2.1 Bean Class Requirements and Changes From 2.x

The following bullets list the 3.1 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.Singleton`
 - `@javax.ejb.MessageDriven`

Note: Oracle Kodo JPA/JDO is deprecated in this release. However, if you still using Oracle Kodo, programming a 3.0 entity bean (`@javax.ejb.Entity`) is discussed in a separate document. See "Java Persistence API" in the Oracle Kodo documentation.

Customers are encouraged to use Oracle TopLink, which supports JPA 2.0. Kodo supports only JPA 1.0. For more information, see [Section 8.3, "Using Oracle TopLink in Oracle WebLogic Server."](#)

- If the bean is a session bean, the bean class can implement either:
 - The no-interface local client view type, which simplifies EJB development by providing local session bean access without requiring a separate local business interface.
 - 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.

5.2.2 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`.

5.3 Programming the Bean File

The sections that follow provide the recommended steps when programming the annotated EJB 3.1 class file.

- [Section 5.3.1, "Typical Steps When Programming the Bean File"](#)
- [Section 5.3.2, "Specifying the Business and Other Interfaces"](#)
- [Section 5.3.3, "Specifying the Bean Type \(Stateless, Singleton, Stateful, or Message-Driven\)"](#)
- [Section 5.3.4, "Injecting Resource Dependency into a Variable or Setter Method"](#)

- [Section 5.3.5, "Invoking a 3.0 Entity"](#)
- [Section 5.3.6, "Specifying Interceptors for Business Methods or Life Cycle Callback Events"](#)
- [Section 5.3.7, "Programming Application Exceptions"](#)
- [Section 5.3.8, "Securing Access to the EJB"](#)
- [Section 5.3.9, "Specifying Transaction Management and Attributes"](#)

5.3.1 Typical Steps When Programming the Bean File

The following procedure describes the typical basic steps when programming the 3.1 bean file for a EJB. The steps you follow depends, of course, on what your EJB does.

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

1. Import the EJB 3.1 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 3.1 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 interface that your EJB is going to implement, either a business interface or a no-interface view, as well as other standard interfaces. You can either explicitly implement the interface, or use an annotation to specify it.

See [Section 5.3.2, "Specifying the Business and Other Interfaces."](#)

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

See [Section 5.3.3, "Specifying the Bean Type \(Stateless, Singleton, Stateful, or Message-Driven\)."](#)

4. Optionally, use dependency injection to use external resources, such as another EJB or other Java EE 6 object.

See [Section 5.3.4, "Injecting Resource Dependency into a Variable or Setter Method."](#)

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

See [Section 5.3.5, "Invoking a 3.0 Entity."](#)

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 [Section 5.3.6, "Specifying Interceptors for Business Methods or Life Cycle Callback Events."](#)

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 [Section 5.3.7, "Programming Application Exceptions"](#) for EJB 3.1 specific information.

8. Optionally, specify the security roles that are allowed to invoke the EJB methods using the security-related metadata annotations.

See [Section 5.3.8, "Securing Access to the EJB."](#)

9. Optionally, change the default transaction configuration in which the EJB runs.

See [Section 5.3.9, "Specifying Transaction Management and Attributes."](#)

5.3.2 Specifying the Business and Other Interfaces

The EJB 3.x local or remote client of a session bean written to the EJB 3.x API accesses a session bean through its business interface. A local client may also access a session bean through a no-interface view that exposes all public methods of the bean class.

5.3.2.1 Specifying the 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 [Section 3.1.2, "Example of a Simple Business Interface Stateless EJB."](#)

5.3.2.2 Specifying the No-interface View

Client access to an enterprise bean that exposes a local, no-interface view is accomplished through either dependency injection or JNDI lookup.

- To obtain a reference to the no-interface view of an enterprise bean through dependency injection, use the `javax.ejb.EJB` annotation and specify the enterprise bean's implementation class:

```
@EJB
ExampleBean exampleBean;
```

- To obtain a reference to the no-interface view of an enterprise bean through JNDI lookup, use the `javax.naming.InitialContext` interface's `lookup` method:

```
ExampleBean exampleBean = (ExampleBean)
InitialContext.lookup("java:module/ExampleBean");
```

Clients do not use the new operator to obtain a new instance of an enterprise bean that uses a no-interface view.

There is an example of using the No-interface client view bundled in the WebLogic Server distribution kit. See [Section 3.2.4, "EJB 3.1: Example of Simplified No-interface Programming and Packaging in a WAR File."](#)

For more detailed information about the implementing the no-interface client view, see "Accessing Local Enterprise Beans Using the No-Interface View" in the "Enterprise Beans" chapter of the Java EE 6 Tutorial at

<http://download.oracle.com/javaee/6/tutorial/doc/gipjf.html>

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

There is only one required metadata annotation in a 3.1 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.Singleton`—Specifies that you are programming a singleton 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.

Note: Oracle Kodo JPA/JDO is deprecated in this release. However, if you still using Oracle Kodo, programming a 3.0 entity bean (`@javax.ejb.Entity`) is discussed in a separate document. See "Java Persistence API" in the Oracle Kodo documentation.

Customers are encouraged to use Oracle TopLink, which supports JPA 2.0. Kodo supports only JPA 1.0. For more information, see [Section 8.3, "Using Oracle TopLink in Oracle WebLogic Server."](#)

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 [Section 3.1.2, "Example of a Simple Business Interface Stateless EJB."](#)

5.3.4 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 [Section A.2.10, "javax.ejb.EJB"](#) and [Section A.5.3, "javax.annotation.Resource"](#) 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 `ServiceBean`, as shown in the `sayHelloFromAccountBean` method implementation in which the `sayHelloFromServiceBean` method is invoked.

5.3.5 Invoking a 3.0 Entity

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

Note: Oracle TopLink, a JPA 2.0 persistence provider, is now the default JPA provider, replacing Kodo, which was the default provider in previous releases. Any application that does not specify a JPA provider in `persistence.xml` will now use TopLink by default. Applications can continue to use Kodo (a JPA 1.0 provider) by explicitly specifying Kodo/OpenJPA as their persistence provider in `persistence.xml`. In addition, a Weblogic Server domain can be configured to use Kodo by default, if desired. For more information, see [Chapter 8, "Configuring the Persistence Provider in Oracle WebLogic Server."](#)

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:

- [Section 5.3.5.1, "Injecting Persistence Context Using Metadata Annotations"](#)
- [Section 5.3.5.2, "Finding an Entity Using the EntityManager API"](#)
- [Section 5.3.5.3, "Creating and Updating an Entity Using EntityManager"](#)

5.3.5.1 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;
    ...
}
```

5.3.5.2 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 [Chapter 8.3, "Using Oracle TopLink in Oracle WebLogic Server."](#) 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.

5.3.5.3 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 [Chapter 8.3, "Using Oracle TopLink in Oracle WebLogic Server."](#)

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.

5.3.6 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:

- [Section 5.3.6.1, "Specifying Business or Life Cycle Interceptors: Typical Steps"](#)
- [Section 5.3.6.2, "Programming the Interceptor Class"](#)

- [Section 5.3.6.3, "Programming Business Method Interceptor Methods"](#)
- [Section 5.3.6.4, "Programming Asynchronous Business Methods"](#)
- [Section 5.3.6.5, "Programming Life Cycle Callback Interceptor Methods"](#)
- [Section 5.3.6.6, "Specifying Default Interceptor Methods"](#)
- [Section 5.3.6.7, "Saving State Across Interceptors With the InvocationContext API"](#)

5.3.6.1 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 [Section 3.1.3, "Example of a Simple Stateful EJB"](#) for an example of specifying interceptors and [Section 3.1.4, "Example of an Interceptor Class"](#) 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 [Section 5.3.6.2, "Programming the Interceptor Class."](#)
 - 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 [Section 5.3.6.3, "Programming Business Method Interceptor Methods."](#)
4. In the bean class or interceptor class (wherever you are programming the interceptor methods), program life cycle callback interceptor methods, as described in [Section 5.3.6.3, "Programming Business Method Interceptor Methods."](#)
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 [Section 5.3.6.6, "Specifying Default Interceptor Methods."](#)

5.3.6.2 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 [Section 3.1.4, "Example of an Interceptor Class."](#)

5.3.6.3 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 [Section 5.3.6.7, "Saving State Across Interceptors With the InvocationContext API"](#) 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 [Section 3.1.4, "Example of an Interceptor Class."](#)

5.3.6.4 Programming Asynchronous Business Methods

Session beans can implement asynchronous methods, business methods where control is returned to the client by the enterprise bean container before the method is invoked on the session bean instance. Clients may then use the Java SE concurrency API to retrieve the result, cancel the invocation, and check for exceptions. Asynchronous methods are typically used for long-running operations, for processor-intensive tasks, for background tasks, to increase application throughput, or to improve application response time if the method invocation result isn't required immediately.

When a session bean client invokes a typical non-asynchronous business method, control is not returned to the client until the method has completed. Clients calling

asynchronous methods, however, immediately have control returned to them by the enterprise bean container. This allows the client to perform other tasks while the method invocation completes. If the method returns a result, the result is an implementation of the `java.util.concurrent.Future<V>` interface, where “V” is the result value type. The `Future<V>` interface defines methods the client may use to check if the computation is completed, wait for the invocation to complete, retrieve the final result, and cancel the invocation.

For detailed instructions on creating an asynchronous business method, see “Asynchronous Method Invocation” in the “Enterprise Beans” chapter of the Java EE 6 Tutorial at <http://download.oracle.com/javaee/6/tutorial/doc/gkkqg.html>.

5.3.6.5 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 [Section 5.3.6.7, "Saving State Across Interceptors With the InvocationContext API"](#) for additional information about the `InvocationContext` API.

See [Section A.2.18, "javax.ejb.PostActivate,"](#) [Section A.2.19, "javax.ejb.PrePassivate,"](#) [Section A.5.1, "javax.annotation.PostConstruct,"](#) and [Section A.5.2, "javax.annotation.PreDestroy"](#) for additional requirements when programming the life cycle interceptor class.

For an example, see [Section 3.1.4, "Example of an Interceptor Class."](#)

5.3.6.6 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/ebj-jar_3_0.xsd">
  ...

  <assembly-descriptor>
  ...
```

```

    <interceptor-binding>
      <ejb-name>*</ejb-name>
      <interceptor-class>org.mycompany.DefaultIC</interceptor-class>
    </interceptors>

  </assembly-descriptor>

</ejb-jar>

```

5.3.6.7 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 [Section 3.1.4, "Example of an Interceptor Class."](#)

5.3.7 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);}
}
```

5.3.8 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 3.1.

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" in the *Oracle WebLogic Server Administration Console Help*.

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.

5.3.9 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 [Section 3.1.3, "Example of a Simple Stateful EJB."](#)

5.4 Complete List of Metadata Annotations By Function

The tables in the sections that follow group the annotations based on what task they perform. [Appendix A, "EJB Metadata Annotations Reference,"](#) provides full reference information about the EJB 3.1 metadata annotations in alphabetical order.

- [Section 5.4.1, "Annotations to Specify the Bean Type"](#)
- [Section 5.4.2, "Annotations to Specify the Local or Remote Interfaces"](#)
- [Section 5.4.3, "Annotations to Support EJB 2.x Client View"](#)
- [Section 5.4.4, "Annotations to Invoke a 3.0 Entity Bean"](#)
- [Section 5.4.5, "Transaction-Related Annotations"](#)
- [Section 5.4.6, "Annotations to Specify Interceptors"](#)
- [Section 5.4.7, "Annotations to Specify Life Cycle Callbacks"](#)
- [Section 5.4.8, "Security-Related Annotations"](#)
- [Section 5.4.9, "Context Dependency Annotations"](#)
- [Section 5.4.10, "Timeout and Exceptions Annotations"](#)

5.4.1 Annotations to Specify the Bean Type

The following summarize the annotations used to specify the bean type.

Table 5–1 Annotations to Specify the Bean Type

Annotation	Description
@javax.ejb.Stateless	Specifies that the bean class is a stateless session bean. For more information, see Section A.2.28 , "javax.ejb.Stateless."
@javax.ejb.Singleton	Specifies that the bean class is a singleton session bean. For more information, see Section A.2.25 , "javax.ejb.Singleton."
@javax.ejb.Stateful	Specifies that the bean class is a stateful session bean. For more information, see Section A.2.26 , "javax.ejb.Startup."
@javax.ejb.Init	Specifies the correspondence of a stateful session bean class method with a create<METHOD> method for an adapted EJB 2.1 EJBHome and/or EJBLocalHome client view. For more information, see Section A.2.12 , "javax.ejb.Init."
@javax.ejb.Remove	Specifies a remove method of a stateful session bean. For more information, see Section A.2.22 , "javax.ejb.Remove."
@javax.ejb.MessageDriven	Specifies that the bean class is a message-driven bean. For more information, see Section A.2.17 , "javax.ejb.MessageDriven."
@javax.ejb.ActivationConfigProperty	Specifies properties used to configure a message-driven bean in its operational environment. For more information, see Section A.2.2 , "javax.ejb.ActivationConfigProperty."

5.4.2 Annotations to Specify the Local or Remote Interfaces

The following summarize the annotations used to specify the local or remote interfaces.

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

Annotation	Description
@javax.ejb.Local	Specifies a local interface of the bean. For more information, see Section A.2.13 , "javax.ejb.Local."
@javax.ejb.Remote	Specifies a remote interface of the bean. For more information, see Section A.2.20 , "javax.ejb.Remote."

5.4.3 Annotations to Support EJB 2.x Client View

The following summarize the annotations used to support EJB 2.x client view.

Table 5–3 Annotations to Support EJB 2.x Client View

Annotation	Description
@javax.ejb.LocalHome	Specifies a local home interface of the bean. For more information, see Section A.2.15 , "javax.ejb.LocalHome."
@javax.ejb.RemoteHome	Specifies a remote home interface of the bean. For more information, see Section A.2.21 , "javax.ejb.RemoteHome."

5.4.4 Annotations to Invoke a 3.0 Entity Bean

The following summarize the annotations used 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. For more information, see Section A.4.1 , " "javax.persistence.PersistenceContext." "
<code>@javax.persistence.PersistenceContexts</code>	Specifies one or more <code>PersistenceContext</code> annotations. For more information, see Section A.4.2 , " "javax.persistence.PersistenceContexts." "
<code>@javax.persistence.PersistenceUnit</code>	Specifies a dependency on an <code>EntityManagerFactory</code> . For more information, see Section A.4.3 , " "javax.persistence.PersistenceUnit." "
<code>@javax.persistence.PersistenceUnits</code>	Specifies one or more <code>PersistenceUnit</code> annotations. For more information, see Section A.4.4 , " "javax.persistence.PersistenceUnits." "

5.4.5 Transaction-Related Annotations

The following summarize the annotations used for transactions.

Table 5–5 Transaction-Related Annotations

Annotation	Description
<code>@javax.ejb.TransactionManagement</code>	Specifies the transaction management demarcation type (container- or bean-managed). For more information, see Section A.2.31 , " "javax.ejb.TransactionManagement." "
<code>@javax.ejb.TransactionAttribute</code>	Specifies whether a business method is invoked within the context of a transaction. For more information, see Section A.2.31 , " "javax.ejb.TransactionManagement." "

5.4.6 Annotations to Specify Interceptors

The following summarize the annotations used 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. For more information, see Section A.3.4 , " "javax.interceptor.Interceptors." "
<code>@javax.interceptor.AroundInvoke</code>	Specifies an interceptor method. For more information, see Section A.3.1 , " "javax.interceptor.AroundInvoke." "
<code>@javax.interceptor.ExcludeClassInterceptors</code>	Specifies that, when the annotated method is invoked, the class-level interceptors should <i>not</i> invoke. For more information, see Section A.3.2 , " "javax.interceptor.ExcludeClassInterceptors." "

Table 5–6 (Cont.) Annotations to Specify Interceptors

Annotation	Description
<code>@javax.interceptor.ExcludeDefaultInterceptors</code>	Specifies that, when the annotated method is invoked, the default interceptors should <i>not</i> invoke. For more information, see Section A.3.3 , " javax.interceptor.ExcludeDefaultInterceptors ."

5.4.7 Annotations to Specify Life Cycle Callbacks

The following summarize the annotations used 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. For more information, see Section A.2.18 , " javax.ejb.PostActivate ."
<code>@javax.ejb.PrePassivate</code>	Designates a method to receive a callback before a stateful session bean is passivated. For more information, see Section A.2.19 , " javax.ejb.PrePassivate ."
<code>@javax.annotation.PostConstruct</code>	Specifies the method that needs to be executed after dependency injection is done to perform any initialization. For more information, see Section A.5.1 , " javax.annotation.PostConstruct ."
<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. For more information, see Section A.5.2 , " javax.annotation.PreDestroy ."

5.4.8 Security-Related Annotations

The following metadata annotations are not specific to EJB 3.1, 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.RolesAllowed</code>	Specifies the references to security roles in the bean class. For more information, see Section A.6.1 , " javax.annotation.security.RolesAllowed ."
<code>@javax.annotation.security.PermitAll</code>	Specifies that all security roles are allowed to invoke the method. For more information, see Section A.6.3 , " javax.annotation.security.PermitAll ."
<code>@javax.annotation.security.DenyAll</code>	Specifies that no security roles are allowed to invoke the method. For more information, see Section A.6.2 , " javax.annotation.security.DenyAll ."

Table 5–8 (Cont.) Security-Related Annotations

Annotation	Description
<code>@javax.annotation.security.RunAs</code>	Specifies the security role which the method is run as. For more information, see Section A.6.5 , " javax.annotation.security.RunAs ."

5.4.9 Context Dependency Annotations

The following summarize the annotations used for context dependency.

Table 5–9 Context Dependency Annotations

Annotation	Description
<code>@javax.ejb.EJB</code>	Specifies a dependency to an EJB business interface or home interface. For more information, see Section A.2.10 , " javax.ejb.EJB ."
<code>@javax.ejb.EJBs</code>	Specifies one or more <code>@EJB</code> annotations. For more information, see Section A.2.11 , " javax.ejb.EJBs ."
<code>@javax.annotation.Resource</code>	Specifies a dependency on an external resource in the bean's environment. For more information, see Section A.5.3 , " javax.annotation.Resource ."
<code>@javax.annotation.Resources</code>	Specifies one or more <code>@Resource</code> annotations. For more information, see Section A.5.4 , " javax.annotation.Resources ."

5.4.10 Timeout and Exceptions Annotations

The following summarize the annotations used for timeout and exceptions.

Table 5–10 Timeout and Exception Annotations

Annotation	Description
<code>@javax.ejb.Timeout</code>	Specifies the timeout method of the bean class. For more information, see Section A.2.29 , " javax.ejb.Timeout ."
<code>@javax.ejb.ApplicationException</code>	Specifies that an exception is an application exception and should be reported to the client directly. For more information, see Section A.2.5 , " javax.ejb.ApplicationException ."

5.4.11 Timer and Scheduling Annotations

The following summarize the annotations used for timers scheduling-specific annotations.

Table 5–11 Timer and Scheduling Annotations

Annotation	Description
<code>@javax.ejb.Timeout</code>	Specifies the timeout method of the bean class. For more information, see Section A.2.29 , " javax.ejb.Timeout ."
<code>@javax.ejb.ApplicationException</code>	Specifies that an exception is an application exception and should be reported to the client directly. For more information, see Section A.2.5 , " javax.ejb.ApplicationException ."

Deployment Guidelines for Enterprise JavaBeans

The following sections contain EJB-specific deployment guidelines. For deployment topics that are common to all deployable application units, you will see cross-references to topics in *Deploying Applications to Oracle WebLogic Server*, a comprehensive guide to deploying WebLogic Server applications and modules.

- [Section 6.1, "Before You Deploy an EJB"](#)
- [Section 6.2, "Understanding and Performing Deployment Tasks"](#)
- [Section 6.3, "Deployment Guidelines for EJBs"](#)

6.1 Before You Deploy an EJB

Before starting the deployment process you should have:

- Functional, tested bean code, in an exploded directory format or packaged in an archive file—a JAR for a stand-alone EJB, an EAR if the EJB is part of an enterprise application, or a WAR if the EJB is part of a Web application—along with the deployment descriptors. For production environments, Oracle recommends that you package your application as an EAR.

Note: EJB 3.1 has removed the restriction that enterprise bean classes must be packaged in an `ejb-jar` file. Therefore, EJB classes can be packaged directly inside a Web application archive (WAR) using the same packaging guidelines that apply to Web application classes. See [Section 6.3.2, "Deploying EJBs as Part of a Web Application."](#)

For an overview of the steps required to create and package an EJB, see [Section 4.1, "Overview of the EJB Development Process."](#)

- Program the required annotated EJB class to specify the type of EJB—either: `@javax.ejb.Stateful`, `@javax.ejb.Stateless`, `@javax.ejb.Singleton`, or `@javax.ejb.MessageDriven`.

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

- Configured the optional, but supported, deployment descriptors—`ejb-jar.xml` and `weblogic-ear.xml`, and, for entity EJBs that use container-managed persistence, `weblogic-cmp-jar.xml`.

To create EJB deployment descriptors, see "Generate Deployment Descriptors" in *Programming Enterprise JavaBeans, Version 2.1, for Oracle WebLogic Server*.

6.2 Understanding and Performing Deployment Tasks

Table 6–1 is a guide to WebLogic Server documentation topics that help you make decisions about deployment strategies and provide instructions for performing deployment tasks. For EJB-specific deployment topics, see [Section 6.3, "Deployment Guidelines for EJBs."](#)

Table 6–1 *Deployment Tasks and Topics*

If You Want To....	See This Topic
Deploy in a development environment	"Deploying and Packaging from a Split Development Directory" in <i>Developing Applications for Oracle WebLogic Server</i> .
Select a deployment tool	"Deployment Tools" in <i>Deploying Applications to Oracle WebLogic Server</i>
Determine appropriate packaging for a deployment	"Preparing Applications and Modules for Deployment" in <i>Deploying Applications to Oracle WebLogic Server</i> .
Organizing EJB components in a split directory structure.	"EJBs" in <i>Developing Applications for Oracle WebLogic Server</i> .
Select staging mode	"Controlling Deployment File Copying with Staging Modes" in <i>Deploying Applications to Oracle WebLogic Server</i> .
Perform specific deployment tasks	"Overview of the Deployment Process" in <i>Deploying Applications to Oracle WebLogic Server</i> .

6.3 Deployment Guidelines for EJBs

The following sections provide guidelines for deploying EJBs.

- [Section 6.3.1, "Deploying Standalone EJBs as Part of an Enterprise Application"](#)
- [Section 6.3.2, "Deploying EJBs as Part of an Web Application"](#)
- [Section 6.3.3, "Deploying EJBs That Call Each Other in the Same Application"](#)
- [Section 6.3.4, "Deploying EJBs That Use Dependency Injection"](#)
- [Section 6.3.5, "Deploying Homogeneously to a Cluster"](#)
- [Section 6.3.6, "Deploying Pinned EJBs to a Cluster"](#)
- [Section 6.3.7, "Redeploying an EJB"](#)
- [Section 6.3.8, "Using FastSwap Deployment to Minimize Deployment"](#)
- [Section 6.3.9, "Understanding Warning Messages"](#)
- [Section 6.3.10, "Disabling EJB Deployment Warning Messages"](#)

6.3.1 Deploying Standalone EJBs as Part of an Enterprise Application

Oracle recommends that you package and deploy your stand-alone EJB applications as part of an Enterprise application. An Enterprise application is a Java EE 6 deployment unit that bundles together Web applications, EJBs, and Resource Adapters into a single deployable unit.

This is an Oracle best practice, which allows for easier application migration, additions, and changes. Also, packaging your applications as part of an Enterprise application allows you to take advantage of the split development directory structure, which provides a number of benefits over the traditional single directory structure.

See "Overview of the Split Development Directory Environment" in *Developing Applications for Oracle WebLogic Server*.

6.3.2 Deploying EJBs as Part of a Web Application

Enterprise beans can also be packaged within a web application module (WAR).

EJB 3.1 has removed the restriction that enterprise bean classes must be packaged in an `ejb-jar` file. Therefore, EJB classes can be packaged directly inside a Web application archive (WAR) using the same packaging guidelines that apply to Web application classes. Simply put your EJB classes in the `WEB-INF/classes` directory or in a JAR file within `WEB-INF/lib` directory. Optionally, if you are also using the EJB deployment descriptor, you can package it as `WEB-INF/ejb-jar.xml`. When you run the appc compiler, a WAR file with the classes required to access the EJB components is generated.

See [Chapter 4.11.2, "Packaging an EJB In a WAR."](#)

6.3.3 Deploying EJBs That Call Each Other in the Same Application

When an EJB in one application calls an EJB in another application, WebLogic Server passes method arguments by value, due to classloading requirements. When EJBs are in the same application, WebLogic Server can pass method arguments by reference; this improves the performance of method invocation because parameters are not copied.

For best performance, package components that call each other in the same application, and set `enable-call-by-reference` in `weblogic-ejb-jar.xml` to `True`. (By default, `enable-call-by-reference` is `False`.)

6.3.4 Deploying EJBs That Use Dependency Injection

When an EJB uses dependency injection, the resource name defined in the class and the superclass must be unique. For example:

```
public class ClientServlet extends HttpServlet {
    @EJB(name = 'DateServiceBean', beanInterface = DateService.class)
    private DateService bean; ... }
public class DerivedClientServlet extends ClientServlet {
    @EJB(name = 'MyDateServiceBean', beanInterface = DateService.class)
    private DateService bean; ... }
```

For more information about dependency injection, see "Using Java EE Annotations and Dependency Injection" in *Developing Applications for Oracle WebLogic Server*.

6.3.5 Deploying Homogeneously to a Cluster

If your EJBs will run on a WebLogic Server cluster, Oracle recommends that you deploy them homogeneously—to each Managed Server in the cluster. Alternatively, you can deploy an EJB to only to a single server in the cluster (that is, "pin" a module to a server). This type of deployment is less common, and should be used only in special circumstances where pinned services are required. For more information, "Understanding Cluster Configuration" in *Using Clusters for Oracle WebLogic Server*.

6.3.6 Deploying Pinned EJBs to a Cluster

There is a known issue with deploying or redeploying EJBs to a single server instance in a cluster—referred to as pinned deployment—if the JAR file contains uncompiled classes and interfaces.

During deployment, the uncompiled EJB is copied to each server instance in the cluster, but it is compiled only on the server instance to which it has been deployed. As a result, the server instances in the cluster to which the EJB was not targeted lack the classes generated during compilation that are necessary to invoke the EJB. When a client on another server instance tries to invoke the pinned EJB, it fails, and an Assertion error is thrown in the RMI layer.

If you are deploying or redeploying an EJB to a single server instance in a cluster, compile the EJB with `appc` before deploying it, to ensure that the generated classes are copied to all server instances available to all nodes in the cluster.

For more information on pinned deployments, see "Deploying to a Single Server Instance (Pinned Deployment)" in *Using Clusters for Oracle WebLogic Server*.

6.3.7 Redeploying an EJB

When you make changes to a deployed EJB's classes, you must redeploy the EJB. If you use automatic deployment, deployment occurs automatically when you restart WebLogic Server. Otherwise, you must explicitly redeploy the EJB.

Redeploying an EJB deployment enables an EJB provider to make changes to a deployed EJB's classes, recompile, and then "refresh" the classes in a running server.

When you redeploy, the classes currently loaded for the EJB are immediately marked as unavailable in the server, and the EJB's classloader and associated classes are removed. At the same time, a new EJB classloader is created, which loads and maintains the revised EJB classes.

When clients next acquire a reference to the EJB, their EJB method calls use the changed EJB classes.

You can redeploy an EJB that is standalone or part of an application, using the `weblogic.Deployer` tool or the Administration Console. For instructions, see "Redeploying Applications in a Production Environment" in *Deploying Applications to Oracle WebLogic Server*.

Production redeployment is not supported for:

- applications that use JTS drivers.
- applications that include EJB 1.1 container-managed persistence (CMP) EJBs. To use production redeployment with applications that include CMP EJBs, use EJB 2.x CMP instead of EJB 1.1 CMP.

For more information on production redeployment limitations, see "Requirements and Restrictions for Production Redeployment" in *Deploying Applications to Oracle WebLogic Server*.

6.3.8 Using FastSwap Deployment to Minimize Deployment

During iterative development of an EJB application, you make many modifications to the EJB implementation class file, typically redeploying an EJB module multiple times during its development.

Java EE 5 introduces the ability to redefine a class at runtime without dropping its ClassLoader or abandoning existing instances. This allows containers to reload altered

classes without disturbing running applications, vastly speeding up iterative development cycles and improving the overall development and testing experiences.

With FastSwap, Java classes are redefined in-place without reloading the ClassLoader, thereby having the decided advantage of fast turnaround times. This means that you do not have to wait for an application to redeploy for your changes to take affect. Instead, you can make your changes, auto compile, and then see the effects immediately.

For more information about FastSwap, see "Using FastSwap Deployment to Minimize Redeployment" in *Deploying Applications to Oracle WebLogic Server*.

6.3.9 Understanding Warning Messages

To get information about a particular warning, use the `weblogic.GetMessage` tool. For example:

```
java weblogic.GetMessage -detail -id BEA-010202
```

6.3.10 Disabling EJB Deployment Warning Messages

You can disable certain WebLogic Server warning messages that occur during deployment. You may find this useful if the messages provide information of which you are already aware.

For example, if the methods in your EJB makes calls by reference rather than by value, WebLogic Server generates this warning during deployment: "Call-by-reference not enabled."

You can use the `disable-warning` element in `weblogic-ejb-jar.xml` to disable certain messages. For a list of messages you can disable, and instructions for disabling the messages, see "disable-warning" in the "weblogic-ejb-jar.xml Deployment Descriptor Reference" chapter of *Programming Enterprise JavaBeans, Version 2.1, for Oracle WebLogic Server*.

Using an Embedded EJB Container in Oracle WebLogic Server

This chapter provides an overview of using an embeddable EJB container in Oracle WebLogic Server. The following topics are covered:

- [Section 7.1, "Overview of the Embeddable EJB Container"](#)
- [Section 7.2, "EJB 3.1 Lite Functionality Supported in the Embedded EJB Container"](#)

7.1 Overview of the Embeddable EJB Container

Unlike traditional Java EE server-based execution, embeddable usage allows client code and its corresponding enterprise beans to run within the same virtual machine and class loader. This provides better support for testing, offline processing (e.g., batch jobs), and the use of the EJB programming model in desktop applications.

Most of the services present in the enterprise bean container in a Java EE server are available in the embedded enterprise bean container, including injection, container-managed transactions, and security. Enterprise bean components execute similarly in both embedded and Java EE environments, and therefore the same enterprise bean can be easily reused in both standalone and networked applications.

For detailed information about using the Embedded Enterprise Bean Container, see "Developing Embeddable Enterprise Bean Applications" in the "Enterprise Beans" chapter of the Java EE 6 Tutorial at <http://download.oracle.com/javaee/6/tutorial/doc/gkcrr.html>

7.2 EJB 3.1 Lite Functionality Supported in the Embedded EJB Container

The EJB Lite subset of the EJB 3.1 API is supported in the Embedded EJB Container. EJB Lite is by definition a subset of functionality and doesn't describe any new feature or functionality. This section outlines the requirements of EJB Lite support as defined by the EJB 3.1 specification.

[Table 7-1](#) represents the official requirements for EJB 3.1 Lite functionality support as defined by the EJB 3.1 specification.

Table 7-1 EJB 3.1 Lite Functionality Supported in Embedded EJB Container

	EJB 3.1 Lite	EJB 3.1 Full
Components		

Table 7-1 EJB 3.1 Lite Functionality Supported in Embedded EJB Container

	EJB 3.1 Lite	EJB 3.1 Full
Session Beans (stateful, stateless, singleton)	Yes	Yes
Message-Driven Beans	No	Yes
2.x/1.1 CMP/BMP Entity Beans	No	Yes
JPA 2.0	Yes	Yes
Session Bean Client Views		
Local/No Interface	Yes	Yes
3.0 Remote	No	Yes
2.x Remote Home/Component	No	Yes
JAX-WS Web Services Endpoint	No	Yes
JAX-RPC Web Services Endpoint	No	Yes
Services		
EJB Timer Service	No	Yes
Asynchronous Session Bean Invocations	No	Yes
Interceptors	Yes	Yes
RMI-IIOP Interoperability	No	Yes
Container-managed Transactions/Bean-managed Transactions	Yes	Yes
Declarative and Programmatic Security	Yes	Yes
Miscellaneous		
Embeddable API	Yes	Yes

Configuring the Persistence Provider in Oracle WebLogic Server

This chapter describes Oracle TopLink, the default persistence provider in Oracle WebLogic Server, and introduces how to use it. This chapter also tells how to set the default persistence provider in Weblogic Server, how to modify existing Kodo applications for use in the current release, and how to upgrade to a newer version of OpenJPA.

The following topics are covered:

- Section 8.1, "Overview of Oracle TopLink"
- Section 8.2, "Specifying a Persistence Provider"
- Section 8.3, "Using Oracle TopLink in Oracle WebLogic Server"
- Section 8.4, "Using Oracle Kodo in Oracle WebLogic Server"
- Section 8.5, "Using a Newer Version of OpenJPA in Oracle WebLogic Server"

8.1 Overview of Oracle TopLink

Oracle TopLink is the default persistence provider in WebLogic Server 12c and later. It is a comprehensive standards-based object-persistence and object-transformation framework that provides APIs, schemas, and run-time services for the persistence layer of an application.

The core component of TopLink is the EclipseLink project's produced libraries and utilities. EclipseLink is the open source implementation of the development framework and the runtime provided in TopLink. EclipseLink implements the following specifications, plus value-added extensions:

- Java Persistence 2.0 (JPA 2.0).

JPA 2.0 is part of Java Platform, Enterprise Edition 6 (Java EE 6). It includes improvements and enhancements to domain modeling, object/relational mapping, `EntityManager` and `Query` interfaces, and the Java Persistence Query Language (JPQL). It includes an API for criteria queries, a metamodel API, and support for validation.

For the JPA 2.0 Specification, see "JSR-000317 Java Persistence 2.0" at <http://jcp.org/aboutJava/communityprocess/final/jsr317/index.html>.

- Java Architecture for XML Binding (JAXB) 2.2. (The EclipseLink JAXB implementation, plus EclipseLink extensions, is called MOXy.)

For the JAXB 2.0 specification, see "JSR-000222 Java Architecture for XML Binding (JAXB) 2.0" at <http://jcp.org/aboutJava/communityprocess/pfd/jsr222/index.html>.

- EclipseLink also includes Database Web Service (DBWS), which provides access to relational database artifacts by using a Java API for XML Web Services (JAX-WS) 2 Web service.

EclipseLink also provides support for Oracle Spatial and Oracle XDB mapping.

For more information about EclipseLink, including other supported services, see the EclipseLink project home at <http://wiki.eclipse.org/EclipseLink> and the EclipseLink Documentation Center at <http://wiki.eclipse.org/EclipseLink/UserGuide>.

In addition to all of EclipseLink, Oracle TopLink includes:

- TopLink Grid, an integration between TopLink and Oracle Coherence that allows TopLink to use Oracle Coherence as a level 2 (L2) cache and persistence layer for entities. For more information, see *Oracle Coherence Developer's Guide* and *Oracle Fusion Middleware Integration Guide for Oracle TopLink with Coherence Grid*.

Note: You must have a license for Oracle Coherence to be able to use TopLink Grid.

- Logging integration with WebLogic Server.
- MBean support in WebLogic Server.

For information about developing, deploying, and configuring Oracle TopLink applications, see the following:

- *Oracle Fusion Middleware Solution Guide for Oracle TopLink*
- *Oracle Fusion Middleware Oracle TopLink Concepts*
- *Oracle Fusion Middleware Java API Reference for EclipseLink*
- EclipseLink Documentation Center at <http://wiki.eclipse.org/EclipseLink/UserGuide>
- EclipseLink examples at <http://wiki.eclipse.org/EclipseLink/Examples>.

8.2 Specifying a Persistence Provider

You can specify what persistence provider to use for a persistence unit in the application code or by accepting the default persistence provider set for the WebLogic Server domain, as described in the following sections:

- [Section 8.2.1, "Setting the Default Provider for the Domain"](#)
- [Section 8.2.2, "Specifying the Persistence Provider in an Application"](#)

8.2.1 Setting the Default Provider for the Domain

Unless you specify otherwise, TopLink is used as the default persistence provider for a WebLogic Server domain. The default provider is used for any entities in an application that are not configured to use a different persistence provider. The default

provider is used for both injected and application-managed entity managers and factories.

You can set the default provider in the WLS Administration Console or by directly setting `JPAMBean.DefaultJPAProvider`. In the Administration Console, the options are Oracle TopLink or Oracle Kodo.

Note: Oracle Kodo JPA/JDO is deprecated in this release. Customers are encouraged to use Oracle TopLink, which supports JPA 2.0. Kodo supports only JPA 1.0.

For instructions on setting the default through the Administration Console, see "Configure the Default JPA Persistence Provider" in *Oracle WebLogic Server Administration Console Help*.

If you change the default provider, you must do the following for any deployed applications that do not specify a JPA provider:

- Restart applications that use application-managed entity manager factories.
- Redeploy applications that use injected entity manager factories or entity managers.

8.2.2 Specifying the Persistence Provider in an Application

A persistence provider specified in an application takes precedence over the default provider set for the WebLogic Server domain.

You can set the provider to use in the following ways:

- Specify the provider in the `<provider>` element for a persistence unit in the `persistence.xml` file, for example:


```
<persistence-unit name="example">
  <provider>org.eclipse.persistence.jpa.PersistenceProvider</provider>
  ...
</persistence-unit>
```
- Specify the provider in the `javax.persistence.provider` property passed to the `Map` parameter of the `javax.persistence.Persistence.createEntityManagerFactory(String, Map)` method.

8.3 Using Oracle TopLink in Oracle WebLogic Server

For detailed information about using Oracle TopLink in WebLogic server, see "Using TopLink with WebLogic Server" in *Oracle Fusion Middleware Solution Guide for Oracle TopLink*.

8.4 Using Oracle Kodo in Oracle WebLogic Server

It is advantageous to use the default Oracle TopLink for your applications, due to its implementation of JPA 2.0, its enhanced set of features, its integration in WebLogic Server, and its integration with TopLink Grid. However, you may want to continue to use Kodo for existing applications that were written for use with Kodo. Although deprecated, Kodo is still included with WebLogic Server: you do not have to install it.

You can set it as the default provider, as described in [Section 8.2, "Specifying a Persistence Provider,"](#) or configure an application to use it.

Note: Switching an application from one JPA provider to another usually involves more than just flipping a switch. Most obviously, if the application uses proprietary features of the original JPA implementation, that usage will have to change when moving to a different implementation.

In addition, differences in the implementations of a standard specification (that is, JPA) can affect the behavior of an application. Although Kodo and TopLink implement their JPA specifications, as do other conforming JPA implementations, each implementation makes its own decisions about unspecified details that can affect the behavior of the application using JPA. When making the switch from one JPA implementation to another, the application should be thoroughly tested with the expectation that some behavioral differences will have to be corrected.

If you use Kodo, be aware of the following:

- [Section 8.4.1, "Using JPA 1.0 APIs"](#)
- [Section 8.4.2, "Using persistence-configuration.xml"](#)
- [Section 8.4.3, "Updating Applications to Overcome Conflicts with JPA 2.0"](#)

8.4.1 Using JPA 1.0 APIs

You cannot use any JPA 2 APIs with Kodo. Because Kodo supports only JPA 1, trying to use any JPA 2 APIs will generate an error.

8.4.2 Using persistence-configuration.xml

The Kodo `persistence-configuration.xml` descriptor is valid only when Kodo is the persistence provider. If you include a `persistence-configuration.xml` in an application that uses injection and you do not specify Kodo as the persistence provider, the `persistence-configuration.xml` descriptor is ignored and a warning is issued.

8.4.3 Updating Applications to Overcome Conflicts with JPA 2.0

Kodo is based on open-source OpenJPA 1.x, the JPA 1.0 framework and provider from the Apache Foundation. Although JPA 1.0 is upwardly compatible with JPA 2.0, JPA 2.0 introduced some methods to existing JPA interfaces that conflict with existing signatures in the OpenJPA interfaces used in Kodo. These conflicts arise because the OpenJPA interfaces extend the JPA interfaces that contain the methods that are revised in JPA 2.0.

These conflicts could cause problems with applications that use Kodo as the persistence provider, because the JPA 2.0 jar (rather than the JPA 1.0 jar) is on the WebLogic Server classpath.

To prevent these conflicts, two methods have been changed in the OpenJPA interfaces that ship with Kodo in WebLogic Server 12c and later. The OpenJPA methods that cause the conflict are:

- `public <T> T OpenJPAEntityManager.detach(T pc)`

The above method conflicts with the following JPA 2.0 method in the `EntityManager` interface:

```
public void detach(Object)
```

- `public Properties OpenJPAEntityManagerFactory.getProperties()`

The above method conflicts with the following JPA 2.0 method in the `EntityManagerFactory` interface:

```
public Map<String, Object> getProperties()
```

The new signatures are:

- `public <T> T OpenJPAEntityManager.detachCopy(T pc)`
- `public Map<String, Object> OpenJPAEntityManagerFactory.getProperties()`
- `public Properties OpenJPAEntityManagerFactory.getPropertiesAsProperties()`

The first two new signatures follow the changes made in OpenJPA 2.x to address these same conflicts. These signature changes do not alter the semantics of the methods. The `getPropertiesAsProperties` method is provided as a convenience to application developers who do not want to accept the different return type of the redesigned `getProperties` method. However, the method `getPropertiesAsProperties` is not currently in the OpenJPA 2.x interface.

Because of these changes, you should update any applications that use these methods and recompile them with the Kodo OpenJPA 1.0 jar (`org.apache.openjpa_n.jar`) and the JPA 1 jar (`javax.persistence_n.jar`) that are shipped with WebLogic Server 12c (or later). This will prevent the old method signatures from being used. After recompiling and deploying, the applications will run without change in behavior.

Applications that use Kodo/JDO are not affected by the conflicts described above. Therefore, you do not have to modify or recompile them.

8.5 Using a Newer Version of OpenJPA in Oracle WebLogic Server

The current release of OpenJPA from the Apache Foundation supports JPA 2, and you can use it as the persistence provider for applications deployed to WebLogic Server.

Note: OpenJPA is a third-party library that you must plug into WebLogic Server. As such, Oracle support for Oracle WebLogic Server does not include support for the use of OpenJPA. These instructions are provided as a courtesy.

To use the newer version of OpenJPA, do the following:

- Configure a filtering classloader to use the Kodo and OpenJPA classes, as follows:
 - In a Web application, include the following stanza in the `weblogic.xml` file:

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

- In an enterprise application, include the `prefer-application-packages` stanza in `weblogic-application.xml`.)

For more information about filtering classloaders, see "Using a Filtering Classloader" in *Developing Applications for Oracle WebLogic Server*.

- Include the OpenJPA jar file in the application's `lib` directory. For more information about library directories, see "Library Directories" in *Developing Applications for Oracle WebLogic Server*.
- Configure the application to use OpenJPA (Kodo) as the provider. You cannot simply set Kodo as the domain's default provider in WebLogic Server. You must configure it in the application, as described in [Section 8.2.2, "Specifying the Persistence Provider in an Application."](#)

Note: This last step applies for any JPA implementation that is packaged with an application. You must specify the implementation to use in the application code and not just accept the default set in WebLogic Server.

EJB Metadata Annotations Reference

The following topics provide reference information about the EJB 3.x metadata annotations:

- [Section A.1, "Overview of EJB 3.x Annotations"](#)
- [Section A.2, "Annotations for Stateless, Stateful, and Message-Driven Beans"](#)
- [Section A.3, "Annotations Used to Configure Interceptors"](#)
- [Section A.4, "Annotations Used to Interact With Entity Beans"](#)
- [Section A.5, "Standard JDK Annotations Used By EJB 3.x"](#)
- [Section A.6, "Standard Security-Related JDK Annotations Used by EJB 3.x"](#)
- [Section A.7, "WebLogic Annotations"](#)

A.1 Overview of EJB 3.x Annotations

The WebLogic Server EJB 3.1 programming model uses the Java EE 6 metadata annotations feature in which you create an annotated EJB 3.1 bean file, and then compile the class with standard Java compiler, which can then be packaged into a target module for deployment. At runtime, WebLogic Server parses the annotations and applies the required behavioral aspects to the bean file.

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 EE 6 components, and are thus in more generic packages, such as `javax.annotation`.

Note: If you are continuing to use deployment descriptors in your EJB implementation, refer to "EJB Deployment Descriptors" in *Programming Enterprise JavaBeans, Version 2.1, for Oracle WebLogic Server*.

A.2 Annotations for Stateless, Stateful, and Message-Driven Beans

This section provides reference information for the following annotations:

- [Section A.2.1, "javax.ejb.AccessTimeout"](#)
- [Section A.2.2, "javax.ejb.ActivationConfigProperty"](#)
- [Section A.2.3, "javax.ejb.AfterBegin"](#)

- Section A.2.4, "javax.ejb.AfterCompletion"
- Section A.2.5, "javax.ejb.ApplicationException"
- Section A.2.6, "javax.ejb.Asynchronous"
- Section A.2.7, "javax.ejb.BeforeCompletion"
- Section A.2.8, "javax.ejb.ConcurrencyManagement"
- Section A.2.9, "javax.ejb.DependsOn"
- Section A.2.10, "javax.ejb.EJB"
- Section A.2.11, "javax.ejb.EJBs"
- Section A.2.12, "javax.ejb.Init"
- Section A.2.13, "javax.ejb.Local"
- Section A.2.14, "javax.ejb.LocalBean"
- Section A.2.15, "javax.ejb.LocalHome"
- Section A.2.16, "javax.ejb.Lock"
- Section A.2.17, "javax.ejb.MessageDriven"
- Section A.2.18, "javax.ejb.PostActivate"
- Section A.2.19, "javax.ejb.PrePassivate"
- Section A.2.20, "javax.ejb.Remote"
- Section A.2.21, "javax.ejb.RemoteHome"
- Section A.2.22, "javax.ejb.Remove"
- Section A.2.23, "javax.ejb.Schedule"
- Section A.2.24, "javax.ejb.Schedules"
- Section A.2.25, "javax.ejb.Singleton"
- Section A.2.26, "javax.ejb.Startup"
- Section A.2.27, "javax.ejb.StatefulTimeout"
- Section A.2.28, "javax.ejb.Stateless"
- Section A.2.29, "javax.ejb.Timeout"
- Section A.2.30, "javax.ejb.TransactionAttribute"
- Section A.2.31, "javax.ejb.TransactionManagement"

A.2.1 javax.ejb.AccessTimeout

New Section - added this section per the 3.1 spec.

The following sections describe the annotation in more detail.

A.2.1.1 Description

Target: Method, Type

Specifies the amount of time in a given time unit that a concurrent access attempt should block before timing out.

This annotation may be applied to a stateful session bean or to a singleton session bean that uses container managed concurrency.

By default, clients are allowed to make concurrent calls to a stateful session object and the container is required to serialize such concurrent requests. The `AccessTimeout` annotation is used to specify the amount of time a stateful session bean request should block in the case that the bean instance is already processing a different request. Use of the `AccessTimeout` annotation with a value of 0 specifies to the container that concurrent client requests to a stateful session bean are prohibited.

The `AccessTimeout` annotation can be specified on a business method or a bean class. If it is specified on a class, it applies to all business methods of that class. If it is specified on both a class and on a business method of the class, the method-level annotation takes precedence for the given method.

Access time-outs for a singleton session bean only apply to methods eligible for concurrency locks. The `AccessTimeout` annotation can be specified on the singleton session bean class or on an eligible method of the class. If `AccessTimeout` is specified on both a class and on a method of that class, the method-level annotation takes precedence for the given method.

For details, see [Section 4.6, "Optionally Program the EJB Timer Service."](#)

A.2.1.2 Attributes

The following table summarizes the attributes.

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

Name	Description	Data Type	Required?
value	<p>Specifies the amount of time in a given time unit that a concurrent access attempt should block before timing out.</p> <ul style="list-style-type: none"> ▪ A value > 0 indicates a timeout value in the units specified by the unit element. ▪ A value of 0 means concurrent access is not permitted. ▪ A value of -1 indicates that the client request will block indefinitely until forward progress it can proceed. <p>Values less than -1 are not valid.</p>	Long	No
unit	<p>Specifies the units used for the specified value.</p> <p>The default value for this attribute is <code>java.util.concurrent.TimeUnit.MILLISECONDS</code>.</p>	TimeUnit	No

A.2.2 `javax.ejb.ActivationConfigProperty`

The following sections describe the annotation in more detail.

A.2.2.1 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. For more information about this annotation, see "Using EJB 3.3.x0 Compliant MDBs" and "Deployment Elements and

Annotations for MDBs" in *Programming Message-Driven Beans for Oracle WebLogic Server*.

A.2.2.2 Attributes

The following table summarizes the attributes.

Table A–2 Attributes of the `javax.ejb.ActivationConfigProperty` Annotation

Name	Description	Data Type	Required?
value	<p>The semantics of the value element are as follows:</p> <ul style="list-style-type: none"> ▪ A value > 0 indicates a timeout value in the units specified by the unit element. ▪ A value of 0 means concurrent access is not permitted. ▪ A value of -1 indicates that the client request will block indefinitely until forward progress it can proceed. <p>Values less than -1 are not valid.</p>	String	Yes
unit	Specifies the value of the activation property.	String	Yes

A.2.3 `javax.ejb.AfterBegin`

New Section - added this section per the 3.1 spec.

The following sections describe the annotation in more detail.

A.2.3.1 Description

Target: Method

Designate a stateful session bean method to receive the after begin session synchronization callback.

The after begin callback notifies a stateful session bean instance that a new transaction has started and that the subsequent business methods on the instance will be invoked in the context of the transaction.

This method executes in the proper transaction context. A bean must have at most one `AfterBegin` method. The signature of this method must observe the following rules:

- The method must not be declared as `final` or `static`.
- The method may have any access type.
- The return type must be `void`.
- The method must take no arguments.

This method executes with no transaction context.

A stateful session bean class may use either the `SessionSynchronization` interface or the session synchronization annotations, but not both.

A.2.4 `javax.ejb.AfterCompletion`

New Section - added this section per the 3.1 spec.

The following sections describe the annotation in more detail.

A.2.4.1 Description

Target: Method

Designate a stateful session bean method to receive the after completion session synchronization callback.

The after completion callback notifies a stateful session bean instance that a transaction commit protocol has completed. A completion status of `true` indicates that the transaction has committed. A status of `false` indicates that a rollback has occurred.

A bean must have at most one `AfterCompletion` method. The signature of this method must observe the following rules:

- The method must not be declared as `final` or `static`.
- The method may have any access type.
- The return type must be `void`.
- The method must take a single argument of type `boolean`.

This method executes with no transaction context.

A stateful session bean class may use either the `SessionSynchronization` interface or the session synchronization annotations, but not both.

A.2.5 `javax.ejb.ApplicationException`

The following sections describe the annotation in more detail.

A.2.5.1 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.

A.2.5.2 Attributes

The following table summarizes the attributes.

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

Name	Description	Data Type	Required?
<code>rollback</code>	Specifies whether the EJB container should rollback the transaction, if the bean is currently being invoked inside of one, if the exception is thrown. 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.	<code>boolean</code>	No

A.2.6 `javax.ejb.Asynchronous`

New Section - added this section per the 3.1 spec.

The following sections describe the annotation in more detail.

A.2.6.1 Description

Target: Method, Type

Used to mark a session bean method as an asynchronous method or to designate all business methods of a session bean class as asynchronous.

An `asynchronous` method must have return type `void` or `Future<V>`, where `V` is the result value type.

Asynchronous method invocation semantics only apply to the no-interface, local business, and remote business client views. Methods exposed through the EJB 2.x local, EJB 2.x remote, and Web service client views must not be designated as asynchronous.

A.2.7 `javax.ejb.BeforeCompletion`

New Section - added this section per the 3.1 spec.

The following sections describe the annotation in more detail.

A.2.7.1 Description

Target: Method

Designate a stateful session bean method to receive the before completion session synchronization callback.

The before completion callback notifies a stateful session bean instance that a transaction is about to be committed.

This method executes in the proper transaction context.

Note: The instance may still cause the container to rollback the transaction by invoking the `setRollbackOnly()` method on the session context or by throwing an exception. A bean must have at most one `BeforeCompletion` method.

The signature of this method must observe the following rules:

- The method must not be declared as `final` or `static`.
- The method may have any access type.
- The return type must be `void`.
- The method must take no arguments.

This method executes with no transaction context.

A stateful session bean class may use either the `SessionSynchronization` interface or the session synchronization annotations, but not both.

A.2.8 `javax.ejb.ConcurrencyManagement`

New Section - added this section per the 3.1 spec.

The following sections describe the annotation in more detail.

A.2.8.1 Description

Target: Type

Declares a singleton session bean's concurrency management type.

If this annotation is not specified, the singleton bean is assumed to have container managed concurrency.

This annotation may be applied to stateful session beans, but doing so has no impact on the semantics of concurrency management for such beans. The concurrency management type for bean-managed concurrency (BEAN) does not apply to stateful session beans.

A.2.8.2 Attributes

The following table summarizes the attributes.

Table A-4 *Attributes of the `javax.ejb.ConcurrencyManagement` Annotation*

Name	Description	Data Type	Required?
value	<p>Specifies the concurrency management type used by the bean class.</p> <p>Valid values for this attribute are:</p> <ul style="list-style-type: none"> ▪ <code>ConcurrencyManagementType.CONTAINER</code> ▪ <code>ConcurrencyManagementType.BEAN</code> <p>The default value for this attribute is <code>javax.ejb.ConcurrencyManagementType.CONTAINER</code>.</p>	String	No.

A.2.9 `javax.ejb.DependsOn`

The following sections describe the annotation in more detail.

A.2.9.1 Description

Target: Type

Used to express an initialization dependency between singleton components.

The container ensures that all singleton beans with which a singleton has a `DependsOn` relationship have been initialized before the singleton's `PostConstruct` method is called.

During application shutdown the container ensures that all singleton beans on which the singleton has a `DependsOn` relationship are still available during the singleton's `PreDestroy` method.

A.2.9.2 Attributes

The following table summarizes the attributes.

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

Name	Description	Data Type	Required?
value	<p>Specifies <code>ejb-names</code> of singleton components whose initialization must occur before this singleton. The order in which these names are listed is not significant.</p>	String	No.

A.2.10 `javax.ejb.EJB`

The following sections describe the annotation in more detail.

Note: This section contains "mapped-name" refs instead of "lookup-name"

A.2.10.1 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.

A.2.10.2 Attributes

The following table summarizes the attributes.

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

Name	Description	Data Type	Required?
<code>name</code>	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	No
<code>beanInterface</code>	Specifies the interface type of the referenced EJB (either a business or home interface). Default value for this attribute is <code>Object.class</code>	Class	No
<code>beanName</code>	Specifies the name of the referenced EJB. This attribute corresponds to the <code>name</code> 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. 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.	String	No

Table A-6 (Cont.) Attributes of the `javax.ejb.EJB` Annotation

Name	Description	Data Type	Required?
<code>mappedName</code>	<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
<code>description</code>	Describes the EJB reference.	String	No

A.2.11 `javax.ejb.EJBs`

The following sections describe the annotation in more detail.

A.2.11.1 Description

Target: Class

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

A.2.11.2 Attribute

The following table summarizes the attributes.

Table A-7 Attribute of the `javax.ejb.EJBs` Annotation

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

A.2.12 `javax.ejb.Init`

The following sections describe the annotation in more detail.

A.2.12.1 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.

A.2.12.2 Attributes

The following table summarizes the attributes.

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

Name	Description	Data Type	Required?
value	Specifies the name of the corresponding <code>createMETHOD</code> method. 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.	String	No

A.2.13 `javax.ejb.Local`

The following sections describe the annotation in more detail.

A.2.13.1 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.

A.2.13.2 Attributes

The following table summarizes the attributes.

Table A-9 Attributes of the `javax.ejb.Local`

Name	Description	Data Type	Required?
value	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: <ul style="list-style-type: none"> ▪ <code>java.io.Serializable</code> ▪ <code>java.io.Externalizable</code> ▪ <code>javax.ejb.*</code> 	Class[]	No

A.2.14 `javax.ejb.LocalBean`

New Section - added this section per the 3.1 spec.

The following sections describe the annotation in more detail.

A.2.14.1 Description**Target:** Type

Designates that a session bean exposes a no-interface view.

A.2.15 javax.ejb.LocalHome

The following sections describe the annotation in more detail.

A.2.15.1 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.x bean. You can also use this annotation with bean classes that have been written to the EJB 2.1 APIs.

A.2.15.2 Attributes

The following table summarizes the attributes.

Table A-10 Attributes of the *javax.ejb.LocalHome* Annotation

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

A.2.16 javax.ejb.Lock**New Section - added this section per the 3.1 spec.**

The following sections describe the annotation in more detail.

A.2.16.1 Description**Target:** Type, Method

Declares a concurrency lock for a singleton session bean with container managed concurrency.

This annotation may be specified on the bean class, the business methods of the bean class or both. Specifying the annotation on a business method overrides the value specified at class level, if any.

If this annotation is not used, a value of `Lock (WRITE)` is assumed.**A.2.16.2 Attributes**

The following table summarizes the attributes.

Table A–11 *Attributes of the javax.ejb.LockType Annotation*

Name	Description	Data Type	Required?
value	<p>Specifies the concurrency lock used by the singleton session bean with container managed concurrency.</p> <p>Valid values for this attribute are:</p> <ul style="list-style-type: none"> ▪ <code>LockType.READ</code> ▪ <code>LockType.WRITE</code> <p>Default value is <code>javax.ejb.LockType.WRITE</code>.</p>	String	No.

A.2.17 javax.ejb.MessageDriven

The following sections describe the annotation in more detail.

Note: This section contains "mapped-name" refs instead of "lookup-name"

A.2.17.1 Description

Target: Class

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

A.2.17.2 Attributes

The following table summarizes the attributes.

Table A–12 *Attributes of the javax.ejb.MessageDriven Annotation*

Name	Description	Data Type	Required?
name	<p>Specifies the name of the message-driven bean.</p> <p>If you do not specify this attribute, the default value is the unqualified name of the bean class.</p>	String	No
messageListenerInterface	<p>Specifies the message-listener interface of the bean class.</p> <p>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.</p> <p>The default value for this attribute is <code>Object.class</code>.</p>	Class	No

Table A-12 (Cont.) 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>	ActivationConfigProperty[]	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

A.2.18 `javax.ejb.PostActivate`

The following sections describe the annotation in more detail.

A.2.18.1 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.

A.2.19 `javax.ejb.PrePassivate`

The following sections describe the annotation in more detail.

A.2.19.1 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.

A.2.20 `javax.ejb.Remote`

The following sections describe the annotation in more detail.

A.2.20.1 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.

A.2.20.2 Attributes

The following table summarizes the attributes.

Table A-13 *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

A.2.21 javax.ejb.RemoteHome

The following sections describe the annotation in more detail.

A.2.21.1 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.x bean. You can also use this annotation with bean classes that have been written to the EJB 2.1 APIs.

A.2.21.2 Attributes

The following table summarizes the attributes.

Table A-14 *Attributes of the javax.ejb.RemoteHome Annotation*

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

A.2.22 javax.ejb.Remove

The following sections describe the annotation in more detail.

A.2.22.1 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.

A.2.22.2 Attributes

The following table summarizes the attributes.

Table A–15 *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

A.2.23 javax.ejb.Schedule

New Section - added this section per the 3.1 spec.

The following sections describe the annotation in more detail.

A.2.23.1 Description

Target: Class

Schedule a timer for automatic creation with a timeout schedule based on a CRON-like time expression. The annotated method is used as the timeout callback method.

All elements of this annotation are optional. If none are specified a persistent timer will be created with callbacks occurring every day at midnight in the default time zone associated with the container in which the application is executing.

There are seven elements that constitute a schedule specification which are listed below. In addition, the `timezone` element may be used to specify a non-default time zone in whose context the schedule specification is to be evaluated; the `persistent` element may be used to specify a non-persistent timer, and the `info` element may be used to specify additional information that may be retrieved when the timer callback occurs.

A.2.23.1.1 Calendar-based Schedule Elements The elements that specify the calendar-based schedule itself are as follows:

- `second` – one or more seconds within a minute.
Allowable values: [0,59]
- `minute` – one or more minutes within an hour.
Allowable values: [0,59]
- `hour` – one or more hours within a day.
Allowable values: [0,23]
- `dayOfMonth` – one or more days within a month.
Allowable values:]
 - 1,31]
 - -7, -1
 - "Last"
 - {"1st", "2nd", "3rd", "4th", "5th", "Last"} {"Sun", "Mon", "Tue", "Wed", "Thu", "Fri", "Sat"}

"Last" means the last day of the month.
-x (where x is in the range [-7, -1]) means x day(s) before the last day of the month

"1st", "2nd", etc. applied to a day of the week identifies a single occurrence of that day within the month.

- `month` – one or more months within a year.
Allowable values:]
 - [1,12]
 - {"Jan", "Feb", "Mar", "Apr", "May", "Jun", "Jul", "Aug", "Sep", "Oct", "Nov", "Dec"}
- `dayOfWeek` – one or more days within a week.
Allowable values:]
 - [0,7]
 - {"Sun", "Mon", "Tue", "Wed", "Thu", "Fri", "Sat"}

"0" and "7" both refer to Sunday
- `year` – a particular calendar year.
Allowable values: a four-digit calendar year

A.2.23.1.2 Forms of Supported Element Values Each element supports values expressed in one of the following forms:

- **Single Value** – Constrains the attribute to only one of its possible values.
Example: `second = "10"`
Example: `month = "Sep"`
- **Wild Card** – "*" Represents all allowable values for a given attribute.
Example: `second = "*"`
Example: `dayOfWeek = "*"`
- **List** – Constrains the attribute to two or more allowable values or ranges, with a comma used as a separator character within the string. Each item in the list must be a single value or range. List items cannot be lists, wild cards, or increments. Duplicate values are ignored.
Example: `second = "10,20,30"`
Example: `dayOfWeek = "Mon,Wed,Fri"`
Example: `minute = "0-10,30,40"`
- **Range** – Constrains the attribute to an inclusive range of values, with a dash separating both ends of the range. Each side of the range must be a single attribute value. Members of a range cannot be lists, wild cards, ranges, or increments. If `x` is larger than `y` in a range "`x-y`", the range is equivalent to "`x-max, min-y`", where `max` is the largest value of the corresponding attribute and `min` is the smallest. The range "`x-x`", where both range values are the same, evaluates to the single value `x`. The day of the week range "`0-7`" is equivalent to "*".
Example: `second = "1-10"`
Example: `dayOfWeek = "Fri-Mon"`
Example: `dayOfMonth = "27-3"` (Equivalent to "`27-Last, 1-3`")
- **Increments** – The forward slash constrains an attribute based on a starting point and an interval, and is used to specify every `N` seconds, minutes, or hours within the minute, hour, or day, respectively. For the expression `x/y`, the attribute is constrained to every `y`th value within the set of allowable values beginning at time `x`. The `x` value is inclusive. The wild card character (`*`) can be used in the `x`

position, and is equivalent to 0. The use of increments is only supported within the `second`, `minute`, and `hour` elements. For the `second` and `minute` elements, `x` and `y` must each be in the range `[0, 59]`. For the `hour` element, `x` and `y` must each be in the range `[0, 23]`.

Example: `minute = "?/5"` (Every five minutes within the hour)

This is equivalent to: `minute = "0,5,10,15,20,25,30,35,40,45,50,55"`

Example: `second = "30/10"` (Every 10 seconds within the minute, starting at second 30)

This is equivalent to: `second = "30,40,50"`

Note that the set of matching increment values stops once the maximum value for that attribute is exceeded. It does not "roll over" past the boundary.

Example : (`minute = "?/14"`, `hour="1,2"`)

This is equivalent to: (`minute = "0,14,28,42,56"`, `hour = "1,2"`) (Every 14 minutes within the hour, for the hours of 1 and 2 a.m.)

A.2.23.1.3 Additional Rules for Schedule Specification Elements The following additional rules apply to the schedule specification elements.

- If the `dayOfMonth` element has a non-wildcard value and the `dayOfWeek` element has a non-wildcard value, then any day matching either the `dayOfMonth` value or the `dayOfWeek` value will be considered to apply.
- Whitespace is ignored, except for string constants and numeric values.
- All string constants (e.g., "Sun", "Jan", "1st", etc.) are case insensitive.

Schedule-based timer times are evaluated in the context of the default time zone associated with the container in which the application is executing. A schedule-based timer may optionally override this default and associate itself with a specific time zone. If the schedule-based timer is associated with a specific time zone, all its times are evaluated in the context of that time zone, regardless of the default time zone in which the container is executing.

The timeout callback method to which the `Schedule` annotation is applied must have one of the following signatures, where `<METHOD>` designates the method name:

```
void <METHOD>()
void <METHOD>(Timer timer)
```

A timeout callback method can have public, private, protected, or package level access. A timeout callback method must not be declared as final or static. Timeout callback methods must not throw application exceptions.

A.2.23.2 Attributes

The following table summarizes the attributes.

Table A-16 *Attributes of the `javax.ejb.Schedule` Annotation*

Name	Description	Data Type	Required?
<code>dayOfMonth</code>	Specifies one or more days within a month. Default value is <code>*</code> .	String	No
<code>dayOfWeek</code>	Specifies one or more days within a week. Default value is <code>*</code> .	String	No

Table A–16 (Cont.) Attributes of the `javax.ejb.Schedule` Annotation

Name	Description	Data Type	Required?
hour	Specifies one or more hours within a day. Default value is 0.	String	No
info	Specifies an information string that is associated with the timer. Default value is 0.	String	No
minute	Specifies one or more minutes within a hour. Default value is 0.	String	No
month	Specifies one or more months within a year. Default value is *.	String	No
persistent	Specifies whether the timer that is created is persistent. Valid values for this attribute are <code>true</code> and <code>false</code> . Default value is <code>true</code> .	Boolean	No
second	Specifies one or more seconds with in a minute. Default value is 0.	String	No
timezone	Specifies the time zone within which the schedule is evaluated. Time zones are specified as an ID string. The set of required time zone IDs is defined by the <code>Zone Name (TZ)</code> column of the public domain <code>zoneinfo</code> database. Default value: If a <code>timezone</code> is not specified, the schedule is evaluated in the context of the default <code>timezone</code> associated with the container in which the application is executing.	String	No
year	Specifies one or more years. Default value is *.	String	No

A.2.24 `javax.ejb.Schedules`

New Section - added this section per the 3.1 spec.

The following sections describe the annotation in more detail.

A.2.24.1 Description

Target: Method

Applied to a timer callback method to schedule multiple calendar-based timers for the method. The method to which the `Schedules` annotation is applied must have one of the following signatures, where `<METHOD>` designates the method name:

```
void <METHOD>()
void <METHOD>(Timer timer)
```

A.2.24.2 Attributes

The following table summarizes the attributes.

Table A-17 Attributes of the `javax.ejb.Schedules` Annotation

Name	Description	Data Type	Required?
value	Specifies one or more calendar-based timer specifications.	Schedule[]	Yes

A.2.25 `javax.ejb.Singleton`

New Section - added this section per the 3.1 spec.

The following sections describe the annotation in more detail.

Note: This section contains "mapped-name" refs instead of "lookup-name"

A.2.25.1 Description

Target: Class

Specifies that the Enterprise JavaBean is a singleton session bean.

A.2.25.2 Attributes

The following table summarizes the attributes.

Table A-18 Attributes of the `javax.ejb.Singleton` Annotation

Name	Description	Data Type	Required?
name	Specifies the name of the singleton 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 singleton session bean should be mapped. You can also use this attribute to specify the JNDI name of this singleton 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 singleton session bean may not be portable.	String	No
description	Describes the singleton session bean.	String	No.

A.2.26 `javax.ejb.Startup`

New Section - added this section per the 3.1 spec.

The following sections describe the annotation in more detail.

A.2.26.1 Description

Target: Class

Specifies that the Enterprise JavaBean is a stateful session bean.

A.2.27 `javax.ejb.StatefulTimeout`

New Section - added this section per the 3.1 spec.

The following sections describe the annotation in more detail.

A.2.27.1 Description

Target: Type

Mark a singleton bean for eager initialization during the application startup sequence.

A.2.27.2 Attributes

The following table summarizes the attributes.

Table A-19 *Attributes of the `javax.ejb.StatefulTimeout` Annotation*

Name	Description	Data Type	Required?
value	<p>Specifies the amount of time the stateful session bean can be idle.</p> <ul style="list-style-type: none"> ▪ A value > 0 indicates a timeout value in the units specified by the unit element. ▪ A value of 0 means concurrent access is not permitted. ▪ A value of -1 indicates that the client request will block indefinitely until forward progress it can proceed. <p>Values less than -1 are not valid.</p>	Long	No
unit	<p>Specifies the units used for the specified value.</p> <p>The default value for this attribute is <code>java.util.concurrent.TimeUnit.MINUTES</code>.</p>	TimeUnit	No

A.2.28 `javax.ejb.Stateless`

The following sections describe the annotation in more detail.

Note: This section contains "mapped-name" refs instead of "lookup-name"

A.2.28.1 Description

Target: Class

Specifies that the Enterprise JavaBean is a stateless session bean.

A.2.28.2 Attributes

The following table summarizes the attributes.

Table A–20 Attributes of the `javax.ejb.Stateless` Annotation

Name	Description	Data Type	Required?
name	<p>Specifies the name of the stateless 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 stateless session bean should be mapped.</p> <p>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:</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 stateless session bean may not be portable.</p>	String	No
description	Describes the stateless session bean.	String	No.

A.2.29 `javax.ejb.Timeout`

The following sections describe the annotation in more detail.

A.2.29.1 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.x, your bean class was required to implement `javax.ejb.TimerService` 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.x of EJB. You no longer are required to implement the `javax.ejb.TimerService` 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.x way of programming the timer service if you want.

For details, see [Section 4.6, "Optionally Program the EJB Timer Service."](#)

This annotation does not have any attributes.

A.2.30 javax.ejb.TransactionAttribute

The following sections describe the annotation in more detail.

A.2.30.1 Description

Target: Class, Method

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

Note: 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.

A.2.30.2 Attributes

The following table summarizes the attributes.

Table A–21 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 trans-attribute element in the "Container-Managed Transactions Elements" table in <i>Programming Enterprise JavaBeans, Version 2.1, for Oracle WebLogic Server</i>.</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>	String	No.

A.2.31 javax.ejb.TransactionManagement

The following sections describe the annotation in more detail.

A.2.31.1 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" in *Programming Enterprise JavaBeans, Version 2.1, for Oracle WebLogic Server*.

A.2.31.2 Attributes

The following table summarizes the attributes.

Table A–22 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>	String	No.

A.3 Annotations Used to Configure Interceptors

This section provides reference information for the following annotations:

- [Section A.3.1, "javax.interceptor.AroundInvoke"](#)
- [Section A.3.2, "javax.interceptor.ExcludeClassInterceptors"](#)
- [Section A.3.3, "javax.interceptor.ExcludeDefaultInterceptors"](#)
- [Section A.3.4, "javax.interceptor.Interceptors"](#)

A.3.1 javax.interceptor.AroundInvoke

The following sections describe the annotation in more detail.

A.3.1.1 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.

A.3.2 `javax.interceptor.ExcludeClassInterceptors`

The following sections describe the annotation in more detail.

A.3.2.1 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.

A.3.3 `javax.interceptor.ExcludeDefaultInterceptors`

The following sections describe the annotation in more detail.

A.3.3.1 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.

A.3.4 `javax.interceptor.Interceptors`

The following sections describe the annotation in more detail.

A.3.4.1 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).

A.3.4.2 Attributes

The following table summarizes the attributes.

Table A-23 *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

A.4 Annotations Used to Interact With Entity Beans

This section provides reference information about the following annotations:

- [Section A.4.1, "javax.persistence.PersistenceContext"](#)
- [Section A.4.2, "javax.persistence.PersistenceContexts"](#)
- [Section A.4.3, "javax.persistence.PersistenceUnit"](#)
- [Section A.4.4, "javax.persistence.PersistenceUnits"](#)

A.4.1 `javax.persistence.PersistenceContext`

The following sections describe the annotation in more detail.

A.4.1.1 Description

Target: Class, Method, Field

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

You use this annotation to interact with a 3.x 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.

A.4.1.2 Attributes

The following table summarizes the attributes.

Table A-24 Attributes of the `javax.persistence.PersistenceContextAnnotation`

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
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>	<code>PersistenceContextType</code>	No

A.4.2 `javax.persistence.PersistenceContexts`

The following sections describe the annotation in more detail.

A.4.2.1 Description

Target: Class

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

A.4.2.2 Attributes

The following table summarizes the attributes.

Table A–25 *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.

A.4.3 `javax.persistence.PersistenceUnit`

The following sections describe the annotation in more detail.

A.4.3.1 Description

Target: Class, Method, Field

Specifies a dependency on an `EntityManagerFactory` object.

You use this annotation to interact with a 3.x 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.

A.4.3.2 Attributes

The following table summarizes the attributes.

Table A–26 *Attributes of the `javax.persistence.PersistenceUnit` Annotation*

Name	Description	Data Type	Required?
name	Specifies the name by which the <code>EntityManagerFactory</code> is to be known within the context of the session or message-driven bean You are not required to specify this attribute if you use dependency injection, only if you also use JNDI to look up information.	String	No

Table A–26 (Cont.) Attributes of the `javax.persistence.PersistenceUnit` Annotation

Name	Description	Data Type	Required?
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

A.4.4 `javax.persistence.PersistenceUnits`

The following sections describe the annotation in more detail.

A.4.4.1 Description

Target: Class

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

A.4.4.2 Attributes

The following table summarizes the attributes.

Table A–27 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

A.5 Standard JDK Annotations Used By EJB 3.x

This section provides reference information about the following annotations:

- [Section A.5.1, "javax.annotation.PostConstruct"](#)
- [Section A.5.2, "javax.annotation.PreDestroy"](#)
- [Section A.5.3, "javax.annotation.Resource"](#)
- [Section A.5.4, "javax.annotation.Resources"](#)

A.5.1 `javax.annotation.PostConstruct`

The following sections describe the annotation in more detail.

A.5.1.1 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.

A.5.2 `javax.annotation.PreDestroy`

The following sections describe the annotation in more detail.

A.5.2.1 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.

A.5.3 `javax.annotation.Resource`

The following sections describe the annotation in more detail.

Note: This section contains "mapped-name" refs instead of "lookup-name"

A.5.3.1 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.

A.5.3.2 Attributes

The following table summarizes the attributes.

Table A–28 *Attributes of the `javax.annotation.Resource` Annotation*

Name	Description	Data Type	Required?
name	Specifies the name of the resource reference. 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.	String	No
type	Specifies the Java data type of the resource. 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.	Class	No
authenticationType	Specifies the authentication type to use for the resource. You specify this attribute only for resources representing a connection factory of any supported type. Valid values for this attribute are: <ul style="list-style-type: none"> ■ <code>AuthenticationType.CONTAINER</code> ■ <code>AuthenticationType.APPLICATION</code> Default value is <code>AuthenticationType.CONTAINER</code>	AuthenticationType	No
shareable	Indicates whether a resource can be shared between this EJB and other EJBs. You specify this attribute only for resources representing a connection factory of any supported type or ORB object instances. Valid values for this attribute are <code>true</code> and <code>false</code> . Default value is <code>true</code> .	boolean	No

Table A–28 (Cont.) Attributes of the `javax.annotation.Resource` Annotation

Name	Description	Data Type	Required?
<code>mappedName</code>	Specifies the global JNDI name of the dependent resource. For example: <code>mappedName="my.Datasource"</code> specifies that the JNDI name of the dependent resources is <code>my.Datasource</code> and is deployed in the WebLogic Server JNDI tree.	String	No
<code>description</code>	Specifies a description of the resource.	String	No

A.5.4 `javax.annotation.Resources`

The following sections describe the annotation in more detail.

A.5.4.1 Description

Target: Class

Specifies an array of `@Resource` annotations.

A.5.4.2 Attributes

The following table summarizes the attributes.

Table A–29 Attributes of the `javax.annotation.Resources` Annotation

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

A.6 Standard Security-Related JDK Annotations Used by EJB 3.x

This section provides reference information about the following annotations:

- [Section A.6.1, "javax.annotation.security.DeclareRoles"](#)
- [Section A.6.2, "javax.annotation.security.DenyAll"](#)
- [Section A.6.3, "javax.annotation.security.PermitAll"](#)
- [Section A.6.4, "javax.annotation.security.RolesAllowed"](#)
- [Section A.6.5, "javax.annotation.security.RunAs"](#)

A.6.1 `javax.annotation.security.DeclareRoles`

The following sections describe the annotation in more detail.

A.6.1.1 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" in the *Oracle WebLogic Server Administration Console Help*.

A.6.1.2 Attributes

The following table summarizes the attributes.

Table A-30 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.

A.6.2 `javax.annotation.security.DenyAll`

The following sections describe the annotation in more detail.

A.6.2.1 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.

A.6.3 `javax.annotation.security.PermitAll`

The following sections describe the annotation in more detail.

A.6.3.1 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.

A.6.4 `javax.annotation.security.RolesAllowed`

The following sections describe the annotation in more detail.

A.6.4.1 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" in the *Oracle WebLogic Server Administration Console Help*.

A.6.4.2 Attributes

The following table summarizes the attributes.

Table A–31 *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.

A.6.5 `javax.annotation.security.RunAs`

The following sections describe the annotation in more detail.

A.6.5.1 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" in the *Oracle WebLogic Server Administration Console Help*.

A.6.5.2 Attributes

The following table summarizes the attributes.

Table A–32 *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.

A.7 WebLogic Annotations

This section provides reference information for the following WebLogic annotations:

- [Section A.7.1, "weblogic.javaee.AllowRemoveDuringTransaction"](#)
- [Section A.7.2, "weblogic.javaee.CallByReference"](#)
- [Section A.7.3, "weblogic.javaee.DisableWarnings"](#)
- [Section A.7.4, "weblogic.javaee.EJBReference"](#)
- [Section A.7.5, "weblogic.javaee.Idempotent"](#)
- [Section A.7.6, "weblogic.javaee.JMSClientID"](#)
- [Section A.7.7, "weblogic.javaee.JNDIName"](#)
- [Section A.7.8, "weblogic.javaee.JNDINames"](#)
- [Section A.7.9, "weblogic.javaee.MessageDestinationConfiguration"](#)
- [Section A.7.10, "weblogic.javaee.TransactionIsolation"](#)
- [Section A.7.11, "weblogic.javaee.TransactionTimeoutSeconds"](#)

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 Enterprise JavaBeans, Version 2.1, for Oracle WebLogic Server*.

A.7.1 `weblogic.javaee.AllowRemoveDuringTransaction`

The following sections describe the annotation in more detail.

A.7.1.1 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 Enterprise JavaBeans, Version 2.1, for Oracle WebLogic Server*.

A.7.2 `weblogic.javaee.CallByReference`

The following sections describe the annotation in more detail.

A.7.2.1 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 Enterprise JavaBeans, Version 2.1, for Oracle WebLogic Server*.

A.7.3 `weblogic.javaee.DisableWarnings`

The following sections describe the annotation in more detail.

A.7.3.1 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 Enterprise JavaBeans, Version 2.1, for Oracle WebLogic Server*.

A.7.3.2 Attributes

The following table summarizes the attributes.

Table A-33 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

A.7.4 weblogic.javaee.EJBReference

The following sections describe the annotation in more detail.

A.7.4.1 Description

Target: Class, Method, Field

Maps EJB reference name to its JNDI name.

A.7.4.2 Attribute

The following table summarizes the attributes.

Table A-34 Attribute of the `weblogic.javaee.EJBReference` 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	Yes
jndiName	Specifies the JNDI name of an actual EJB, resource, or reference available in WebLogic Server.	String	Yes

A.7.5 `weblogic.javaee.Idempotent`

The following sections describe the annotation in more detail.

A.7.5.1 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 Enterprise JavaBeans, Version 2.1, for Oracle WebLogic Server*.

A.7.5.2 Attributes

The following table summarizes the attributes.

Table A-35 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

A.7.6 `weblogic.javaee.JMSClientID`

The following sections describe the annotation in more detail.

A.7.6.1 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 [Section A.7.9, "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 [Section A.7.9, "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 Enterprise JavaBeans, Version 2.1, for Oracle WebLogic Server*.

A.7.6.2 Attributes

The following table summarizes the attributes.

Table A-36 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

A.7.7 `weblogic.javaee.JNDIName`

The following sections describe the annotation in more detail.

A.7.7.1 Description

Target: Class (Stateful or stateless session EJBs only)

Specifies a custom JNDI name that can be applied to a bean class for a certain client view. When applied to a bean class to indicate the JNDI name of a no-interface view, the `className` is optional.

Note: This annotation is overridden by the `jndi-binding` element in the `weblogic-ejb-jar.xml` deployment descriptor. For more information, see "weblogic-ejb-jar.xml Deployment Descriptor Reference" in *Programming Enterprise JavaBeans, Version 2.1, for Oracle WebLogic Server*.

A.7.7.2 Attributes

The following table summarizes the attributes.

Table A-37 Attributes of the `weblogic.javaee.JNDIName`

Name	Description	Data Type	Required?
classname	Class name of the client view.	String	No
value	JNDI name of the client view.	String	No

A.7.8 `weblogic.javaee.JNDINames`

The following sections describe the annotation in more detail.

A.7.8.1 Description

Target: Class (Stateful or stateless session EJBs only)

Specifies the multiple, custom JNDI names that can be applied to an EJB.

A.7.8.2 Attributes

The following table summarizes the attributes.

Table A–38 *Attributes of the `weblogic.javaee.JNDINames`*

Name	Description	Data Type	Required?
value	Multiple, custom JNDI names for the EJB.	JNDIName	No

A.7.9 `weblogic.javaee.MessageDestinationConfiguration`

The following sections describe the annotation in more detail.

A.7.9.1 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.

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 Enterprise JavaBeans, Version 2.1, for Oracle WebLogic Server*.

A.7.9.2 Attributes

The following table summarizes the attributes.

Table A–39 *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

A.7.10 `weblogic.javaee.TransactionIsolation`

The following sections describe the annotation in more detail.

A.7.10.1 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 Enterprise JavaBeans, Version 2.1, for Oracle WebLogic Server*.

A.7.10.2 Attributes

The following table summarizes the attributes.

Table A-40 *Attributes of the `weblogic.javaee.Idempotent`*

Name	Description	Data Type	Required?
IsolationLevel	Isolation level. Valid values include: <ul style="list-style-type: none"> ▪ READ_COMMITTED—Transaction can view only committed updates from other transactions. ▪ READ_UNCOMMITTED—Transactions can view uncommitted updates from other transactions. ▪ REPEATABLE_READ—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. ▪ SERIALIZABLE—Simultaneously executing this transaction multiple times has the same effect as executing the transaction multiple times in a serial fashion. 	int	No
This attribute defaults to DEFAULT.			

A.7.11 `weblogic.javaee.TransactionTimeoutSeconds`

The following sections describe the annotation in more detail.

A.7.11.1 Description

Target: Class

Defines the timeout for transactions in seconds.

A.7.11.2 Attributes

The following table summarizes the attributes.

Table A-41 *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

Using Oracle Kodo with Oracle WebLogic Server

This appendix provides an overview of developing, deploying, and configuring an Oracle Kodo application using WebLogic Server.

Kodo Deprecation Notes: Oracle Kodo JPA/JDO is deprecated in this release. Customers are encouraged to use Oracle TopLink, a JPA 2.0 provider. For more information about Oracle TopLink, see [Chapter 8, "Configuring the Persistence Provider in Oracle WebLogic Server."](#) Just as Kodo is deprecated, the Kodo documentation is also deprecated. For current information about using Kodo in WebLogic Server, see [Section 8.4, "Using Oracle Kodo in Oracle WebLogic Server."](#)

This release of WebLogic Server runs with the JPA 2.0 JAR in the server's classpath. Although JPA 2.0 is upwardly compatible with JPA 1.0, JPA 2.0 introduced some methods to existing JPA interfaces that conflict with existing signatures in OpenJPA interfaces. As a result, applications that continue to use Kodo/JPA as the persistence provider with WebLogic Server 12.1.1 must be recompiled. For more information, see [Section 8.4.3, "Updating Applications to Overcome Conflicts with JPA 2.0."](#)

Kodo does not support JPA 2.0. Trying to use JPA 2.0 APIs with Kodo in WebLogic Server will produce an error. However, using Kodo as the persistence provider for a specific persistence unit or as the default provider for the domain does not prevent you from using TopLink for a different persistence unit. For more information, see [Section 8.4, "Using Oracle Kodo in Oracle WebLogic Server."](#)

- [Section B.1, "Overview of Oracle Kodo"](#)
- [Section B.2, "Creating an Oracle Kodo Application"](#)
- [Section B.3, "Using Different Oracle Kodo Versions"](#)
- [Section B.4, "Configuring Persistence"](#)
- [Section B.5, "Deploying an Oracle Kodo Application"](#)
- [Section B.6, "Configuring an Oracle Kodo Application"](#)

B.1 Overview of Oracle Kodo

Oracle Kodo is an implementation of the 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 Server's overall Enterprise JavaBean 3.0 persistence implementation.

For general information on creating an application using JPA and JDO, see the *Kodo 4.2.0 Developers Guide for JPA/JDO*.

B.2 Creating an Oracle 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.

B.3 Using Different Oracle Kodo Versions

If you choose to use a different version of Oracle Kodo than the one provided by default within WebLogic Server, you must use the `FilteringClassLoader` to specify the packages—in this case Oracle Kodo and OpenJPA—that you want to be loaded from the application, and not from the system classloader.

The following example shows how to load the Oracle Kodo and OpenJPA packages from the application 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 for Oracle WebLogic Server*.

Then, you can package the Oracle Kodo and OpenJPA packages with your application using the library directory, as described in "Library Directories" in *Developing Applications for Oracle WebLogic Server*.

B.4 Configuring Persistence

The following sections describe how to configure persistence.

- [Section B.4.1, "Editing the Configuration Property Files"](#)
- [Section B.4.2, "Using the Configuration Files Together"](#)
- [Section B.4.3, "Configuring Plug-ins"](#)

B.4.1 Editing the Configuration Property Files

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

Table B-1 Persistence Configuration Files

Configuration File	Description
<code>persistence.xml</code>	<p>Oracle 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 4.2.0 Developers Guide for JPA/JDO</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://xmlns.oracle.com/weblogic/persistence-configuration/1.0/persistence-configuration.xsd.</p> <p>Note: The <code>weblogic.jar</code> file must be in the <code>CLASSPATH</code> 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

B.4.2 Using the Configuration Files Together

The following provide considerations for using the configuration files together.

- When using the `persistence-configuration.xml` file, all Oracle 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 Oracle Kodo persistence unit and provide advanced configuration and tunables support.

- If Oracle 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 Oracle Kodo (`kodo`) or OpenJPA (`openjpa`) properties can be specified in the `<properties>` tag of the `persistence.xml` file for that persistence unit.

B.4.3 Configuring Plug-ins

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

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

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

B.5 Deploying an Oracle Kodo Application

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

- Create a Oracle Kodo application, as described in [Section B.2, "Creating an Oracle Kodo Application."](#)
- Configure persistence, as described in [Section B.4, "Configuring Persistence."](#)
- Created an archive for your application (`.ear` or `.war`).

Once completed, you are ready to deploy your application on WebLogic Server. Once your application is configured, an Oracle Kodo application is deployed just like any other application. For complete information on deploying applications, see *Deploying Applications to Oracle WebLogic Server*.

B.6 Configuring an Oracle 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 Oracle Kodo application, you can alter the configuration parameters defined in `persistence.xml` and `persistence-configuration.xml`.

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

B.6.1 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 Oracle Kodo persistence parameters that can be edited from the Administration Console.

B.6.2 Configuring Oracle 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 Deployment for Oracle WebLogic Server*.
- Use the `WLST loadApplication()` method to load and update the application deployment plan. For more information, see "Updating the Deployment Plan" in *Oracle 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 Oracle WebLogic Server*.

Oracle Kodo Persistence Configuration Schema Reference

This appendix describes the namespace, schema location, file structure, and elements in the Kodo-specific deployment descriptor, `persistence-configuration.xml`.

Kodo Deprecation Notes: Oracle Kodo JPA/JDO is deprecated in this release of WebLogic Server. Customers are encouraged to consider using Oracle TopLink. For more information about Oracle TopLink, see ["Configuring the Persistence Provider in Oracle WebLogic Server"](#).

WebLogic Server does not support JPA 2.0 when using Kodo. For more information, see [Section 8.3, "Using Oracle TopLink in Oracle WebLogic Server."](#)

- [Section C.1, "persistence-configuration.xml Namespace Declaration and Schema Location"](#)
- [Section C.2, "persistence-configuration.xml Deployment Descriptor File Structure"](#)
- [Section C.3, "persistence-configuration.xml Deployment Descriptor Elements"](#)

C.1 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://xmlns.oracle.com/weblogic/persistence-configuration"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="http://xmlns.oracle.com/weblogic/persistence-configuration
http://xmlns.oracle.com/weblogic/persistence-configuration/1.0/persistence-configu
ration.xsd">
...
</persistence-configuration>
```

C.2 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](#) | [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](#)
 * [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

C.3 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.

C.4 abstract-store-broker-factory

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.4.1 Function

Specifies a plugin string (see "Plugin Configuration" in *Kodo 4.2.0 Developers Guide for JPA/JDO*) describing the `kodo.kernel.BrokerFactory` type to use, in this case abstract store. For more information, see "kodo.BrokerFactory" in the *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.4.2 Example

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

C.5 access-dictionary

Range of values: N/A

Default value: N/A

Parent elements:

persistence-configuration
 persistence-configuration-unit

C.5.1 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 *Oracle WebLogic Server Administration Console Help*.

C.5.2 Example

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

```
<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-distinct-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</supprots-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-multiple-nontransac
tional-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-f
or-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-join>
<supports-correlated-subselect>>true</supports-correlated-subselect>

<supports-null-table-for-get-imported-keys>>false</supports-null-table-for-get-impo

```

```

rted-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-i
nfo>
  <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-join>
  <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-for-get-prim
ary-keys>
  <system-schemas></system-schemas>
  <requires-cast-for-math-functions>false</requires-cast-for-math-functions>
  <supports-null-delete-action>true</supports-null-delete-action>
  <requires-auto-commit-for-meta-data>false</requires-auto-commit-for-meta-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>

```

C.6 access-unloaded

Range of values: true | false

Default value: true

Parent elements:

```

persistence-configuration
  persistence-configuration-unit

```

```

detach-options-all

persistence-configuration
  persistence-configuration-unit
    detach-options-fetch-groups

persistence-configuration
  persistence-configuration-unit
    detach-options-loaded

```

C.6.1 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 4.2.0 Developers Guide for JPA/JDO*.

C.6.2 Example

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

C.7 action

Range of values: warn | exception

Default value: N/A

Parent elements:

```

persistence-configuration
  persistence-configuration-unit
    inverse-manager

```

C.7.1 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 4.2.0 Developers Guide for JPA/JDO*.

C.7.2 Example

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

C.8 addresses

Range of values: Valid IP addresses

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.8.1 Function

Specifies a semicolon separated list of IP addresses to which notifications should be sent. For more information, see "TCP" in the *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.8.2 Example

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

C.9 advanced-sql

Range of values: String

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.9.1 Function

Configures advanced SQL features. For more information, see "kodo.jdbc.SQLFactory" in the *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.10 aggregate-listeners

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.10.1 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 4.2.0 Developers Guide for JPA/JDO*.

C.10.2 Example

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

C.11 allocate

Range of values: Integer

Default value: 50

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
    class-table-jdbc-seq
```

```
persistence-configuration
  persistence-configuration-unit
    native-jdbc-seq
```

```
persistence-configuration
  persistence-configuration-unit
    table-jdbc-seq
```

```
persistence-configuration
  persistence-configuration-unit
    value-table-jdbc-seq
```

C.11.1 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 4.2.0 Developers Guide for JPA/JDO*.

C.11.2 Example

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

C.12 assert-allowed-type

Range of values: true | false

Default value: false

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.12.1 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 4.2.0 Developers Guide for JPA/JDO*.

C.12.2 Example

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

C.13 auto-clear

Range of values:

`persistence-configuration-unit`: datastore | all

`kodo-broker`: 0 | 1

Default value: datastore (0)

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

```
persistence-configuration
  persistence-configuration-unit
    kodo-broker
```

C.13.1 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.

C.13.2 Example

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

C.14 auto-detach

Range of values:

`auto-detaches`: close | commit | nontx-read

`kodo-broker`: 0 | 1 | 2

Default value: close (0)

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
    auto-detaches
```

```
persistence-configuration
  persistence-configuration-unit
    kodo-broker
    auto-detach
```

C.14.1 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 4.2.0 Developers Guide for JPA/JDO*.

C.14.2 Example

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

C.15 auto-detaches

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.15.1 Function

Specifies a list of events on which the managed instances are automatically detached. Specifies one or more [Section C.14, "auto-detach"](#) elements. For more information, see "Automatic Detachment" in *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.15.2 Example

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

C.16 base-name

Range of values: String

Default value: N/A

Parent elements:

```
profiling
  export-profiling
```

C.16.1 Function

Defines the basename of the exported data file to create. See "Dumping Profiling Data to Disk from a Batch Process" in *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.16.2 Example

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

C.17 batching-operation-order-update-manager

Range of values: N/A

Default value: N/A

Parent elements:

profiling

C.17.1 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 4.2.0 Developers Guide for JPA/JDO*.

C.17.2 Example

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

C.18 buffer-size

Range of values: Integer

Default value: 10

Parent elements:

persistence-configuration
 persistence-configuration-unit
 cluster-remote-commit-provider

C.18.1 Function

Specifies the buffer size of the remote commit provider. For more information, see "Remote Commit Provider Configuration" in the *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.18.2 Example

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

C.19 cache-map

Range of values: N/A

Default value: N/A

Parent elements:

persistence-configuration
 persistence-configuration-unit

C.19.1 Function

Defines the cache map used to associate query strings and their parsed form. For more information, see "Query Compilation Cache" in the *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.19.2 Example

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

C.20 cache-size

Range of values: Integer

Default value: 1000

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
    kodo-concurrent-data-cache
```

```
persistence-configuration
  persistence-configuration-unit
    lru-data-cache
```

```
persistence-configuration
  persistence-configuration-unit
    kodo-query-cache
```

```
persistence-configuration
  persistence-configuration-unit
    lru-query-cache
```

C.20.1 Function

Set the data or query cache size. For more information, see "Configuring the Data Cache Size" and "Query Cache" in *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.20.2 Example

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

C.21 channel

Range of values: String

Default value: openjpa.Runtime

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.21.1 Function

Specifies the channel to which you want to log. For more information, see "Orphaned Keys" in the *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.21.2 Example

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

C.22 class-table-jdbc-seq

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.22.1 Function

Specifies a plugin string (see "Plugin Configuration" in *Kodo 4.2.0 Developers Guide for JPA/JDO*) 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 4.2.0 Developers Guide for JPA/JDO*.

C.22.2 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>
```

C.23 classname

Range of values: Valid classname

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
    aggregate-listener
      custom-aggregate-listener

persistence-configuration
  persistence-configuration-unit
```

```

    custom-broker-factory

persistence-configuration
  persistence-configuration-unit
    custom-broker-impl

persistence-configuration
  persistence-configuration-unit
    custom-class-resolver

persistence-configuration
  persistence-configuration-unit
    custom-compatibility

persistence-configuration
  persistence-configuration-unit
    custom-connection-decorator

persistence-configuration
  persistence-configuration-unit
    custom-data-cache

persistence-configuration
  persistence-configuration-unit
    custom-dictionary

```

C.23.1 Function

Name of the class associated with the custom feature.

C.24 classpath-scan

Range of values: true | false

Default value: true

Parent elements:

```

persistence-configuration
  persistence-configuration-unit
    deprecated-jdo-meta-data-factory

persistence-configuration
  persistence-configuration-unit
    extension-deprecated-jdo-mapping-factory

persistence-configuration
  persistence-configuration-unit
    jdo-meta-data-factory

persistence-configuration
  persistence-configuration-unit
    jdor-mapping-factory

persistence-configuration
  persistence-configuration-unit
    kodo-persistence-mapping-factory

persistence-configuration
  persistence-configuration-unit

```

```
kodo-persistence-meta-data-factory

persistence-configuration
  persistence-configuration-unit
    mapping-file-deprecated-jdo-mapping-factory

persistence-configuration
  persistence-configuration-unit
    orm-file-jdor-mapping-factory

persistence-configuration
  persistence-configuration-unit
    table-deprecated-jdo-mapping-factory
```

C.24.1 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 4.2.0 Developers Guide for JPA/JDO*.

C.24.2 Example

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

C.25 clear-on-close

Range of values: true | false

Default value: false

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
    tangosol-data-cache
```

C.25.1 Function

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

C.25.2 Example

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

C.26 client-broker-factory

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.26.1 Function

A plugin string (see "Plugin Configuration" in *Kodo 4.2.0 Developers Guide for JPA/JDO*) describing the `kodo.kernel.BrokerFactory` type to use, in this case `remote`. For more information, see "kodo.BrokerFactory" in the *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.26.2 Example

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

C.27 close-on-managed-commit

Range of values: true | false

Default value: true

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
    compatibility
```

C.27.1 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 4.2.0 Developers Guide for JPA/JDO*.

C.27.2 Example

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

C.28 cluster-remote-commit-provider

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.28.1 Function

Configures Kodo to share commit notifications among persistenceManagerFactories in a cluster. For more information, see "Remote Commit Provider Configuration" in the *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.28.2 Example

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

C.29 commons-log-factory

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.29.1 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 4.2.0 Developers Guide for JPA/JDO*.

C.29.2 Example

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

C.30 compatibility

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.30.1 Function

Configure backwards compatibility settings for certain runtime behaviors. For more information, see "Compatibility Configuration" in *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.30.2 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>

```

C.31 concurrent-hash-map

Range of values: N/A

Default value: N/A

Parent elements:

```

persistence-configuration
  persistence-configuration-unit

```

C.31.1 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 4.2.0 Developers Guide for JPA/JDO*.

C.31.2 Example

```

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

```

C.32 connection-decorators

Range of values: N/A

Default value: N/A

Parent elements:

```

persistence-configuration
  persistence-configuration-unit

```

C.32.1 Function

A plugin string (see "Plugin Configuration" in *Kodo 4.2.0 Developers Guide for JPA/JDO*) describing `org.apache.openjpa.lib.jdbc.ConnectionDecorator` implementations to install on the connection factory. Specify one or more [Section C.57](#), "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 4.2.0 Developers Guide for JPA/JDO*.

C.32.2 Example

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

C.33 connection-driver-name

Range of values: Valid classname

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

```
persistence-configuration
  persistence-configuration-unit
  kodo-pooling-data-source
```

```
persistence-configuration
  persistence-configuration-unit
  simple-driver-data-source
```

C.33.1 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 4.2.0 Developers Guide for JPA/JDO*.

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.jsql.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.jtdbcon.JXDBConDriver|
postgres95.PGDriver
postgresql.Driver
solid.jdbc.SolidDriver
sun.jdbc.odbc.JdbcOdbcDriver
weblogic.jdbc.mssqlserver4.Driver
weblogic.jdbc.pool.Driver

```

C.33.2 Example

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

C.34 connection-factory-mode

Range of values: local | managed

Default value: local

Parent elements:

```

persistence-configuration
  persistence-configuration-unit

```

C.34.1 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 4.2.0 Developers Guide for JPA/JDO*.

C.34.2 Example

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

C.35 connection-factory-name

Range of values: JNDI name

Default value: N/A

Parent elements:

persistence-configuration
 persistence-configuration-unit

C.35.1 Function

The JNDI location of a `javax.sql.DataSource` to use to connect to the database. For more information, see "JDBC" in *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.35.2 Example

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

C.36 connection-factory-properties

Range of values: N/A

Default value: N/A

Parent elements:

persistence-configuration
 persistence-configuration-unit

C.36.1 Function

A plugin string (see "Plugin Configuration" in *Kodo 4.2.0 Developers Guide for JPA/JDO*) listing properties for configuration of the `javax.sql.DataSource`, described at <http://download.oracle.com/javase/6/docs/api/javax/sql/DataSource.html>, in use. For more information, see "JDBC" in *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.36.2 Example

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

C.37 connection-factory2-name

Range of values: JNDI name

Default value: N/A

Parent elements:

persistence-configuration
 persistence-configuration-unit

C.37.1 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 4.2.0 Developers Guide for JPA/JDO*.

C.37.2 Example

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

C.38 connection-factory2-properties

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.38.1 Function

Properties used to configure the `javax.sql.DataSource`, described at <http://download.oracle.com/javase/6/docs/api/javax/sql/DataSource.html>, used as the unmanaged `ConnectionFactory`. For more information, see "Managed and XA DataSources" in *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.38.2 Example

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

C.39 connection-password

Range of values: Valid password

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.39.1 Function

Password for the user specified by `connection-user-name`. For more information, see "JDBC" in *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.39.2 Example

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

C.40 connection-properties

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.40.1 Function

A plugin string (see "Plugin Configuration" in *Kodo 4.2.0 Developers Guide for JPA/JDO*) 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 4.2.0 Developers Guide for JPA/JDO*.

C.40.2 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>
```

C.41 connection-retain-mode

Range of values: always | on-demand | persistence-manager | transaction

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.41.1 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 4.2.0 Developers Guide for JPA/JDO*.

C.41.2 Example

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

C.42 connection-url

Range of values: Valid URL

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

```
persistence-configuration
  persistence-configuration-unit
  kodo-pooling-data-source
```

```
persistence-configuration
  persistence-configuration-unit
  simple-driver-data-source
```

C.42.1 Function

The JDBC URL for the database. For more information, see "JDBC" in *Kodo 4.2.0 Developers Guide for JPA/JDO*.

The following provides possible values:

```
jdbc:JSQConnect://<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=<data
base>
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=cu
rsor
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
```

C.42.2 Example

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

C.43 connection-user-name

Range of values: Valid username

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

```
persistence-configuration
  persistence-configuration-unit
    kodo-pooling-data-source
```

```
persistence-configuration
  persistence-configuration-unit
    simple-driver-data-source
```

C.43.1 Function

Username to use when connecting to the data source specified by [Section C.42](#), "[connection-url](#)." For more information, see "JDBC" in *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.43.2 Example

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

C.44 connection2-driver-name

Range of values: Valid username

Default value: N/A

Parent elements:

persistence-configuration
persistence-configuration-unit

C.44.1 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 4.2.0 Developers Guide for JPA/JDO*.

C.44.2 Example

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

C.45 connection2-password

Range of values: Valid password

Default value: N/A

Parent elements:

persistence-configuration
persistence-configuration-unit

C.45.1 Function

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

C.45.2 Example

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

C.46 connection2-properties

Range of values: N/A

Default value: N/A

Parent elements:

persistence-configuration
persistence-configuration-unit

C.46.1 Function

Properties for the unmanaged data source specified by [connection2-driver-name](#). For more information, see "Managed and XA DataSources" in *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.46.2 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>
```

C.47 connection2-url

Range of values: Valid URL

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.47.1 Function

The JDBC URL for the unmanaged datasource. For more information, see "Managed and XA DataSources" in *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.47.2 Example

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

C.48 connection2-user-name

Range of values: Valid username

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.48.1 Function

Username to use when connecting to the data source specified by [Section C.44](#), "connection2-driver-name."

C.48.2 Example

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

C.49 constraint-names

Range of values: true | false

Default value: false

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
    table-jdor-mapping-factory
```

C.49.1 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 4.2.0 Developers Guide for JPA/JDO*.

C.49.2 Example

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

C.50 constraint-update-manager

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.50.1 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 4.2.0 Developers Guide for JPA/JDO*.

C.50.2 Example

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

C.51 copy-object-ids

Range of values: true | false

Default value: false

Parent elements:

```
persistence-configuration
```

```
persistence-configuration-unit
  compatibility
```

C.51.1 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 4.2.0 Developers Guide for JPA/JDO*.

C.51.2 Example

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

C.52 custom-aggregate-listener

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
    aggregate-listeners
```

C.52.1 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 4.2.0 Developers Guide for JPA/JDO*.

C.52.2 Example

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

C.53 custom-broker-factory

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.53.1 Function

Enables you to specify a custom broker factory. For more information, see "kodo.BrokerFactory" in *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.53.2 Example

```
<custom-broker-factory>
  <classname>...</classname>
  <properties>...</properties>
</custom-broker-factory>
```

C.54 custom-broker-impl

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.54.1 Function

Enables you to define a custom broker implementation. For more information, see "kodo.BrokerImpl" in *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.54.2 Example

```
<custom-broker-impl>
  <classname>...</classname>
  <properties>...</properties>
</custom-broker-impl>
```

C.55 custom-class-resolver

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.55.1 Function

Enables you to define a custom class resolver for class name resolution. For more information, see "kodo.ClassResolver" in *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.55.2 Example

```
<custom-class-resolver>
  <classname>...</classname>
  <properties>...</properties>
</custom-class-resolver>
```

C.56 custom-compatibility

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.56.1 Function

Enables you to define your own custom compatibility flag. For more information, see "Compatibility Configuration" in *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.56.2 Example

```
<custom-compatibility>
  <classname>...</classname>
  <properties>...</properties>
</custom-compatibility>
```

C.57 custom-connection-decorator

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
    connection-decorators
```

C.57.1 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 4.2.0 Developers Guide for JPA/JDO*.

C.57.2 Example

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

C.58 custom-data-cache

Range of values: N/A

Default value: N/A

Parent elements:

```

persistence-configuration
  persistence-configuration-unit
    data-caches

```

C.58.1 Function

Enables you to define a custom data cache. For more information, see "kodo.DataCache" in *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.58.2 Example

```

<data-caches>
  <custom-data-cache>
    <name>...</name>
    <classname>...</classname>
    <properties>...</properties>
  </custom-data-cache>
</data-caches>

```

C.59 custom-data-cache-manager

Range of values: N/A

Default value: N/A

Parent elements:

```

persistence-configuration
  persistence-configuration-unit

```

C.59.1 Function

Enables you to define a custom data cache manager. For more information, see "kodo.DataCacheManager" in *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.59.2 Example

```

<custom-data-cache-manager>
  <classname>...</classname>
  <properties>...</properties>
</custom-data-cache-manager>

```

C.60 custom-detach-state

Range of values: N/A

Default value: N/A

Parent elements:

```

persistence-configuration
  persistence-configuration-unit

```

C.60.1 Function

Enables you to define a custom detach state. For more information, see "kodo.DetachState" in *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.60.2 Example

```
<custom-detach-state>
  <classname>...</classname>
  <properties>...</properties>
</custom-detach-state>
```

C.61 custom-dictionary

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.61.1 Function

Enables you to define a custom dictionary. For more information, see "Custom Dictionary Configuration" in the *Oracle WebLogic Server Administration Console Help*.

C.61.2 Example

```
<custom-dictionary>
  <name>custom-dictionary</name>
  <classname>...</classname>
  <properties>...</properties>
</custom-dictionary>
```

C.62 custom-driver-data-source

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.62.1 Function

Enables you to define a custom datasource driver. For more information, see "kodo.jdbc.DriverDataSource" in *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.62.2 Example

```
<custom-driver-data-source>
  <classname>...</classname>
  <properties>...</properties>
</custom-driver-data-source>
```

C.63 custom-filter-listener

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
    filter-listeners
```

C.63.1 Function

Enables you to define a custom filter listener. For more information, see "kodo.FilterListeners" in *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.63.2 Example

```
<filter-listeners>
  <custom-filter-listener>
    <classname>...</classname>
    <properties>...</properties>
  </custom-filter-listener>
</filter-listeners>
```

C.64 custom-jdbc-listener

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
    jdbc-listeners
```

C.64.1 Function

Enables you to define a custom JDBC listener. For more information, see "kodo.JDBCListeners" in *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.64.2 Example

```
<jdbc-listeners>
  <custom-jdbc-listener>
    <classname>...</classname>
    <properties>...</properties>
  </custom-jdbc-listener>
</jdbc-listeners>
```

C.65 custom-lock-manager

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.65.1 Function

Enables you to define a custom lock manager. For more information, see "Lock Manager" in *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.65.2 Example

```
<custom-lock-manager>
  <classname>...</classname>
  <properties>...</properties>
</custom-lock-manager>
```

C.66 custom-log

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.66.1 Function

Enables you to define a custom log file. For more information, see "Custom Log" in *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.66.2 Example

```
<custom-log>
  <classname>...</classname>
  <properties>...</properties>
</custom-log>
```

C.67 custom-mapping-defaults

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.67.1 Function

Enables you to define custom mapping defaults. For more information, see "Mapping Factory" in *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.67.2 Example

```
<custom-mapping-defaults>
  <classname>...</classname>
  <properties>...</properties>
</custom-mapping-defaults>
```


C.68 custom-mapping-factory

Range of values: N/A

Default value: N/A

Parent elements:

persistence-configuration
persistence-configuration-unit

C.68.1 Function

Specifies a custom mapping factory. For more information, see "Mapping Factory" in *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.68.2 Example

```
<custom-mapping-factory/>
```

C.69 custom-meta-data-factory

Range of values: N/A

Default value: N/A

Parent elements:

persistence-configuration
persistence-configuration-unit

C.69.1 Function

Enables you to specify a custom metadata factory. For more information, see "Metadata Factory" in the *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.69.2 Example

```
<custom-meta-data-factory>
  <classname>...</classname>
  <properties>...</properties>
</custom-meta-data-factory>
```

C.70 custom-meta-data-repository

Range of values: N/A

Default value: N/A

Parent elements:

persistence-configuration
persistence-configuration-unit

C.70.1 Function

Enables you to specify a custom metadata repository. For more information, see "Metadata Factory" in the *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.70.2 Example

```
<custom-meta-data-repository>
  <classname>...</classname>
  <properties>...</properties>
</custom-meta-data-repository>
```

C.71 custom-orphaned-key-action

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.71.1 Function

Enables you to define a custom orphaned key action. For more information, see "Orphaned Keys" in the *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.71.2 Example

```
<custom-orphaned-key-action>
  <channel>openjpa.Runtime</channel>
  <level>4</level>
  <classname>...</classname>
  <properties>...</properties>
</custom-orphaned-key-action>
```

C.72 custom-persistence-server

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.72.1 Function

Enables you to define a custom persistence server. For more information, see "Standalone Persistence Server" in the *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.72.2 Example

```
<custom-persistence-server>
  <classname>...</classname>
  <properties>...</properties>
</custom-persistence-server>
```

C.73 custom-proxy-manager

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.73.1 Function

Enables you to define a custom proxy manager. For more information, see "Custom Proxies" in the *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.73.2 Example

```
<custom-proxy-manager>
  <classname>...</classname>
  <properties>...</properties>
</custom-proxy-manager>
```

C.74 custom-query-compilation-cache

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.74.1 Function

Enables you to define a custom query compilation cache. For more information, see "Query Compilation Cache" in the *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.74.2 Example

```
<custom-query-compilation-cache>
  <classname>...</classname>
  <properties>...</properties>
</custom-query-compilation-cache>
```

C.75 custom-remote-commit-provider

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.75.1 Function

Enables you to specify a custom remote commit provider. For more information, see "JMS" in the *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.75.2 Example

```
<custom-remote-commit-provider>
  <name>kodo.RemoteCommitProvider</name>
  <classname>...</classname>
  <properties>...</properties>
</custom-remote-commit-provider>
```

C.76 custom-savepoint-manager

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.76.1 Function

Enables you to specify a custom savepoint manager. For more information, see "Savepoints" in the *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.76.2 Example

```
<custom-savepoint-manager>
  <classname>...</classname>
  <properties>...</properties>
</custom-savepoint-manager>
```

C.77 custom-schema-factory

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.77.1 Function

Specifies a custom schema factory. For more information, see "Schema Factory" in the *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.77.2 Example

```
<custom-schema-factory>
  <classname>...</classname>
  <properties>...</properties>
</custom-schema-factory>
```

C.78 custom-seq

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.78.1 Function

Enables you to define a custom generator. For more information, see "Generators" in the *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.78.2 Example

```
<custom-seq>
  <classname>...</classname>
  <properties>...</properties>
</custom-seq>
```

C.79 custom-sql-factory

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.79.1 Function

Enables you to define a custom SQL factory. For more information, see "kodo.jdbc.SQLFactory" in the *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.79.2 Example

```
<custom-sql-factory>
  <classname>...</classname>
  <properties>...</properties>
</custom-sql-factory>
```

C.80 custom-update-manager

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.80.1 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 4.2.0 Developers Guide for JPA/JDO*.

C.80.2 Example

```
<custom-update-manager>
  <classname>...</classname>
  <properties>...</properties>
</custom-update-manager>
```

C.81 data-caches

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.81.1 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 4.2.0 Developers Guide for JPA/JDO*.

C.81.2 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>
```

C.82 data-cache-manager-impl

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.82.1 Function

Enables the `kodo.datacache.DataCacheManager` that manages the system data caches. For more information, see "kodo.DataCacheManager" in *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.82.2 Example

```
<data-cache-manager-impl/>
```

C.83 data-cache-timeout

Range of values: Integer

Default value: -1

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.83.1 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 4.2.0 Developers Guide for JPA/JDO*.

C.83.2 Example

```
<data-cache-timeout>-1</data-cache-timeout>
```

C.84 db2-dictionary

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.84.1 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 *Oracle WebLogic Server Administration Console Help*

C.84.2 Example

The `db2-dictionary` element shares the same subelements as "[access-dictionary](#)" on page C-5.

C.85 default-access-type

Range of values: FIELD | PROPERTY

Default value: FIELD

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
    kodo-persistence-mapping-factory
```

```
persistence-configuration
  persistence-configuration-unit
    kodo-persistence-meta-data-factory
```

C.85.1 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 4.2.0 Developers Guide for JPA/JDO*.

C.85.2 Example

```
<default-access-type>FIELD</default-access-type>
```

C.86 default-broker-factory

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.86.1 Function

Specifies a plugin string (see "Plugin Configuration" in *Kodo 4.2.0 Developers Guide for JPA/JDO*) 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 4.2.0 Developers Guide for JPA/JDO*.

C.86.2 Example

```
<default-broker-factory/>
```

C.87 default-broker-impl

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.87.1 Function

Specifies a plugin string (see "Plugin Configuration" in *Kodo 4.2.0 Developers Guide for JPA/JDO*) describing the `kodo.kernel.Broker` type to use at runtime, in this case

the default. For more information, see "kodo.BrokerImpl" in *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.87.2 Example

```
<default-broker-impl/>
```

C.88 default-class-resolver

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.88.1 Function

Specifies a plugin string (see "Plugin Configuration" in *Kodo 4.2.0 Developers Guide for JPA/JDO*) 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 4.2.0 Developers Guide for JPA/JDO*.

C.88.2 Example

```
<default-class-resolver/>
```

C.89 default-compatibility

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.89.1 Function

Sets the default compatibility settings. For more information, see "Compatibility Configuration" in *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.89.2 Example

```
<default-compatibility/>
```

C.90 default-data-cache

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
    data-caches
```

C.90.1 Function

Defines the default data cache. For more information, see "Data Cache" in *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.90.2 Example

```
<data-caches>
  <default-data-cache>
    <name>kodo.DataCache</name>
  </default-data-cache>
</data-caches>
```

C.91 default-detach-state

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.91.1 Function

Specifies the default detached state, in this case `loaded`. For more information, see "Defining the Detached Object Graph" in *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.91.2 Example

```
<default-detach-state/>
```

C.92 default-data-cache-manager

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.92.1 Function

Specifies a plugin string (see "Plugin Configuration" in *Kodo 4.2.0 Developers Guide for JPA/JDO*) describing the `kodo.datacache.DataCacheManager` that manages the system data caches, in this case the default. For more information, see "Data Cache" in *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.92.2 Example

```
<default-data-cache-manager>...</default-data-cache-manager>
```

C.93 default-driver-data-source

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.93.1 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 4.2.0 Developers Guide for JPA/JDO*.

C.93.2 Example

```
<default-driver-data-source/>
```

C.94 default-level

Range of values: TRACE | DEBUG | INFO | WARN | ERROR

Default value: INFO

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
    log-factory-impl
```

C.94.1 Function

Specifies the default logging level of unconfigured channels. For more information, see "Kodo Logging" in *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.94.2 Example

```
<default-level>INFO</default-level>
```

C.95 default-lock-manager

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.95.1 Function

Specifies to use the default lock manager. For more information, see "Configuring Default Locking" in *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.95.2 Example

```
<default-lock-manager/>
```

C.96 default-mapping-defaults

Range of values: String

Standard implementations include: `jdo` and `jpa`

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.96.1 Function

Specifies a plugin string (see "Plugin Configuration" in *Kodo 4.2.0 Developers Guide for JPA/JDO*) 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 4.2.0 Developers Guide for JPA/JDO*.

C.96.2 Example

```
<default-mapping-defaults>jpa</default-mapping-defaults>
```

C.97 default-meta-data-factory

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.97.1 Function

Specifies a plugin string (see "Plugin Configuration" in *Kodo 4.2.0 Developers Guide for JPA/JDO*) 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 4.2.0 Developers Guide for JPA/JDO*.

C.97.2 Example

```
<default-meta-data-factory/>
```

C.98 default-meta-data-repository

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.98.1 Function

Specifies a plugin string (see "Plugin Configuration" in *Kodo 4.2.0 Developers Guide for JPA/JDO*) 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 4.2.0 Developers Guide for JPA/JDO*.

C.98.2 Example

```
<default-meta-data-repository/>
```

C.99 default-orphaned-key-action

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.99.1 Function

Specifies a plugin string (see "Plugin Configuration" in *Kodo 4.2.0 Developers Guide for JPA/JDO*) 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 4.2.0 Developers Guide for JPA/JDO*.

C.99.2 Example

```
<default-orphaned-key-action>
  <channel>openjpa.Runtime</channel>
  <level>4</level>
</default-orphaned-key-action>
```

C.100 default-proxy-manager

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.100.1 Function

Specifies a plugin string (see "Plugin Configuration" in *Kodo 4.2.0 Developers Guide for JPA/JDO*) 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 4.2.0 Developers Guide for JPA/JDO*.

C.100.2 Example

```
<default-proxy-manager/>
```

C.101 default-query-compilation-cache

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.101.1 Function

Specifies a plugin string (see "Plugin Configuration" in *Kodo 4.2.0 Developers Guide for JPA/JDO*) 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 4.2.0 Developers Guide for JPA/JDO*.

C.101.2 Example

```
<default-query-compilation-cache/>
```

C.102 default-savepoint-manager

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.102.1 Function

Specifies a plugin string (see "Plugin Configuration" in *Kodo 4.2.0 Developers Guide for JPA/JDO*) 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 4.2.0 Developers Guide for JPA/JDO*.

C.102.2 Example

```
<default-savepoint-manager/>
```

C.103 default-schema-factory

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.103.1 Function

Specifies a plugin string (see "Plugin Configuration" in *Kodo 4.2.0 Developers Guide for JPA/JDO*) 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 4.2.0 Developers Guide for JPA/JDO*.

C.103.2 Example

```
<default-schema-factory/>
```

C.104 default-sql-factory

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.104.1 Function

Specifies a plugin string (see "Plugin Configuration" in *Kodo 4.2.0 Developers Guide for JPA/JDO*) 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 4.2.0 Developers Guide for JPA/JDO*.

C.104.2 Example

```
<default-meta-data-factory/>
```

C.105 default-update-manager

Range of values: N/A

Default value: N/A

Parent elements:

persistence-configuration
 persistence-configuration-unit

C.105.1 Function

Specifies to use the default update manager, `kodo.jdbc.kernel.ConstraintUpdateManager`. For more information, see "kodo.jdbc.UpdateManager" in *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.105.2 Example

```
<default-update-manager/>
```

C.106 deprecated-jdo-mapping-defaults

Range of values: N/A

Default value: N/A

Parent elements:

persistence-configuration
 persistence-configuration-unit

C.106.1 Function

Specifies the mapping defaults for JDO 1.0. For more information, see "Mapping Defaults" in *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.106.2 Example

The `deprecated-jdo-mapping-defaults` element shares the same subelements as [mapping-defaults-impl](#).

C.107 deprecated-jdo-meta-data-factory

Range of values: N/A

Default value: N/A

Parent elements:

persistence-configuration
 persistence-configuration-unit

C.107.1 Function

Specifies a plugin string (see "Plugin Configuration" in *Kodo 4.2.0 Developers Guide for JPA/JDO*) describing the `kodo.meta.MetaDataFactory` type to use, in this case

`kodo.jdo.DeprecatedJDOMetaDataFactory`. For more information, see "Metadata Factory" in the *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.107.2 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>
```

C.108 derby-dictionary

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.108.1 Function

Configures the Derby Dictionary. For a complete description of each of the values that you can specify, see "Derby Dictionary Configuration" in the *Oracle WebLogic Server Administration Console Help*.

C.108.2 Example

The `derby-dictionary` element shares the same subelements as, plus the following subelement (default shown):

```
<derby-dictionary>
...
  <shutdown-on-close>true</shutdown-on-close>
</derby-dictionary>
```

C.109 detach-options-all

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.109.1 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 4.2.0 Developers Guide for JPA/JDO*.

C.109.2 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>
```

C.110 detach-options-fetch-groups

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.110.1 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 4.2.0 Developers Guide for JPA/JDO*. For more information about the detach state, see "Defining the Detached Object Graph" in *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.110.2 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>
```

C.111 detach-options-loaded

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.111.1 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 4.2.0 Developers Guide for JPA/JDO*.

C.111.2 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>
```

C.112 detach-state

Range of values: 1 | 2 | 3

Default value: 1

Parent elements:

```
persistence-configuration
  persistence-configuration-unit]
    kodo-broker
```

C.112.1 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 4.2.0 Developers Guide for JPA/JDO*.

C.112.2 Example

```
<detach-state>1</detach-state>
```

C.113 detached-state-field

Range of values: true | false

Default value: true

Parent elements:

```
persistence-configuration
```

```
persistence-configuration-unit  
  detach-options-all
```

```
persistence-configuration  
  persistence-configuration-unit  
    detach-options-fetch-groups
```

```
persistence-configuration  
  persistence-configuration-unit  
    detach-options-loaded
```

C.113.1 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 4.2.0 Developers Guide for JPA/JDO*.

For more information about the detach state, see "Defining the Detached Object Graph" in *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.113.2 Example

```
<detached-state-field>true</detached-state-field>
```

C.114 detached-state-manager

Range of values: true | false

Default value: true

Parent elements:

```
persistence-configuration  
  persistence-configuration-unit  
    detach-options-all
```

```
persistence-configuration  
  persistence-configuration-unit  
    detach-options-fetch-groups
```

```
persistence-configuration  
  persistence-configuration-unit  
    detach-options-loaded
```

C.114.1 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 [Section C.113, "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 4.2.0 Developers Guide for JPA/JDO*.

C.114.2 Example

```
<detached-state-manager>true</detached-state-manager>
```

C.115 detached-state-transient

Range of values: true | false

Default value: true

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
    detach-options-all
```

```
persistence-configuration
  persistence-configuration-unit
    detach-options-fetch-groups
```

```
persistence-configuration
  persistence-configuration-unit
    detach-options-loaded
```

C.115.1 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 4.2.0 Developers Guide for JPA/JDO*.

C.115.2 Example

```
<detached-state-transient>false</detached-state-transient>
```

C.116 detached-new

Range of values: true | false

Default value: true

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
    kodo-broker
```

C.116.1 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 4.2.0 Developers Guide for JPA/JDO*.

C.116.2 Example

```
<detached-new>true</detached-new>
```

C.117 diagnostic-context

Range of values: String

Default value: `kodo.ID` property, if set

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
    log-factory-impl
```

C.117.1 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 4.2.0 Developers Guide for JPA/JDO*.

C.117.2 Example

```
<diagnostic-context>KodoID</diagnostic-context>
```

C.118 dynamic-data-structs

Range of values: `true` | `false`

Default value: `false`

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.118.1 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 4.2.0 Developers Guide for JPA/JDO*.

C.118.2 Example

```
<dynamic-data-structs>>false</dynamic-data-structs>
```

C.119 dynamic-schema-factory

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.119.1 Function

Specifies a plugin string (see "Plugin Configuration" in *Kodo 4.2.0 Developers Guide for JPA/JDO*) 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 4.2.0 Developers Guide for JPA/JDO*.

C.119.2 Example

```
<dynamic-schema-factory/>
```

C.120 eager-fetch-mode

Range of values: join | multiple | none | parallel | single

Default value: parallel

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.120.1 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 4.2.0 Developers Guide for JPA/JDO*.

You can set this value to one of the following:

- `join`—In this mode, Kodo joins 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 4.2.0 Developers Guide for JPA/JDO*.

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 `TelephoneNumber`, 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 4.2.0 Developers Guide for JPA/JDO*.

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`.

C.120.2 Example

```
<eager-fetch-mode>parallel</eager-fetch-mode>
```

C.121 `empress-dictionary`

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.121.1 Function

Configures the Empress Dictionary. For a complete description of each of the values that you can specify, see "Empress Dictionary Configuration" in the *Oracle WebLogic Server Administration Console Help*.

C.121.2 Example

The `empress-dictionary` element shares the same subelements as [Section C.5, "access-dictionary,"](#) plus the following sub-element (default shown):

```
<empress-dictionary>
...
  <allow-concurrent-read>false</allow-concurrent-read>
</empress-dictionary>
```

C.122 `EnableLogMBean`

Range of values: `true` | `false`

Default value: N/A

Parent elements:

```
jmx
  local-jmx
```

C.122.1 Function

Flag that specifies whether the `LogMBean` should be registered. For more information, see "Optional Parameters in Management Group" in *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.122.2 Example

```
<EnableLogMBean>true</EnableLogMBean>
```

C.123 EnableRuntimeMBean

Range of values: true | false

Default value: N/A

Parent elements:

```
jmx
  local-jmx
```

C.123.1 Function

Flag that specifies whether the `RuntimeMBean` should be registered. For more information, see "Optional Parameters in Management Group" in *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.123.2 Example

```
<EnableRuntimeMBean>true</EnableRuntimeMBean>
```

C.124 evict-from-data-cache

Range of values: true | false

Default value: false

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
    kodo-broker
```

C.124.1 Function

Flag that specifies whether to evict data from the data cache at specific times. You defined the eviction schedule using the [Section C.125, "eviction-schedule"](#) element. For more information, see "Data Cache" in *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.124.2 Example

```
<evict-from-data-cache>true</evict-from-data-cache>
```

C.125 eviction-schedule

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
    data-caches
      gem-fire-data-cache
```

```
persistence-configuration
  persistence-configuration-unit
    data-caches
      kodo-current-data-cache
```

```
persistence-configuration
  persistence-configuration-unit
    data-caches
      lru-data-cache
```

```
persistence-configuration
  persistence-configuration-unit
    data-caches
      tangosol-data-cache
```

C.125.1 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 4.2.0 Developers Guide for JPA/JDO*.

C.125.2 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>
```

C.126 exception-orphaned-key-action

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.126.1 Function

Specifies a plugin string (see "Plugin Configuration" in *Kodo 4.2.0 Developers Guide for JPA/JDO*) 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 4.2.0 Developers Guide for JPA/JDO*.

C.126.2 Example

```
<exception-orphaned-key-action/>
```

C.127 exception-reconnect-attempts

Range of values: Integer

Default value: 0

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
    jms-remote-commit-provider
```

C.127.1 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 4.2.0 Developers Guide for JPA/JDO*.

C.127.2 Example

```
<exception-reconnect-attempts>0</exception-reconnect-attempts>
```

C.128 execution-context-name-provider

Range of values: N/A

Default value: N/A

Parent elements: N/A

C.128.1 Function

Configure the execution context name provider. For more information, see "kodo.ExecutionContextNameProvider" in *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.128.2 Example

```
<execution-context-name-provider>
  <stack-execution-context-name-provider>...
</stack-execution-context-name-provider>
  <transaction-name-execution-context-name-provider>...
```

```

    </transaction-name-execution-context-name-provider>
    <user-object-execution-context-name-provider>...
  </user-object-execution-context-name-provider>
</execution-context-name-provider>

```

C.129 export-profiling

Range of values: N/A

Default value: N/A

Parent elements:

```

persistence-configuration
  persistence-configuration-unit

profiling

```

C.129.1 Function

Configure export profiling data. See "Dumping Profiling Data to Disk from a Batch Process" in *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.129.2 Example

```

<profiling>
  <export-profiling>
    <interval-millis>-1</interval-millis>
    <base-name>name</base-name>
  </export-profiling>
</profiling>

```

C.130 extension-deprecated-jdo-mapping-factory

Range of values: N/A

Default value: N/A

Parent elements:

```

persistence-configuration
  persistence-configuration-unit

```

C.130.1 Function

A plugin string (see "Plugin Configuration" in *Kodo 4.2.0 Developers Guide for JPA/JDO*) 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 4.2.0 Developers Guide for JPA/JDO*.

C.130.2 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>
```

C.131 fetch-batch-size

Range of values: Integer

Default value: -1

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.131.1 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 4.2.0 Developers Guide for JPA/JDO*.

C.131.2 Example

```
<fetch-batch-size>-1</fetch-batch-size>
```

C.132 fetch-direction

Range of values: forward | reverse | unknown

Default value: true

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.132.1 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 4.2.0 Developers Guide for JPA/JDO*.

C.132.2 Example

```
<fetch-direction>forward</fetch-direction>
```

C.133 fetch-group

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
    fetch-group
```

C.133.1 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 4.2.0 Developers Guide for JPA/JDO*.

C.133.2 Example

```
<fetch-groups>
  <fetch-group>detail</fetch-group>
</fetch-groups>
```

C.134 fetch-groups

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.134.1 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 4.2.0 Developers Guide for JPA/JDO*.

C.134.2 Example

```
<fetch-groups>
  <fetch-group>detail</fetch-group>
```

```
</fetch-groups>
```

C.135 field-override

Range of values: true | false

Default value: true

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
    kodo-persistence-mapping-factory

persistence-configuration
  persistence-configuration-unit
    kodo-persistence-meta-data-factory
```

C.135.1 Function

Flag that specifies whether to enable field override.

C.135.2 Example

```
<field-override>true</field-override>
```

C.136 file

Range of values: String

Default value: Described below

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
    file-schema-factory

persistence-configuration
  persistence-configuration-unit
    log-factory-impl
```

C.136.1 Function

If the parent element is [Section C.138, "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 4.2.0 Developers Guide for JPA/JDO*.

If the parent element is [Section C.186, "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 4.2.0 Developers Guide for JPA/JDO*.

C.136.2 Example

```
<file>stdout</file>
```


C.137 file-name

Range of values: String

Default value: `package.schema`

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
    file-schema-factory
```

C.137.1 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 4.2.0 Developers Guide for JPA/JDO*.

C.137.2 Example

```
<file-name>package.schema</file-name>
```

C.138 file-schema-factory

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.138.1 Function

Specifies a plugin string (see "Plugin Configuration" in *Kodo 4.2.0 Developers Guide for JPA/JDO*) 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 [Section C.280, "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 4.2.0 Developers Guide for JPA/JDO*.

C.138.2 Example

```
<file-schema-factory>
  <file-name>package.schema</file-name>
  <file>package.schema</file>
</file-schema-factory>
```

C.139 files

Range of values: N/A

Default value: N/A

Parent elements:

```

persistence-configuration
  persistence-configuration-unit
    deprecated-jdo-meta-data-factory

persistence-configuration
  persistence-configuration-unit
    extension-deprecated-jdo-mapping-factory

persistence-configuration
  persistence-configuration-unit
    jdo-meta-data-factory

persistence-configuration
  persistence-configuration-unit
    jdor-mapping-factory

persistence-configuration
  persistence-configuration-unit
    kodo-persistence-mapping-factory

persistence-configuration
  persistence-configuration-unit
    kodo-persistence-meta-data-factory

persistence-configuration
  persistence-configuration-unit
    mapping-file-deprecated-jdo-mapping-factory

persistence-configuration
  persistence-configuration-unit
    orm-file-jdor-mapping-factory

persistence-configuration
  persistence-configuration-unit
    table-deprecated-jdo-mapping-factory

```

C.139.1 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 4.2.0 Developers Guide for JPA/JDO*.

C.139.2 Example

```
<files>com/file1;com/file2</files>
```

C.140 filter-listeners

Range of values: N/A

Default value: N/A

Parent elements:

persistence-configuration
persistence-configuration-unit

C.140.1 Function

Defines one or more full plugin strings (see "Plugin Configuration" in *Kodo 4.2.0 Developers Guide for JPA/JDO*) 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 4.2.0 Developers Guide for JPA/JDO*.

For more information, see "kodo.FilterListeners" in *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.140.2 Example

```
<filter-listeners>
  <custom-filter-listener>detail</custom-filter-listener>
</filter-listeners>
```

C.141 foreign-keys

Range of values: true | false

Default value: false

Parent elements:

persistence-configuration
persistence-configuration-unit
lazy-schema-factory

C.141.1 Function

Flag that specifies whether to read automatically foreign key information during schema validation. For more information, see "Schema Factory" in the *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.141.2 Example

```
<foreign-keys>>false</foreign-keys>
```

C.142 format

Range of values: String

Default value: N/A

Parent elements:

persistence-configuration
persistence-configuration-unit
native-jdbc-seq

C.142.1 Function

Specifies the format. For more information, see "Generators" in the *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.143 foxpro-dictionary

Range of values: N/A

Default value: N/A

Parent elements:

persistence-configuration
persistence-configuration-unit

C.143.1 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 *Oracle WebLogic Server Administration Console Help*.

C.143.2 Example

The foxpro-dictionary element shares the same subelements as [Section C.5, "access-dictionary."](#)

C.144 flush-before-queries

Range of values: true | false | with-connection

Default value: true

Parent elements:

persistence-configuration
persistence-configuration-unit

C.144.1 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 4.2.0 Developers Guide for JPA/JDO*.

C.144.2 Example

```
<flush-before-queries>true</flush-before-queries>
```

C.145 hsql-dictionary

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.145.1 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 *Oracle WebLogic Server Administration Console Help*.

C.145.2 Example

The `hsql-dictionary` element shares the same subelements as [Section C.5](#), "`access-dictionary`."

C.146 gem-fire-data-cache

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
    data-caches
```

C.146.1 Function

Integrates with GemStone's GemFire v5.0.1 caching system. For more information, see "GemStone GemFire Integration" in *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.146.2 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>
```

C.147 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
```

C.147.1 Function

GemFire region name. For more information, see "GemStone GemFire Integration" in *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.147.2 Example

```
<gem-fire-data-cache-name>root/kodo-data-cache</gem-fire-data-cache-name>
```

C.148 gui-jmx

Range of values: N/A

Default value: N/A

Parent elements:

```
jmx
```

C.148.1 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 4.2.0 Developers Guide for JPA/JDO*.

C.148.2 Example

```
<gui-jmx>
  <MBeanServerStrategy>any-create</MBeanServerStrategy>
  <EnableLogMBean>true</EnableLogMBean>
  <EnableRuntimeMBean>true</EnableRuntimeMBean>
</gui-jmx>
```

C.149 gui-profiling

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit

profiling
```

C.149.1 Function

Turns on the local profiling GUI. See "Profiling in an Embedded GUI" in the Kodo Developer's Guide for more information.

C.149.2 Example

```
<profiling>
  <gui-profiling/>
</profiling>
```

C.150 Host

Range of values: N/A

Default value: N/A

Parent elements:

```
jmx
  mx4j1-jmx
```

C.150.1 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 4.2.0 Developers Guide for JPA/JDO*.

C.150.2 Example

```
<Host>localhost</Host>
```

C.151 host

Range of values: Valid hostname

Default value: `localhost`

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
    tcp-transport
```

C.151.1 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 4.2.0 Developers Guide for JPA/JDO*.

C.151.2 Example

```
<host>localhost</host>
```

C.152 http-transport

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.152.1 Function

Specifies the URL of the remote server. For more information, see "Client Managers" in *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.152.2 Example

```
<http-transport>
  <url>servlet-url</url>
</http-transport>
```

C.153 ignore-changes

Range of values: true | false

Default value: false

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
    kodo-broker
```

C.153.1 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 in-memory or flush before running it against the datastore. For more information, see "kodo.IgnoreChanges" in *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.153.2 Example

```
<ignore-changes>false</ignore-changes>
```

C.154 ignore-unmapped

Range of values: true | false

Default value: false

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
    class-table-jdbc-seq
```


C.154.1 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 4.2.0 Developers Guide for JPA/JDO*.

C.154.2 Example

```
<ignore-unmapped>>false</ignore-unmapped>
```

C.155 ignore-virtual

Range of values: true | false

Default value: true

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
    class-table-jdbc-seq
```

C.155.1 Function

Flag that specifies whether to ignore virtual classes. For more information, see "Generators" in the *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.155.2 Example

```
<ignore-virtual>>false</ignore-virtual>
```

C.156 in-memory-savepoint-manager

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.156.1 Function

Specifies a plugin string (see "Plugin Configuration" in *Kodo 4.2.0 Developers Guide for JPA/JDO*) 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 4.2.0 Developers Guide for JPA/JDO*.

C.156.2 Example

```
<in-memory-savepoint-manager/>
```

C.157 increment

Range of values: Integer

Default value: 1

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
    class-table-jdbc-seq
```

```
persistence-configuration
  persistence-configuration-unit
    native-jdbc-seq
```

```
persistence-configuration
  persistence-configuration-unit
    table-jdbc-seq
```

```
persistence-configuration
  persistence-configuration-unit
    time-seeded-seq
```

```
persistence-configuration
  persistence-configuration-unit
    value-table-jdbc-seq
```

C.157.1 Function

Specifies the amount of sequence increments. For more information, see "Generators" in the *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.157.2 Example

```
<increment>1</increment>
```

C.158 indexes

Range of values: true | false

Default value: false

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
    lazy-schema-factory
```

C.158.1 Function

Flag that specifies whether to read automatically index information during schema validation. For more information, see "Schema Factory" in the *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.158.2 Example

```
<indexes>>false</indexes>
```

C.159 informix-dictionary

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.159.1 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 *Oracle WebLogic Server Administration Console Help*.

C.159.2 Example

The `informix-dictionary` element shares the same subelements as [Section C.5, "access-dictionary,"](#) 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>
```

C.160 initial-value

Range of values: N/A

Default value: 1

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
    native-jdbc-seq
```

C.160.1 Function

Specifies the initial sequence value. For more information, see "Generators" in the *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.160.2 Example

```
<initial-value>1</intial-value>
```

C.161 interval-millis

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
```

```
persistence-configuration-unit
  export-profiling
```

```
profiling
  export-profiling
```

C.161.1 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 4.2.0 Developers Guide for JPA/JDO*.

C.161.2 Example

```
<interval-millis>-1</interval-millis>
```

C.162 inverse-manager

Range of values: true | false

Default value: false

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.162.1 Function

A plugin string (see "Plugin Configuration" in *Kodo 4.2.0 Developers Guide for JPA/JDO*) describing a `kodo.kernel.InverseManager` to use for managing bidirectional relations upon a flush. For more information, see "Managed Inverses" in *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.162.2 Example

```
<inverse-manager>
  <action>warn</action>
  <manage-lrs>false</manage-lrs>
</inverse-manager>
```

C.163 jdatastore-dictionary

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.163.1 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 *Oracle WebLogic Server Administration Console Help*.

C.163.2 Example

The `jdbc3-savepoint-manager` element shares the same subelements as [Section C.5, "access-dictionary."](#)

C.164 jdbc-broker-factory

Range of values: N/A

Default value: N/A

Parent elements:

`persistence-configuration`
`persistence-configuration-unit`

C.164.1 Function

A plugin string (see "Plugin Configuration" in *Kodo 4.2.0 Developers Guide for JPA/JDO*) describing the `kodo.kernel.BrokerFactory` type to use, in this case JDBC. For more information, see "kodo.BrokerFactory" in the *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.164.2 Example

```
<jdbc-broker-factory/>
```

C.165 jdbc-listeners

Range of values: N/A

Default value: N/A

Parent elements:

`persistence-configuration`
`persistence-configuration-unit`

C.165.1 Function

Defines one or more full plugin strings (see "Plugin Configuration" in *Kodo 4.2.0 Developers Guide for JPA/JDO*) 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 4.2.0 Developers Guide for JPA/JDO*.

C.165.2 Example

```
<jdbc-listeners>detail</jdbc-listeners>
```

C.166 jdbc3-savepoint-manager

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.166.1 Function

Specifies a plugin string (see "Plugin Configuration" in *Kodo 4.2.0 Developers Guide for JPA/JDO*) 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 4.2.0 Developers Guide for JPA/JDO*.

C.166.2 Example

```
<jdbc3-savepoint-manager/>
```

C.167 jdo-meta-data-factory

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.167.1 Function

Specifies a plugin string (see "Plugin Configuration" in *Kodo 4.2.0 Developers Guide for JPA/JDO*) describing the `kodo.meta.MetaDataFactory` type to use, in this case `kodo.jdo.JDOMetaDataFactory`. For more information, see "Metadata Factory" in the *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.167.2 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>
```

C.168 jms-remote-commit-provider

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
```

persistence-configuration-unit

C.168.1 Function

Configures the JMS remote commit provider. For more information, see "JMS" in the *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.168.2 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>
```

C.169 jmx

Range of values: N/A

Default value: N/A

Parent elements:

C.169.1 Function

Enables for the configuration of management capabilities. For more information, see "Management and Monitoring" in the *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.169.2 Example

```
<jmx>
  <none-jmx/>
  <local-jmx>...</local-jmx>
  <gui-jmx>...</gui-jmx>
  <jmx2-jmx>...</jmx2-jmx>
  <mx4j1-jmx>...</mx4j1-jmx>
  <wls81-jmx>...</wls81jmx>
</jmx>
```

C.170 jmx2-jmx

Range of values: N/A

Default value: N/A

Parent elements:

jmx

C.170.1 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 4.2.0 Developers Guide for JPA/JDO*, respectively.

C.170.2 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>
```

C.171 JNDIName

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
    JNDIName
```

C.171.1 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 4.2.0 Developers Guide for JPA/JDO*.

C.171.2 Example

```
<JNDIName>jrmp</JNDIName>
```

C.172 kodo-broker

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.172.1 Function

Configures the `kodo.BrokerImpl`. For more information, see "kodo.BrokerImpl" in *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.172.2 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>false</nontransactional-write>
    <sync-with-managed-transactions>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</order-dirty-objects>
  </kodo-broker>

```

C.173 kodo-concurrent-data-cache

Range of values: N/A

Default value: N/A

Parent elements:

```

persistence-configuration
  persistence-configuration-unit
    data-caches

```

C.173.1 Function

Enables the basic concurrent data cache. For more information, see "Data Cache" in *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.173.2 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>

```

C.174 kodo-data-cache-manager

Range of values: N/A

Default value: N/A

Parent elements:

```

persistence-configuration
  persistence-configuration-unit

```

C.174.1 Function

Enables the `kodo.datacache.DataCacheManager` that manages the system data caches. For more information, see "kodo.DataCacheManager" in *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.174.2 Example

```
<kodo-data-cache-manager/>
```

C.175 kodo-mapping-repository

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.175.1 Function

Specifies a plugin string (see "Plugin Configuration" in *Kodo 4.2.0 Developers Guide for JPA/JDO*) describing the `kodo.meta.MetadataRepository` to use. For more information, see "kodo.MetadataRepository" in the *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.175.2 Example

```
<kodo-mapping-repository>
  <resolve>3</resolve>
  <validate>7</validate>
  <source-mode>7</source-mode>
</kodo-mapping-repository>
```

C.176 kodo-persistence-mapping-factory

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.176.1 Function

A plugin string (see "Plugin Configuration" in *Kodo 4.2.0 Developers Guide for JPA/JDO*) 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 4.2.0 Developers Guide for JPA/JDO*.

C.176.2 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>

```

C.177 kodo-persistence-meta-data-factory

Range of values: N/A

Default value: N/A

Parent elements:

```

persistence-configuration
  persistence-configuration-unit

```

C.177.1 Function

Specifies a plugin string (see "Plugin Configuration" in *Kodo 4.2.0 Developers Guide for JPA/JDO*) 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 4.2.0 Developers Guide for JPA/JDO*.

C.177.2 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>

```

C.178 kodo-pooling-data-source

Range of values: N/A

Default value: N/A

Parent elements:

```

persistence-configuration
  persistence-configuration-unit

```

C.178.1 Function

Configures the `kodo.jdbc.schema.KodoPoolingDataSource` which by default performs connection pooling. For more information, see "JDBC" in *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.178.2 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>
```

C.179 kodo-sql-factory

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.179.1 Function

Specifies a plugin string (see "Plugin Configuration" in *Kodo 4.2.0 Developers Guide for JPA/JDO*) describing the `kodo.jdbc.SQLFactory` to use to abstract common SQL constructs. For more information, see "kodo.jdbc.SQLFactory" in the *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.179.2 Example

```
<kodo-sql-factory>
  <advanced-sql>...</advanced-sql>
</kodo-sql-factory>
```

C.180 large-transaction

Range of values: true | false

Default value: true

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
    kodo-broker
```

C.180.1 Function

Flag that specifies whether the transaction will generate a large number of objects.

C.180.2 Example

```
<large-transaction>>false</large-transaction>
```

C.181 lazy-schema-factory

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.181.1 Function

Specifies a plugin string (see "Plugin Configuration" in *Kodo 4.2.0 Developers Guide for JPA/JDO*) 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 4.2.0 Developers Guide for JPA/JDO*.

C.181.2 Example

```
<lazy-schema-factory>
  <foreign-keys>>false</foreign-keys>
  <indexes>>false</indexes>
  <primary-keys>>false</primary-keys>
</lazy-schema-factory>
```

C.182 level

Range of values: 1 | 2 | 3 | 4 | 5

Default value: 4

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.182.1 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 4.2.0 Developers Guide for JPA/JDO*.

C.182.2 Example

```
<level>4</level>
```

C.183 local-jmx

Range of values: N/A

Default value: N/A

Parent elements:

jmx

C.183.1 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 4.2.0 Developers Guide for JPA/JDO*.

C.183.2 Example

```
<local-jmx>
  <MBeanServerStrategy>any-create</MBeanServerStrategy>
  <EnableLogMBean>true</EnableLogMBean>
  <EnableRuntimeMBean>true</EnableRuntimeMBean>
</local-jmx>
```

C.184 local-profiling**Range of values:** N/A**Default value:** N/A**Parent elements:**

```
persistence-configuration
  persistence-configuration-unit

profiling
```

C.184.1 Function

Enable profiling without export or GUI. Useful when trying to access the `ProfilingAgent` programmatically. For more information, see "Profiling" in *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.184.2 Example

```
<profiling>
  <local-profiling/>
</profiling>
```

C.185 lock-timeout**Range of values:** Integer**Default value:** -1**Parent elements:**

```
persistence-configuration
  persistence-configuration-unit
```

C.185.1 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 4.2.0 Developers Guide for JPA/JDO*.

C.185.2 Example

```
<lock-timeout>-1</lock-timeout>
```

C.186 log-factory-impl

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.186.1 Function

Configures the default logging system. For more information, see "Kodo Logging" in *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.186.2 Example

```
<log-factory-impl>
  <diagnostic-context>ID</diagnostic-context>
  <default-level>INFO</default-level>
  <file>stdout</file>
</log-factory-impl>
```

C.187 log-orphaned-key-action

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.187.1 Function

Specifies a plugin string (see "Plugin Configuration" in *Kodo 4.2.0 Developers Guide for JPA/JDO*) 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 4.2.0 Developers Guide for JPA/JDO*.

C.187.2 Example

```
<log-orphaned-key-action>
  <channel>openjpa.Runtime</channel>
```

```
<level>4</level>
</log-orphaned-key-action>
```

C.188 log4j-log-factory

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.188.1 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* at <http://jakarta.apache.org/log4j/docs/manual.html>. For more information, see "Log4J" in *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.188.2 Example

```
<log4j-log-factory/>
```

C.189 login-timeout

Range of values: Valid classname

Default value:

Section C.178, "kodo-pooling-data-source": 10

Section C.258, "simple-driver-data-source": 0

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
    kodo-pooling-data-source
```

```
persistence-configuration
  persistence-configuration-unit
    simple-driver-data-source
```

C.189.1 Function

Specifies the login timeout. For more information, see "Using the Kodo Data Source" in *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.189.2 Example

```
<login-timeout>0</login-timeout>
```


C.190 lrs-size

Range of values: query | last | unknown

Default value: query

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.190.1 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 4.2.0 Developers Guide for JPA/JDO*.

C.190.2 Example

```
<lrs-size>query</lrs-size>
```

C.191 lru-data-cache

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
    data-caches
```

C.191.1 Function

Sets the least-recently-used (LRU) data caching option. For more information, see "Configuring the LRU Caching Option" in *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.191.2 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>
```

C.192 manage-lru

Range of values: true | false

Default value: false

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
    inverse-manager
```

C.192.1 Function

Flag that specifies whether large result set fields are included from management. For more information, see "Managed Inverses" in *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.192.2 Example

```
<manage-lrs>>false</manage-lrs>
```

C.193 mapping

Range of values: String

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

```
persistence-configuration
  persistence-configuration-unit
    orm-file-mapping-factory
```

C.193.1 Function

When a child element of [Section C.225, "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 4.2.0 Developers Guide for JPA/JDO*.

When a child element of [Section C.221, "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 4.2.0 Developers Guide for JPA/JDO*.

C.193.2 Example

```
<mapping>org/mag/Magazine-hsql.orm</mapping>
```

C.194 mapping-column

Range of values: String

Default value: MAPPING_DEF

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
    table-deprecated-jdo-mapping-factory
```

```
persistence-configuration
  persistence-configuration-unit
    table-jdor-mapping-factory
```

C.194.1 Function

Specifies the name of the column that holds the XML mapping. For more information, see "Mapping Factory" in *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.194.2 Example

```
<mapping-column>MAPPING_DEF</mapping-column>
```

C.195 mapping-defaults-impl

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.195.1 Function

Configures the mapping defaults if you set [Section C.96, "default-mapping-defaults"](#) to jdo. For a complete description of each of the values that you can specify, see "Mapping Defaults" in *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.195.2 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>
```

C.196 mapping-file-deprecated-jdo-mapping-factory

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.196.1 Function

A plugin string (see "Plugin Configuration" in *Kodo 4.2.0 Developers Guide for JPA/JDO*) 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 4.2.0 Developers Guide for JPA/JDO*.

C.196.2 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>
```

C.197 max-active

Range of values: Integer

Default value: 2

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

tcp-remote-commit-provider

C.197.1 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 4.2.0 Developers Guide for JPA/JDO*.

C.197.2 Example

```
<max-active>2</max-active>
```

C.198 max-idle

Range of values: Integer

Default value: 2

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
    tcp-remote-commit-provider
```

C.198.1 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 4.2.0 Developers Guide for JPA/JDO*.

C.198.2 Example

```
<max-idle>2</max-idle>
```

C.199 max-size

Range of values: Integer

Default value: 2147483647

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.199.1 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 4.2.0 Developers Guide for JPA/JDO*.

C.199.2 Example

```
<max-size>2147483647</max-size>
```

C.200 maximize-batch-size

Range of values: true | false

Default value: true

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
    constraint-update-manager
```

```
persistence-configuration
  persistence-configuration-unit
    batching-operation-order-update-manager
```

```
persistence-configuration
  persistence-configuration-unit
    table-lock-update-manager
```

C.200.1 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 4.2.0 Developers Guide for JPA/JDO*.

C.200.2 Example

```
<maximize-batch-size>true</maximize-batch-size>
```

C.201 MBeanServerStrategy

Range of values: any-create | create | agentID:<agentID>

Default value: any-create

Parent elements:

```
jmx
  local-jmx
```

C.201.1 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 4.2.0 Developers Guide for JPA/JDO*.

C.201.2 Example

```
<MBeanServerStrategy>any-create</MBeanServerStrategy>
```

C.202 multithreaded

Range of values: true | false

Default value: false

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

```
persistence-configuration
  persistence-configuration-unit
    kodo-broker
```

C.202.1 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 4.2.0 Developers Guide for JPA/JDO*.

C.202.2 Example

```
<multithreaded>>false</multithreaded>
```

C.203 mx4j1-jmx

Range of values: N/A

Default value: N/A

Parent elements:

```
jmx
```

C.203.1 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 4.2.0 Developers Guide for JPA/JDO*, respectively.

C.203.2 Example

```
<mx4j1-jmx>
  <MBeanServerStrategy>any-create</MBeanServerStrategy>
  <EnableLogMBean>true</EnableLogMBean>
  <EnableRuntimeMBean>true</EnableRuntimeMBean>
  <Host>localhost</Host>
  <Port>1099</Port>
  <JNDIName>jrmp</JNDIName>
</mx4j1-jmx>
```

C.204 mysql-dictionary

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.204.1 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 *Oracle WebLogic Server Administration Console Help*.

C.204.2 Example

The `mysql-dictionary` element shares the same subelements as [Section C.5, "access-dictionary,"](#) 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>
```

C.205 name

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
    data-cache
```

```
persistence-configuration
  persistence-configuration-unit
    db-dictionary
```

```
persistence-configuration
  persistence-configuration-unit
    jms-remote-commit-provider
```

```
persistence-configuration
  persistence-configuration-unit
    property-type
```

```
persistence-configuration
  persistence-configuration-unit
    query-cache
```

```
persistence-configuration
  persistence-configuration-unit
    single-jvm-remote-commit-provider
```

```
persistence-configuration
  persistence-configuration-unit
    tcp-remote-commit-provider
```



```

persistence-configuration
  persistence-configuration-unit
    cluster-remote-commit-provider

```

C.205.1 Function

Specifies the name of the corresponding element.

C.206 name-column

Range of values: String

Default value: Name

Parent elements:

```

persistence-configuration
  persistence-configuration-unit
    table-deprecated-jdo-mapping-factory

```

```

persistence-configuration
  persistence-configuration-unit
    table-jdor-mapping-factory

```

C.206.1 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 4.2.0 Developers Guide for JPA/JDO*.

C.206.2 Example

```
<name-column>NAME</name-column>
```

C.207 NamingImpl

Range of values: Valid classname

Default value: N/A

Parent elements:

```

jmx
  jmx2-jmx

```

C.207.1 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 4.2.0 Developers Guide for JPA/JDO*.

C.207.2 Example

```
<NamingImpl>mx4j.tools.naming.NamingService</NamingImpl>
```

C.208 native-jdbc-seq

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.208.1 Function

Specifies a plugin string (see "Plugin Configuration" in *Kodo 4.2.0 Developers Guide for JPA/JDO*) 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 4.2.0 Developers Guide for JPA/JDO*.

C.208.2 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>
```

C.209 none-jmx

Range of values: N/A

Default value: N/A

Parent elements:

```
jmx
```

C.209.1 Function

No management is turned on. This is the default. For more information, see "Configuration" in *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.209.2 Example

```
<none-jmx/>
```

C.210 none-lock-manager

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.210.1 Function

Specifies the `kodo.kernel.NoneLockManager` which does not perform any locking. For more information, see "Lock Manager" in *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.210.2 Example

```
<none-lock-manager/>
```

C.211 none-log-factory

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.211.1 Function

Disable logging. For more information, see "Disable Logging" in *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.211.2 Example

```
<none-log-factory/>
```

C.212 none-orphaned-key-action

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.212.1 Function

Specifies a plugin string (see "Plugin Configuration" in *Kodo 4.2.0 Developers Guide for JPA/JDO*) 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 4.2.0 Developers Guide for JPA/JDO*.

C.212.2 Example

```
<none-orphaned-key-action/>
```

C.213 none-profiling

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit

profiling
```

C.213.1 Function

Turns off profiling of your code. For more information, see "Profiling" in *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.213.2 Example

```
<profiling>
  <none-profiling/>
</profiling>
```

C.214 nontransactional-read

Range of values: true | false

Default value:

```
persistence-configuration: false
kodo-broker: true
```

Parent elements:

```
persistence-configuration
  persistence-configuration-unit

persistence-configuration
  persistence-configuration-unit
  kodo-broker
```

C.214.1 Function

Flag that specifies whether the Kodo runtime allows you to read data outside of a transaction. For more information, see "kodo.NontransactionalRead" in *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.214.2 Example

```
<nontransactional-read>true</nontransactional-read>
```

C.215 nontransactional-write

Range of values: true | false

Default value: false

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

```
persistence-configuration
  persistence-configuration-unit
    kodo-broker
```

C.215.1 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.NontransactionalWrite" in *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.215.2 Example

```
<nontransactional-write>>false</nontransactional-write>
```

C.216 num-broadcast-threads

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
    tcp-remote-commit-provider
```

C.216.1 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 4.2.0 Developers Guide for JPA/JDO*. 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 4.2.0 Developers Guide for JPA/JDO*.

C.216.2 Example

```
<num-broadcast-threads>2</num-broadcast-threads>
```

C.217 operation-order-update-manager

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.217.1 Function

Defines an update manager that writes SQL in object-level operation order. For more information, see "kodo.jdbc.UpdateManager" in *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.217.2 Example

```
<operation-order-update-manager/>
```

C.218 optimistic

Range of values: true | false

Default value: true

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

```
persistence-configuration
  persistence-configuration-unit
    kodo-broker
```

C.218.1 Function

Flag that specifies whether the datastore transactional mode is optimistic. For more information, see "kodo.Optimistic" in *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.218.2 Example

```
<optimistic>true</optimistic>
```

C.219 oracle-dictionary

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.219.1 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 *Oracle WebLogic Server Administration Console Help*.

C.219.2 Example

The `oracle-dictionary` element shares the same subelements as [Section C.5, "access-dictionary,"](#) 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>
```

C.220 oracle-savepoint-manager

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.220.1 Function

Specifies a plugin string (see "Plugin Configuration" in *Kodo 4.2.0 Developers Guide for JPA/JDO*) 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 [Section C.166, "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 4.2.0 Developers Guide for JPA/JDO*.

C.220.2 Example

```
<oracle-savepoint-manager/>
```

C.221 orm-file-jdor-mapping-factory

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.221.1 Function

A plugin string (see "Plugin Configuration" in *Kodo 4.2.0 Developers Guide for JPA/JDO*) describing the `kodo.jdo.jdbc.ORMFileJDORMappingFactory` that stores mapping metadata in `.orm` files. This setting is valid for JDO 1.0. For more information, see "Mapping Factory" in *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.221.2 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>
```

C.222 order-dirty-objects

Range of values: true | false

Default value: false

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
    kodo-broker
```

C.222.1 Function

Flag that specifies whether Kodo maintains the order that changes were made to dirty objects when writing the changes back to the database.

C.222.2 Example

```
<order-dirty-objects>>false</order-dirty-objects>
```

C.223 Password

Range of values: N/A

Default value: N/A

Parent elements:

```
jmx
  wls81-jmx
```


C.223.1 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 4.2.0 Developers Guide for JPA/JDO*.

C.223.2 Example

```
<Password>admin</Password>
```

C.224 persistence-configuration

Range of values: N/A

Default value: N/A

Parent elements: N/A

C.224.1 Function

Root element of the persistence deployment descriptor.

C.225 persistence-configuration-unit

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
```

C.225.1 Function

Contains the deployment information for a persistence unit that is available in WebLogic Server.

C.226 pessimistic-lock-manager

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.226.1 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 4.2.0 Developers Guide for JPA/JDO*.

C.226.2 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>
```

C.227 persistence-mapping-defaults

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration  
  persistence-configuration-unit
```

C.227.1 Function

Configures the mapping defaults if you set [Section C.96, "default-mapping-defaults"](#) to `jpa`. For a complete description of each of the values that you can specify, see "Mapping Defaults" in *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.227.2 Example

The `persistence-mapping-defaults` element shares the same subelements as [Section C.195, "mapping-defaults-impl."](#)

C.228 populate-data-cache

Range of values: `true` | `false`

Default value: `true`

Parent elements:

```
persistence-configuration  
  persistence-configuration-unit  
    kodo-broker
```

C.228.1 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 data cache so that the transaction does not fill the entire data cache with objects that will not be accessed again.

C.228.2 Example

```
<populate-data-cache>true</populate-data-cache>
```

C.229 Port

Range of values: N/A

Default value: N/A

Parent elements:

```
jmx
  mx4j1-jmx
```

C.229.1 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 4.2.0 Developers Guide for JPA/JDO*.

C.229.2 Example

```
<Port>1099</Port>
```

C.230 port

Range of values: Valid port

Default value: See below (depends on parent element)

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
    tcp-transport
```

```
persistence-configuration
  persistence-configuration-unit
    tcp-remote-commit-provider
```

C.230.1 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 [Section C.285, "tcp-transport,"](#) the element defaults to 5637. For more information, see "Standalone Persistence Server" in *Kodo 4.2.0 Developers Guide for JPA/JDO*.

If the parent element is [Section C.284, "tcp-remote-commit-provider,"](#) the element defaults to 5636. For more information, see "TCP" in *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.230.2 Example

```
<port>5637</port>
```

C.231 postgres-dictionary

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.231.1 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 *Oracle WebLogic Server Administration Console Help*.

C.231.2 Example

The `postgres-dictionary` element shares the same subelements as [Section C.5, "access-dictionary,"](#) 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>
```

C.232 primary-key-column

Range of values: String

Default value: ID

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
    class-table-jdbc-seq
```

```
persistence-configuration
  persistence-configuration-unit
    table-jdbc-seq
```

```
persistence-configuration
  persistence-configuration-unit
    table-schema-factory
```

```
persistence-configuration
  persistence-configuration-unit
    value-table-jdbc-seq
```

C.232.1 Function

Specifies the name of the table's numeric primary key column. For more information, see "Generators" or "Schema Factory" in the *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.232.2 Example

```
<primary-key-column>ID</primary-key-column>
```

C.233 primary-key-value

Range of values: String

Default value: DEFAULT

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
    value-table-jdbc-seq
```

C.233.1 Function

Specifies the primary key that is used by this instance. For more information, see "Generators" in the *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.233.2 Example

```
<primary-key-value>DEFAULT</primary-key-value>
```

C.234 primary-keys

Range of values: true | false

Default value: false

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
    lazy-schema-factory
```

C.234.1 Function

Flag that specifies whether to read automatically primary key information during schema validation. For more information, see "Schema Factory" in the *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.234.2 Example

```
<primary-keys>>false</primary-keys>
```

C.235 profiling

Range of values: N/A

Default value: N/A

Parent elements: N/A

C.235.1 Function

Enables you to configure profiling of your application code. For more information, see "Profiling" in *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.235.2 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>
```

C.236 profiling-proxy-manager

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.236.1 Function

Configures the profiling proxy manager. For more information, see "Custom Proxies" in the *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.236.2 Example

```
<profiling-proxy-manager>
  <assert-allowed-type>>false</assert-allowed-type>
  <track-changes>true</track-changes>
</profiling-proxy-manager>
```

C.237 properties

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
    aggregate-listener
      custom-aggregate-listener
```

persistence-configuration
 persistence-configuration-unit
 custom-broker-factory

persistence-configuration
 persistence-configuration-unit
 custom-broker-impl

persistence-configuration
 persistence-configuration-unit
 custom-class-resolver

persistence-configuration
 persistence-configuration-unit
 custom-compatibility

persistence-configuration
 persistence-configuration-unit
 custom-connection-decorator

persistence-configuration
 persistence-configuration-unit
 custom-data-cache

persistence-configuration
 persistence-configuration-unit
 custom-dictionary

persistence-configuration
 persistence-configuration-unit
 custom-driver-data-source

persistence-configuration
 persistence-configuration-unit
 custom-filter-listener

persistence-configuration
 persistence-configuration-unit
 custom-jdbc-listener

persistence-configuration
 persistence-configuration-unit
 custom-lock-manager

persistence-configuration
 persistence-configuration-unit
 custom-logType

persistence-configuration
 persistence-configuration-unit
 custom-mapping-defaults

persistence-configuration
 persistence-configuration-unit
 custom-meta-data-respository

persistence-configuration
 persistence-configuration-unit
 custom-orphaned-key-action

```
persistence-configuration
  persistence-configuration-unit
    custom-persistence-server

persistence-configuration
  persistence-configuration-unit
    custom-proxy-manager

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

persistence-configuration
  persistence-configuration-unit
    custom-query-compilation

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

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

persistence-configuration
  persistence-configuration-unit
    custom-seq

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

persistence-configuration
  persistence-configuration-unit
    custom-update-manager
```

C.237.1 Function

Specifies one or more [Section C.238](#), "property" elements.

C.238 property

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
    properties
```

C.238.1 Function

Enables you to specify the name and value of a property.

C.238.2 Example

```
<property>
  <name>timeout</name>
```



```
<value>1000</value>
</property>
```

C.239 proxy-manger-impl

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.239.1 Function

Configures the default proxy manager. For more information, see "Custom Proxies" in the *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.239.2 Example

```
<proxy-manager-impl>
  <assert-allowed-type>false</assert-allowed-type>
  <track-changes>true</track-changes>
</proxy-manager-impl>
```

C.240 query-caches

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.240.1 Function

A plugin string (see "Plugin Configuration" in *Kodo 4.2.0 Developers Guide for JPA/JDO*) 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 4.2.0 Developers Guide for JPA/JDO*.

C.240.2 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>
```

C.241 quoted-numbers-in-queries

Range of values: true | false

Default value: false

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
    compatibility
```

C.241.1 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 4.2.0 Developers Guide for JPA/JDO*.

C.241.2 Example

```
<quoted-numbers-in-queries>>false</quoted-numbers-in-queries>
```

C.242 read-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
```

C.242.1 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 4.2.0 Developers Guide for JPA/JDO*.

C.242.2 Example

```
<read-lock-level>read</read-lock-level>
```

C.243 recover-action

Range of values: none | clear

Default value: none

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
    cluster-remote-commit-provider
```

C.243.1 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 4.2.0 Developers Guide for JPA/JDO*.

C.243.2 Example

```
<recover-action>none</recover-action>
```

C.244 recovery-time-millis

Range of values: Integer

Default value: 15000

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
    tcp-remote-commit-provider
```

C.244.1 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 4.2.0 Developers Guide for JPA/JDO*.

C.244.2 Example

```
<recovery-time-millis>15000</recovery-time-millis>
```

C.245 resources

Range of values: Valid resource path

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
    tangosol-data-cache
```

C.245.1 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 4.2.0 Developers Guide for JPA/JDO*.

C.245.2 Example

```
<resources>com/aaa/package.jdo;com/bbb/package.jdo</resources>
```

C.246 restore-state

Range of values: true

persistence-configuration-unit: all | immutable | none | true | false

kodo-broker: 1 | 2 | 3 | 4 | 5

Default value: none

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

```
persistence-configuration
  persistence-configuration-unit
    kodo-broker
```

C.246.1 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 4.2.0 Developers Guide for JPA/JDO*.

C.246.2 Example

```
<restore-state>none</restore-state>
```

C.247 result-set-type

Range of values: forward--only | scroll-sensitive | scroll-insensitive

Default value: forward-only

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.247.1 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 4.2.0 Developers Guide for JPA/JDO*.

C.247.2 Example

```
<result-set-type>forward-only</result-set-type>
```

C.248 retain-state

Range of values: true | false

Default value: See below (depends on parent)

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

```
persistence-configuration
  persistence-configuration-unit
    kodo-broker
```

C.248.1 Function

Flag that specifies whether persistent fields retain their values on transaction commit. This element defaults to `true` when a child element of the [Section C.225](#), "[persistence-configuration-unit](#)" and `false` when a child element of the [Section C.172](#), "[kodo-broker](#)" element.

For more information, see "kodo.RetainState" in *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.248.2 Example

```
<retain-state>true</retain-state>
```

C.249 retry-class-registration

Range of values: `true` | `false`

Default value: `false`

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.249.1 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 4.2.0 Developers Guide for JPA/JDO*.

C.249.2 Example

```
<retry-class-registration>false</retry-class-registration>
```

C.250 scan-top-down

Range of values: `true` | `false`

Default value: `false`

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
    deprecated-jdo-meta-data-factory

persistence-configuration
  persistence-configuration-unit
    extension-deprecated-jdo-mapping-factory

persistence-configuration
  persistence-configuration-unit
    jdo-meta-data-factory
```

```

persistence-configuration
  persistence-configuration-unit
    jdor-mapping-factory

persistence-configuration
  persistence-configuration-unit
    mapping-file-deprecated-jdo-mapping-factory

persistence-configuration
  persistence-configuration-unit
    orm-file-jdor-mapping-factory

persistence-configuration
  persistence-configuration-unit
    table-deprecated-jdo-mapping-factory

```

C.250.1 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 4.2.0 Developers Guide for JPA/JDO*.

C.250.2 Example

```
<scan-top-down>false</scan-top-down>
```

C.251 schema

Range of values: String

Default value: N/A

Parent elements:

```

persistence-configuration
  persistence-configuration-unit

```

C.251.1 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 4.2.0 Developers Guide for JPA/JDO*.

C.251.2 Example

```
<schema>SCHEMA</schema>
```

C.252 schema-column

Range of values: String

Default value: SCHEMA_DEF

Parent elements:

```
persistence-configuration
```

```
persistence-configuration-unit  
  table-schema-factory
```

C.252.1 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 4.2.0 Developers Guide for JPA/JDO*.

C.252.2 Example

```
<schema-column>SCHEMA_DEF</schema-column>
```

C.253 schemas

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration  
  persistence-configuration-unit
```

C.253.1 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 4.2.0 Developers Guide for JPA/JDO*.

C.253.2 Example

```
<schemas>BUSOBSJS, GENERAL.ADDRESS, .SYSTEM_INFO </schemas>
```

C.254 sequence

Range of values: N/A

Default value: OPENJPA_SEQUENCE

Parent elements:

```
persistence-configuration  
  persistence-configuration-unit  
    native-jdbc-seq
```

C.254.1 Function

Specifies the name of the database sequence. For more information, see "Generators" in the *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.254.2 Example

```
<sequence>OPENJPA_SEQUENCE</sequence>
```


C.255 sequence-column

Range of values: String

Default value: SEQUENCE_VALUE

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
    class-table-jdbc-seq
```

```
persistence-configuration
  persistence-configuration-unit
    table-jdbc-seq
```

```
persistence-configuration
  persistence-configuration-unit
    value-table-jdbc-seq
```

C.255.1 Function

Specifies the name of the column that holds the current sequence value. For more information, see "Generators" in the *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.255.2 Example

```
<sequence-column>SEQUENCE_VALUE</sequence-column>
```

C.256 sequence-name

Range of values: String

Default value: OPENJPA_SEQUENCE

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
    native-jdbc-seq
```

C.256.1 Function

Specifies the name of the database sequence. For more information, see "Generators" in the *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.256.2 Example

```
<sequence-name>OPENJPA_SEQUENCE</sequence-name>
```

C.257 ServiceURL

Range of values: Valid service URL

Default value: service:jmx:rmi://localhost/jndi/jmxservice

Parent elements:

```
jmx
```

jmx2-jmx

C.257.1 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 4.2.0 Developers Guide for JPA/JDO*.

C.257.2 Example

```
<ServiceURL>service:jmx:rmi://localhost/jndi/jmxservice</ServiceURL>
```

C.258 simple-driver-data-source

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration  
  persistence-configuration-unit
```

C.258.1 Function

Configures the Kodo datasource implementation. For more information, see "Using the Kodo DataSource" in *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.258.2 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>
```

C.259 single-file

Range of values: true | false

Default value: false

Parent elements:

```
persistence-configuration  
  persistence-configuration-unit  
    mapping-file-deprecated-jdo-mapping-factory
```

C.259.1 Function

Flag specifying whether or not to use a single file.

C.259.2 Example

```
<single-file>>false</single-file>
```

C.260 single-jvm-exclusive-lock-manager

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.260.1 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 4.2.0 Developers Guide for JPA/JDO*.

C.260.2 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>
```

C.261 single-jvm-remote-commit-provider

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.261.1 Function

Configures Kodo to share commit notifications among `persistenceManagerFactories` in the same JVM. For more information, see "Remote Commit Provider Configuration" in the *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.261.2 Example

```
<single-jvm-remote-commit-provider>
  <name>kodo.RemoteCommitProvider</name>
</single-jvm-remote-commit-provider>
```

C.262 soft-reference-size

Range of values: Integer

Default value: -1

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
    kodo-concurrent-data-cache
```

```
persistence-configuration
  persistence-configuration-unit
    kodo-concurrent-query-cache
```

```
persistence-configuration
  persistence-configuration-unit
    lru-data-cache
```

```
persistence-configuration
  persistence-configuration-unit
    lru-query-cache
```

```
persistence-configuration
  persistence-configuration-unit
    query-compilation-cache
```

C.262.1 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 4.2.0 Developers Guide for JPA/JDO*.

C.262.2 Example

```
<soft-reference-size>-1</soft-reference-size>
```

C.263 so-timeout

Range of values: Integer

Default value: 0

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
    http-transport
```

C.263.1 Function

Specifies the socket read timeout in milliseconds. For more information, see "Standalone Persistence Server" in *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.263.2 Example

```
<so-timeout>0</so-timeout>
```

C.264 sql-server-dictionary

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.264.1 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 *Oracle WebLogic Server Administration Console Help*.

C.264.2 Example

The `sql-server-dictionary` element shares the same subelements as [Section C.5, "access-dictionary,"](#) plus the following subelements (defaults shown):

```
<sql-server-dictionary>
...
  <unique-identifier-as-varbinary>true</unique-identifier-as-varbinary>
</sql-server-dictionary>
```

C.265 stack-execution-context-name-provider

Range of values: N/A

Default value: N/A

Parent elements:

```
execution-context-name-provider
```

C.265.1 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 4.2.0 Developers Guide for JPA/JDO*.

C.265.2 Example

```
<stack-execution-context-name-provider>
  <style>full</style>
</stack-execution-context-name-provider>
```

C.266 store-mode

Range of values: Integer

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
    deprecated-jdo-meta-data-factory

persistence-configuration
  persistence-configuration-unit
    extension-deprecated-jdo-mapping-factory

persistence-configuration
  persistence-configuration-unit
    jdo-meta-data-factory

persistence-configuration
  persistence-configuration-unit
    jdor-mapping-factory

persistence-configuration
  persistence-configuration-unit
    kodo-persistence-mapping-factory

persistence-configuration
  persistence-configuration-unit
    kodo-persistence-meta-data-factory

persistence-configuration
  persistence-configuration-unit
    mapping-file-deprecated-jdo-mapping-factory

persistence-configuration
  persistence-configuration-unit
    orm-file-jdor-mapping-factory

persistence-configuration
  persistence-configuration-unit
    table-deprecated-jdo-mapping-factory

persistence-configuration
  persistence-configuration-unit
    table-jdor-mapping-factory
```

C.266.1 Function

Specifies the data store mode.

C.266.2 Example

```
<store-mode>1</store-mode>
```

C.267 strict

Range of values: true | false

Default value: false

Parent elements:

```

persistence-configuration
  persistence-configuration-unit
    deprecated-jdo-meta-data-factory

persistence-configuration
  persistence-configuration-unit
    extension-deprecated-jdo-mapping-factory

persistence-configuration
  persistence-configuration-unit
    jdo-meta-data-factory

persistence-configuration
  persistence-configuration-unit
    jdor-mapping-factory

persistence-configuration
  persistence-configuration-unit
    kodo-persistence-mapping-factory

persistence-configuration
  persistence-configuration-unit
    kodo-persistence-meta-data-factory

persistence-configuration
  persistence-configuration-unit
    mapping-file-deprecated-jdo-mapping-factory

persistence-configuration
  persistence-configuration-unit
    orm-file-jdor-mapping-factory

persistence-configuration
  persistence-configuration-unit
    table-deprecated-jdo-mapping-factory

persistence-configuration
  persistence-configuration-unit
    table-jdor-mapping-factory

```

C.267.1 Function

Flag that specifies whether strict mode is enabled.

C.267.2 Example

```
<strict>true</strict>
```

C.268 strict-identity-values

Range of values: true | false

Default value: false

Parent elements:

```

persistence-configuration
  persistence-configuration-unit
    compatibility

```

C.268.1 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 4.2.0 Developers Guide for JPA/JDO*.

C.268.2 Example

```
<strict-identity-values>false</strict-identity-values>
```

C.269 style

Range of values: line | partial | full

Default value: N/A

Parent elements:

```
execution-context-name-provider  
stack-execution-context-name-provider
```

C.269.1 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 4.2.0 Developers Guide for JPA/JDO*.

C.269.2 Example

```
<style>full</style>
```

C.270 subclass-fetch-mode

Range of values: join | multiple | none | parallel | single

Default value: join

Parent elements:

```
execution-context-name-provider  
stack-execution-context-name-provider
```

C.270.1 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 [Section C.120, "eager-fetch-mode."](#) For complete details, see "Eager Fetching" in *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.270.2 Example

```
<subclass-fetch-mode>join</subclass-fetch-mode>
```

C.271 sybase-dictionary

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.271.1 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 *Oracle WebLogic Server Administration Console Help*.

C.271.2 Example

The `sybase-dictionary` element shares the same subelements as [Section C.5, "access-dictionary,"](#) 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>
```

C.272 sync-with-managed-transactions

Range of values: true | false

Default value: false

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
    kodo-broker
```

C.272.1 Function

Flag that specifies whether to sync with managed transactions. For more information, see "kodo.BrokerImpl" in *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.272.2 Example

```
<sync-with-managed-transactions>false</sync-with-managed-transactions>
```

C.273 synchronize-mappings

Range of values: true | false

Default value: false

Parent elements:

persistence-configuration

C.273.1 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.

C.273.2 Example

```
<synchronize-mappings>false</synchronize-mappings>
```

C.274 table

Range of values: String

Default value: See below

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
    class-table-jdbc-seq
```

```
persistence-configuration
  persistence-configuration-unit
    table-jdbc-seq
```

```
persistence-configuration
  persistence-configuration-unit
    table-jdor-mapping-factory
```

```
persistence-configuration
  persistence-configuration-unit
    table-schema-factory
```

```
persistence-configuration
  persistence-configuration-unit
    value-table-jdbc-seq
```

C.274.1 Function

Specifies the name of the table. For more information, see "Mapping Factory" in *Kodo 4.2.0 Developers Guide for JPA/JDO*. 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

C.274.2 Example

```
<table>KODO_JDO_MAPPING</table>
```

C.275 table-deprecated-jdo-mapping-factory

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.275.1 Function

A plugin string (see "Plugin Configuration" in *Kodo 4.2.0 Developers Guide for JPA/JDO*) 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 4.2.0 Developers Guide for JPA/JDO*.

C.275.2 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>
```

C.276 table-jdbc-seq

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.276.1 Function

Specifies a plugin string (see "Plugin Configuration" in *Kodo 4.2.0 Developers Guide for JPA/JDO*) 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 4.2.0 Developers Guide for JPA/JDO*.

C.276.2 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>
```

C.277 table-jdor-mapping-factory

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.277.1 Function

A plugin string (see "Plugin Configuration" in *Kodo 4.2.0 Developers Guide for JPA/JDO*) 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 4.2.0 Developers Guide for JPA/JDO*.

C.277.2 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>
```

C.278 table-lock-update-manager

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
```

persistence-configuration-unit

C.278.1 Function

Defines an update manager capable of statement batching, that ignores foreign keys and auto-increment, 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 4.2.0 Developers Guide for JPA/JDO*.

C.278.2 Example

```
<table-lock-update-manager>
  <maximize-batch-size>true</maximize-batch-size>
</table-lock-update-manager>
```

C.279 table-name

Range of values: String

Default value: See below

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
    class-table-jdbc-seq
```

```
persistence-configuration
  persistence-configuration-unit
    native-jdbc-seq
```

```
persistence-configuration
  persistence-configuration-unit
    table-deprecated-jdo-mapping-factory
```

```
persistence-configuration
  persistence-configuration-unit
    table-jdbc-seq
```

```
persistence-configuration
  persistence-configuration-unit
    table-schema-factory
```

```
persistence-configuration
  persistence-configuration-unit
    value-table-jdbc-seq
```

C.279.1 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 4.2.0 Developers Guide for JPA/JDO*.
- [Section C.208, "native-jdbc-seq"](#): `DUAL`. For more information, see "Generators" in the *Kodo 4.2.0 Developers Guide for JPA/JDO*.

- [Section C.275, "table-deprecated-jdo-mapping-factory"](#): JDO_MAPPING. For more information, see "Mapping Factory" in *Kodo 4.2.0 Developers Guide for JPA/JDO*.
- [Section C.276, "table-jdbc-seq"](#): OPENJPA_SEQUENCE_TABLE. For more information, see "Generators" in the *Kodo 4.2.0 Developers Guide for JPA/JDO*.
- [Section C.280, "table-schema-factory"](#): OPENJPA_SCHEMA. For more information, see "Schema Factory" in the *Kodo 4.2.0 Developers Guide for JPA/JDO*.
- [Section C.305, "value-table-jdbc-seq"](#): OPENJPA_SEQUENCE_TABLE. For more information, see "Generators" in the *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.279.2 Example

```
<table-name>JDO_MAPPING</table-name>
```

C.280 table-schema-factory

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.280.1 Function

Specifies a plugin string (see "Plugin Configuration" in *Kodo 4.2.0 Developers Guide for JPA/JDO*) 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 4.2.0 Developers Guide for JPA/JDO*.

C.280.2 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>
```

C.281 tangosol-cache-name

Range of values: String

Default value: kodo

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
    tangosol-data-cache
```

C.281.1 Function

Name of the Tangosol Coherence cache to use. For more information, see "Tangosol Integration" in *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.281.2 Example

```
<tangosol-cache-name>kodo</tangosol-cache-name>
```

C.282 tangosol-cache-type

Range of values: named | distributed | replicated

Default value: named

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
    tangosol-data-cache
```

```
persistence-configuration
  persistence-configuration-unit
    tangosol-query-cache
```

C.282.1 Function

Type of Tangosol Coherence cache to use. Valid values include:

- `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.
- `distributed`
- `replicated`

For more information, see "Tangosol Integration" in *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.282.2 Example

```
<tangosol-cache-type>name</tangosol-cache-type>
```

C.283 tangosol-data-cache

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
    data-caches
```

C.283.1 Function

Integrates with Tangosol's Coherence caching system. For more information, see "Tangosol Integration" in *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.283.2 Example

```
<data-caches>
  <tangosol-data-cache>
    <name>kodo.dataCache</name>
    <clear-on-close>false</clear-on-close>
    <tangosol-cache-type>name</tangosol-cache-type>
    <tangosol-cache-name>kodo</tangosol-cache-name>
    <eviction-schedule>15,45 15 * * 1</eviction-schedule>
  </tangosol-data-cache>
</data-caches>
```

C.284 tcp-remote-commit-provider

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.284.1 Function

Configures the TCP remote commit provider. For more information, see "TCP" in the *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.284.2 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>
```

C.285 tcp-transport

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```


C.285.1 Function

Specifies the transport layer for the remote communication. For more information, see "Standalone Persistence Server" in *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.285.2 Example

```
<tcp-transport>
  <so-timeout>0</so-timeout>
  <host>localhost</host>
  <port>5637</port>
</tcp-transport>
```

C.286 time-seeded-seq

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.286.1 Function

Specifies a plugin string (see "Plugin Configuration" in *Kodo 4.2.0 Developers Guide for JPA/JDO*) 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 4.2.0 Developers Guide for JPA/JDO*.

C.286.2 Example

```
<time-seeded-seq>
  <type>0</type>
  <increment>1</increment>
</time-seeded-seq>
```

C.287 topic

Range of values: Valid topic

Default value: `topic/KodoCommitProviderTopic`

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
    jms-remote-commit-provider
```

C.287.1 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 4.2.0 Developers Guide for JPA/JDO*.

C.287.2 Example

```
<topic>topic/KodoCommitProviderTopic</topic>
```

C.288 topic-connection-factory

Range of values: Valid connection factory

Default value: java:/ConnectionFactory

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
    jms-remote-commit-provider
```

C.288.1 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 4.2.0 Developers Guide for JPA/JDO*.

C.288.2 Example

```
<topic-connection-factory>java:/ConnectionFactory</topic-connection-factory>
```

C.289 track-changes

Range of values: true | false

Default value: true

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.289.1 Function

Flag that specifies whether to use smart proxies. For more information, see "Custom Proxies" in the *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.289.2 Example

```
<track-changes>true</track-changes>
```

C.290 transaction-isolation

Range of values: default | none | read-committed | read-uncommitted | repeatable-read | serializable

Default value: default

Parent elements:

persistence-configuration
persistence-configuration-unit

C.290.1 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.

C.290.2 Example

```
<transaction-isolation>default</transaction-isolation>
```

C.291 transaction-mode

Range of values: local | managed

Default value: local

Parent elements:

persistence-configuration
persistence-configuration-unit

C.291.1 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.

C.291.2 Example

```
<transaction-mode>local</transaction-mode>
```

C.292 transaction-name-execution-context-name-provider

Range of values: N/A

Default value: N/A

Parent elements:

```
execution-context-name-provider
```

C.292.1 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 4.2.0 Developers Guide for JPA/JDO*.

C.292.2 Example

```
<transacation-name-execution-context-name-provider/>
```

C.293 type

Range of values: N/A

Default value: 0

Parent elements:

```
persistence-configuration  
  persistence-configuration-unit  
    class-table-jdbc-seq
```

```
persistence-configuration  
  persistence-configuration-unit  
    native-jdbc-seq
```

```
persistence-configuration  
  persistence-configuration-unit  
    table-jdbc-seq
```

```
persistence-configuration  
  persistence-configuration-unit  
    time-seeded-seq
```

```
persistence-configuration  
  persistence-configuration-unit  
    value-table-jdbc-seq
```

C.293.1 Function

Specifies the type of generator. For more information, see "Generators" in the *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.293.2 Example

```
<type>0</type>
```

C.294 type-column

Range of values: 0 | 1 | 2 | 3

Default value: 0

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
    table-jdor-mapping-factory
```

C.294.1 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 4.2.0 Developers Guide for JPA/JDO*.

C.294.2 Example

```
<type-column>0</type-column>
```

C.295 types

Range of values: Valid persistent class names

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
    tangosol-data-cache
```

C.295.1 Function

Specifies a semicolon-separated list of fully-qualified persistent class names. For more information, see "Metadata Factory" in *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.295.2 Example

```
<types>classname1; classname2</types>
```

C.296 URL

Range of values: Valid URL

Default value: `t3://localhost:7001`

Parent elements:

```
jmx
  wls81-jmx
```

C.296.1 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 4.2.0 Developers Guide for JPA/JDO*.

C.296.2 Example

```
<URL>t3://localhost:7001</URL>
```

C.297 url

Range of values: Valid URL

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
    http-transport
```

C.297.1 Function

URL to which Kodo should connect to access the remote server. For more information, see "Client Managers" in *Kodo 4.2.0 Developers Guide for JPA/JDO*.

C.297.2 Example

```
<url>t3://localhost:7001/remote-server</url>
```

C.298 urls

Range of values: Valid URL

Default value: `t3://localhost:7001`

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
    deprecated-jdo-meta-data-factory

persistence-configuration
  persistence-configuration-unit
    extension-deprecated-jdo-mapping-factory

persistence-configuration
```

```

persistence-configuration-unit
  jdo-meta-data-factory

persistence-configuration
  persistence-configuration-unit
  jdor-mapping-factory

persistence-configuration
  persistence-configuration-unit
  kodo-persistence-mapping-factory

persistence-configuration
  persistence-configuration-unit
  kodo-persistence-meta-data-factory

persistence-configuration
  persistence-configuration-unit
  mapping-file-deprecated-jdo-mapping-factory

persistence-configuration
  persistence-configuration-unit
  orm-file-jdor-mapping-factory

persistence-configuration
  persistence-configuration-unit
  table-deprecated-jdo-mapping-factory

```

C.298.1 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 4.2.0 Developers Guide for JPA/JDO*.

C.298.2 Example

```
<urls>t3://localhost:7001/metadata.jar</urls>
```

C.299 use-aliases

Range of values: true | false

Default value: false

Parent elements:

```

persistence-configuration
  persistence-configuration-unit
    class-table-jdbc-seq

```

C.299.1 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 4.2.0 Developers Guide for JPA/JDO*.

C.299.2 Example

```
<use-aliases>false</use-aliases>
```

C.300 use-schema-validation

Range of values: true | false

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
    deprecated-jdo-meta-data-factory

persistence-configuration
  persistence-configuration-unit
    extension-deprecated-jdo-mapping-factory

persistence-configuration
  persistence-configuration-unit
    jdo-meta-data-factory

persistence-configuration
  persistence-configuration-unit
    jdor-mapping-factory

persistence-configuration
  persistence-configuration-unit
    mapping-file-deprecated-jdo-mapping-factory

persistence-configuration
  persistence-configuration-unit
    orm-file-jdor-mapping-factory

persistence-configuration
  persistence-configuration-unit
    table-deprecated-jdo-mapping-factory

persistence-configuration
  persistence-configuration-unit
    table-jdor-mapping-factory
```

C.300.1 Function

Flag that specifies whether to use schema validation.

C.300.2 Example

```
<use-schema-validation>true</use-schema-validation>
```

C.301 user-object-execution-context-name-provider

Range of values: N/A

Default value: N/A

Parent elements:


```
execution-context-name-provider
```

C.301.1 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 4.2.0 Developers Guide for JPA/JDO*.

C.301.2 Example

```
<user-object-execution-context-name-provider/>
```

C.302 UserName

Range of values: N/A

Default value: N/A

Parent elements:

```
jmx
  wls81-jmx
```

C.302.1 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 4.2.0 Developers Guide for JPA/JDO*.

C.302.2 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>
```

C.303 validate-false-returns-hollow

Range of values: true | false

Default value: true

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
    compatibility
```

C.303.1 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 4.2.0 Developers Guide for JPA/JDO*.

C.303.2 Example

```
<validate-false-returns-hollow>true</validate-false-returns-hollow>
```

C.304 validate-true-checks-store

Range of values: `true` | `false`

Default value: `false`

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
    compatibility
```

C.304.1 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 4.2.0 Developers Guide for JPA/JDO*.

C.304.2 Example

```
<validate-true-checks-store>false</validate-true-checks-store>
```

C.305 value-table-jdbc-seq

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.305.1 Function

Specifies a plugin string (see "Plugin Configuration" in *Kodo 4.2.0 Developers Guide for JPA/JDO*) 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 [Section C.276, "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 4.2.0 Developers Guide for JPA/JDO*.

C.305.2 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>
```

C.306 version-check-on-read-lock

Range of values: true | false

Default value: false

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
    pessimistic-lock-manager
```

C.306.1 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 4.2.0 Developers Guide for JPA/JDO*.

C.306.2 Example

```
<version-check-on-read-lock>false</version-check-on-read-lock>
```

C.307 version-check-on-write-lock

Range of values: true | false

Default value: false

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
    pessimistic-lock-manager
```

C.307.1 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 4.2.0 Developers Guide for JPA/JDO*.

C.307.2 Example

```
<version-check-on-write-lock>false</version-check-on-write-lock>
```

C.308 version-lock-manager

Range of values: N/A

Default value: N/A

Parent elements:

```
persistence-configuration
  persistence-configuration-unit
```

C.308.1 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 4.2.0 Developers Guide for JPA/JDO*.

C.308.2 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>
```

C.309 wls81-jmx

Range of values: N/A

Default value: N/A

Parent elements:

```
jmx
```

C.309.1 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 4.2.0 Developers Guide for JPA/JDO*.

C.309.2 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>
```

C.310 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
```

C.310.1 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 4.2.0 Developers Guide for JPA/JDO*.

C.310.2 Example

```
<write-lock-level>read</write-lock-level>
```

