

# OJB

## Table of contents

1 OJB.....	8
1.1 Apache ObJectRelationalBridge - OJB.....	8
1.1.1 Summary.....	8
1.1.1.1 flexibility.....	8
1.1.1.2 scalability.....	8
1.1.1.3 functionality.....	8
1.2 News.....	8
1.2.1 .....	9
1.3 OJB - Features.....	9
1.4 OJB - API Status.....	10
1.5 OJB - References and Testimonials.....	11
1.5.1 References and Testimonials.....	11
1.5.1.1 projects using OJB.....	11
1.5.1.2 user testimonials .....	12
1.6 OJB - Mailing Lists.....	13
1.6.1 Mailing Lists.....	13
1.7 OJB - Mailing Lists Archives.....	13
1.7.1 Mailing Lists Archives.....	13
1.8 OJB - Links and further readings.....	13
1.8.1 Summary.....	13
1.8.2 OJB - Tutorials and Howto.....	14
1.8.2.1 OJB and Spring.....	14
1.8.3 Design.....	14
1.8.4 Further readings on O/R mapping.....	15
1.8.5 Patterns.....	15
1.8.6 Books covering OJB.....	15
2 Download.....	15
3 Development.....	16
3.1 Coding Standards.....	16
4 Index.....	17
4.1 Site Linkmap.....	17
4.1.1 Table of Contents.....	17
5 Documentation.....	20

5.1 Documentation - Introduction.....	20
5.2 Frequently Asked Questions.....	20
5.2.1 Questions.....	20
5.2.1.1 1. General.....	20
5.2.1.2 2. Getting Started.....	22
5.2.1.3 3. OJB APIs.....	24
5.2.1.4 4. Howto.....	25
5.3 Getting Started.....	34
5.3.1 Acquiring ojb-blank.....	34
5.3.2 Contents of ojb-blank.....	35
5.3.2.1 Sample project.....	36
5.3.3 The build files.....	36
5.3.3.1 Configuration via build.properties.....	36
5.3.3.2 Building via build.xml.....	37
5.3.3.3 Sample project.....	38
5.3.4 The runtime configuration files.....	41
5.3.4.1 Configuring the OJB runtime.....	41
5.3.4.2 Configuring the database connection.....	41
5.3.4.3 Configuring the repository.....	41
5.3.4.4 Sample project.....	41
5.3.5 Learning More.....	43
5.4 Tutorials.....	43
5.4.1 Tutorial Summary.....	43
5.4.2 Mapping Tutorial.....	44
5.4.2.1 What is the Object-Relational Mapping Metadata?.....	44
5.4.2.2 Advanced Topics.....	47
5.4.3 Persistence Broker Tutorial.....	47
5.4.3.1 The PersistenceBroker API.....	47
5.4.3.2 Exception Handling.....	52
5.4.4 The ODMG API.....	53
5.4.4.1 Introduction.....	53
5.4.4.2 Initializing ODMG.....	54
5.4.4.3 Persisting New Objects.....	54
5.4.4.4 Querying Persistent Objects.....	55
5.4.4.5 Updating Persistent Objects.....	55
5.4.4.6 Deleting Persistent Objects.....	55
5.4.5 JDO Tutorial.....	56
5.4.5.1 Using the ObJectRelationalBridge JDO API.....	56

5.4.5.2 Using the JDO API in the UseCase Implementations.....	57
5.4.5.3 Conclusion.....	61
5.4.6 Object Transaction Manager Tutorial.....	61
5.4.6.1 The OTM API.....	61
5.4.6.2 Notes on the Object Transaction Manager.....	66
5.5 Reference Guides.....	67
5.5.1 Reference Guides Summary.....	67
5.5.2 PB-api Guide.....	67
5.5.2.1 Introduction.....	68
5.5.2.2 How to access the PB-api?.....	68
5.5.2.3 Notes on Using the PersistenceBroker API.....	68
5.5.2.4 Questions.....	69
5.5.3 ODMG-api Guide.....	70
5.5.3.1 Introduction.....	70
5.5.3.2 Specific Metadata Settings.....	70
5.5.3.3 How to access ODMG-api.....	71
5.5.3.4 Configuration Properties.....	71
5.5.3.5 OJB Extensions of ODMG.....	73
5.5.3.6 Notes on Using the ODMG API.....	74
5.5.3.7 ODMG Named Objects.....	75
5.5.3.8 ODMG's DCollections.....	78
5.5.3.9 Foreign Keys Constraints and ODMG-api.....	78
5.5.3.10 Questions and Tips.....	78
5.5.4 Platforms.....	80
5.5.4.1 How to use OJB with a specific relational database.....	80
5.5.4.2 Basic Concepts.....	80
5.5.4.3 The setup process.....	83
5.5.5 OJB.properties Configuration File.....	85
5.5.5.1 OJB Configuration.....	85
5.5.5.2 OJB.properties File.....	85
5.5.6 JDBC Types.....	85
5.5.6.1 Mapping of JDBC Types to Java Types.....	85
5.5.6.2 Type and Value Conversions - The FieldConversion Interface.....	86
5.5.7 Repository File.....	89
5.5.7.1 Introduction - repository syntax.....	89
5.5.7.2 descriptor-repository.....	89
5.5.7.3 jdbc-connection-descriptor.....	90
5.5.7.4 connection-pool.....	92

5.5.7.5 sequence-manager.....	96
5.5.7.6 object-cache.....	97
5.5.7.7 custom attribute.....	98
5.5.7.8 class-descriptor.....	99
5.5.7.9 extent-class.....	100
5.5.7.10 field-descriptor.....	101
5.5.7.11 reference-descriptor.....	102
5.5.7.12 foreignkey.....	104
5.5.7.13 collection-descriptor.....	104
5.5.7.14 order-by.....	106
5.5.7.15 inverse-foreignkey.....	106
5.5.7.16 fk-pointing-to-this-class.....	106
5.5.7.17 fk-pointing-to-element-class.....	107
5.5.7.18 query-customizer.....	107
5.5.7.19 index-descriptor.....	107
5.5.7.20 index-column.....	107
5.5.7.21 Stored Procedure Support.....	107
5.5.8 Basic O/R Mapping Technique.....	110
5.5.8.1 Mapping 1:1 associations.....	110
5.5.8.2 Mapping 1:n associations.....	112
5.5.8.3 Mapping m:n associations.....	115
5.5.8.4 Setting Load, Update, and Delete Cascading.....	120
5.5.8.5 Using Proxy Classes.....	122
5.5.8.6 Type and Value Conversions.....	128
5.5.9 Advanced O/R Mapping Technique.....	128
5.5.9.1 Extents and Polymorphism.....	128
5.5.9.2 Mapping Inheritance Hierarchies.....	131
5.5.9.3 Using interfaces with OJB.....	140
5.5.9.4 Change PersistentField Class.....	143
5.5.9.5 How do anonymous keys work?.....	144
5.5.9.6 Using Rowreader.....	145
5.5.9.7 Nested Objects.....	147
5.5.9.8 Instance Callbacks.....	149
5.5.9.9 Manageable Collection.....	150
5.5.9.10 Customizing collection queries.....	152
5.5.9.11 Metadata runtime changes.....	153
5.5.10 OJB Queries.....	153
5.5.10.1 Introduction.....	153

5.5.10.2 Query by Criteria.....	153
5.5.10.3 ODMG OQL.....	162
5.5.10.4 JDO queries.....	162
5.5.11 Metadata handling.....	162
5.5.11.1 Introduction.....	162
5.5.11.2 When does OJB read metadata.....	163
5.5.11.3 Connection metadata.....	163
5.5.11.4 Persistent object metadata.....	164
5.5.11.5 Questions.....	167
5.5.12 Deployment.....	168
5.5.12.1 Introduction.....	168
5.5.12.2 Things needed for deploying OJB.....	168
5.5.12.3 Deployment in standalone applications.....	170
5.5.12.4 Deployment in servlet based applications.....	170
5.5.12.5 Deployment in managed environment (e.g. EJB based).....	170
5.5.13 Connection Handling.....	181
5.5.13.1 Introduction.....	181
5.5.13.2 ConnectionFactory.....	181
5.5.13.3 ConnectionManager.....	182
5.5.13.4 Questions and Answers.....	183
5.5.14 The Object Cache.....	184
5.5.14.1 Introduction.....	184
5.5.14.2 Why a cache and how it works?.....	185
5.5.14.3 How to declare and change the used ObjectCache implementation.....	185
5.5.14.4 Shipped cache implementations:.....	187
5.5.14.5 Distributed ObjectCache?.....	194
5.5.14.6 Implement your own cache.....	194
5.5.14.7 Future prospects.....	194
5.5.15 Sequence Manager.....	194
5.5.15.1 The OJB Sequence Manager.....	194
5.5.16 OJB logging configuration.....	206
5.5.16.1 Logging in OJB.....	206
5.5.16.2 Logging configuration within OJB.....	207
5.5.16.3 Logging configuration via configuration files.....	208
5.5.16.4 Logging configuration at runtime.....	210
5.5.16.5 Defining your own logger.....	210
5.5.17 Locking.....	210
5.5.17.1 Introduction.....	210

5.5.17.2 Optimistic Locking.....	211
5.5.17.3 Pessimistic-Locking.....	212
5.5.17.4 ODMG-api Locking.....	216
5.5.17.5 Locking in distributed environment.....	217
5.5.17.6 Pluggin own locking classes.....	217
5.5.18 XDoclet OJB module documentation.....	218
5.5.18.1 Acquiring and building.....	218
5.5.18.2 Usage.....	219
5.5.18.3 Tag reference.....	221
5.5.18.4 Interfaces and Classes.....	221
5.5.18.5 Fields and Bean properties.....	230
5.5.18.6 References.....	235
5.5.18.7 Collections.....	237
5.5.18.8 Nested objects.....	240
5.5.19 OJB Performance.....	243
5.5.19.1 Introduction.....	243
5.5.19.2 The Performance Test Suite.....	243
5.5.19.3 How OJB compares to other O/R mapping tools?.....	247
5.5.19.4 What are the best settings for maximal performance?.....	248
5.6 Howto's.....	248
5.6.1 Howto's Summary.....	248
5.6.2 HOWTO - Build O/R Mapping Files.....	248
5.6.2.1 How to build O/R mapping files.....	249
5.6.2.2 classification of O/R related transformations.....	249
5.6.2.3 Forward engineering from XMI.....	249
5.6.2.4 Forward engineering from Torque.....	251
5.6.2.5 Forward engineering from repository.xml.....	252
5.6.2.6 XDoclet transformation from Java code.....	252
5.6.2.7 Reverse engineering from database.....	252
5.6.3 HOWTO - Use Anonymous Keys.....	253
5.6.3.1 Why Do We Need Anonymous Keys?.....	253
5.6.3.2 How it works.....	254
5.6.3.3 Using Anonymous Keys.....	254
5.6.3.4 Benefits and Drawbacks.....	258
5.6.4 HOWTO - Use DB Sequences.....	258
5.6.4.1 Introduction.....	258
5.6.4.2 The Sample Database.....	258
5.6.4.3 Using OJB.....	259

5.6.5 HOWTO - Work with LOB Data Types.....	261
5.6.5.1 Using Oracle LOB Data Types with OJB.....	261
5.6.5.2 Backgrounder: Large objects in databases.....	261
5.6.5.3 Large Objects in OJB.....	263
5.6.6 HOWTO - Use OJB in clustered environments.....	263
5.6.6.1 How to use OJB in clustered environments.....	263
5.6.6.2 Three steps to clustering your OJB application.....	264
5.6.6.3 Notes.....	265
5.6.7 HOWTO - Stored Procedure Support.....	266
5.6.7.1 Introduction.....	266
5.6.7.2 Repository entries.....	266
5.6.7.3 Common attributes.....	267
5.6.7.4 insert-procedure.....	267
5.6.7.5 update-procedure.....	267
5.6.7.6 delete-procedure.....	268
5.6.7.7 Argument descriptors.....	268
5.6.7.8 A simple example.....	269
5.6.7.9 A complex example.....	273
5.6.8 HOWTO - Spring with OJB.....	275
5.6.8.1 Spring with OJB.....	275
5.7 Testing.....	275
5.7.1 Testing Summary.....	276
5.7.2 JUnit Test Suite.....	276
5.7.2.1 Introduction.....	276
5.7.2.2 How to run the Test Suite.....	276
5.7.2.3 What about known issues?.....	277
5.7.2.4 Donate own tests for OJB Test Suite.....	277
5.7.3 Write Tests.....	277
5.7.3.1 Introduction.....	277
5.7.3.2 How to write a new Test.....	278

## 1. OJB

### 1.1. Apache ObJectRelationalBridge - OJB

#### 1.1.1. Summary

Apache ObJectRelationalBridge (OJB) is an Object/Relational mapping tool that allows transparent persistence for Java Objects against relational databases.

##### 1.1.1.1. flexibility

OJB supports multiple persistence APIs to provide users with their API of choice:

- A **PersistenceBroker** API which serves as the OJB persistence kernel. The OTM-, ODMG- and JDO-implementations are built on top of this kernel. This API can also be used directly by applications that don't need full fledged object level transactions. (See the [Persistence Broker Tutorial](#) for details.)
- A full featured **ODMG 3.0** compliant API. (See the [ODMG Tutorial](#) for an introduction.)
- A **JDO** compliant API. We currently provide a plugin to the JDO Reference Implementation (RI). Combining the JDO RI and our plugin provides a JDO 1.0 compliant o/r solution. A full JDO implementation is scheduled for OJB 2.0. (See the [JDO tutorial](#) for an introduction to the JDO programming model.)
- An Object Transaction Manager (OTM) layer that contains all features that JDO and ODMG have in common. (See the [OTM tutorial](#) for details).

See the [FAQ](#) for a detailed view of the OJB layering. Get the latest information on each API's [status](#).

##### 1.1.1.2. scalability

OJB has been designed for a large range of applications, from embedded systems to rich client application to multi-tier J2EE based architectures.

OJB integrates smoothly into J2EE Application servers. It supports JNDI lookup of datasources. It ships with full JTA and JCA integration. OJB can be used within JSPs, Servlets and SessionBeans. OJB provides special support for Bean Managed EntityBeans (BMP).

##### 1.1.1.3. functionality

OJB uses an XML based Object/Relational mapping. The mapping resides in a dynamic MetaData layer, which can be manipulated at runtime through a simple Meta-Object-Protocol (MOP) to change the behaviour of the persistence kernel.

OJB provides several advanced O/R features like [Object Caching](#), lazy materialization through [virtual proxies](#) and distributed [lock-management](#) with configurable Transaction-Isolation levels. Optimistic and pessimistic locking is supported.

OJB provides a flexible configuration and plugin mechanism that allows to select from set of predefined components or to implement your own extensions and plugins.

A more complete [featurelist can be found here](#).

Learn more about the OJB design principles [in this document](#).

## 1.2. News



### 1.2.1.

12/2005 - OJB 1.0.4.rc1 released

Contains bug fixes and new features. For more details see [release-notes](#).

11/2005 - Work for OJB 1.x started

The 1.x release of OJB will be a big step forward in usability and integration in existing frameworks (like Spring).

08/2005 - Linguine Maps supports OJB

[Linguine Maps](#) is an open-source (LGPL, Java 1.4+) utility that will automatically produce easy to read UML-style entity-relation diagrams from OJB mapping files.

## 1.3. OJB - Features

- Supports both standard- and non-standard APIs:
  - PB API (non-standard)
  - ODMG API (standard)
  - OTM API (non-standard)
  - JDO API (standard)
- The PersistenceBroker kernel API and all top-level APIs (ODMG, OTM, JDO) allows Java Programmers to store and retrieve Java Objects in/from (any) JDBC-compliant RDBMS
- Transparent persistence: classes does not have to inherit from any OJB base class nor implement a special interface. OJB delivers pure transparent persistence for POJOs.
- Scalable architecture that allows to build massively distributed and clustered systems.
- [Configurable persistence by reachability](#): All Objects associated to a persistent object by references can made persitent too.
- Extremely flexible design with pluggable implementation of most service classes like *PersistenceBroker*, *ObjectCache*, *SequenceManager*, *RowReader*, *ConnectionFactory*, *ConnectionManager*, *IndirectionHandler*, *SQLGenerator*, *JdbcAccess*, ... and so on.
- Quality assurance taken seriously: More than **800 [JUnit Test Cases](#)** for regression tests. JUnit tests are integrated into the build scripts and used as quality assurance for daily development.
- Mapping support for 1:1, 1:n and m:n [associations](#).
- Configurable collection queries to control loading of relationships. See [QueryCustomizer](#).
- [Automatic](#) and [manual](#) assignment of foreign key values.
- The Object / Relational mapping is defined in an XML Repository. The mapping is completely dynamic and can be [manipulated at runtime](#) for maximum flexibility
- Easy use of [multiple databases](#).
- Configurable Lazy Materialization through [Proxy](#) support in the PersistenceBroker. The user can implement specific Proxy classes or let OJB generate dynamic Proxies.
- Support for [Polymorphism and Extents](#). You can use Interface-types and abstract classes as attribute types in your persistent classes. Queries are also aware of extents: A query against a baseclass or interface will return matches from derived classes, even if they are mapped to different DB-tables
- Support for Java Array- and Collection-attributes in persistent classes. The [attribute-types](#) can be Arrays, java.util.Collection or may be user defined collections that implement the interface `ojb.broker.ManageableCollection`.
- [Sequence-Managing](#) . The SequenceManager is aware of "extents" and maintains uniqueness of ids accross any number of tables. Sequence Numbering can be declared in the mapping repository.  
Native Database based Sequence Numbering is also supported.
- Reusing Prepared Statements, internal [connection pooling](#).
- Integrates smoothly in controlled environments like [EJB containers](#)

- Full JTA and JCA (in progress) Integration.
- Support for [prefetched relationships](#) to minimize the number of queries.
- ODMG compliant API, [a Tutorial](#), and TestCases are included.
- JDO 1.0.1 compliant API (based on jdori, native implementation in progress), [a Tutorial](#), and TestCases are included.
- The [Lockmanagement](#) supporting four pessimistic Transaction Isolation Levels (uncommitted or "dirty" reads, committed reads, repeatable reads, serializable transactions) - distributed locking is possible.
- [Optimistic locking](#) support. Users may declare `int` or `long` fields as version attributes or `java.sql.Timestamp` fields as timestamp attributes.
- Support for persistent object caching. Different caching strategies and [distributed caches](#).
- Comes along with fully functional demo applications running against HSQLDB.
- Provides [Commons-Logging](#) and [Log4J](#) logging facilities.
- 100%: pure Java, Open Source, Apache License

**Note:**

- OQL is currently not fully implemented (Aggregations and Method Invocations)  
 - ODMG implicit locking is partly implemented but does currently not maintain transaction isolation properly. To achieve safe transaction isolation client applications must use explicit lock acquisition.

## 1.4. OJB - API Status

Status reports for the different APIs in Apache OJB.

Any known issues, including for the parts of OJB that are *stable*, can be found in [release-notes.txt](#).

- **PB API**

The PB (Persistence Broker) API implementation is **stable**.

- **ODMG**

The ODMG API implementation is **stable**.

**Note:**

OQL is currently not fully implemented (Aggregations and Method Invocations). Workaround is to use the *PB-API* [within the ODMG implementation](#).

- **JDO**

By providing a plugin to the SUN JDO Reference Implementation we provide a complete JDO 1.0.1 *prototype* O/R mapping tool.

Integration in managed environments (in particular JTA integration) is not supported for the JDO 1.0 prototype.

A complete Apache licensed JDO 2.0 implementation is scheduled for OJB 2.0.

- **OTM**

The OTM (Object Transaction Manager) API implementation is in *early beta*.

**Note:**

The active development on the OTM API implementation is currently stopped due to lack of developer resources in that area. The future of the OTM layer will be discussed on the [OJB developer mailing list](#).

- **SODA**

The legacy S.O.D.A. (Simple Object Database Access) API should be considered *deprecated*

and might be removed in a future release.

For more information about S.O.D.A. Query API, see the project's [SourceForge Website](#).

## 1.5. OJB - References and Testimonials

### 1.5.1. References and Testimonials

#### 1.5.1.1. projects using OJB

##### **Jakarta JetSpeed**

Jetspeed is an Open Source implementation of an Enterprise Information Portal, using Java and XML.

OJB will be the default persistence model within Jetspeed 2.

##### **BIT**

The swiss federal office for information technology and telecommunications (BIT) uses OJB 1.0.1 as data access layer in their framework for webbased applications.

The BIT extended OJB with a complex history-mechanism by simply replacing the JdbcAccessImpl their own class.

##### **The Tammi project**

Tammi is a JMX-based Java application development framework and run-time environment providing a service architecture for J2EE server side Internet applications that are accessible from any device that supports HTTP including mobile (wireless) handsets.

Future plans include integration of Apache OJB based persistence services to the framework.

##### **The Object Console project**

The Object Console is an open web based application meant for the administration of objects via the web. Any object that is persistable by the ObJectRelationalBridge (OJB) framework can be managed through this tool. In addition, this tool provides administration functionality for the ObJectRelationalBridge (OJB) framework itself.

Object Console uses Struts and OJB. It ships with full sourcecode and is thus a great source for learning Struts + OJB techniques.

##### **The IntAct project**

The IntAct project establishes a knowledgebase for protein-protein interaction data. It's hosted at EBI - European Bioinformatics Institute, Cambridge.

IntAct uses OJB as its persistence layer.

##### **Network for Earthquake Engineering Simulation**

The NEES program will provide an unprecedented infrastructure for research and education, consisting of networked and geographically distributed resources for experimentation, computation, model-based simulation, data management, and communication.

OJB is used as the O/R mapping layer.

##### **The OJB.NET project**

OJB.NET is an object-to-relational persistence tool for the .NET platform. It enables applications to transparently store and retrieve .NET objects using relational databases.

OJB.NET is a port ojb Apache OJB to the .NET platform

##### **The OpenEMed project**

OpenEMed is a set of distributed healthcare information service components built around the OMG distributed object specifications and the HL7 (and other) data standards and is written in Java for platform portability.

OpenEMed uses ODMG as its persistence API. OJB is used as ODMG compliant O/R tool.

### 1.5.1.2. user testimonials

"At the [BIT](#) some stress-test were performed simulating 3000 parallel users accessing tables containing more than 1.8 million rows per table. These test were run on Websphere 4.1 and DB2 on IBM z/OS (Host). The PB-API of OJB 1.RC1 was used without problems. The ODMG-API of this release then had too many bugs (deadlocks, parallel threads, etc.)."

"We're using OJB in two production applications at the Northwest Alliance for Computational Science and Engineering (NACSE). One is a data mining toolset, and the other is a massive National Science Foundation project that involves huge amounts of data, and about 20 or 25 universities and research groups like mine.

In fact, I've begun making OJB sort of a de-facto standard for NACSE java/database development. I've thrown out EJB's for the most part and I've tried JDO from Castor, but I'm sticking with OJB. Maybe we'll reconsider JDO when the OJB implementation is more complete."

"We are planning a November 2003 production deployment with OJB and WE LOVE IT!! We have been in development on a very data-centric application in the power industry for about 5 months now and OJB has undoubtedly saved us countless hours of development time. We have received benefits in the following areas:

-> Easily adapts to any data model that we've thrown at it. No problems mapping tables with compound keys, tables mapping polymorphic relationships, identity columns, etc.

-> Seemlesly switches between target DB platforms. We develop and unit test on our local workstations with HSQLDB and PostgreSQL, and deploy to DB2 using the Type 4 JDBC driver from IBM. Works great!

-> Makes querying a breeze with the PersistenceBroker API

Overall we have found OJB to be very stable (and we've really tested it out quite a bit). The only issues we've got outstanding at the moment is support for connections to multiple databases, but I've noticed in CVS that the OJB guys are already fixing this for OJB 0.9.9."

"We've been using it in "production" for a long time now, from about version 0.9.4, I believe. It has been very robust. We don't use all of its features. We've only see to failures of our persistent store in about 9 months, and I'm not sure they were due to OJB."

"So yes, we have made a quite useful mediumsized production website based on OJB (with JBoss, Jakarta Jetspeed, Jakarta Turbine and Jakarta Jelly, three Tomcats, OpenSymphony OSCache and for the database MSSQL server, all running on Win2000.) It is attracting between 600 and 9000 (peak) users a day, and runs smoothly for extended periods of time. And no, I can not actually show you the wonders of the editorial interface of the content management system, because it is hidden behind a firewall.

I feel OJB is quite useful in production, but you certainly have to know what you are doing and what you are trying to achieve with it. And there have been some tricky aspects, but these could be solved by simple workarounds and small hacks.

The main thing about OJB is that AFAIK it has an overall clean design, and it far beats making your own database abstraction layer and object/relational mapper. We certainly do not use all of it, only the Persistence Broker parts, so there was less to learn. We love the virtual proxy and collection proxy concepts, the criteria objects for building queries, and the nice little hidden features that you find when you start to learn the system."

"My Company is building medium to large scale, mission critical applications (100 - 5.000 concurrent users) for our customers. Our largest customer is KarstadtQuelle, Europes largest retail company. The next big system that will go in production (in June) is the new logistics system for the stationary logistics of Karstadt.

Of course we are using OJB in those Systems! We have several OJB based systems now in

production for over a year. We never had any OJB related problems in production. Most problems we faced during development were related to the learning curve developers had to face who were new to O/R mapping."

"I've also worked with OJB on high-load situations in J2EE environments. We're using JRun and/or Orion with OJB in a clustered/distributed environment. This is a National Science Foundation project called the Network for Earthquake Engineering Simulation (NEES).

The only major problem that we ran into was the cache. JCS just isn't good, and hasn't seemed to get much better over the last year. We ended up plugging in Tangosol's Coherence Clustered Cache into the system. We can also do write-behinds, and buffered data caching that is queued for transaction. That's important to us because we're dealing with very expensive scientific data that can't get lost if a db goes down. Some of these Tsunami experiments can get pretty expensive. Otherwise, we use mostly the PersistenceBroker, and a little of the ODMG. Performance seems better on PB, but less functional. It's not really that much of a problem anyway, because we can cheaply and quickly add app-servers to the cluster."

## 1.6. OJB - Mailing Lists

### 1.6.1. Mailing Lists

These are the mailing lists that have been established for this project. For each list, there is a subscribe, unsubscribe, and an archive link.

The user and dev list are subscriber only lists, this means you have to subscribe before you can post to the list.

List Name	Subscribe	Unsubscribe	Archive
OJB User List	<a href="#">Subscribe</a>	<a href="#">Unsubscribe</a>	<a href="#">Archive</a>
OJB Developer List	<a href="#">Subscribe</a>	<a href="#">Unsubscribe</a>	<a href="#">Archive</a>

## 1.7. OJB - Mailing Lists Archives

### 1.7.1. Mailing Lists Archives

Follow the links below to browse through or search in the mailing list archives, through a number of different providers and interfaces.

Archive Provider	OJB User List	OJB Developer List	Searchable?	Remarks
The Mail Archive	<a href="#">ojb-user</a>	<a href="#">ojb-dev</a>	yes	--
GMANE	<a href="#">gmane.comp.jakart</a>	<a href="#">gmane.comp.jakart</a>	yes	Latest 600 postings available via web access. Unlimited access through NNTP (news reader)
Apache	<a href="#">ojb-user@db.apach</a>	<a href="#">ojb-dev@db.apach</a>	yes	--
MARC	<a href="#">ojb-user</a>	<a href="#">ojb-dev</a>	yes	--

## 1.8. OJB - Links and further readings

### 1.8.1. Summary

This page contains interesting Links and recommended readings that will help to learn more about OJB concepts, related projects, didactic material, research reports, integration in frameworks etc.

## 1.8.2. OJB - Tutorials and Howto

- [The famous Beer4All Struts/OJB tutorial by Chuck Cavaness](#)
- [A presentation on OJB held at the Atlanta Java Users Group by Chuck Cavaness](#)
- [An extensive tutorial on OJB by John Carnell](#)
- [Roberto Ghizzioli's tutorial on Struts, OJB, and nested tags](#)
- [An introductory tutorial on the O'Reilly site.](#)

### 1.8.2.1. OJB and Spring

- [Spring, OJB, and Struts, getting started](#) (June 23, 2005)
- [Spring, OJB, and Struts. Version 2 with DAO](#) (June 23, 2005)
- [Spring, OJB, and Struts. Version 3: JNDI Datasources and Caching](#) (June 27, 2005)
- [Springframework and OJB](#) (Nov 21, 2004)

## 1.8.3. Design

OJB is based on a variety of conceptual sources. In this section I'll give a summary about the most prominent influences.

1. **Craig Larmans Applying UML and Patterns**
2. **The Siemens Guys "Pattern-Oriented Software Architecture"**
3. **Scott Amblers classic papers on O/R mapping**
4. **The "Crossing Chasms" paper from Brown et. al.**
5. **The GOF Design Patterns**

(sorted by relevance)

1. The most important input came from [Applying UML and Patterns](#). It contains a chapter describing the design of a PersistenceBroker based approach persistence layer. His presentation contains a lot of other good ideas (e.g. usage of Proxies, caching etc.) I implemented a lot of his things 1:1. This book is a must have for all OJB developers !

2. Larman does not cover the dynamic metadata concept. He mentions that such a thing would be possible, but does not go into details. As I had been a fan of MetaLevel architectures for quite a while I wanted to have such a thing in OJB too !!!



[mop-gif](#)

I took the concepts from the book [Pattern-Oriented Software Architecture](#). They have a chapter on the **Reflection** pattern (aka [Open Implementation, Meta-Level Architecture](#)).

They even provide an example how to apply this pattern to a persistence layer.

There is another Architectural pattern from this book that I am using: The **Microkernel** pattern. My idea was to have a kernel (the PersistenceBroker) that does all the hard work (O/R mapping, JDBC access, etc.)

High Level object transaction frameworks like a ODMG or JDO implementations are clients to the PersistenceBroker kernel in this concept!

3. I read [Scott Amblers](#) papers before starting OJB. Sure! There are several things in OJB that are from his classic [The design of a robust persistence layer](#) and from his [Mapping Objects To Relational Databases](#). Most prominent: The **PersistenceBroker** concept.

I incorporated the Query API from the OpenSource project [COBRA](#) that applies Amblers PersistentCriteria concept.

**Reading Amblers paper on these topics is a must.**

But IMO these are the only aspects of Amblers presentation that map directly to OJB. Here are the concepts that differ:

- Amblers concept relies on a persistent base class.
- caching is not covered by his design
- his concept of OID does not fit for legacy databases with compound PKs.
- The OJB proxy concept is quite different (Ambler has proxy functionality in his PersistentObject base class.)
- OJB does not use Insert- and UpdateCriteria
- OJB uses a different mapping approach (A full metadata layer)

4. For several detail questions (like mapping inheritance hierarchies) I consulted [crossing chasms](#). This is also a very good source for all O/R implementors.

5. For all the "small things" I'm using the [common GOF patterns](#) like Factory, Observer, Singleton, Proxy, Adaptor, State, Command, etc.

[Here is a thesis describing concepts very similar to OJB.](#)

As I read this paper I saw a lot of thing inspired by OJB. It's giving a nice introduction into the PersistenceBroker pattern and related topics.

[The PARC software design area](#) pioneering in Metalevel computation, aspect oriented programming etc.

#### 1.8.4. Further readings on O/R mapping

- [ObjectArchitects O/R pattern page](#)
- [JavaSkyLine page on database integration](#)
- [Barry and Associates page on O/R mapping](#)
- [Portland Pattern Repository page on O/R](#)
- [Martin Fowlers book "Pattern of Enterprise Application Architecture" covers many O/R patterns that can be found in OJB. Here you will find an online catalog of these patterns.](#)

#### 1.8.5. Patterns

- [The Hillside Pattern page](#)
- [The Portland Pattern Repository](#)

#### 1.8.6. Books covering OJB

- [The O'Reilly book on Struts programming by Chuck Cavaness has a whole chapter about how to build an applications model layers based on OJB. A must reading for everyone intending to use Struts and OJB.](#) All source code from the book can be found here: [Struts Programming sources](#).
- [There's also a WROX book on Struts + OJB](#) All source code from the book can be found here: [Professional Struts and OJB sources](#).
- [Enterprise Java Development on a Budget](#)

## 2. Download

## 3. Development

### 3.1. Coding Standards

This document describes a list of coding conventions that are required for code submissions to the project. By default, the coding conventions for most Open Source Projects should follow the existing coding conventions in the code that you are working on. For example, if the bracket is on the same line as the if statement, then you should write all your code to have that convention.

**If you commit code that does not follow these conventions, you are responsible for also fixing your own code.**

Below is a list of coding conventions that are specific to OJB, everything else not specifically mentioned here should follow the official [Sun Java Coding Conventions](#).

1. Brackets should begin and end on a new line and should exist even for one line statements.

Examples:

```

if (foo)
{
    // code here
}

try
{
    // code here
}
catch (Exception bar)
{
    // code here
}
finally
{
    // code here
}

while (true)
{
    // code here
}

```

2. Though it's considered okay to include spaces inside parens, the preference is to not include them. Both of the following are okay:

```

if (foo)

or

if ( foo )

```

3. Use 4 space indent. **NO tabs**. Period. We understand that many developers like to use tabs, but the fact of the matter is that in a distributed development environment where diffs are sent to the mailing lists by both developers and the version control system (which sends commit log messages), the use tabs makes it impossible to preserve legibility.

In Emacs-speak, this translates to the following command:

```
(setq-default tab-width 4 indent-tabs-mode nil)
```

4. Unix linefeeds for all .java source code files. Other platform specific files should have the



platform specific linefeeds.

5. JavaDoc **MUST** exist on all methods. If your code modifications use an existing class/method/variable which lacks JavaDoc, it is required that you add it. This will improve the project as a whole.
6. The ASF license **MUST** be placed at the top of each and every file.
7. All .java files should have a @version tag with CVS Id keyword expansion, like the one below.

```
@version $Id: code-standards.xml 365208 2005-12-14 01:19:23Z arminw $
```

**Note:**

To add the keyword to a new file, either use \$Id: code-standards.xml 365208 2005-12-14 01:19:23Z arminw \$ or copy an existing expanded id-string from another file (all the parameters will be replaced by CVS). Just watch out not to type \$Id \$, since that extra space will signal to CVS that keyword expansion already took place.

8. Import statements must be fully qualified for clarity.

```
import java.util.ArrayList;
import java.util.Hashtable;

import org.apache.foo.Bar;
import org.apache.bar.Foo;
```

And not

```
import java.util.*;
import org.apache.foo.*;
import org.apache.bar.*;
```

Emacs/XEmacs users might appreciate the following in their .emacs file.

```
(defun apache-db-mode ()
  "The Java mode specialization for Apache DB projects."
  (if (not (assoc "apache-db" c-style-alist))
      ;; Define the Apache DB cc-mode style.
      (c-add-style "apache-db" '("java" (indent-tabs-mode . nil))))

  (c-set-style "apache-db")
  (c-set-offset 'substatement-open 0 nil)
  (setq mode-name "Apache DB")

  ;; Turn on syntax highlighting when X is running.
  (if (boundp 'window-system)
      (progn (setq font-lock-support-mode 'lazy-lock-mode)
             (font-lock-mode t))))

  ;; Activate Apache DB-mode for JDE.
  (if (fboundp 'jde-mode)
      (add-hook 'jde-mode-hook 'apache-db-mode)
      (add-hook 'java-mode-hook 'apache-db-mode)))
```

Thanks for your cooperation.

## 4. Index

### 4.1. Site Linkmap

#### 4.1.1. Table of Contents

- OJB \_\_\_\_\_ *site*
  - OJB \_\_\_\_\_ *ojb*
    - Home \_\_\_\_\_ *index*
    - News \_\_\_\_\_ *news*
    - Features \_\_\_\_\_ *features*
    - Status \_\_\_\_\_ *status*
    - Release Notes \_\_\_\_\_ *release-notes*
    - References \_\_\_\_\_ *references*
    - Wiki \_\_\_\_\_ *wiki*
    - Mailing Lists \_\_\_\_\_ *mail-lists*
    - List Archives \_\_\_\_\_ *mail-archives*
    - Issue Tracking \_\_\_\_\_ *bugs*
    - Links \_\_\_\_\_ *links*
  - Download \_\_\_\_\_ *download*
    - Binaries \_\_\_\_\_ *binaries*
    - Source \_\_\_\_\_ *source*
  - Development \_\_\_\_\_ *development*
    - OJB SVN \_\_\_\_\_ *todo*
    - Apache SVN \_\_\_\_\_ *todo*
    - License \_\_\_\_\_ *license*
    - Coding Standards \_\_\_\_\_ *coding*
  - Index \_\_\_\_\_ *all*
    - Sitemap \_\_\_\_\_ *linkmap*
    - Site as PDF \_\_\_\_\_ *whole\_site\_pdf*
  - Documentation \_\_\_\_\_ *documentation*
    - Summary \_\_\_\_\_ *index*
    - FAQ's \_\_\_\_\_ *faq*
    - JavaDocs \_\_\_\_\_ *javadoc*
    - Getting Started \_\_\_\_\_ *getting-started*
    - Tutorials \_\_\_\_\_ *tutorials*
      - Summary \_\_\_\_\_ *summary*
      - Mapping Tutorial \_\_\_\_\_ *mapping-tutorial*
      - The PB API \_\_\_\_\_ *pb-tutorial*
      - The ODMG API \_\_\_\_\_ *odmg-tutorial*
      - The JDO API \_\_\_\_\_ *jdo-tutorial*
      - The OTM API \_\_\_\_\_ *otm-tutorial*

- Reference Guides \_\_\_\_\_ *guides*
  - Summary \_\_\_\_\_ *summary*
  - PB API guide \_\_\_\_\_ *pb-guide*
  - ODMG API guide \_\_\_\_\_ *odmg-guide*
  - Platform settings \_\_\_\_\_ *platform*
  - OJB propertyfile \_\_\_\_\_ *ojb-properties*
  - JDBC Types \_\_\_\_\_ *jdbc-types*
  - Repository file \_\_\_\_\_ *repository*
  - Basic mapping \_\_\_\_\_ *basic-technique*
  - Advanced mapping \_\_\_\_\_ *advanced-technique*
  - OJB queries \_\_\_\_\_ *query*
  - Metadata handling \_\_\_\_\_ *metadata*
  - Deployment \_\_\_\_\_ *deployment*
  - Connection handling \_\_\_\_\_ *connection*
  - Caching \_\_\_\_\_ *object-cache*
  - Sequence manager \_\_\_\_\_ *sequence-manager*
  - Logging \_\_\_\_\_ *logging*
  - Locking \_\_\_\_\_ *lock-manager*
  - XDoclet module \_\_\_\_\_ *xdoclet-module*
  - Performance \_\_\_\_\_ *performance*
- Howto's \_\_\_\_\_ *howto*
  - Summary \_\_\_\_\_ *summary*
  - Build large Metadata Mappings \_\_\_\_\_ *large-metadata*
  - Anonymous Keys \_\_\_\_\_ *anonymous-keys*
  - Using database sequences \_\_\_\_\_ *db-sequences*
  - Use Oracle LOB's \_\_\_\_\_ *use-lobs*
  - Work in clustered environment \_\_\_\_\_ *clustering*
  - Work with Stored Procedures \_\_\_\_\_ *stored-procedures*
  - Using Spring with OJB \_\_\_\_\_ *use-spring*
- Testing \_\_\_\_\_ *testing*
  - Summary \_\_\_\_\_ *summary*
  - Test Suite \_\_\_\_\_ *test-suite*
  - Write Tests \_\_\_\_\_ *test-write*

## 5. Documentation

### 5.1. Documentation - Introduction

This section contains the Apache OJB documentation.

Some additional information exists on our [Wiki](#).

If you are new to OJB, we recommend that you start with reading the [Getting Started](#) section and the [FAQ](#).

There are tools for building the metadata mapping files used by OJB. Information about them can be [found here](#).

- [Tutorials](#)  
Tutorials for the API's supported by OJB.
- [Reference Guides](#)  
OJB reference guides.
- [Howto's](#)  
Practical hands-on HOWTO documents provided by both OJB users and committers.
- [Testing](#)  
Info about OJB's quality assurance and test writing.

### 5.2. Frequently Asked Questions

#### 5.2.1. Questions

##### 5.2.1.1. 1. General

###### 1.1. Why OJB? Why do we need another O/R mapping tool?

here are some outstanding OJB features:

- It's fully ODMG 3.0 compliant
- It will have a full JDO implementation
- It's highly scalable (Loadbalanced Multiserver scenario)
- It provides multiple APIs:
  - The full fledged ODMG-API,
  - The JDO API (planned)
  - and the PersistenceBroker API. This API provides a O/R persistence kernel which can be used to build higher level APIs (like the ODMG and JDO Implementations)
- It's able to handle multiple RDBMS simultaneously.
- it has a slick MetaLevel Architecture: By changing the MetaData at runtime you can change the O/R mapping behaviour. (E.G. turning on/off usage of Proxies.)
- It has a simple CacheMechanisms that is fully garbage collectable by usage of weak references.
- It has a simple and clean pattern based design.
- It uses a configurable plugin concept. This allows to replace components (e.g. the ObjectCache) by user defined Replacements.
- It has a modular architecture (you can quite easily reuse some components in your own applications if you don't want to use the whole thing:
  - The PersistenceBroker (e.g. to build your own PersistenceManager)
  - The Query Interface as an abstract query syntax
  - The OQL Parser
  - The MetaData Layer

- The JDBC Accesslayer
- It has a very sharp focus: It's concerned with O/R mapping and nothing else.

Before making OJB an OpenSource project I had a look around at the emerging OpenSource O/R scene and was asking myself if there is really a need for yet another O/R tool. I came to the conclusion that there was a need for OJB because:

- There was no ODMG/JDO compliant opensource tool available
- There was no scalable opensource O/R tool available
- there was no tool available with the idea of a PersistenceBroker Kernel that could be easily extended
- The tools available had no dynamic MetaData architectures.
- The tools available were not as clearly designed as I hoped, thus extending one of them would have been very difficult.

## 1.2. How is OJB related to ODMG and JDO?

ODMG is a standard API for Object Persistence specified by the ODMG consortium ([www.odmg.org](http://www.odmg.org)). JDO is Sun's API specification for Object Persistence. ODMG may well be regarded as a Precursor to JDO. In fact JDO incorporates many ideas from ODMG and several people who have been involved in the ODMG spec are now in the JDO team.

I assume JDO will have tremendous influence on OODBMS-, RDBMS-, J2EE-server and O/R-tool-vendors to provide compliant products.

OJB wants to provide first class support for JDO and ODMG APIs.

OJB currently contains of four main layers, each with its own API:

1. A low-level **PersistenceBroker** API which serves as the OJB persistence kernel. The PersistenceBroker also provides a scalable multi-server architecture that allows to use it in heavy-duty app-server scenarios.  
This API can also be used directly by applications that don't need full fledged object level transactions (see [PB tutorial](#) for details).
2. An Object Transaction Manager (OTM) layer that contains all features that JDO and ODMG have in common as Object level transactions, lock-management, instance lifecycle etc. (See [OTM tutorial](#) for details.) The OTM is work in progress.
3. A full featured **ODMG 3.0** compliant API. (See [ODMG tutorial](#) for an introduction.)  
Currently this API is implemented on top the PersistenceBroker. Once the OTM layer is finished ODMG will be implemented on top of OTM.
4. A **JDO** compliant API. This is work in progress. (See [JDO tutorial](#) for an introduction.)  
Currently this API is implemented on top the PersistenceBroker. Once the OTM layer is finished JDO will be implemented on top of OTM.

The following graphics shows the layering of these APIs. Please note that the layers coloured in yellow are not yet implemented.

OJB Layer

## 1.3. What are the OJB design principals?

OJB has a "pattern driven" design. [Please refer to this document for more details](#)

## 1.4. Where can I learn more about Object/Relational mapping in general?

[We have a link list pointing to further readings.](#)

## 1.5. How OJB performance compares to native JDBC programming?

See page [Performance](#).

#### 1.6. How OJB performance compares to other O/R mapping tools?

See page [Performance](#).

#### 1.7. Is OJB ready for production environments?

Depends on your production environment. If you want to program an aeroplane autopilot system you should not use Java at all. (according to the official disclaimer).

But I assume we are talking about enterprise business applications, aren't we? And for such applications it's a clear **yes**. OJB is used in production application since version 0.5. We have about 6.000 downloads each month (and growing) and a large user base using it in a wide spectrum of production scenarios.

We provide a regression test suite for Quality Assurance. You can use this testsuite to check if OJB works smoothly in your target environment. (see [supported platforms documentation](#))

We also provide a performance testsuite that compares OJB performance against native JDBC. This test will give you an impression of the performance impact OJB will have in your target environment. (see [Performance testsuite documentation](#))

OJB is also the persistence layer of choice in several books on programming J2EE based enterprise business systems. (see [our links and references section](#))

Reference projects and user testimonials are listed [here](#).

#### 1.8. Does OJB supports caching?

Short answer is yes. To get a detailed answer, please read the [caching guide](#).

### 5.2.1.2. 2. Getting Started

#### 2.1. Help! I'm having problems installing and using OJB!

Please read the [Getting Started](#) document. OJB is a powerful and complex system - installing and configuring OJB is not a trivial task. Be sure to follow **all the steps** mentioned in that document - don't skip any steps when first installing OJB on your systems.

If you are having problems running OJB against your target database, read the respective [platform documentation](#). Before you try to deploy OJB to your environment, read the [deployment guide](#).

#### 2.2. Help! I still have serious problems installing OJB!

The following answer is quoted from the OJB user-list. It is from a reply to a user who had serious problems getting started with OJB.

I would say it was stupid not to understand OJB. How can you know what another programmer wrote. I've been a Java programmer for quite some time and I could show you stuff I wrote that I know you wouldn't understand. I'll just break it down the best I can on what, where and why.

OJB is a data persistence layer for Java. I'll just use an example of how I use it. I have an RDMS. I would like to save Java object states to this database and I would like to be able to search this information as well. If you serialize objects it's hard to search and if you use SQL it won't work with any different database. Plus it's a mess having to work with all that SQL in your code. And by using SQL you don't get to work with just Java objects. But, with OJB your separated from having to work outside the object world and unlike serialization you can preform SQL like searches on your data. Also, there's things like caching and connection pooling in OJB that help with performance. After setting up OJB you will use either PB-API or ODMG or JDO to access your

information in a object centric manner. PB API is a non-standard O/R mapping API with many features and great flexibility. All top-level API's like ODMG or JDO build on top of the PB-api. ODMG is a standard for the api for accessing your data. That means you can use any ODMG compliant api if you don't want to use OJB. The JDO part is like ODMG except it's the SUN JDO standard. I use ODMG because the JDO interface is not ready yet.

OJB is easy to use. I'll just break it down into two sides. There's the side your writing your code for your application and there's the side that you configure to make OJB connect to your database. Starting with your application side, all that is needed is to use the interface you wish. I use ODMG because JDO is not complete yet. Here's a [link to the ODMG part](#) with some code for examples. That's all you need on the application side. Next there's the configuration side. This is the one your fighting with. Here you need to setup the core tables for OJB and you will define the classes you wish to store in your database.

First thing to do is to build the cvs's with the default database HSQL, because you know it will work. If you get past this point you should have a working OJB compiled. Now if your using JDK 1.4 you will need to set in build.properties JDBC=+JDBC30 and do a *ant preprocess* first. Next you will do a *ant junit* and this will build OJB and test everything for you. If you get a build successful then your in business. Then you will want to run *ant jar* to create the OJB jar to put in your /lib. You will need a couple other jars in you /lib directory to make it all work. See this page for those. <http://jakarta.apache.org/ojb/deployment.html>

Next you will need some xml and configuration files in your class path for OJB. You will find those files under `{$OJB_base_dir}/target/test/ojb`. All the repository.xml's and OJB.properties for sure. With all these files in place with your application you should be ready to use OJB and start writing your application.

Finally you will want to setup your connection to your database and define your classes you will be storing in your database. In the repository.xml file you can configure your JDBC parameters so OJB can connect to your database. You will also need your JDBC jar somewhere in your class path. Then you will define your classes in the repository\_user.xml file. Look here for examples. <http://jakarta.apache.org/ojb/tutorial1.html> Note you will want to comment out the junit part in repository.xml because it's just for testing.

The final thing to do is to make sure the OJB core tables are in your database. Look on this page for [the core tables](#). These core tables are used by OJB to store internal data while it's running. It needs these. Then there's the tables you define. The ones you mapped in the repository\_user.xml file.

Sorry if any of this is off. OJB is growing so fast that it's hard to keep up with all changes. The order I gave the steps in is just how I would think it's understood better. You can go in any order you want. The steps I've shown are mostly for deployment. Hope this helps you understand OJB a little better. I'm not sure if this is what your wanting or not.

### 2.3. OJB does not start?

If you carefully attended the [installing hints](#) there may be something wrong with your metadata mapping defined in the [repository file](#) or one the included sub files.

- Are you included all configuration files in classpath?
- On update to a new release, make sure you replaced all configuration files
- Check your metadata mapping - typos,... ?

If something going wrong while OJB read the metadata files you can enable *debug* log level for `org.apache.ojb.broker.metadata.RepositoryXmlHandler` and `org.apache.ojb.broker.metadata.ConnectionDescriptorXmlHandler` to get more detailed information.

**Note:**

If OJB default logging was used, change entries for these classes in [OJB.properties](#) file (this may change in future).

**2.4. Does OJB support my RDBMS?**

[please refer to this document.](#)

**2.5. What are the OJB internal tables for?**

[Please refer to this document.](#)

**2.6. What does the exception Could not borrow connection from pool mean?**

There can be several reasons

**2.7. Any tools help to generate the metadata files?**

[Please refer to this document.](#)

**5.2.1.3. 3. OJB APIs****3.1. What are the differences between the different OJB APIs? Which one should I use in my applications?**

The PersistenceBroker (PB) provides a minimal API for transparent persistence:

- O/R mapping
- Retrieval of objects with a simple query interface from RDBMS
- storing (insert, update) of objects to RDBMS
- deleting of objects from RDBMS

This is all you need for simple applications as in tutorial1.

The OJB ODMG implementation uses the PB as its persistence kernel. But it provides much more functionality to the application developer. ODMG is a full fledged API for Object Persistence, including:

- OQL Query interface
- real Object Transactions
- A Locking Mechanism for management of concurrent threads (apps) accessing the same objects
- predefined persistent capable Collections and Hashtables

Some examples explaining the implications of these functional differences:

1. Say you use the PB to query an object O that has a collection attribute col with five elements a,b,c,d,e. Next you delete Objects d and e from col and store O again with PersistenceBroker.store(O);  
PB will store the remaining objects a,b,c. But it will not delete d and e ! If you then requery object O it will again contain a,b,c,d,e !!!  
The PB keeps no transactional state of the persistent Objects, thus it does not know that d and e have to be deleted. (as a side note: deletion of d and e could also be an error, as there might be references to them from other objects !!!)  
Using ODMG for the above scenario will eliminate all trouble: Objects are registered to a transaction so that on commit of the transaction it knows that d and e do not longer belong to the collection. the ODMG collection will not delete the objects d and e but only the REFERENCES from the collection to those objects!
2. Say you have two threads (applications) that try to access and modify the same object O. The



PB has no means to check whether objects are used by concurrent threads. Thus it has no locking facilities. You can get all kind of trouble by this situation. The ODMG implementation has a Lockmanager that is capable of synchronizing concurrent threads. You can even use four transaction isolation levels:

read-uncommitted, read-committed, repeatable-read, serializable.

In my eyes the PB is a persistence kernel that can be used to build high-level PersistenceManagers like an ODMG or JDO implementation. It can also be used to write simple applications, but you have to do all management things (locking, tracking objects state, object transactions) on your own.

### 3.2. I don't like OQL, can I use the PersistenceBroker Queries within ODMG?

Please refer to the [ODMG-guide](#).

### 3.3. The OJB JDO implementation is not finished, how can I start using OJB?

I recommend to not use JDO now, but to use the existing ODMG api for the time being.

Migrating to JDO later will be smooth if you follow the following steps. I recommend to first divide your model layer into Activity- (or Process-) classes and Entity classes.

Entity classes represent classes that must be made persistent at some point in time, say a "Customer" or a "Order" object. These persistent classes and the respective O/R mapping in repository.xml will remain unchanged.

Activities are classes that perform business tasks and work upon entities, e.g. "edit a Customer entry", "enter a new Order"... They implement (parts of) use cases.

Activities are driving transactions against the persistent storage.

I recommend to have a Transaction interface that your Activities can use. This Transaction interface can be implemented by ODMG or by JDO Transactions (which are quite similar). The implementation should be made configurable to allow to switch from ODMG to JDO later.

The most obvious difference between ODMG and JDO are the query languages: ODMG uses OQL, JDO define JDOQL. As an OO developer you won't like both of them. I recommend to use the ojb Query objects that allow an abstract syntax representation of queries. It is possible to use these queries within ODMG transactions and it will also be possible to use them within JDO Transactions. (this is contained in the FAQ too).

Using your own Transaction interface in conjunction with the OJB query api will provide a simple but powerful abstraction of the underlying persistence layer.

We are using this concept to provide an abstract layer above OJB-ODMG, TopLink and LDAP servers in my company. Making it work with OJB-JDO will be easy!

## 5.2.1.4. 4. Howto

### 4.1. How to use OJB with my RDBMS?

[please refer to this document](#).

### 4.2. How to use OJB in an web app?

If you follow these rules, then OJB works fine in web apps:

- Don't put OJB's jars into one of the servers directories but rather put them into the WEB-INF/lib folder of your web app.
- OJB searches for its configuration files (OJB.properties, repository.xml) in the

classpath. Therefore, it is easiest if you put them in the `WEB-INF/classes` folder which is automatically in the classpath of the web app

- Don't hold onto the `PersistenceBroker` instances, rather get one whenever you want to do something, and close it once you're done.

See [deployment doc](#) for more information.

#### 4.3. What are the best settings for maximal performance?

See [performance section](#).

#### 4.4. How to page and sort?

Sorting can be configured by

```
org.apache.ojb.broker.query.Criteria::orderBy(column_name).
```

There is no paging support in OJB. OJB is concerned with Object/Relational mapping and not with application specific presentation details like presenting a scrollable page of items.

OJB returns query results as Collections or Iterators.

You can easily implement your partial display of result data by using an Iterator as returned by `ojb.broker.PersistenceBroker::getIteratorByQuery(...)`.

#### 4.5. What about performance and memory usage if thousands of objects matching a query are returned as a Collection?

You can do two things to enhance performance if you have to process queries that produce thousands of result objects:

1. Use `getIteratorByQuery()` rather than `getCollectionByQuery()`. The returned Iterator is lazy and does not materialize Objects in advance. Objects are only materialized if you call the Iterators `next()` method. Thus you have total control about when and how many Objects get materialized! Please see [here for proper handling](#).
2. You can define Proxy Objects as placeholder for your persistent business objects. Proxies are lightweight objects that contain only primary key information. Thus their materialization is not as expensive as a full object materialization. In your case this would result in a collection containing 1000 lightweight proxies. Materialization of the full objects does only occur if the objects are accessed directly. Thus you can build similar lazy paging as with the Iterator. You will find examples in the OJB test suite (src-distribution only: `[db-ojb]/src/test`). More info about [Proxy object here](#).

The Performance of 1. will be better than 2. This approach will also work for VERY large resultsets, as there are no references to result objects that would prevent their garbage collectability.

#### 4.6. When is it helpful to use Proxy Classes?

Proxy classes can be used for "lazy loading" aka "lazy materialization". Using Proxy classes can help you in reducing unnecessary db lookups. Example:

Say you load a `ProductGroup` object from the db which contains a collection of 15 `Article` objects.

Without proxies all 15 `Article` objects are immediately loaded from the db, even if you are not interested in them but just want to lookup the description-attribute of the `ProductGroup` object.

With a proxy class, the collection is filled with 15 proxy objects, that implement the same interface as the "real objects" but contain only an OID and a void reference.

Once you access such a proxy object it loads its "real subject" by OID and delegates the method

call to it.

have a look at section [proxy usage](#) of page [basic technique](#).

#### 4.7. How can I convert data between RDBMS and OJB?

For Example I have a DB column of type INTEGER but a class attribute of type boolean. How can I provide an automatic mapping with OJB?

OJB provides a concept of ConversionStrategies that can be used for such conversion tasks. [Have a look at the respective document](#).

#### 4.8. How can I trace and/or profile SQL statements executed by OJB?

OJB ships with out of the box support for [P6Spy](#). P6Spy is a JDBC proxy which delegates all JDBC calls to the real JDBC driver and traces all calls to a log file.

P6Spy is contained in the `p6spy.jar`, which you'll find in the `lib` folder of your OJB distribution. Add this to the classpath of your app (if you're using the [ojb-blank](#) project, then simply copy the jar into the `lib` folder of the project and if you're using Eclipse then also add it to the project build path).

Now the only other thing left is to configure OJB to use P6Spy, and P6Spy to use your database's driver. To achieve this, change the database driver in your `jdbc-connection-descriptor` (in your repository file) to

```
<jdbc-connection-descriptor
  ...
  driver="com.p6spy.engine.spy.P6SpyDriver"
  ...
/>
```

In `ojb-blank` this setting is changed in the `build.properties` instead.

Also copy the file `spy.properties` which can be found in the `src/test/org/apache/ojb` folder into your classpath (e.g. in the same place where your `OJB.properties` file is). In this file you'll find a line starting with `realdriver` where you should put the name of the jdbc driver of your database, e.g.

```
realdriver=org.hsqldb.jdbcDriver
```

Also, here you can influence to where P6Spy will output the SQL statements. The appender defines how the logging is performed, e.g. to the console or to a file. The `logfile` setting defines into which file the statements will be printed (when a file appender is used). For instance, these settings will write to a file `spy.log`:

```
logfile = spy.log
appender = com.p6spy.engine.logging.appender.FileLogger

# This would be logging to the console
#appender = com.p6spy.engine.logging.appender.StdoutLogger
```

That's all there is to it, no recompile or other change of your app is necessary. Btw, P6Spy also measures the time needed to execute each statement!

#### 4.9. How does OJB manage foreign keys?

Automatically! you just define 1:1, 1:n or m:n associations in the `repository_user.xml` file. OJB does the rest!

Please refer to [basic technique](#) and [xml-metadata repository](#) for details.

#### 4.10. How does OJB manage 'null' for primitive primary key?

Primitive values (int, long, ...) can't be null, so OJB interpret '0' as null for primitive PK/FK fields in persistent objects. Thus primitive PK fields of persistent objects should never be represented by a '0' value in DB and never used as a [sequence key](#) value.

This is only true for primitive PK/FK fields (e.g. Integer(0) is allowed). All other fields have 'normal' behavior.

#### 4.11. How to lookup object by primary key?

Please see [PB tutorial section](#).

#### 4.12. Difference between getIteratorByQuery() and getCollectionByQuery()?

The first one returns an `org.apache.ojb.broker.OJBIterator` instance. The returned Iterator instance is lazy and does not materialize Objects in advance. Objects are only materialized from the underlying query result set if you call the Iterators `next()` method. If all objects materialized or the calling `org.apache.ojb.broker.PersistenceBroker` instance was closed or transaction demarcations ends the Iterator instance release all used resources (e.g. used Statement and ResultSet instances).

Method `getCollectionByQuery()` use an Iterator to materialize all objects first and then return the materialized objects within the `java.util.Collection` instance.

#### Note:

If method `getIteratorByQuery()` was used keep in mind that the used Iterator instance is only valid as long as the used `org.apache.ojb.broker.PersistenceBroker` instance ends transaction or be closed. So it is NOT possible to get an Iterator, close the PersistenceBroker and pass the Iterator instance to a servlet or client. In that case use `getCollectionByQuery()`.

#### 4.13. How can Collections of primitive typed elements be mapped?

The first thing to ask is: How are these primitive typed elements (Strings are also treated as primitive types here) stored in the database.

- 1) are they treated as ordinary domain objects and stored in a separate table?
- 2) are they serialized into a Varchar field?
- 3) are they stored as a comma separated varchar field?
- 4) is each element of the vector or array stored in a separate column? (this solution does only work for a fixed number of elements!)

Follow these steps for solution 3):

- a) simply define ordinary collection-descriptors as for every other collection of domain objects.
- b) use the `Object2ByteArrFieldConversion`. See [jdbc-types.html](#) for details on conversion strategies.
- c) use the `StringVector2VarcharFieldConversion`. See [jdbc-types.html](#) for details on conversion strategies.
- d) provide a field-descriptor for each element.

#### 4.14. How could class 'myClass' represent a collection of 'myClass' objects

OJB can handle such recursive associations without problems.

- add a collection attribute 'myClasses' to the class `myClass` this collection will hold the associated `myClass` objects.
- you have to decide wether this association is 1:n or m:n.  
for 1:n you just need an additional foreignkey attribute in the `MY_CLASS` table. Of course

you'll also need a matching attribute in the class `myClass`.

For a m:n association you'll have to define an intermediary table to hold the mapping entries.

- define a `collection-descriptor` tag in the `class-descriptor` of `myClass` in `repository.xml`. Follow the steps in [basic technique](#) on 1:n and m:n.

#### 4.15. How to lookup PersistenceBroker instances?

Please refer to [PB-guide](#).

#### 4.16. How to access ODMG?

Please refer to [ODMG-guide](#).

#### 4.17. Needed to put user/password of database connection in repository file?

There is no need to put user/password in the repository file (more exact in the `jdbc-connection-descriptor`). You can pass this information at runtime. See [Many different database user - How do they login?](#)

Only if you want to use convenience `PersistenceBroker` lookup method of `PersistenceBrokerFactory`, OJB needs all database connection information in the configuration files. More details see [repository file doc - section jdbc-connection-descriptor default-connection attribute](#)

See [lookup PB api](#).

See [lookup ODMG api](#).

```
PBKey pbKey = new PBKey(jcdAlias, user, passwd);
PersistenceBroker broker =
PersistenceBrokerFactory.createPersistenceBroker(pbKey);
// or using a convenience (when default-connection was set in
jdbc-connection-descriptor)
PersistenceBroker broker = PersistenceBrokerFactory.defaultPersistenceBroker();
```

#### 4.18. Many different database user - How do they login?

There are two ways to do that. Define for each user a `jdbc-connection-descriptor` (unattractive way, because we have to add each new user to repository file), or let OJB handle this for you.

For it define **one** `jdbc-connection-descriptor`, now you can use the same `jcdAlias` name with different User/Password. OJB **copy** the defined `jdbc-connection-descriptor` and replace the username and password with the given User/Password.

PersistenceBroker-api example:

```
PBKey user_1 = new PBKey(jcdAlias,username, passwd);
PersistenceBroker broker =
PersistenceBrokerFactory.createPersistenceBroker(user_1);
...
```

ODMG-api example:

```
Implementation odmg = OJB.getInstance();
Database db = odmg.newDatabase();
db.open("jcdAlias#username#passwd", Database.OPEN_READ_WRITE);
...
```

Keep in mind, when the `connection-pool` element enables connection pooling, every user get its separate pool. See [How does OJB handle connection pooling?](#)

#### 4.19. How do I use multiple databases within OJB?

Define for each database a `jdbc-connection-descriptor`, use the different `jcdAlias` names in the [repository file](#) to match the according database.

```
<jdbc-connection-descriptor
  jcd-alias="myFirstDb"
  ...
>
...
</jdbc-connection-descriptor>
<jdbc-connection-descriptor
  jcd-alias="mySecondDb"
  ...
>
...
</jdbc-connection-descriptor>
```

Specific notes related to the [PB-api here](#).

Specific notes related to the [ODMG-api here](#).

#### Note:

OJB does not provide distributed transactions by itself. To use distributed transactions, OJB have to be [integrated in an j2ee conform environment](#) (or made work with an JTA/JTS implementation).

#### 4.20. How does OJB handle connection pooling?

Please have a look in section [Connection Handling](#).

#### 4.21. Can I directly obtain a `java.sql.Connection` within OJB?

Please have a look in section [Connection Handling](#).

#### 4.22. Is it possible to perform my own `sql-queries` in OJB?

There are several ways in OJB to do that.

If you completely want to bypass the OJBquery-api see [Can I directly obtain a `java.sql.Connection` within OJB?](#)

A more elegant way is to use a `QueryBySQL` object:

```
String sql =
"SELECT A.Artikel_Nr FROM Artikel A, Kategorien PG"
+ " WHERE A.Kategorie_Nr = PG.Kategorie_Nr"
+ " AND PG.Kategorie_Nr = 2";
// get the QueryBySQL
Query q2 = QueryFactory.newQuery(Article.class, sql);

Iterator iter2 = broker.getIteratorByQuery(q2);
// or
Collection col2 = broker.getCollectionByQuery(q2);
```

#### 4.23. When does OJB open/close a connection?

Please see [Connection handling guide](#).

#### 4.24. Start OJB without a repository file?

See section [Metadata Handling](#).

#### 4.25. Connect to database at runtime?

See section [Metadata Handling](#).

#### 4.26. Hook into OJB - How to add Listener, callback interface?

See *Listener/Callback* section in [PB-Guide](#).

#### 4.27. Add new persistent objects metadata ( class-descriptor) at runtime?

See section [Metadata Handling](#).

#### 4.28. Global metadata changes at runtime?

Please see section [Metadata Handling](#).

#### 4.29. Per thread metadata changes at runtime?

Please see section [Metadata Handling](#).

#### 4.30. Is it possible to use OJB within EJB's?

Yes, see [deployment](#) instructions in the docs. Additional you can find some EJB example beans in package `org.apache.ojb.ejb` under `[jakarta-ojb]/src/ejb`.

#### 4.31. Can OJB handle ternary (or higher) associations?

Yes, that's possible. Here is an example. With a ternary relationship there are three (or more) entities 'related' to each other. An example would be Developer, Language and Project.

Each entity is mapped to one table ( DEVELOPER, LANGUAGE and PROJECT). To represent the combinations of these entities we need an additional bridge table ( PROJECTRELATIONSHIP)with three columns holding the foreign keys to the other three tables (just like an m:n association is represented by an intermediary table with 2 columns).

To handle this table with OJB we have to define a class that is mapped on it. This Relationship class can then be used to perform queries/updates as with any other persistent class. Here is the layout of this class:

```
public class ProjectRelationship {
    Integer developerId;
    Integer languageId;
    Integer projectId;

    Developer developer;
    Language lanuage;
    Project project;

    /** setters and getters not shown for brevity**/
}
```

Here is the respective extract from the repository :

```
<class-descriptor
  class="ProjectRelationship"
  table="PROJECTRELATIONSHIP"
>
  <field-descriptor
    name="developerId"
    column="DEVELOPER_ID"
    jdbc-type="INTEGER"
    primarykey="true"
  />
  <field-descriptor
```

```

        name="languageId"
        column="LANGUAGE_ID"
        jdbc-type="INTEGER"
        primarykey="true"
    />
    <field-descriptor
        name="projectId"
        column="PROJECT_ID"
        jdbc-type="INTEGER"
        primarykey="true"
    />
    <reference-descriptor
        name="developer"
        class-ref="Developer"
    >
        <foreignkey field-id-ref="developerId" />
    </reference-descriptor>
    <reference-descriptor
        name="language"
        class-ref="Language"
    >
        <foreignkey field-id-ref="languageId" />
    </reference-descriptor>
    <reference-descriptor
        name="project"
        class-ref="Project"
    >
        <foreignkey field-ref="projectId" />
    </reference-descriptor>
</class-descriptor>

```

Here is some sample code for storing a relationship :

```

Developer dev = ... ; // create or retrieve
Project proj = ... ; // create or retrieve
Language lang = ... ; // create or retrieve

ProjectRelationship rel = new ProjectRelationship();
rel.setDeveloper(dev);
rel.setLanguage(lang);
rel.setProject(proj);

broker.store(r);

```

In the next code sample we are looking up all Projects that Developer "Bob" has done in "Java".

```

Criteria criteria = new Criteria();
criteria.addEqualTo("developer.name", "Bob");
criteria.addEqualTo("language.name", "Java");

Query q = new QueryByCriteria(ProjectRelationship.class, criteria, true);
Iterator iter = Broker.getIteratorByQuery(q);

// now iterate over the collection and retrieve all projects:
while (iter.hasNext())
{
    ProjectRelationship rel = (ProjectRelationship) iter.next();
    System.out.println(rel.getProject().toString());
}

```

You could also have on the Project class-descriptor a collection-descriptor that returns all relationships associated with the Project. If it was call "projectRelationships" the following would give you all projects that have a relationship with "bob" and the language "java".

```

Criteria criteria = new Criteria();
criteria.addEqualTo("projectRelationships.developer.name", "bob");
criteria.addEqualTo("projectRelationships.language.name", "java");

Query q = new QueryByCriteria(Project.class, criteria, true);
Collection projects = Broker.getCollectionByQuery(q);

```



This is the layout of the Project class:

```
public class Project {
    Integer id;
    String name;
    Collection projectRelationships;

    /** setters and getters not shown for brevity**/
}
```

This is the class-descriptor of the Project class:

```
<class-descriptor
  class="Project"
  table="PROJECT"
>
  <field-descriptor
    name="id"
    column="ID"
    jdbc-type="INTEGER"
    primarykey="true"
  />
  <field-descriptor
    name="name"
    column="NAME"
    jdbc-type="VARCHAR"
  />
  <collection-descriptor
    name="projectRelationships"
    element-class-ref="ProjectRelationship"
  >
    <inverse-foreignkey field-ref="projectId" />
  </collection-descriptor>
</class-descriptor>
```

#### 4.32. How to map a list of Strings

You can not map a list of Strings with a collection descriptor. A collection descriptor can only be used if the element class is a persistent class too. But `element-class-ref="java.lang.String"` won't work, because it's no persistent entity class!

Follow these steps to provide a mapping for an attribute holding a list of Strings. Let's assume your persistent class has an attribute `listOfStrings` holding a list of Strings:

```
protected Collection listOfStrings;
```

The database table mapped to the persistent class has a column `LIST_OF_STRINGS` of type `VARCHAR` that is used to hold all strings.

```
<field-descriptor
  name="listOfStrings"
  column="LIST_OF_STRINGS"
  jdbc-type="VARCHAR"
  conversion=
"o.a.ojb.broker.accesslayer.conversions.StringVector2VarcharFieldConversion"
/>
```

#### 4.33. How to set up Optimistic Locking

Please see [locking section](#).

#### 4.34. How to use OJB in a cluster

Q: I'm running a web site in a load-balanced/cluster environment. Multiple servlet engines (different VMs/HTTP sessions), each running an OJB instance, against a single shared database. How should OJB be configured to get the concurrent servlet engines synchronized properly?

### transactional isolation and locking

If you are using the PersistenceBroker API use [optimistic locking](#) (OL) to let OJB handle write conflicts. To use OL define a `TIMESTAMP` or `INTEGER` column and the respective Java attribute for it. In the field-descriptor of this attribute set the attribute `locking="true"`.

If you are working with the ODMG API [distributed pessemistic locking](#) should be used, by setting the respective flag in `OJB.properties`.

### sequence numbers

Use a [SequenceManager](#) that is safe across multiple JVMs. The NextVal based SequenceManagers or any other SequenceManager based on database mechanisms will be fine.

### caching

You could use different [caching implementations](#)

1. Use the EmptyCacheImpl to avoid any dirty reads. (But: The EmptyCache cannot handle cyclic structures on load!)
2. Use the PerBrokerCache Implementation to avoid dirty reads.
3. Use the OSCache cache implementation as distributed object cache.

There is also a [complete howto document](#) available that covers these topics.

#### 4.35. How to turn of caching?

Declare an no-op implementation of the `ObjectCache` interface as cache. See [detailed description here](#).

#### 4.36. JDO - Why must my persisten class implement `javax.jdo.spi.PersistenceCapable`?

As specified by JDO all persistent classe must implement the interface `javax.jdo.spi.PersistenceCapable`. If a class does not implement this interface a JDO implementation does not know how to handle it.

On the other hand the JDO spec claims to provide transparent persistence. That is no persistence class is required to implement a specific interface or to be derived from a special base class.

Sounds like a contradiction? It is! The JDO spec resolves this contradiction by stating that a JDO implementation is responsible to add the methods required by

`javax.jdo.spi.PersistenceCapable` to the the user classes. This "injection" could be achieved by Pre- or Post-processing. The strategy most implementations use is called "bytecode-enhancement". This is a postprocesing step that adds the required methods to the .class files of the persistent user classes.

The JDO Reference implementation also uses bytecode-enhancement. In order to enhance the Product class to implement the `javax.jdo.spi.PersistenceCapable` interface use the ant target "enhance-jdori" before launching the tutorial5 application. This is documented in the first section of tutorial4.html.

## 5.3. Getting Started

This document will guide you through the very first steps of setting up a project with OJB. To make this easier, OJB comes with a blank project template called `obj-blank` which you're encouraged to use. You can download it [here](#).

For the purpose of this guide, we'll be showing you how to setup the project for a simple application that handles products and uses MySQL. This is continued later on in the next [tutorial parts](#).

### 5.3.1. Acquiring `obj-blank`

First off, OJB uses [Ant](#) to build, so please install it prior to using OJB. In addition, please make sure that the environment variables `ANT_HOME` and `JAVA_HOME` are correctly set to the top-level folders of your Ant distribution and your JDK installation, respectively.

Next download the latest [ojb-blank and OJB binary distributions](#). You can also start with the [source distribution](#) rather than the binary as the unit tests provide excellent sample code and you can build the ojb-blank project on your own with it.

The ojb-blank project contains all libraries necessary to get running. However, there may be additional libraries required when you venture deeper into OJB's APIs. See [here](#) for a list of additional libraries.

Most notably, you'll probably want to add the jdbc driver for your database unless you plan to use the embedded [Hsqldb database](#) for which the ojb-blank project is pre-configured (including all necessary jars).

### 5.3.2. Contents of ojb-blank

Copy the `ojb-blank.jar` file to your project directory and unpack it via the command

```
jar xvf ojb-blank.jar
```

This will unpack it into the `ojb-blank` directory under wherever you unpacked it from. You can move things out of that directory into your project directory, or, more simply, rename the `ojb-blank` directory to be whatever you want your project directory to be named.

After you unpacked the jar, you'll get the following directory layout:

```
\ojb-blank
  .classpath
  .project
  build.properties
  build.xml
  \lib
  \src
    \java
    \resources
    \schema
    \test
```

Here's a quick rundown on what the individual directories and files are:

#### **.classpath, .project**

An [Eclipse](#) project for your convenience. You can simply import it into Eclipse via **File -> Import... -> Existing Project into Workspace**.

#### **build.xml, build.properties**

The Ant build script and the build properties. These are described in more detail [below](#).

#### **lib**

Contains the libraries necessary to compile and run your project. If you want to use a different database than Hsqldb, then put the jars of your jdbc driver in here.

#### **src/java**

Put your java source code here.

#### **src/resources**

Contains the runtime configuration files for OJB. For more detail see [below](#).

#### **src/schema**

Here you will find a schema containing tables that are required by certain components of OJB such as clustered locking and OJB managed sequences. More information on these tables is available in the [platform documentation](#). The schema is in a database-independent format that can be used by [Torque](#) or [commons-sql](#) to create the database.

The ojb-blank project contains the runtime files of Torque 3.0.2, and provides a build target that can be invoked on your schema (see [below](#) for details). Therefore, this directory also contains the build script of Torque, but you won't need to invoke it directly.

### src/java

Place your unit tests in here.

#### 5.3.2.1. Sample project

For our sample project, we should rename the directory to something more fitting, like `productmanager`.

Also, since we're using MySQL, we put the MySQL jar of the [jdbc driver](#), which is called something like `mysql-connector-java-[version]-stable-bin.jar`, into the `lib` subdirectory.

The only other thing missing is the source code, but since that's what the other tutorials are dealing with, we will silently assume that it is already present in the `src/java` subdirectory.

If you don't want to write the code yourself, you can use the code from one of the [tutorials](#) which you can download [here](#).

#### Warning:

Note that if you do not intend to use JDO, then you should delete the files in the `ojb.apache.ojb.tutorial5`, otherwise you'll get compilation errors.

### 5.3.3. The build files

#### 5.3.3.1. Configuration via build.properties

The next step is to adapt the build files, especially the `build.properties` file to your environment. It basically contains two sets of information, the database settings and the build configuration. While you shouldn't have to change the latter, the database settings probably need to be adapted to suit your needs:

Property	Purpose
<code>jcdAlias</code>	The name of the connection. You should leave the default value, which is <code>default</code> .
<code>databaseName</code>	This is the name of the database, per default <code>ojb_blank</code> .
<code>databaseUser</code>	The user name for accessing the database (default: <code>sa</code> ). If you're using Torque to create the database, then this user also requires sufficient rights to create databases and tables.
<code>databasePassword</code>	Password for the user, per default empty.
<code>dbmsName</code>	The type of database, which is one of the following: <b>Db2, Firebird, Hsqldb, Informix, MaxDB, MsAccess, MsSQL, MySQL, Oracle</b> (pre-9i versions), <b>Oracle9i, WLOracle9i</b> (Oracle 9i or above used from WebSphere), <b>PostgreSQL, Sapdb, Sybase</b> (generic), <b>SybaseASA, SybaseASE</b> . Please note that this setting is case-sensitive. Per default, <a href="#">Hsqldb</a> is used, which is an

	embedded database. All files required for this database come with the ojb-blank project.
<code>jdbcRuntimeDriver</code>	The fully-qualified classname of the jdbc driver. For Hsqldb this is <code>org.hsqldb.jdbcDriver</code> .
<code>jdbcLevel</code>	The jdbc level that the driver conforms to. Please check the documentation of your jdbc driver for this value, though most jdbc drivers conform to version 2.0 at least. For the Hsqldb jdbc driver this is 2.0.
<code>urlProtocol</code>	The protocol of the database url (see below), usually <code>jdbc</code> .
<code>urlSubprotocol</code>	The sub-protocol of the database url which is database- and driver-specific. For Hsqldb, you're using <code>hsqldb</code> .
<code>urlDbalias</code>	This is the address that points the jdbc driver to the database. For Hsqldb this is per default the database name.
<code>torque.database</code>	If you're using Torque to create the database, then you have to set the database here (again). Unfortunately, this value is different from the <code>dbmsName</code> which defines the database for OJB. Currently, these values are defined: <b>axion, cloudscape, db2, db2400, hypersonic</b> (which is Hsqldb), <b>interbase</b> (use for Firebird), <b>mssql, mysql, oracle, postgresql, sapdb</b> , and <b>sybase</b> . Default value is <code>hypersonic</code> for use with Hsqldb.
<code>torque.database.createUrl</code>	This specifies the url that Torque will use in order to create the database. Depending on the database, this may be the same as the normal access url (the default value), but for some database this is different. Please check the manual of your database for this url.

If you know how the jdbc url for connecting to your database looks like, then you can derive the settings `databaseName`, `databaseName`, `databaseName` and `databaseName` easily: Assume this url is given as:

```
jdbc:mysql://localhost:3306/myDatabase
```

then these properties are

Property	Value
<code>databaseName</code>	<code>myDatabase</code>
<code>urlProtocol</code>	<code>jdbc</code>
<code>urlSubprotocol</code>	<code>mysql</code>
<code>urlDbalias</code>	<code>//localhost/myDatabase</code>

### 5.3.3.2. Building via build.xml

After setting up the build you're probably eager to actually build the project. Here's the actions that you can perform using the Ant build file `build.xml`:

Action (target in the build.xml file)	What it does
clean	Cleans up all files from the previous build.
compile	Compiles your java source files to <code>build/classes</code> . Usually, you don't run this target, but rather the next one which includes the compilation step.
build	Compiles your java sources files (using the compile action), and prepares the runtime configuration files using the settings that you specified in the <code>build.properties</code> file, most notably the <a href="#">repository database.xml</a> which will be located in the <code>build/resources</code> directory after the build. After you run this action, your application is ready to go (if the action ran successfully, of course).
jar	A convenience action that packs your successfully build application into a jar.
xdoclet	Creates the runtime configuration files that describe the repository, from javadoc comments embedded in your java source files. Details on how to this are given in the <a href="#">tutorials</a> and in the documentation of the <a href="#">XDoclet OJB module</a> .
setup-db	Creates the database and tables from a database-independent schema using Torque. You'll find more info on this schema in the documentation of the <a href="#">XDoclet OJB module</a> and on the <a href="#">Torque homepage</a> .
enhance-jdori	This is a sample target that shows how a class meant to be persistent with JDO, is processed by the JDO bytecode enhancer from the <a href="#">JDO reference implementation</a> . It uses the <code>Product</code> class from the <a href="#">JDO tutorial</a> (tutorial 5).

So, a typical build would be achieved with this Ant call:

```
ant build
```

If you want to create the database as well, and you have javadoc comments in your source code that describe the repository, then you would call Ant this way:

```
ant build setup-db
```

This will perform in that order the actions `build`, `xdoclet` (invoked automatically from the next action) and `setup-db`.

Of course, you do not need to use Torque to setup your database, but it is a convenient way to do so.

### 5.3.3.3. Sample project

First we change the database properties to these values (assuming that Torque will be used to setup the database):

Property	Value
<code>jcdAlias</code>	We leave the default value of <code>default</code> .
<code>databaseName</code>	Since the application manages products, we call

	the database <code>productmanager</code> .
<code>databaseUser</code>	This depends on your setup. For the purposes of this guide, let's call him <code>steve</code> .
<code>databasePassword</code>	Again depending on your setup. How about <code>secret</code> (you know that you should not use this password in reality ?!).
<code>dbmsName</code>	MySQL
<code>jdbcRuntimeDriver</code>	Its called <code>com.mysql.jdbc.Driver</code> .
<code>jdbcLevel</code>	For the newer Mysql drivers this is 3.0.
<code>urlProtocol</code>	The default of <code>jdbc</code> will do.
<code>urlSubprotocol</code>	For MySQL, we're using <code>mysql</code> .
<code>urlDbalias</code>	Assuming that the database runs locally on the default port, we have <code>//localhost/\${databaseName}</code> .
<code>torque.database</code>	We want to use Torque, so we put <code>mysql</code> here.
<code>torque.database.createUrl</code>	MySQL allows to create a database via <code>jdbc</code> . The url that we should use to do so, is the normal url used to access the database minus the database name. So the value here is: <code>\${urlProtocol}:\${urlSubProtocol}://localhost/</code> . Please note that the trailing slash is important.

Ok, now we have everything configured for building. The `build.properties` file now looks like this (the comments have been removed for brevity):

```
jcdAlias=default
databaseName=productmanager
databaseUser=steve
databasePassword=secret

dbmsName=MySQL
jdbcLevel=3.0
jdbcRuntimeDriver=com.mysql.jdbc.Driver
urlProtocol=jdbc
urlSubprotocol=mysql
urlDbalias="//localhost/${databaseName}"

torque.database=mysql
torque.database.createUrl=${urlProtocol}:${urlSubProtocol}://localhost/

jar.name=projectmanager.jar

source.dir=src
source.java.dir=${source.dir}/java
source.resource.dir=${source.dir}/resources
source.test.dir=${source.dir}/test
source.schema.dir=${source.dir}/schema

build.dir=build
build.lib.dir=lib
build.classes.dir=${build.dir}/classes/
build.resource.dir=${build.dir}/resources/

target.dir=target
```

Looks like we're ready for building. Again, we're assuming that the source code is already present. So we're invoking Ant now in the top-level folder `productmanager`:

```
ant build setup-db
```

which should (assuming five java classes) produce an output like this

```
Buildfile: build.xml

compile:
  [mkdir] Created dir: /home/steve/projects/productmanager/build
  [mkdir] Created dir: /home/steve/projects/productmanager/build/classes
  [javac] Compiling 5 source files to
/home/steve/projects/productmanager/build/classes

build:
  [copy] Copying 10 files to
/home/steve/projects/productmanager/build/resources

xdoclet:
[ojbdoclet] (XDocletMain.start                47 ) Running <ojbrepository/>
[ojbdoclet] Generating ojb repository descriptor
(build/resources//repository_user.xml)
[ojbdoclet] Type test.Project
[ojbdoclet] Processed 5 types
[ojbdoclet] Processed 5 types
[ojbdoclet] (XDocletMain.start                47 ) Running <torqueschema/>
[ojbdoclet] Generating torque schema (build/resources//project-schema.xml)
[ojbdoclet] Processed 5 types

setup-db:

check-use-classpath:

check-run-only-on-schema-change:

sql-check:

sql:
  [echo] +-----+
  [echo] |           |
  [echo] | Generating SQL for YOUR Torque project! |
  [echo] |           |
  [echo] |           |
  [echo] |           |
  [echo] +-----+

sql-classpath:
[torque-sql] Using contextProperties file:
/home/steve/projects/productmanager/build.properties
[torque-sql] Using classpath
[torque-sql] Generating to file
/home/steve/projects/productmanager/build/resources/report.productmanager.sql.generation
[torque-sql] Parsing file: 'ojbcore-schema.xml'
[torque-sql] (transform.DTDResolver            128 ) Resolver: used
database.dtd from
org.apache.torque.engine.database.transform package
[torque-sql] Parsing file: 'project-schema.xml'
[torque-sql] (transform.DTDResolver            140 ) Resolver: used
http://jakarta.apache.org/turbine/dtd/database.dtd

sql-template:

create-db-check:

create-db:
[torque-data-model] Using classpath
[torque-data-model] Generating to file
/home/steve/projects/productmanager/build/resources/create-db.sql
[torque-data-model] Parsing file: 'ojbcore-schema.xml'
[torque-data-model] (transform.DTDResolver            128 ) Resolver: used
database.dtd from
org.apache.torque.engine.database.transform package
[torque-data-model] Parsing file: 'project-schema.xml'
```



```
[torque-data-model] (transform.DTDResolver 140 ) Resolver: used
                    http://jakarta.apache.org/turbine/dtd/database.dtd
[echo]
[echo]      Executing the create-db.sql script ...
[echo]
[sql]  Executing file:
        /home/steve/projects/productmanager/build/resources/create-db.sql
[sql]  2 of 2 SQL statements executed successfully

insert-sql:
[torque-sql-exec] Our new url -> jdbc:mysql://localhost/productmanager
[torque-sql-exec] Executing file:
/home/steve/projects/productmanager/build/resources/project-schema.sql
[torque-sql-exec] Executing file:
/home/steve/projects/productmanager/build/resources/ojbcore-schema.sql
[torque-sql-exec] 50 of 50 SQL statements executed successfully

BUILD SUCCESSFUL
```

That was it. You now have your database setup properly. Go on, have a look:

```
mysql -u steve productmanager
mysql> show tables;
```

There, all tables for your project, as well as the tables required for some OJB functionality which we also used in the above process (you can recognize them by their names which start with `objb_`).

### 5.3.4. The runtime configuration files

The last thing missing for actually running your project is to adapt the runtime configuration files used by OJB. There are basically three sets of configuration that need to be provided: configuration of the OJB runtime, description of the database connection, and description of the repository.

#### 5.3.4.1. Configuring the OJB runtime

With the [OJB.properties](#) file and [OJB-logging.properties](#) (both located in `src/resources`), you configure and finetune the runtime aspects of OJB. For a simple application you'll probably won't have to change anything in them, though.

#### 5.3.4.2. Configuring the database connection

For projects that use OJB, you configure the connections to the database via [jdbc connection descriptors](#). These are usually defined in a file called `repository_database.xml` (located in `src/resources`). In the `objb-blank` project, the build file will setup this file for you and place it in the `build/resources` directory.

#### 5.3.4.3. Configuring the repository

Finally you need to configure the repository. It consists of descriptors that define which java classes are mapped in what way to which database tables, and it is typically contained in the `repository_user.xml` file. This is the most complicated configuration part which will be explained in much more detail in the rest of the [tutorials](#).

An convenient way of creating the repository metadata is to use the [XDoclet OJB module](#). Basically, you put specific Javadoc comments into your source code, which are then processed by the build file (`xdoclet` and `setup-db` targets) and the repository metadata and the database schema are generated.

#### 5.3.4.4. Sample project

Actually, there is not much to do here. For our simple sample application the default properties of

OJB work just fine, so we leave `OJB.properties` and `OJB-logging.properties` untouched.

Also, the build file generated the connection descriptor for us, and we were using the XDoclet OJB module and Torque to generate the repository metadata and database for us. For instance, the processed connection descriptor (file `build/resources/repository_database.xml`) looks like this:

```
<jdbc-connection-descriptor
  jcd-alias="default"
  default-connection="true"
  platform="MySQL"
  jdbc-level="3.0"
  driver="com.mysql.jdbc.Driver"
  protocol="jdbc"
  subprotocol="mysql"
  dbalias="//localhost/productmanager"
  username="steve"
  password="secret"
  eager-release="false"
  batch-mode="false"
  useAutoCommit="1"
  ignoreAutoCommitExceptions="false"
>
  <object-cache class="org.apache.ojb.broker.cache.ObjectCacheDefaultImpl">
    <attribute attribute-name="timeout" attribute-value="900"/>
    <attribute attribute-name="autoSync" attribute-value="true"/>
  </object-cache>
  <connection-pool
    maxActive="21"
    validationQuery="" />
  <sequence-manager
className="org.apache.ojb.broker.util.sequence.SequenceManagerHighLowImpl">
    <attribute attribute-name="grabSize" attribute-value="20"/>
    <attribute attribute-name="autoNaming" attribute-value="true"/>
    <attribute attribute-name="globalSequenceId" attribute-value="false"/>
    <attribute attribute-name="globalSequenceStart "
attribute-value="10000"/>
  </sequence-manager>
</jdbc-connection-descriptor>
```

If you're curious as to what this stuff means, check this [reference guide](#).

The repository metadata (file `build/resources/repository_user.xml`) starts like:

```
<class-descriptor
  class="productmanager.Product"
  table="Product"
>
  <field-descriptor
    name="name"
    column="name"
    jdbc-type="VARCHAR"
    length="32"
  >
  </field-descriptor>
  <field-descriptor
    name="price"
    column="price"
    jdbc-type="FLOAT"
  >
  </field-descriptor>
  <field-descriptor
    name="stock"
    column="stock"
    jdbc-type="INTEGER"
  >
  </field-descriptor>
```

```

    <field-descriptor
      name="id"
      column="id"
      jdbc-type="INTEGER"
      primaryKey="true"
    >
  </field-descriptor>
</class-descriptor>
...

```

Now you should be able to run your application:

```

cd build/resources
java productmanager.Main

```

Of course, you'll need to setup the CLASSPATH before running your application. You'll should add all jars from the lib folder except the ones for Torque (torque-[version].jar, velocity-[version].jar and commons-collections-[version].jar) and for the XDoclet OJB module (xdoclet-[version].jar, xjavadoc-[version].jar and xdoclet-ojb-module-[version].jar).

It is important to note that OJB per default assumes the OJB.properties and OJB-logging.properties files in the directory where you're starting the application. Hence, we changed to the build/resources directory before running the application. This of course requires the compiled classes to be on the classpath, as well (directory build/classes).

Per default, the same applies to the other configuration files (repository\*.xml) but you can change this in the OJB.properties file.

### 5.3.5. Learning More

After you've have learned about building and configuring projects that use OJB, you should check out the [tutorials](#) to learn how to specify your persistent classes and how to use OJB's APIs to perform database operations. The [Mapping Tutorial](#) in particular shows you how to map your classes to tables in an RDBMS.

## 5.4. Tutorials

### 5.4.1. Tutorial Summary

Summary of all OJB tutorials.

- [Object-Relational Mapping](#)  
The Object-Relational Mapping tutorial walks though a basic metadata mapping for an object to a relational database.
- [The Persistence Broker API](#)  
The PB tutorial demonstrates how to use the PersistenceBroker API which forms an object persistence kernel for OJB. While it is the lowest level API provided by OJB it is also exceptionally easy to use.
- [The ODMG API](#)  
The ODMG API tutorial steps though using the ODMG 3.0 API provided by OJB. This is an industry standard API designed for Object Databases.
- [The JDO API](#)  
JDO is a standard API for accessing persistent objects in Java. This tutorial steps through how to use OJB's JDO plugin.
- [The Object Transaction Manager](#)  
The OTM is OJB's implementation of object level transactions. These are transactions independent of the underlying relational database providing more efficient resource utilisation

and extremely flexible locking semantics.

Further strongly recommended documentation for all beginners:

- [OJB Queries](#)  
This document explains the usage of the query syntax.
- [Basic O/R Technique](#)  
This tutorial explains basic object-relational mapping technique in OJB like 1:1, 1:n and m:n relations, the auto-xxx settings for references and proxy objects/collections.
- [Tools to build large metadata mappings](#)  
Explains how to build large metadata mapping and present useful tools.

## 5.4.2. Mapping Tutorial

### 5.4.2.1. What is the Object-Relational Mapping Metadata?

The O/R mapping metadata is the specific configuration information that specifies how to map classes to relational tables. In OJB this is primarily accomplished through an xml document, the [repository.xml](#) file, which contains all of the initial mapping information.

#### The Product Class

This tutorial looks at mapping a simple class with no relations:

```
package org.apache.ojb.tutorials;

public class Product
{
    /** product name */
    private String name;

    /** price per item */
    private Double price;

    /** stock of currently available items */
    private int stock;

    ...
}
```

This class has three fields, `price`, `stock`, and `name`, that need to be mapped to the database. Additionally, we will introduce one artificial field used by the database that has no real meaning to the class, an artificial key primary id:

```
/** Artificial primary-key */
private Integer id;
```

Including the primary-key attribute in the class definition is mandatory, but under certain conditions [anonymous keys](#) can also be used to keep this database artifact hidden in the database. However, as access to an artificial unique identifier for a particular object instance can be useful, particularly in web-based applications, this tutorial will expose it

#### The Database

OJB is very flexible in terms of how it can map classes to database tables, however the simplest technique for mapping a single class to a relational database is to map the class to a single table, and each attribute on the class to a single column. Each row will then represent a unique instance of that class.

The DDL for such a table, for the `Product` class might look like:

```
CREATE TABLE Product
(
  id INTEGER PRIMARY KEY,
  name VARCHAR(100),
  price DOUBLE,
  stock INTEGER
)
```

The individual field names in the database and class definition match here, but this is no requirement. They may vary independently of each other as the metadata will specify what maps to what.

### The Metadata

The `repository.xml` document is split into several physical documents. The `repository_user.xml` xml file is used to contain user-defined mappings. OJB uses the other ones for managing other metadata, such as database information.

In general each class will be defined within a `class-descriptor` element with `field-descriptor` child elements for each field. In addition the mapping of references and collections is described in the [basic technique section](#). This tutorial sticks to mapping a single, simplistic, class.

The complete mapping for the `Product` class is as follows:

```
<class-descriptor
  class="org.apache.ojb.tutorials.Product"
  table="Product"
>
  <field-descriptor
    name="id"
    column="id"
    primarykey="true"
    autoincrement="true"
  />
  <field-descriptor
    name="name"
    column="name"
  />
  <field-descriptor
    name="price"
    column="price"
  />
  <field-descriptor
    name="stock"
    column="stock"
  />
</class-descriptor>
```

Examine the `class-descriptor` element. It has two attributes:

- `class` - This attribute is used to specify the fully-qualified Java class name for this mapping.
- `table` - This attribute specifies which table is used to store instances of this class.

Other information can be specified here, such as proxies and custom row-readers as specified in the [repository.xml documentation](#).

Examine now the first `field-descriptor` element. This is used to describe the `id` field of the `Product` class. Two required attributes are specified:

- `name` - This specifies the name of the instance variable in the Java class.
- `column` - This specifies the column in the table specified for this class used to store the value.

In addition to those required attributes, notice that the first element specifies two optional

attributes:

- `primary-key` - This attribute specifies that this field is the primary key for this class.
- `autoincrement` - The `autoincrement` attribute specifies that the value will be automatically assigned by OJB [sequence manager](#). This might use a database supplied sequence, or, by default, an OJB generated value.

### Using the XDoclet module

OJB provides an XDoclet module to make generating the repository descriptor and the corresponding table schema easier. An XDoclet module basically processes custom JavaDoc tags in the source code, and generates files from them. In the case of OJB, two types of files can be generated: the repository descriptor (`repository_user.xml`) and a Torque schema which can be used to create the tables in the database. This provides one important benefit: the descriptor and the database schema are much more likely in sync with the code thus avoiding errors that are usually hard to find. Furthermore, the XDoclet module contains some checks that find common mapping errors.

In the above example, the source code for Product class with JavaDoc tags would look like:

```
package org.apache.ojb.tutorials;

/**
 * @ojb.class
 */
public class Product
{
    /**
     * Artificial primary-key
     *
     * @ojb.field primaryKey="true"
     *           autoincrement="ojb"
     */
    private Integer id;

    /**
     * product name
     *
     * @ojb.field length="100"
     */
    private String name;

    /**
     * price per item
     *
     * @ojb.field
     */
    private Double price;

    /**
     * stock of currently available items
     *
     * @ojb.field
     */
    private int stock;
}
```

As you can see, much of the stuff that is present in the descriptor (and the DDL) is generated automatically by the XDoclet module, e.g. the table/column names and the jdbc-types. Of course, you can also specify them in the JavaDoc tags, e.g. if they differ from the java names.

For details on OJB's JavaDoc tags and how to generate and use the mapping files please see the [OJB XDoclet Module documentation](#).

## 5.4.2.2. Advanced Topics

### Relations

As most object models have relationships between objects, mapping specific types of relationships (1:1, 1:Many, Many:Many) is important in mapping objects into a relational database. The [basic technique tutorial](#) discusses this in great detail.

It is important to note that this metadata mapping can be [modified at runtime](#) through the [org.apache.ojb.metadata.MetadataManager](#) class.

### Inheritance

OJB can map inheritance hierarchies using a variety of techniques discussed in the [Extents and Polymorphism](#) section of the [Advanced O/R Documentation](#)

### Anonymous Keys

This tutorial uses explicit keys mapped into the Java class. It is also possible to keep artificial keys completely hidden within the database. The [Anonymous Keys HOWTO](#) explains how this is accomplished.

### Large Projects

Projects with small numbers of persistent classes can be mapped by hand, however, many projects can have hundreds, or even thousands, of distinct classes which must be mapped. In these circumstances managing the class-database mapping by hand is not viable. The [How To Build Mappings HOWTO](#) explores different tools which can be used for managing large-scale mapping.

### Custom JDBC Mapping

OJB maps Java types to JDBC types according to the [JDBC Types](#) table. You can, however, define custom JDBC -> Java type mappings via custom [field conversions](#).

## 5.4.3. Persistence Broker Tutorial

### 5.4.3.1. The PersistenceBroker API

#### Introduction

The *PersistenceBroker API* provides the lowest level access to OJB's persistence engine. While it is a low-level API compared to the OTM, ODMG, or JDO API's it is still very straightforward to use.

The core class in the PersistenceBroker API is the `org.apache.ojb.broker.PersistenceBroker` class. This class provides the point of access for all persistence operations in this API.

More detailed information can be found in the [PB-guide](#) and in the other [reference guides](#).

This tutorial operates on a simple example class:

```
package org.apache.ojb.tutorials;

public class Product
{
    /* Instance Properties */
}
```

```

private Double price;
private Integer stock;
private String name;

/* artificial property used as primary key */

private Integer id;

/* Getters and Setters */
...
}

```

The metadata descriptor for mapping this class is described in the [mapping tutorial](#)

The source code for all tutorials is available in the separate `tutorials-src.jar` which you can download [here](#). If you're eager to try them out, you can use them with the `ojb-blank` project which can be downloaded from the same place. It is described in the [Getting started](#) section.

Further information about the OJB PB-api implementation can be found in the [PB guide](#).

### A First Look - Persisting New Objects

The most basic operation is to persist an object. This is handled very easily by just

1. obtaining a `PersistenceBroker`
2. begin the PB-transaction
3. storing the object via the `PersistenceBroker`
4. commit transaction
5. closing the `PersistenceBroker`

For example, the following function stores a single object of type `Product`.

```

public static void storeProduct(Product product)
{
    PersistenceBroker broker = null;
    try
    {
        broker = PersistenceBrokerFactory.defaultPersistenceBroker();
        broker.beginTransaction();
        broker.store(product);
        broker.commitTransaction();
    }
    catch(PersistenceBrokerException e)
    {
        if(broker != null) broker.abortTransaction();
        // do more exception handling
    }
    finally
    {
        if (broker != null) broker.close();
    }
}

```

Two OJB classes are used here, the `PersistenceBrokerFactory` and the `PersistenceBroker`. The `PersistenceBrokerFactory` class manages the lifecycles of `PersistenceBroker` instances: it creates them, pools them, and destroys them as needed. The exact behavior is very configurable.

In this case we used the static

`PersistenceBrokerFactory.defaultPersistenceBroker()` method to obtain an instance of a `PersistenceBroker` to the default data source. This is most often how it is used if there is only one database for an application. If there are multiple data sources, a broker may be obtained by name (using a `PBKey` instance as argument in `PersistenceBrokerFactory.createPersistenceBroker(pbKey)`).



It is worth noting that the `broker.close()` call is made within a `finally {...}` block. This ensures that the broker will be closed, and returned to the broker pool, even if the function throws an exception.

To use this function, we just create a `Product` and pass it to the function:

```
Product product = new Product();
product.setName("Sprocket");
product.setPrice(1.99);
product.setStock(10);
storeProduct(product);
```

Once a `PersistenceBroker` has been obtained, its `PersistenceBroker.store(Object)` method is used to make an object persistent.

Maybe you have noticed that there has not been an assignment to `product.id`, the primary-key attribute. Upon storing `product` OJB detects that the attribute is not properly set and assigns a unique id. This automatic assignment of unique Ids for the attribute `id` has been explicitly declared in the [XML repository](#) file, as we discussed in the .

If several objects need to be stored, this can be done within a transaction, as follows.

```
public static void storeProducts(Product[] products)
{
    PersistenceBroker broker = null;
    try
    {
        broker = PersistenceBrokerFactory.defaultPersistenceBroker();
        broker.beginTransaction();
        for (int i = 0; i < products.length; i++)
        {
            broker.store(products[i]);
        }
        broker.commitTransaction();
    }
    catch(PersistenceBrokerException e)
    {
        if(broker != null) broker.abortTransaction();
        // do more exception handling
    }
    finally
    {
        if (broker != null) broker.close();
    }
}
```

This contrived example stores all of the passed `Product` instances within a single transaction via the `PersistenceBroker.beginTransaction()` and `PersistenceBroker.commitTransaction()`. These are database level transactions, not object level transactions.

### Querying Persistent Objects

Once objects have been stored to the database, it is important to be able to get them back. The `PersistenceBroker` API provides two mechanisms for [building queries](#), by using a template object, or by using specific criteria.

```
public static Product findByTemplate(Product template)
{
    PersistenceBroker broker = null;
    Product result = null;
    try
    {
        broker = PersistenceBrokerFactory.defaultPersistenceBroker();
```

```

        QueryByCriteria query = new QueryByCriteria(template);
        result = (Product) broker.getObjectByQuery(query);
    }
    finally
    {
        if (broker != null) broker.close();
    }
    return result;
}

```

This function finds a `Product` by building a query against a template `Product`. The template should have any properties set which should be matched by the query. Building on the previous example where a product was stored, we can now query for that same product:

```

Product product = new Product();
product.setName("Sprocket");
product.setPrice(new Double(1.99));
product.setStock(new Integer(10));
storeProduct(product);

Product template = new Product();
template.setName("Sprocket");
Product sameProduct = findByTemplate(template);

```

In the above code snippet, `product` and `sameProduct` will reference the same object (assuming there are no additional products in the database with the name "Sprocket").

The template `Product` has only one of its properties set, the name property. The others are all null. Properties with null values are not used to match.

An alternate, and more flexible, way to have specified a query via the `PersistenceBroker` API is by constructing the criteria on the query by hand. The following function does this.

```

public static Collection getExpensiveLowStockProducts()
{
    PersistenceBroker broker = null;
    Collection results = null;
    try
    {
        broker = PersistenceBrokerFactory.defaultPersistenceBroker();

        Criteria criteria = new Criteria();
        criteria.addLessOrEqualThan("stock", new Integer(20));
        criteria.addGreaterOrEqualThan("price", new Double(100000.0));

        QueryByCriteria query = new QueryByCriteria(Product.class, criteria);
        results = broker.getCollectionByQuery(query);
    }
    finally
    {
        if (broker != null) broker.close();
    }
    return results;
}

```

This function builds a `Criteria` object and uses it to set more complex query parameters - in this case greater-than and less-than constraints. Looking at the first constraint put on the criteria, `criteria.addLessOrEqualThan("stock", new Integer(10));` notice the arguments. The first is the property name on the object being queried for. The second is an `Integer` instance to be used for the comparison.

After the `Criteria` has been built, the `QueryByCriteria` constructor used is also different from the previous example. In this case the criteria does not know the type of the object it is being used against, so the `Class` must be specified to the query.

Finally, notice that this example uses the

`PersistenceBroker.getCollectionByQuery(...)` method instead of the `PersistenceBroker.getObjectByQuery(...)` method used previously. This is used because we want all of the results. Either form can be used with either method of constructing queries. In the case of the `PersistenceBroker.getObjectByQuery(...)` style query, the first matching object is returned, even if there are multiple matching objects.

### Updating Persistent Objects

The same mechanism, and method, is used for updating persistent objects as for inserting persistent objects. The same `PersistenceBroker.store(Object)` method is used to store a modified object as to insert a new one - the difference between new and modified objects is irrelevant to OJB.

This can cause some confusion for people who are very used to working in the stricter confines of SQL inserts and updates. Basically, OJB will insert a new object into the relational store if the primary key, as specified in the O/R metadata is not in use. If it is in use, it will update the existing object rather than create a new one.

This allows programmers to treat every object the same way in an object model, whether it has been newly created and made persistent, or materialized from the database.

Typically, making changes to a persistent object first requires retrieving a reference to the object, so the typical update cycle, unless the application caches objects, is to query for the object to modify, modify the object, and then store the object. The following function demonstrates this behavior by "selling" a Product.

```
public static boolean sellOneProduct(Product template)
{
    PersistenceBroker broker = null;
    boolean isSold = false;
    try
    {
        broker = PersistenceBrokerFactory.defaultPersistenceBroker();
        QueryByCriteria query = new QueryByCriteria(template);
        Product result = (Product) broker.getObjectByQuery(query);

        if (result != null)
        {
            broker.beginTransaction();
            result.setStock(new Integer(result.getStock().intValue() - 1));
            broker.store(result);
            // alternative, more performant
            // broker.store(result, ObjectModificationDefaultImpl.UPDATE);
            broker.commitTransaction();
            isSold = true;
        }
    }
    catch(PersistenceBrokerException e)
    {
        if(broker != null) broker.abortTransaction();
        // do more exception handling
    }
    finally
    {
        if (broker != null) broker.close();
    }
    return isSold;
}
```

This function uses the same query-by-template and `PersistenceBroker.store()` API's examined previously, but it uses the store method to store changes to the object it retrieved. It is worth noting that the entire operation took place within a transaction.

## Deleting Persistent Objects

Deleting persistent objects from the repository is accomplished via the `PersistenceBroker.delete()` method. This removes the persistent object from the repository, but does not affect any change on the object itself. For example:

```
public static void deleteProduct(Product product)
{
    PersistenceBroker broker = null;
    try
    {
        broker = PersistenceBrokerFactory.defaultPersistenceBroker();
        broker.beginTransaction();
        broker.delete(product);
        broker.commitTransaction();
    }
    catch(PersistenceBrokerException e)
    {
        if(broker != null) broker.abortTransaction();
        // do more exception handling
    }
    finally
    {
        if (broker != null) broker.close();
    }
}
```

This method simply deletes an object from the database.

## Find object by primary key

In some cases only the primary key values (single field or n-fields for composed primary keys) of an object are known. In OJB you have several ways to request the whole object. It is possible to build a query as [shown above](#), but the smarter solution is to use `PersistenceBroker#getObjectByIdentity(Identity oid)`. An [Identity](#) object is a unique representation of a persistence capable object based on the object primary key values and the *top-level class* (abstract class, interface or the class itself, depending on the [extent metadata mapping](#)).

For example, to find an *Product* with an single primary key of '23' do

```
Identity oid = broker.serviceIdentity().buildIdentity(Product.class, new
Integer(23));
Product product = (Product) broker.getObjectByIdentity(oid);
```

### 5.4.3.2. Exception Handling

Most `PersistenceBroker` operations throw a `org.apache.ojb.broker.PersistenceBrokerException`, which is derived from `java.lang.RuntimeException` if an error occurs. This means that no try/catch block is **required** but does not mean that it should not be used. This tutorial specifically does not catch exceptions all in order to focus more tightly on the specifics of the API, however, best usage would be to include a try/catch/finally block around persistence operations using the `PersistenceBroker` API.

Additionally, the closing of `PersistenceBroker` instances is best handled in `finally` blocks in order to guarantee that it is run, even if an exception occurs. If the `PersistenceBroker.close()` is not called then the application will leak broker instances. The best way to ensure that it is always called is to always retrieve and use `PersistenceBroker` instances within a `try { ... }` block, and always close the broker in a `finally { ... }` block attached to the `try { ... }` block.

A better designed `getExpensiveLowStockProducts()` method is presented here.

```
public static Collection betterGetExpensiveLowStockProducts()
{
    PersistenceBroker broker = null;
    Collection results = null;
    try
    {
        broker = PersistenceBrokerFactory.defaultPersistenceBroker();

        Criteria criteria = new Criteria();
        criteria.addLessOrEqualThan("stock", new Integer(20));
        criteria.addGreaterOrEqualThan("price", new Double(100000.0));

        QueryByCriteria query = new QueryByCriteria(Product.class, criteria);
        results = broker.getCollectionByQuery(query);
    }
    catch (PersistenceBrokerException e)
    {
        // Handle exception
    }
    finally
    {
        if (broker != null) broker.close();
    }
    return results;
}
```

Notice first that the `PersistenceBroker` is retrieved and used within the confines of a `try { ... }` block. Assuming nothing goes wrong the entire operation will execute there, all the way to the `return results;` line. Java guarantees that `finally { ... }` blocks will be called before a method returns, so the `broker.close()` method is only included once, in the `finally` block. As an exception may have occurred while attempting to retrieve the broker, a not-null test is first performed before closing the broker.

## 5.4.4. The ODMG API

### 5.4.4.1. Introduction

The *ODMG API* is an implementation of the [ODMG 3.0 Object Persistence API](#). The ODMG API provides a higher-level API and query language based interface over the [PersistenceBroker API](#).

More detailed information can be found in the [ODMG-guide](#) and in the other [reference guides](#).

This tutorial operates on a simple example class:

```
package org.apache.ojb.tutorials;

public class Product
{
    /* Instance Properties */

    private Double price;
    private Integer stock;
    private String name;

    /* artificial property used as primary key */

    private Integer id;

    /* Getters and Setters */
    ...
}
```

The metadata descriptor for mapping this class is described in the [mapping tutorial](#)

When using 1:1, 1:n and m:n references (the example doesn't use it) the ODMG-api need specific metadata settings on relationship definition, the mandatory settings are listed in the [ODMG-Guide](#) - additional info see [auto-xxx settings](#) and [repository file settings](#).

As with the other tutorials, the source code for this tutorial is contained in the `tutorials-src.jar` which can be downloaded [here](#). The source files are contained in the `org/apache/ojb/tutorial2/` directory.

You can try it out with the ojb-blank project which can be downloaded from the same place and is described in the [Getting started](#) section.

Further information about the OJB odmg-api implementation can be found in the [ODMG guide](#).

#### 5.4.4.2. Initializing ODMG

The ODMG implementation needs to have a database opened for it to access. This is accomplished via the following code:

```
Implementation odmg = OJB.getInstance();
Database db = odmg.newDatabase();
db.open("default", Database.OPEN_READ_WRITE);

/* ... use the database ... */

db.close();
```

With method call `OJB.getInstance()` always a **new** [org.odmg.Implementation](#) instance will be created and `odmg.newDatabase()` returns a new `Database` instance.

Call `db.open(...)` opens an ODMG Database using the name specified in [metadata for the database](#) -- "default" in this case. Notice the Database is opened in read/write mode. It is possible to open it in *read-only* or *write-only* modes as well.

Once a `Implementation` instance is created and a `Database` has been opened it is available for use. Unlike [PersistenceBroker instances](#), ODMG `Implementation` and `Database` instances are threadsafe and can typically be used for the entire lifecycle of an application. There is no need to call the `Database.close()` method until the database is truly no longer needed.

The `OJB.getInstance()` function provides the ODMG `Implementation` instance required for using the ODMG API. From here on out it is straight ODMG code that should work against any compliant ODMG implementation.

#### 5.4.4.3. Persisting New Objects

Persisting an object via the ODMG API is handled by writing it to the persistence store within the context of a transaction:

```
public static void storeNewProduct(Product product)
{
    // get the used Implementation instance
    Implementation odmg = ...;
    Transaction tx = odmg.newTransaction();
    tx.begin();
    // get current used Database instance
    Database db = odmg.getDatabase(null);
    // make persistent new object
    db.makePersistent(product);
    tx.commit();
}
```

Once the ODMG implementation has been obtained it is used to begin a transaction, obtain a write lock on the `Product`, and commit the transaction. It is very important to note that all changes

need to be made within transactions in the ODMG API. When the transaction is committed the changes are made to the database. Until the transaction is committed the database is unaware of any changes -- they exist solely in the object model.

#### 5.4.4.4. Querying Persistent Objects

The ODMG API uses the OQL query language for obtaining references to persistent objects. OQL is very similar to SQL, and using it is very similar to use JDBC. The ODMG implementation is used to create a query, the query is specified, executed, and a list of results is returned:

```
public static Product findProductByName(String name) throws Exception
{
    // get the used Implementation instance
    Implementation odmg = ...;
    Transaction tx = odmg.newTransaction();
    tx.begin();

    OQLQuery query = odmg.newOQLQuery();
    query.create("select products from "
                + Product.class.getName()
                + " where name = $1");
    query.bind(name);
    List results = (List) query.execute();
    Product product = (Product) results.iterator().next();

    tx.commit();
    return product;
}
```

#### 5.4.4.5. Updating Persistent Objects

Updating a persistent object is done by modifying it in the context of a transaction, and then committing the transaction:

```
public static void sellProduct(Product product, int number)
{
    // get the used Implementation instance
    Implementation odmg = ...;
    Transaction tx = odmg.newTransaction();
    tx.begin();

    tx.lock(product, Transaction.WRITE);
    product.setStock(new Integer(product.getStock().intValue() - number));

    tx.commit();
}
```

The sample code obtains a write lock on the object (**before** the changes are made), binding it to the transaction, changes the object, and commits the transaction. The newly modified `Product` now has a new stock value.

#### 5.4.4.6. Deleting Persistent Objects

Deleting persistent objects requires directly addressing the Database which contains the persistent object. This can be obtained from the ODMG Implementation by asking for it. Once retrieved, just ask the Database to delete the object. Once again, this is all done in the context of a transaction.

```
public static void deleteProduct(Product product)
{
    // get the used Implementation instance
    Implementation odmg = ...;
    Transaction tx = odmg.newTransaction();
```

```

tx.begin();
// get current used Database instance
Database db = odmgt.getDatabase(null);
db.deletePersistent(product);
tx.commit();
}

```

It is important to note that the `Database.deletePersistent()` call does not delete the object itself, just the persistent representation of it. The transient object still exists and can be used however desired -- it is simply no longer persistent.

## 5.4.5. JDO Tutorial

### 5.4.5.1. Using the ObjectRelationalBridge JDO API

#### Introduction

This document demonstrates how to use `ObjectRelationalBridge` and the JDO API in a simple application scenario. The tutorial application implements a product catalog database with some basic use cases. The source code for the tutorial application is shipped in the `tutorials-src.jar` which can be downloaded [here](#). The source for this tutorial is found in the directory `org/apache/ojb/tutorial5`.

This document is not meant as a complete introduction to JDO. For more information see: [Sun's JDO site](#).

#### Note:

OJB does not provide it's own JDO implementation yet. A full JDO implementation is in the scope of the 2.0 release. For the time being we provide a plugin to the JDO reference implementation called `OjbStore`. The `OjbStore` plugin resides in the package `org.apache.ojb.jdori.sql`.

#### Running the Tutorial Application

To install and run the demo application with the `ojb-blank` sample project (which is described in more detail [here](#)) please follow the following steps:

1. Extract the `tutorial-src.jar` that you downloaded from [here](#) into the `src/java` subdirectory of the `ojb-blank` project.

The JDO tutorial source files are contained in the `org/apache/ojb/tutorial5` subdirectory, and you can safely erase the subdirectories of the other tutorials.

2. Download the JDO Reference Implementation from [Sun's JDO site](#).

Extract the archiv to a local directory and copy the files:

- `jdori.jar`
- `jdo.jar`

into the `lib` directory of the project.

3. Now you can run the test application with these commands:

```
ant build enhance-jdori
```

from the toplevel project directory. The latter of these commands will enhance the jdo tutorial classes. Note that due to some limitations in the JDO reference implementation, the ant target will only work for the JDO tutorial, so if you want to create you own JDO application using the `ojb-blank` project, you have to adapt the build file accordingly.

To setup the test database you can issue this command

```
ant setup-db
```

4. Now you can start the tutorial application by executing



```
cd build/resources
java org.apache.ojb.tutorial5.Main
```

from the project toplevel directory.

#### 5.4.5.2. Using the JDO API in the UseCase Implementations

As shown [here](#) OJB supports four different API's. The PersistenceBroker, the OTM layer, the ODMG implementation, and the JDO implementation.

The [PB tutorial](#) implemented the sample application's use cases with the PersistenceBroker API. This tutorial will show how the same use cases can be implemented using the JDO API.

You can get more information about the JDO API at [JDO javadocs](#).

#### Obtaining the JDO PersistenceManager Object

In order to access the functionalities of the JDO API you have to deal with a special facade object that serves as the main entry point to all JDO operations. This facade is specified by the Interface `javax.jdo.PersistenceManager`.

A Vendor of a JDO compliant product must provide a specific implementation of the `javax.jdo.PersistenceManager` interface. JDO also specifies that a JDO implementation must provide a `javax.jdo.PersistenceManagerFactory` implementation that is responsible for generating `javax.jdo.PersistenceManager` instances.

So if you know how to use the JDO API you only have to learn how to obtain the OJB specific `PersistenceManagerFactory` object. Ideally this will be the only vendor specific operation.

In our tutorial application the `PersistenceManagerFactory` object is obtained in the constructor of the Application class and reached to the use case implementations for further usage:

```
public Application()
{
    factory = null;
    manager = null;
    try
    {
        // create OJB specific factory:
        factory = new OjbStorePMF();
    }
    catch (Throwable t)
    {
        System.out.println("ERROR: " + t.getMessage());
        t.printStackTrace();
    }
    useCases = new Vector();
    useCases.add(new UCListAllProducts(factory));
    useCases.add(new UCEnterNewProduct(factory));
    useCases.add(new UCEditProduct(factory));
    useCases.add(new UCDeleteProduct(factory));
    useCases.add(new UCQuitApplication(factory));
}
```

The class `org.apache.ojb.jdori.sql.OjbStorePMF` is the OJB specific `javax.jdo.PersistenceManagerFactory` implementation.

##### TODO: Put information about the .jdo files #####

The `PersistenceManagerFactory` object is reached to the constructors of the UseCases. These constructors store it in a protected attribute `factory` for further usage.

#### Retrieving collections

The next thing we need to know is how this Implementation instance integrates into our persistence operations.

In the use case `UCListAllProducts` we have to retrieve a collection containing all product entries from the persistent store. To retrieve a collection containing objects matching some criteria we can use the JDOQL query language as specified by the JDO spec. In our use case we want to select *all* persistent instances of the class `Products`. In this case the query is quite simple as it does not need any limiting search criteria.

We use the factory to create a `PersistenceManager` instance in step one. In the second step we ask the `PersistenceManager` to create a query returning all `Product` instances.

In the third step we perform the query and collect the results in a collection.

In the fourth step we iterate through the collection to print out each product matching our query.

```
public void apply()
{
    // 1. get a PersistenceManager instance
    PersistenceManager manager = factory.getPersistenceManager();
    System.out.println("The list of available products:");

    try
    {
        // clear cache to provoke query against database
        PersistenceBrokerFactory
            .defaultPersistenceBroker().clearCache();

        // 2. start tx and form query
        manager.currentTransaction().begin();
        Query query = manager.newQuery(Product.class);

        // 3. perform query
        Collection allProducts = (Collection)query.execute();

        // 4. now iterate over the result to print each
        // product and finish tx
        java.util.Iterator iter = allProducts.iterator();
        if (! iter.hasNext())
        {
            System.out.println("No Product entries found!");
        }
        while (iter.hasNext())
        {
            System.out.println(iter.next());
        }
        manager.currentTransaction().commit();
    }
    catch (Throwable t)
    {
        t.printStackTrace();
    }
    finally
    {
        manager.close();
    }
}
```

### Storing objects

Now we will have a look at the use case `UCEnterNewProduct`. It works as follows: first create a new object, then ask the user for the new product's data (productname, price and available stock). These data is stored in the new object's attributes. This part is no different from the [PB tutorial](#) implementation. (Steps 1. and 2.)

Now we will store the newly created object in the persistent store by means of the JDO API. With JDO, all persistence operations must happen within a transaction. So the third step is to ask the PersistenceManager object for a fresh `javax.jdo.Transaction` object to work with. The `begin()` method starts the transaction.

We then have to ask the PersistenceManager to make the object persistent in step 4.

In the last step we commit the transaction. All changes to objects touched by the transaction are now made persistent. As you will have noticed there is no need to explicitly store objects as with the PersistenceBroker API. The Transaction object is responsible for tracking which objects have been modified and to choose the appropriate persistence operation on commit.

```
public void apply()
{
    // 1. this will be our new object
    Product newProduct = new Product();
    // 2. now read in all relevant information and fill the new object:
    System.out.println("please enter a new product");
    String in = readLineWithMessage("enter name:");
    newProduct.setName(in);
    in = readLineWithMessage("enter price:");
    newProduct.setPrice(Double.parseDouble(in));
    in = readLineWithMessage("enter available stock:");
    newProduct.setStock(Integer.parseInt(in));

    // 3. create PersistenceManager and start transaction
    PersistenceManager manager = factory.getPersistenceManager();

    Transaction tx = null;
    tx = manager.currentTransaction();
    tx.begin();

    // 4. mark object as persistent
    manager.makePersistent(newProduct);

    // 5. commit transaction
    tx.commit();

    manager.close();
}
```

## Updating Objects

The UseCase `UCEditProduct` allows the user to select one of the existing products and to edit it.

The user enters the products unique id. The object to be edited is looked up by this id. (Steps 1., 2. and 3.) This lookup is necessary as our application does not hold a list of all product objects.

The product is then edited (Step 4.).

In step five the transaction is committed. All changes to objects touched by the transaction are now made persistent. Because we modified an existing object an update operation is performed against the backend database.

```
public void apply()
{
    PersistenceManager manager = null;

    // ask user which object should edited
    String in = readLineWithMessage("Edit Product with id:");
    int id = Integer.parseInt(in);

    Product toBeEdited;
    try
```

```

{
    // 1. start transaction
    manager = factory.getPersistenceManager();
    manager.currentTransaction().begin();

    // We don't have a reference to the selected Product.
    // So we have to look it up first,

    // 2. Build a query to look up product by the id
    Query query = manager.newQuery(Product.class, "id == " + id);

    // 3. execute query
    Collection result = (Collection) query.execute();
    toBeEdited = (Product) result.iterator().next();

    if (toBeEdited == null)
    {
        System.out.println("did not find a matching instance...");
        manager.currentTransaction().rollback();
        return;
    }

    // 4. edit the existing entry
    System.out.println("please edit the product entry");
    in =
        readLineWithMessage(
            "enter name (was " + toBeEdited.getName() + "):");
    toBeEdited.setName(in);
    in =
        readLineWithMessage(
            "enter price (was " + toBeEdited.getPrice() + "):");
    toBeEdited.setPrice(Double.parseDouble(in));
    in =
        readLineWithMessage(
            "enter available stock (was "
                + toBeEdited.getStock()
                + "):");
    toBeEdited.setStock(Integer.parseInt(in));

    // 5. commit changes
    manager.currentTransaction().commit();
}
catch (Throwable t)
{
    // rollback in case of errors
    manager.currentTransaction().rollback();
    t.printStackTrace();
}
finally
{
    manager.close();
}
}

```

## Deleting Objects

The UseCase `UCDeleteProduct` allows the user to select one of the existing products and to delete it from the persistent storage.

The user enters the products unique id. The object to be deleted is looked up by this id. (Steps 1., 2. and 3.) This lookup is necessary as our application does not hold a list of all product objects.

In the fourth step we check if a `Product` matching to the id could be found. If no entry is found we print a message and quit the work.

If a `Product` entry was found we delete it in step 5 by calling the `PersistenceManager` to delete the persistent object. On transaction commit all changes to objects touched by the transaction are made

persistent. Because we marked the Product entry for deletion, a delete operation is performed against the backend database.

```
public void apply()
{
    PersistenceManager manager = null;
    Transaction tx = null;
    String in = readLineWithMessage("Delete Product with id:");
    int id = Integer.parseInt(in);

    try
    {
        // 1. start transaction
        manager = factory.getPersistenceManager();
        tx = manager.currentTransaction();
        tx.begin();

        // 2. Build a query to look up product by the id
        Query query = manager.newQuery(Product.class, "id == " + id);

        // 3. execute query
        Collection result = (Collection) query.execute();

        // 4. if no matching product was found, print a message
        if (result.size() == 0)
        {
            System.out.println("did not find a Product with id=" + id);
            tx.rollback();
            manager.close();
            return;
        }
        // 5. if a matching product was found, delete it
        else
        {
            Product toBeDeleted = (Product) result.iterator().next();
            manager.deletePersistent(toBeDeleted);
            tx.commit();
            manager.close();
        }
    }
    catch (Throwable t)
    {
        // rollback in case of errors
        //broker.abortTransaction();
        tx.rollback();
        t.printStackTrace();
    }
}
```

### 5.4.5.3. Conclusion

In this tutorial you learned to use the standard JDO API as implemented by the OJB system within a simple application scenario. I hope you found this tutorial helpful. Any comments are welcome.

## 5.4.6. Object Transaction Manager Tutorial

### 5.4.6.1. The OTM API

#### Introduction

The Object Transaction Manager (OTM) is written as a tool on which to implement other high-level object persistence APIs. It is, however, very usable directly. It supports API's similar to the [ODMG](#) and [PersistenceBroker](#) API's in OJB. Several of its idioms are designed around the fact that it is meant to have additional, client-oriented, API's built on top of it, however.

The `OTMKit` is the initial access point to the OTM interfaces. The kit provides basic configuration information to the OTM components used in your system. This tutorial will use the `SimpleKit` which will work well under most circumstances for local transaction implementations.

This tutorial operates on a simple example class:

```
package org.apache.ojb.tutorials;

public class Product
{
    /* Instance Properties */

    private Double price;
    private Integer stock;
    private String name;

    /* artificial property used as primary key */

    private Integer id;

    /* Getters and Setters */
    ...
}
```

The metadata descriptor for mapping this class is described in the [mapping tutorial](#).

As always the source code for this tutorial can be found in the `tutorials-src.jar` available from [here](#), more specifically in the `org/apache/ojb/tutorials/` directory.

### Persisting New Objects

The starting point for using the OTM directly is to look at making a transient object persistent. This code will use three things, an `OTMKit`, an `OTMConnection`, and a `Transaction`. The connection and transaction objects are obtained from the kit.

Initial access to the OTM client API's is through the `OTMKit` interface. We'll use the `SimpleKit`, an implementation of the `OTMKit` suitable for most circumstances using local transactions.

```
public static void storeProduct(Product product) throws LockingException
{
    OTMKit kit = SimpleKit.getInstance();
    OTMConnection conn = null;
    Transaction tx = null;
    try
    {
        conn = kit.acquireConnection(PersistenceBrokerFactory.getDefaultKey());
        tx = kit.getTransaction(conn);
        tx.begin();
        conn.makePersistent(product);
        tx.commit();
    }
    catch (LockingException e)
    {
        if (tx.isInProgress()) tx.rollback();
        throw e;
    }
    finally
    {
        conn.close();
    }
}
```

A kit is obtained and is used to obtain a connection. Connections are against a specific JCD alias. In this case we use the default, but a named datasource could also be used, as configured in the metadata repository. A transaction is obtained from the kit for the specific connection. Because

multiple connections can be bound to the same transaction in the OTM, the transaction needs to be acquired from the kit instead of the connection itself. The `SimpleKit` uses the commonly seen transaction-per-thread idiom, but other kits do not need to do this.

Every persistence operation within the OTM needs to be executed within the context of a transaction. The JDBC concept of implicit transactions doesn't exist in the OTM -- transactions must be explicit.

Locks, on the other hand, are implicit in the OTM (though explicit locks are available). The `conn.makePersistent(...)` call obtains a write lock on `product` and will commit (insert) the object when the transaction is committed.

The `LockingException` will be thrown if the object cannot be write-locked in this transaction. As it is a transient object to begin with, this will probably only ever happen if it has been write-locked in another transaction already -- but this depends on the transaction semantics configured in the repository metadata.

Finally, connections maintain resources so it is important to make sure they are closed when no longer needed.

### Deleting Persistent Objects

Deleting a persistent object from the backing store (making a persistent object transient) is almost identical to making it persistent -- the difference is just in the `conn.deletePersistent(product)` call instead of the `conn.makePersistent(product)` call. The same notes about transactions and resources apply here.

```
public static void storeProduct(Product product) throws LockingException
{
    OTMKit kit = SimpleKit.getInstance();
    OTMConnection conn = null;
    Transaction tx = null;
    try
    {
        conn = kit.acquireConnection(PersistenceBrokerFactory.getDefaultKey());
        tx = kit.getTransaction(conn);
        tx.begin();
        conn.deletePersistent(product);
        tx.commit();
    }
    catch (LockingException e)
    {
        if (tx.isInProgress()) tx.rollback();
        throw e;
    }
    finally
    {
        conn.close();
    }
}
```

### Querying for Objects

The OTM implements a transaction system, not a new client API. As such it supports two styles of query at present -- an `PersistenceBroker` like query-by-criteria style querying system, and an ODMG OQL query system.

Information on constructing these types of queries is available in the [PersistenceBroker](#) and [ODMG](#) tutorials respectively. Using those queries with the OTM is examined here.

A PB style query can be handled as follows:

```
public Iterator findByCriteria(Query query)
{
    OTMKit kit = SimpleKit.getInstance();
    OTMConnection conn = null;
    Transaction tx = null;
    try
    {
        conn = kit.acquireConnection(PersistenceBrokerFactory.getDefaultKey());
        tx = kit.getTransaction(conn);
        tx.begin();
        Iterator results = conn.getIteratorByQuery(query);
        tx.commit();
        return results;
    }
    finally
    {
        conn.close();
    }
}
```

Where, by default, a read lock is obtained on the returned objects. If a different lock is required it may be specified specifically:

```
public Iterator findByCriteriaWithLock(Query query, int lock)
{
    OTMKit kit = SimpleKit.getInstance();
    OTMConnection conn = null;
    Transaction tx = null;
    try
    {
        conn = kit.acquireConnection(PersistenceBrokerFactory.getDefaultKey());
        tx = kit.getTransaction(conn);
        tx.begin();
        Iterator results = conn.getIteratorByQuery(query, lock);
        tx.commit();
        return results;
    }
    finally
    {
        conn.close();
    }
}
```

The int lock argument is one of the integer constants on `org.apache.ojb.otm.lock.LockType`:

```
LockType.NO_LOCK
LockType.READ_LOCK
LockType.WRITE_LOCK
```

OQL queries are also supported, as this somewhat more complex example demonstrates:

```
public Iterator findByOQL(String query, Object[] bindings) throws Exception
{
    OTMKit kit = SimpleKit.getInstance();
    OTMConnection conn = null;
    Transaction tx = null;
    try
    {
        conn = kit.acquireConnection(PersistenceBrokerFactory.getDefaultKey());
        tx = kit.getTransaction(conn);
        OQLQuery oql = conn.newOQLQuery();
        oql.create(query);
        for (int i = 0; i < bindings.length; ++i)
        {
            oql.bind(bindings[i]);
        }
    }
}
```



```

        tx.begin();
        Iterator results = conn.getIteratorByOQLQuery(oql);
        tx.commit();
        return results;
    }
    catch (QueryInvalidException e)
    {
        if (tx.isInProgress()) tx.rollback();
        throw new Exception("Invalid OQL expression given", e);
    }
    catch (QueryParameterCountInvalidException e)
    {
        if (tx.isInProgress()) tx.rollback();
        throw new Exception("Incorrect number of bindings given", e);
    }
    catch (QueryParameterTypeInvalidException e)
    {
        if (tx.isInProgress()) tx.rollback();
        throw new Exception("Incorrect type of object given as binding", e);
    }
    finally
    {
        conn.close();
    }
}

```

This function is, at its core, doing the same thing as the PB style queries, except that it constructs the OQL query, which supports binding values in a manner similar to JDBC prepared statements.

The OQL style queries also support specifying the lock level the same way:

```

Iterator results = conn.getIteratorByOQLQuery(query, lock);

```

### More Sophisticated Transaction Handling

These examples are a bit simplistic as they begin and commit their transactions all in one go -- they are only good for retrieving data. More often data will need to be retrieved, used, and committed back.

Only changes to persistent objects made within the bounds of a transaction are persisted. This means that frequently a query will be executed within the bounds of an already established transaction, data will be changed on the objects obtained, and the transaction will then be committed back.

A very convenient way to handle transactions in many applications is to start a transaction and then let any downstream code be executed within the bounds of the transaction automatically. This is straightforward to do with the OTM using the SimpleKit! Take a look at a very slightly modified version of the query by criteria function:

```

public Iterator moreRealisticQueryByCriteria(Query query, int lock)
{
    OTMKit kit = SimpleKit.getInstance();
    OTMConnection conn = null;
    Transaction tx = null;
    try
    {
        conn = kit.acquireConnection(PersistenceBrokerFactory.getDefaultKey());
        tx = kit.getTransaction(conn);
        boolean auto = ! tx.isInProgress();
        if (auto) tx.begin();
        Iterator results = conn.getIteratorByQuery(query, lock);
        if (auto) tx.commit();
        return results;
    }
    finally
    {

```

```

        conn.close();
    }
}

```

In this case the function looks to see if a transaction is already in progress and sets a boolean flag if it is, `auto`. It then handles transactions itself, or allows the already opened transaction to maintain control.

Because connections can be attached to existing transactions the `SimpleKit` can attach the new connection to the already established transaction, allowing this function to work as expected whether there is a transaction in progress or not!

Client code using this function could then open a transaction, query for products, change them, and commit the changes back. For example:

```

public void renameWidgetExample()
{
    OTMKit kit = SimpleKit.getInstance();
    OTMConnection conn = null;
    Transaction tx = null;
    try
    {
        conn = kit.acquireConnection(PersistenceBrokerFactory.getDefaultKey());
        tx = kit.getTransaction(conn);
        tx.begin();
        Product sample = new Product();
        sample.setName("Wonder Widget");
        Query query = QueryFactory.newQueryByExample(sample);
        Iterator wonderWidgets
            = moreRealisticQueryByCriteria(query, LockType.WRITE_LOCK);
        while (wonderWidgets.hasNext())
        {
            Product widget = (Product) wonderWidgets.next();
            widget.setName("Improved Wonder Widget");
        }
        tx.commit();
    }
    finally
    {
        conn.close();
    }
}

```

This sample renames a whole bunch of products from "Wonder Widget" to "Improved Wonder Widget" and stores them back. It must make the changes within the context of the transaction it obtained for those changes to be stored back to the database. If the same iterator were obtained outside of a transaction, and the changes made, the changes would be made on the objects in memory, but not in the database. You can think of non-transaction objects as free immutable transfer objects.

This example also demonstrates two connections bound to the same transaction, as the `renameWidgetExample(...)` function obtains a connection, and the `moreRealisticQueryByCriteria(...)` function obtains an additional connection to the same transaction!

#### 5.4.6.2. Notes on the Object Transaction Manager

##### Transactions

The Object Transaction Manager (OTM) is a transaction management layer for Java objects. It typically maps 1:1 to database transactions behind the scenes, but this is not actually required for the OTM to work correctly.

The OTM supports a wide range of transactional options, delimited in the [LockManager](#)

documentation. While the lock manager is written to the ODMG API, the same locking rules apply at the OTM layer.

## 5.5. Reference Guides

### 5.5.1. Reference Guides Summary

Summary and explanation of the OJB reference guides.

- [PB guide](#)  
This document explains specific usage of the *PB API*.
- [ODMG guide](#)  
This document explains specific usage of the *ODMG API*.
- [OJB Queries](#)  
This document explains the usage of the query syntax.
- [Basic O/R Technique](#)  
This tutorial explains basic object-relational mapping technique in OJB like 1:1, 1:n and m:n relations, the auto-xxx settings for references and proxy objects/collections.
- [Platforms](#)  
What OJB requires from relational databases, and how to let it know which database to use.
- [Logging configuration](#)  
Detailed information about configuring the logging within OJB.
- [OJB.properties configuration](#)  
The details on how to modify OJB's behaviour. This includes changing pluggable components.
- [JDBC Types](#)  
This document explains the standard mapping of JDBC types to Java classes.
- [Repository Metadata](#)  
The specific details of OJB metadata.
- [Advanced O/R Technique](#)  
This document explains some advanced O/R techniques like Polymorphism and "OJB Extents", Mapping Inheritance Hierarchies, Nested Objects and so on.
- [Metadata Handling](#)  
This document explains how the metadata xml file work and how the metadata information can be modified at runtime.
- [Deployment](#)  
Specifics on what is required to deploy OJB, including deployment to EJB containers.
- [Connection Handling](#)  
This document explains how OJB handles the `Connection` instances and how this can be user-configured.
- [Caching](#)  
Documentation on the different object caching implementations and strategies included with OJB.
- [The Sequence Manager](#)  
How to use different sequence management strategies with OJB.
- [Locking](#)  
The optimistic and pessimistic locking capabilities of OJB.
- [OJB XDoclet Module](#)  
Documentation for the OJB XDoclet module. The module can build mappings and schema.
- [OJB Performance](#)  
A look at how OJB performs and how to use OJB's performance tests.

### 5.5.2. PB-api Guide

### 5.5.2.1. Introduction

The *PersistenceBroker API (PB-api)* provides the lowest level access to OJB's persistence engine. While it is a low-level API compared to the standardised ODMG or JDO API's it is still very straightforward to use.

The core class in the PersistenceBroker API is the [org.apache.ojb.broker.PersistenceBroker](#) class. This class provides the point of access for all persistence operations in this API.

This document is not a [PB tutorial](#) (newbies please read the tutorial first) rather than a guide showing the specific usage and possible pitfalls in handling the PB-api.

If you don't find an answer for a specific question, please have a look at the [FAQ](#) and the other [reference guides](#).

### 5.5.2.2. How to access the PB-api?

The `org.apache.ojb.broker.PersistenceBrokerFactory` make several methods available:

```
public PersistenceBroker createPersistenceBroker(PBKey key) throws
PBFactoryException;

public PersistenceBroker createPersistenceBroker(String jcdAlias, String user,
String password)
    throws PBFactoryException;

public PersistenceBroker defaultPersistenceBroker() throws PBFactoryException;
```

Method `defaultPersistenceBroker()` can be used if the attribute [default-connection](#) is set *true* in *jdbc-connection-descriptor*. It's a convenience method, useful when only one database is used.

The standard way to lookup a broker instance is via `org.apache.ojb.broker.PBKey` by specify *jcdAlias* (defined in the *jdbc-connection-descriptor* of the [repository file or sub file](#)), *user* and *passwd*. If the user and password is already set in *jdbc-connection-descriptor* it is possible to lookup the broker instance only be specify the *jcdAlias* in `PBKey`:

```
PBKey pbKey = new PBKey("myJcdAliasName", "user", "password");
// alternative if user/passwd set in configuration file
PBKey pbKey = new PBKey("myJcdAliasName");
PersistenceBroker broker =
PersistenceBrokerFactory.createPersistenceBroker(pbKey);
```

See further in FAQ ["Needed to put user/password of database connection in repository file?"](#).

### 5.5.2.3. Notes on Using the PersistenceBroker API

#### Exception Handling

The exception handling is described in the PB-tutorial [exception handling section](#).

#### Management of PersistenceBroker instances

There is no need to cache or pool the used [PersistenceBroker](#) instances, because OJB itself use a PB-pool. The configuration of the PB-pool is adjustable in the [OJB.properties](#) file.

Using the `PersistenceBroker.close()` method releases the broker back to the pool under

the default implementation. For this reason the examples in the [PB tutorial](#) all retrieve, use, and close a new broker for each logical transaction.

Apart from the pooling management `PersistenceBroker.close()` force the internal cleanup of the used broker instance - e.g. removing of temporary `PersistenceBrokerListener` instances, release of used connection if needed, internal used object registration lists, ... Therefore it's not recommended always refer to the same PB instance without closing it.

## Transactions

Transactions in the PersistenceBroker API are *database level transactions*. This differs from *object level transactions* used by e.g. the [odmg-api](#). The broker does not maintain a collection of modified, created, or deleted objects until a commit is called -- it operates on the database using the databases transaction mechanism. If object level transactions are required, one of the higher level API's (ODMG, JDO, or OTM) should be used.

## 5.5.2.4. Questions

### How to use multiple Databases

For each database define a [jdbc-connection-descriptor](#) same way as described in the [FAQ](#).

Now each database will be accessible via the `PersistenceBrokerFactory` using a `PBKey` matching the defined `jcdAlias` name as shown in section [How to access the PB-api?](#).

### Hook into OJB - PB-Listener and Instance Callbacks

All *Listener* and instance callback interfaces supported by the *PB-api* can be used in the top-level API (like [ODMG-api](#)) as well.

The OJB Kernel supports three types of "hook" into OJB:

- [PersistenceBrokerAware](#)

A callback interface used by persistence capable objects (the object class is declared in OJB [metadata mapping](#)) to be aware of `PersistenceBroker` operations on itself.

More detailed information can be found in the [Advanced-Technique Guide](#).

- [PBStateListener](#)

The listener interface for receiving [PersistenceBroker](#) state changes.

- [PBLifeCycleListener](#)

The listener interface for receiving *persistent object* life cycle information. This interface is intended for non persistent objects which want to track persistent object life cycle. Persistence capable objects can implement [PersistenceBrokerAware](#) - see above.

To add a *PBListener* use one of the following [PersistenceBroker](#) methods:

```
public void addListener(PBListener listener) throws PersistenceBrokerException;
public void addListener(PBListener listener, boolean permanent) throws PersistenceBrokerException;
```

The first method adds a **temporary** `org.apache.ojb.broker.PBListener` to the current `PersistenceBroker` instance - on `PersistenceBroker.close()` call the listener was removed.

The second one allows to add **permanent** `org.apache.ojb.broker.PBListener` when the method argument is set *true*. If set *false* the listener only be temporary added.

**Note:**

Be carefully when adding **permanent** listener, keep in mind you don't know which instance was returned next time from the pool, with a permanent listener or without!

To guarantee that any pooled broker instance use the permanent listener, best way is to extend the used `org.apache.ojb.broker.core.PersistenceBrokerFactoryIF` implementation and add the listener at creation of the `PersistenceBroker` instances.

Or add each time you lookup a `PersistenceBroker` instance the listener as a temporary listener.

### 5.5.3. ODMG-api Guide

#### 5.5.3.1. Introduction

The *ODMG API* is an implementation of the [ODMG 3.0 Object Persistence API](#). The ODMG API provides a higher-level API and [OQL query](#) language based interface over the [PersistenceBroker API](#).

This document is not a [ODMG tutorial](#) (newbies please read the tutorial first) rather than a guide showing the specific usage and possible pitfalls in handling the ODMG-api and the proprietary extensions by OJB.

If you don't find an answer for a specific question, please have a look at the [FAQ](#) and the other [reference guides](#).

Additionally the OJB's ODMG implementation has several extensions described [below](#).

#### 5.5.3.2. Specific Metadata Settings

To make OJB's *ODMG-api* implementation proper work, some specific metadata settings needed in the [repository mapping files](#).

All defined [reference-descriptor](#) and [collection-descriptor](#) need specific *auto-xxx* settings:

- `auto-retrieve="true"`
- `auto_update="none"`
- `auto-delete="none"` or `auto-delete="object"` (to enable cascading delete, since OJB 1.0.4!)

**Note:**

These settings are mandatory for proper work of the odmg-api!

So an example object mapping [class-descriptor](#) look like:

```
<class-descriptor
  class="org.apache.ojb.odmg.shared.Master"
  table="MDTEST_MASTER"
  >
  <field-descriptor
    name="masterId"
    column="MASTERID"
    jdbc-type="INTEGER"
    primarykey="true"
    autoincrement="true"
  />
  <field-descriptor
    name="masterText"
    column="MASTER_TEXT"
    jdbc-type="VARCHAR"
  />
  <collection-descriptor
    name="collDetailFKinPK"
    element-class-ref="org.apache.ojb.odmg.shared.DetailFKinPK"
```

```

        proxy="false"
        auto-retrieve="true"
        auto-update="none"
        auto-delete="none"
    >
    <inverse-foreignkey field-ref="masterId"/>
</collection-descriptor>
...
</class-descriptor>

```

A lot of mapping samples can be found in mappings for the [OJB test suite](#). All mappings for the ODMG unit test are in `repository_junit_odmg.xml` file, which can be found under the `src/test` directory.

### 5.5.3.3. How to access ODMG-api

Obtain a `org.odmg.Implementation` instance first, then create a new `org.odmg.Database` instance and open this instance by setting the used [jcd-alias](#) name:

```

Implementation odmg = OJB.getInstance();
Database database = odmg.newDatabase();
database.open("jcdAliasName#user#password", Database.OPEN_READ_WRITE);

```

The `user` and `password` separated by `#` hash only needed, when the user/passwd is not specified in the connection metadata (`jdbc-connection-descriptor`).

The [jdb-connection-descriptor](#) may look like:

```

<jdb-connection-descriptor
    jcd-alias="jcdAliasName"
    ...
    username="user"
    password="password"
    ...
>
...
</jdb-connection-descriptor>

```

With method call `OJB.getInstance()` always a **new** [org.odmg.Implementation](#) instance will be created and `odmg.newDatabase()` returns a new `Database` instance.

For best performance it's recommended to share the [Implementation](#) instance across the application. To get the current open database from the `Implementation` instance, use method `Implementation.getDatabase(null)`

```

Implementation odmg = ....
// get current used database
Database database = odmg.getDatabase(null);

```

Or share the open `Database` instance as well.

See further in FAQ "[Needed to put user/password of database connection in repository file?](#)".

### 5.5.3.4. Configuration Properties

The OJB *ODMG-api* implementation has some adjustable properties and pluggable components. All configuration properties can be set in the [OJB.properties](#) file.

Here are all properties used by OJB's *ODMG-api* implementation:

Property Name	Description
<code>OqlCollectionClass</code>	This entry defines the collection type returned from <a href="#">OQL queries</a> . By default this value is set to a List implementation. This will be suffice in

	<p>most situations.</p> <p>If you want to use the additional features of the DList interface (DList itself is persistable, support of <i>predicates</i>) directly on query results, change setting to the DList implementation (See also property 'DListClass' entry).</p> <p>But this will affect the performance - especially for large result sets. So recommended way is create <code>DCollection</code> instances only when needed (e.g. by converting a List result set to a DList).</p> <p>Important note: The collection class to be used MUST implement the interface <a href="http://org.apache.ojb.broker.ManageableCollection">org.apache.ojb.broker.ManageableCollection</a>. More info about <a href="#">implementing OJB collection types here</a>.</p>
ImplementationClass	Specifies the used base class for the <i>ODMG API</i> implementation. In managed environments a specific class is needed to potentiate JTA integration of OJB's ODMG implementation.
OJBTxManagerClass	Specifies the class for transaction management. In managed environments a specific class is needed to potentiate JTA integration of OJB's ODMG implementation.
ImplicitLocking	<p>This property defines the <i>implicit locking</i> behavior. If set to <i>true</i> OJB implicitly locks objects to ODMG transactions after performing OQL queries or when do a single lock on an object using <code>Transaction#lock(...)</code> method.</p> <p>If implicit locking is used locking objects is recursive, that is associated objects are also locked.</p> <p>If <code>ImplicitLocking</code> is set to <i>false</i>, no locks are obtained in OQL queries and there is also no recursive locking when do single lock on an object.</p>
LockAssociations	<p>This property was only used when <i>ImplicitLocking</i> is enabled. It defines the behaviour for the OJB implicit locking feature. If set to <i>true</i> acquiring a write-lock on a given object x implies write locks on all objects associated to x.</p> <p>If set to <i>false</i>, in any case implicit read-locks are acquired. Acquiring a read- or write lock on x thus allways results in implicit read-locks on all associated objects.</p>
Ordering	<p>Enable/Disable OJB's persistent object ordering algorithm on commit of a transaction. If enabled OJB try to calculate a valid order for all new/modified objects (and referenced objects).</p> <p>If the used databases support 'deferred checks' it's recommended to use this feature and to disable OJB's object ordering.</p>



	<div style="border: 1px solid black; background-color: #e0e0e0; padding: 5px;"> <p><b>Note:</b> This setting can be changed at runtime using OJB's <a href="#">ODMG extensions</a>.</p> </div>
<i>ImplicitLockingBackward</i>	A <b>@deprecated</b> property only for backward compatibility with older versions (before 1.0.4). If set <i>true</i> the behavior of method <code>ImplementationImpl#setImplicitLocking(...)</code> will be the same as in OJB in 1.0.3 or earlier (set the implicit locking behavior of the current used transaction) and disable the new possibility of global 'implicit locking' setting at runtime with <code>ImplementationExt#setImplicitLocking</code> . This is only for backward compatibility and will be removed at a later date.
DListClass	The used <code>org.odmg.DList</code> implementation class.
DArrayClass	The used <code>org.odmg.DArray</code> implementation class.
DMapClass	The used <code>org.odmg.DMap</code> implementation class.
DBagClass	The used <code>org.odmg.DBag</code> implementation class.
DSetClass	The used <code>org.odmg.DSet</code> implementation class.

### 5.5.3.5. OJB Extensions of ODMG

This section describes the proprietary extension of the *ODMG-api* provided by OJB.

#### The ImplementationExt Interface

The OJB extension of the odmg [Implementation](#) interface is called [ImplementationExt](#) and provide additional methods missed in the standard class definition.

- `get/setOqlCollectionClass`  
Use this methods to change the used OQL query result class at runtime. Description can be found in [Configuration Properties](#) section and in javadoc of [ImplementationExt](#).
- `is/setImpliciteWriteLocks`  
Use this methods to global change the associated locking type at runtime when implicit locking is used. Description can be found in [Configuration Properties](#) section and in javadoc of [ImplementationExt](#).
- `is/setOrdering`  
Use this methods to global enable/disable OJB's object ordering algorithm. Description can be found in [Configuration Properties](#) section and in javadoc of [ImplementationExt](#).

#### The TransactionExt Interface

The OJB extension of the odmg [Transaction](#) interface is called [TransactionExt](#) and provide additional methods missed in the standard class definition.

- `markDelete`  
Description can be found in javadoc of [TransactionExt](#).

- `markDirty`  
Description can be found in javadoc of [TransactionExt](#).
- `flush`  
Description can be found in javadoc of [TransactionExt](#).
- `is/setImplicitLocking`  
Description can be found in javadoc of [TransactionExt](#).
- `is/setOrdering`  
Description can be found in javadoc of [TransactionExt](#).
- `setCascadingDelete`  
Description can be found in javadoc of [TransactionExt](#).
- `getBroker()`  
Returns the current used broker instance. Usage example is [here](#).

### The EnhancedOQLQuery Interface

The OJB extension of the odmng [OQLQuery](#) interface is called [EnhancedOQLQuery](#) and provide additional methods missed in the standard class definition.

- `create(String queryString, int startAtIndex, int endAtIndex)`  
Description can be found in javadoc of [EnhancedOQLQuery](#).

### Access the PB-api within ODMG

As the [PB-api](#) was used by OJB's *ODMG-api* implementation, thus it is possible to get access of the used `PersistenceBroker` instance using the extended Transaction interface class [TransactionExt](#):

```
Implementation odmng = ...;
TransactionExt tx = (TransactionExt) odmng.newTransaction();
tx.begin();
...
PersistenceBroker broker = tx.getBroker();
// do work with broker
...
tx.commit();
```

It's mandatory that the used `PersistenceBroker` instance **never** be closed with a `PersistenceBroker.close()` call or be committed with `PersistenceBroker.commitTransaction()`, this will be done internally by the ODMG implementation.

### 5.5.3.6. Notes on Using the ODMG API

#### Transactions

The ODMG API uses *object-level transactions*, compared to the `PersistenceBroker` *database-level transactions*. An ODMG [Transaction](#) instance contains all of the changes made to the object model within the context of that transaction, and will not commit them to the database until the ODMG `Transaction` is committed. At that point it will use a database transaction (the underlying PB-api) to ensure atomicity of its changes.

#### Locks

The ODMG specification includes several levels of locks and isolation. These are explained in much more detail in the [Locking](#) documentation.

In the ODMG API, locks obtained on objects are locked within the context of a transaction. Any object modified within the context of a transaction will be stored with the transaction, other

changes made to the same object instance by other threads, ignoring the lock state of the object, will also be stored - so take care of locking conventions.

The ODMG locking conventions (obtain a write lock before do any modifications on an object) ensure that an object can only be modified within the transaction.

It's possible to configure OJB's ODMG implementation to support implicit locking with *WRITE* locks. Then a write lock on an object forces OJB to obtain implicit write locks on all referenced objects. See [configuration properties](#).

### Persisting Non-Transactional Objects

Frequently, objects will be modified outside of the context of an ODMG transaction, such as a data access object in a web application. In those cases a persistent object can still be modified, but not directly through the *OMG ODMG specification*. OJB provides an extension to the ODMG specification for instances such as this. Examine this code:

```
public static void persistChanges(Product product)
{
    Implementation impl = OJB.getInstance();
    TransactionExt tx = (TransactionExt) impl.newTransaction();

    tx.begin();
    tx.markDirty(product);
    tx.commit();
}
```

In this function the product is modified outside the context of the transaction, and is then the changes are persisted within a transaction. The `TransactionExt.markDirty()` method indicates to the Transaction that the passed object has been modified, even if the Transaction itself sees no changes to the object.

### 5.5.3.7. ODMG Named Objects

Using *named objects* allows to persist all serializable objects under a specified name. The methods to handle *named objects* are:

```
/**
 * Associate a name with an object and make it persistent.
 * An object instance may be bound to more than one name.
 * Binding a previously transient object to a name makes that object persistent.
 * @param object The object to be named.
 * @param name The name to be given to the object.
 * @exception org.odmg.ObjectNameNotUniqueException
 * If an attempt is made to bind a name to an object and that name is already
bound
 * to an object.
 */
public void bind(Object object, String name) throws
ObjectNameNotUniqueException;

/**
 * Lookup an object via its name.
 * @param name The name of an object.
 * @return The object with that name.
 * @exception ObjectNameNotFoundException There is no object with the specified
name.
 * @see ObjectNameNotFoundException
 */
public Object lookup(String name) throws ObjectNameNotFoundException;

/**
 * Disassociate a name with an object
 * @param name The name of an object.
```

```
* @exception ObjectNameNotFoundException No object exists in the database with
that name.
*/
public void unbind(String name) throws ObjectNameNotFoundException;
```

To use this feature a internal table and metadata mapping is madatory (by default these settings are enabled in OJB). More information about the needed internal tables see in [Platform Guide](#).

If the object to bind is a persistence capable object (the object class is declared in OJB [metadata mapping](#)), then the object will be persisted (if needed) dependent on the declared [metadata mapping](#) and the *named object* will be a **link** to the real persisted object.

On unbind of the *named object* only the link of the persistent object will be removed, the persistent object itself will be untouched.

If the object to bind is a serializable non-persistence capable object, the object will be serialized and persisted under the specified name.

On unbind the serialized object will be removed.

## Examples

In [OJB test-suite](#) is a test case called `org.apache.ojb.odmg.NamedRootsTest` which shown similar examples as below, but more detailed.

### 1. Persist a serializable object as named object

We want to persist a name list of all planets:

```
Transaction tx = odmg.newTransaction();
tx.begin();
List planets = new ArrayList();
example.add("Mercury");
example.add("Venus");
example.add("Earth");
...
database.bind(planets, "planet-list");
tx.commit();
```

The specified `List` with all planet names will be serialized and persisted as VARBINARY object.

To lookup the persisted list of the solar system planets:

```
Transaction tx = odmg.newTransaction();
tx.begin();
List planets = (List) database.lookup("planet-list");
tx.commit();
```

To remove the persistent list do:

```
Transaction tx = odmg.newTransaction();
tx.begin();
database.unbind("planet-list");
tx.commit();
```

### 2. Persist a persistence capable object as named object

We want to create a *named object* representing a persistence capable `Article` object (`Article` class is declared in OJB [metadata mapping](#)):

```
Transaction tx = odmg.newTransaction();
tx.begin();
// get existing or a new Article object
Article article = ....
database.bind(article, "my-article");
tx.commit();
```

OJB first checks if the specified `Article` object is already persisted - if not it will be persisted.

Then based on the `Article` object [Identity](#) the *named object* will be persisted. So the persistent *named object* is a link to the persistent real `Article` object.

On lookup of the *named object* the real `Article` instance will be returned:

```
Transaction tx = odmgt.newTransaction();
tx.begin();
Article article = (Article) database.lookup("my-article");
tx.commit();
```

On unbind of the *named object* only the link to the real `Article` object will be removed, the `Article` itself will not be touched.

To remove the *named object* and the `Article` instance do:

```
tx.begin();
// this only remove the named object link, the Article object
// itself will not be touched
database.unbind("my-article");
// thus delete the object itself too
database.deletePersistent(article);
tx.commit();
```

### 3. Persist a collection of persistence capable object as named object

We want to persist a list of the last shown `Article` objects. The `Article` class is a persistence capable object (declared in OJB [metadata mapping](#)). Thus we don't want to persist a serialized List of `Article` objects (because the real `Article` object may change), as shown in [example 1](#), rather we want to persist a List that links to the real persistent `Article` objects.

This is possible when the ODMG [DCollections](#) are used:

```
// get the list with last shown Article objects
List lastArticles = ...
Transaction tx = odmgt.newTransaction();
tx.begin();
// obtain new DList instance from Implementation class
DList namedArticles = odmgt.newDList();
// push Articles to DList
namedArticles.addAll(lastArticles);
database.bind(namedArticles, "last-shown");
tx.commit();
```

In this case OJB first checks for transient `Article` objects and make these new objects persistent, then based on the `Article` object [Identity](#) the *named object* will be persisted. So the persistent *named object* is in this case a list of links to persistent `Article` objects.

On `database.lookup("last-shown")` the `DList` will be returned and when access the list entries the `Article` objects will be materialized.

To remove the *named object* some more attention is needed:

```
tx.begin();
DList namedArticles = ...
// we want to completely remove the named object
// the persisted DList with all DList entries,
// but the Article objects itself shouldn't be deleted:
// 1. mandatory, clear the list to remove all entries
namedArticles.clear();
// 2. unbind named object
database.unbind("last-shown");
tx.commit();
```

After this the *named object* will be completely removed, but all `Article` object will be untouched.

### 5.5.3.8. ODMG's DCollections

The ODMG api declare some specific extensions of the `java.util.Collection` interface:

- `org.odmg.DList`
- `org.odmg.DSet`
- `org.odmg.DBag`
- `org.odmg.DMap`
- `org.odmg.DArray`

The ODMG [Implementation](#) class provide methods to get new instances of these classes.

In OJB all associations between persistence capable classes are declared in the [mapping files](#) and *1:n* and *m:n* relations can use any collection type class which implement the specific interface [ManageableCollection](#).

So there is no need to use the ODMG specific collection classes in object relations or when oql-queries are performed (more detailed info see '[oql collection class setting](#)').

One difference to *normal* collection classes is that `DCollection` implementation classes are persistence capable classes itself. This means that they can be persisted - e.g. see [named objects example](#). Mandatory is that all containing objects are persistence capable itself.

When persisting a `DCollection` object OJB first lock the collection entries, then the collection itself was locked. On commit the collection entries will be handled in a *normal* way and for each entry a *link object* (containing the [Identity](#) of the persistence capable object) is persisted.

When lookup the persisted `DCollection` object the *link objects* are materialized and on access the collection entry will be materialized by the identity.

### 5.5.3.9. Foreign Keys Constraints and ODMG-api

If cross-referenced database tables are used it's recommended to set *foreign key constraints* to guarantee database consistency. The consequence of using *foreign key constraints* is that the order of the persistence capable objects on *insert* and *delete* operations will become crucial.

Some databases support *deferred constraint checks*, this can help to avoid foreign key issues.

On transaction commit (using standard settings) OJB try to order the objects by itself. If this doesn't suffice it's possible to determine the [object order "by hand"](#).

If foreign key constraint violations arise when using [1:1 references](#) and [circular/bidirectional 1:1 references](#) it's possible to use a workaround introduced in version 1.0.4 to specify the database FK constraint in OJB using a [custom attribute](#) named '*constraint*':

```
<reference-descriptor name="refAA"
  class-ref="org.apache.ojb.odmg.CircularTest$ObjectAA"
  proxy="false"
  auto-retrieve="true"
  auto-update="none"
  auto-delete="none"
>
  <foreignkey field-ref="fkId"/>
  <attribute attribute-name="constraint" attribute-value="true"/>
</reference-descriptor>
```

### 5.5.3.10. Questions and Tips

**Disable OJB's object ordering, determine object order "by hand"**

By default OJB try to order all persistent objects on transaction commit call to avoid ordering problems. If this is not needed or helpful it can be disabled in two ways.

In most cases it's needed to disable *implicite locking* too, because it will lock/register dependend objects (e.g. 1:n references) automatically. First in [OJB.properties](#) file:

```
# Enable/Disable OJB's persistent object ordering algorithm on commit
# of a transaction. If enabled OJB try to calculate a valid order for
# all new/modified objects (and referenced objects).
# If the used databases support 'deferred checks' it's recommended to use this
# feature and to disable OJB's object ordering.
# This setting can be changed at runtime using OJB's ODMG extensions.
Ordering=false
```

Second at runtime, using OJB's ODMG extension classes [ImplementationExt](#) (global setting) and [TransactionExt](#) (per tx setting).

```
TransactionExt tx = (TransactionExt) odmg.newTransaction();
tx.begin();
...
/*
we want to manually insert new object, so we disable
OJB's ordering and implicit object locking
*/
tx.setOrdering(false);
tx.setImplicitLocking(false);
...
tx.commit();
```

### Circular- and Bidirectional References

The good news, OJB can handle *bidirectional*- and *circular*- references. When using [foreign key constraints](#) for referential integrity in these cases you have to pay attention.

In OJB [test-suite](#) a unit test called `org.apache.ojb.odmg.CircularTest` can be found. The tests show the handling of circular/bidirectional references and the possibilities how to handle object insert/update/delete ordering on transaction commit.

### I don't like OQL, can I use the PersistenceBroker Queries within ODMG

Yes you can! The ODMG implementation relies on PB Queries internally! Several users (including myself) are doing this.

If you have a look at the simple example below you will see how OJB Query objects can be used withing ODMG transactions.

The most important thing is to lock all objects returned by a query to the current transaction before starting manipulating these objects.

Further on do not commit or close the obtained PB-instance, this will be done by the ODMG transaction on `tx.commit()` / `tx.rollback()`.

```
TransactionExt tx = (TransactionExt) odmg.newTransaction();
tx.begin();
...
// cast to get intern used PB instance
PersistenceBroker broker = tx.getBroker();
...
// build query
QueryByCriteria query = ...
// perform PB-query
Collection result = broker.getCollectionByQuery(query);
// use result
...
```

```
tx.commit();
...
```

Note: Don't close or commit the used broker instance, this will be done by the odmng-api.

### How to use multiple Databases

For each database define a [jdbc-connection-descriptor](#) same way as described in the [FAQ](#).

Now it is possible to

- access the databases one after another, by closing the current used Database instance and by open a new one.

```
// get current used database instance
Database database = ...;
// close it
database.close();
// open a new one
database = odmng.newDatabase();
database.open("jcdAliasName#user#password", Database.OPEN_READ_WRITE);
...
```

The Database.close() call close the current used Database instance, after this it is possible to open a new database instance.

- use multiple databases in parallel, by creating a separate Implementation and Database instance for each [jdbc-connection-descriptor](#) defined in the mapping metadata.

```
Implementation odmng_1 = OJB.getInstance();
Database database_1 = odmng.newDatabase();
database.open("db_1#user#password", Database.OPEN_READ_WRITE);

Implementation odmng_2 = OJB.getInstance();
Database database_2 = odmng.newDatabase();
database.open("db_2#user#password", Database.OPEN_READ_WRITE);
```

Now it's possible to use both databases in parallel.

#### Note:

OJB does not provide distributed transactions by itself. To use distributed transactions, OJB have to be [integrated in an j2ee conform environment](#) (or made work with an JTA/JTS implementation).

## 5.5.4. Platforms

### 5.5.4.1. How to use OJB with a specific relational database

OJB has been designed to smoothly integrate with any relational database that provides JDBC support. OJB can be configured to use only JDBC 1.0 API calls to avoid problems with restrictions of several JDBC drivers.

It uses a limited SQL subset to avoid problems with restrictions of certain RDBMS. This design allows to keep the OJB code generic and free from database specifics.

This document explains basic concepts and shows how OJB can be configured to run against a specific RDBMS.

If you not already have done so, then you also might want to have a look at the [Getting Started](#) section which presents a sample skeleton project.

### 5.5.4.2. Basic Concepts



## OJB internal tables

For certain features OJB relies on several internal tables that must be present in the target RDBMS to allow a proper functioning. The associated internal object [metadata mapping](#) of these internal used tables can be found in [repository internal.xml](#) file.

If those features are not needed/used OJB can be safely run without any internal tables and metadata mapping.

The following table lists all tables and their specific purpose.

Table	Purpose
<pre>OJB_HL_SEQ &lt;table name="OJB_HL_SEQ" description="HIGH/LOW SequenceManager table"&gt;   &lt;column name="TABLENAME" required="true" primaryKey="true" type="VARCHAR" size="250" /&gt;   &lt;column name="MAX_KEY" type="BIGINT" /&gt;   &lt;column name="GRAB_SIZE" type="INTEGER" /&gt;   &lt;column name="VERSION" type="INTEGER" /&gt; &lt;/table&gt;</pre>	<p>Table for the high/low sequence manager. The column <i>TABLENAME</i> was used to persist the sequence name (may be re-named in further versions of OJB).</p> <p>If the built-in OJB sequence manager is not used, this table is not needed.</p>
<pre>OJB_NRM &lt;table name="OJB_NRM" description="OJB NAMED ROOTS Table"&gt;   &lt;column name="NAME" required="true" primaryKey="true" type="VARCHAR" size="250" /&gt;   &lt;column name="OID_" type="LONGVARBINARY" /&gt; &lt;/table&gt;</pre>	<p>The "Named Roots Map". ODMG allows to bind persistent objects to a user defined name - called <a href="#">named objects</a>.</p> <p>The Named roots map is used to store these bindings. It has NAME (String of arbitrary length) as primary key and keeps the serialized OID of a persistent object or an arbitrary serialized object in the field OID (String of arbitrary length).</p> <p>If Database.bind(...) and Database.lookup(...) are not used in client apps, this table is not needed.</p>
<pre>OJB_DLIST &lt;table name="OJB_DLIST" description="DLIST IMPLEMENTATION"&gt;   &lt;column name="ID" required="true" primaryKey="true" type="INTEGER" /&gt;   &lt;column name="SIZE_" type="INTEGER" /&gt; &lt;/table&gt;</pre>	<p>The table used for the ODMG persistent DList collections.</p> <p>If ODMG DLists are not used, this table is not needed.</p>
<pre>OJB_DLIST_ENTRIES &lt;table name="OJB_DLIST_ENTRIES" description="DList entry table"&gt;   &lt;column name="ID" required="true" primaryKey="true" type="INTEGER" /&gt;   &lt;column name="DLIST_ID" required="true" type="INTEGER" /&gt;   &lt;column name="POSITION_" type="INTEGER" /&gt;   &lt;column name="OID_" type="LONGVARBINARY" /&gt; &lt;/table&gt;</pre>	<p>stores the entries of DLists (a wrapper to objects stored in the DList)</p> <p>If ODMG DLists are not used, this table is not needed.</p>

<p>OJB_DSET</p> <pre>&lt;table name="OJB_DSET" description="DSET IMPLEMENTATION"&gt;   &lt;column name="ID" required="true" primaryKey="true"   type="INTEGER"/&gt;   &lt;column name="SIZE_" type="INTEGER"/&gt; &lt;/table&gt;</pre>	<p>The table used to store ODMG persistent DSET collections If ODMG DSets are not used, this table is not needed.</p>
<p>OJB_DSET_ENTRIES</p> <pre>&lt;table name="OJB_DSET_ENTRIES" description="DSet entry table"&gt;   &lt;column name="ID" required="true" primaryKey="true"   type="INTEGER"/&gt;   &lt;column name="DLIST_ID" required="true" type="INTEGER"/&gt;   &lt;column name="POSITION_" type="INTEGER"/&gt;   &lt;column name="OID_" type="LONGVARBINARY"/&gt; &lt;/table&gt;</pre>	<p>This table stores the entries of DSets. If ODMG DSets are not used, this table is not needed.</p>
<p>OJB_DMAP</p> <pre>&lt;table name="OJB_DMAP" description="DMap table"&gt;   &lt;column name="ID" required="true" primaryKey="true"   type="INTEGER"/&gt;   &lt;column name="SIZE_" type="INTEGER"/&gt; &lt;/table&gt;</pre>	<p>The table use to store the ODMG persistent DMap tables If ODMG DMaps are not used, this table is not needed.</p>
<p>OJB_DMAP_ENTRIES</p> <pre>&lt;table name="OJB_DMAP_ENTRIES" description="DMap entry table"&gt;   &lt;column name="ID" required="true" primaryKey="true"   type="INTEGER"/&gt;   &lt;column name="DMAP_ID" required="true" type="INTEGER"/&gt;   &lt;column name="KEY_OID" type="VARBINARY"/&gt;   &lt;column name="VALUE_OID" type="VARBINARY"/&gt; &lt;/table&gt;</pre>	<p>The table containing the DMap entries. The Keys and Values of the map can be arbitrary persistent objects. If ODMG DMaps are not used, this table is not needed.</p>

OJB uses [Torque](#) to create all required tables and data. Thus there is no SQL DDL file, but an XML file describing the tables in format readable by Torque. The Torque DDL information for the internal tables resides in the file `src/schema/ojbcore-schema.xml`.

The o/r mappings for these tables are contained in the file `repository_internal.xml`.

If you want to have a look at how these files could be used, have a look at the the [ojb-blank](#) sample project which is already prepared to use these files.

#### Tables for the regression testbed

It is recommended to run the [OJB test-suite](#) against your target database. Thus you will have to provide several more tables, filled with the proper testdata.

The DDL information for these tables resides in the file

The testdata is defined in the file `src/schema/objbtest-data.xml`.

The o/r mappings for these tables are contained in the file `repository_junit.xml`.

### Tables for the tutorial applications

If you intend to run the OJB tutorial applications against your target database you will have to provide one extra table.

The DDL information for this table also resides in the file `src/schema/objbtest-schema.xml`.

The testdata is also defined in the file `src/schema/objbtest-data.xml`.

The o/r mappings for this table is contained in the file `repository_user.xml`.

### 5.5.4.3. The setup process

OJB provides a setup routine to generate the target database and to fill it with the required testdata. This routine is based on Torque scripts and is driven from the `build.xml` file. This section describes how to use it.

#### Selecting a platform profile

OJB ships with support for several popular database platforms. The target platform is selected by the switch `profile` in the file `build.properties`. You can choose one out of the predefined profiles:

```
# With the 'profile' property you can choose the RDBMS platform OJB is using
# implemented profiles:
#
profile=hsqldb
# use the mssqldb-JSQLConnect profile for Microsoft SQL Server and
# you will automatically JSQLConnect driver, from http://www.j-netdirect.com/
# MBAIRD: This is my driver of preference for MS SQL Server, I find the OEM'd
# MS driver to have some problems.
#profile=mssqldb-JSQLConnect
#profile=mssqldb-Opta2000
#profile=mssqldb-ms
#profile=mysql
#profile=db2
#profile=oracle
#profile=oracle9i
#profile=oracle9i-Seropto
#profile=msaccess
#profile=postgresql
#profile=informix
#profile=sybase
#profile=sapdb
#profile=maxdb
```

The profile switch activated in `build.properties` is used to select a profile file from the `profile` directory.

If you set `profile=db2`, then the file `profile/db2.profile` is selected.

This file is used by the Torque scripts to set platform specific properties and to perform platform specific SQL operations.

#### editing the profile to point to your target db

The platform specific file `profile/xxx.profile` contains lots of information used by Torque. You can ignore most of it. The only important part in this file is the section where the url to the

target db is assembled, here is an snip of the DB2 profile:

```
# -----
#
#   D A T A B A S E   S E T T I N G S
#
# -----
# JDBC connection settings. This is used by the JDBCToXML task
# that will create an XML database schema from JDBC metadata.
# These settings are also used by the SQL Ant task to initialize
# your Turbine system with the generated SQL.
# -----

dbmsName = Db2
jdbcLevel = 1.0
urlProtocol = jdbc
urlSubprotocol = db2
urlDbalias = OJB

createDatabaseUrl = ${urlProtocol}:${urlSubprotocol}:${urlDbalias}
buildDatabaseUrl = ${urlProtocol}:${urlSubprotocol}:${urlDbalias}
databaseUrl = ${urlProtocol}:${urlSubprotocol}:${urlDbalias}
databaseDriver = COM.ibm.db2.jdbc.app.DB2Driver
databaseUser = admin
databasePassword = db2
databaseHost = 127.0.0.1
```

These settings result in a database URL `jdbc:db2:OJB`. If your production database is registered with the name `MY_PRODUCTION_DB` you have to edit the entry `urlDbalias` to:

```
urlDbalias = MY_PRODUCTION_DB.
```

In this section you can also set application user name and password. You can also enter a different jdbc driver class, to activate a different driver.

Before progressing, please check that the jdbc driver class, named in the `databaseDriver` entry is located on the classpath! You can either edit the global environment variable `CLASSPATH` or place the jdbc driver jar file into the `jakarta-ojb-xxx/lib` directory.

### Executing the build script

Now everything should be prepared to launch the setup routine. This routine can be invoked by calling `ant prepare-testdb`.

If you are prompted with a `BUILD SUCCESSFUL` message after some time, everything is OK.

If you are prompted with a `BUILD FAILED` message after some time, something went wrong. This may have several reasons:

- You entered some incorrect settings. Please check the log messages to see what went wrong.
- Torque does not work properly against your target database. Torque is very flexible and should be able to work against a wide range of databases. But the code templates for each database may not be accurate. Please check the ojb-user mailinglist archive if there are any failure reports for your specific database. Please also check if some contributed a fix already. If you don't find anything please post your problem to the ojb user-list.

As a last resort you can try the following: Switch back to the default `hsqldb` profile and execute `ant prepare-testdb` This will setup the default `hsqldb` database. And it will also generate SQL scripts that you may use to generate your database manually.

The SQL scripts are generated to `jakarta-ojb-xxx/target/src/sql`. You can touch these scripts to match your database specifics and execute them manually against your platform.

### Verifying the installation

Now everything is setup to run the junit regression tests against your target database.

Execute

```
ant junit
```

to see if everything works as expected. more information about the [OJB Test Suite here](#). If you did not manage to set up the target database with the `ant prepare-testdb` you can use `ant junit-no-compile-no-prepare` to run the testsuite without generation of the test database.

## 5.5.5. OJB.properties Configuration File

### 5.5.5.1. OJB Configuration

OJB provides two different configuration mechanisms:

1. An XML based `repository.xml` is used to define the Object/Relational Mapping. This Mapping is translated into a metadata dictionary at runtime. The metadata layer may also be manipulated at runtime through OJB API calls. [Follow this link to learn more about the XML repository](#).
2. A properties file `OJB.properties` that is responsible for the configuration of the OJB runtime environment. It contains information that does not change at runtime and does not contain O/R mapping related information.

The rest of this document details on this properties file.

### 5.5.5.2. OJB.properties File

By default this file is named `OJB.properties` and is loaded from the `classpath` by a J2EE compliant resource lookup:

```
Thread.currentThread().getContextClassLoader().getResource(getFilename());
```

The filename of the properties file can be changed by setting a Java system property. This can be done programmatically:

```
System.setProperty("OJB.properties", "myOwnPropertiesFile.props");
```

or by setting a `-D` option to the JVM:

```
java -DOJB.properties=myOwnPropertiesFile.props my.own.ojb.Application
```

All things that can be configured by `OJB.properties` are commented in the file itself. [Have a look at the default version of this file](#).

## 5.5.6. JDBC Types

### 5.5.6.1. Mapping of JDBC Types to Java Types

OJB implements the mapping conversions for JDBC and Java types as specified by the JDBC 3.0 specification (see JDBC 3.0 specification *Appendix B, Data Type Conversion Tables*). See the [table below](#) for details.

If a sql-java type mapping is needed, that doesn't match the java types defined by the specification, e.g. a field in the persistent object class is of type `int[]` and the DB type is `VARCHAR` or a `List` field have to be mapped to `VARCHAR` a [field-conversion](#) class can be used.

JDBC Type	Java Type
-----------	-----------

CHAR	String
VARCHAR	String
LONGVARCHAR	String
NUMERIC	java.math.BigDecimal
DECIMAL	java.math.BigDecimal
BIT	boolean
BOOLEAN	boolean
TINYINT	byte
SMALLINT	short
INTEGER	int
BIGINT	long
REAL	float
FLOAT	double
DOUBLE	double
BINARY	byte[]
VARBINARY	byte[]
LONGVARBINARY	byte[]
DATE	java.sql.Date
TIME	java.sql.Time
TIMESTAMP	java.sql.Timestamp
CLOB	Clob
BLOB	Blob
ARRAY	Array
DISTINCT	mapping of underlying type
STRUCT	Struct
REF	Ref
DATALINK	java.net.URL
JAVA_OBJECT	underlying Java class

### 5.5.6.2. Type and Value Conversions - The FieldConversion Interface

#### Introduction

A typical problem with O/R tools is mismatching datatypes: a class from the domain model has an attribute of type *boolean* but the corresponding database table stores this attribute in a column of type *BIT* or *int*.

This example explains how OJB allows you to define **FieldConversions** that do the proper translation of types and values.

The source code of this example is included in the OJB source distribution and resides in the test package `org.apache.ojb.broker`.

### The problem

The test class `org.apache.ojb.broker.Article` contains an attribute `isSelloutArticle` of type `boolean`:

```
public class Article implements InterfaceArticle
{
    protected int articleId;
    protected String articleName;

    // maps to db-column Auslaufartikel of type int
    protected boolean isSelloutArticle;

    ...
}
```

The corresponding table uses an `int` column (`Auslaufartikel`) to store this attribute:

```
CREATE TABLE Artikel (
    Artikel_Nr          INT PRIMARY KEY,
    Artikelname        CHAR(60),
    Lieferanten_Nr     INT,
    Kategorie_Nr       INT,
    Liefereinheit      CHAR(30),
    Einzelpreis        DECIMAL,
    Lagerbestand       INT,
    BestellteEinheiten INT,
    MindestBestand     INT,
    Auslaufartikel     INT
)
```

### The Solution

OJB allows to use predefined (or self-written) `FieldConversions` that do the appropriate mapping. The `FieldConversion` interface declares two methods: `javaToSql(...)` and `sqlToJava(...)`:

```
/**
 * FieldConversion declares a protocol for type and value
 * conversions between persistent classes attributes and the columns
 * of the RDBMS.
 * The default implementation does not modify its input.
 * OJB users can use predefined implementation and can also
 * build their own conversions that perform arbitrary mappings.
 * the mapping has to defined in the xml repository
 * in the field-descriptor.
 *
 * @author Thomas Mahler
 */
public interface FieldConversion extends Serializable
{
    /**
     * convert a Java object to its SQL
     * pendant, used for insert & update
     */
    public abstract Object javaToSql(Object source) throws ConversionException;

    /**
     * convert a SQL value to a Java Object, used for SELECT
     */
    public abstract Object sqlToJava(Object source) throws ConversionException;
}
```

The method `FieldConversion.sqlToJava()` is a callback that is called within the OJB

broker when Object attributes are read in from JDBC result sets. If OJB detects that a FieldConversion is declared for a persistent classes attributes, it uses the FieldConversion to do the marshalling of this attribute.

For the above mentioned problem of mapping an int column to a boolean attribute we can use the predefined FieldConversion Boolean2IntFieldConversion. Have a look at the code to see how it works:

```
public class Boolean2IntFieldConversion implements FieldConversion
{
    private static Integer I_TRUE = new Integer(1);
    private static Integer I_FALSE = new Integer(0);

    private static Boolean B_TRUE = new Boolean(true);
    private static Boolean B_FALSE = new Boolean(false);

    /**
     * @see FieldConversion#javaToSql(Object)
     */
    public Object javaToSql(Object source)
    {
        if (source instanceof Boolean)
        {
            if (source.equals(B_TRUE))
            {
                return I_TRUE;
            }
            else
            {
                return I_FALSE;
            }
        }
        else
        {
            return source;
        }
    }

    /**
     * @see FieldConversion#sqlToJava(Object)
     */
    public Object sqlToJava(Object source)
    {
        if (source instanceof Integer)
        {
            if (source.equals(I_TRUE))
            {
                return B_TRUE;
            }
            else
            {
                return B_FALSE;
            }
        }
        else
        {
            return source;
        }
    }
}
```

There are other helpful standard conversions defined in the package `org.apache.ojb.broker.accesslayer.conversions`: Of course it is possible to map between `java.sql.date` and `java.util.date` by using a Conversion. A very interesting Conversion is the `Object2ByteArrayFieldConversion` it allows to store inlined objects in varchar columns!



Coming back to our example, there is only one thing left to do: we must tell OJB to use the proper *FieldConversion* for the *Article* class. This is done in the XML repository file. The [field-descriptor](#) allows to define a *conversion* attribute declaring the fully qualified *FieldConversion* class:

```
<!-- Definitions for test.ojb.broker.Article -->
  <class-descriptor
    class="org.apache.ojb.broker.Article"
    proxy="dynamic"
    table="Artikel"
  >
    ...
    <field-descriptor
      name="isSelloutArticle"
      column="Auslaufartikel"
      jdbc-type="INTEGER"
      conversion="org.apache.ojb.broker.accesslayer.conversions.Boolean2IntFieldConversion"
    />
    ...
  </class-descriptor>
```

## 5.5.7. Repository File

### 5.5.7.1. Introduction - repository syntax

The syntax of the OJB repository xml files is defined by the *repository.dtd*.

An overview of all [repository.dtd-elements can be found here](#). The [repository.dtd can be found here](#).

The actual repository metadata declaration is split up into several separate files, here is an excerpt of the most important files:

1. [the repository.xml](#). Main file for metadata declaration. This file is split into several sub files using xml-Entity references.
2. [the repository\\_database.xml](#). This file contains the mapping information for database/connection handling.
3. [the repository\\_internal.xml](#). This file contains the mapping information for the OJB internal tables. These tables are used for implementing SequenceManagers and persistent collections.
4. [the repository\\_user.xml](#). This file contains mappings for the tutorial applications and may be used to hold further user defined class mappings.
5. [the repository\\_junit.xml](#). This file contains mapping information for common OJB JUnit regression test suite. In production environments these tables are not needed.
6. other repository\_junit\_XYZ.xml  
More specific junit test mapping. In production environments these tables are not needed.
7. There are some more files, for more information see comment in appropriate xml-file.

### 5.5.7.2. descriptor-repository

The *descriptor-repository* is the root element of a repository.xml file. It consists of one or more *jdbc-connection-descriptor* and at least one *class-descriptor* element. But it's also possible to startup OJB without any of these elements and [add them at runtime](#).

#### Elements

```
<!ELEMENT descriptor-repository (documentation?, attribute*,
  jdbc-connection-descriptor*, class-descriptor*)>
```

The *documentation* element can be used to store arbitrary information.

The *attribute* element allows to add [custom attributes](#), e.g. for passing arbitrary properties.

The *jdbc-connection-descriptor* element specifies a jdbc connection for the repository.

The *class-descriptor* element specify o/r mapping information for persistent class.

```
<!ELEMENT descriptor-repository (
  documentation?,
  attribute*,
  jdbc-connection-descriptor*,
  class-descriptor* )
>
```

### Attributes

The *version* attribute is used to bind a repository.xml file to a given version of this dtd. A given OJB release will work properly only with the repository version shipped with that release. This strictness maybe inconvenient but it does help to avoid the most common version conflicts.

The *isolation-level* attribute defines the default locking isolation level used by OJB's [pessimistic locking](#) api. All jdbc-connection-descriptor or class-descriptor that do not define a specific isolation level will use this.

Note: This does NOT touch the jdbc-level of the connection.

The *proxy-prefetching-limit* attribute specifies a default value to be applied to all proxy instances. If none is specified a default value of 50 is used. Proxy prefetching specifies how many instances of a proxied class should be loaded in a single query when the proxy is first accessed.

```
<!ATTLIST descriptor-repository
  version (1.0) #REQUIRED
  isolation-level (read-uncommitted | read-committed | repeatable-read |
                 serializable | optimistic | none) "read-uncommitted"
  proxy-prefetching-limit CDATA "50"
>
```

### 5.5.7.3. jdbc-connection-descriptor

The *jdbc-connection-descriptor* element specifies a jdbc connection for the repository. It is allowed to define more than one *jdbc-connection-descriptor*. All *class-descriptor* elements are independent from the *jdbc-connection-descriptors*. More info about [connection handling here](#).

### Elements

The *object-cache* element specifies the [object-cache implementation](#) class associated with this class.

A *connection-pool* element may be used to define connection pool properties for the specified JDBC connection.

Further a *sequence-manager* element may be used to define which sequence manager implementation should be used within the defined connection.

Use the [custom-attribute](#) element to pass implementation specific properties.

```
<!ELEMENT jdbc-connection-descriptor (documentation?, attribute*,
  object-cache?, connection-pool?, sequence-manager?)>
```

### Attributes

The *jdbc-connection-descriptor* element contains a bunch of required and implied attributes:

The *jcdAlias* attribute is a shortcut name for the defined connection descriptor. OJB uses the jcd alias as key for the defined connections.

The *default-connection* attribute used to define if this connection should be used as the default connection with OJB. You could define only one connection as the default connection. It is also possible to set the default connection at runtime using *PersistenceBrokerFactory#setDefaultKey(...)* method. If set *true* you can use a PB-api [shortcut-method](#) of the *PersistenceBrokerFactory* to lookup *PersistenceBroker* instances.

**Note:**

If *default-connection* is not set at runtime, it is mandatory that [username and password](#) is set in the repository file.

The *platform* attribute is used to define the specific RDBMS Platform. This attribute corresponds to a *org.apache.obj.broker.platforms.PlatformXXXImpl* class. Supported databases [see here](#). Default is *Hsqldb*.

The *jdbc-level* attribute is used to specify the Jdbc compliance level of the used Jdbc driver. Allowed values are: *1.0, 2.0, 3.0*. Default is *1.0*.

**DEPRECATED!**. The *eager-release* attribute is used to solve a problem that occurs when using OJB within JBoss (3.0 <= version < 3.2.2, seems to be fixed in jboss 3.2.2 and higher). Only use within JBoss. **DEPRECATED** attribute.

The *batch-mode* attribute allows to enable JDBC connection batch support (if supported by used database), 'true' value allows to enable per-session batch mode, whereas 'false' prohibits it. *PB.serviceConnectionManager.setBatchMode(...)* method can be used to switch on/off batch mode, if batch-mode is enabled. On *PB.close()* OJB switches off batch mode, thus you have to do '*...setBatchMode(true)*' on each obtained PB instance again.

**Note:**

OJB 1.0.4 and earlier:

When using [database identity columns](#) it's not allowed to enable *batch mode* for insert operations.

When using [optimistic locking](#) the version check will always succeed for update operations when *batch-mode* is enabled - take care!!.

This will be fixed and automatically handled by OJB till the next major release.

The *useAutoCommit* attribute allows to set how OJB uses the autoCommit state of the used connections. The default mode is 1. When using mode 0 or 2 with the PB-api, you must use PB transaction demarcation.

- 0 - OJB ignores the autoCommit setting of the connection and does not try to change it. This mode could be helpful if the connection won't let you set the autoCommit state (e.g. using datasources within an application server).
- 1 - [default mode] set the connection's *autoCommit* state temporary to 'false' if needed (when using transaction demarcation in a non-managed environment) and restore the 'old' state after use. In versions before *OJB 1.0.4* the autoCommit state was explicitly set 'true' when the connection was created, now OJB expects that this is done by the jdbc-driver/DataSource configuration. To enable the old behavior set a [custom attribute](#) *initializationCheck* to 'true'.

```
<attribute attribute-name="initializationCheck"
  attribute-value="false" />
```

Then OJB sets the *autoCommit* state explicitly to 'true' when the connection is created by the [ConnectionFactory](#).

- 2 - Set the connection's *autoCommit* explicitly to *false* when a connection is created.

If the *ignoreAutoCommitExceptions* attribute is set to *true*, all exceptions caused by setting the autoCommit state will be ignored. Default mode is *false*.

If a *jndi-datasource-name* for JNDI based lookup of Jdbc connections is specified, the following four attributes *driver*, *protocol*, *subprotocol*, and *dbalias* used for Jdbc DriverManager based

construction of Jdbc Connections must not be declared.

If a *jndi-datasource-name* is specified, OJB always assume that a JNDI based datasource connection lookup was expected (so take care that this attribute is empty or absent on *driver based* connection handling).

The *username* and *password* attributes are used as credentials for obtaining a jdbc connections. If users don't want to keep user/password information in the repository.xml file, they can pass user/password using a *PBKey* to obtain a PersistenceBroker. More info [see FAQ](#).

```
<!ATTLIST jdbc-connection-descriptor
  jcd-alias CDATA #REQUIRED
  default-connection (true | false) "false"
  platform ( Db2 | Hsqldb | Informix | MsAccess | MsSQLServer |
             MySQL | Oracle | PostgreSQL | Sybase | SybaseASE |
             SybaseASA | Sapdb | Firebird | Axion | NonstopSql |
             Oracle9i | MaxDB ) "Hsqldb"
  jdbc-level (1.0 | 2.0 | 3.0) "1.0"
  eager-release (true | false) "false"
  batch-mode (true | false) "false"
  useAutoCommit (0 | 1 | 2) "1"
  ignoreAutoCommitExceptions (true | false) "false"

  jndi-datasource-name CDATA #IMPLIED

  driver CDATA #IMPLIED
  protocol CDATA #IMPLIED
  subprotocol CDATA #IMPLIED
  dbalias CDATA #IMPLIED

  username CDATA #IMPLIED
  password CDATA #IMPLIED
>
```

#### Custom attributes

The *JdbcConnectionDescriptor* supports specific configuration properties via [custom-attributes](#).

Attribute *initializationCheck* is an attribute to support backward compatibility with OJB versions before 1.0.4.

In older versions OJB change the 'autoCommit' state dependent of the used 'useAutoCommit' attribute setting at connection initialization. This doesn't work in all situations/environments, thus for useAutoCommit="1" the ConnectionFactory does no longer set autoCommit to *true* on connection creation.

To use the old behavior (OJB version 1.0.3 or earlier) set this property to *true*, then OJB change the 'autoCommit' state (if needed) of new obtained connections at connection initialization.

If *false* or this property is removed, OJB dosen't try to change connection 'autoCommit' state at connection initialization.

Usage example of supported custom attributes:

```
<jdbc-connection-descriptor
  ...
>
  <attribute attribute-name="initializationCheck"
    attribute-value="false" />
  ...
</jdbc-connection-descriptor>
```

#### 5.5.7.4. connection-pool

The *connection-pool* element specifies the connection pooling and low-level JDBC driver parameters. Read more about OJB [connection handling](#).

## Elements

The *documentation* element can be used to store arbitrary information.

Use the *attribute* element to set JDBC-level properties or to enable DBCP PreparedStatement pooling if your JDBC driver does not have a PreparedStatement cache already.

See section [custom attributes](#) below for more information.

### Note:

When using an external DataSource, OJB cannot configure any JDBC-properties.

```
<!ELEMENT connection-pool ( documentation?, attribute* )>
```

## Attributes

*maxActive* (default=21) The maximum number of active connections that can be allocated from this pool at the same time, or zero for no limit.

*maxIdle* (default=-1) The maximum number of active connections that can remain idle in the pool, without extra ones being released, or zero for no limit.

*minIdle* (Since OJB 1.0.4, default=0) The minimum number of active connections that can remain idle in the pool, without extra ones being created, or zero to create none.

*maxWait* (default=5000) The maximum number of milliseconds that the pool will wait (when there are no available connections) for a connection to be returned before throwing an exception, or -1 to wait indefinitely.

Must be > 0 for timeout to actually happen in DBCP PoolingDataSource.

*whenExhaustedAction* (default=0)

- 0 - fail when pool is exhausted
- 1 - block when pool is exhausted
- 2 - grow when pool is exhausted

*validationQuery* (default=not specified) The SQL query that will be used to validate connections from this pool according to *testOnBorrow*/*testOnReturn*/*testWhileIdle*. If specified, this query **must** be an SQL SELECT statement that returns at least one row.

If not specified, only *connection.isClosed()* checks will be performed according to *testOnBorrow*/*testOnReturn*/*testWhileIdle*.

### Note:

Many database servers will discard idle connections after some time of inactivity. This timespan is usually configurable by the DBA and can range from anything between one hour and several days. Consider specifying a validation query that fits your database server and set at least *testOnBorrow=true*.

Example validation queries:

```
Oracle      SELECT 1 FROM DUAL
PostgreSQL  SELECT 1
MySQL       SELECT 1
```

*testOnBorrow* (default=true) The indication of whether connections will be validated before being borrowed from the pool. If the connection fails to validate, it will be dropped from the pool, and OJB will attempt to borrow another.

*testOnReturn* (default=false) The indication of whether connections will be validated before being

returned to the pool.

*testWhileIdle* (default=false) The indication of whether connections will be validated by the idle object evictor (if any). If a connection fails to validate, it will be dropped from the pool.

*timeBetweenEvictionRunsMillis* (default=-1) The number of milliseconds to sleep between runs of the idle object evictor thread. When non-positive, no idle object evictor thread will be run.

*numTestsPerEvictionRun* (default=10) The number of objects to examine during each run of the idle object evictor thread (if any).

Has no meaning if *timeBetweenEvictionRunsMillis* is non-positive.

*minEvictableIdleTimeMillis* (default=1800000) The minimum amount of time a connection may sit idle in the pool before it is eligible for eviction by the idle object evictor (if any).

When non-positive, no connection will be dropped from the pool due to idle time alone.

Has no meaning if *timeBetweenEvictionRunsMillis* is non-positive.

*removeAbandoned* [ConnectionFactoryDBCPIImpl] (default=false) Flag to remove abandoned connections if they exceed the *removeAbandonedTimeout*. If set to true a connection is considered abandoned and eligible for removal if it has been idle longer than the *removeAbandonedTimeout*. Setting this to true can recover db connections from poorly written applications which fail to close a connection.

If you have enabled "removeAbandoned" then it is possible that a connection is reclaimed by the pool because it is considered to be abandoned. This mechanism is triggered on *borrowObject* (ie in OJB when a PersistenceBroker gets a Connection) when:

(*numIdle* < 2) and (*numActive* > *maxActive* - 3)

For example *maxActive*=20, 18 active connections and 1 idle connection would trigger the

"removeAbandoned". But only the active connections that aren't used for more than

*removeAbandonedTimeout* seconds are removed. Traversing a resultset doesn't count as being used.

The abandoned object eviction takes place before normal *borrowObject* logic (there is no asynch evictor thread like for *testWhileIdle*).

*removeAbandonedTimeout* [ConnectionFactoryDBCPIImpl] (default=300) Timeout in seconds before an abandoned connection can be removed.

Has no meaning if *removeAbandoned* is false.

*logAbandoned* [ConnectionFactoryDBCPIImpl] (default=false) Flag to log stack traces for application code which abandoned a Statement or Connection.

**Note:**

Logging of abandoned Statements and Connections adds overhead for every Connection open or new Statement because a stack trace has to be generated.

```
<!ATTLIST connection-pool
  maxActive           CDATA #IMPLIED
  minIdle             CDATA #IMPLIED
  maxIdle             CDATA #IMPLIED
  maxWait             CDATA #IMPLIED
  minEvictableIdleTimeMillis CDATA #IMPLIED
  numTestsPerEvictionRun CDATA #IMPLIED
  testOnBorrow        ( true | false ) #IMPLIED
  testOnReturn        ( true | false ) #IMPLIED
  testWhileIdle       ( true | false ) #IMPLIED
  timeBetweenEvictionRunsMillis CDATA #IMPLIED
  whenExhaustedAction ( 0 | 1 | 2 ) #IMPLIED
  validationQuery     CDATA #IMPLIED

  removeAbandoned    ( true | false ) #IMPLIED
  removeAbandonedTimeout CDATA #IMPLIED
```

```
logAbandoned ( true | false ) #IMPLIED
>
```

### Custom attributes

OJB itself and the *ConnectionFactory* implementation classes support specific connection configuration properties, these properties can be set by using *custom-attributes*.

Usage example of supported custom attributes:

```
<connection-pool
  maxActive="30"
  validationQuery="@VALIDATION_QUERY@"
  testOnBorrow="@TEST_ON_BORROW@"
  testOnReturn="@TEST_ON_RETURN@"
  whenExhaustedAction="0"
  maxWait="10000">

  <!-- Set fetchSize to 0 to use driver's default. -->
  <attribute attribute-name="fetchSize" attribute-value="0"/>

  <!-- Attributes with name prefix "jdbc." are passed directly to the JDBC
driver. -->
  <!-- Example setting (used by Oracle driver when Statement batching is
enabled) -->
  <attribute attribute-name="jdbc.defaultBatchValue" attribute-value="5"/>

  <!-- Attributes determining if ConnectionFactoryDBCPImpl
should also pool PreparedStatement. This is programmatically disabled
when using platform=Oracle9i since Oracle statement caching will
conflict
with DBCP ObjectPool-based PreparedStatement caching (ie setting true
here has no effect for Oracle9i platform). -->
  <attribute attribute-name="dbcp.poolPreparedStatements"
attribute-value="true"/>
  <attribute attribute-name="dbcp.maxOpenPreparedStatements"
attribute-value="60"/>
  <!-- Attribute determining if the Commons DBCP connection wrapper will allow
JDBC-driver
access to the underlying concrete Connection instance from the
(normally this is not allowed, like in J2EE-containers using wrappers).
-->
  <attribute attribute-name="dbcp.accessToUnderlyingConnectionAllowed"
attribute-value="false"/>
</connection-pool>
```

#### jdbc.\*

Since OJB 1.0.4, custom attributes with names starting with "jdbc." will be passed (without the "jdbc." prefix) to the JDBC DriverManager when creating new Connection objects.

Use this attribute to set driver-specific customized tuning options. For example, to set Oracle-batching to 5 statements:

```
<attribute attribute-name="jdbc.defaultBatchValue" attribute-value="5"/>
```

#### fetchSize

(default=0, unspecified) Sets a hint in the JDBC driver not to fetch more than specified number of rows per server roundtrip for any ResultSet.

Settings different than the default (0) are especially useful to reduce memory footprint when using drivers that default to not using server-side cursors and retrieves all rows to the JDBC client-side driver buffer. PostgreSQL JDBC driver is a well-known example of this.

#### Note:

\* Many JDBC drivers will silently ignore the *fetchSize* hint.  
 \* Also note that *fetchSize* has nothing to do with max rows returned by a ResultSet, only number of rows retrieved per JDBC- driver network roundtrip to the database server (if the driver cares about the hint at all, that is).

#### dbcp.poolPreparedStatements

Only valid for *ConnectionFactoryDBCPImpl* (default=false) Enable prepared statement pooling.

#### Note:

PreparedStatement pooling with Commons DBCP is programmatically disabled when using `platform=Oracle9i` in OJB, since the platform implementation activates Oracle-specific statement caching that conflicts with DBCP ObjectPool-based caching. Ie, for a descriptor with `platform="Oracle9i"` there is no effect in setting:

```
<attribute attribute-name="dbcp.poolPreparedStatements" attribute-value="true"/>
```

#### dbcp.maxOpenPreparedStatements

Only valid for *ConnectionFactoryDBCPImpl* (default=0, unlimited) The maximum number of open statements that can be allocated from the statement pool at the same time, or zero for no limit.

#### dbcp.accessToUnderlyingConnectionAllowed

Only valid for *ConnectionFactoryDBCPImpl* (default=false) Controls if the DBCP "PoolGuard" connection wrapper allows access to the underlying Connection instance from the JDBC-driver.

Only use when you need direct access to driver-specific extentions. It is generally **not** needed to change this setting in OJB.

#### Note:

\* Do not close the underlying connection, only the original one.  
 \* If using P6Spy, the underlying connection in DBCP will still be wrapped by P6Spy and you will have to continue unwrapping to the innermost delegate and Connection of JDBC-driver specific class.

### 5.5.7.5. sequence-manager

The *sequence-manager* element specifies the sequence manager implementation used for key generation. All sequence manager implementations shipped with OJB can be found in the *org.apache.ojb.broker.util.sequence* package. If no sequence manager is defined, OJB uses the default one. More info about [sequence key generation here](#).

#### Elements

Use the *custom-attribute* element to pass implementation specific properties.

```
<!ELEMENT sequence-manager (
  documentation?,
  attribute* )
>
```

#### Attributes

The *className* attribute represents the full qualified class name of the desired sequence manager implementation - it is mandatory when using the *sequence-manager* element. All sequence manager implementations you find will under *org.apache.ojb.broker.util.sequence* package named as





Use the [custom-attribute](#) element to pass implementation specific properties.

```
<!ELEMENT object-cache (documentation?, attribute*)>
```

#### Attributes

Attribute 'class' specifies the full qualified class name of the used ObjectCache implementation.

```
<!ATTLIST object-cache class CDATA #REQUIRED>
```

#### Custom Attributes

Many *ObjectCache* implementation classes support specific configuration properties, these properties can be set by using [custom-attributes](#).

The description of the properties can be found in [object cache docs](#).

Usage example of supported custom attributes:

```
<object-cache class="org.apache.ojb.broker.cache.ObjectCacheTwoLevelImpl">
  <!-- meaning of attributes, please see docs section "Caching" -->
  <!-- common attributes -->
  <attribute attribute-name="cacheExcludes" attribute-value="" />

  <!-- ObjectCacheTwoLevelImpl attributes -->
  <attribute attribute-name="applicationCache"
    attribute-value="org.apache.ojb.broker.cache.ObjectCacheDefaultImpl" />
  <attribute attribute-name="copyStrategy"
    attribute-value="org.apache.ojb.broker.cache.ObjectCacheTwoLevelImpl$CopyStrategyImpl" />
  <attribute attribute-name="forceProxies" attribute-value="false" />

  <!-- ObjectCacheDefaultImpl attributes -->
  <attribute attribute-name="timeout" attribute-value="900" />
  <attribute attribute-name="autoSync" attribute-value="true" />
  <attribute attribute-name="cachingKeyType" attribute-value="0" />
  <attribute attribute-name="useSoftReferences" attribute-value="true" />
</object-cache>
```

#### 5.5.7.7. custom attribute

An *attribute* element allows arbitrary name/value pairs to be represented in the repository. See the [repository.dtd](#) for details on which elements support it (e.g. [class-descriptor](#), [object-cache](#), ...).

```
<!ELEMENT attribute EMPTY>
```

The *attribute-name* identifies the name of the attribute.

The *attribute-value* identifies the value of the attribute.

```
<!ATTLIST attribute
  attribute-name CDATA #REQUIRED
  attribute-value CDATA #REQUIRED
>
```

To get access of the defined attribute use methods of `org.apache.ojb.broker.metadata.AttributeContainer`. All classes supporting *custom attributes* have to implement this interface.

Here you can see an example how to define an *custom attribute* within the [class-descriptor](#) element:

```
<class-descriptor
  class="my.TestClass"
  table="OJB_TEST_CLASS"
  >
  <field-descriptor
    name="id"
```

```

        column="ID"
        jdbc-type="INTEGER"
        primaryKey="true"
        autoincrement="true"
    />
    ...
    <attribute attribute-name="myAttribute" attribute-value="myValue"/>
</class-descriptor>

```

To access the attribute you have to know the associated `AttributeContainer` class. Here it was [ClassDescriptor](#). To read the attribute at runtime do:

```

// get the ClassDescriptor
ClassDescriptor cld = broker.getClassDescriptor(TestClass.class);
String value = cld.getAttribute("myAttribute");

```

#### 5.5.7.8. class-descriptor

A *class-descriptor* and the associated java class [ClassDescriptor](#) encapsulate metadata information of an interface, abstract or concrete class.

##### Elements

For *interfaces* or *abstract* classes a *class-descriptor* holds a sequence of *extent-class* elements which specify the [types extending this class](#).

Concrete base classes may specify a sequence of *extent-class* elements, naming the derived classes.

For *concrete* classes it must have *field-descriptors* that describe primitive typed instance variables.

References to other persistent entity classes are specified by *reference-descriptor* elements.

Collections or arrays attributes that contain other persistent entity classes are specified by *collection-descriptor* elements

A *class-descriptor* may contain user defined custom attribute elements.

Use the [custom-attribute](#) element to pass implementation specific properties.

```

<!ELEMENT class-descriptor (
  (
    documentation?,
    extent-class+,
    attribute* ) |
  (
    documentation?,
    object-cache?,
    extent-class*,
    field-descriptor+,
    reference-descriptor*,
    collection-descriptor*,
    index-descriptor*,
    attribute*,
    insert-procedure?,
    update-procedure?,
    delete-procedure? )
)
>

```

##### Attributes

The *class* attribute contains the full qualified name of the specified class. As this attribute is of the XML type ID there can only be one *class-descriptor* per class.

The *isolation-level* attribute defines the locking isolation level of the specified class (used by OJB's [pessimistic locking](#) api).

##### Note:

The *isolation-level* does not touch the jdbc-connection isolation level. It's completely independent from the database connection setting and only important when [pessimistic locking](#) was used.

If the *proxy* attribute is set, proxies are used for all loading operations of instances of this class. If set to *dynamic*, dynamic proxies are used. If set to another value this value is interpreted as the full-qualified name of the proxy class to use. More info about [using of proxies here](#).

The *proxy-prefetching-limit* attribute specifies a limit to the number of elements loaded on a proxied reference. When the first proxied element is loaded, a number up to the proxy-prefetch-limit will be loaded in addition.

The *schema* attribute may contain the database schema owning the table mapped to this class.

The *table* attribute specifies the table name this class is mapped to.

The *row-reader* attribute may contain a full qualified class name. This class will be used as the [RowReader](#) implementation used to materialize instances of the persistent class.

The *extends* attribute is **deprecated** and will be removed or reintroduced with changed functionality in future. DON'T USE IT!

The *accept-locks* attribute specifies whether implicit locking should propagate to this class. Currently relevant for the ODMG layer only.

The optional *initialization-method* specifies a no-argument instance method that is invoked after reading an instance from a database row. It can be used to do initialization and validations.

The optional *factory-class* specifies a factory class that that is to be used instead of a no argument constructor when new objects are created. If the factory class is specified, then the *factory-method* also must be defined. It refers to a static no-argument method of the factory class that returns a new instance.

The *refresh* attribute can be set to *true* to force OJB to refresh instances when loaded from cache. Means all field values (except references) will be replaced by values retrieved from the database. It's set to *false* by default.

```
<!ATTLIST class-descriptor
  class ID #REQUIRED
  isolation-level (read-uncommitted | read-committed |
    repeatable-read | serializable | optimistic | none) "read-uncommitted"
  proxy CDATA #IMPLIED
  proxy-prefetching-limit CDATA #IMPLIED
  schema CDATA #IMPLIED
  table CDATA #IMPLIED
  row-reader CDATA #IMPLIED
  extends IDREF #IMPLIED
  accept-locks (true | false) "true"
  initialization-method CDATA #IMPLIED
  factory-class CDATA #IMPLIED
  factory-method CDATA #IMPLIED
  refresh (true | false) "false"
>
```

### 5.5.7.9. extent-class

An extent-class element is used to specify an implementing class or a derived class that belongs to the extent of all instances of the interface or base class.

```
<!ELEMENT extent-class EMPTY>
```

The *class-ref* attribute must contain a fully qualified classname and the repository file must contain a class-descriptor for this class.

```
<!ATTLIST extent-class class-ref IDREF #REQUIRED>
```

#### 5.5.7.10. field-descriptor

A field descriptor contains mapping info for a primitive typed attribute of a persistent class. A field descriptor may contain [custom attribute](#) elements.

Use the [custom-attribute](#) element to pass implementation specific properties.

```
<!ELEMENT field-descriptor (documentation?, attribute*)>
```

**The *id* attribute is optional.** If not specified, OJB internally sorts field-descriptors according to their order of appearance in the repository file.

If a different sort order is intended the *id* attribute may be used to hold a unique number identifying the descriptors position in the sequence of field-descriptors.

#### Note:

The order of the numbers for the field-descriptors must correspond to the order of columns in the mapped table.

The *name* attribute holds the name of the persistent classes attribute. More info about [persistent field handling](#).

The *table* attribute may specify a table different from the mapped table for the persistent class. (**currently not implemented**).

The *column* attribute specifies the column the persistent classes field is mapped to.

The *jdbc-type* attribute specifies the JDBC type of the column. If not specified OJB tries to identify the JDBC type by inspecting the Java attribute by reflection - OJB use the java/jdbc mapping described [here](#).

The *primarykey* specifies if the column is a primary key column, default value is *false*. It's possible to auto assign primary key fields, more info see [autoincrement section](#)

The *nullable* attribute specifies if the column may contain null values.

The *indexed* attribute specifies if there is an index on this column

The *autoincrement* attribute specifies if the values for the persistent attribute should be automatically generated by OJB. More info about [sequence key generation here](#).

The *sequence-name* attribute can be used to state explicitly a sequence name used by the sequence manager implementations. Check the [javadocs](#) of the used sequence manager implementation to get information if this is a mandatory attribute. OJB standard sequence manager implementations build a sequence name by its own, if the attribute is not set. More info about [sequence key generation here](#).

The *locking* attribute is set to *true* if the persistent attribute is used for *optimistic locking*. More about [optimistic locking](#). The default value is *false*.

The *updatelock* attribute is set to *false* if the persistent attribute is used for optimistic locking AND the dbms should update the lock column itself. The default is *true* which means that when locking is true then OJB will update the locking fields. Can only be set for *TIMESTAMP* and *INTEGER* columns.

The *default-fetch* attribute specifies whether the persistent attribute belongs to the JDO default fetch group.

The *conversion* attribute contains a fully qualified class name. This class must implement the

interface `org.apache.ojb.accesslayer.conversions.FieldConversion`. A `FieldConversion` can be used to implement conversions between Java- attributes and database columns. More about [field conversion](#).

The *length* attribute can be used to specify a length setting if required by the jdbc-type of the underlying database column.

The *precision* attribute can be used to specify a precision setting, if required by the jdbc-type of the underlying database column.

The *scale* attribute can be used to specify a scale setting, if required by the jdbc-type of the underlying database column.

The *access* attribute specifies the accessibility of the field. Fields marked as *readonly* are not to be modified. *readwrite* marks fields that may be read and written to. *anonymous* marks anonymous fields.

An anonymous field has a database representation (column) but no corresponding Java attribute. Hence the name of such a field does not refer to a Java attribute of the class, but is used as a unique identifier only. More info about [anonymous keys here](#).

```
<!ATTLIST field-descriptor
  id CDATA #IMPLIED
  name CDATA #REQUIRED
  table CDATA #IMPLIED
  column CDATA #REQUIRED
  jdbc-type (BIT | TINYINT | SMALLINT | INTEGER | BIGINT | DOUBLE |
            FLOAT | REAL | NUMERIC | DECIMAL | CHAR | VARCHAR |
            LONGVARCHAR | DATE | TIME | TIMESTAMP | BINARY |
            VARBINARY | LONGVARBINARY | CLOB | BLOB) #REQUIRED
  primarykey (true | false) "false"
  nullable (true | false) "true"
  indexed (true | false) "false"
  autoincrement (true | false) "false"
  sequence-name CDATA #IMPLIED
  locking (true | false) "false"
  update-lock (true | false) "true"
  default-fetch (true | false) "false"
  conversion CDATA #IMPLIED
  length CDATA #IMPLIED
  precision CDATA #IMPLIED
  scale CDATA #IMPLIED
  access (readonly | readwrite | anonymous) "readwrite"
>
```

#### 5.5.7.11. reference-descriptor

A reference-descriptor contains mapping info for an attribute of a persistent class that is not primitive but references another persistent entity Object. More about [1:1 references here](#).

A *foreignkey* element contains information on foreign key columns that implement the association on the database level.

```
<!ELEMENT reference-descriptor ( foreignkey+)>
```

The *name* attribute holds the name of the persistent classes attribute. Depending on the used [PersistendField](#) implementation, there must be e.g. an attribute in the persistent class with this name or a JavaBeans compliant property of this name.

The *class-ref* attribute contains a fully qualified class name. This class is the Object type of the persistent reference attribute. As this is an IDREF there must be a class-descriptor for this class in the repository too.

The *proxy* attribute can be set to *true* to specify that proxy based lazy loading should be used for

this attribute.

The *proxy-prefetch-limit* attribute specifies a limit to the number of elements loaded on a proxied reference. When the first proxied element is loaded, a number up to the *proxy-prefetch-limit* will be loaded in addition.

The *refresh* attribute can be set to *true* to force OJB to refresh the object reference when the object is loaded from cache. If *true* OJB try to retrieve the reference (dependent on the [auto-xxx settings](#)) again when the main object is loaded from cache (normally only make sense for 1:n and m:n relations).

This could be useful if the [ObjectCache implementation](#) cache full object graphs without synchronize the referenced objects.

**Note:**

This does not mean that all referenced objects will be read from database. It only means that the reference will be refreshed, the objects itself may provided by the cache. To refresh the object fields itself set the *refresh* attribute in [class-descriptor](#) of the referenced object or disable [caching](#) (to always read objects from the persistent storage).

The *auto-retrieve* attribute specifies whether OJB automatically retrieves this reference attribute on loading the persistent object. If set to *false* the reference attribute is set to null. In this case the user is responsible to fill the reference attribute.

More info about [auto-retrieve here](#).

The *auto-update* attribute specifies whether OJB automatically stores this reference attribute on storing the persistent object.

More info about the [auto-XXX settings here](#).

**Note:**

This attribute must be set to *false* if using the OTM or JDO layer.  
For ODMG-api *none* is mandatory (since OJB 1.0.2).

The *auto-delete* attribute specifies whether OJB automatically deletes this reference attribute on deleting the persistent object.

More info about the [auto-XXX settings here](#).

**Note:**

This attribute must be set to *false* if using the OTM or JDO layer.  
For ODMG-api *none* is mandatory (since OJB 1.0.2).

The *otm-dependent* attribute specifies whether the OTM layer automatically creates the referred object or deletes it if the reference field is set to null. Also *otm-dependent* references behave as if *auto-update* and *auto-delete* were set to true, but the *auto-update* and *auto-delete* attributes themselves must be always set to false for use with OTM layer.

```
<!ATTLIST reference-descriptor
  name CDATA #REQUIRED
  class-ref IDREF #REQUIRED

  proxy (true | false) "false"
  proxy-prefetching-limit CDATA #IMPLIED
  refresh (true | false) "false"

  auto-retrieve (true | false) "true"
  auto-update (none | link | object | true | false) "false"
  auto-delete (none | link | object | true | false) "false"
  otm-dependent (true | false) "false"
>
```

### 5.5.7.12. foreignkey

A *foreignkey* element contains information on a foreign-key persistent attribute that implement the association on the database level.

```
<!ELEMENT foreignkey EMPTY>
```

The *field-ref* and *field-id-ref* attributes contain the name and the id attributes of the field-descriptor used as a foreign key.

#### Note:

Exactly one of these attributes must be specified.

```
<!ATTLIST foreignkey
  field-id-ref CDATA #IMPLIED
  field-ref CDATA #IMPLIED
>
```

### 5.5.7.13. collection-descriptor

A collection-descriptor contains mapping info for a Collection- or Array-attribute of a persistent class that contains persistent entity Objects. See more about [1:n](#) and [m:n](#) references.

The *orderby* element(s) allow to specify the order the collection objects. It's allowed to specify multiple order fields.

The *inverse-foreignkey* elements contains information on foreign-key attributes that implement the association on the database level.

The *fk-pointing-to-this-class* and *fk-pointing-to-element-class* elements are only needed if the Collection or array implements a m:n association. In this case they contain information on the foreign-key columns of the intermediary table.

Use the *custom-attribute* element to pass implementation specific properties.

```
<!ELEMENT collection-descriptor (
  documentation?,
  orderby*,
  inverse-foreignkey*,
  fk-pointing-to-this-class*,
  fk-pointing-to-element-class*,
  attribute*)>
```

The *name* attribute holds the name of the persistent classes attribute. More info about [persistent field handling](#).

The *collection-class* may hold a fully qualified class name. This class must be the Java type of the Collection attribute. This attribute must only specified if the attribute type is not a `java.util.Collection` (or subclass) or Array type. It is also possible to use non Collection or Array type user defined "collection" classes. More info see section [manageable collection](#).

The *element-class-ref* attribute contains a fully qualified class name. This class is the Object type of the elements of persistent collection or Array attribute. As this is an IDREF there must be a class-descriptor for this class in the repository too.

DEPRECATED, please use the '*orderby*'-element. The *orderby* attribute may specify a field of the element class. The Collection or Array will be sorted according to the specified attribute. The sort attribute may be used to specify ascending or descending order for this operation.

The *indirection-table* must specify the name of an intermediary table, if the persistent collection



attribute implements a m:n association.

The *proxy* attribute can be set to true to specify that proxy based lazy loading should be used for this attribute. More about [using proxy here](#).

The *proxy-prefetch-limit* attribute specifies a limit to the number of elements loaded on a proxied reference. When the first proxied element is loaded, a number up to the proxy-prefetch-limit will be loaded in addition.

The *refresh* attribute can be set to *true* to force OJB to refresh the object reference when the object is loaded from cache. If *true* OJB try to retrieve the reference (dependent on the [auto-xxx settings](#)) again when the main object is loaded from cache (normally only make sense for 1:n and m:n relations).

This could be useful if the [ObjectCache implementation](#) cache full object graphs without synchronize the referenced objects.

**Note:**

This does not mean that all referenced objects will be read from database. It only means that the reference will be refreshed, the objects itself may provided by the cache. To refresh the object fields itself set the *refresh* attribute in [class-descriptor](#) of the referenced object or disable [caching](#) (to always read objects from the persistent storage).

The *auto-retrieve* attribute specifies whether OJB automatically retrieves this reference attribute on loading the persistent object. If set to *false* the reference attribute is set to null. In this case the user is responsible to fill the reference attribute.

More info about [auto-retrieve here](#).

The *auto-update* attribute specifies whether OJB automatically stores this reference attribute on storing the persistent object.

More info about the [auto-XXX settings here](#).

**Note:**

This attribute must be set to *false* if using the OTM or JDO layer.  
For ODMG-api *none* is mandatory (since OJB 1.0.2).

The *auto-delete* attribute specifies whether OJB automatically deletes this reference attribute on deleting the persistent object.

More info about the [auto-XXX settings here](#).

**Note:**

This attribute must be set to *false* if using the OTM or JDO layer.  
For ODMG-api *none* is mandatory (since OJB 1.0.2).

The *otm-dependent* attribute specifies whether the OTM layer automatically creates collection elements that were included into the collection, and deletes collection elements that were removed from the collection. Also otm-dependent references behave as if auto-update and auto-delete were set to true, but the auto-update and auto-delete attributes themselves must be always set to false for use with OTM layer.

```
<!ATTLIST collection-descriptor
  name CDATA #IMPLIED
  collection-class CDATA #IMPLIED
  element-class-ref IDREF #REQUIRED
  orderby CDATA #IMPLIED
  sort (ASC | DESC) "ASC"

  indirection-table CDATA #IMPLIED

  proxy (true | false) "false"
```

```

proxy-prefetching-limit CDATA #IMPLIED
refresh (true | false) "false"

auto-retrieve (true | false) "true"
auto-update (none | link | object | true | false) "false"
auto-delete (none | link | object | true | false) "false"
otm-dependent (true | false) "false"
>

```

#### 5.5.7.14. order-by

A *order-by* element contains an attribute name and a sort order.

```
<!ELEMENT orderby (documentation?)>
```

The *name* attribute specifies the field or the column (full qualified column name) the order based on. The *sort* attribute specifies the order direction.

```

<!ATTLIST orderby
  name CDATA #REQUIRED
  sort (ASC | DESC) "ASC"
>

```

Here is an examples of how to use ordering for one side of a m:n reference:

```

<collection-descriptor
  name="actors"
collection-class="org.apache.ojb.broker.util.collections.ManageableArrayList"
  element-class-ref="org.apache.ojb.broker.M2NTest$Actor"
  auto-retrieve="false"
  auto-update="false"
  auto-delete="false"
  indirection-table="M2N_ROLE"
>
  <!-- Check the use of order by element for fields and plain columns -->
  <orderby name="name" sort="ASC"/>
  <orderby name="M2N_ROLE.MOVIE_ID_INT" sort="DESC"/>

  <fk-pointing-to-this-class column="MOVIE_ID_INT"/>
  <fk-pointing-to-this-class column="MOVIE_ID2_INT"/>
  <fk-pointing-to-this-class column="MOVIE_ID_STR"/>
  <fk-pointing-to-element-class column="ACTOR_ID"/>
  <fk-pointing-to-element-class column="ACTOR_ID2"/>
</collection-descriptor>

```

#### 5.5.7.15. inverse-foreignkey

A *inverse-foreignkey* element contains information on a foreign-key persistent attribute that implement the association on the database level.

```
<!ELEMENT inverse-foreignkey EMPTY>
```

The *field-ref* and *field-id-ref* attributes contain the name and the id attributes of the field-descriptor used as a foreign key. Exactly one of these attributes must be specified.

```

<!ATTLIST inverse-foreignkey
  field-id-ref CDATA #IMPLIED
  field-ref CDATA #IMPLIED
>

```

#### 5.5.7.16. fk-pointing-to-this-class

A *fk-pointing-to-this-class* element contains information on a foreign-key column of an intermediary table in a m:n scenario.

```
<!ELEMENT fk-pointing-to-this-class EMPTY>
```

The *column* attribute specifies the foreign-key column in the intermediary table that points to the class holding the collection.

```
<!ATTLIST fk-pointing-to-this-class
  column CDATA #REQUIRED
>
```

#### 5.5.7.17. fk-pointing-to-element-class

A *fk-pointing-to-element-class* element contains information on a foreign-key column of an intermediary table in a m:n scenario.

```
<!ELEMENT fk-pointing-to-element-class EMPTY>
```

The *column* attribute specifies the foreign-key column in the intermediary table that points to the class of the collection elements.

```
<!ATTLIST fk-pointing-to-element-class
  column CDATA #REQUIRED
>
```

#### 5.5.7.18. query-customizer

A query enhancer element to enhance the 1:n query, e.g. to modify the result objects of a query. More info about [customizing collection queries](#).

Use the *custom-attribute* element to pass implementation specific properties.

```
<!ELEMENT query-customizer (
  documentation?,
  attribute*)>

<!ATTLIST query-customizer
  class CDATA #REQUIRED
>
```

#### 5.5.7.19. index-descriptor

An *index-descriptor* describes an index by listing its columns. It may be unique or not.

```
<!ELEMENT index-descriptor (documentation?, index-column+)>

<!ATTLIST index-descriptor
  name CDATA #REQUIRED
  unique (true | false) "false">
```

#### 5.5.7.20. index-column

An *index-column* is just the name of a column in an index.

```
<!ELEMENT index-column (documentation?)>

<!ATTLIST index-column
  name CDATA #REQUIRED>
```

#### 5.5.7.21. Stored Procedure Support

OJB supports stored procedures for insert, update and delete operations. [How to use stored procedures within OJB can be found here](#).

#### insert-procedure

Identifies the procedure/function that should be used to handle insertions for a specific class-descriptor.

The nested *argument* elements define the argument list for the procedure/function as well as the source for each argument.

Use the [custom-attribute](#) element to pass implementation specific properties.

```
<!ELEMENT insert-procedure
  (documentation?, (runtime-argument | constant-argument)?, attribute*)>
```

The *name* attribute identifies the name of the procedure/function to use

The *return-field-ref* identifies the field-descriptor that will receive the value that is returned by the procedure/function. If the procedure/ function does not include a return value, then do not specify a value for this attribute.

The *include-all-fields* attribute indicates if all field-descriptors in the corresponding class-descriptor are to be passed to the procedure/ function. If include-all-fields is 'true', any nested 'argument' elements will be ignored. In this case, values for all field-descriptors will be passed to the procedure/function. The order of values that are passed to the procedure/function will match the order of field-descriptors on the corresponding class-descriptor. If include-all-fields is false, then values will be passed to the procedure/function based on the information in the nested 'argument' elements.

```
<!ATTLIST insert-procedure
  name CDATA #REQUIRED
  return-field-ref CDATA #IMPLIED
  include-all-fields (true | false) "false"
>
```

### update-procedure

Identifies the procedure/function that should be used to handle updates for a specific class-descriptor.

The nested *argument* elements define the argument list for the procedure/function as well as the source for each argument.

Use the [custom-attribute](#) element to pass implementation specific properties.

```
<!ELEMENT update-procedure
  (documentation?, (runtime-argument | constant-argument)?, attribute*)>
```

The *name* attribute identifies the name of the procedure/function to use

The *return-field-ref* identifies the field-descriptor that will receive the value that is returned by the procedure/function. If the procedure/ function does not include a return value, then do not specify a value for this attribute.

The *include-all-fields* attribute indicates if all field-descriptors in the corresponding class-descriptor are to be passed to the procedure/ function. If include-all-fields is 'true', any nested 'argument' elements will be ignored. In this case, values for all field-descriptors will be passed to the procedure/function. The order of values that are passed to the procedure/function will match the order of field-descriptors on the corresponding class-descriptor. If include-all-fields is false, then values will be passed to the procedure/function based on the information in the nested 'argument' elements.

```
<!ATTLIST update-procedure
  name CDATA #REQUIRED
```

```

return-field-ref CDATA #IMPLIED
include-all-fields (true | false) "false"
>

```

### delete-procedure

Identifies the procedure/function that should be used to handle deletions for a specific class-descriptor.

The nested *runtime-argument* and *constant-argument* elements define the argument list for the procedure/function as well as the source for each argument.

Use the [custom-attribute](#) element to pass implementation specific properties.

```

<!ELEMENT delete-procedure
  (documentation?, (runtime-argument | constant-argument)?, attribute*)>

```

The *name* attribute identifies the name of the procedure/function to use

The *return-field-ref* identifies the field-descriptor that will receive the value that is returned by the procedure/function. If the procedure/ function does not include a return value, then do not specify a value for this attribute.

The *include-pk-only* attribute indicates if all field-descriptors in the corresponding class-descriptor that are identified as being part of the primary key are to be passed to the procedure/function. If *include-pk-only* is 'true', any nested 'argument' elements will be ignored. In this case, values for all field-descriptors that are identified as being part of the primary key will be passed to the procedure/function. The order of values that are passed to the procedure/function will match the order of field-descriptors on the corresponding class-descriptor. If *include-pk-only* is false, then values will be passed to the procedure/ function based on the information in the nested 'argument' elements.

```

<!ATTLIST delete-procedure
  name CDATA #REQUIRED
  return-field-ref CDATA #IMPLIED
  include-pk-only (true | false) "false"
>

```

### runtime-argument

Defines an argument that is passed to a procedure/function. Each argument will be set to a value from a field-descriptor or null.

Use the [custom-attribute](#) element to pass implementation specific properties.

```

<!ELEMENT runtime-argument
  (documentation?, attribute*)>

```

The *field-ref* attribute identifies the field-descriptor in the corresponding class-descriptor that provides the value for this argument. If this attribute is unspecified, then this argument will be set to null.

```

<!ATTLIST runtime-argument
  field-ref CDATA #IMPLIED
  return (true | false) "false"
>

```

### constant-argument

Defines a constant value that is passed to a procedure/function.

Use the [custom-attribute](#) element to pass implementation specific properties.

```
<!ELEMENT constant-argument
  (documentation?, attribute*)>
```

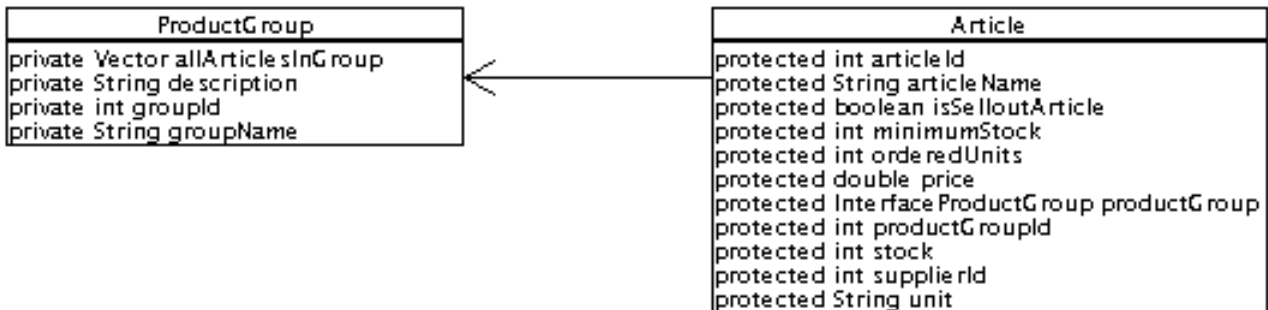
The *value* attribute identifies the value that is passed to the procedure/ function.

```
<!ATTLIST constant-argument
  value CDATA #REQUIRED
>
```

## 5.5.8. Basic O/R Mapping Technique

### 5.5.8.1. Mapping 1:1 associations

As a sample for a simple association we take the reference from an article to its productgroup. This association is navigable only from the article to its productgroup. Both classes are modelled in the following class diagram. This diagram does not show methods, as only attributes are relevant for the O/R mapping process.



1:1 association

The association is implemented by the attribute `productGroup`. To automatically maintain this reference OJB relies on foreignkey attributes. The foreign key containing the `groupId` of the referenced `productgroup` is stored in the attribute `productGroupId`. To avoid FK attribute in persistent object class see section about [anonymous keys](#).

This is the DDL of the underlying tables:

```
CREATE TABLE Artikel
(
  Artikel_Nr          INT NOT NULL PRIMARY KEY,
  Artikelname        VARCHAR(60),
  Lieferanten_Nr     INT,
  Kategorie_Nr       INT,
  Liefereinheit      VARCHAR(30),
  Einzelpreis        FLOAT,
  Lagerbestand       INT,
  BestellteEinheiten INT,
  MindestBestand     INT,
  Auslaufartikel     INT
)

CREATE TABLE Kategorien
(
  Kategorie_Nr       INT NOT NULL PRIMARY KEY,
  KategorieName      VARCHAR(20),
  Beschreibung       VARCHAR(60)
)
```

To declare the foreign key mechanics of this reference attribute we have to add a reference-descriptor to the class-descriptor of the Article class. This descriptor contains the following information:

- The attribute implementing the association ( name="productGroup" ) is productGroup.
- The referenced object is of type ( class-ref="org.apache.ojb.broker.ProductGroup" ) org.apache.ojb.broker.ProductGroup.
- A reference-descriptor contains one or more foreignkey elements. These elements define foreign key attributes. The element

```
<foreignkey field-ref="productGroupId"/>
```

contains the name of the field-descriptor describing the foreignkey fields. The FieldDescriptor with the name "productGroupId" describes the foreignkey attribute productGroupId:

```
<field-descriptor
  name="productGroupId"
  column="Kategorie_Nr"
  jdbc-type="INTEGER"
/>
```

See the following extract from the repository.xml file containing the Article ClassDescriptor:

```
<!-- Definitions for org.apache.ojb.ojb.broker.Article -->
<class-descriptor
  class="org.apache.ojb.broker.Article"
  proxy="dynamic"
  table="Artikel"
>
  <extent-class class-ref="org.apache.ojb.broker.BookArticle" />
  <extent-class class-ref="org.apache.ojb.broker.CdArticle" />
  <field-descriptor
    name="articleId"
    column="Artikel_Nr"
    jdbc-type="INTEGER"
    primarykey="true"
    autoincrement="true"
  />
  <field-descriptor
    name="articleName"
    column="Artikelname"
    jdbc-type="VARCHAR"
  />
  <field-descriptor
    name="supplierId"
    column="Lieferanten_Nr"
    jdbc-type="INTEGER"
  />
  <field-descriptor
    name="productGroupId"
    column="Kategorie_Nr"
    jdbc-type="INTEGER"
  />
  ...
  <reference-descriptor
    name="productGroup"
    class-ref="org.apache.ojb.broker.ProductGroup"
  >
    <foreignkey field-ref="productGroupId"/>
  </reference-descriptor>
</class-descriptor>
```

This example provides unidirectional navigation only. Bidirectional navigation may be added by including a reference from a ProductGroup to a single Article (for example, a sample article for the productgroup). To accomplish this we need to perform the following steps:

1. Add a private Article attribute named sampleArticle to the class ProductGroup.
2. Add a private int attribute named sampleArticleId to the *ProductGroup* class representing the foreign key. To avoid FK attribute in persistent object class see section about [anonymous keys](#).

3. Add a column `SAMPLE_ARTICLE_ID INT` to the table `Kategorien`.
4. Add a `FieldDescriptor` for the foreignkey attribute to the `ClassDescriptor` of the Class `ProductGroup`:

```
<field-descriptor
  name="sampleArticleId"
  column="SAMPLE_ARTICLE_ID"
  jdbc-type="INTEGER"
/>
```

1. Add a `ReferenceDescriptor` to the `ClassDescriptor` of the Class `ProductGroup`:

```
<reference-descriptor
  name="sampleArticle"
  class-ref="org.apache.ojb.broker.Article"
>
  <foreignkey field-ref="sampleArticleId" />
</reference-descriptor>
```

**Note:**

When using primitive primary key fields, please pay attention on [how OJB manage null for primitive PK/FK](#)

### 1:1 auto-xxx setting

General info about the `auto-xxx` and `proxy` attributes can be found [here](#)

#### auto-retrieve

See [here](#)

#### auto-update

- **none** On updating or inserting of the main object with `PersistenceBroker.store(...)`, the referenced object will NOT be updated by default. The reference will not be *inserted* or *updated*, the link to the reference (foreign key value to the reference) on the main object will not be assigned automatically. The user has to link the main object and to store the reference *before* the main object to avoid violation of referential integrity.
- **link** On updating or inserting of the main object with `PersistenceBroker.store(...)`, the FK assignment on the main object was done automatic. OJB reads the PK from the referenced object and sets these values as FK in main object. But the referenced object remains untouched. If no referenced object is found, the FK will be nullified. (On insert it is allowed to set the FK without populating the referenced object)
- **object** On updating or inserting of the main object with `PersistenceBroker.store(...)`, the referenced object will be stored first, then OJB does the same as in *link*.
- **false** Is equivalent to *link*.
- **true** Is equivalent to *object*.

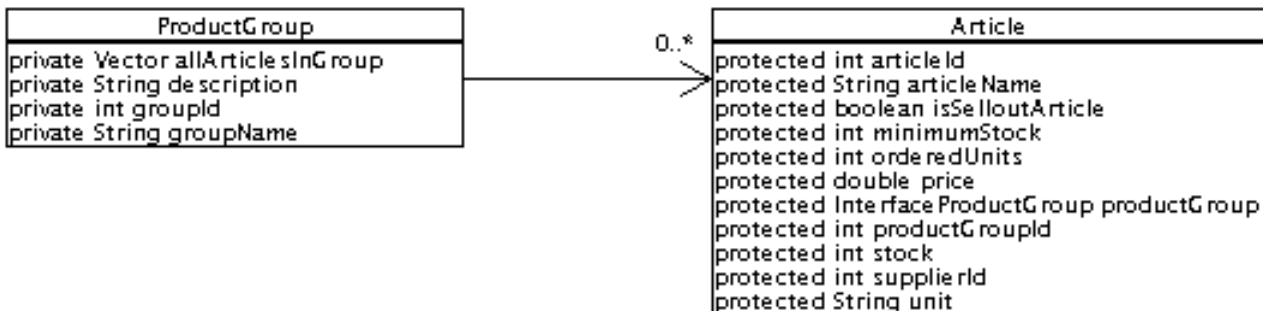
#### auto-delete

- **none** On deleting an object with `PersistenceBroker.delete(...)` the referenced object will NOT be touched.
- **link** Is equivalent to *none*.
- **object** On deleting an object with `PersistenceBroker.delete(...)` the referenced object will be deleted too.
- **false** Is equivalent to *none*.
- **true** Is equivalent to *object*.

### 5.5.8.2. Mapping 1:n associations



We will take a different perspective from the previous example for a 1:n association. We will associate multiple Articles with a single ProductGroup. This association is navigable only from the ProductGroup to its Article instances. Both classes are modelled in the following class diagram. This diagram does not show methods, as only attributes are relevant for the O/R mapping process.



1:n association

The association is implemented by the `Vector` attribute `allArticlesInGroup` on the `ProductGroup` class. As in the previous example, the `Article` class contains a foreignkey attribute named `productGroupId` that identifies an `Article`'s `ProductGroup`. The Database table are the same as above.

To declare the foreign key mechanics of this collection attribute we must add a `CollectionDescriptor` to the `ClassDescriptor` of the `ProductGroup` class. This descriptor contains the following information:

1. The attribute implementing the association (`name="allArticlesInGroup"`)
2. The class of the elements in the collection (`element-class-ref="org.apache.ojb.broker.Article"`)
3. The name of field-descriptor of the element class used as foreign key attributes are defined in inverse-foreignkey elements:

```
<inverse-foreignkey field-ref="productGroupId"/>
```

This is again pointing to the field-descriptor for the attribute `productGroupId` in class `Article`.

4. optional attributes to define the sort order of the retrieved collection:  
`orderby="articleId" sort="DESC".`

See the following extract from the `repository.xml` file containing the `ProductGroup` `ClassDescriptor`:

```

<!-- Definitions for org.apache.ojb.broker.ProductGroup -->
<class-descriptor
  class="org.apache.ojb.broker.ProductGroup"
  table="Kategorien"
>
  <field-descriptor
    name="groupId"
    column="Kategorie_Nr"
    jdbc-type="INTEGER"
    primarykey="true"
    autoincrement="true"
  />
  <field-descriptor
    name="groupName"
    column="KategorieName"
    jdbc-type="VARCHAR"
  />
  <field-descriptor
    name="description"
    column="Beschreibung"
    jdbc-type="VARCHAR"
  />

```

```

/>
<collection-descriptor
  name="allArticlesInGroup"
  element-class-ref="org.apache.ojb.broker.Article"
  order-by="articleId"
  sort="DESC"
>
  <inverse-foreignkey field-ref="productGroupId"/>
</collection-descriptor>
</class-descriptor>

```

With the mapping shown above OJB has two possibilities to load the Articles belonging to a ProductGroup:

1. loading all Articles of the ProductGroup immediately after loading the ProductGroup. This is done with **two** SQL-calls: one for the ProductGroup and one for all Articles.
2. if Article is a proxy ( [using proxy classes](#) ), OJB will only load the keys of the Articles after the ProductGroup. When accessing an Article-proxy OJB will have to materialize it with another SQL-Call. Loading the ProductGroup and all it's Articles will thus produce **n+2** SQL-calls: one for the ProductGroup, one for keys of the Articles and one for each Article.

Both approaches have their benefits and drawbacks:

- **A.** is suitable for a small number of related objects that are easily instantiated. It's efficient regarding DB-calls. The major drawback is the amount of data loaded. For example to show a list of ProductGroups the Articles may not be needed.
- **B.** is best used for a large number of related heavy objects. This solution loads the objects when they are needed ("lazy loading"). The price to pay is a DB-call for each object.

Further down a third solution [using a single proxy for a whole collection](#) will be presented to circumvent the described drawbacks.

OJB supports different Collection types to implement 1:n and m:n associations. OJB detects the used type automatically, so there is no need to declare it in the repository file. But in some cases the *default* behaviour of OJB is undesired. Please [read here for more information](#).

#### Note:

When using primitive primary key fields, please pay attention on [how OJB manage null for primitive PK/FK](#)

### 1:n auto-xxx setting

General info about the *auto-xxx* and *proxy* attributes can be found [here](#).

#### auto-retrieve

See [here](#)

#### auto-update

- **none** On updating or inserting of the main object with `PersistenceBroker.store(...)`, the referenced objects are NOT updated by default. The referenced objects will not be *inserted* or *updated*, the referenced objects will not be linked (foreign key assignment on referenced objects) to the main object automatically. The user has to link and to store the referenced objects *after* storing the main object to avoid violation of referential integrity.
- **link** On updating or inserting of the main object with `PersistenceBroker.store(...)`, the referenced objects are NOT updated by default. The referenced objects will not be *inserted* or *updated*, but the referenced objects will be linked automatically (FK assignment) the main object.
- **object** On updating or inserting of the main object with `PersistenceBroker.store(...)`, the referenced objects will be linked and stored

automatically.

- **false** Is equivalent to *link*.
- **true** Is equivalent to *object*.

#### auto-delete


- **none** On deleting an object with `PersistenceBroker.delete(...)` the referenced objects are NOT touched. This may lead to violation of referential integrity if the referenced objects are childs of the main object. In this case the referenced objects have to be deleted manually first.
- **link** Is equivalent to *none*.
- **object** On deleting an object with `PersistenceBroker.delete(...)` the referenced objects will be deleted too.
- **false** Is equivalent to *none*.
- **true** Is equivalent to *object*.

### 5.5.8.3. Mapping m:n associations

OJB provides support for manually decomposed m:n associations as well as an automated support for non decomposed m:n associations.

#### Manual decomposition into two 1:n associations

Have a look at the following class diagram:



m:n association

We see a two classes with a m:n association. A Person can work for an arbitrary number of Projects. A Project may have any number of Persons associated to it.

Relational databases don't support m:n associations. They require to perform a manual decomposition by means of an intermediary table. The DDL looks like follows:

```

CREATE TABLE PERSON (
    ID          INT NOT NULL PRIMARY KEY,
    FIRSTNAME  VARCHAR(50),
    LASTNAME   VARCHAR(50)
);


CREATE TABLE PROJECT (
    ID          INT NOT NULL PRIMARY KEY,
    TITLE      VARCHAR(50),
    DESCRIPTION VARCHAR(250)
);

CREATE TABLE PERSON_PROJECT (
    PERSON_ID  INT NOT NULL,
    PROJECT_ID INT NOT NULL,
    PRIMARY KEY (PERSON_ID, PROJECT_ID)
);
  
```

This intermediary table allows to decompose the m:n association into two 1:n associations. The intermediary table may also hold additional information. For example, the role a certain person plays for a project:

```
CREATE TABLE PERSON_PROJECT (
  PERSON_ID INT NOT NULL,
  PROJECT_ID INT NOT NULL,
  ROLENAME VARCHAR(20),
  PRIMARY KEY (PERSON_ID, PROJECT_ID)
);
```

The decomposition is mandatory on the ER model level. On the object model level it is not mandatory, but may be a valid solution. It is mandatory on the object level if the association is qualified (as in our example with a rolename). This will result in the introduction of an association class. A class-diagram reflecting this decomposition looks like:



A Person object has a Collection attribute roles containing Role entries. A Project has a Collection attribute roles containing Role entries. A Role has reference attributes to its Person and to its Project.

Handling of 1:n mapping has been explained above. Thus we will finish this section with a short look at the repository entries for the classes `org.apache.ojb.broker.Person`, `org.apache.ojb.broker.Project` and `org.apache.ojb.broker.Role`:

```
<!-- Definitions for org.apache.ojb.broker.Person -->
<class-descriptor
  class="org.apache.ojb.broker.Person"
  table="PERSON"
>
  <field-descriptor
    name="id"
    column="ID"
    jdbc-type="INTEGER"
    primarykey="true"
    autoincrement="true"
  />
  <field-descriptor
    name="firstname"
    column="FIRSTNAME"
    jdbc-type="VARCHAR"
  />
  <field-descriptor
    name="lastname"
    column="LASTNAME"
    jdbc-type="VARCHAR"
  />
  <collection-descriptor
    name="roles"
    element-class-ref="org.apache.ojb.broker.Role"
  >
    <inverse-foreignkey field-ref="person_id"/>
  </collection-descriptor>
  ...
</class-descriptor>

<!-- Definitions for org.apache.ojb.broker.Project -->
<class-descriptor
  class="org.apache.ojb.broker.Project"
  table="PROJECT"
>
  <field-descriptor
    name="id"
    column="ID"
  />
```

```

        jdbc-type="INTEGER"
        primaryKey="true"
        autoincrement="true"
    />
    <field-descriptor
        name="title"
        column="TITLE"
        jdbc-type="VARCHAR"
    />
    <field-descriptor
        name="description"
        column="DESCRIPTION"
        jdbc-type="VARCHAR"
    />
    <collection-descriptor
        name="roles"
        element-class-ref="org.apache.ojb.broker.Role"
    >
        <inverse-foreignkey field-ref="project_id"/>
    </collection-descriptor>
    ...
</class-descriptor>
<!-- Definitions for org.apache.ojb.broker.Role -->
<class-descriptor
    class="org.apache.ojb.broker.Role"
    table="PERSON_PROJECT"
    >
    <field-descriptor
        name="person_id"
        column="PERSON_ID"
        jdbc-type="INTEGER"
        primaryKey="true"
    />
    <field-descriptor
        name="project_id"
        column="PROJECT_ID"
        jdbc-type="INTEGER"
        primaryKey="true"
    />
    <field-descriptor
        name="roleName"
        column="ROLENAME"
        jdbc-type="VARCHAR"
    />
    <reference-descriptor
        name="person"
        class-ref="org.apache.ojb.broker.Person"
    >
        <foreignkey field-ref="person_id"/>
    </reference-descriptor>
    <reference-descriptor
        name="project"
        class-ref="org.apache.ojb.broker.Project"
    >
        <foreignkey field-ref="project_id"/>
    </reference-descriptor>
</class-descriptor>

```

### Support for Non-Decomposed m:n Mappings

If there is no need for an association class at the object level (we are not interested in role information), OJB can be configured to do the m:n mapping transparently. For example, a `Person` does not have a collection of `Role` objects but only a `Collection` of `Project` objects (held in the attribute `projects`). `Projects` also are expected to contain a collection of `Person` objects (held in attribute `persons`).

To tell OJB how to handle this m:n association the `CollectionDescriptors` for the `Collection` attributes `projects` and `roles` need additional information on the intermediary table and the

foreign key columns pointing to the PERSON table and the foreign key columns pointing to the PROJECT table:

**Note:**

OJB supports a [multiplicity of collection implementations](#), inter alia `org.apache.ojb.broker.util.collections.Re movalAwareCollection` and `org.apache.ojb.broker.util.collections.Re movalAwareList`. By default the removal aware collections were used. This cause problems in m:n relations when `auto-update="true"` or `"object"` and `auto-delete="false"` or `"none"` is set, because objects deleted in the collection will be deleted on update of main object. Thus it is recommended to use a NOT removal aware collection class in m:n relations using the [collection-class](#) attribute.

Example for setting a collection class in the collection-descriptor:

```
collection-class="org.apache.ojb.broker.util.collections.ManageableArrayList"
```

An full example for a non-decomposed m:n relation looks like:

```
<class-descriptor
  class="org.apache.ojb.broker.Person"
  table="PERSON"
>
  <field-descriptor
    name="id"
    column="ID"
    jdbc-type="INTEGER"
    primarykey="true"
    autoincrement="true"
  />
  <field-descriptor
    name="firstname"
    column="FIRSTNAME"
    jdbc-type="VARCHAR"
  />
  <field-descriptor
    name="lastname"
    column="LASTNAME"
    jdbc-type="VARCHAR"
  />
  ...
  <collection-descriptor
    name="projects"
    collection-class="org.apache.ojb.broker.util.collections.ManageableArrayList"
    element-class-ref="org.apache.ojb.broker.Project"
    auto-retrieve="true"
    auto-update="true"
    indirection-table="PERSON_PROJECT"
  >
    <fk-pointing-to-this-class column="PERSON_ID"/>
    <fk-pointing-to-element-class column="PROJECT_ID"/>
  </collection-descriptor>
</class-descriptor>

<!-- Definitions for org.apache.ojb.broker.Project -->
<class-descriptor
  class="org.apache.ojb.broker.Project"
  table="PROJECT"
>
  <field-descriptor
    name="id"
    column="ID"
    jdbc-type="INTEGER"
    primarykey="true"
    autoincrement="true"
  />
  <field-descriptor
    name="title"
    column="TITLE"
    jdbc-type="VARCHAR"
  />
```

```

/>
<field-descriptor
  name="description"
  column="DESCRIPTION"
  jdbc-type="VARCHAR"
/>
...
<collection-descriptor
  name="persons"
  collection-class="org.apache.ojb.broker.util.collections.ManageableArrayList"
  element-class-ref="org.apache.ojb.broker.Person"
  auto-retrieve="true"
  auto-update="false"
  indirection-table="PERSON_PROJECT"
>
  <fk-pointing-to-this-class column="PROJECT_ID"/>
  <fk-pointing-to-element-class column="PERSON_ID"/>
</collection-descriptor>
</class-descriptor>

```

That is all that needs to be configured! See the code in class `org.apache.ojb.broker.MtoNMapping` for JUnit testmethods using the classes `Person`, `Project` and `Role`.

#### Note:

When using primitive primary key fields, please pay attention on [how OJB manage null for primitive PK/FK](#)

### m:n auto-xxx setting

General info about the `auto-xxx` and `proxy` attributes can be found [here](#)

#### auto-retrieve

See [here](#)

#### auto-update

- **none** On updating or inserting of the main object with `PersistenceBroker.store(...)`, the referenced objects are NOT updated by default. The referenced objects will not be *inserted* or *updated*, the referenced objects will not be linked (creation of FK entries in the indirection table) automatically. The user has to store the main object, the referenced objects and to link the m:n relation after storing of all objects. establishing the m:n relationship *before* storing main and referenced objects may violate referential integrity.
- **link** On updating or inserting of the main object with `PersistenceBroker.store(...)`, the referenced objects are NOT updated by default. The referenced objects will not be *inserted* or *updated*, but the m:n relation will be linked automatically (creation of FK entries in the indirection table).

#### Note:

Make sure that the referenced objects exist in database before storing the main object with `auto-update` set *link* to avoid violation of referential integrity.

- **object** On updating or inserting of the main object with `PersistenceBroker.store(...)`, the referenced objects will be linked and stored automatically.
- **false** Is equivalent to *link*.
- **true** Is equivalent to *object*.

#### auto-delete

- **none** On deleting an object with `PersistenceBroker.delete(...)` the referenced

objects are NOT touched. The corresponding entries of the main object in the indirection table will not be removed. This may lead to violation of referential integrity depending on the definition of the indirection table.

- **link** On deleting an object with `PersistenceBroker.delete(...)` the m:n relation will be *unlinked* (all entries of the main object in the indirection table will be removed).
- **object** On deleting an object with `PersistenceBroker.delete(...)` all referenced objects will be deleted too.
- **false** Is equivalent to *link*.
- **true** Is equivalent to *object*.

#### 5.5.8.4. Setting Load, Update, and Delete Cascading

As shown in the sections on 1:1, 1:n and m:n mappings, OJB manages associations (or object references in Java terminology) by declaring special Reference and Collection Descriptors. These Descriptor may contain some additional information that modifies OJB's behaviour on object materialization, updating and deletion.

The behaviour depends on specific attributes

- *auto-retrieve* - possible settings are *false*, *true*. If not specified in the descriptor the default value is *true*
- *auto-update* - possible settings are *none*, *link*, *object* and deprecated [*false*, *true*]. If not specified in the descriptor the default value is *false*
- *auto-delete* - possible settings are *none*, *link*, *object* and deprecated [*false*, *true*]. If not specified in the descriptor the default value is *false*

#### Warning:

When using a top-level api (ODMG, OTM, JDO) it is **mandatory to use specific auto-xxx settings**.

For OTM- and JDO-api the settings are:

```
- auto-retrieve="true"
- auto-update="false"
- auto-delete="false"
```

This are at the same time the *default* auto-XXX settings (so don't specify any of this attributes will have the same effect).

For the ODMG-api the mandatory settings are (since OJB 1.0.2):

```
- auto-retrieve="true"
- auto-update="none"
- auto-delete="none"
```

The attribute *auto-update* and *auto-delete* are described in detail in the corresponding sections for [1:1](#), [1:n](#) and [m:n](#) references. The *auto-retrieve* setting is described below:

#### auto-retrieve setting

The `auto-retrieve` attribute used in `reference-descriptor` or `collection-descriptor` elements handles the loading behaviour of references (1:1, 1:n and m:n):

- **false** If set *false* the referenced objects will not be materialized on object materialization. The user has to materialize the n-side objects (or single object for 1:1) by hand using one of the following service methods of the `PersistenceBroker` class:

```
PersistenceBroker.retrieveReference(Object obj, String attributeName);
// or
PersistenceBroker.retrieveAllReferences(Object obj);
```

The first method load only the specified reference, the second one loads all references declared for the given object.

#### Note:

Be careful when using "opposite" settings, e.g. if you declare a 1:1 reference with `auto-retrieve="false"` BUT `auto-update="object"` (or `"true"` or `"link"`).



Before you can perform an update on the main object, you have to "retrieve" the 1:1 reference. Otherwise you will end up with a nullified reference entry in main object, because OJB doesn't find the referenced object on update and assume the reference was removed.

- **true** If set *true* the referenced objects (single reference or all n-side objects) will be automatic loaded by OJB when the main object was materialized.

If OJB is configured to use proxies, the referenced objects are not materialized immediately, but lazy loading proxy objects are used instead.

In the following code sample, a reference-descriptor and a collection-descriptor are configured to use cascading retrieval ( `auto-retrieve="true"` ), cascading insert/update ( `auto-update="object"` or `auto-update="true"` ) and cascading delete ( `auto-delete="object"` or `auto-delete="true"` ) operations:

```
<reference-descriptor
  name="productGroup"
  class-ref="org.apache.ojb.broker.ProductGroup"
  auto-retrieve="true"
  auto-update="object"
  auto-delete="object"
>
  <foreignkey field-ref="productGroupId"/>
</reference-descriptor>

<collection-descriptor
  name="allArticlesInGroup"
  element-class-ref="org.apache.ojb.broker.Article"
  auto-retrieve="true"
  auto-update="object"
  auto-delete="object"
  orderby="articleId"
  sort="DESC"
>
  <inverse-foreignkey field-ref="productGroupId"/>
</collection-descriptor>
```

### Link references

If in `reference-descriptor` or `collection-descriptor` the *auto-update* or *auto-delete* attributes are set to *none*, OJB does not touch the referenced objects on insert, update or delete operations of the main object. The user has to take care of the correct handling of referenced objects. When using referential integrity (who does not ?) it's essential that insert and delete operations are done in the correct sequence.

One important thing is assignment of the FK values. The assign of the FK values is transcribed with **link references** in OJB. In 1:1 references the main object has a FK to the referenced object, in 1:n references the referenced objects have FK pointing to the main object and in non-decomposed m:n relations a indirection table containing FK values from both sides of the relationship is used.

OJB provides some helper methods for linking references manually (assignment of the FK) in `org.apache.ojb.broker.util.BrokerHelper` class.

```
public void link(Object obj, boolean insert)
public void unlink(Object obj)
public boolean link(Object obj, String attributeName, boolean insert)
public boolean unlink(Object obj, String attributeName)
```

These methods are accessible via `org.apache.ojb.broker.PersistenceBroker`:

```
BrokerHelper bh = broker.serviceBrokerHelper();
```

### Note:

The *link/unlink* methods are only useful if you set auto-update/delete to *none*. In all other cases OJB handles the link/unlink of references internally. It is also possible to set all FK values by hand without using the link/unlink service methods.

## Examples

Now we prepared for some example. Say class `Movie` has an m:n reference with class `Actor` and we want to store an `Movie` object with a list of `Actor` objects. The auto-update setting of collection-descriptor for `Movie` is *none*:

```
broker.beginTransaction();
// store main object first
broker.store(movie);
//now we store the right-side objects
Iterator it = movie.getActors().iterator();
while(it.hasNext())
{
    Object actor = it.next();
    broker.store(actor);
}
// now both side exist and we can link the references
broker.serviceBrokerHelper().link(movie, "actors", true);
/*
alternative call
broker.serviceBrokerHelper().link(movie, true);
*/
broker.commitTransaction();
```

First store the main object and the references, then use `broker.serviceBrokerHelper().link(movie, "actors", true)` to link the main object with the references. In case of a m:n relation linking create all FK entries in the indirection table.

In the next examples we want to manually delete a `Project` object with a 1:n relation to class `SubProject`. In the example, the `Project` object has load all `SubProject` objects and we want to delete the `Project` but **don't** want to delete the referenced `SubProjects` too (don't ask if this make sense ;-)). `SubProject` has an FK to `Project`, so we first have to *unlink* the reference from the main object to the references to avoid integrity constraint violation. Then we can delete the main object:

```
broker.beginTransaction();
// first unlink the n-side references
broker.serviceBrokerHelper().unlink(project, "subProjects");

// update the n-side references, store SubProjects with nullified FK
Iterator it = project.getSubProjects().iterator();
while(it.hasNext())
{
    SubProject subProject = (SubProject) it.next();
    broker.store(subProject);
}

// now delete the main object
broker.delete(project);
broker.commitTransaction();
```

### 5.5.8.5. Using Proxy Classes

Proxy classes can be used for "lazy loading" aka "lazy materialization". Using Proxy classes can help you in reducing unnecessary database lookups.

There are two kind of proxy mechanisms available:

1. **Dynamic proxies** provided by OJB. They can simply be activated by setting certain switches in `repository.xml`. This is the solution recommended for **most** cases.
2. **User defined proxies**. User defined proxies allow the user to write proxy implementations.

As it is important to understand the mechanics of the proxy mechanism I highly recommend to read this section before turning to the next sections "using dynamic proxies", "using a single proxy for a whole collection" and "using a proxy for a reference", covering dynamic proxies.

As a simple example we take a ProductGroup object pg which contains a collection of fifteen Article objects. Now we examine what happens when the ProductGroup is loaded from the database:

Without using proxies all fifteen associated Article objects are immediately loaded from the db, even if you are not interested in them and just want to lookup the description-attribute of the ProductGroup object.

If proxies are used, the collection is filled with fifteen proxy objects, that implement the same interface as the "real objects" but contain only an OID and a void reference. The fifteen article objects are not instantiated when the ProductGroup is initially materialized. Only when a method is invoked on such a proxy object will it load its "real subject" and delegate the method call to it. Using this dynamic delegation mechanism instantiation of persistent objects and database lookups can be minimized.

To use proxies, the persistent class in question (in our case the Article class) must implement an interface (for example InterfaceArticle). This interface is needed to allow replacement of the proper Article object with a proxy implementing the same interface. Have a look at the code:

```
public class Article implements InterfaceArticle
{
    /** maps to db-column "Artikel-Nr"; PrimaryKey*/
    protected int articleId;
    /** maps to db-column "Artikelname"*/
    protected String articleName;
    ...

    public int getArticleId()
    {
        return articleId;
    }

    public java.lang.String getArticleName()
    {
        return articleName;
    }
    ...
}

public interface InterfaceArticle
{
    public int getArticleId();
    public java.lang.String getArticleName();
    ...
}
```

```
public class ArticleProxy extends VirtualProxy implements InterfaceArticle
{
    public ArticleProxy(ojb.broker.Identity uniqueId, PersistenceBroker broker)
    {
        super(uniqueId, broker);
    }

    public int getArticleId()
    {
        return realSubject().getArticleId();
    }

    public java.lang.String getArticleName()
    {
```

```

        return realSubject().getArticleName();
    }

    private InterfaceArticle realSubject()
    {
        try
        {
            return (InterfaceArticle) getRealSubject();
        }
        catch (Exception e)
        {
            return null;
        }
    }
}

```

The proxy is constructed from the identity of the real subject. All method calls are delegated to the object returned by `realSubject()`.

This method uses `getRealSubject()` from the base class `VirtualProxy`:

```

public Object getRealSubject() throws PersistenceBrokerException
{
    return indirectionHandler.getRealSubject();
}

```

The proxy delegates the the materialization work to its `IndirectionHandler`. If the real subject has not yet been materialized, a `PersistenceBroker` is used to retrieve it by its `OID`:

```

public synchronized Object getRealSubject()
    throws PersistenceBrokerException
{
    if (realSubject == null)
    {
        materializeSubject();
    }
    return realSubject;
}

private void materializeSubject()
    throws PersistenceBrokerException
{
    realSubject = broker.getObjectByIdentity(id);
}

```

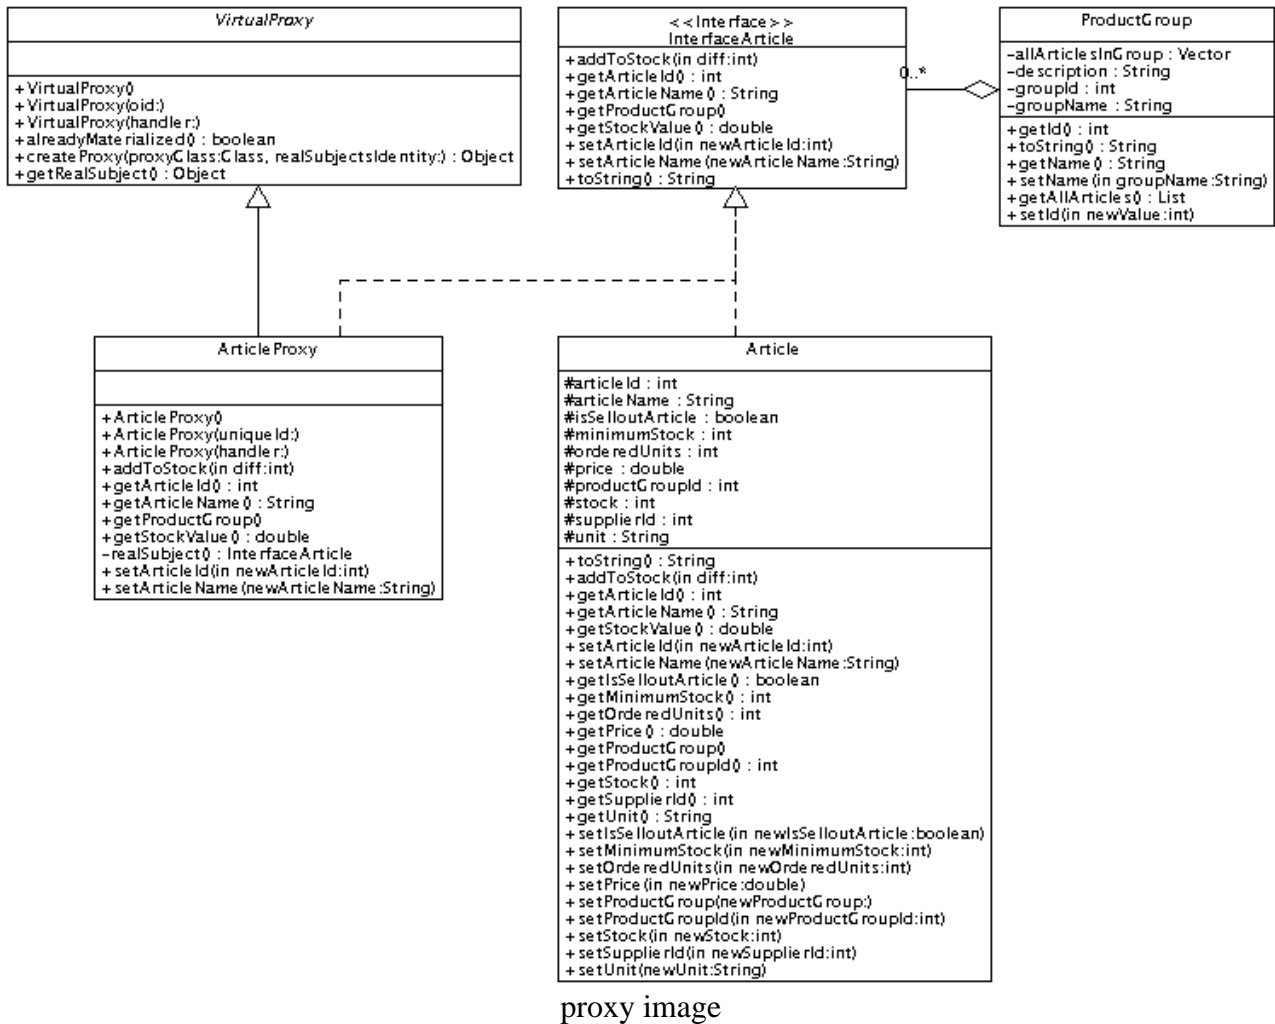
To tell OJB to use proxy objects instead of materializing full `Article` objects we have to add the following section to the XML repository file:

```

<class-descriptor
  class="org.apache.ojb.broker.Article"
  proxy="org.apache.ojb.broker.ArticleProxy"
  table="Artikel"
>
...

```

The following class diagram shows the relationships between all above mentioned classes:



## Using Dynamic Proxies

The implementation of a proxy class is a boring task that repeats the same delegation scheme for each new class. To liberate the developer from this unproductive job OJB provides a dynamic proxy solution based on the JDK 1.3 dynamic proxy concept. (For JDK1.2 we ship a replacement for the required `java.lang.reflect` classes. Credits for this solution to ObjectMentor.) The basic idea of the dynamic proxy concept is to catch all method invocations on the not-yet materialized (loaded from database) object. When a method is called on the object, Java directs this call to the invocation handler registered for it (in OJB's case a class implementing the `org.apache.ojb.broker.core.proxy.IndirectionHandler` interface). This handler then materializes the object from the database and replaces the proxy with the real object. By default OJB uses the class `org.apache.ojb.broker.core.proxy.IndirectionHandlerDefaultImpl`. If you are interested in the mechanics have a look at this class.

To use a dynamic proxy for lazy materialization of Article objects we have to declare it in the `repository.xml` file.

```

<class-descriptor
  class="org.apache.ojb.broker.Article"
  proxy="dynamic"
  table="Artikel"
>
...

```

Just as with normal proxies, the persistent class in question (in our case the Article class) must

implement an interface (for example `InterfaceArticle`) to be able to benefit from dynamic proxies.

#### Note:

As of OJB 1.0.4, a facility is now present to allow the generation of dynamic proxies that **do not** require the persistent class to implement an interface. Previous versions generated Proxies using the JDK proxy pattern. That has been extracted into a new configuration setting named 'ProxyFactoryClass'.

Two implementations of this ProxyClass have been provided: the previous JDK-based version (default), and a new CGLIB-based implementation. Since the CGLIB version does byte-code manipulation to generate the proxy, your class is not required to implement any interface. All generated Proxies will automatically be sub-classes of your persistent class.

See below in the section "Customizing the proxy mechanism" for how to enable the new CGLIB Proxy generation.

### Using a Single Proxy for a Whole Collection

A collection proxy represents a whole collection of objects, where as a proxy class represents a single object.

The advantage of this concept is a reduced number of db-calls compared to using proxy classes. A collection proxy only needs a **single** db-call to materialize all it's objects. This happens the first time its content is accessed (ie: by calling `iterator()`). An additional db-call is used to calculate the size of the collection if `size()` is called *before* loading the data. So collection proxy is mainly used as a deferred execution of a query.

OJB uses three specific proxy classes for collections:

1. **List proxies** are specific `java.util.List` implementations that are used by OJB to replace lists. The default set proxy class is  
`org.apache.ojb.broker.core.proxy.ListProxyDefaultImpl`
2. **Set proxies** are specific `java.util.Set` implementations that are used by OJB to replace sets. The default set proxy class is  
`org.apache.ojb.broker.core.proxy.SetProxyDefaultImpl`
3. **Collection proxies** are collection classes implementing the more generic `java.util.Collection` interface and are used if the collection is neither a list nor a set. The default collection proxy class is  
`org.apache.ojb.broker.core.proxy.CollectionProxyDefaultImpl`

Which of these proxy class is actually used, is determined by the `collection-class` setting of this collection. If none is specified in the repository descriptor, or if the specified class does not implement `java.util.List` nor `java.util.Set`, then the generic collection proxy is used.

The following mapping shows how to use a collection proxy for a relationship:

```
<!-- Definitions for
org.apache.ojb.broker.ProductGroupWithCollectionProxy -->
<class-descriptor
  class="org.apache.ojb.broker.ProductGroupWithCollectionProxy"
  table="Kategorien"
>
  <field-descriptor
    name="groupId"
    column="Kategorie_Nr"
    jdbc-type="INTEGER"
    primarykey="true"
  />
  ...
  <collection-descriptor
    name="allArticlesInGroup"
    element-class-ref="org.apache.ojb.broker.Article"
    proxy="true"
  >
    <inverse-foreignkey field-ref="productGroupId"/>
  </collection-descriptor>
</class-descriptor>
```

The classes participating in this relationship do not need to implement a special interface to be used in a collection proxy.

Although it is possible to mix the collection proxy concept with the proxy class concept, it is not recommended because it increases the number of database calls.

### Using a Proxy for a Reference

A proxy reference is based on the original proxy class concept. The main difference is that the `ReferenceDescriptor` defines when to use a proxy class and not the `ClassDescriptor`. In the following mapping the class `ProductGroup` is **not** defined to be a proxy class in its `ClassDescriptor`. Only for shown relationship a proxy of `ProductGroup` should be used:

```
<!-- Definitions for org.apache.ojb.broker.ArticleWithReferenceProxy -->
<class-descriptor
  class="org.apache.ojb.broker.ArticleWithReferenceProxy"
  table="Artikel"
>
  <field-descriptor
    name="articleId"
    column="Artikel_Nr"
    jdbc-type="INTEGER"
    primarykey="true"
    autoincrement="true"
  />
  ...
  <reference-descriptor
    name="productGroup"
    class-ref="org.apache.ojb.broker.ProductGroup"
    proxy="true"
  >
    <foreignkey field-ref="productGroupId"/>
  </reference-descriptor>
</class-descriptor>
```

Because a proxy reference is only about the location of the definition, the referenced class must implement a special interface (see [using proxy classes](#)).

### Customizing the proxy mechanism

Both the dynamic and the collection proxy mechanism can be customized by supplying a user-defined implementation.

For dynamic proxies, you can select a `ProxyFactory`, as well as provide your own indirection handler. Two default indirection handler implementations have been provided that correspond to the appropriate `ProxyFactory` (`IndirectionHandlerJDKImpl` and `IndirectionHandlerCGLIBImpl`).

**Note:** All indirection handlers must implement the appropriate base indirection handler class, depending on what `ProxyFactory` is being used. For example: when using `ProxyFactoryJDKImpl`, the specified indirection handler **must** implement the `IndirectionHandlerJDK` interface.

Each of the three collection proxy classes can be replaced by a user-defined class. The only requirement is that such a class implements both the corresponding interface (`java.util.Collection`, `java.util.List`, or `java.util.Set`) as well as the `org.apache.ojb.broker.ManageableCollection` interface.

Proxy implementations are configured in the `ojb` properties file. These are the relevant settings:

```
...
#-----
# ProxyFactory and IndirectionHandler
#-----
```

```

# The ProxyFactoryClass entry defines which ProxyFactory implementation
is to be used.
# By default, a 1.0 compatible, JDK-based version is used. However, a
the CGLIB
# based entry is available.
#
# - ProxyFactoryCGLIBImpl: Reference proxies are generated using
bytecode manipulation
#           from the CGLIB library. Any class can become a dynamic
proxy, and not
#           just ones that implement an interface.
# - ProxyFactoryJDKImpl: OJB 1.0 compatible Proxy implementation.
Proxies in this method
#           can only be generated from classes that implement an
interface, and
#           the generated Proxy will implement all methods of that
interface.
#
# NOTE: The appropriate cooresponding IndirectionHandler must be chosen
as well
#
#ProxyFactoryClass=org.apache.ojb.broker.core.proxy.ProxyFactoryCGLIBImpl
ProxyFactoryClass=org.apache.ojb.broker.core.proxy.ProxyFactoryJDKImpl
#
# The IndirectionHandlerClass entry defines the class to be used by
OJB's proxies to
# handle method invocations
#
#IndirectionHandlerClass=org.apache.ojb.broker.core.proxy.IndirectionHandlerCGLIBImpl
IndirectionHandlerClass=org.apache.ojb.broker.core.proxy.IndirectionHandlerJDKImpl
#-----
#
# ListProxy
#-----
# The ListProxyClass entry defines the proxy class to be used for
collections that
# implement the java.util.List interface.
#
ListProxyClass=org.apache.ojb.broker.core.proxy.ListProxyDefaultImpl
#-----
#
# SetProxy
#-----
# The SetProxyClass entry defines the proxy class to be used for
collections that
# implement the java.util.Set interface.
#
SetProxyClass=org.apache.ojb.broker.core.proxy.SetProxyDefaultImpl
#-----
#
# CollectionProxy
#-----
# The CollectionProxyClass entry defines the proxy class to be used for
collections that
# do not implement java.util.List or java.util.Set.
#
CollectionProxyClass=org.apache.ojb.broker.core.proxy.CollectionProxyDefaultImpl
...

```

#### 5.5.8.6. Type and Value Conversions

Say your database column contains INTEGER values but you have to use boolean attributes in your Domain objects. You need a type and value mapping described by a [FieldConversion!](#)

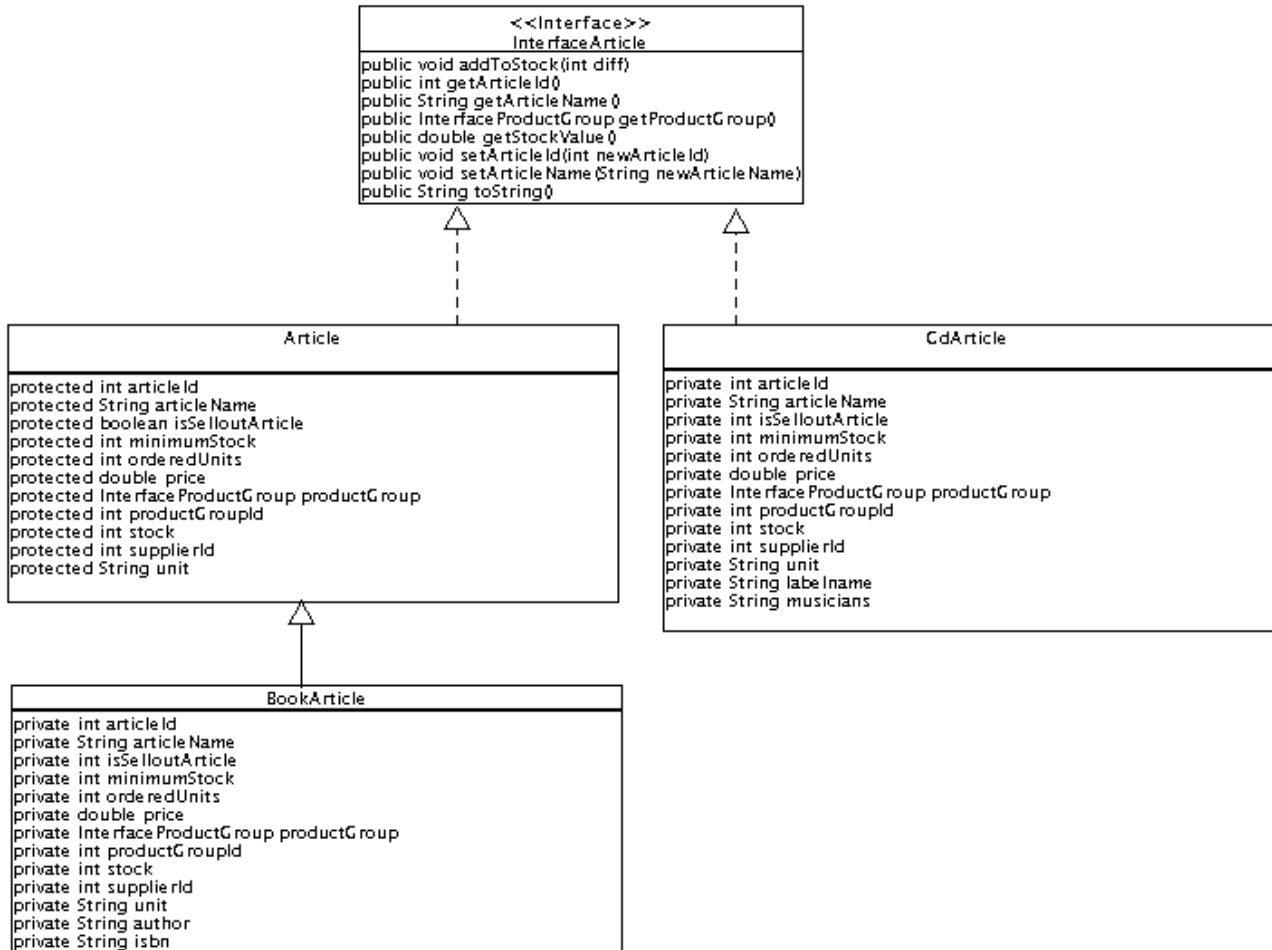
### 5.5.9. Advanced O/R Mapping Technique

#### 5.5.9.1. Extents and Polymorphism



Working with inheritance hierarchies is a common task in object oriented design and programming. Of course, any serious Java O/R tool must support inheritance and interfaces for persistent classes. There are many example classes for polymorphism in [OJB's JUnit TestSuite](#).

To demonstrate/explain *Extents* and *Polymorphism* we will look at a simple class hierarchy: There is a primary interface `InterfaceArticle`. This interface is implemented by `Article` and `CdArticle`. There is also a class `BookArticle` derived from `Article`. (See the following class diagram for details)



polymorphism.gif

## Polymorphism

OJB allows us to use interfaces, abstract or concrete base classes in queries, or in [type definitions of reference attributes](#). A Query against the interface `InterfaceArticle` must not only return objects of type `Article` but also of `CdArticle` and `BookArticle`!

The following [example](#) method searches for all objects implementing `InterfaceArticle` with an `articleName` equal to *Hamlet* (provided that the object mapping is correct, details will be described later). The Collection is e.g. filled with one matching `BookArticle` object.

```

public void testCollectionByQuery() throws Exception
{
    Criteria crit = new Criteria();
    crit.addEqualTo("articleName", "Hamlet");
    Query q = QueryFactory.newQuery(InterfaceArticle.class, crit);

    Collection result = broker.getCollectionByQuery(q);
}
  
```

Of course it is also possible to define [reference attributes](#) of an interface or baseclass type. The [example](#) class `Article` has a reference attribute ([1:1 reference](#)) of type `ProductGroup` and this can be a concrete/abstract class or interface.

## Extents

The query in the last example returned just one object. Now, imagine a query against the `InterfaceArticle` interface with no selecting criteria. OJB returns all the objects implementing `InterfaceArticle`. E.g. all `Article`, `BookArticle` and `CdArticles` objects.

In the following [example](#) the method prints out the collection of all `InterfaceArticle` objects:

```
public void testExtentByQuery() throws Exception
{
    // no criteria signals to omit a WHERE clause
    Query q = QueryFactory.newQuery(InterfaceArticle.class, null);
    Collection result = broker.getCollectionByQuery(q);

    System.out.println(
        "The InterfaceArticle Extent objects: " +result);
}
```

### Note:

The set of all instances of a class (whether living in memory or stored in a persistent medium) is called an **Extent** in ODMG and JDO terminology. OJB extends this notion slightly, as all objects which are subclasses of a concrete/abstract base class or implementing a given interface can be regarded as members of the base class or interface extent.

In our class diagram we find:

1. two simple *one-class-only* extents, `BookArticle` and `CdArticle`.
2. A compound extent `Article` containing all `Article` and `BookArticle` instances.
3. An interface extent containing all `Article`, `BookArticle` and `CdArticle` instances.

There is no extra coding necessary to define *extents*, but they have to be declared in the [metadata mapping file](#). The classes from the above [example](#) require the following declarations:

1. *one-class-only* extents require no declaration
2. A declaration for the base class `Article`, defining which classes are subclasses of `Article`:

```
<!-- Definitions for org.apache.ojb.ojb.broker.Article -->
<class-descriptor
  class="org.apache.ojb.broker.Article"
  proxy="false"
  table="Artikel"
  ...
>
  <extent-class class-ref="org.apache.ojb.broker.BookArticle" />
  ...
</class-descriptor>
```

3. A declaration for `InterfaceArticle`, defining which classes implement this interface:

```
<!-- Definitions for org.apache.ojb.broker.InterfaceArticle -->
<class-descriptor class="org.apache.ojb.broker.InterfaceArticle">
  <extent-class class-ref="org.apache.ojb.broker.Article" />
  <extent-class class-ref="org.apache.ojb.broker.CdArticle" />
  <!-- not needed to declare -->
  <!--<extent-class class-ref="org.apache.ojb.broker.BookArticle" />-->
</class-descriptor>
```

No need to declare `BookArticle` here, because it's a declared sub class of `Article`, so it's implicit declared by `Article` extent.

Why is it necessary to explicitly declare which classes implement an interface and which classes are derived from a base class?

Of course it is quite simple in Java to check whether a class implements a given interface or extends some other class. But sometimes it may not be appropriate to treat special implementors (e.g. proxies) as proper implementors.

Other problems might arise because a class may implement multiple interfaces, but is only allowed to be regarded as member of one extent.

In other cases it may be necessary to treat certain classes as implementors of an interface or as derived from a base even if they are not (we don't recommend to use this feature it's bad design, but if you don't have an alternative...).

As an example, you will find that the [ClassDescriptor](#) of abstract test class `org.apache.ojb.broker.CollectionTest$BookShelfItem` in the [OJB's Test Suite](#) contains an entry declaring class `org.apache.ojb.broker.CollectionTest$Candie` as a derived class:

```
<class-descriptor class="org.apache.ojb.broker.CollectionTest$BookShelfItem">
  <extent-class class-ref="org.apache.ojb.broker.CollectionTest$Book" />
  <extent-class class-ref="org.apache.ojb.broker.CollectionTest$DVD" />
  <!-- This class isn't a subclass of Book or DVD or a implementation of
    BookShelfItem, anyway it's possible to declare it as extent (but not
    recommended) -->
  <extent-class class-ref="org.apache.ojb.broker.CollectionTest$Candie" />
</class-descriptor>
```

#### Performance Tip

When using *extents* OJB will produce some overhead for each declared extent (e.g. execute a separate select-query for each extent or using complex table joins).

Thus it's important to avoid unnecessary *extent* declarations. If in the above [example](#) class `InterfaceArticle` is never used in queries, don't declare the extents for the implementing classes (`Article`, `CdArticle`). It's always possible to add additional *extents* in [mapping files](#).

#### 5.5.9.2. Mapping Inheritance Hierarchies

In the literature on object/relational mapping the problem of mapping inheritance hierarchies to RDBMS has been widely covered. In the following sections we will use a simple inheritance example to show the different inheritance mapping strategies.

Assume we have a base class `Employee` and class `Executive` extends `Employee`. Further on class `Manager` extends `Executive`.

mapping-inheritance.png

If we have to define database tables that have to contain these classes we have to choose one of the following solutions:

1. [Map each class of a hierarchy to a distinct table](#) and have all attributes from the base class in the derived class.
2. [Map class hierarchy onto one table](#).
3. [Map subclass fields of a hierarchy to a distinct table](#), but do not map super class fields to derived classes. Use joins to materialize over all tables to materialize objects.

OJB provides direct support for all three approaches.

#### Note:

But it's currently not recommended to mix mapping strategies within the same hierarchy !

In the following we demonstrate how these mapping approaches can be implemented by using OJB.

### Mapping Each Class of a Hierarchy to a Distinct Table (table per class)

This is the most simple solution. Just write a complete [ClassDescriptor](#) with [FieldDescriptors](#) for all of the attributes, including inherited attributes.

The classes of our [mapping example](#) would look like:

```
public class Employee implements Serializable
{
    private Integer id;
    private String name;

    public Employee()
    {
    }
    ....
    // getter/setter for id and ojbConcreteClass
}

public class Executive extends Employee
{
    private String department;
    ....
    // getter/setter
}

public class Manager extends Executive
{
    private int consortiumKey;
    ....
    // getter/setter
}
```

The [ClassDescriptors](#) include all fields of the representing java-class and each descriptor points to a different table:

```
<class-descriptor
  class="Employee"
  table="EMPLOYEE"
>
  <extent-class class-ref="Executive" />
  <field-descriptor
    name="id"
    column="ID"
    jdbc-type="INTEGER"
    primaryKey="true"
    autoincrement="true"
  />
  <field-descriptor
    name="name"
    column="NAME"
    jdbc-type="VARCHAR"
  />
</class-descriptor>

<class-descriptor
  class="Executive"
  table="EXECUTIVE"
>
  <extent-class class-ref="Manager" />
  <field-descriptor
    name="id"
    column="ID"
    jdbc-type="INTEGER"
```

```

        primaryKey="true"
        autoincrement="true"
    />
    <field-descriptor
        name="name"
        column="NAME"
        jdbc-type="VARCHAR"
    />
    <field-descriptor
        name="department"
        column="DEPARTMENT"
        jdbc-type="VARCHAR"
    />
</class-descriptor>

<class-descriptor
    class="Manager"
    table="MANAGER"
>
    <field-descriptor
        name="id"
        column="ID"
        jdbc-type="INTEGER"
        primaryKey="true"
        autoincrement="true"
    />
    <field-descriptor
        name="name"
        column="NAME"
        jdbc-type="VARCHAR"
    />
    <field-descriptor
        name="department"
        column="DEPARTMENT"
        jdbc-type="VARCHAR"
    />
    <field-descriptor
        name="consortiumKey"
        column="CONSORTIUM_KEY"
        jdbc-type="INTEGER"
    />
</class-descriptor>

```

The [extent-class element](#) is needed to declare the inheritance between the classes.

The DDL for the tables would look like:

```

CREATE TABLE EMPLOYEE
(
    ID          INT NOT NULL PRIMARY KEY,
    NAME       VARCHAR(150)
)
CREATE TABLE EXECUTIVE
(
    ID          INT NOT NULL PRIMARY KEY,
    NAME       VARCHAR(150),
    DEPARTMENT  VARCHAR(150)
)
CREATE TABLE MANAGER
(
    ID          INT NOT NULL PRIMARY KEY,
    NAME       VARCHAR(150),
    DEPARTMENT  VARCHAR(150),
    CONSORTIUM_KEY  INT
)

```

### Mapping Class Hierarchy on the Same Table (table per hierarchy)

Mapping several classes on one table works well under OJB. There is only one special situation

that needs some attention:

Storing `Employee`, `Executive` and `Manager` objects to this table works fine. But now consider a Query against the baseclass `Employee`. How can the correct type of the stored objects be determined?

OJB needs a *discriminator column* of type `CHAR` or `VARCHAR` that contains the class name to be used for instantiation. This column must be mapped on a special attribute `objConcreteClass`. On loading objects from the table, OJB checks this attribute and instantiates objects of this type.

**Note:**

The criterion for `objConcreteClass` is statically added to the query in class `QueryFactory` and it therefore appears in the select-statement for each extent. This means that mixing mapping strategies should be avoided.

The classes of our [mapping example](#) would look like:

```
public class Employee implements Serializable
{
    private Integer id;
    /**
     * This special attribute telling OJB which concrete class
     * this Object has.
     * NOTE: this attribute MUST be called objConcreteClass
     */
    private String objConcreteClass;
    private String name;

    public Employee()
    {
        // this guarantee that always the correct class name will be set
        this.objConcreteClass = this.getClass().getName();
    }
    ....
    // getter/setter for id and objConcreteClass
}

public class Executive extends Employee
{
    private String department;

    public Executive()
    {
        super();
    }
    ....
    // getter/setter
}

public class Manager extends Executive
{
    private int consortiumKey;

    public Manager()
    {
        super();
    }
    ....
    // getter/setter
}
```

**Note:**

Getter/setter for attribute `objConcreteClass` in base class `Employee` are only needed if OJB is forced to use [getter/setter for field access](#).

Here are the metadata mappings of our [mapping example](#):

```

<class-descriptor
  class="Employee"
  table="MANPOWER"
>
  <extent-class class-ref="Executive" />

  <field-descriptor
    name="id"
    column="ID"
    jdbc-type="INTEGER"
    primaryKey="true"
    autoincrement="true"
  />
  <field-descriptor
    name="objConcreteClass"
    column="CLASS_NAME"
    jdbc-type="VARCHAR"
  />
  <field-descriptor
    name="name"
    column="NAME"
    jdbc-type="VARCHAR"
  />
</class-descriptor>

<class-descriptor
  class="Executive"
  table="MANPOWER"
>
  <extent-class class-ref="Manager" />

  <field-descriptor
    name="id"
    column="ID"
    jdbc-type="INTEGER"
    primaryKey="true"
    autoincrement="true"
  />
  <field-descriptor
    name="objConcreteClass"
    column="CLASS_NAME"
    jdbc-type="VARCHAR"
  />
  <field-descriptor
    name="name"
    column="NAME"
    jdbc-type="VARCHAR"
  />
  <field-descriptor
    name="department"
    column="DEPARTMENT"
    jdbc-type="VARCHAR"
  />
</class-descriptor>

<class-descriptor
  class="Manager"
  table="MANPOWER"
>
  <field-descriptor
    name="id"
    column="ID"
    jdbc-type="INTEGER"
    primaryKey="true"
    autoincrement="true"
  />
  <field-descriptor
    name="objConcreteClass"
    column="CLASS_NAME"
    jdbc-type="VARCHAR"
  />

```

```

/>
<field-descriptor
  name="name"
  column="NAME"
  jdbc-type="VARCHAR"
/>
<field-descriptor
  name="department"
  column="DEPARTMENT"
  jdbc-type="VARCHAR"
/>
<field-descriptor
  name="consortiumKey"
  column="CONSORTIUM_KEY"
  jdbc-type="INTEGER"
/>
</class-descriptor>

```

The column CLASS\_NAME is used to store the concrete type of each object.

The [extent-class element](#) is needed to declare the inheritance between the classes.

The DDL for the table would look like:

```

CREATE TABLE MANPOWER
(
  ID          INT NOT NULL PRIMARY KEY,
  CLASS_NAME  VARCHAR(150)
  NAME       VARCHAR(150),
  DEPARTMENT  VARCHAR(150),
  CONSORTIUM_KEY INT
)

```

#### Implement your own Discriminator Handling

If you cannot provide such an additional column, but need to use some other means of indicating the type of each object you will require some additional programming:

You have to derive a Class from

`org.apache.ojb.broker.accesslayer.RowReaderDefaultImpl` and override the method `RowReaderDefaultImpl.selectClassDescriptor()` to implement your specific type selection mechanism. The code of the default implementation looks like follows:

```

protected ClassDescriptor selectClassDescriptor(Map row)
    throws PersistenceBrokerException
{
    // check if there is an attribute which tells us
    // which concrete class is to be instantiated
    ClassDescriptor result = m_cld;
    Class ojbConcreteClass = (Class) row.get(OJB_CONCRETE_CLASS_KEY);
    if(ojbConcreteClass != null)
    {
        result = m_cld.getRepository().getDescriptorFor(ojbConcreteClass);
        // if we can't find class-descriptor for concrete
        // class, something wrong with mapping
        if (result == null)
        {
            throw new PersistenceBrokerException(
                "Can't find class-descriptor for ojbConcreteClass '"
                + ojbConcreteClass + "', the main class was "
                + m_cld.getClassNameOfObject());
        }
    }
    return result;
}

```

After implementing your own [RowReader](#) you must edit the ClassDescriptor for the respective class in the XML repository to specify the usage of your RowReader Implementation:



```

<class-descriptor
  class="my.Object"
  table="MY_OBJECT"
  ...
  row-reader="my.own.RowReaderImpl"
  ...
>
...

```

You will learn more about RowReaders in [this section](#).

### Mapping Each Subclass to a Distinct Table (table per subclass)

This mapping strategy maps all subclass fields of a hierarchy to a distinct table (but do not map super class fields to derived class tables - except the [primary key](#) fields) and use joins to materialize over all tables to materialize the objects.

The classes of the inheritance hierarchy don't need any specific fields or settings, thus our [mapping example](#) java-classes look would look like the classes for the [table-per-class mapping](#).

The next code block contains the [class-descriptors](#) of our [mapping example](#).

```

<class-descriptor
  class="Employee"
  table="EMPLOYEE"
>
  <field-descriptor
    name="id"
    column="ID"
    jdbc-type="INTEGER"
    primaryKey="true"
    autoincrement="true"
  />
  <field-descriptor
    name="name"
    column="NAME"
    jdbc-type="VARCHAR"
  />
</class-descriptor>

<class-descriptor
  class="Executive"
  table="EXECUTIVE"
>
  <field-descriptor
    name="id"
    column="ID"
    jdbc-type="INTEGER"
    primaryKey="true"
  />
  <field-descriptor
    name="department"
    column="DEPARTMENT"
    jdbc-type="VARCHAR"
  />

  <reference-descriptor name="super"
    class-ref="Employee"
  >
    <foreignkey field-ref="id"/>
  </reference-descriptor>
</class-descriptor>

<class-descriptor
  class="Manager"
  table="MANAGER"
>

```

```

<field-descriptor
  name="id"
  column="ID"
  jdbc-type="INTEGER"
  primaryKey="true"
/>
<field-descriptor
  name="consortiumKey"
  column="CONSORTIUM_KEY"
  jdbc-type="INTEGER"
/>

<reference-descriptor name="super"
  class-ref="Executive"
>
  <foreignkey field-ref="id"/>
</reference-descriptor>
</class-descriptor>

```

The mapping for base class `Employee` is ordinary and we using a [autoincrement](#) primary key field.

In the subclasses `Executive` and `Manager` it's not allowed to use [autoincrement](#) primary keys, because OJB will automatically copy the primary keys of the base class to all subclasses.

As you can see this mapping needs a special [reference-descriptor](#) in the subclasses `Executive` and `Manager` that advises OJB to load the values for the inherited attributes from the super-class by a *JOIN* using the foreign key reference.

The `name="super"` attribute is not used to address an actual attribute of the super-class but as a marker keyword defining the *JOIN* to the super-class.

**Note:**

1. The [auto-xxx attributes](#) and the [proxy attribute](#) will be ignored when using the `super` keyword.
2. Be aware that this sample does not declare `Executive` or `Manager` to be an extent of `Employee`. Using `extends` here will lead to problems (instantiate the wrong class) because the primary key is not unique within the hierarchy defined in the [repository](#).

The DDL for the tables would look like:

```

CREATE TABLE EMPLOYEE
(
  ID          INT NOT NULL PRIMARY KEY,
  NAME       VARCHAR(150)
)
CREATE TABLE EXECUTIVE
(
  ID          INT NOT NULL PRIMARY KEY,
  DEPARTMENT VARCHAR(150)
)
CREATE TABLE MANAGER
(
  ID          INT NOT NULL PRIMARY KEY,
  CONSORTIUM_KEY INT
)

```

Attributes from the base- or superclasses can be used the same way as attributes of the target class when querying - e.g. for `Executive` or `Manager`. No [path-expression](#) is needed in this case. The following examples returns all `Executive` and `Manager` matching the criteria:

```

Criteria c = new Criteria();
// attribute defined in base class Employee
c.addEqualTo("name", "Kent");
// attribute defined in Executive
c.addEqualTo("department", "press");
Query q = QueryFactory.newQuery(Executive.class, c);
// returns all matching Executive and Manager instances
Collection result = broker.getCollectionByQuery(q);

```

### Table Per Subclass via Foreign Key

The above example is based on the assumption that the [primary key](#) attribute `Employee.id` and its underlying column `EMPLOYEE.ID` is also used as the foreign key attribute in the the subclasses.

Now let us consider a case where this is not possible, then it's possible to use an additional foreign key field/column in the subclass referencing the base-/superclass.

In this case the layout for class `Executive` would need an additional field `employeeFk` to store the foreign key reference to `Employee`.

To avoid the additional field in the subclass (if desired) we can use OJB's [anonymous field feature](#) to get everything working without the `employeeFk` attribute in subclass `Employee` (thus the [java classes](#) of our [mapping example](#)). We keep the [field-descriptor](#) for `employeeFk`, but declare it as an *anonymous field*. We just have to add an attribute `access="anonymous"` to the new field-descriptor `employeeFk`:

```
<class-descriptor
  class="Employee"
  table="EMPLOYEE"
>
  <field-descriptor
    name="id"
    column="ID"
    jdbc-type="INTEGER"
    primarykey="true"
    autoincrement="true"
  />
  <field-descriptor
    name="name"
    column="NAME"
    jdbc-type="VARCHAR"
  />
</class-descriptor>

<class-descriptor
  class="Executive"
  table="EXECUTIVE"
>
  <field-descriptor
    name="id"
    column="ID"
    jdbc-type="INTEGER"
    primarykey="true"
    autoincrement="true"
  />
  <field-descriptor
    name="department"
    column="DEPARTMENT"
    jdbc-type="VARCHAR"
  />
  <field-descriptor
    name="employeeFk"
    column="EMPLOYEE_FK"
    jdbc-type="INTEGER"
    access="anonymous"
  />
  <reference-descriptor name="super"
    class-ref="Employee"
  >
    <foreignkey field-ref="employeeFk" />
  </reference-descriptor>
</class-descriptor>

<class-descriptor
  class="Manager"
```

```

table="MANAGER"
>
<field-descriptor
  name="id"
  column="ID"
  jdbc-type="INTEGER"
  primarykey="true"
  autoincrement="true"
/>
<field-descriptor
  name="consortiumKey"
  column="CONSORTIUM_KEY"
  jdbc-type="INTEGER"
/>

<field-descriptor
  name="executiveFk"
  column="EXECUTIVE_FK"
  jdbc-type="INTEGER"
  access="anonymous"
/>

<reference-descriptor name="super"
  class-ref="Executive"
>
  <foreignkey field-ref="executiveFk"/>
</reference-descriptor>
</class-descriptor>

```

Now it's possible to use *autoincrement* primary key fields in all classes of the hierarchy (because they are decoupled from the inheritance references).

The *foreignkey*-element have to refer the new (anonymous) foreign-key field.

#### Warning:

The used primary keys (compound or single) have to be unique over the mapped class hierarchy to avoid object identity conflicts. Else it could happen e.g. when searching for a `Employee` with `id="42"` OJB may find a `Employee` and a `Executive` object with `id="42"`!

Thus it's problematic to use a [database identity columns](#) based [sequence-manager](#). In this case it's mandatory to use a different value scope (start index of identity column) for each class in hierarchy (e.g. 1 for `Employee`, 1000000000 for `Executive`, ...).

### 5.5.9.3. Using interfaces with OJB

Sometimes you may want to declare class descriptors for interfaces rather than for concrete classes. With OJB this is no problem, but there are a couple of things to be aware of, which are detailed in this section.

Consider this example hierarchy :

```

public interface A
{
    String getDesc();
}

public class B implements A
{
    /** primary key */
    private Integer id;
    /** sample attribute */
    private String desc;

    public String getDesc()
    {
        return desc;
    }
    public void setDesc(String desc)
    {
        this.desc = desc;
    }
}

```

```

}
}

public class C
{
    /** primary key */
    private Integer id;
    /** foreign key */
    private Integer aId;
    /** reference */
    private A obj;

    public void test()
    {
        String desc = obj.getDesc();
    }
}

```

Here, class C references the interface A rather than B. In order to make this work with OJB, four things must be done:

- All features common to all implementations of A are declared in the class descriptor of A. This includes references (with their foreignkeys) and collections.
- Since interfaces cannot have instance fields, it is necessary to use bean properties instead. This means that for every field (including collection fields), there must be accessors (a get method and, if the field is not marked as `access="readonly"`, a set method) declared in the interface.
- Since we're using bean properties, the appropriate `org.apache.ojb.broker.metadata.fieldaccess.PersistentField` implementation must be used (see [below](#)). This class is used by OJB to access the fields when storing/loading objects. Per default, OJB uses a direct access implementation (`org.apache.ojb.broker.metadata.fieldaccess.PersistentFieldDirectImpl`) which requires actual fields to be present. In our case, we need an implementation that rather uses the accessor methods. Since the `PersistentField` setting is (currently) global, you have to check whether there are accessors defined for every field in the metadata. If yes, then you can use the `org.apache.ojb.broker.metadata.fieldaccess.PersistentFieldIntrospectorImpl` otherwise you'll have to resort to the `org.apache.ojb.broker.metadata.fieldaccess.PersistentFieldAutoProxyImpl`, which determines for every field what type of field it is and then uses the appropriate strategy.
- If at some place OJB has to create an object of the interface, say as the result type of a query, then you have to specify `factory-class` and `factory-method` for the interface. OJB then uses the specified class and (static) method to create an uninitialized instance of the interface.

In our example, this would result in:

```

public interface A
{
    void setId(Integer id);
    Integer getId();
    void setDesc(String desc);
    String getDesc();
}

public class B implements A
{
    /** primary key */
    private Integer id;
    /** sample attribute */
    private String desc;

    public String getId()

```

```

    {
        return id;
    }
    public void setId(Integer id)
    {
        this.id = id;
    }
    public String getDesc()
    {
        return desc;
    }
    public void setDesc(String desc)
    {
        this.desc = desc;
    }
}

public class C
{
    /** primary key */
    private Integer id;
    /** foreign key */
    private Integer aId;
    /** reference */
    private A obj;

    public void test()
    {
        String desc = obj.getDesc();
    }
}

public class AFactory
{
    public static A createA()
    {
        return new B();
    }
}

```

The class descriptors would look like:

```

<class-descriptor
  class="A"
  table="A_TABLE"
  factory-class="AFactory"
  factory-method="createA"
>
  <extent-class class-ref="B"/>
  <field-descriptor
    name="id"
    column="ID"
    jdbc-type="INTEGER"
    primarykey="true"
    autoincrement="true"
  />
  <field-descriptor
    name="desc"
    column="DESC"
    jdbc-type="VARCHAR"
    length="100"
  />
</class-descriptor>

<class-descriptor
  class="B"
  table="B_TABLE"
>
  <field-descriptor
    name="id"
    column="ID"

```

```

        jdbc-type="INTEGER"
        primaryKey="true"
        autoincrement="true"
    />
    <field-descriptor
        name="desc"
        column="DESC"
        jdbc-type="VARCHAR"
        length="100"
    />
</class-descriptor>

<class-descriptor
    class="C"
    table="C_TABLE"
>
    <field-descriptor
        name="id"
        column="ID"
        jdbc-type="INTEGER"
        primaryKey="true"
        autoincrement="true"
    />
    <field-descriptor
        name="aId"
        column="A_ID"
        jdbc-type="INTEGER"
    />
    <reference-descriptor name="obj"
        class-ref="A">
        <foreignkey field-ref="aId" />
    </reference-descriptor>
</class-descriptor>

```

One scenario where you might run into problems is the use of interfaces for [nested objects](#). In the above example, we could construct such a scenario if we remove the descriptors for A and B, as well as the foreign key field aId from class C and change its class descriptor to:

```

<class-descriptor
    class="C"
    table="C_TABLE"
>
    <field-descriptor
        name="id"
        column="ID"
        jdbc-type="INTEGER"
        primaryKey="true"
        autoincrement="true"
    />
    <field-descriptor
        name="obj::desc"
        column="DESC"
        jdbc-type="VARCHAR"
        length="100"
    />
</class-descriptor>

```

The access to desc will work because of the usage of bean properties, but you will get into trouble when using [dynamic proxies](#) for C. Upon materializing an object of type C, OJB will try to create the instance for the field obj which is of type A. Of course, this is an interface but OJB won't check whether there is class descriptor for the type of obj (in fact there does not have to be one, and usually there isn't) because obj is not defined as a reference. As a result, OJB tries to instantiate an interface, which of course fails.

Currently, the only way to handle this is to write a [custom invocation handler](#) that knows how to create an object of type A.

#### 5.5.9.4. Change PersistentField Class

OJB supports a pluggable strategy to read and set the persistent attributes in the persistence capable classes. All strategy implementation classes have to implement the interface `org.apache.ojb.broker.metadata.fieldaccess.PersistentField`. OJB provide a few implementation classes which can be set in [OJB.properties](#) file:

```
# The PersistentFieldClass property defines the implementation class
# for PersistentField attributes used in the OJB MetaData layer.
# By default the best performing attribute/reflection based implementation
# is selected (PersistentFieldDirectAccessImpl).
#
# - PersistentFieldDirectAccessImpl
#   is a high-speed version of the access strategies.
#   It does not cooperate with an AccessController,
#   but accesses the fields directly. Persistent
#   attributes don't need getters and setters
#   and don't have to be declared public or protected
# - PersistentFieldPrivilegedImpl
#   Same as above, but does cooperate with AccessController and do not
#   suppress the java language access check (but is slow compared with direct
#   access).
# - PersistentFieldIntrospectorImpl
#   uses JavaBeans compliant calls only to access persistent attributes.
#   No Reflection is needed. But for each attribute xxx there must be
#   public getXxx() and setXxx() methods.
# - PersistentFieldDynaBeanAccessImpl
#   implementation used to access a property from a
#   org.apache.commons.beanutils.DynaBean.
# - PersistentFieldAutoProxyImpl
#   for each field determines upon first access how to access this particular
#   field
#   (directly, as a bean, as a dyna bean) and then uses that strategy
#
#PersistentFieldClass=org.apache.ojb.broker.metadata.fieldaccess.PersistentFieldDirectImpl
#PersistentFieldClass=org.apache.ojb.broker.metadata.fieldaccess.PersistentFieldPrivileged
#PersistentFieldClass=org.apache.ojb.broker.metadata.fieldaccess.PersistentFieldIntrospect
#PersistentFieldClass=org.apache.ojb.broker.metadata.fieldaccess.PersistentFieldDynaBeanIn
#PersistentFieldClass=org.apache.ojb.broker.metadata.fieldaccess.PersistentFieldAutoProxyI
#(DynaBean implementation does not support nested fields)
#
```

E.g. if the `PersistentFieldDirectImpl` is used there must be an attribute in the persistent class with this name, if the `PersistentFieldIntrospectorImpl` is used there must be a JavaBeans compliant property of this name. More info about the individual implementation can be found in [javadoc](#).

#### 5.5.9.5. How do anonymous keys work?

To play for safety it is mandatory to understand how this feature is working. In the HOWTO section is detailed described [how to use anonymous keys](#).

All involved classes can be found in `org.apache.ojb.broker.metadata.fieldaccess` package. The classes used for *anonymous keys* start with a `AnonymousXYZ.java` prefix.

Main class used for provide anonymous keys is

`org.apache.ojb.broker.metadata.fieldaccess.AnonymousPersistentField`.

Current implementation use an object identity based weak `HashMap`. The persistent object identity is used as key for the anonymous key value. The `(Anonymous)PersistentField` instance is associated with the *FieldDescriptor* declared in the repository.

This means that all anonymous key information will be lost when the object identity change, e.g. the persistent object will be de-/serialized or copied. In conjunction with 1:1 references this will be no problem, because OJB can use the referenced object to re-create the anonymous key information (FK to referenced object).

#### Warning:



The use of anonymous keys in 1:n references (FK to main object) or for PK fields is only valid when object identity does not change, e.g. use in single JVM without persistent object serialization and without persistent object copying.

### 5.5.9.6. Using Rowreader

RowReaders provide a callback mechanism that allows to interact with the OJB load mechanism. All implementation classes have to implement [interface RowReader](#).

You can specify the RowReader implementation in

- the [OJB.properties](#) file to set the standard used RowReader implementation

```
#-----
# RowReader
#-----
# Set the standard RowReader implementation. It is also possible to specify the
# RowReader on class-descriptor level.
RowReaderDefaultClass=org.apache.ojb.broker.accesslayer.RowReaderDefaultImpl
```

- within the [class-descriptor](#) to set the RowReader for a specific class.

RowReader setting on *class-descriptor* level will override the standard reader set in `OJB.properties` file. If neither a RowReader was set in `OJB.properties` file nor in *class-descriptor* was set, OJB use an default implementation.

To understand how to use them we must know some of the details of the load mechanism. To materialize objects from a rdbms OJB uses `RsIterators`, that are essentially wrappers to JDBC `ResultSets`. `RsIterators` are constructed from queries against the Database.

The method `RsIterator.next()` is used to materialize the next object from the underlying `ResultSet`. This method first checks if the underlying `ResultSet` is not yet exhausted and then delegates the construction of an Object from the current `ResultSet` row to the method `getObjectFromResultSet()`:

```
protected Object getObjectFromResultSet() throws PersistenceBrokerException
{
    if (getItemProxyClass() != null)
    {
        // provide m_row with primary key data of current row
        getQueryObject().getClassDescriptor().getRowReader()
            .readPkValuesFrom(getRsAndStmt().m_rs, getRow());
        // assert: m_row is filled with primary key values from db
        return getProxyFromResultSet();
    }
    else
    {
        // 0. provide m_row with data of current row
        getQueryObject().getClassDescriptor().getRowReader()
            .readObjectArrayFrom(getRsAndStmt().m_rs, getRow());
        // assert: m_row is filled from db

        // 1.read Identity
        Identity oid = getIdentityFromResultSet();
        Object result = null;

        // 2. check if Object is in cache. if so return cached version.
        result = getCache().lookup(oid);
        if (result == null)
        {
            // 3. If Object is not in cache
            // materialize Object with primitive attributes filled from
            // current row
            result = getQueryObject().getClassDescriptor()
                .getRowReader().readObjectFrom(getRow());
            // result may still be null!
        }
    }
}
```

```

        if (result != null)
        {
            synchronized (result)
            {
                getCache().enableMaterializationCache();
                getCache().cache(oid, result);
                // fill reference and collection attributes
                ClassDescriptor cld = getQueryObject().getClassDescriptor()
                .getRepository().getDescriptorFor(result.getClass());
                // don't force loading of reference
                final boolean unforced = false;
                // Maps ReferenceDescriptors to HashSets of owners
                getBroker().getReferenceBroker().retrieveReferences(result,
                cld, unforced);
                getBroker().getReferenceBroker().retrieveCollections(result,
                cld, unforced);
                getCache().disableMaterializationCache();
            }
        }
    }
    else // Object is in cache
    {
        ClassDescriptor cld = getQueryObject().getClassDescriptor()
        .getRepository().getDescriptorFor(result.getClass());
        // if refresh is required, update the cache instance from the db
        if (cld.isAlwaysRefresh())
        {
            getQueryObject().getClassDescriptor()
            .getRowReader().refreshObject(result,
            getRow());
        }
        getBroker().refreshRelationships(result, cld);
    }
    return result;
}
}
}

```

This method first uses a RowReader to instantiate a new object array and to fill it with primitive attributes from the current ResultSet row.

The RowReader to be used for a Class can be configured in the XML repository with the attribute [row-reader](#). If no RowReader is specified, the standard RowReader is used. The method `readObjectArrayFrom(...)` of this class looks like follows:

```

public void readObjectArrayFrom(ResultSet rs, ClassDescriptor cld, Map row)
{
    try
    {
        Collection fields = cld.getRepository().
            getFieldDescriptorsForMultiMappedTable(cld);
        Iterator it = fields.iterator();
        while (it.hasNext())
        {
            FieldDescriptor fmd = (FieldDescriptor) it.next();
            FieldConversion conversion = fmd.getFieldConversion();
            Object val = JdbcAccess.getObjectFromColumn(rs, fmd);
            row.put(fmd.getColumnName(), conversion.sqlToJava(val));
        }
    }
    catch (SQLException t)
    {
        throw new PersistenceBrokerException("Error reading from result set",t);
    }
}

```

In the second step OJB checks if there is already a cached version of the object to materialize. If so the cached instance is returned. If not, the object is fully materialized by first reading in primary attributes with the RowReader method `readObjectFrom(Map row, ClassDescriptor descriptor)` and in a second step by retrieving reference- and collection-attributes. The fully

materialized Object is then returned.

```
public Object readObjectFrom(Map row, ClassDescriptor descriptor)
    throws PersistenceBrokerException
{
    // allow to select a specific classdescriptor
    ClassDescriptor cld = selectClassDescriptor(row, descriptor);
    return buildWithReflection(cld, row);
}
```

By implementing your own RowReader you can hook into the OJB materialization process and provide additional features.

### Rowreader Example

Assume that for some reason we do not want to map a 1:1 association with a foreign key relationship to a different database table but read the associated object 'inline' from some columns of the master object's table. This approach is also called 'nested objects'. The section [nested objects](#) contains a different and much simpler approach to implement nested fields.

The class `org.apache.ojb.broker.ArticleWithStockDetail` has a `stockDetail` attribute, holding a reference to a `StockDetail` object. The class `StockDetail` is not declared in the XML repository. Thus OJB is not able to fill this attribute by ordinary mapping techniques.

We have to define a RowReader that does the proper initialization. The Class `org.apache.ojb.broker.RowReaderTestImpl` extends the `RowReaderDefaultImpl` and overrides the `readObjectFrom(...)` method as follows:

```
public Object readObjectFrom(Map row, ClassDescriptor cld)
{
    Object result = super.readObjectFrom(row, cld);
    if (result instanceof ArticleWithStockDetail)
    {
        ArticleWithStockDetail art = (ArticleWithStockDetail) result;
        boolean sellout = art.isSelloutArticle();
        int minimum = art.minimumStock();
        int ordered = art.orderedUnits();
        int stock = art.stock();
        String unit = art.unit();
        StockDetail detail = new StockDetail(sellout, minimum,
                                           ordered, stock, unit, art);
        art.stockDetail = detail;
        return art;
    }
    else
    {
        return result;
    }
}
```

To activate this RowReader the ClassDescriptor for the class `ArticleWithStockDetail` contains the following entry:

```
<class-descriptor
  class="org.apache.ojb.broker.ArticleWithStockDetail"
  table="Artikel"
  row-reader="org.apache.ojb.broker.RowReaderTestImpl"
>
```

### 5.5.9.7. Nested Objects

In the last section we discussed the usage of a user written RowReader to implement nested objects. This approach has several disadvantages.

1. It is necessary to write code and to have some understanding of OJB internals.

2. The user must take care that all nested fields are written back to the database on store.

This section shows that nested objects can be implemented without writing code, and without any further trouble just by a few settings in the repository.xml file.

The class `org.apache.ojb.broker.ArticleWithNestedStockDetail` has a `stockDetail` attribute, holding a reference to a `StockDetail` object. The class `StockDetail` is not declared in the XML repository as a first class entity class.

```
public class ArticleWithNestedStockDetail implements java.io.Serializable
{
    /**
     * this attribute is not filled through a reference lookup
     * but with the nested fields feature
     */
    protected StockDetail stockDetail;
    ...
}
```

The *StockDetail* class has the following layout:

```
public class StockDetail implements java.io.Serializable
{
    protected boolean isSelloutArticle;
    protected int minimumStock;
    protected int orderedUnits;
    protected int stock;
    protected String unit;
    ...
}
```

Only precondition to make things work is that *StockDetail* needs a default constructor. The nested fields semantics can simply declared by the following class- descriptor:

```
<class-descriptor
  class="org.apache.ojb.broker.ArticleWithNestedStockDetail"
  table="Artikel"
>
  <field-descriptor
    name="articleId"
    column="Artikel_Nr"
    jdbc-type="INTEGER"
    primarykey="true"
    autoincrement="true"
  />
  <field-descriptor
    name="articleName"
    column="Artikelname"
    jdbc-type="VARCHAR"
  />
  <field-descriptor
    name="supplierId"
    column="Lieferanten_Nr"
    jdbc-type="INTEGER"
  />
  <field-descriptor
    name="productGroupId"
    column="Kategorie_Nr"
    jdbc-type="INTEGER"
  />
  <field-descriptor
    name="stockDetail::unit"
  />
```

```

        column="Liefereinheit"
        jdbc-type="VARCHAR"
    />
    <field-descriptor
        name="price"
        column="Einzelpreis"
        jdbc-type="FLOAT"
    />
    <field-descriptor
        name="stockDetail::stock"
        column="Lagerbestand"
        jdbc-type="INTEGER"
    />
    <field-descriptor
        name="stockDetail::orderedUnits"
        column="BestellteEinheiten"
        jdbc-type="INTEGER"
    />
    <field-descriptor
        name="stockDetail::minimumStock"
        column="MindestBestand"
        jdbc-type="INTEGER"
    />
    <field-descriptor
        name="stockDetail::isSelloutArticle"
        column="Auslaufartikel"
        jdbc-type="INTEGER"
        conversion="org.apache.ojb.broker.accesslayer.conversions.Boolean2IntFieldConversion"
    />
</class-descriptor>

```

That's all! Just add nested fields by using `::` to specify attributes of the nested object. All aspects of storing and retrieving the nested object are managed by OJB.

#### 5.5.9.8. Instance Callbacks

OJB does provide transparent persistence. That is, persistent classes do not need to implement an interface or extend a persistent baseclass.

For certain situations it may be necessary to allow persistent instances to interact with OJB. This is supported by a simple instance callback mechanism.

The interface [org.apache.ojb.PersistenceBrokerAware](#) provides a set of methods that are invoked from the PersistenceBroker during operations on persistent instances:

#### Example

If you want that all persistent objects take care of CRUD operations performed by the PersistenceBroker you have to do the following steps:

1. let your persistent entity class implement the interface `PersistenceBrokerAware`.
2. provide empty implementations for all required methods.
3. implement the method `afterUpdate(PersistenceBroker broker)`, `afterInsert(PersistenceBroker broker)` and `afterDelete(PersistenceBroker broker)` to perform your intended logic.

In the following "for demonstration only code" you see a class `BaseObject` (all persistent objects extend this class) that does send a notification using a messenger object after object state change.

```

public abstract class BaseObject implements PersistenceBrokerAware
{
    private Messenger messenger;

    public void afterInsert(PersistenceBroker broker)
    {
        if(messenger != null)

```

```

        {
            messenger.send(this.getClass + " Object insert");
        }
    }
    public void afterUpdate(PersistenceBroker broker)
    {
        if(messenger != null)
        {
            messenger.send(this.getClass + " Object update");
        }
    }
    public void afterDelete(PersistenceBroker broker)
    {
        if(messenger != null)
        {
            messenger.send(this.getClass + " Object deleted");
        }
    }

    public void afterLookup(PersistenceBroker broker){}
    public void beforeDelete(PersistenceBroker broker){}
    public void beforeStore(PersistenceBroker broker){}

    public void setMessenger(Messenger messenger)
    {
        this.messenger = messenger;
    }
}

```

#### 5.5.9.9. Manageable Collection

In [1:n](#) or [m:n](#) relations, OJB can handle `java.util.Collection` as well as user defined collection classes as collection attributes in persistent classes. See [collection-descriptor.collection-class](#) attribute for more information.

In order to collaborate with the OJB mechanisms these collection must provide a minimum protocol as defined by this interface `org.apache.ojb.broker.ManageableCollection`.

```

public interface ManageableCollection extends java.io.Serializable
{
    /**
     * add a single Object to the Collection. This method is used during reading
     * Collection elements from the database. Thus it is is save to cast
     * anObject
     * to the underlying element type of the collection.
     */
    void ojbAdd(Object anObject);

    /**
     * adds a Collection to this collection. Used in reading Extents from the
     * Database. Thus it is save to cast otherCollection to this.getClass().
     */
    void ojbAddAll(ManageableCollection otherCollection);

    /**
     * returns an Iterator over all elements in the collection. Used during
     * store and
     * delete Operations.
     * If the implementor does not return an iterator over ALL elements, OJB
     * cannot
     * store and delete all elements properly.
     */
    Iterator ojbIterator();

    /**
     * A callback method to implement 'removal-aware' (track removed objects and
     * delete

```

```

    * them by its own) collection implementations.
    */
    public void afterStore(PersistenceBroker broker) throws
PersistenceBrokerException;
}

```

The methods have a prefix "obj" that indicates that these methods are "technical" methods, required by OJB and not to be used in business code.

In package `org.apache.obj.broker.util.collections` can be found a bunch of pre-defined implementations of `org.apache.obj.broker.ManageableCollection`.

More info about [which collection class to used here](#).

### Types Allowed for Implementing 1:n and m:n Associations

OJB supports different Collection types to implement 1:n and m:n associations. OJB detects the used type automatically, so there is no need to declare it in the repository file. There is also no additional programming required. The following types are supported:

1. `java.util.Collection`, `java.util.List`, `java.util.Vector` as in the example above. Internally OJB uses `java.util.Vector` to implement collections.
2. Arrays (see the file `ProductGroupWithArray`).
3. User-defined collections (see the file `ProductGroupWithTypedCollection`). A typical application for this approach are typed Collections.

Here is some sample code from the Collection class `ArticleCollection`. This Collection is typed, i.e. it accepts only `InterfaceArticle` objects for adding and will return `InterfaceArticle` objects with `get(int index)`. To let OJB handle such a user-defined Collection it **must** implement the callback interface `ManageableCollection` and the typed collection class must be declared in the *collection-descriptor* using the *collection-class* attribute.

`ManageableCollection` provides hooks that are called by OJB during object materialization, updating and deletion.

```

public class ArticleCollection implements ManageableCollection,
                                         java.io.Serializable
{
    private Vector elements;

    public ArticleCollection()
    {
        super();
        elements = new Vector();
    }

    public void add(InterfaceArticle article)
    {
        elements.add(article);
    }

    public InterfaceArticle get(int index)
    {
        return (InterfaceArticle) elements.get(index);
    }

    /**
     * add a single Object to the Collection. This method is
     * used during reading Collection elements from the
     * database. Thus it is is save to cast anObject
     * to the underlying element type of the collection.
     */
    public void objAdd(java.lang.Object anObject)
    {
        elements.add((InterfaceArticle) anObject);
    }
}

```

```

/**
 * adds a Collection to this collection. Used in reading
 * Extends from the Database.
 * Thus it is save to cast otherCollection to this.getClass().
 */
public void ojbAddAll(
    ojb.broker.ManageableCollection otherCollection)
{
    elements.addAll(
        ((ArticleCollection) otherCollection).elements);
}

/**
 * returns an Iterator over all elements in the collection.
 * Used during store and delete Operations.
 */
public java.util.Iterator ojbIterator()
{
    return elements.iterator();
}
}

```

And the collection-descriptor have to declare this class:

```

<collection-descriptor
name="allArticlesInGroup"
element-class-ref="org.apache.ojb.broker.Article"
collection-class="org.apache.ojb.broker.ArticleCollection"
auto-retrieve="true"
auto-update="false"
auto-delete="true"
>
<inverse-foreignkey field-ref="productGroupId"/>
</collection-descriptor>

```

#### Which collection-class type should be used?

[Earlier in this section](#) the `org.apache.ojb.broker.ManageableCollection` was introduced. Now we talk about which type to use.

By default OJB use a *removal-aware* collection implementation. These implementations (classes prefixed with *Removal...*) track removal and addition of elements.

This tracking allow the PersistenceBroker to **delete elements** from the database that have been removed from the collection before a `PB.store()` operation occurs.

This default behaviour is **undesired** in some cases:

- In [m:n relations](#), e.g. between *Movie* and *Actor* class. If an Actor was removed from the Actor collection of a Movie object expected behaviour was that the Actor be removed from the [indirection table](#), but not the Actor itself. Using a removal aware collection will remove the Actor too. In that case a simple manageable collection is recommended by set e.g. `collection-class="org.apache.ojb.broker.util.collections.ManageableArray"` in collection-descriptor.
- In [1:n relations](#) when the n-side objects be removed from the collection of the main object, but we don't want to remove them itself (be careful with this, because the FK entry of the main object still exists - more info about [linking here](#)).

#### 5.5.9.10. Customizing collection queries

Customizing the query used for collection retrieval allows a **developer** to take full control of collection mechanism. For example only children having a certain attribute should be loaded. This is achieved by a QueryCustomizer defined in the collection-descriptor of a relationship:

```

<collection-descriptor

```



```

    name="allArticlesInGroup"
    ...
  >
  <inverse-foreignkey field-ref="productGroupId"/>

  <query-customizer
    class="org.apache.ojb.broker.accesslayer.QueryCustomizerDefaultImpl">
    <attribute
      attribute-name="attr1"
      attribute-value="value1"
    />
  </query-customizer>
</collection-descriptor>

```

The query customizer must implement the interface `org.apache.ojb.broker.accesslayer.QueryCustomizer`. This interface defines the single method below which is used to customize (or completely rebuild) the query passed as argument. The interpretation of attribute-name and attribute-value read from the collection-descriptor is up to your implementation.

```

/**
 * Return a new Query based on the original Query, the
 * originator object and the additional Attributes
 *
 * @param anObject the originator object
 * @param aBroker the PersistenceBroker
 * @param aCod the CollectionDescriptor
 * @param aQuery the original 1:n-Query
 * @return Query the customized 1:n-Query
 */
public Query customizeQuery(Object anObject,
                           PersistenceBroker aBroker,
                           CollectionDescriptor aCod, Query aQuery);

```

The class `org.apache.ojb.broker.accesslayer.QueryCustomizerDefaultImpl` provides a default implementation without any functionality, it simply returns the query.

### 5.5.9.11. Metadata runtime changes

This was described in [metadata section](#).

## 5.5.10. OJB Queries

### 5.5.10.1. Introduction

This tutorial describes the use of the different queries mechanisms. The sample code shown here is taken mainly from JUnit test classes. The junit test source can be found under `[db-ojb]/src/test` in the source distribution.

### 5.5.10.2. Query by Criteria

In this section you will learn how to use the query by criteria. The classes are located in the package `org.apache.ojb.broker.query`. Using query by criteria you can either [query for whole objects](#) (ie. person) or you can use [report queries](#) returning row data.

A query consists mainly of the following parts:

1. the class of the objects to be retrieved
2. a list of criteria
3. a DISTINCT flag
4. additional ORDER BY and GROUP BY

OJB offers a `QueryFactory` to create a new `Query`. Although the constructors of the query classes are public using the `QueryFactory` is the preferred way to create a new query.

```
Query q = QueryFactory.newQuery(Person.class, crit);
```

To create a `DISTINCT-Query`, simply add **true** as third parameter.

```
Query q = QueryFactory.newQuery(Person.class, crit, true);
```

Each criterion stands for a column in the `SQL-WHERE`-clause.

```
Criteria crit = new Criteria();
crit.addEqualTo("upper(firstname)", "TOM");
crit.addEqualTo("lastname", "hanks");
Query q = QueryFactory.newQuery(Person.class, crit);
```

This query will generate an SQL statement like this:

```
SELECT ... FROM PERSON WHERE upper(FIRSTNAME) = "TOM" AND LASTNAME = "hanks";
```

OJB supports **functions** in field criteria ie. `upper(firstname)`. When converting a field name to a database column name, the function is added to the generated sql. OJB does not and can not verify the correctness of the specified function, an illegal function will produce an `SQLException`.

### Query Criteria

OJB provides selection criteria for almost any SQL-comparator. In most cases you do not have to deal directly with the implementing classes like `EqualToCriteria`. The `Criteria` class provides factory methods for the appropriate classes. There are four kinds of factory methods:

- create criteria to compare a field to a value: ie. `addEqualTo("firstname", "tom");`
- create criteria to compare a field to another field: ie. `addEqualToField("firstname", "other_field");`
- create criteria to check null value: ie. `addIsNull("firstname");`
- create a raw sql criteria: ie: `addSql("REVERSE(name) like 're%'");`

The following list shows some of the factory methods to compare a field to a value:

- `addEqualTo`
- `addLike`
- `addGreaterOrEqualThan`
- `addGreaterThan`
- `addLike`
- `addBetween`, this methods has two value parameters
- `addIn`, this method uses a `Collection` as value parameter
- and of course there negative forms

This list shows some factory methods to compare a field to another field, all those methods end on `...field`:

- `addEqualToField`
- `addGreaterThanField`
- and of course there negative forms

### in / not in

Some databases limit the number of parameters in an `IN`-statement.

If the limit is reached OJB will split up the `IN`-Statement into multiple Statements, the limit is set to 3 for the following sample:

```
SELECT ... FROM Artikel A0 WHERE A0.Kategorie_Nr IN ( ? , ? , ? )
OR A0.Kategorie_Nr IN ( ? , ? ) ORDER BY 7 DESC
```

The IN-limit for prefetch can be defined in OJB.properties:

```
...
# The SqlInLimit entry limits the number of values in IN-sql
# statement, -1 for no limits. This hint is used in Criteria.
SqlInLimit=200
...
```

#### and / or

All selection criteria added to a criteria set using the above factory methods will be **ANDed** in the WHERE-clause. To get an **OR** combination two criteria sets are needed. These sets are combined using `addOrCriteria`:

```
Criteria crit1 = new Criteria();
crit1.addLike("firstname", "%o%");
crit1.addLike("lastname", "%m%");
Criteria crit2 = new Criteria();
crit2.addEqualTo("firstname", "hank");

crit1.addOrCriteria(crit2);
Query q = QueryFactory.newQuery(Person.class, crit1);

Collection results = broker.getCollectionByQuery(q);
```

This query will generate an SQL statement like this:

```
SELECT ... WHERE (FIRSTNAME LIKE "%o%") AND LASTNAME
LIKE "%m%" OR FIRSTNAME = "hank"
```

#### negating the criteria

A criteria can be negated to obtain **NOT** in the WHERE-clause:

```
Criteria crit1 = new Criteria();
crit1.addLike("firstname", "%o%");
crit1.addLike("lastname", "%m%");
crit1.setNegative(true);

Collection results = broker.getCollectionByQuery(q);
```

This query will generate an SQL statement like this:

```
SELECT ... WHERE NOT (FIRSTNAME LIKE "%o%" AND LASTNAME LIKE "%m%")
```

#### ordering and grouping

The following methods of `QueryByCriteria` are used for ordering and grouping:

- `addOrderByAscending(String anAttributeName);`
- `addOrderByDescending(String anAttributeName);`
- `addGroupBy(String anAttributeName);` this method is used for [report queries](#)

You can of course have multiple order by and group by clauses, simply repeat the `addOrderBy`.

```
crit = new Criteria();
query = new QueryByCriteria(Person.class, crit);
query.addOrderByDescending("id");
query.addOrderByAscending("lastname");
broker.getCollectionByQuery(query);
```

The code snippet will query all Persons and order them by **attribute** "id" descending and "lastname" ascending. The query will produce the following SQL-statement using column numbers

in the ORDER BY clause:

```
SELECT A0.ID,A0.FIRSTNAME,A0.LASTNAME FROM
PERSON A0 ORDER BY 1 DESC, 3
```

When you use the **column** name "LASTNAME" instead of the **attribute** name "lastname" (query.addOrderBy("LASTNAME");), an additional column named "LASTNAME" without alias will be added.

```
SELECT A0.ID,A0.FIRSTNAME,A0.LASTNAME, LASTNAME FROM
PERSON A0 ORDER BY 1 DESC, 4
```

If there are multiple tables with a column "LASTNAME" the SQL-Statement will produce an error, so it's better to always use attribute names.

## subqueries

Subqueries can be used instead of values in selection criteria. The subquery should thus be a ReportQuery.

The following example queries all articles having a price greater or equal than the average price of articles named 'A%':

```
ReportQueryByCriteria subQuery;
Criteria subCrit = new Criteria();
Criteria crit = new Criteria();

subCrit.addLike("articleName", "A%");
subQuery = QueryFactory.newReportQuery(Article.class, subCrit);
subQuery.setAttributes(new String[] { "avg(price)" });

crit.addGreaterOrEqualThan("price", subQuery);
Query q = QueryFactory.newQuery(Article.class, crit);

Collection results = broker.getCollectionByQuery(q);
```

It's also possible to build a subquery with attributes referencing the enclosing query. These attributes have to use a special prefix **Criteria.PARENT\_QUERY\_PREFIX**.

The following example queries all product groups having more than 10 articles:

```
ReportQueryByCriteria subQuery;
Criteria subCrit = new Criteria();
Criteria crit = new Criteria();

subCrit.addEqualToField("productGroupId", Criteria.PARENT_QUERY_PREFIX +
"groupId");
subQuery = QueryFactory.newReportQuery(Article.class, subCrit);
subQuery.setAttributes(new String[] { "count(productGroupId)" });

crit.addGreaterThan(subQuery, "10"); // MORE than 10 articles
crit.addLessThan("groupId", new Integer(987654));
Query q = QueryFactory.newQuery(ProductGroup.class, crit);

Collection results = broker.getCollectionByQuery(q);
```

### Note:

Subqueries are not extent aware. Thus it's not possible to use an abstract class or an interface as search class of a subquery.

## joins

Joins resulting from **path expressions** ("relationship.attribute") in criteria are automatically handled by OJB. Path expressions are supported for all relationships 1:1, 1:n and m:n (decomposed and non-decomposed) and can be nested.

The following sample looks for all articles belonging to the product group "Liquors". Article and product group are linked by the relationship "productGroup" in class Article:

```
<!-- Definitions for org.apache.ojb.ojb.broker.Article -->
  <class-descriptor
    class="org.apache.ojb.broker.Article"
    proxy="dynamic"
    table="Artikel"
  >
    ...
    <reference-descriptor
      name="productGroup"
      class-ref="org.apache.ojb.broker.ProductGroup"
    >
      <foreignkey field-ref="productGroupId"/>
    </reference-descriptor>
  </class-descriptor>

  <class-descriptor
    class="org.apache.ojb.broker.ProductGroup"
    proxy="org.apache.ojb.broker.ProductGroupProxy"
    table="Kategorien"
  >
    ...
    <field-descriptor
      name="groupName"
      column="KategorieName"
      jdbc-type="VARCHAR"
    />
  </class-descriptor>
```

The path expression includes the 1:1 relationship "productGroup" and the attribute "groupName":

```
Criteria crit = new Criteria();
crit.addEqualTo("productGroup.groupName", "Liquors");
Query q = QueryFactory.newQuery(Article.class, crit);
Collection results = broker.getCollectionByQuery(q);
```

If path expressions refer to a class having **extents**, the tables of the extent classes participate in the JOIN and the criteria is ORed. The shown sample queries all ProductGroups having an Article named 'F%'. The path expression 'allArticlesInGroup' refers to the class Articles which has two extents: Books and CDs.

```
Criteria crit = new Criteria();
crit.addLike("allArticlesInGroup.articleName", "F%");
QueryByCriteria q = QueryFactory.newQuery(ProductGroup.class, crit, true);
Collection results = broker.getCollectionByQuery(q);
```

This sample produces the following SQL:

```
SELECT DISTINCT A0.KategorieName,A0.Kategorie_Nr,A0.Beschreibung
FROM Kategorien A0
INNER JOIN Artikel A1 ON A0.Kategorie_Nr=A1.Kategorie_Nr
LEFT OUTER JOIN BOOKS A1E0 ON A0.Kategorie_Nr=A1E0.Kategorie_Nr
LEFT OUTER JOIN CDS A1E1 ON A0.Kategorie_Nr=A1E1.Kategorie_Nr
WHERE A1.Artikelname LIKE 'F%' OR
A1E0.Artikelname LIKE 'F%' OR
A1E1.Artikelname LIKE 'F%'
```

OJB tries to do it's best to automatically use **outer** joins where needed. This is currently the case for classes having extents and ORed criteria. But you can force the SQLGenerator to use outer joins where you find it useful.

This is done by the method `QueryByCriteria#setPathOuterJoin(String)`.

```

ReportQueryByCriteria query;
Criteria crit;
Iterator result1, result2;

crit = new Criteria();

query = new ReportQueryByCriteria(Person.class, crit);
query.setAttributes(new String[] { "id", "name", "vorname", "sum(konti.saldo)"
});
query.addGroupBy(new String[]{ "id", "name", "vorname" });

result1 = broker.getReportQueryIteratorByQuery(query);

query.setPathOuterJoin("konti");
result2 = broker.getReportQueryIteratorByQuery(query);

```

The first query will use an inner join for relationship "konti", the second an outer join.

The whole path is evaluated, thus for a multi segment path (ie. owner.address) you'll have to set an outer join path for each segment that needs an outer join.

Given 'Account' - 'Owner' - 'Adress' and a query looking for 'Account'. To force an outer join from 'Owner' to 'Adress' setPathOuterJoin('owner.address') does the trick. In case you also need an outer join between 'Account' and 'Owner' an additional setPathOuterJoin('owner') is needed.

#### user defined alias

This feature allows to have **multiple** aliases for the same table. The standard behaviour of OJB is to build one alias for one relationship.

Suppose you have two classes Issue and Keyword and there is a 1:N relationship between them. Now you want to retrieve Issues by querying on Keywords. Suppose you want to retrieve all Issues with keywords 'JOIN' and 'ALIAS'. If these values are stored in the attribute 'value' of Keyword, OJB generates a query that contains " A1.value = 'JOIN' AND A1.value = 'ALIAS' " in the where-clause. Obviously, this will not work, no hits will occur because A1.value can not have more than 1 value at the time !

For the examples below, suppose you have the following classes (pseudo-code):

```

class Container
    int id
    Collection allAbstractAttributes

class AbstractAttribute
    int id
    int ref_id
    String name
    String value
    Collection allAbstractAttributes

```

OJB maps these classes to separate tables where it maps allAbstractAttributes using a collectiondescriptor to AbstractAttribute using ref\_id as inverse foreignkey on Container for the collection descriptor.

For demo purposes : AbstractAttribute also has a collection of abstract attributes.

```

Criteria crit1 = new Criteria();
crit1.setAlias("company"); // set an alias
crit1.addEqualTo("allAbstractAttributes.name", "companyName");
crit1.addEqualTo("allAbstractAttributes.value", "iBanx");

Criteria crit2 = new Criteria();
crit2.setAlias("contact"); // set an alias
crit2.addEqualTo("allAbstractAttributes.name", "contactPerson");
crit2.addLike("allAbstractAttributes.value", "janssen");

```

```
Criteria crit3 = new Criteria();
crit3.addEqualTo("allAbstractAttributes.name", "size");
crit3.addGreaterThan("allAbstractAttributes.value", new Integer(500));

crit1.addAndCriteria(crit2);
crit1.addAndCriteria(crit3);

q = QueryFactory.newQuery(Container.class, crit1);
q.addOrderBy("company.value"); // user alias
```

The generated query will be as follows. Note that the alias name 'company' does not show up in the SQL.

```
SELECT DISTINCT A0.ID, A1.VALUE
FROM CONTAINER A0 INNER JOIN ABSTRACT_ATTRIBUTE A1
  ON A0.ID=A1.REF_ID
  INNER JOIN ABSTRACT_ATTRIBUTE A2
  ON A0.ID=A2.REF_ID
  INNER JOIN ABSTRACT_ATTRIBUTE A3
  ON A0.ID=A3.REF_ID
WHERE (( A0.NAME = 'companyName' ) AND (A0.VALUE = 'iBanx' )) AND
      (( A1.NAME = 'contactPerson' ) AND (A1.VALUE LIKE '%janssen%' )) AND
      (( A2.NAME = 'size' ) AND (A2.VALUE = '500' ))
ORDER BY 2
```

The next example uses a report query.

```
Criteria crit1 = new Criteria();
crit1.setAlias("ALIAS1");
crit1.addEqualTo("allAbstractAttributes.allAbstractAttributes.name", "xxxx");
crit1.addEqualTo("allAbstractAttributes.allAbstractAttributes.value", "hello");

Criteria crit2 = new Criteria();
crit2.setAlias("ALIAS2");
crit2.addEqualTo("allAbstractAttributes.name", "yyyy");
crit2.addLike("allAbstractAttributes.value", "");

crit1.addAndCriteria(crit2);

q = QueryFactory.newReportQuery(Container.class, crit1);

String[] cols = { id, "ALIAS2.name", "ALIAS2.name", "ALIAS1.name", "ALIAS1.name"
};
q.setAttributes(cls);
```

The generated query will be:

```
SELECT DISTINCT A0.ID, A1.NAME, A1.VALUE, A2.NAME, A2.VALUE
FROM CONTAINER A0 INNER JOIN ABSTRACT_ATTRIBUTE A1
  ON A0.ID=A1.REF_ID
  INNER JOIN ABSTRACT_ATTRIBUTE A2
  ON A1.ID=A2.REF_ID
WHERE (( A2.NAME = 'xxxx' ) AND (A2.VALUE = 'hello' )) AND
      (( A1.NAME = 'yyyy' ) AND (A2.VALUE LIKE '%%' )) AND
ORDER BY 2
```

#### Note:

When you define an alias for a criteria, you have to make sure that *all* attributes used in this criteria belong to the *same* class. If you break this rule OJB will probably use a wrong ClassDescriptor to resolve your attributes !

#### class hints

This feature allows the user to specify which class of an extent to use for a path-segment. The standard behaviour of OJB is to use the base class of an extent when it resolves a path-segment.

In the following sample the path **allArticlesInGroup** points to class Article, this is defined in the repository.xml. Assume we are only interested in ProductGroups containing CdArticles performed

by Eric Clapton or Books authored by Eric Clapton, a class hint can be defined for the path. This hint is defined by:

```
Criteria#addPathClass("allArticlesInGroup", CdArticle.class);
```

```
//
// find a ProductGroup with a CD or a book by a particular artist
//
String artistName = new String("Eric Clapton");
crit1 = new Criteria();
crit1.addEqualTo("allArticlesInGroup.musicians", artistName);
crit1.addPathClass("allArticlesInGroup", CdArticle.class);

crit2 = new Criteria();
crit2.addEqualTo("allArticlesInGroup.author", artistName);
crit2.addPathClass("allArticlesInGroup", BookArticle.class);

crit1.addOrCriteria(crit2);

query = new QueryByCriteria(ProductGroup.class, crit1);
broker.getObjectByQuery(query);
```

#### Note:

This feature is also available in class QueryByCriteria but using it on Criteria-level provides additional flexibility. QueryByCriteria#addPathClass is only useful for ReportQueries to limit the class of the selected columns.

### prefetched relationships

This feature can help to minimize the number of queries when reading objects with relationships. In our Testcases we have ProductGroups with a one to many relationship to Articles. When reading the ProductGroups one query is executed to get the ProductGroups and for **each** ProductGroup another query is executed to retrieve the Articles.

With prefetched relationships OJB tries to read all Articles belonging to the ProductGroups in **one** query. See further down why one query is not always possible.

```
Criteria crit = new Criteria();
crit.addLessOrEqualThan("groupId", new Integer(5));

QueryByCriteria q = QueryFactory.newQuery(ProductGroup.class, crit);
q.addOrderByDescending("groupId");
q.addPrefetchedRelationship("allArticlesInGroup");

Collection results = broker.getCollectionByQuery(q);
```

The first query reads all matching ProductGroups:

```
SELECT ... FROM Kategorien A0 WHERE
A0.Kategorie_Nr <= ? ORDER BY 3 DESC
```

The second query retrieves Articles belonging to the ProductGroups read by the first query:

```
SELECT ... FROM Artikel A0 WHERE A0.Kategorie_Nr
IN ( ? , ? , ? , ? , ? ) ORDER BY 7 DESC
```

After reading all Articles they are associated with their ProductGroup.

#### Note:

This function is not yet supported for relationships using Arrays.

Some databases limit the number of parameters in an IN-statement. If the limit is reached OJB will split up the second query into multiple queries, the limit is set to 3 for the following sample:

```
SELECT ... FROM Artikel A0 WHERE A0.Kategorie_Nr
```



```
IN ( ? , ? , ? ) ORDER BY 7 DESC
SELECT ... FROM Artikel A0 WHERE A0.Kategorie_Nr
IN ( ? , ? ) ORDER BY 7 DESC
```

The IN-limit for prefetch can be defined in OJB.properties [SqlInLimit](#).

### querying for objects

OJB queries return **complete** objects, that means all instance variables are filled and all 'auto-retrieve' relationships are loaded. Currently there's no way to retrieve partially loaded objects (ie. only first- and lastname of a person).

More info about [manipulation of metadata setting here](#).

### Report Queries

Report queries are used to retrieve row data, not 'real' business objects. A row is an array of Object. With these queries you can define what attributes of an object you want to have in the row. The attribute names may also contain path expressions like 'owner.address.street'. To define the attributes use ReportQuery `#setAttributes(String[] attributes)`.

The following ReportQuery retrieves the name of the ProductGroup, the name of the Article etc. for all Articles named like "C%":

```
Criteria crit = new Criteria();
Collection results = new Vector();
crit.addLike("articleName", "C%");
ReportQueryByCriteria q = QueryFactory.newReportQuery(Article.class, crit);
q.setAttributes(new String[] { "productGroup.groupName", "articleId",
"articleName", "price" });

Iterator iter = broker.getReportQueryIteratorByQuery(q);
```

The ReportQuery returns an Iterator over a Collection of Object[4] ([String, Integer, String, Double]).

### Limitations of Report Queries

ReportQueries should not be used with columns referencing classes with extents. Assume we want to select all ProductGroups and summarize the amount and prize of items in stock per group. The class Article referenced by `allArticlesInGroup` has the extents Books and CDs.

```
Criteria crit = new Criteria();
Collection results = new Vector();
ReportQueryByCriteria q = QueryFactory.newReportQuery(ProductGroup.class, crit);
q.setAttributes(new String[] { "groupName", "sum(allArticlesInGroup.stock)",
"sum(allArticlesInGroup.price)" });
q.addGroupBy("groupName");

Iterator iter = broker.getReportQueryIteratorByQuery(q);
```

The ReportQuery looks quite reasonable, but it will produce an SQL not suitable for the task:

```
SELECT A0.KategorieName, sum(A1.Lagerbestand), sum(A1.Einzelpreis)
FROM Kategorien A0
LEFT OUTER JOIN artikel A1 ON A0.Kategorie_Nr=A1.Kategorie_Nr
LEFT OUTER JOIN books A1E2 ON A0.Kategorie_Nr=A1E2.Kategorie_Nr
LEFT OUTER JOIN cds A1E1 ON A0.Kategorie_Nr=A1E1.Kategorie_Nr
GROUP BY A0.KategorieName
```

This SQL will select the columns "Lagerbestand" and "Einzelpreis" from one extent only, and for ProductGroups having Articles, Books and CDs the result is wrong!

As a workaround the query can be "reversed". Instead of selection the ProductGroup we go for the

Articles:

```
Criteria crit = new Criteria();
Collection results = new Vector();
ReportQueryByCriteria q = QueryFactory.newReportQuery(Article.class, crit);
q.setAttributes(new String[] { "productGroup.groupName", "sum(stock)",
"sum(price)" });
q.addGroupBy("productGroup.groupName");
```

This ReportQuery executes the following three selects (one for each concrete extent) and produces better results.

```
SELECT  A1.KategorieName, sum(A0.Lagerbestand), sum(A0.Einzelpreis)
FROM artikel A0
INNER JOIN Kategorien A1 ON A0.Kategorie_Nr=A1.Kategorie_Nr
GROUP BY A1.KategorieName

SELECT  A1.KategorieName, sum(A0.Lagerbestand), sum(A0.Einzelpreis)
FROM cds A0
INNER JOIN Kategorien A1 ON A0.Kategorie_Nr=A1.Kategorie_Nr
GROUP BY A1.KategorieName

SELECT  A1.KategorieName, sum(A0.Lagerbestand), sum(A0.Einzelpreis)
FROM books A0
INNER JOIN Kategorien A1 ON A0.Kategorie_Nr=A1.Kategorie_Nr
GROUP BY A1.KategorieName
```

Of course there's also a drawback here: the same ProductGroup may be selected several times, so to get the correct sum, the results of the ProductGroup has to be added. In our sample the ProductGroup "Books" will be listed three times.

After listing so many drawbacks and problems, here's the SQL the produces the desired result. This is a manually created SQL, not generated by OJB. Unfortunately it's not fully supported by some DBMS because of "union" and sub-selects.

```
select KategorieName, sum(lagerbestand), sum(einzelpreis)
from
(
    SELECT  A1.KategorieName,A0.Lagerbestand,A0.Einzelpreis
    FROM artikel A0
    INNER JOIN Kategorien A1 ON A0.Kategorie_Nr=A1.Kategorie_Nr

    union

    SELECT  A1.KategorieName,A0.Lagerbestand,A0.Einzelpreis
    FROM books A0
    INNER JOIN Kategorien A1 ON A0.Kategorie_Nr=A1.Kategorie_Nr

    union

    SELECT  A1.KategorieName,A0.Lagerbestand,A0.Einzelpreis
    FROM cds A0
    INNER JOIN Kategorien A1 ON A0.Kategorie_Nr=A1.Kategorie_Nr
)
group by kategorieName
```

### 5.5.10.3. ODMG OQL

### 5.5.10.4. JDO queries

## 5.5.11. Metadata handling

### 5.5.11.1. Introduction

To make OJB proper work information about the used databases (more info see [connection handling](#)) and [sequence managers](#) is needed. Henceforth these metadata information is called **connection metadata**.

Further on OJB needs information about the persistent objects and object relations, henceforth this information is called **(persistent) object metadata**.

All metadata information need to be stored in the OJB [repository file](#).

The *connection metadata* are completely decoupled from the *persistent object metadata*. Thus it is possible to use the same *object metadata* on different databases.

But it is also possible to use different [object metadata profiles](#) .

In OJB there are several ways to make metadata information available:

- using xml configuration files parsed at start up by OJB
- set metadata instances at runtime by building metadata class instances at runtime
- parse additional xml configuration files (additional repository files) and merge at runtime

All classes used for managing metadata stuff can be find under `org.apache.ojb.broker.metadata.*`-package.

The main class for metadata handling and entry point for metadata manipulation at runtime is [org.apache.ojb.broker.metadata.MetadataManager](#) .

### 5.5.11.2. When does OJB read metadata

By default all metadata is read at startup of OJB, when the first call to `PersistenceBrokerFactory` (directly or by a top-level api) or `MetadataManager` class was done.

OJB expects a [repository file](#) at startup, but it is also possible to start OJB [without an repository file](#) or only [load connection metadata](#) and [object metadata at runtime](#) or what ever combination fit your requirements.

### 5.5.11.3. Connection metadata

The *connection metadata* encapsulate all information referring to used database and must be declared in [OJB repository file](#).

For each database a [jdbc-connection-descriptor](#) must be declared. This element encapsaltes the connection specific metadata information.

The `JdbcConnectionDescriptor` instances are managed by [org.apache.ojb.broker.metadata.ConnectionRepository](#)

### Load and merge connection metadata

It is possible to load additional connection metadata at runtime and merge it with the existing one. The used repository files have to be valid against the [repository.dtd](#):

```
<?xml version="1.0" encoding="UTF-8"?>
<!DOCTYPE descriptor-repository SYSTEM "repository.dtd">

<descriptor-repository version="1.0" isolation-level="read-uncommitted">
  <jdbc-connection-descriptor
    jcd-alias="runtime"
    platform="Hsqldb"
    jdbc-level="2.0"
    driver="org.hsqldb.jdbcDriver"
    protocol="jdbc">
```

```

        subprotocol="hsqldb"
        dbalias="../OJB_FarAway"
        username="sa"
        password=""
        batch-mode="false"
    >

    <object-cache
class="org.apache.ojb.broker.cache.ObjectCacheDefaultImpl">
        <attribute attribute-name="timeout" attribute-value="900"/>
        <attribute attribute-name="autoSync" attribute-value="true"/>
    </object-cache>

    <connection-pool
        maxActive="5"
        whenExhaustedAction="0"
        validationQuery="select count(*) from OJB_HL_SEQ"
    />

    <sequence-manager
class="org.apache.ojb.broker.util.sequence.SequenceManagerHighLowImpl">
        <attribute attribute-name="grabSize" attribute-value="5"/>
    </sequence-manager>
</jdbc-connection-descriptor>

<!-- user/passwd at runtime required -->
<jdbc-connection-descriptor
    jcd-alias="minimal"
    platform="Hsqldb"
    jdbc-level="2.0"
    driver="org.hsqldb.jdbcDriver"
    protocol="jdbc"
    subprotocol="hsqldb"
    dbalias="../OJB_FarAway"
    >
</jdbc-connection-descriptor>
</descriptor-repository>

```

In the above additional repository file two new *jdbc-connection-descriptor* (new databases) *runtime* and *minimal* are declared, to load and merge the additional connection metadata the *MetadataManager* was used:

```

// get MetadataManager instance
MetadataManager mm = MetadataManager.getInstance();

// read connection metadata from repository file
ConnectionRepository cr = mm.readConnectionRepository("valid path/url to
repository file");

// merge new connection metadata with existing one
mm.mergeConnectionRepository(cr);

```

After the merge the access to the new databases is ready for use.

#### 5.5.11.4. Persistent object metadata

The *object metadata* encapsulate all information referring to the persistent capable java objects and the associated tables in database. *Object metadata* must be declared in [OJB repository file](#). Each persistence capable java object must be declared in a corresponding [class-descriptor](#).

The *ClassDescriptor* instances are managed by [org.apache.ojb.broker.metadata.DescriptorRepository](#). Per default OJB use only **one global instance** of this class - it's the repository file read at startup of OJB. But it is possible to change the global use repository:

```

// get MetadataManager instance
MetadataManager mm = MetadataManager.getInstance();

```

```
mm.setDescriptor(myGlobalRepository, true);
```

### Load and merge object metadata

It is possible to load additional object metadata at runtime and merge it with the existing one. The used repository files have to be valid against the [repository.dtd](#):

#### Note:

When using the dynamic mapping technique described below, all objects in the structure must implement `java.io.Serializable` for OJB to be able to create cloned copies. OJB currently uses [SerializationUtils](#) from Commons Lang Core Language Utilities for all deep-cloning operations.

An additional repository file may look like:

```
<?xml version="1.0" encoding="UTF-8"?>
<!DOCTYPE descriptor-repository SYSTEM "repository.dtd">

<descriptor-repository version="1.0" isolation-level="read-uncommitted">

  <class-descriptor
    class="org.my.MyObject"
    table="MY_OBJ"
  >
    <field-descriptor
      name="id"
      column="OBJ_ID"
      jdbc-type="INTEGER"
      primaryKey="true"
      autoincrement="true"
    />

    <field-descriptor
      name="name"
      column="NAME"
      jdbc-type="VARCHAR"
    />
  </class-descriptor>
</descriptor-repository>
```

To load and merge the object metadata of the additional repository files first read the metadata using the [MetadataManager](#).

```
// get MetadataManager instance
MetadataManager mm = MetadataManager.getInstance();

// read the additional repository file
DescriptorRepository dr = mm.readDescriptorRepository("valid path/url to
repository file");

// merge the new class-descriptor with existing object metadata
mm.mergeDescriptorRepository(dr);
```

It is also possible to keep the different *object metadata* for the same classes parallel by using [metadata profiles](#).

### Global object metadata changes

The [MetadataManager](#) provide several methods to read/set and manipulate object metadata.

Per default OJB use a [global instance](#) of class [DescriptorRepository](#) to manage all *object metadata*. This means that all *PersistenceBroker* instances (kernel component used by all top-level api) use the same object metadata.

So changes of the object metadata (e.g. remove of a *CollectionDescriptor* instance from a

ClassDescriptor) will be seen immediately by all *PersistenceBroker* instances. This is in most cases not the favoured behaviour and OJB supports [per thread changes of object metadata](#).

### Per thread metadata changes

Per default the manager handle [one global \*DescriptorRepository\*](#) for all calling threads (keep in mind PB-api is not threadsafe, thus each thread use it's own *PersistenceBroker* instance), but it is ditto possible to use different *metadata profiles* in a per thread manner - profiles means different instances of *DescriptorRepository* objects. Each thread/*PersistenceBroker* instance can be associated with a specific *DescriptorRepository* instance. All made object metadata changes only will be seen by the *PersistenceBroker* instances using the same *DescriptorRepository* instance. In theory each *PersistenceBroker* instance could be associated with a separate instance of object metadata, but the recommended way is to use [metadata profiles](#).

To enable the use of different *DescriptorRepository* instances for each thread do:

```
MetadataManager mm = MetadataManager.getInstance();
// tell the manager to use per thread mode
mm.setEnablePerThreadChanges(true);
...
```

This can be done e.g. at start up or at runtime when it's needed. If method `setEnablePerThreadChanges` is set *false* only the [global \*DescriptorRepository\*](#) was used. Now it's possible to use dedicated *DescriptorRepository* instances per thread:

```
// e.g get a copy of the global repository
DescriptorRepository dr = mm.copyOfGlobalRepository();
// now we can manipulate the persistent object metadata of the copy
.....

// set the changed repository for current thread
mm.setDescriptor(dr);

// now let this thread lookup a PersistenceBroker instance
// with the modified metadata
// all other threads use still the global object metadata
PersistenceBroker broker =
PersistenceBrokerFactory.createPersistenceBroker(myKey)
```

#### Note:

Set object metadata (setting of the *DescriptorRepository*) before lookup the *PersistenceBroker* instance for current thread, because the metadata was bound to the *PersistenceBroker* instance at lookup.

### Object metadata profiles

*MetadataManager* was shipped with a simple mechanism to add, remove and load different *persistent objects metadata profiles* (different *DescriptorRepository* instances) in a per thread manner. Use method `addProfile` to add different persistent object metadata profiles, method `removeProfile` to remove profiles and `loadProfile` load a profile for the calling thread.

```
// get MetadataManager instance
MetadataManager mm = MetadataManager.getInstance();

// enable per thread mode if not done before
mm.setEnablePerThreadChanges(true);

// Load additional object metadata by parsing an repository file
DescriptorRepository dr_1 = mm.readDescriptorRepository("pathOrURLtoFile_1");
DescriptorRepository dr_2 = mm.readDescriptorRepository("pathOrURLtoFile_2");

// add profiles
```

```
mm.addProfile("global", mm.copyOfGlobalRepository());
mm.addProfile("guest", dr_1);
mm.addProfile("admin", dr_2);

// now load a specific profile
mm.loadProfile("admin");
broker = PersistenceBrokerFactory.defaultPersistenceBroker();
```

After the *loadProfile* call all PersistenceBroker instances will be associated with the *admin* profile.

#### Note:

Method *loadProfile* only proper work if the [per thread mode](#) is enabled.

### Reference runtime changes on per query basis

#### FIXME (arminw):

Changes of reference settings on a per query basis will be supported with next upcoming release 1.1

### Pitfalls

OJB's flexibility of *metadata handling* demanded specific attention on object caching. If a [global cache](#) (shared permanent cache) was used, be aware of side-effects caused by runtime metadata changes.

For example, using two metadata profiles *A* and *B*. In profile *A* all fields of a class are showed, in profile *B* only the 'name filed' is showed. Thread 1 use profile *A*, thread 2 use profile *B*. It is obvious that a global shared cache will cause trouble.

### 5.5.11.5. Questions

#### Start OJB without a repository file?

It is possible to start OJB without any repository file. In this case you have to declare the `jdbc-connection-descriptor` and `class-descriptor` at runtime. See [Connect to database at runtime?](#) and [Add new persistent objects \(class-descriptors\) at runtime?](#) for more information.

#### Connect to database at runtime?

There are two possibilities to connect your database at runtime:

- load connection metadata by parsing additional repository files
- create the *JdbcConnectionDescriptor* at runtime

The first one is described in section [load and merge connection metadata](#). For the second one a new instance of class [org.apache.ojb.broker.metadata.JdbcConnectionDescriptor](#) is needed. The prepared instance will be passed to class *ConnectionRepository*:

```
ConnectionRepository cr = MetadataManager.getInstance().connectionRepository();
JdbcConnectionDescriptor jcd = new JdbcConnectionDescriptor();
jcd.setJcdAlias("testConnection");
jcd.setUserName("sa");
jcd.setPassword("sa");
jcd.setDbAlias("aAlias");
jcd.setDbms("aDatabase");
// ... the other required setter
// add new descriptor
```

```
cr.addDescriptor(jcd);

// Now it's possible to obtain a PB-instance
PBKey key = new PBKey("testConnection", "sa", "sa");
PersistenceBroker broker = PersistenceBrokerFactory.
createPersistenceBroker(key);
```

Please [read this section from beginning](#) for further information.

### Add new persistent objects metadata ( class-descriptor) at runtime?

There are two possibilities to add new *object metadata* at runtime:

- load object metadata by parsing additional repository files
- create new metadata objects at runtime

The first one is described in section [load object metadata](#).

To create and add new metadata objects at runtime we create new [org.apache.ojb.broker.metadata.ClassDescriptor](#) instances at runtime and using the `MetadataManager` to add them to OJB:

```
DescriptorRepository dr = MetadataManager.getInstance().getRepository();

ClassDescriptor cld = new ClassDescriptor(dr);
cld.setClassOfObject(A.class);
//... other setter

// add the fields of the class
FieldDescriptor fd = new FieldDescriptor(cld, 1);
fd.setPersistentField(A.class, "someAField");
cld.addFieldDescriptor(fd);

// now we add the the class descriptor
dr.setClassDescriptor(cld);
```

Please [read this section from beginning](#) for further information.

## 5.5.12. Deployment

### 5.5.12.1. Introduction

This section enumerates all things needed to deploy OJB in standalone or servlet based applications and j2ee-container.

### 5.5.12.2. Things needed for deploying OJB

#### 1. The OJB binary jar archive

You need a `db-ojb-<version>.jar` file containing the compiled OJB library. This jar files contains all OJB code necessary in production level environments. It does not contain any test code. It also does not contain any configuration data. You'll find this file in the `lib` directory of the binary distribution. If you are working with the source distribution you can assemble the binary jar archive By calling

```
ant jar
```

This ant task generates the binary jar to the `dist` directory.

#### 2. Configuration data

OJB needs two kinds of configuration data:



1. Configuration of the OJB runtime environment. This data is stored in a file named [OJB.properties](#). [Learn more about this file here.](#)
2. Configuration of the MetaData layer. This data is stored in file named [repository.xml](#) (and several included files). [Learn more about this file here.](#)

**Note:**

These configuration files are read in through ClassLoader resource lookup and must therefore be placed on the classpath.

**3. External dependencies that do not come with OJB**

Some components of OJB depend on external libraries and components that cannot be shipped with OJB. You'll also need these if you want to compile OJB from source. Here is a list of these dependencies:

**j2ee.jar**

This is the main archive of the [J2EE SDK](#).

**jdo.jar, jdori\*.jar**

The [JDO Reference implementation](#) is required if you plan to use the JDO Api.

**4. Optional jar archives that come with OJB**

Some of jar files in the `lib` folder are only used during build-time or are only required by certain components of OJB, and so they might need not to be needed in runtime environments. Apart from wasting disk space they do no harm. If you don't care about disk space you just take all jars from the `lib` folder.

If you do care, here is the list of jars you might omit during runtime:

**ant-\*.jar**

These are the [Apache Ant 1.6](#) jars.

**antlr-[version].jar**

[ANTLR](#) is a parser generator which is used in the ODMG component of OJB. If you only use the PB Api, then you don't need this.

**junit.jar**

[JUnit](#) for running the unit tests. You'll need this only if you're also writing unit tests for you app.

**xerces.jar, xml-apis.jar**

The [Xerces](#) XML parser. Since most newer JDK's ship with an XML parser, it is likely that you do not need these files.

**xalan.jar**

[Xalan](#) is used to generate the unit test report, so you'll probably don't need this.

**jakarta-regexp-[version].jar**

The [Jakarta Regular Expression library](#) is only used when building OJB from source.

**torque-xxx.jar, velocity-xxx.jar**

[Torque](#) is used to generate concrete databases from database-independent schema files. OJB uses it internally to setup databases for the unit tests.

**xdoclet-[version].jar, xjavadoc-[version].jar,**

**xdoclet-obj-module-[version].jar,**

**commons-collections-[version].jar**

The [XDoclet OJB module](#) can be used to generate the repository metadata and

Torque schema files from Javadoc comments in the Java source files. It is however not required at runtime, so you can safely ignore these files then.

**5. Don't forget the JDBC driver**

The `repository.xml` defines JDBC Connections to your runtime databases. To use the declared

JDBC drivers the respective jar archives must also be present in the classpath. Refer to the documentation of your databases.

In the following sections I will describe how to deploy these items for specific runtime environments.

#### 5.5.12.3. Deployment in standalone applications

Deploying OJB for standalone applications is most simple. If you follow these four steps your application will be up in a few minutes.

1. Add `db-ojb-<version>.jar` to the classpath
2. place `OJB.properties` and `repository.xml` files on the classpath
3. Add the additional runtime jar archives to the classpath.
4. Add your JDBC drivers jar archive to the classpath.

#### 5.5.12.4. Deployment in servlet based applications

Generally speaking the four steps described in the previous section have to be followed also in Servlet / JSP based environments.

The exact details may differ for your specific Servlet container, but the general concepts should be quite similar.

1. Deploy `db-ojb-<version>.jar` with your servlet applications WAR file.  
The WAR format specifies that application specific jars are to be placed in a directory `WEB-INF/lib`. Place `db-ojb-<version>.jar` to this directory.
2. Deploy [OJB.properties](#) and [repository.xml](#) with your servlet applications WAR file.  
The WAR format specifies that Servlet classes are to be placed in a directory `WEB-INF/classes`. The OJB configuration files have to be in this directory.
3. Add the additional runtime jar archives to `WEB-INF/lib` too.
4. Add your JDBC drivers jar archive to `WEB-INF/lib`.

By executing `ant war` you can generate a sample servlet application assembled to a valid WAR file. The resulting `ojb-servlet.war` file is written to the `dist` directory. You can deploy this WAR file to your servlet engine or unzip it to have a look at its directory structure. you can also use the target `war` as a starting point for your own deployment scripts.

#### 5.5.12.5. Deployment in managed environment (e.g. EJB based)

The above mentioned guidelines concerning jar files and placing of the `OJB.properties` and the `repository.xml` are valid for managed/EJB environments as well.

But apart from these basic steps you'll have to perform some additional configurations to integrate OJB into a managed environment.

*Managed environment: Using of OJB in a managed environment means primarily the cooperation of OJB with the application server JTA service (via JCA or by using JTA classes).*

The instructions to make OJB running within your application server should be similar for all server. So the following instructions for JBoss should be useful for all user. E.g. most [OJB.properties](#) file settings are the same for all application server.

There are some topics you should examine very carefully:

- **Connection handling:** Lookup DataSource from your AppServer, only these connections can be enlisted in running transactions (JTA)
- **Caching:** Do you need caching? Do you need distributed caching?
- **Locking:** Do you need distributed locking (when using odmng-api in clustered environments)?

## Configure OJB for managed environments considering as JBoss example

The following steps describe how to configure OJB for managed environments and deploy on a ejb conform Application Server on the basis of the [shipped ejb-examples](#). In managed environments OJB needs some specific properties:

### 1. Adapt OJB.properties file

If the PB-api is the only persistence API being used (no ODMG nor JDO) and it is **only** being used in a managed environment, it is strongly recommended to use a special `PersistenceBrokerFactory` class, which enables `PersistenceBroker` instances to participate in the running JTA transaction - e.g. this makes `PBStateListener` proper work in managed environments and enables OJB to synchronize the persistent [caches](#) (e.g. the two-level cache):

```
PersistenceBrokerFactoryClass=org.apache.ojb.broker.core.PersistenceBrokerFactorySyncImpl
```

#### Note:

Don't use this setting in conjunction with any other top-level api (e.g. ODMG-api).  
If **no permanent caching** (only the "empty" cache implementation or the "per broker cache") is used and the `PBStateListener` is not used to detect tx demarcation, it's possible to use the default `PersistenceBrokerFactory` implementation, because OJB doesn't need to synchronize anything.

Your [OJB.properties](#) file need the following additional settings to work within managed environments (apply to **all** used api):

```
...
# only needed when using OJB 1.0.3 or earlier in managed environments. Since
version
# 1.0.4 OJB detects datasources from managed environments automatically.
ConnectionFactoryClass=
org.apache.ojb.broker.accesslayer.ConnectionFactoryManagedImpl
...
# set used application server TM access class
JTATransactionManagerClass=
org.apache.ojb.otm.transaction.factory.JBossTransactionManagerFactory
```

A specific [ConnectionFactory](#) implementation is used in version before 1.0.4 to by-pass all forbidden method calls in managed environments. Since OJB 1.0.4 datasources from managed environments are detected automatically by checking the *JTA-TxManager*.

The `JTATransactionManagerClass` property specify the used implementation class to lookup the *transaction manager* used by the application server. The `javax.transaction.TransactionManager` is needed to make it possible for OJB to participate in running *JTA transaction* via `javax.transaction.Synchronization` interface.

The ODMG-api needs some additional settings for use in managed environments (only needed when `odmg-api` was used):

```
...
# only needed for odmg-api
ImplementationClass=org.apache.ojb.odmg.ImplementationJTAImpl
...
# only needed for odmg-api
OJBTxManagerClass=org.apache.ojb.odmg.JTATxManager
```

The `ImplementationClass` specify the ODMG base class implementation. In managed environments a specific implementation is used, able to participate in *JTA transactions*.

The *OJBTxManagerClass* specify the used OJBTxManager implementation to manage the transaction synchronization in managed environments.

**Note:**

Currently OJB integrate in managed environments via `javax.transaction.Synchronization` interface. When the *JCA adapter* is finished (work in progress) integration will be more smooth.

**2. Declare datasource in the repository (repository\_database) file and do additional configuration**

Do only use `DataSource` from the application server to connect to your database (Local used connections can not participate in JTA transaction).

**Note:**

We strongly recommend to use JBoss 3.2.2 or higher of the 3.x series of JBoss. With earlier versions of 3.x we got Statement/Connection resource problems when running some ejb stress tests. As workaround we introduce a jboss specific attribute *eager-release* for version before 3.2.2, but it seems that this attribute can cause side-effects. Again, this problem seems to be fixed in 3.2.2.

Define OJB to use a `DataSource`:

```
<!-- Datasource example -->
<jdbc-connection-descriptor
  jcd-alias="default"
  default-connection="true"
  platform="Sapdb"
  jdbc-level="2.0"
  jndi-datasource-name="java:DefaultDS"
  username="sa"
  password=""
  eager-release="false"
  batch-mode="false"
  useAutoCommit="0"
  ignoreAutoCommitExceptions="false"
>
  <object-cache class="org.apache.ojb.broker.cache.ObjectCacheDefaultImpl">
    <attribute attribute-name="timeout" attribute-value="900"/>
    <attribute attribute-name="autoSync" attribute-value="true"/>
  </object-cache>

  <sequence-manager
className="org.apache.ojb.broker.util.sequence.SequenceManagerNextValImpl">
  </sequence-manager>
</jdbc-connection-descriptor>
```

In OJB versions before *1.0.4* the attribute `useAutoCommit="0"` is mandatory in managed environments, because it's in most cases not allowed to change the connection's *autoCommit* state.

**Note:**

In managed environments you can't use the *default* sequence manager implementation (`SequenceManagerHighLowImpl`) of OJB. For alternative sequence manager implementation [see here](#).

**[2b. How to deploy ojb test hsqldb database to jboss]**

If you use hsql database for testing you can easy setup the DB on jboss. After creating the database in OJB test directory with `ant prepare-testdb`, take the generated `.../target/test/OJB.script` file and rename it to `default.script`. Then replace the jboss `default.script` file in `.../jboss-3.x.y/server/default/db/hypersonic` with this file.

### 3. Include all OJB configuration files in classpath

Include the all needed OJB configuration files in your classpath:

- OJB.properties
- repository.dtd
- repository.xml
- repository\_internal.xml
- repository\_database.xml,
- repository\_ejb.xml (if you want to run the ejb examples)

To deploy the ejb-examples beans we include all configuration files in a ejb jar file - more info about this see [below](#).

The repository.xml for the [ejb-example beans](#) look like:

```
<?xml version="1.0" encoding="UTF-8"?>
<!-- This is a sample metadata repository for the ObjectBridge
System. Use this file as a template for building your own
mappings-->

<!-- defining entities for include-files -->
<!DOCTYPE descriptor-repository SYSTEM "repository.dtd" [
<!ENTITY database SYSTEM "repository_database.xml">
<!ENTITY internal SYSTEM "repository_internal.xml">
<!ENTITY ejb SYSTEM "repository_ejb.xml">
]>

<descriptor-repository version="1.0"
    isolation-level="read-uncommitted">

    <!-- include all used database connections -->
        &database;

    <!-- include ojb internal mappings here -->
        &internal;

    <!-- include mappings for the EJB-examples -->
        &ejb;

</descriptor-repository>
```

### 4. Enclose all libraries OJB depend on

In most cases it is recommended to include all libraries OJB depend on in the application .ear/.sar or ejb .jar file to make OJB run and (re-)deployable. Here are the libraries needed to make the ojb sample session beans run on JBoss:

- The jakarta commons libraries files (all commons-xxx.jar) from OJB /lib directory
- The antlr jar file (antlr-xxx.jar) from OJB /lib directory
- jakarta-regexp-xxx.jar from OJB /lib directory
- [jakarta turbine jcs.jar from OJB /lib directory, only if ObjectCacheJCSImpl was used]

(This was tested with jboss 3.2.2)

### 5. Take care of caching

Very important thing is cache synchronization with the database. When using the ODMG-api or PB-api (with [special PBF \(see 1.\)](#) setting) it's possible to use all [ObjectCache](#) implementations as long as OJB doesn't run in a clustered mode. When the ObjectCacheDefaultImpl cache implementation was used it's recommended to enable the *autoSync* mode.

In clustered environments (OJB run on different AppServer nodes) you need a [distributed ObjectCache](#) or you should use a local/empty cache like

```
ObjectCacheClass=org.apache.ojb.broker.cache.ObjectCachePerBrokerImpl
```

or

```
ObjectCacheClass=org.apache.ojb.broker.cache.ObjectCacheEmptyImpl
```

The cache is pluggable, so you can write your own ObjectCache implementation to accomplish your expectations.

More info you can find in [clustering](#) and [ObjectCache](#) topic.

## 6. Take care of locking

If the used api supports *Object Locking* (e.g. ODMG-api, PB-api does not), in clustered environments (OJB run on different AppServer nodes) a distributed [lock management](#) is mandatory.

## 7. Put all together

Now put all files together. We keep the examples as simple as possible, thus we deploy only a `ejb.jar` file. [Below](#) you can find a short instruction how to pack an `ejb` application `.ear` file including OJB.

Generate the `ejb`-examples [described below](#) or build your own `ejb.jar` file including all beans, `ejb-jar.xml` and appServer dependend files. Then add all OJB configuration files, the `db-ojb.jar` file and all libraries OJB depends on into this `ejb.jar` file.

The structure of the `ejb.jar` file should now look like this:

```
/OJB.properties
/repository.dtd
/repository.xml
/all used repository-XYZ.xml
/META-INF
../Manifest.mf
../ejb-jar.xml
../jboss.xml

/all ejb classes

/db-ojb-1.X.jar
/all used libraries
```

### 7b. Example: Deployable jar

For example the `jar`-file used to test the [ejb-examples](#) shipped with OJB, base on the `db-ojb-XY-beans.jar` file. This `jar` was created when the [ejb-examples](#) target was called.

The generated `jar` contains only the `ejb`-classes and the deployment-descriptor. We have to add additional `jars` (all libraries used by OJB) and files (all configuration files) to make it deployable. The deployable `db-ojb-XY-beans.jar` should look like this:

```
/OJB.properties
/repository.dtd
/repository.xml
/repository_database.xml
/repository_ejb.xml
/repository_internal.xml
/META-INF
../Manifest.mf
../ejb-jar.xml
```

```

../jboss.xml

/org
../apache (all ejb classes)

/db-ojb-1.X.jar

/antlr-XXX.jar
/commons-beanutils-XXX.jar
/commons-collections-XXX.jar
/commons-dbcp-XXX.jar
/commons-lanf-XXX.jar
/commons-logging-XXX.jar
/commons-pool-XXX.jar
/jakarta-regexp-XXX.jar

```

Please pay attention on the [configuration settings](#) to make OJB work in managed environments (especially the OJB.properties settings).

**Note:**

This example isn't a real world production example. Normally you will setup one or more enterprise archive files (.ear files) to bundle one or more complete J2EE (web) applications. More about how to build an *J2EE application* using OJB [see here](#).

The described example should be re-deployable/hot-deployable in JBoss.

**If you will get any problems, please let me know. All suggestions are welcome!**

#### 8. How to access OJB API?

In managed environments it is possible to access OJB in same way used in non-managed environments:

```

// PB-api
PersistenceBroker broker = PersistenceBrokerFactory.create...;

//ODMG-api
Implementation odmg = OJB.getInstance();

```

But it is recommended to bind OJB api access classes to JNDI and lookup the the api entry [classes via JNDI](#).

#### 9. OJB logging within JBoss

Jboss use [log4j](#) as standard logging api.

In summary, to use log4j logging with OJB within jBoss:

1) in OJB.properties set

```
LoggerClass=org.apache.ojb.broker.util.logging.Log4jLoggerImpl
```

There is no need for a separate log4j.properties file of OJB-specific log4j settings (in fact the OJB.properties setting LoggerConfigFile is ignored). Instead, the jBoss log4j configuration file must be used:

2) in JBOSS\_HOME/server/default/conf/log4j.xml,

define appenders and add categories to add or filter logging of desired OJB packages, following the numerous examples in that file. For example,

```

<category name="org.apache.ojb">
  <priority value="DEBUG" />
  <appender-ref ref="CONSOLE" />
  <appender-ref ref="FILE" />
</category>

<category name="org.apache.ojb.broker.metadata.RepositoryXmlHandler">

```

```
<priority value="ERROR" />
<appender-ref ref="CONSOLE" />
<appender-ref ref="FILE" />
</category>
```

## Example Session Beans

### Introduction

The OJB source distribution was shipped with a bunch of sample session beans and client classes for testing. Please recognize that we don't say that these examples show "best practices" of using OJB within enterprise java beans - it's only one way to make it work.

To keep the examples as simple as possible we directly use the OJB main classes via static lookup or helper classes on each *ejbCreate()* call. But we recommend to [bind the OJB main classes](#) in JNDI instead of direct use in the session beans.

### Generate the sample session beans

The source code of the sample beans is stored in directory  
`[db-ojb]/src/ejb/org/apache/ojb/ejb`  
 To generate the sample beans call

```
ant ejb-examples
```

This ant target copies the bean sources to `[db-ojb]/target/src/ejb` generates all needed bean classes and deployment descriptor ( [by using xdoclet](#)) to the same directory, compiles the sources and build an `ejb.jar` file called `[db-ojb]/dist/db-ojb-XXX-beans.jar`. Test clients for the generated beans included in the `[db-ojb]/dist/db-ojb-XXX-client.jar`.

To run xdoclet properly the following xdoclet jar files needed in `[db-ojb]/lib` directory (xdoclet version 1.2xx or higher):

```
xdoclet-xxx.jar
xdoclet-ejb-module-xxx.jar
xdoclet-jboss-module-xxx.jar
xdoclet-jmx-module-xxx.jar
xdoclet-web-module-xxx.jar
xdoclet-xjavadoc-module-xxx.jar
```

If you using a different application server than JBoss, you have to modify the *xdoclet* ant target in `[db-ojb]/build-ejb-examples.xml` to force xdoclet to generate the appServer specific files. See xdoclet documentation for further information.

### How to run test clients for PB / ODMG - api

If the "extended `ejb.jar`" file was successfully deployed we need a test client to invoke the `ejb-examples`. As said above, the *ejb-examples* target generates a test client jar too. It's called `[db-ojb]/dist/db-ojb-XXX-client.jar` and contains junit based test clients for the PB-/ODMG-api.

The main test classes are:

- `org.apache.ojb.ejb.AllODMGTests`
- `org.apache.ojb.ejb.AllPBTests`

OJB provide an ant target to run the client side bean tests. Include all needed appServer libraries in `[db-ojb]/lib` (e.g. for JBoss `jbossall-client.jar` do the job, beside the "j2ee jars"). To run the PB-api test clients (access running JBoss server with default settings) call

```
ant ejb-examples-run -Dclient.class=org.apache.ojb.ejb.AllPBTests
```



To run the test clients on an arbitrary appServer pass the JNDI properties for naming context initialisation too, e.g.

- `-Djava.naming.factory.initial="org.jnp.interfaces.NamingContextFactory"`
- `-Djava.naming.provider.url="jnp://localhost:1099"`
- `-Djava.naming.factory.url.pkgs="org.jboss.naming:org.jnp.interfaces"`

Then the target call may look like

```
ant ejb-examples-run -Dclient.class=org.apache.ojb.ejb.AllPBTests
-Djava.naming.factory.initial="org.jnp.interfaces.NamingContextFactory"
-Djava.naming.provider.url="jnp://localhost:1099"
-Djava.naming.factory.url.pkgs="org.jboss.naming:org.jnp.interfaces"
```

### Packing an .ear file

Here is an example of the .ear package structure. It is redeployable without having to restart JBoss.

#### The Package Structure

The package structure of the .ear file should look like:

```
/ejb.jar/
...EJBs
...META-INF/
.....ejb-jar.xml
.....jboss.xml
.....MANIFEST.MF

/web-app.war/
...JSP
...WEB-INF/
.....web.xml

/META-INF/
...application.xml
/ojb.jar
/[ojb required runtime jar]

/OJB.properties
/OJB-logging.properties
/repository.dtd
/repository_internal.xml
/repository.xml
/repository_database1.xml
/repository_app1.xml
/repository_database2.xml
/repository_app2.xml
```

#### Make OJB API Resources available

There are two approaches to use OJB api in the ejb.jar file:

**1.** To create a Manifest.mf file with classpath attribute that include all the runtime jar required by OJB (Very important to include all required jar). The sample below works fine (replace [version] with distributed JAR versions):

```
Class-Path: db-ojb-[version].jar antlr-[version].jar
commons-beanutils-[version].jar
commons-collections-[version].jar commons-dbc-[version].jar
commons-lang-[version].jar
commons-logging-[version].jar commons-pool-[version].jar
jakarta-regexp-[version].jar
```

#### Note:

If you to include the jar file under a directory of the ear file, says `/lib/db-objb-[version].jar` and etc. At the classpath attribute it will be something like: `Class-Path: ./lib/db-objb-[version].jar` and etc (The "." in front is important)

## 2. To add the required jar file as a "java" element in the application.xml file:

```
<module>
  <java>antlr-[version].jar</java>
</module>
<module>
  <java>commons-beanutils-[version].jar</java>
</module>
<module>
  <java>commons-collections-[version].jar</java>
</module>
<module>
  <java>commons-dbcpr-[version].jar</java>
</module>
<module>
  <java>commons-lang-[version].jar</java>
</module>
<module>
  <java>commons-logging-[version].jar</java>
</module>
<module>
  <java>commons-pool-[version].jar</java>
</module>
<module>
  <java>db-objb-[version].jar</java>
</module>
```

### Note:

To use this approach, all the library had to be in the root of the ear.

(This was tested on Jboss 3.2.3)

### Make OJB accessible via JNDI

Current bean examples do directly use OJB main classes, but it's also possible to make OJB accessible via JNDI and use a JNDI-lookup to access OJB api's in your beans. To make the OJB api's accessible via JNDI, bind main/access classes to JNDI. How to do this depends on the used environment. The main classes/methods to bind are:

- PB-api:  
Method  
`org.apache.objb.broker.core.PersistenceBrokerFactoryFactory#instance()`  
returns the used  
`org.apache.objb.broker.core.PersistenceBrokerFactoryIF`. Make this instance accessible via JNDI.
- ODMG-api:  
Method `org.apache.objb.odmg.OJB#getInstance()` returns a new instance of the `org.odmg.Implementation` instance. Open a new Database and make this instance and the Database instance accessible via JNDI.

### JBoss

In JBoss you can write *mbean* classes to bind OJB main/access classes to JNDI, similar to the [Weblogic example](#) below.

Let JBoss know about the new mbeans, so declare them in a `jboss-service.xml` file. Please see JBoss documentation how to write mbeans and bind objects to JNDI.

## Other Application Server

In other application server you can do similar steps to bind OJB main api classes to JNDI. For example in Weblogic you can use *startup class* implementation to bind OJB main/access classes to JNDI (see [below](#)).

## Instructions for Weblogic

1. Add the OJB jar files and dependencies into the Weblogic classpath
2. As usual create the connection pool and the datasource.
3. Prepare the OJB.properties file. Should be similar to [jboss](#). Expect the following entry:

```
...
# Weblogic Transaction Manager Factory
JTATransactionManagerClass=
org.apache.ojb.broker.transaction.tm.WeblogicTransactionManagerFactory
```

4. Modify the connection information in the repository.xml (specify the datasource name). SequenceManager implementation depends on the used DB, more info [see here](#):

```
<jdbc-connection-descriptor
jcd-alias="default"
default-connection="true"
platform="Sapdb"
jdbc-level="2.0"
jndi-datasource-name="datasource_demodb"
eager-release="false"
batch-mode="false"
useAutoCommit="0"
ignoreAutoCommitExceptions="false"
>

<sequence-manager
className="org.apache.ojb.broker.util.sequence.SequenceManagerNextValImpl">
<attribute attribute-name="grabSize" attribute-value="20"/>
</sequence-manager>
</jdbc-connection-descriptor>
```

### Note:

The following step is only necessary if you want to bind OJB main api classes to JNDI.

[5.] Compile the following classes (see at the [end of this section](#)) and add them to the weblogic classpath. This allows to access the PB-api via JNDI lookup. Register via the weblogic console the startup class (see `OjbPbStartup` class below). The JNDI name and the OJB.properties file path can be specified as parameters in this startup class.

To use the ODMG-api you have to write a similar startup class. This shouldn't be too complicated. Take a look in `org.apache.ojb.jboss` package (dir `src/connector/main`). Here you could find the `jboss mbeans`. All you have to do is bound a similar class to JNDI in weblogic. Implement `ODMGJ2EEFactory` Interface in your class bound this class to JNDI (in the `ejb-examples` the beans try to lookup the `Implementation` instance via `"java:/obj/defaultODMG"`). Your `ODMGFactory` class should implement this method

```
public Implementation getInstance()
{
    return OJB2EE_2.getInstance();
}
```

Write a session bean similar to those provided for the JBOSS samples. It is also possible to use the `ejb-example` beans (doing minor modifications when the JNDI lookup should be used).

*Weblogic startup class*

Write an OJB startup class to make OJB accessible via JNDI can look like (I couldn't test this sample class, so don't know if it will work ;-)):

```

package org.apache.ojb.weblogic;

import javax.naming.*;

import org.apache.ojb.broker.core.PersistenceBrokerFactoryFactory;
import org.apache.ojb.broker.core.PersistenceBrokerFactoryIF;

import weblogic.common.T3ServicesDef;
import weblogic.common.T3StartupDef;
import java.util.Hashtable;

/**
 * This startup class created and binds an instance of a
 * PersistenceBrokerFactoryIF into JNDI.
 */
public class OjbPbStartup
    implements T3StartupDef, OjbPbFactory, Serializable
{
    private String defaultPropsFile = "org/apache/ojb/weblogic/OJB.properties";

    public void setServices(T3ServicesDef services)
    {
    }

    public PersistenceBrokerFactoryIF getInstance()
    {
        return PersistenceBrokerFactoryFactory.instance();
    }

    public String startup(String name, Hashtable args)
        throws Exception
    {
        try
        {
            String jndiName = (String) args.get("jndiname");
            if(jndiName == null || jndiName.length() == 0)
                jndiName = OjbPbFactory.DEFAULT_JNDI_NAME;

            String propsFile = (String) args.get("propsfile");
            if(propsFile == null || propsFile.length() == 0)
            {
                System.setProperty("OJB.properties", defaultPropsFile);
            }
            else
            {
                System.setProperty("OJB.properties", propsFile);
            }

            InitialContext ctx = new InitialContext();
            bind(ctx, jndiName, this);

            // return a message for logging
            return "Bound OJB PersistenceBrokerFactoryIF to " + jndiName;
        }
        catch(Exception e)
        {
            e.printStackTrace();
            // return a message for logging
            return "Startup Class error: impossible to bind OJB PB factory";
        }
    }

    private void bind(Context ctx, String name, Object val)
        throws NamingException

```

```

    {
        Name n;
        for(n = ctx.getNameParser("").parse(name); n.size() > 1; n =
n.getSuffix(1))
        {
            String ctxName = n.get(0);
            try
            {
                ctx = (Context) ctx.lookup(ctxName);
            }
            catch(NameNotFoundException namenotfoundexception)
            {
                ctx = ctx.createSubcontext(ctxName);
            }
        }
        ctx.bind(n.get(0), val);
    }
}

```

The used OjbPbFactory interface:

```

package org.apache.ojb.weblogic;

import org.apache.ojb.broker.core.PersistenceBrokerFactoryIF;

public interface OjbPbFactory
{
    public static String DEFAULT_JNDI_NAME = "PBFactory";
    public PersistenceBrokerFactoryIF getInstance();
}

```

### 5.5.13. Connection Handling

#### 5.5.13.1. Introduction

In this section the connection handling within OJB is described. The connection management is implemented through two OJB interfaces:

- org.apache.ojb.broker.accesslayer.ConnectionFactory
- org.apache.ojb.broker.accesslayer.ConnectionManagerIF

#### 5.5.13.2. ConnectionFactory

The org.apache.ojb.broker.accesslayer.ConnectionFactory interface implementation is a pluggable component (via the [OJB.properties](#) file - more about [the OJB.properties file here](#)) responsible for creation/lookup and release of connections.

```

public interface ConnectionFactory
{
    Connection lookupConnection(JdbcConnectionDescriptor jcd) throws
LookupException;

    void releaseConnection(JdbcConnectionDescriptor jcd, Connection con);

    void releaseAllResources();
}

```

To enable a specific *ConnectionFactory* implementation class in the *OJB.properties* file, set property *ConnectionFactoryClass*. Default:

```
ConnectionFactoryClass=org.apache.ojb.broker.accesslayer.ConnectionFactoryPooledImpl
```

OJB is shipped with several different implementation classes for use in different situations. The default implementation for example, will pool created Connection instances for increased performance (since instance creation normally makes a database server roundtrip and thus is

costly).

To make it more easier to implement your own *ConnectionFactory* class, an abstract base class called `org.apache.obj.broker.accesslayer.ConnectionFactoryAbstractImpl` exists, most shipped implementation classes inherit from this class.

**Note:**

All shipped implementations of *ConnectionFactory* with support for connection pooling will only use object pools for connections obtained directly from the JDBC DriverManager. If you are using a *DataSource* configuration, the JNDI DataSource is responsible for pooling.

### ConnectionFactoryPooledImpl

A *ConnectionFactory* implementation using [commons-pool](#) to pool the Connection instances. On `lookupConnection` a Connection instance is borrowed from the object pool, and returned on the `releaseConnection` call. This implementation is used as *default* setting in the [OJB.properties](#) file.

This implementation allows a wide range of different settings, more info about the configuration properties can be found in the metadata repository [connection-pool](#) section.

### ConnectionFactoryNotPooledImpl

Implementation that creates a new Connection instance on each `lookupConnection` call and closes (destroys) it on `releaseConnection`. All [connection-pool](#) settings are ignored by this implementation.

### ConnectionFactoryManagedImpl

[@deprecated since OJB 1.0.4, now OJB automatic detect the running JTA-transaction and suppress critical method calls on the used connection]

Implementation specifically for use in *managed environments* like J2EE conformant application servers. In managed environments it is **mandatory** to use *DataSource* configuration, with Connection objects provided by the application server. OJB will not control Connection properties or transaction handling when using this implementation.

All [connection-pool](#) settings are ignored by this implementation.

### ConnectionFactoryDBCPImpl

Implementation using [commons-dbc](#) to pool the Connection instances. Since DBCP is using [commons-pool](#) internally, this implementation is very similar to *ConnectionFactoryPooledImpl*, but permits additional configuration for logging abandoned Connection instances (usable under development for detecting bad programming patterns).

This implementation allows a wide range of different settings, more info about the configuration properties can be found in the metadata repository [connection-pool](#) section.

### 5.5.13.3. ConnectionManager

The `org.apache.obj.broker.accesslayer.ConnectionManagerInterface` interface implementation is a pluggable component (via the [OJB.properties](#) file - more about [the OJB.properties file here](#)) responsible for managing the connection usage lifecycle and connection status (commit/rollback of connections).

```
public interface ConnectionManagerInterface
```

```

{
    JdbcConnectionDescriptor getConnectionDescriptor();
    Platform getSupportedPlatform();
    boolean isAlive(Connection conn);
    Connection getConnection() throws LookupException;
    boolean isInLocalTransaction();
    void localBegin();
    void localCommit();
    void localRollback();
    void releaseConnection();
    void setBatchMode(boolean mode);
    boolean isBatchMode();
    void executeBatch();
    void executeBatchIfNecessary();
    void clearBatch();
}

```

The *ConnectionManager* is used by the *PersistenceBroker* to handle connection usage lifecycle.

#### 5.5.13.4. Questions and Answers

##### How does OJB handle connection pooling?

OJB does connection pooling per default, except for datasources that are never pooled internally by OJB. Pooling of Connection instances when configuring OJB with DataSource lookup must be configured and performed in the DataSource provider.

The implementations of the

`org.apache.ojb.broker.accesslayer.ConnectionFactory.java` interface are responsible for managing the connections in OJB. There are several implementations shipped with OJB called

`org.apache.ojb.broker.accesslayer.ConnectionFactoryImpl.java`. There is, among others, a non-pooling implementation and an implementation using Commons DBCP API.

Configuration of the connection pooling is specified using the *connection-pool* element for each *jdbc-connection-descriptor*. The *connection-pool* element can be configured with properties for the specific *ConnectionFactory* implementation or JDBC driver used. For general information about the configuration, see the [repository section](#) or read the comments in [repository.dtd](#).

##### Can I directly obtain a java.sql.Connection within OJB?

It is possible to obtain a Connection using the PB API and a PersistenceBroker instance. Example:

```

PersistenceBroker broker =
PersistenceBrokerFactory.createPersistenceBroker(myKey);
broker.beginTransaction();
// do something

Connection con = broker.serviceConnectionManager().getConnection();
// perform your connection action and do more

```

```
// close the created statement and result set
broker.commitTransaction();
broker.close();
```

After obtaining a Connection with `broker.getServiceConnectionManager().getConnection()`, the connection can be used for any JDBC operations (except for transaction handling, more on this below). The user is responsible for cleanup of created Statement and ResultSet instances, so be sure to guard your call with a finally clause and close resources after use.

For read-only operations there is no need to start a PB transaction as in the example.

**Note:**

Do not commit any transactions on the Connection level, this should be left to OJB's PB API and will be performed automatically by calling PersistenceBroker `commit-/abortTransaction` methods.

**Note:**

Do not call `Connection.close()` on the obtained Connection, this should be left to OJB's [ConnectionFactory](#) and will be performed automatically when calling `broker.close()`.

If **no** transaction is running, it is possible to release a connection "by hand" after use by calling:

```
broker.getServiceConnectionManager().releaseConnection();
```

This call performs cleanup operations on the used connection and pass the instance to the release method of [ConnectionFactory](#) (this will e.g. return the connection to pool or close it).

If you do not do any connection cleanup, the connection will at the latest be released when calling `broker.close()`.

Users who are interested in this section might also be interested in ['Is it possible to perform my own sql-queries in OJB?'](#).

#### When does OJB open/close a connection

This is dependent on the used OJB api. Generally OJB try to obtain a connection as late as possible and close (if a connection pool was used, OJB return the connection to the pool) the connection as soon as possible.

Using the [PB-api](#) the connection is obtained when

[PersistenceBroker.beginTransaction\(\)](#) was called or a query is executed.

On `PersistenceBroker.commitTransaction()` or

`PersistenceBroker.abortTransaction()` call the connection was released. If no PB-tx is running, the connection will be released on `PersistenceBroker.close()` call.

Using the [ODMG-api](#) the connection is obtained when a query is executed or when the transaction commit. On leaving the commit method, the connection will be released.

All other top-level API should behave similar.

## 5.5.14. The Object Cache

### 5.5.14.1. Introduction

OJB supports several caching strategies and allow to [pluggin own](#) caching solutions by implementing the [ObjectCache](#) interface. All implementations shipped with OJB can be found in package `org.apache.ojb.broker.cache`. The naming convention of the implementation classes is `ObjectCacheXXXImpl`.



To classify the different implementations we differ *local/session cache* and *shared/global/application cache* implementations (we use the different terms synonymous). The [ObjectCacheTwoLevelImpl](#) use both characteristics.

- Local cache implementation mean that each instance use its own map to manage cached objects.
- Shared/global cache implementations share one (in most cases static) map to manage cached objects.

A [distributed object cache](#) implementation supports caching of objects across different JVM.

#### 5.5.14.2. Why a cache and how it works?

OJB provides a pluggable object cache provided by the [ObjectCache](#) interface:

```
public interface ObjectCache
{
    /**
     * Write to cache.
     */
    public void cache(Identity oid, Object obj);

    /**
     * Lookup object from cache.
     */
    public Object lookup(Identity oid);

    /**
     * Removes an Object from the cache.
     */
    public void remove(Identity oid);

    /**
     * Clear the ObjectCache.
     */
    public void clear();
}
```

Each [PersistenceBroker](#) instance (PersistenceBroker is a standalone api and the basic layer for all top-level api's like ODMG) use it's own ObjectCache instance. The ObjectCache instances are created by the ObjectCacheFactory class on PersistenceBroker instantiation.

Each cache implementation holds objects previously loaded or stored by the PersistenceBroker - dependend on the implementation.

Using a Cache has several advantages:

- It increases performance as it reduces database lookups or/and object materialization. If an object is looked up by Identity the associated PersistenceBroker instance does not perform a SELECT against the database immediately but first looks up the cache if the requested object is already loaded. If the object is cached it is returned as the lookup result. If it is not cached a SELECT is performed. Other queries were performed against the database, but before an object from the ResultSet was materialized the object identity was looked up in cache. If not found the whole object was materialized.
- It allows to perform circular lookups (as by crossreferenced objects) that would result in non-terminating loops without such a cache (Note: Since OJB 1.0.2 this is handled internally by OJB and does not depend on the used cache implementation).

#### 5.5.14.3. How to declare and change the used ObjectCache implementation

The object-cache element can be used to specify the ObjectCache implementation used by

OJB. If no *object-cache* is declared in configuration files (see below), OJB use by default a [noop-implementation](#) of the `ObjectCache` interface.

There are two levels of declaration:

- jdbc-connection-descriptor level
- class-descriptor level

and the possibility to exclude all persistent objects of [specified package names](#).

Use a **jdbc-connection-descriptor level** declaration to declare `ObjectCache` implementation on a per connection/user level. Additional configuration properties can be passed by using [custom attributes](#) entries:

```
<jdbc-connection-descriptor ...>
...
<object-cache class="org.apache.ojb.broker.cache.ObjectCacheDefaultImpl">
  <attribute attribute-name="timeout" attribute-value="900"/>
  <attribute attribute-name="useAutoSync" attribute-value="true"/>
</object-cache>
...
</jdbc-connection-descriptor>
```

Set an *object-cache* tag on [class-descriptor level](#) , to declare `ObjectCache` implementation on a per class level:

```
<class-descriptor
  class="org.apache.ojb.broker.util.sequence.HighLowSequence"
  table="OJB_HL_SEQ"
>
  <object-cache class="org.apache.ojb.broker.cache.ObjectCacheEmptyImpl">
  </object-cache>
...
</class-descriptor>
```

Additional configuration properties can be passed by using [custom attributes](#) entries.

#### Note:

If [polymorphism](#) was used it's only possible to declare the *object-cache* element in the [class-descriptor](#) of the top-level class/interface (root class), all *object-cache* declarations in the sub-classes will be ignored by OJB.

### Priority of Cache Level

Since it is possible to mix the different levels of *object-cache* element declaration a ordering of priority is needed:

#### Note:

The order of priority of declared *object-cache* elements in metadata are:  
per class > [excluded packages](#) > per jdbc-connection-descriptor

E.g. if you declare `ObjectCache` 'OC1' on connection level and set `ObjectCache` 'OC2' in class-descriptor of class A. Then OJB use 'OC2' as `ObjectCache` for class A instances and 'OC1' for all other classes.

### Exclude classes from being cached

If it's undesirable to cache an persistent object (e.g. persistent objects with BLOB fields or large binary fields) declare an `object-cache` descriptor with the *noop-cache* implementation called [ObjectCacheEmptyImpl](#).

```
<class-descriptor
```

```

class="org.apache.ojb.broker.util.sequence.HighLowSequence"
table="OJB_HL_SEQ"
>
  <object-cache class="org.apache.ojb.broker.cache.ObjectCacheEmptyImpl">
  </object-cache>
</class-descriptor>

```

**Note:**

If [polymorphism](#) was used and the class to exclude is part of an inheritance hierarchy **and** it's declared in in OJB metadata, it's not possible to exclude it. Only for the top-level class/interface (root class) it's allowed to specify the *object-cache* element in metadata. So it's only possible to exclude all sub-classes of the top-level class/interface (root class). More info [see here](#).

**Exclude packages from being cached**

To exclude all persistent objects of a whole package from being cached use the [custom attribute](#) *cacheExcludes* on connection level within the object-cache declaration. To declare several packages use a comma separated list.

```

<jdbc-connection-descriptor
  jcd-alias="myDefault"
  ...>
  <object-cache class="org.apache.ojb.broker.cache.ObjectCacheTwoLevelImpl">
    <attribute attribute-name="cacheExcludes"
      attribute-value="my.core, my.persistent.local"/>
    ... more attributes
  </object-cache>
</jdbc-connection-descriptor>

```

To include a persistent class of a excluded package, simply declare an object-cache descriptor on class-descriptor level of the class to include, object cache declarations on class-descriptor level have a higher priority as the excluded packages - [see more](#).

**Turn off caching**

If you don't declare a *object-cache* element in configuration files (see [here](#)), OJB doesn't cache persistent objects by default.

To explicitly turn off caching declare a *no-op* implementation of the [ObjectCache](#) interface as caching implementation. OJB was shipped with such a class called [ObjectCacheEmptyImpl](#). To explicitly turn off caching for a used database look like this:

```

<jdbc-connection-descriptor ...>
  ...
  <object-cache class="org.apache.ojb.broker.cache.ObjectCacheEmptyImpl">
  </object-cache>
  ...
</jdbc-connection-descriptor>

```

To get more detailed info about the different level of cache declaration, please see [here](#).

**5.5.14.4. Shipped cache implementations:****ObjectCacheDefaultImpl**

Per default OJB use a shared reference based [ObjectCache](#) implementation - [ObjectCacheDefaultImpl](#). It's a really fast cache but there are a few drawbacks:

- There is no transaction isolation, when thread one modify an object, thread two will see the modification when lookup the same object or use a reference of the same object, so "dirty-reads" can happen.
- If you rollback/abort a transaction the modified/corrupted objects will **not** be removed from the

cache by default(when using PB-api, top-level api may support automatic cache synchronization). You have to do this by your own using a service method to remove cached objects or enable the [autoSync](#) property.

```
broker.removeFromCache(obj);
// or (using Identity object)
ObjectCache cache = broker.serviceObjectCache();
cache.remove(oid);
```

- This implementation cache full object graphs (the object with all referenced objects) and does **not** synchronize the references. So if cached object *ProductGroup* has a 1:n reference to *Article*, e.g. article1, article2, article3 and another thread delete article2, the *ProductGroup* still has a reference to article2. To avoid such a behavior you can use the [collection-descriptor 'refresh' attribute](#) to force OJB to query the referenced objects when the main object is loaded from cache or use another *ObjectCache* implementation supporting synchronization of references (e.g. [ObjectCacheTwoLevelImpl](#)).

This implementation use by default *SoftReference* to wrap all cached objects. If the cached object was not longer referenced by your application but only by the cache, it can be reclaimed by the garbage collector.

As we don't know when the garbage collector reclaims the freed objects, it is possible to set a timeout property. So an cached object was only returned from cache if it was not garbage collected and was not timed out.

To enable this *ObjectCache* implementation declare

```
<object-cache class="org.apache.ojb.broker.cache.ObjectCacheDefaultImpl">
  <attribute attribute-name="cacheExcludes" attribute-value=""/>
  <attribute attribute-name="timeout" attribute-value="900"/>
  <attribute attribute-name="autoSync" attribute-value="true"/>
  <attribute attribute-name="cachingKeyType" attribute-value="0"/>
  <attribute attribute-name="useSoftReferences" attribute-value="true"/>
</object-cache>
```

Implementation configuration properties:

Property Key	Property Values
timeout	Lifetime of the cached objects in seconds. If expired, the cached object was discarded - default was 900 sec. When set to -1 the lifetime of the cached object never expire.
autoSync	<p>If set <i>true</i> all cached/looked up objects within a PB-transaction are traced. If the the PB-transaction was aborted all traced objects will be removed from cache. Default is <i>false</i>.</p> <p><b>NOTE:</b> This does not prevent "dirty-reads" by concurrent threads (more info see above).</p> <p>It's not a smart solution for keeping cache in sync with DB but should do the job in most cases. E.g. if OJB read 1000 objects from the database within a transaction, one object was modified and the transaction will be aborted, then 1000 objects will be passed to the cache on lookup, 1000 objects will be traced and all 1000 objects will be removed from cache on abort.</p> <p>Read these objects without running tx or in a former tx and then modify one object in a tx and abort the tx, only one object was traced/removed. Keep in mind</p>

	that this property counteract the <i>useSoftReferences</i> property as long as the PB-transaction is running, because all traced objects will have strong references.
cachedKeyType	Determines how the key was build for the cached objects: 0 - Identity object was used as key, this was the <i>default</i> setting. 1 - Identity + jcdAlias name was used as key. Useful when the same object metadata model (DescriptorRepository instance) are used for different databases (JdbcConnectionDescriptor), because different databases should use separated caches (persistent object instances). 2 - Identity + model (DescriptorRepository) was used as key. Useful when different metadata model (DescriptorRepository instance) are used for the same database. Keep in mind that there was no synchronization between cached objects with same Identity but different metadata model. E.g. when the same database use different metadata versions of the same persistent object class. 3 - all together (Identity + jcdAlias + model) If possible '0' is recommended, because it will be the best performing setting.
useSoftReferences	If set <i>true</i> this class use {@link java.lang.ref.SoftReference} to cache objects. Default value is <i>true</i> . If set <i>true</i> and the cached object was not longer referenced by your application but only by the cache, it can be reclaimed by the garbage collector. If set <i>false</i> it's strongly recommended to the <i>timeout</i> property to prevent memory problems of the JVM.

**Recommendation:**

If you take care of cache synchronization (or use autoSync property) and be aware of dirty reads, this implementation is useful for read-only or less update centric classes.

**ObjectCacheTwoLevelImpl**

[ObjectCacheTwoLevelImpl](#) is a two level [ObjectCache](#) implementation with a transactional session- and a shared application-cache part.

The first level is a transactional session cache that cache objects till [PersistenceBroker](#)#close() or if a PB-tx is running till #abortTransaction() or #commitTransaction() was called. On commit all objects reside in the session cache will be pushed to the application cache. If objects be new materialized from the database (e.g. when achieve a query), the full materialized objects will be pushed immediately to the application cache (more precisely, if the application cache doesn't contain the "new materialized" objects).

The second level cache can be specified with the *applicationCache* property. Properties of the specified application cache are allowed too. Here is an example how to use the two level cache with [ObjectCacheDefaultImpl](#) as second level cache.

```
<object-cache class="org.apache.ojb.broker.cache.ObjectCacheTwoLevelImpl">
  <!-- meaning of attributes, please see docs section "Caching" -->
  <!-- common attributes -->
```

```

<attribute attribute-name="cacheExcludes" attribute-value="" />

<!-- ObjectCacheTwoLevelImpl attributes -->
<attribute attribute-name="applicationCache"
  attribute-value="org.apache.ojb.broker.cache.ObjectCacheDefaultImpl" />
<attribute attribute-name="copyStrategy"
  attribute-value="org.apache.ojb.broker.cache.ObjectCacheTwoLevelImpl$CopyStrategyImpl" />
<attribute attribute-name="forceProxies" attribute-value="true" />

<!-- ObjectCacheDefaultImpl attributes -->
<attribute attribute-name="timeout" attribute-value="900" />
<attribute attribute-name="autoSync" attribute-value="true" />
<attribute attribute-name="cachingKeyType" attribute-value="0" />
<attribute attribute-name="useSoftReferences" attribute-value="true" />
</object-cache>

```

The most important characteristic of the two-level cache is that all objects put to or read from the application cache are copies of the target object, so the cached objects never could be corrupted by the user when changing fields, because all operations done on copies of objects cached in the application cache (in contrast to [ObjectCacheDefaultImpl](#)).

The strategy to make copies of the persistent objects is pluggable and can be specified by the *copyStrategy* property which expects an implementation of the `ObjectCacheTwoLevelImpl.CopyStrategy` interface.

The *default* `ObjectCacheTwoLevelImpl.CopyStrategy` implementation make copies based on the [field-descriptors](#) of the cached object and set these values in a new instance of the cached object. If you lookup a cached object with *1:n* or *m:n* relation a query is needed to get the ID's of the referenced objects, because in application cache only "flat" objects without references/reference-info will be cached.

#### Note:

This two-level cache implementation does not guarantee that cache and persistent storage (e.g. database) are always consistent, because the session cache push the persistent objects to application cache *after* the PB-tx was committed.

Let us assume that thread 1 (using broker 1) update objects A1, A2, ... within a transaction and does commit the tx. Now before OJB could execute the after commit call on thread 1 to force session cache to push the objects to the application cache, thread 2 (using broker 2) lookup and update object A2 too (improbably but could happen, because thread 1 has already committed the objects A1, A2,... to the persistent storage) and push A2 to application cache. After this thread 1 was able to perform the after commit call and the 'outdated' version of A2 was pushed to the application cache overwriting the actual version of A2 in cache - cache and persistent storage are out of synchronization.

To avoid writing of outdated data to the persistence storage [optimistic locking](#) can be used. OL will not prevent the above scenario, but if it happens and e.g. broker 3 read the outdated object A1 from the cache and try to perform an update of A1, an optimistic locking exception will be thrown. So it is guaranteed that the persistent storage is always consistent.

A possibility to completely prevent synchronization problems of cache and persistent storage is the usage of [pessimistic locking](#) (if the used api supports it) with an adequate locking isolation level. If only one thread/broker could modify an object at the same time and the lock will be released after all work is done, the above scenario can't happen.

To avoid corrupted data, all objects cached by users (using the methods of the *ObjectCache* interface) will never be pushed to the application cache, they will be buffered in the session cache till it was cleared.

Implementation configuration properties:

Property Key	Property Values
applicationCache	Specifies the <a href="#">ObjectCache</a> implementation used as application cache (second level cache). By default <a href="#">ObjectCacheDefaultImpl</a> was used. It's recommended to use a shared cache implementation (all used PB instances should access the same pool of objects - e.g. by using a static Map in cache implementation).
copyStrategy	Specifies the implementation class of the <code>ObjectCacheTwoLevelImpl.CopyStrategy</code>

	interface, which was used to copy objects on read and write operations to application cache. If not set, a default implementation was used ( <code>ObjectCacheTwoLevelImpl.CopyStrategyImpl</code> make field-descriptor based copies of the cached objects).
forceProxies	<p>If <i>true</i> on materialization of cached objects, all referenced objects will be represented by proxy objects (independent from the proxy settings in <a href="#">reference-</a> or <a href="#">collection-descriptor</a>).</p> <p><b>Note:</b> To use this feature all persistence capable objects have to be interface based <b>or</b> the <code>ProxyFactory</code> and <code>IndirectionHandler</code> implementation classes supporting dynamic proxy enhancement for all classes (see <a href="#">OJB.properties</a>, find more information about <a href="#">proxy settings here</a>).</p>

### ObjectCachePerBrokerImpl

[ObjectCachePerBrokerImpl](#) is a local/session cache implementation allows to have dedicated caches per [PersistenceBroker](#) instance.

Note: When the used broker instance was closed (returned to pool) the cache was cleared.

This cache implementation is not synchronized with the other `ObjectCache` instances, there will be no automatic refresh of objects modified/updated by other threads (*PersistenceBroker* instances).

So, objects modified by other threads will not influence the cached objects, because for each broker instance the objects will be cached separately and each thread should use it's own *PersistenceBroker* instance.

### ObjectCacheEmptyImpl

This is an *no-op* `ObjectCache` implementation. Useful when caching was not desired.

#### Note:

This implementation supports *circular references* as well (since OJB 1.0.2, materialization of object graphs with circular references will be handled internally by OJB).

### ObjectCacheJCSImpl

A shared `ObjectCache` implementation using a JCS region for each classname. More info see [turbine-JCS](#).

### ObjectCacheOSCacheImpl

You're basically in good shape at this point. Now you've just got to set up `OSCache` to work with OJB. Here are the steps for that:

- Download `OSCache` from [OSCache](#). Add the `oscache-2.0.x.jar` to your project so that it is in your classpath (for Servlet/J2EE users place in your WEB-INF/lib directory).
- Download `JavaGroups` from [JavaGroups](#). Add the `javagroups-all.jar` to your classpath (for Servlet/J2EE users place in your WEB-INF/lib directory).
- Add `oscache.properties` from your `OSCache` distribution to your project so that it is in the classpath (for Servlet/J2EE users place in your WEB-INF/classes directory). Open the file and

make the following changes:

1. Add the following line to the CACHE LISTENERS section of your oscache.properties file:  
*cache.event.listeners=com.opensymphony.oscache.plugins.clustersupport.JavaGroupsBroadcastingListe*
  2. Add the following line at the end of the oscache.properties file (your network must support multicast):  
*cache.cluster.multicast.ip=231.12.21.132*
- Add the following class to your project (feel free to change package name, but make sure that you specify the full qualified class name in configuration files). You can find source of this class under *db-ojb/contrib/src/ObjectCacheOSCacheImpl* or copy this source:

```

public class ObjectCacheOSCacheImpl implements ObjectCacheInternal
{
    private Logger log = LoggerFactory.getLogger(ObjectCacheOSCacheImpl.class);
    private static GeneralCacheAdministrator admin = new
GeneralCacheAdministrator();
    private static final int REFRESH_PERIOD =
com.opensymphony.oscache.base.CacheEntry.INDEFINITE_EXPIRY;

    public ObjectCacheOSCacheImpl()
    {
    }

    public ObjectCacheOSCacheImpl(PersistenceBroker broker, Properties prop)
    {
    }

    public void cache(Identity oid, Object obj)
    {
        try
        {
            /*
            Actually, OSCache sends notifications (Events) only on flush
            events. The putInCache method do not flush the cache, so no event is
            sent.
            The ObjectCacheOSCacheInternalImpl should force OSCache to flush the
            entry
            in order to generate an event. This guarantee that other nodes
            always
            in sync with the DB.
            Alternative a non-indefinite refresh-period could be used in
            conjunction
            with optimistic-locking for persistent objects.
            */
            remove(oid);
            admin.putInCache(oid.toString(), obj);
        }
        catch(Exception e)
        {
            log.error("Error while try to cache object: " + oid, e);
        }
    }

    public void doInternalCache(Identity oid, Object obj, int type)
    {
        cache(oid, obj);
    }

    public boolean cacheIfNew(Identity oid, Object obj)
    {
        boolean result = false;
        Cache cache = admin.getCache();
        try
        {
            cache.getFromCache(oid.toString());
        }
        catch(NeedsRefreshException e)
        {
        }
    }
}

```



```

        try
        {
            cache.putInCache(oid.toString(), obj);
            result = true;
        }
        catch(Exception e1)
        {
            cache.cancelUpdate(oid.toString());
            log.error("Error while try to cache object: " + oid, e);
        }
    }
    return result;
}

public Object lookup(Identity oid)
{
    Cache cache = admin.getCache();
    try
    {
        return cache.getFromCache(oid.toString(), REFRESH_PERIOD);
    }
    catch(NeedsRefreshException e)
    {
        // not found in cache
        if(log.isDebugEnabled()) log.debug("Not found in cache: " + oid);
        cache.cancelUpdate(oid.toString());
        return null;
    }
    catch(Exception e)
    {
        log.error("Unexpected error when lookup object from cache: " + oid,
e);
        cache.cancelUpdate(oid.toString());
        return null;
    }
}

public void remove(Identity oid)
{
    try
    {
        if(log.isDebugEnabled()) log.debug("Remove from cache: " + oid);
        admin.flushEntry(oid.toString());
    }
    catch(Exception e)
    {
        throw new RuntimeException("Unexpected error when remove object
from cache: " + oid, e);
    }
}

public void clear()
{
    try
    {
        if(log.isDebugEnabled()) log.debug("Clear cache");
        admin.flushAll();
    }
    catch(Exception e)
    {
        throw new RuntimeException("Unexpected error while clear
cache", e);
    }
}
}

```

To allow usage of this implementation as application cache level in the [two-level cache](#) implement the internal object cache interface instead of the standard one.

Now *OSCache* can be used by OJB as standalone cache (by declaring the implementation class on

[connection- or class-level](#)) or as application cache in the [two-level cache](#).

#### More implementations ...

Additional *ObjectCache* implementations can be found in *org.apache.ojb.broker.cache* package.

#### 5.5.14.5. Distributed ObjectCache?

If OJB was used in a clustered environment it is mandatory to distribute all shared cached objects across different JVM. OJB does not support distributed caching "out of the box", to do this an external caching library is needed, e.g. the [OSCache implementation](#) supports distributed caching. More information how to setup OJB in clustered environments see [clustering howto](#).

#### 5.5.14.6. Implement your own cache

The OJB cache implementations are quite simple but should do a good job for most scenarios. If you need a more sophisticated cache or need to plug in a proprietary caching library you'll write your own implementation of the [ObjectCache](#) interface. Integration of your implementation in OJB is easy since the object cache is a pluggable component. All you have to do, is to declare it on [connection- or class-level](#). Here an example how to declare the new implementation on connection level:

```
<jdbc-connection-descriptor
  jcd-alias="myDefault"
  ...
>
  <object-cache class="my.ObjectCacheMyImpl">
    <attribute attribute-name="cacheExcludes" attribute-value="" />
    ... additional attributes of the cache
  </object-cache>
</jdbc-connection-descriptor
```

If interested to get more detailed information about the "type" of the objects to cache (objects written to DB, new materialized objects,...) implement the [ObjectCacheInternal](#) interface (For an implementation example see source for [ObjectCacheTwoLevelImpl](#)).

#### Note:

Of course we are interested in your solutions! If you have implemented something interesting, just contact us.

#### 5.5.14.7. Future prospects

In OJB 1.1 the caching part will be rewritten to get rid of static classes, factories and member variables.

### 5.5.15. Sequence Manager

#### 5.5.15.1. The OJB Sequence Manager

All [sequence manager](#) implementations shipped with OJB you can find under the *org.apache.ojb.broker.util.sequence* package using the following naming convention *SequenceManagerXXXImpl*.

#### Automatic assignment of unique values

As mentioned in [mapping tutorial](#) OJB provides a mechanism to automatically assign unique values for primary key attributes. You just have to enable the *autoincrement* attribute in the respective *field-descriptor* of the XML repository file as follows:

```
<class-descriptor
  class="my.Article"
  table="ARTICLE"
>
  <field-descriptor
    name="articleId"
    column="ARTICLE_ID"
    jdbc-type="INTEGER"
    primaryKey="true"
    autoincrement="true"
  />
  ...
</class-descriptor>
```

This definitions contains the following information:

The attribute `articleId` is mapped on the table's column `ARTICLE_ID`. The JDBC Type of this column is `INTEGER`. This is a primary key column and OJB shall automatically assign unique values to this attribute.

This mechanism works for all whole-numbered column types like `BIGINT`, `INTEGER`, `SMALLINT`,... and for `CHAR`, `VARCHAR` columns. This mechanism helps you to keep your business logic free from code that computes unique ID's for primary key attributes.

### Force computation of unique values

By default OJB triggers the computation of unique ids during calls to `PersistenceBroker.store(...)`. Sometimes it will be necessary to have the ids computed in advance, before a new persistent object was written to database. This can be done by simply obtaining the `Identity` of the respective object as follows:

```
Identity oid = broker.serviceIdentity().buildIdentity(Object
newPersistentObject);
```

This creates an `Identity` object for the new persistent object and set all primary key values of the new persistent object - But it only works if `autoincrement` is enabled for the primary key fields.

#### Warning:

Force computation of unique values is not allowed when using *database based Identity columns* for primary key generation (e.g via `Identity column supporting sequence manager`), because the *real* PK value is at the earliest available after database insert operation. If you nevertheless force PK computing, OJB will use a temporary dummy PK value in the `Identity` object and this may lead to unexpeted behavior.

Info about lookup persistent objects by primary key fields [see here](#).

### How to change the sequence manager?

To enable a specific `SequenceManager` implementation declare an *sequence-manager* attribute within the `jdbc-connection-descriptor` element in the `repository file`.

If no *sequence-manager* was specified in the `jdbc-connection-descriptor`, OJB use a default sequence manager implementation (default was `SequenceManagerHighLowImpl`).

Further information you could find in the `repository.dtd` section `sequence-manager` element.

Example `jdbc-connection-descriptor` using a *sequence-manager* tag:

```
<jdbc-connection-descriptor
  jcd-alias="farAway"
  platform="Hsqldb"
  jdbc-level="2.0"
  driver="org.hsqldb.jdbcDriver"
```

```

        protocol="jdbc"
        subprotocol="hsqldb"
        dbalias="../OJB_FarAway"
        username="sa"
        password=""
        batch-mode="false"
    >

    <connection-pool
        maxActive="5"
        whenExhaustedAction="0"
        validationQuery="select count(*) from OJB_HL_SEQ"
    />

    <sequence-manager
className="org.apache.ojb.broker.util.sequence.SequenceManagerHighLowImpl">
        <attribute attribute-name="seq.start" attribute-value="10000"/>
        <attribute attribute-name="grabSize" attribute-value="20"/>
    </sequence-manager>
</jdbc-connection-descriptor>

```

The mandatory *className* attribute needs the **full-qualified class name** of the desired sequence-manager implementation. If a implementation needs configuration properties you pass them using [custom attribute](#) tags with *attribute-name* represents the property name and *attribute-value* the property value. Each sequence manager implementation shows all properties on the according javadoc page.

### SequenceManager implementations

Source code of all [SequenceManager](#) implementations can be found in `org.apache.ojb.broker.util.sequence` package.

If you still think something is missing, you can just write your [own](#) sequence manager implementation.

#### High/Low sequence manager

The [sequence manager](#) implementation class

`ojb.broker.util.sequence.SequenceManagerHighLowImpl` and is able to generate ID's unique to a given object and all [extent objects](#) declared in the objects class descriptor.

If you ask for an ID using an interface with several implementor classes, or a baseclass with several subclasses the returned ID have to be unique accross all tables representing objects of the interface or base class (more see [here](#)).

It's also possible to use this implementation in a *global mode*, generate global unique id's.

This implementation **needs an internal database table and object mapping declaration** to persist the used sequences. The table structure can be found in [in platform guide](#) the object metadata mapping can be found in [OJB internal mapping file \(called repository internal.xml\)](#).

To declare this *sequence manager* implementation specify a `sequence-manager` element within the [jdbc-connection-descriptor](#):

```

<sequence-manager className=
    "org.apache.ojb.broker.util.sequence.SequenceManagerHighLowImpl">

    <attribute attribute-name="seq.start" attribute-value="5000"/>
    <attribute attribute-name="grabSize" attribute-value="20"/>
    <attribute attribute-name="autoNaming" attribute-value="true"/>
    <attribute attribute-name="globalSequenceId" attribute-value="false"/>
    <!-- deprecated settings -->
    <attribute attribute-name="sequenceStart" attribute-value="5000"/>
</sequence-manager>

```

The property `seq.start` (or deprecated `sequenceStart`) define the start value of the id

generation (default was '1'). It's recommended to use start values greater than '0' to avoid problems with primitive primary key fields when used as foreign key in references.

With property `grabSize` you set the size of the assigned ID's kept in memory for each *autoincrement* field. If the assigned ID's are exhausted a database call is made to lookup the next bunch of ID's (default `grabSize` is 20).

If OJB was shutdown/redeployed all unused assigned ID's are lost.

If property `globalSequenceId` was set `true` you will get global unique ID's over all persistent objects. Default was `false`.

NOTE: If the database is already populated or the global sequence name in `OJB_HL_SEQ` database table was removed (by accident), the `seq.start` value must be greater than the biggest PK value in database.

This sequence manager implementation supports *user defined sequence-names* as well as *automatic generated sequence-names* to manage the sequences - more about [sequence-names here](#).

The attribute `autoNaming` can be used to enable *auto-generation* of *sequence-names*, default value is `true`.

More info about attribute [autoNaming here](#).

Limitations:

- do **not** use in **managed environments** when connections were enlisted in running transactions, e.g. when using `DataSources` of an application server
- if set connection-pool attribute 'whenExhaustedAction' to 'block' (wait for connection if connection-pool is exhausted), under heavy load this sequence manager implementation can block application.
- superfluously to mention, do not use if other non-OJB applications insert objects too

#### In-Memory sequence manager

Another sequence manager implementation is a *In-Memory* version called `org.apache.ojb.broker.util.sequence.SequenceManagerInMemoryImpl`.

Only the first time an UID was requested for a object, the manager query the database for the max value of the target column - all following request were performed in memory. This implementation ditto generate unique ID's across all *extents*, using the same mechanism as the High/Low sequence manager implementation.

To declare this *sequence manager* implementation specify a `sequence-manager` element within the [jdbc-connection-descriptor](#):

```
<sequence-manager
className="org.apache.ojb.broker.util.sequence.SequenceManagerInMemoryImpl">
  <attribute attribute-name="seq.start" attribute-value="0"/>
  <attribute attribute-name="autoNaming" attribute-value="true"/>
</sequence-manager>
```

The property `seq.start` (or deprecated `sequenceStart`) define the start value of the id generation (default was '1'). It's recommended to use start values greater than '0' to avoid problems with primitive primary key fields when used as foreign key in references.

This sequence manager implementation supports *user defined sequence-names* as well as *automatic generated sequence-names* to manage the sequences - more about [sequence-names](#).

The attribute `autoNaming` can be used to enable *auto-generation* of *sequence-names*, default value is `true`.

More info about [autoNaming](#).

The specified sequences will only be used *in memory*. First time a sequence was used OJB does a `select max`-query to find the latest/greatest value for the *autoincrement* field and use this as



The auto-generated name will be set as *sequence-name* in the *field-descriptor*.

If set *false* OJB throws an exception if none *sequence-name* was found in *field-descriptor*, also OJB does NOT try to create a database sequence when for a given sequence name (specified in *field-descriptor*) no database sequence can be found.

The table below show additional *sequence properties*. To specify the properties use *custom-attributes* within the *sequence-manager* element.

The database sequence specific properties are generally speaking, see database user guide for detailed description.

Name	Value	Description	Supported By
seq.as	database specific, e.g. <i>INTEGER</i>	Database sequence specific property. Specifies the datatype of the sequence, the allowed datatypes depend on the used database implementation.	DB2
seq.start	1 ... max INTEGER	Database sequence specific property. Specifies the first sequence number to be generated. Allowed: 1 or greater.	Oracle, PostgreSQL, MaxDB/SapDB, DB2
seq.incrementBy	>=1	Database sequence specific property. Specifies the interval between sequence numbers. This value can be any positive or negative integer, but it cannot be 0. <i>Decrement sequences are currently not supported</i>	Oracle, PostgreSQL, MaxDB/SapDB, DB2
seq.maxValue	1 ... max INTEGER	Database sequence specific property. Set max value for sequence numbers.	Oracle, PostgreSQL, MaxDB/SapDB, DB2
seq.minValue	min INTEGER	Database sequence specific property. Set min value for sequence numbers. <i>Negative sequences are not tested as yet.</i>	Oracle, PostgreSQL, MaxDB/SapDB, DB2
seq.cycle	true/false	Database sequence specific property. If <i>true</i> , specifies that the sequence continues to generate values after reaching either its maximum or minimum value. If <i>false</i> , specifies that the sequence cannot generate more values	Oracle, PostgreSQL, MaxDB/SapDB, DB2

		after reaching its maximum or minimum value.	
seq.cache	>= 2	Database sequence specific property. Specifies how many values of the sequence Oracle preallocates and keeps in memory for faster access. Allowed values: 2 or greater. If set 0, an explicit <i>nocache</i> expression will be set.	Oracle, PostgreSQL, MaxDB/SapDB, DB2
seq.order	true/false	Database sequence specific property. If set <i>true</i> , guarantees that sequence numbers are generated in order of request. If <i>false</i> , a <i>no order</i> expression will be set.	Oracle, MaxDB/SapDB, DB2

Limitations:

- none known

#### Database sequences based high/low implementation

Based on the sequence manager implementation described [above](#), but use a high/low algorithm to avoid database access.

```
<sequence-manager
className="org.apache.ojb.broker.util.sequence.SequenceManagerSeqHiLoImpl">
  <attribute attribute-name="grabSize" attribute-value="20"/>
  <attribute attribute-name="autoNaming" attribute-value="true"/>
</sequence-manager>
```

With property `grabSize` you set the size of the assigned ID's kept in memory for each *autoincrement* field. If the assigned ID's are exhausted a database call is made to lookup the next bunch of ID's using the next database sequence (default *grabSize* is 20).

If OJB was shutdown/redeployed all unused assigned ID's are lost.

**Note:**

Keep in mind that the database sequence value does not correspond with the used value in the *autoincrement*-field (table column value).

Attribute [autoNaming](#) is the same as for [SequenceManagerNextValImpl](#).

This sequence manager implementation supports user defined *sequence-names* to manage the sequences (see [more](#)) or if not set in `field-descriptor` it is done automatic when [autoNaming](#) is enabled.

Limitations:

- superfluously to mention, do not use (or handle with care) if other non-OJB applications insert objects too

#### Database Identity-column based sequence manager



This sequence manager implementation supports database *Identity columns* (supported by MySQL, MsSQL, HSQL, ...). When using identity columns we have to do a trick to make the sequence manager work.

OJB identify each persistence capable object by a unique [ojb-Identity object](#). These ojb-Identity objects were created using the sequence manager instance to get UID's. Often these ojb-Identity objects were created before the persistence capable object was written to database.

When using Identity columns it is not possible to retrieve the next valid UID before the object was written to database. As recently as the real object was written to database, you can ask the DB for the last generated UID. Thus in SequenceManagerNativeImpl we have to do a trick and use a 'temporary' UID till the object was written to database.

So for best compatibility try to avoid using *Identity columns* in your database model. If this is not possible, use this sequence manager implementation to work with database *Identity columns*.

To enable this sequence manager implementation set in your `jdbc-connection-descriptor`:

```
<sequence-manager
  className="org.apache.ojb.broker.util.sequence.SequenceManagerNativeImpl">
</sequence-manager>
```

To declare the identity column in the persistent class mapping [class-descriptor](#), add the following attributes to the primary key/identity key [field-descriptor](#):

`primaryKey="true"`, `autoincrement="true"` and `access="readonly"`

The first and second attributes are the same as all sequence manager implementations use to support autoincrement PK fields, the third one is mandatory for database *Identity columns* only.

```
<field-descriptor
  name="identifier"
  column="NATIVE_ID"
  jdbc-type="BIGINT"
  primaryKey="true"
  autoincrement="true"
  access="readonly"/>
```

Limitations:

- The Identity columns have to **start with value greater than '0'** and should never be negative.
- Use of Identity columns is **not extent aware** (This may change in further versions). More info [here](#).

#### Stored Procedures based (Oracle-style) sequencing

(By Ryan Vanderwerf et al.)

"This solution will give those seeking an oracle-style sequence generator a final answer (Identity columns really suck). If you are using multiple application servers in your environment, and your database does not support read locking like Microsoft SQL Server, this is the only safe way to guarantee unique keys (HighLowSequenceManager WILL give out duplicate keys, and corrupt your data)".

The SequenceManagerStoredProcedureImpl implementation enabled database sequence key generation in a *Oracle-style* for all databases (e.g. MSSQL, MySQL, DB2, ...).

To declare this *sequence manager* implementation specify a `sequence-manager` element within the [jdbc-connection-descriptor](#):

```
<sequence-manager
className="org.apache.ojb.broker.util.sequence.SequenceManagerStoredProcedureImpl">
  <attribute attribute-name="autoNaming" attribute-value="true"/>
</sequence-manager>
```

For attribute [autoNaming](#) see.

This sequence manager implementation supports user defined [sequence-names](#) to manage the sequences or if not set in `field-descriptor` it is done automatic when [autoNaming](#) is enabled.

- Add a new table OJB\_NEXTVAL\_SEQ to your database.
- You will also need a stored procedure called `objb_nextval_proc` that will take care of giving you a guaranteed unique sequence number.

Below you can find the stored procedures you need to use sequencing for MSSQL server and Informix.

You have to adapt the scripts for other databases (We are interested in scripts for other databases).

Here you can find the currently supported databases and the statements to create the sql functions:

Database	Table Statement	Stored Procedure
MSSQL	<pre>DROP TABLE OJB_NEXTVAL_SEQ; CREATE TABLE OJB_NEXTVAL_SEQ (     SEQ_NAME VARCHAR(150) NOT NULL,     MAX_KEY INTEGER,     CONSTRAINT SYS_PK_OJB_NEXTVAL PRIMARY KEY(SEQ_NAME) );</pre>	<pre>CREATE PROCEDURE OJB_NEXTVAL_PROC @SEQ_NAME varchar(150) AS declare @MAX_KEY BIGINT -- return an error if -- sequence does not exist -- so we will know if someone -- truncates the table set @MAX_KEY = 0  UPDATE OJB_NEXTVAL_SEQ SET @MAX_KEY = MAX_KEY = MAX_KEY + 1 WHERE SEQ_NAME = @SEQ_NAME  if @MAX_KEY = 0 select 1/0 else select @MAX_KEY RETURN @MAX_KEY</pre>
Informix	<pre>drop table objb_nextval_seq; create table objb_nextval_seq (     seq_name varchar(250,0) not null,     max_key      int8,     primary key(seq_name) );</pre>	<pre>create function objb_nextval_proc (out arg1 int8, arg2 varchar(250,250)) returns int8; let arg1 = 0; update objb_nextval_seq set max_key = max_key + 1 where seq_name = arg2; select max_key into arg1 from objb_nextval_seq where seq_name = arg2; return arg1; end function;</pre>
Oracle	TODO	TODO

Limitations:

- currently none known

#### Microsoft SQL Server 'uniqueidentifier' type (GUID) sequencing

For those users you are using SQL Server 7.0 and up, the uniqueidentifier was introduced, and allows for your rows Primary Keys to be GUID's that are guaranteed to be unique in time and space.

However, this type is different than the Identity field type, whereas there is no way to programmatically retrieve the inserted value. Most implementations when using the u.i. field type set a default value of "newid()". The SequenceManagerMSSQLGuidImpl class manages this process for you as if it was any normal generated sequence/identity field.

Assuming that your PK on your table is set to 'uniqueidentifier', your field-description would be the same as using any other SequenceManager:

```
<field-descriptor
  name="guid"
  column="document_file_guid"
  jdbc-type="VARCHAR"
  primarykey="true"
  autoincrement="true"
/>
```

Note that the jdbc-type is a VARCHAR, and thus the attribute (in this case 'guid') on your class should be a String (SQL Server does the conversion from the String representation to the binary representation when retrieved/set).

You also need to turn on the SequenceManager in your jdbc-connection-descriptor like this:

```
<sequence-manager
className="org.apache.ojb.broker.util.sequence.SequenceManagerMSSQLGuidImpl"
/>
```

Limitations:

-This will only work with SQL Server 7.0 and higher as the uniqueidentifier type was not introduced until then.

This works well in situations where other applications might be updated the database as well, because it guarantees (well, as much as Microsoft can guarantee) that there will be no collisions between the Guids generated.

#### The sequence-name attribute

Several [SequenceManager](#) implementations using *sequences* (synonyms: *sequence objects*, *sequence generators*) to manage the ID generation. Sequences are *entities* which generate unique ID's using e.g. database table per sequence, database row per sequence or an in-memory java-object.

To address the sequences, each *sequence* has an unique *sequence-name*.

In OJB the sequence-name of an autoincrement field is declared in a *sequence-name* attribute within the [field-descriptor](#).

```
<class-descriptor
class="org.greatest.software.Person"
table="GS_PERSON"
>
<field-descriptor
  name="id"
  column="ID_PERSON"
  jdbc-type="INTEGER"
  sequence-name="ID_PERSON"
/>
```

```

    primaryKey="true"
    autoincrement="true"
    sequence-name="PERSON_SEQUENCE"
  />
  ...
</class-descriptor>

```

The *sequence-name* attribute in the *field-descriptor* is only needed if the used sequence manager supports sequences, the field should be *autoincremented* and the auto-assign of a sequence-name is not desired.

#### Note:

Each *sequence-name* has be [extent-aware](#).

If you don't specify a sequence name in the *field-descriptor* it is possible to auto-assign a sequence-name by OJB if [autoNaming](#) is supported by the used sequence manager implementation.

### The autoNaming property

All shipped [SequenceManager](#) implementations using *sequences* for ID generation support a property called *autoNaming* which can be declared as a *custom attribute* within the *sequence-manager* element:

```

<sequence-manager
  className="org.apache.ojb.broker.util.sequence.SequenceManagerNextValImpl">
  <attribute attribute-name="autoNaming" attribute-value="true" />
</sequence-manager>

```

If set *true* OJB try to build a *sequence name* by it's own (a simple algorithm was used to auto-generate the sequence name - more details how it works in [pitfalls section](#)) and set this name as *sequence-name* in the [field-descriptor](#) of the autoincrement field if no [sequence name](#) is specified.

If set *false* the sequence manager throw an exception if a *sequence name* can't be found or was not declared in the *field-descriptor* of the autoincrement field. In this case OJB expects a valid *sequence-name* in the *field-descriptor*.

If the attribute *autoNaming* is set *false* the sequence manager never try to auto-generate a *sequence-name* (more detailed info [here](#)). If set *true* and a *sequence-name* is set in the *field-descriptor*, the *SequenceManager* will use this one and does **not** override the existing one.

The default setting is *true*.

### How to write my own sequence manager?

Very easy to do, just write a implementation class of the interface `org.apache.ojb.broker.util.sequence.SequenceManager`. OJB use a factory (`SequenceManagerFactory`) to obtain sequence manager instances.

This Factory can be configured to generate instances of your specific implementation by adding a *sequence-manager* tag in the *jdbc-connection-descriptor*.

```

<sequence-manager className="my.SequenceManagerMYImpl">
</sequence-manager>

```

That's it!

If your sequence manager implementation was derived from `org.apache.ojb.broker.util.sequence.AbstractSequenceManager` it's easy to pass configuration properties to your implementation using [custom attributes](#).

```
<sequence-manager className="my.SequenceManagerMYImpl">
  <attribute attribute-name="myProperty" attribute-value="test" />
</sequence-manager>
```

With

```
public String getConfigurationProperty(String key, String defaultValue)
```

method get the properties in your implementation class.

#### Note:

Of course we interested in your solutions! If you have implemented something interesting, just contact us.

## Questions

### When using sequence-name attribute in field-descriptor?

Most `SequenceManager` implementations based on [sequence names](#). If you want retain control of sequencing use your own `sequence-name` attribute in the `field-descriptor`. In that case you are responsible to use the same name across extents, we call it [extent-aware](#) (see more info about [extents and polymorphism](#)). Per default the sequence manager build its own *extent aware* sequence name with a simple algorithm (see `org.apache.ojb.broker.util.sequence.SequenceManagerHelper#buildSequenceName` if necessary).

In most cases this should be sufficient. If you have a very complex data model and you will do many metadata changes in the repository file in future, then it could be better to explicit use `sequence-names` in the `field-descriptor`. See more [avoid pitfalls](#).

### What to hell does extent aware mean?

Say we have an abstract base class `Animal` and two classes `Dog` and `Cat` which extend `Animal`. For each non-abstract class we create a separate database table and [declare the inheritance](#) in OJB. Now it is possible to do a query like *give me all animals* and OJB will return all `Cat` and `Dog` objects. To make this working in OJB the ID's of `Dog` and `Cat` objects must be unique across the tables of both classes or else you may not get a valid query result (e.g. you can't query for the `Animal` with `id=23`, because in both tables such an id can exist).

The reason for this behaviour is the [org.apache.ojb.broker.Identity](#) class implementation (more details see javadoc/source of this class).

### How could I prevent auto-build of the sequence-name?

All shipped `SequenceManager` implementations which using sequence names for UID generation, support by default auto-build (autoNaming) of the sequence name if none was found in the `field-descriptor`.

To prevent this, all relevant SM implementations support a `autoNaming` property - set via attribute element. If set `false` OJB doesn't try to build sequence names automatic.

```
<sequence-manager
className="org.apache.ojb.broker.util.sequence.SequenceManagerNextValImpl">
  <attribute attribute-name="autoNaming" attribute-value="false" />
</sequence-manager>
```

Keep in mind that user defined sequence names have to be [extent-aware](#).

### Sequence manager handling using multiple databases

If you use multiple databases you have to declare a sequence manager in each

`jdbc-connection-descriptor`. If you don't specify a sequence manager OJB use a default one (currently `ojb.broker.util.sequence.SequenceManagerHighLowImpl`).

#### One sequence manager with multiple databases?

OJB was intended to use a sequence manager per database. But it shouldn't be complicated to realize a global sequence manager solution by writing your own `SequenceManager` implementation.

#### Can I get direct access to the sequence manager?

That's no problem:

```
PersistenceBroker broker =
PersistenceBrokerFactory.createPersistenceBroker(myPBKey);
SequenceManager sm = broker.getServiceSequenceManager();
...
broker.close();
```

If you use `autoincrement=true` in your `field-descriptor`, there is no reason to obtain UID directly from the sequence manager or to handle UID in your object model.

Except when using *user-defined* sequence manager implementations, in this case it could be needed.

#### Note:

Don't use `SequenceManagerFactory#getSequenceManager(PersistenceBroker broker)`, this method returns a new sequence manager instance for the given broker instance and not the current used SM instance of the given `PersistenceBroker` instance]

#### Any known pitfalls?

- When using *sequences* based sequence manager implementations it's possible to enable auto-generation of *sequence names* - see [autoNaming section](#). To build the *sequence name* a simple algorithm was used. The algorithm try to get the *top-level class* of the field's (the *autoincrement field-descriptor*) enclosing class, if no top-level class was found, the table name of the field's enclosing class was used. If a top-level class was found, the first found *extent class* table name was used as sequence name. The algorithm can be found in `org.apache.ojb.broker.util.sequence.SequenceManagerHelper#buildSequenceName`. When using base classes/interfaces with *extent classes* (declared in the *class-descriptor*) based on different database tables and the *extent-class* entries in repository often change (e.g. add new top-level class, change top-level class), the algorithm could be corrupted after restart of OJB, because the first found extent class's table name could be change, hence the used sequence-name. Now the ID generation start over and could clash with existing ID's. To avoid this, remove the implementation specific internal sequence name entry (e.g. from table `OJB_HL_SEQ` when using the Hi/Lo implementation, or remove the database sequence entry when using the 'Nextval' implementation) or use custom sequence name attributes in the field descriptor.

## 5.5.16. OJB logging configuration

### 5.5.16.1. Logging in OJB

For generating log messages, OJB provides its own, simplistic logging component [PoorMansLoggerImpl](#), but is also able to use the two most common Java logging libraries, [commons-logging](#) (which is actually a wrapper around several logging components) and [Log4j](#). In addition, it is also possible to define your own logging implementation.

Per default, OJB uses its own [PoorMansLoggerImpl](#) which does not require configuration and prints to `stdout`.

### 5.5.16.2. Logging configuration within OJB

#### How and when OJB determines what kind of logging to use

Logging is the first component of OJB that is initialized. If you access any component of OJB, logging will be initialized first before that component is doing anything else. Therefore, you'll have to provide for the configuration of logging before you access OJB in your program (this is mostly relevant if you plan to initialize OJB at runtime as is described [below](#)). Please note that logging configuration is independent of the configuration of other parts of OJB, namely the runtime (via [OJB.properties](#)) and the database/repository (via [repository.xml](#)).

These are the individual steps OJB performs in order to initialize the logging component:

1. First, OJB checks whether the system property `org.apache.ojb.broker.util.logging.Logger.class` is set. If specified, this property gives the fully qualified class name of the logger class (a class implementing the [Logger](#) interface). Along with this property, another property is then read which may specify a properties file for this logger class, `org.apache.ojb.broker.util.logging.Logger.configFile`.
2. If this property is not set, then OJB tries to read the file `OJB-logging.properties`. The name and path of this file can be changed by setting the runtime property of the same name. See [below](#) for the contents of this file.
3. For backwards compatibility, OJB next tries to read the logging settings from the file [OJB.properties](#) which is the normal runtime configuration file of OJB. Again, the name and path of this file can be changed by setting the runtime property of the same name. This file may contain the same entries as the `OJB-logging.properties` file.
4. If the the `OJB.properties` file does not contain logging settings, next it is checked whether the commons-logging log property `org.apache.commons.logging.Log` or the commons-logging log factory system property `org.apache.commons.logging.LogFactory` is set. If that's the case, OJB will use commons-logging for its logging purposes.
5. Next, OJB checks for the presence of the Log4j properties file `log4j.properties`. If it is found, the OJB uses Log4j directly (without commons-logging).
6. Finally, OJB tries to find the commons-logging properties file `commons-logging.properties` which when found directs OJB to use commons-logging for its logging.
7. If none of the above is true, or if the specified logger class could not be found or initialized, then OJB defaults to its `PoorMansLoggerImpl` logger which simply logs to `stdout`.

The only OJB component whose logging is not initialized this way, is the boot logger which is used by logging component itself and a few other core components. It will (for obvious reasons) always use [PoorMansLoggerImpl](#) and therefore log to `stdout`. You can define the log level of the boot logger via the `OJB.bootLogLevel` system property. Per default, **WARN** is used.

#### Configuration of logging for the individual components

Regardless of the logging implementation that is used by OJB, the configuration is generally similar. The individual logging implementations mainly differ in the syntax and in the configuration of the format of the output and of the output target (where to log to). See below for specific details and examples.

In general, you specify a default log level and for every component (usually a class) that should log differently, the amount and level of detail that is logged about that component. These are the levels:

**DEBUG**

Messages that express what OJB is currently doing. This is the most detailed debugging level

**INFO**

Informational messages

**WARN**

Warnings that may denote potential problems (this is the default level)

**ERROR**

As the name says, this level is for errors which means that some action could not be completed successfully

**FATAL**

Fatal errors which usually prevent an application from continuing

The levels **DEBUG** and **INFO** usually result in a lot of log messages which will reduce the performance of the application. Therefore these levels should only be used when necessary.

There are two special loggers to be aware of. The **boot logger** is the logger used by the logging component itself as well as a few other core components. It will therefore always use the [PoorMansLoggerImpl](#) logging implementation. You can configure its logging level via the `OJB.bootLogLevel` system property.

The **default logger** is denoted in the `OJB-logging.properties` file by the keyword `DEFAULT` instead of the class name. It is used by components that don't require their own logging configuration (usually because they are rather small components).

### 5.5.16.3. Logging configuration via configuration files

#### `OJB-logging.properties`

This file usually specifies which logging implementation to use using the `org.apache.ojb.broker.util.logging.Logger.class` property, and which properties file this logger has (if any) using the `org.apache.ojb.broker.util.logging.Logger.configFile` property. You should also use this file to specify log levels for OJB's components if you're not using Log4j or commons-logging (which have their own configuration files).

A typical `OJB-logging.properties` file looks like this:

```
# Which logger to use
org.apache.ojb.broker.util.logging.Logger.class=org.apache.ojb.broker.util.logging.PoorMan

# Configuration file of the logger
#org.apache.ojb.broker.util.logging.Logger.configFile=

# Global default log level used for all logging entities if not specified
ROOT.LogLevel=ERROR

# The log level of the default logger
DEFAULT.LogLevel=WARN

# Logger for PersistenceBrokerImpl class
org.apache.ojb.broker.core.PersistenceBrokerImpl.LogLevel=WARN

# Logger for RepositoryXmlHandler, useful for debugging parsing of
repository.xml!
org.apache.ojb.broker.metadata.RepositoryXmlHandler.LogLevel=WARN
```

#### `commons-logging.properties`

This file is used by [commons-logging](#). For details on its structure see [here](#).



An example `commons-logging.properties` file would be:

```
# Use Log4j
org.apache.commons.logging.Log=org.apache.commons.logging.impl.Log4JLogger

# Configuration file of the log
log4j.configuration=log4j.properties
```

**Note:**

Since commons-logging provides the same function as the logging component of OJB, it will likely be used as OJB's logging component in the near future.

## log4j.properties

The [commons-logging](#) configuration file. Details can be found [here](#).

A sample log4j configuration is:

```
# Root logging level is WARN, and we're using two logging targets
log4j.rootCategory=WARN, A1, A2

# A1 is set to be ConsoleAppender sending its output to System.out
log4j.appender.A1=org.apache.log4j.ConsoleAppender

# A1 uses PatternLayout
log4j.appender.A1.layout=org.apache.log4j.PatternLayout
log4j.appender.A1.layout.ConversionPattern=%-5r %-5p [%t] %c{2} - %m%n

# Appender A2 writes to the file "org.apache.ojb.log".
log4j.appender.A2=org.apache.log4j.FileAppender
log4j.appender.A2.File=org.apache.ojb.log

# Truncate the log file if it already exists.
log4j.appender.A2.Append=false

# A2 uses the PatternLayout.
log4j.appender.A2.layout=org.apache.log4j.PatternLayout
log4j.appender.A2.layout.ConversionPattern=%-5r %-5p [%t] %c{2} - %m%n

# Special logging directives for individual components
log4j.logger.org.apache.ojb.broker.metadata.RepositoryXmlHandler=DEBUG
log4j.logger.org.apache.ojb.broker.accesslayer.ConnectionManager=INFO
log4j.logger.org.apache.ojb.odmg=INFO
```

## Where to put the configuration files

OJB and the different logging implementations usually look up their configuration files in the classpath. So for instance, OJB searches for the `OJB-logging.properties` file directly in any of the entries of the classpath, directories and jar files. If the classpath contains in that order `some-library.jar`, `db-ojb.jar`, and `.`, then it will first search in the two jars (which themselves contain a directory structure in which OJB will search only in the root), and lastly in the current directory (which only happens if `.` is part of the classpath) but not in sub directories of it.

For applications, this classpath can easily be set either as an environment variable `CLASSPATH` or by using the commandline switch `-classpath` when invoking the java executable.

For web applications however, the server will define the classpath. There are specific folders in the webapp structure that are always part of the webapp's classpath. The one that is normally used to store configuration files, is the `classes` folder:

```
[folder containing webapps]\
  mywebapp\
    WEB-INF\
```

```
lib\
classes\  <-- Put your configuration files here
```

#### 5.5.16.4. Logging configuration at runtime

Sometimes you want to configure OJB completely at runtime (within your program). How to do that for logging depends on the used logging implementation, but you can usually configure them via system properties. The only thing to keep in mind is that logging in OJB is initialized as soon as you use one of its components, so you'll have to define the properties prior to using any OJB parts.

With system properties (which are accessible via `System.getProperty()` from within a Java program) you can always define the following OJB logging settings:

**org.apache.ojb.broker.util.logging.Logger.class**

Which logger OJB shall use

**org.apache.ojb.broker.util.logging.Logger.configFile**

The config file of the logger

**OJB-logging.properties**

The path to the logging properties file, default is `OJB-logging.properties`

**OJB.properties**

The path to the OJB properties file (which may contain logging settings), default is

`OJB.properties`

**org.apache.commons.logging.Log**

Use commons-logging with the specified log implementation

**org.apache.commons.logging.LogFactory**

Use commons-logging with the specified log factory

**log4j.configuration**

When using Log4j directly or via commons-logging, this is the Log4j configuration file (default is `log4j.properties`)

In addition, all Log4j properties (e.g. `log4j.rootCategory`) can be specified as system properties.

#### 5.5.16.5. Defining your own logger

It is rather easy to use your own logger. All you need to do is to provide a class that implements the interface [Logger](#). Besides the actual log methods (`debug`, `info`, `warn`, `error`, `fatal`) this interface defines a method `void configure(Configuration)` which is used to initialize the logger with the logging properties (as contained in `OJB-logging.properties`).

##### Note:

Because commons-logging performs a similar function to the OJB logging component, it is likely that it will be used as such in the near future. Therefore you're encouraged to also implement the [Log](#) interface which is nearly the same as the [Logger](#) interface.

### 5.5.17. Locking

#### 5.5.17.1. Introduction

Lock management is needed to synchronize concurrent access to objects from multiple transactions (possibly in clustered environments).

An example:

Assume there are two transactions `tx1` and `tx2` running. The first transaction `tx1` modify object A and perform an update. At the same time transaction `tx2` modify an object A' with the *same* identity `oidA`, so both objects represent the same row in DB table and both operate on the "same" row at the same time, thus the state of object with identity `oidA` is inconsistent.

Assume that `tx1` was committed, now the modified object `A'` in `tx2` based on outdated data (state before `A` changed). If now `tx2` commits object `A'` the changes of `tx1` will be overwritten with the "illegal" object `A'`.

The OJB lock manager is responsible for detecting such a conflict and e.g. doesn't allow `tx2` to read or modify objects with identity `oidA` as long as `tx1` commit or rollback (*pessimistic locking*). In other words, if in a running transaction an object in a with identity `oidA` has a *write lock*, the lock manager doesn't allow other transactions to acquire a *read* or *write lock* on the same identity `oidA` objects (for the sake of completeness: dependent on the used locking isolation level).

OJB supports two kind of locking strategies:

- [optimistic locking](#)
- [pessimistic locking](#)

OJB provide an pluggable low-level locking-api (located in `org.apache.ojb.broker.locking`) for *pessimistic locking*, which can be used by the top-level api's like [ODMG](#). The [PB-api](#) itself does not support *pessimistic locking* out of the box.

The base classes of the locking-api can be found in `org.apache.ojb.broker.locking` and the entry point is class [LockManager](#).

Object locking helps to guarantee data consistency without the need of database locks. During a transaction objects can be locked without the use a database connection, e.g the [ODMG](#) implementation lookup a database connection not until the transaction commit was called. If database locks are used, a connection is needed during the whole transaction.

### 5.5.17.2. Optimistic Locking

To control concurrent access to objects *optimistic locking* uses a version field on each persistent object.

Optimistic locking is supported by all API's (PB-api, ODMG-api, *JDO when it's done*).

Optimistic locking use an additional field/column for each persistent-object/table (*Long, Integer or Timestamp*) which is incremented each time changes are committed to the object, and is utilized to determine whether an optimistic transaction should succeed or fail. Optimistic locking is fast, because it checks data integrity only at update time.

1. In your table you need a dedicated column of type BIGINT, INTEGER or TIMESTAMP. Say the column is typed as INTEGER and named `VERSION_MAINTAINED_BY_OJB`.
2. You then need a (possibly private) attribute in your java class corresponding to the column. Say the attribute is defined as:

```
private int versionMaintainedByOjb;
```

3. in `repository.xml` you need a [field-descriptor](#) for this attribute. This field-descriptor must specify attribute `locking="true"`
4. The resulting field-descriptor will look as follows:

```
<field-descriptor
  name="versionMaintainedByOjb"
  column="VERSION_MAINTAINED_BY_OJB"
  jdbc-type="INTEGER"
  locking="true"
/>
```

#### Note:

Using of *TIMESTAMP* as optimistic locking field could cause problems, because dependent of the used operating system and database the precision of timestamp values differ (e.g. new value only after 10 ms or 1000 ms). In high concurrency applications this will cause problems.

### 5.5.17.3. Pessimistic-Locking

To control concurrent access to objects *pessimistic locking* uses shared and exclusive locks on persistent object (more precisely, on the identity object of the persistent object).

Pessimistic locking is currently used by the [ODMG-api](#) implementation. The [PB-api](#) does not support PL out of the box.

### Supported Isolation Levels

The OJB locking package supports four different [isolation level](#).

- read-uncommitted
- read-committed
- repeatable-read
- serializable
- (none)
- (optimistic)

The object locking isolation levels can be simply characterized as follows:

#### Uncommitted Reads

Obtaining two concurrent write locks on a given object is not allowed (case 14). Obtaining read locks is allowed even if another transaction is writing to that object (case 13). (That's why this level is also called *dirty reads*, because you can read lock objects with an existing write lock).

#### Committed Reads

Obtaining two concurrent write locks on a given object is not allowed. Obtaining read locks is allowed only if there is no write lock on the given object (case 13).

#### Repeatable Reads

Same as committed reads, but obtaining a write lock on an object that has been locked for reading by another transaction is not allowed (case 7).

#### Serializable transactions

As Repeatable Reads, but it is even not allowed to have multiple read locks on a given object (case 6).

The isolation level *none* and *optimistic* are self-explanatory:

*none* - don't lock objects associated with this isolation level

*optimistic* - don't lock objects associated with this isolation level, because optimistic locking was used instead.

Thus the lock manager will ignore all objects associated with these isolation level.

#### Note:

It's not needed to declare the *optimistic* isolation level in all persistent objects [class-descriptor](#) using this isolation level, because OJB will automatically detect an enabled optimistic locking and will bypass pessimistic locking. Only the proper settings for [optimistic locking](#) are mandatory.

#### Note:

The locking isolation levels named similar to the database transaction isolation level, but the definitions are different from it, so take care when comparing database transaction isolation level with object locking isolation level.

The proper behaviour of the different locking isolation level is checked by JUnit TestCases that implement test methods for each of the 17 cases specified in the above table. (See code for classes in package `org.apache.ojb.broker.locking` in OJB [test suite](#)).

The semantics of the strategies are defined by the following table:

Case	Name of TestCase	Transaction		Transaction-Isolationlevel			
		Tx1	Tx2	ReadUncon	ReadComm	Repeatable	Serializable
1	SingleRead	R		True	True	True	True
18	ReadThenF	R		True	True	True	True
		R					
2	UpgradeRe	R		True	True	True	True
		U					
3	ReadThenV	R		True	True	True	True
		W					
4	SingleWrite	W		True	True	True	True
5	WriteThenF	W		True	True	True	True
		R					
6	MultipleRea	R	R	True	True	True	False
7	UpgradeWit	R	U	True	True	False	False
8	WriteWithE	R	W	True	True	False	False
9	UpgradeWit	R	R	True	True	False	False
			U				
10	WriteWithM	R	R	True	True	False	False
			W				
11	UpgradeWit	R	R	True	True	False	False
		W					
12	WriteWithM	R	R	True	True	False	False
		W					
13	ReadWithE	W	R	True	False	False	False
14	MultipleWrit	W	W	False	False	False	False
15	ReleaseRea	R		True	True	True	True
		Rel	W				

16	ReleaseUp	U		True	True	True	True
		Rel	W				
17	ReleaseWri	W		True	True	True	True
		Rel	W				
	Acquire ReadLock	R					
	Acquire WriteLock	W					
	Upgrade Lock	U					
	Release Lock	Rel					

The table is to be read as follows. The acquisition of a single read lock on a given object (case 1) is allowed (returns True) for all isolationlevels. To upgrade a single read lock (case 2) is also allowed for all isolationlevels. If there is already a write lock on a given object for tx1, it is not allowed (returns False) to obtain a write lock from tx2 for all isolationlevels (case 14).

**Note:**

If the low-level locking api was used by hand:  
Not all [LockManager](#) implementation support the `LockManager#upgrade(...)` method (e.g. upgrade was delegated to write lock) or behavior of this method is a wee bit other than shown above. More detail see javadoc comment of the used `LockManager` implementation.

### How to specify locking isolation level

The locking isolation level can be specified *global* or *per class*.

The global setting is done in the [descriptor-repository](#) element:

```
<descriptor-repository version="1.0" isolation-level="read-uncommitted"
  proxy-prefetching-limit="50">
  ...
</descriptor-repository>
```

The isolation level of a class can be configured with the following attribute to a [class-descriptor](#):

```
<ClassDescriptor isolation-level="read-uncommitted" ...>
  ...
</ClassDescriptor>
```

If no *isolation-level* was specified a default isolation level was used - see interface [IsolationLevels](#). The semantics of isolation levels are described in [isolation level](#) section.

### Specify the LockManager Implementation

To specify the used lock manager implementation set the `LockManagerClass` property in [OJB.properties](#) file. By default an [in memory lock manager](#) is enabled.

```
LockManagerClass=org.apache.ojb.broker.locking.LockManagerInMemoryImpl
...
```

## The LockManager Implementations

Below all [LockManager](#) implementations shipped with OJB are listed.

The *LockManager* implementation can optionally support

- lock timeout: The locked objects of an *owner* will be released after a specified time
- block timeout: The maximal time to wait for acquire a lock (e.g. when an object was locked by another thread). Implementations which do not support this feature are called *non-blocking*

### LockManagerInMemoryImpl

A *non-blocking*, single JVM, in-memory *LockManager* implementation. All `LockManager.upgradeLock(...)` calls are delegated to write locks. It's a simple and fast implementation.

The timeout of locks is supported. The block timeout is ignored, because it's non-blocking.

### LockManagerCommonsImpl

This implementation use the locking part of apache's [commons-transaction](#) api. The timeout of locks is currently (OJB 1.0.2) not supported, maybe in further versions. This implementation supports *blocking* (with deadlock detection) and *non-blocking* of acquired locks.

### LockManagerRemoteImpl

Supports locking in distributed environments based on a servlet. The *LockManagerRemoteImpl* class delegates all locking calls to a remote servlet (`LockManagerServlet`). The URL to contact the servlet have to be set in [OJB.properties](#) file using the *LockServletUrl* property, e.g.

```
LockServletUrl=http://127.0.0.1:8080/ojb-lockserver
```

To make deployment of the `LockManagerServlet` on a servlet container easier an Ant target *lockserver-war* exist, which will build an example *.war* file containing all needed files (maybe some useless files) for deployment.

The generated `web.xml` file look like:

```
<!DOCTYPE web-app
  PUBLIC "-//Sun Microsystems, Inc.//DTD Web Application 2.3//EN"
  "http://java.sun.com/dtd/web-app_2_3.dtd">
<web-app>

  <display-name>OJB ODMG Lock Server</display-name>
  <description>
    OJB ODMG Lock Server
  </description>

  <servlet>
    <servlet-name>lockserver</servlet-name>
  <servlet-class>org.apache.ojb.broker.locking.LockManagerServlet</servlet-class>
  <init-param>
    <param-name>lockManager</param-name>
  <param-value>org.apache.ojb.broker.locking.LockManagerInMemoryImpl</param-value>
  </init-param>
  <init-param>
    <param-name>lockTimeout</param-name>
    <param-value>80000</param-value>
  </init-param>
  <init-param>
    <param-name>blockTimeout</param-name>
    <param-value>1000</param-value>
  </init-param>
```

```

        <!--load-on-startup>1</load-on-startup-->
    </servlet>

    <!-- The mapping for the webdav servlet -->
    <servlet-mapping>
        <servlet-name>lockserver</servlet-name>
        <url-pattern>/</url-pattern>
    </servlet-mapping>

    <!-- Establish the default list of welcome files -->
    <welcome-file-list>
        <welcome-file>index.jsp</welcome-file>
        <welcome-file>index.html</welcome-file>
        <welcome-file>index.htm</welcome-file>
    </welcome-file-list>
</web-app>

```

It's possible to use each *LockManager* implementation as backend of the lock manager servlet - only adapt the *lockManager* init-param entry in the web.xml file.

#### 5.5.17.4. ODMG-api Locking

The OJB ODMG implementation provides object level transactions as specified by the ODMG. This includes features like registering objects to transactions, persistence by reachability (a toplevel object is registered to a transaction, and also all its associated objects become registered implicitly) and as a very important aspect: object level locking.

The ODMG locking implementation is located in `org.apache.ojb.odmg.locking` and base on the OJB kernel locking code in `org.apache.ojb.broker.locking`. The odmg implementation use it's own internal locking interface `org.apache.ojb.odmg.locking.LockManager` with specific methods to handle transactions as owner of a lock and persistent object [Identity objects](#) as resources to lock..

#### What it does

The ODMG-Api allows transactions to lock an object obj as follows:

```
org.odmg.Transaction.lock(Object obj, int lockMode)
```

where lockMode defines the locking modes:

```
org.odmg.Transaction.READ
org.odmg.Transaction.UPGRADE
org.odmg.Transaction.WRITE
```

A sample session could look as follows:

```

// get odmg facade instance
Implementation odmg = ...

//get open database
Database db = ...

// start a transaction
Transaction tx = odmg.newTransaction();
tx.begin();

MyClass myObject = ... ;

// lock object for write access
tx.lock(myObject, Transaction.WRITE);

// now perform write access on myObject ...

// finally commit transaction to make changes to myObject persistent

```



```
tx.commit();
```

The ODMG specification does not say if locks must be acquired explicitly by client applications or may be acquired implicitly. OJB provides implicit locking for the application programmers convenience: On commit of a transaction all read-locked objects are checked for modifications. If a modification is detected, a write lock is acquired for the respective object. If automatic acquisition of read- or write-lock fails, the transaction is aborted.

On locking an object to a transaction, OJB automatically locks all associated objects (as part of the *persistence by reachability feature*) with the same locking level. If application use large object nets which are shared among several transactions acquisition of write-locks may be very difficult. Thus OJB can be configured to acquire only read-locks for associated objects.

You can change this behaviour by modifying the file [OJB.properties](#) and changing the entry `LockAssociations=WRITE` to `LockAssociations=READ`.

The ODMG specification does not prescribe transaction isolation levels or locking strategies to be used. Thus there are no API calls for setting isolation levels. OJB provides [four different isolation levels](#) that can be configured [global or for each persistent class](#) in the configuration files.

### 5.5.17.5. Locking in distributed environment

In distributed or clustered environments the object level locking ([pessimistic locking](#)) have to be consistent over several JVM. The [optimistic locking](#) works in clustered/distributed environments without any modifications.

Currently OJB was shipped was simple servlet based [LockManager](#) implementation called [LockManagerRemoteImpl](#).

Here is a description how to use it:

1. Change `LockManagerClass` entry in [OJB.properties](#) file to the remote implementation: `org.apache.ojb.broker.locking.LockManagerRemoteImpl` and the `LockServletUrl` to the servlet engine where the lock-server servlet will be deployed:

```
LockManagerClass=org.apache.ojb.broker.locking.LockManagerRemoteImpl
...
LockServletUrl=http://127.0.0.1:8080/ojb-lockserver
```

2. Run the ant `lockserver-war` target to generate the lock-server servlet `.war` application file. The generated file will be found in `[db-ojb]/dist`.
3. Check that all needed libraries be copied in `lockserver-war` file.

This implementation has some drawbacks, e.g. it uses one servlet node to deploy the `LockMap` servlet.

A much better solution will be a JMS- or JavaGroups-based [LockManager](#) implementation (hope we can start working on such a implementation some day).

### 5.5.17.6. Pluggin own locking classes

OJB was shipped with several locking classes implementations.

This may not be viable in some environments. Thus OJB allows to plug in user defined [LockManager](#) implementations.

To specify specific implementations change the following entry in the [OJB.properties](#) configuration file:

```
LockManagerClass=my.ojb.LockManagerMyImpl
```

**Note:**

Of course we are interested in your solutions! If you have implemented something interesting, just contact us.

## 5.5.18. XDoclet OJB module documentation

### 5.5.18.1. Acquiring and building

The XDoclet OJB module is part of OJB source. As such, the source of the module is part of the OJB source tree and can be found in directory `src/xdoclet`. Likewise, binary versions of the module and the required libraries (`xjavadoc`, `xdoclet`) are to be found in the `lib` folder.

In order to build the XDoclet OJB module from source, you'll need a source distribution of XDoclet version 1.2, either a source distribution from the sourceforge download site or a CVS checkout/drop. See the XDoclet website at <http://xdoclet.sourceforge.net/install.html> for details.

#### Building with a XDoclet source distribution

Unpack the source distribution of XDoclet which is contained in a file `xdoclet-src-<version>.<archive-format>` somewhere. If you unpacked it side-by-side of OJB, you'll get a directory layout similar to:

```
\xdoclet-1.2
  \config
  \core
  \lib
  ...
\db-ojb
  \bin
  \contrib
  ...
```

The XDoclet OJB module is then build using the `build-xdoclet-module.xml` ant script:

```
ant -Dxdoclet.src.dir=../xdoclet-1.2 -f build-xdoclet-module.xml
```

The build process will take some time, and after successful compilation the three jars `xjavadoc-<version>.jar`, `xdoclet-<version>.jar`, and `xdoclet-ojb-module-<version>.jar` are copied to the library directory of OJB.

#### Building with a XDoclet CVS checkout

When checking out from CVS (the `xdoclet-all` target), you'll get a directory like:

```
\xdoclet-all
  \xdoclet
    \config
    \core
    ...
  \xdocletgui
  \xjavadoc
\db-ojb
  \bin
  \contrib
  ...
```

Building the XDoclet OJB module is performed by calling:

```
ant -Dxdoclet.src.dir=../xdoclet-all/xdoclet -f build-xdoclet-module.xml
```

Since this is the default structure assumed by the build script, this can be shortened to:

```
ant -f build-xdoclet-module.xml
```

### Other build options

The build script for the XDoclet OJB module uses the OJB build properties so the following line added to the `build.properties` file in the OJB root directory allows to omit the `-Dxdoclet.src.dir=<xdoclet src dir>` commandline option:

```
xdoclet.src.dir=<xdoclet src dir>
```

### 5.5.18.2. Usage

Using the XDoclet OJB module is rather easy. Put the module jar along with the `xdoclet` and `xjavadoc` jars in a place where `ant` will find it, and then invoke it in your build file like:

```
<target name="repository-files">
  <taskdef name="ojbdoclet"
    classname="xdoclet.modules.ojb.OjbDocletTask"
    classpathref="build-classpath">
  <ojbdoclet destdir="./build">
    <fileset dir="./src"/>
    <ojbrepository destinationFile="repository_user.xml"/>
    <torqueschema databaseName="test" destinationFile="project-schema.xml"/>
  </ojbdoclet>
</target>
```

The XDoclet OJB module has two sub tasks, `ojbrepository` and `torqueschema`, which generate the OJB repository part containing the user descriptors and the torque table schema, respectively. Please note that the XDoclet OJB module (like all `xdoclet` tasks) expects the directory structure of its input java source files to match their package structure. In this regard it is similar to the `javac` `ant` task.

Due to a bug in XDoclet, you should not call the `ojbdoclet` task more than once in the same `taskdef` scope. So, each `ojbdoclet` call should be in its own target with a leading `taskdef`.

The main `ojbdoclet` task has two attributes:

#### **destdir**

The destination directory where generated files will be placed.

#### **checks : none | basic | strict (default)**

The amount of the checks performed. Per default, `strict` checks are performed which means that for instance classes specified in an attribute (e.g. `collection-class`, `row-reader` etc.) are loaded from the classpath and checked. So in this mode it is necessary to have OJB as well as the processed classes on the classpath (using the `classpathref` attribute of the `taskdef` `ant` task above). If this is for some reason not possible, then use `basic` which performs most of the checks but does not load classes from the classpath. `none` does not perform any checks so use it with care and only if really necessary (in this case it would be helpful if you would post the problem to the `ojb-user` mailing list).

The `ojbrepository` subtask has the following attributes:

#### **destinationFile**

Specifies the output file. The default is `repository_user.xml`.

#### **verbose : true | false (default)**

Whether the task should output some information about its progress.

The `torqueschema` subtask has these attributes:

#### **databaseName**

This attribute gives the name of the database for torque (required).

#### **destinationFile**

The output file, default is `project-schema.xml`.

#### **dtdUrl**

Allows to specify the url of the torque dtd. This is necessary e.g. for XML parsers that have problems with the default dtd url (<http://jakarta.apache.org/turbine/dtd/database.dtd>), or when using a newer version of torque.

#### **generateForeignkeys : true (default) | false**

Whether foreignkey tags are generated in the torque database schema.

#### **verbose : true | false (default)**

Whether the task outputs some progress information.

The `classpathref` attribute in the taskdef can be used to define the classpath for xdoclet (containing the xdoclet and ojb module jars), e.g. via:

```
<path id="build-classpath">
  <fileset dir="lib">
    <include name="**/*.jar"/>
  </fileset>
</path>
```

Using the generated torque schema is a bit more tricky. The easiest way is to use the `build-torque.xml` script which is part of OJB. Include the `lib` subdirectory of the OJB distribution which also includes torque (e.g. in `build-classpath` as shown above). You will also want to use your OJB settings (if you're using the [ojb-blank](#) project, then only `build.properties`), so include them at the beginning of the build script if they are not already there:

```
<property file="build.properties"/>
```

Now you can create the database with ant calls similar to these:

```
<target name="init-db" depends="repository-files">
  <!-- Torque's build file -->
  <property name="torque.buildFile"
    value="build-torque.xml"/>

  <!-- The name of the database which we're taking from the profile -->
  <property name="torque.project"
    value="${databaseName}"/>

  <!-- Where the schemas (your project and, if required, ojb's internal
tables) are -->
  <property name="torque.schema.dir"
    value="src/schema"/>

  <!-- Build directory of Torque -->
  <property name="torque.output.dir"
    value="build"/>

  <!-- Torque will put the generated sql here -->
  <property name="torque.sql.dir"
    value="${torque.output.dir}"/>

  <!-- Torque shall use the classpath (to find the jdbc driver etc.) -->
  <property name="torque.useClasspath"
    value="true"/>

  <!-- Which jdbc driver to use (again from the profile) -->
  <property name="torque.database.driver"
    value="${jdbcRuntimeDriver}"/>

  <!-- The url used to build the database; note that this may be different
from the url to access the database (e.g. for MySQL) -->
  <property name="torque.database.buildUrl"
    value="${urlProtocol}:${urlSubprotocol}:${urlDbalias}"/>
```

```

<!-- Now we're generating the database sql -->
<ant dir="."
      antfile="${torque.buildFile}"
      target="sql">
</ant>
<!-- Next we create the database -->
<ant dir="."
      antfile="${torque.buildFile}"
      target="create-db">
</ant>
<!-- And the tables -->
<ant dir="."
      antfile="${torque.buildFile}"
      target="insert-sql">
</ant>
</target>

```

As you can see, the major problem of using Torque is to correctly setup Torque's build properties.

One important thing to note here is that the latter two calls modify the database and in the process remove any existing data, so use them with care. Similar to the above targets, you can use the additional targets `datadump` for storing the data currently in the database in an XML file, and `datasql` for inserting the data from an XML file into the database.

Also, these steps are only valid for the torque that is delivered with OJB, but probably not for newer or older versions.

### 5.5.18.3. Tag reference

#### Interfaces and Classes

[ojb.class](#)

[ojb.extent-class](#)

[ojb.modify-inherited](#)

[ojb.object-cache](#)

[ojb.index](#)

[ojb.delete-procedure](#)

[ojb.insert-procedure](#)

[ojb.update-procedure](#)

[ojb.constant-argument](#)

[ojb.runtime-argument](#)

#### Fields and Bean properties

[ojb.field](#)

#### References

[ojb.reference](#)

#### Collections

[ojb.collection](#)

#### Nested objects

[ojb.nested](#)

[ojb.modify-nested](#)

### 5.5.18.4. Interfaces and Classes

#### **ojb.class**

The **ojb.class** tag marks interfaces and classes that shall be present in the repository descriptor. This includes types that are used as reference targets or as collection elements, but for instance not abstract base classes not used elsewhere.

*Attributes:***attributes**

Optionally contains attributes of the class as a comma-separated list of name-value pairs.

**determine-extents : true (default) | false**

When set to `true`, then the XDoclet OJB module will automatically determine all extents (ojb-relevant sub types) of this type. If set to `false`, then extents need to be specified via the [ojb.extent-class](#) class tag (see below).

**documentation**

Optionally contains documentation on the class. If no `table-documentation` attribute is specified, then the value is also used for the table documentation in the database schema.

**generate-repository-info : true (default) | false**

Setting this to `false` prevents the generation of field/reference/collection descriptors in the repository XML file, and also automatically enforces `generate-table-info = false`.

Note that there is one case where the XDoclet module will still generate field descriptors. If the type is referenced by a reference or collection, then the corresponding foreign key fields (if 1:n collection) or primary keys (if reference or m:n collection) will be automatically included in the class descriptor, even if they are only defined in subtypes.

**generate-table-info : true (default) | false**

This attribute controls whether the type has an associated table. If set to `true`, a torque table descriptor will be created in the database schema. Otherwise, no table will be in the database schema for this type.

**include-inherited : true (default) | false**

Determines whether base type fields/references/collections with the appropriate tags ([ojb.field](#), [ojb.reference](#), [ojb.collection](#)) will be included in the descriptor and table definition of this class. Note that all base type fields/references/collections with an appropriate tag are included regardless of whether the base types have the **ojb.class** tag or not.

**table**

The name of the table used for this type. Is only used when table info is generated. If not specified, then the short name of the type is used.

**table-documentation**

Optionally contains documentation for the table in the database schema.

The following `class-descriptor` attributes are also supported in the **ojb.class** tag and will be written directly to the generated class descriptor (see the [repository.dtd](#) for their meaning):

- **accept-locks**
- **factory-class**
- **factory-method**
- **initialization-method**
- **isolation-level**
- **proxy**
- **proxy-prefetching-limit**
- **refresh**
- **row-reader**

*Example:* (from the unit tests)

```
/**
```

```

* @ojb.class generate-table-info="false"
*/
public abstract class AbstractArticle implements InterfaceArticle,
java.io.Serializable
...

/**
 * @ojb.class table="ARTICLE"
 *             proxy="dynamic"
 *             include-inherited="true"
 *             documentation="This is important documentation on the Article
class."
 *             table-documentation="And this is important documentation on the
ARTICLE table."
 *             attributes="color=blue,size=big"
 */
public class Article extends AbstractArticle implements InterfaceArticle,
java.io.Serializable
...

```

The `AbstractArticle` class will have an class descriptor in the repository file, but no field, reference or collection descriptors. The `Article` class however will not only have descriptors for its own fields/references/collections but also for those inherited from `AbstractArticle`. Also, its table definition in the torque file will be called "Artikel", not "Article". The resulting class descriptors look like:

```

<class-descriptor
  class="org.apache.ojb.broker.AbstractArticle"
>
  <extent-class class-ref="org.apache.ojb.broker.Article"/>
</class-descriptor>

<class-descriptor
  class="org.apache.ojb.broker.Article"
  proxy="dynamic"
  table="ARTICLE"
>
  <documentation>This is important documentation on the Article
class.</documentation>
  ...
  <attribute attribute-name="color" attribute-value="blue"/>
  <attribute attribute-name="size" attribute-value="big"/>
</class-descriptor>
...

```

### ojb.extent-class

Use the **ojb.extent-class** to explicitly specify extents (direct persistent sub types) of the current type. The **class-ref** attribute contains the fully qualified name of the class. However, these tags are only evaluated if the **determine-extents** attribute of the [ojb.class](#) tag is set to false.

*Attributes:*

**class-ref**

The fully qualified name of the sub-class (required).

*Example:*

```

/**
 * @ojb.class determine-extents="false"
 *             generate-table-info="false"
 * @ojb.extent-class class-ref="org.apache.ojb.broker.CdArticle"
 */
public abstract class AbstractCdArticle extends Article implements
java.io.Serializable
...

```

which results in:

```
<class-descriptor
  class="org.apache.ojb.broker.AbstractCdArticle"
>
  <extent-class class-ref="org.apache.ojb.broker.CdArticle"/>
</class-descriptor>
```

### ojb.modify-inherited

Allows to modify attributes of inherited fields/references/collections (normally, all attributes are used without modifications) for this and all sub types. One special case is the specification of an empty value which leads to a reset of the attribute value. As a result the default value is used for this attribute.

*Attributes:* All of [ojb.field](#), [ojb.reference](#), and [ojb.collection](#) (with the exception of the attributes related to indirection tables ( **indirection-table**, **remote-foreignkey**, **indirection-table-primarykeys**, **indirection-table-documentation**, **foreignkey-documentation**, **remote-foreignkey-documentation**), and also:

**ignore : true | false (default)**

Specifies that this feature will be ignored in this type (but only in the current type, not in subtypes).

**name**

The name of the field/reference/collection to modify (required).

*Example:*

```
/**
 * @ojb.class table="Artikel"
 * @ojb.modify-inherited name="productGroup"
 *                       proxy="true"
 *                       auto-update="object"
 */
public class ArticleWithReferenceProxy extends Article
```

produces the class descriptor

```
<class-descriptor
  class="org.apache.ojb.broker.ArticleWithReferenceProxy"
  table="Artikel"
>
  ...
  <reference-descriptor
    name="productGroup"
    class-ref="org.apache.ojb.broker.ProductGroup"
    proxy="true"
    auto-update="object"
  >
    <documentation>this is the reference to an articles
productgroup</documentation>
    <attribute attribute-name="color" attribute-value="red"/>
    <attribute attribute-name="size" attribute-value="tiny"/>
    <foreignkey field-ref="productGroupId"/>
  </reference-descriptor>
</class-descriptor>
```

### ojb.object-cache

The **ojb.object-cache** tag allows to specify the ObjectCache implementation that OJB uses for objects of this class (instead of the one defined in the jdbc connection descriptor or in the `ojb.properties` file). Classes specified with this tag have to implement the `org.apache.ojb.broker.cache.ObjectCache` interface. Note that object cache



specifications are not inherited.

*Attributes:*

**attributes**

Optionally contains attributes of the object cache as a comma-separated list of name-value pairs.

**class**

The fully qualified name of the object cache class (required).

**documentation**

Optionally contains documentation on the object cache specification.

*Example:*

```
/**
 * @ojb.class
 * @ojb.object-cache
 * class="org.apache.ojb.broker.cache.ObjectCachePerBrokerImpl"
 * documentation="Some important documentation"
 */
public class SomeClass implements Serializable
{
    ...
}
```

and the class descriptor

```
<class-descriptor
  class="SomeClass"
  table="SomeClass"
>
  <object-cache class="org.apache.ojb.broker.cache.ObjectCachePerBrokerImpl">
    <documentation>Some important documentation</documentation>
  </object-cache>
  ...
</class-descriptor>
```

## ojb.index

The **ojb.index** tag is used to define possibly unique indices for the class. An index consists of at least one field of the class (either locally defined or inherited, anonymous or explicit). There is an default index (without a name) that is made up by all fields that have the **indexed** attribute set to `true`. All other indices have to be defined via the **ojb.index** tag. In contrast to the **indexed** attribute, indices defined via the **ojb.index** tag are not inherited.

*Attributes:*

**documentation**

Optionally contains documentation on the index.

**fields**

The fields that make up the index separated by commas (required).

**name**

The name of the index (required). If there are multiple indices with the same name, then only the first one is used (all others are ignored).

**unique : true | false (default)**

Whether the index is unique or not.

*Example:*

```
/**
 * @ojb.class table="SITE"
 * @ojb.index name="NAME_UNIQUE"
```

```

*         unique="true"
*         fields="name"
*/
public class Site implements Serializable
{
    /**
     * @ojb.field indexed="true"
     */
    private Integer nr;
    /**
     * @ojb.field column="NAME"
     *         length="100"
     */
    private String name;
    ...
}

```

the class descriptor

```

<class-descriptor
  class="org.apache.ojb.odmg.shared.Site"
  table="SITE"
>
  <field-descriptor
    name="nr"
    column="nr"
    jdbc-type="INTEGER"
    indexed="true"
  >
</field-descriptor>
  <field-descriptor
    name="name"
    column="NAME"
    jdbc-type="VARCHAR"
    length="100"
  >
</field-descriptor>
  ...
  <index-descriptor
    name="NAME_UNIQUE"
    unique="true"
  >
    <index-column name="NAME" />
  </index-descriptor>
</class-descriptor>

```

and the torque table schema

```

<table name="SITE">
  <column name="nr"
    javaName="nr"
    type="INTEGER"
  />
  <column name="NAME"
    javaName="name"
    type="VARCHAR"
    size="100"
  />
  ...
  <index>
    <index-column name="nr" />
  </index>
  <unique name="NAME_UNIQUE">
    <unique-column name="NAME" />
  </unique>
</table>

```

### ojb.delete-procedure

Declares a database procedure that is used for deleting persistent objects.

*Attributes:***arguments**

A comma-separated list of the names of [constant](#) or [runtime](#) arguments specified in the same class.

**attributes**

Optionally contains attributes of the procedure as a comma-separated list of name-value pairs.

**documentation**

Optionally contains documentation on the procedure.

**include-pk-only : true | false (default)**

Whether all fields of the class that make up the primary key, shall be passed to the procedure. If set to `true` then the **arguments** value is ignored.

**name**

The name of the procedure (required).

**return-field-ref**

Identifies a field of the class that will receive the return value of the procedure. Use only if the procedure has a return value.

*Example:*

```
/**
 * @ojb.class
 * @ojb.delete-procedure name="DELETE_PROC"
 *                       arguments="arg1,arg2"
 *                       return-field-ref="attr2"
 *                       documentation="Some important documentation"
 * @ojb.constant-argument name="arg1"
 *                       value="0"
 * @ojb.runtime-argument name="arg2"
 *                       field-ref="attr1"
 */
public class SomeClass
{
    /** @ojb.field */
    private Integer attr1;
    /** @ojb.field */
    private String attr2;
    ...
}
```

leads to the class descriptor

```
<class-descriptor
  class="SomeClass"
  table="SomeClass"
>
  <field-descriptor
    name="attr1"
    column="attr1"
    jdbc-type="INTEGER"
  >
</field-descriptor>
  <field-descriptor
    name="attr2"
    column="attr2"
    jdbc-type="VARCHAR"
    length="254"
  >
</field-descriptor>
  ...
  <delete-procedure
    name="DELETE_PROC"
    return-field-ref="attr2"
  >
```

```

>
  <documentation>Some important documentation</documentation>
  <constant-argument
    value="0"
  >
  </constant-argument>
  <runtime-argument
    field-ref="attr2"
  >
  </runtime-argument>
</delete-procedure>
</class-descriptor>

```

### obj.insert-procedure

Identifies the database procedure that shall be used for inserting objects into the database.

*Attributes:*

**arguments**

Comma-separated list of names of [constant](#) or [runtime](#) arguments that are specified in the same class.

**attributes**

Contains optional attributes of the procedure in a comma-separated list of name-value pairs.

**documentation**

Contains optional documentation on the procedure.

**include-all-fields : true | false (default)**

Specifies whether all persistent fields of the class shall be passed to the procedure. If so, then the **arguments** value is ignored.

**name**

The name of the procedure (required).

**return-field-ref**

The persistent field that receives the return value of the procedure (should only be used if the procedure returns a value).

For an example see [constant argument](#).

### obj.update-procedure

The database procedure that will be used for updating persistent objects in the database.

*Attributes:*

**arguments**

A comma-separated list of names of [constant](#) or [runtime](#) arguments in the same class.

**attributes**

The optional attributes of the procedure in a comma-separated list of name-value pairs.

**documentation**

Optional documentation on the procedure.

**include-all-fields : true | false (default)**

Whether all persistent fields of the class shall be passed to the procedure in which case the **arguments** value is ignored.

**name**

Name of the database procedure (required).

**return-field-ref**

A persistent field that will receive the return value of the procedure (only to be used if the procedure returns a value).

For an example see [runtime argument](#).

### ojb.constant-argument

A constant argument for a database procedure. These arguments are referenced by the procedure tags in the **arguments** attribute via their names.

*Attributes:*

**attributes**

Optionally contains attributes of the argument.

**documentation**

Optionally contains documentation on the argument.

**value**

The constant value.

**name**

The identifier of the argument to be used the **arguments** attribute of a procedure tag (required).

*Example:*

```
/**
 * @ojb.class
 * @ojb.insert-procedure name="INSERT_PROC"
 *                       arguments="arg"
 * @ojb.constant-argument name="arg"
 *                       value="Some value"
 *                       attributes="name=value"
 */
public class SomeClass
{
    ...
}
```

will result in the class descriptor

```
<class-descriptor
  class="SomeClass"
  table="SomeClass"
>
  ...
  <insert-procedure
    name="INSERT_PROC"
  >
    <constant-argument
      value="Some value"
    >
      <attribute attribute-name="name" attribute-value="value"/>
    </constant-argument>
  </insert-procedure>
</class-descriptor>
```

### ojb.runtime-argument

An argument for a database procedure that is backed by a persistent field. Similar to constant arguments the name is important for referencing by the procedure tags in the **arguments** attribute.

*Attributes:*

**attributes**

Contains optionally attributes of the argument.

**documentation**

Optionally contains documentation on the argument.

**field-ref**

The persistent field that delivers the value. If unspecified, then in the procedure call `null` will be used.

**name**

Identifier of the argument for using it in the **arguments** attribute of a procedure tag (required).

**return**

If the field receives a value (?).

*Example:*

```
/**
 * @ojb.class
 * @ojb.update-procedure name="UPDATE_PROC"
 *                       arguments="arg"
 * @ojb.runtime-argument name="arg"
 *                       field-ref="attr"
 *                       documentation="Some documentation"
 */
public class SomeClass
{
    /** @ojb.field */
    private Integer attr;
    ...
}
```

will result in the class descriptor

```
<class-descriptor
  class="SomeClass"
  table="SomeClass"
>
  <field-descriptor
    name="attr"
    column="attr"
    jdbc-type="INTEGER"
  >
</field-descriptor>
  ...
  <update-procedure
    name="UPDATE_PROC"
  >
    <runtime-argument
      value="attr"
    >
      <documentation>Some documentation</documentation>
    </runtime-argument>
  </update-procedure>
</class-descriptor>
```

### 5.5.18.5. Fields and Bean properties

#### **ojb.field**

Fields or accessor methods (i.e. `get/is` and `set` methods) for properties are marked with the **ojb.field** tag to denote a persistent field. When a method is marked, then the corresponding bean property is used for naming purposes (e.g. "value" for a method `getValue()`). The XDoclet OJB module ensures that a field is not present more than once, therefore it is safe to mark both fields and their accessors. However, in that case these **ojb.field** tags are required to have the same attributes.

Due to a bug in XDoclet, it is currently not possible to process `final` or `transient` fields.

Marked fields are used for descriptor generation in the same type (if it has an [ojb.class](#) tag) and all sub types with the [ojb.class](#) tag having the **include-inherited** attribute set to `true`.

It is also possible to use the **obj.field** tag at the class level (i.e. in the JavaDoc comment of the class). In this case, the tag is used to define an *anonymous* field, e.g. a "field" that has no counterpart in the class but exists in the database. For anonymous fields, both the **name** and the **jdbc-type** attributes are required, and the **access** attribute is ignored (it defaults to the value *anonymous*). Beside these differences, anonymous fields are handled like other fields, e.g. they result in field-descriptor entries in the repository descriptor, and in columns in the table schema, and they are inherited and can be modified via the [obj.modify-inherited](#) tag.

The XDoclet OJB module orders the fields in the repository descriptor and table schema according to the following rules:

1. Fields (anonymous and non-anonymous) from base types/nested objects and from the current file that have an id, sorted by the id value. If fields have the same id, then they are sorted following the rules for fields without an id.
2. Fields (anonymous and non-anonymous) from base types/nested objects and from the current file that have no id, in the order they appear starting with the farthest base type. Per class, the anonymous fields come first, followed by the non-anonymous fields.

*Attributes:*

**access : readonly | readwrite (default)**

Specifies the accessibility of the field. *readonly* marks fields that are not to be modified. *readwrite* marks fields that may be read and written to. Anonymous fields do not have to be marked (i.e. *anonymous* value) as the position of the **obj.field** tag in the class JavaDoc comment suffices.

**attributes**

Optionally contains attributes of the field as a comma-separated list of name-value pairs.

**autoincrement : none (default) | obj | database**

Defines whether this field gets its value automatically. If *obj* is specified, then OJB fills the value using a sequence manager. If the value is *database*, then the column is also defined as *autoIncrement* in the torque schema (i.e. the database fills the field), and in the repository descriptor, the field is marked as *access='readonly'* (if it isn't an anonymous field). The *database* value is intended to be used with the `org.apache.obj.broker.util.sequence.SequenceManagerNativeImpl` sequence manager. For details, see the [Sequence Manager documentation](#). The default value is *none* which means that the field is not automatically filled.

**column**

The name of the database column for this field. If not given, then the name of the attribute is used.

**column-documentation**

Optionally contains documentation on the column in the database schema.

**conversion**

The name of the class to be used for conversion between the java type of the field (e.g. `java.lang.Boolean` or `java.util.Date`) and the java type for the JDBC type (e.g. `java.lang.Integer` or `java.sql.Date`). Conversion classes must implement the

`org.apache.obj.broker.accesslayer.conversions.FieldConversion` interface. If no explicit JDBC type is defined and the java type has no defined conversion (see below), then per default the

`org.apache.obj.broker.accesslayer.conversions.Object2ByteArrayFieldConversion` conversion class is used.

Default conversion is also used for the following java types when no jdbc type is given (default type is used instead), and no conversion is specified:

Java type	Default conversion
<code>org.apache.objbroker.util.GUID</code>	<code>org.apache.objbroker.accesslayer.conversions.GUID2</code>

### documentation

Optionally contains documentation on the field. If no `column-documentation` attribute value is specified, then this value is also used for the documentation of the column in the database schema.

### id

An integer specifying the position in the repository descriptor and table schema. For the placement rules see [above](#).

**jdbc-type** : BIT | TINYINT | SMALLINT | INTEGER | BIGINT | DOUBLE | FLOAT | REAL | NUMERIC | DECIMAL | CHAR | VARCHAR | LONGVARCHAR | DATE | TIME | TIMESTAMP | BINARY | VARBINARY | LONGVARBINARY | CLOB | BLOB | STRUCT | ARRAY | REF | BOOLEAN | DATALINK

The JDBC type for the column. The XDoclet OJB module will automatically determine a jdbc type for the field if none is specified. This means that for anonymous fields, the **jdbc-type** attribute is required. The automatic mapping performed by the XDoclet OJB module from java type to jdbc type is as follows:

Java type	JDBC type
<code>boolean</code>	BIT
<code>byte</code>	TINYINT
<code>short</code>	SMALLINT
<code>int</code>	INTEGER
<code>long</code>	BIGINT
<code>char</code>	CHAR
<code>float</code>	REAL
<code>double</code>	FLOAT
<code>java.lang.Boolean</code>	BIT
<code>java.lang.Byte</code>	TINYINT
<code>java.lang.Short</code>	SMALLINT
<code>java.lang.Integer</code>	INTEGER
<code>java.lang.Long</code>	BIGINT
<code>java.lang.Character</code>	CHAR
<code>java.lang.Float</code>	REAL
<code>java.lang.Double</code>	FLOAT
<code>java.lang.String</code>	VARCHAR
<code>java.util.Date</code>	DATE
<code>java.sql.Date</code>	DATE
<code>java.sql.Time</code>	TIME
<code>java.sql.Timestamp</code>	TIMESTAMP
<code>java.sql.Blob</code>	BLOB



java.sql.Clob	CLOB
java.sql.Ref	REF
java.sql.Struct	STRUCT
java.math.BigDecimal	DECIMAL
org.apache.ojb.broker.util.GUID	VARCHAR

For any other type (including array types) the default mapping is to LONGVARBINARY using the Object2ByteArrayFieldConversion conversion (see **conversion** attribute above).

### length

The length of the column which might be required by the jdbc type in some databases. This is the reason that for some jdbc types, the XDoclet OJB module imposes default lengths if no length is specified:

Jdbc type	Default length
CHAR	1
VARCHAR	254

### name

The name of the field. This attribute is required for anonymous fields, otherwise it is ignored.

### precision

### scale

The precision and scale of the column if required by the jdbc type. They are usually used in combination with the DECIMAL and NUMERIC types, and then specify the number of digits before ( **precision**) and after ( **scale**) the comma (excluding the plus/minus sign). Due to restrictions in some databases (e.g. MySQL), the XDoclet OJB module imposes default values for some types if none are specified:

Jdbc type	Default values for precision, scale
DECIMAL	20,0 (this corresponds to the range of long where the longest number is -9223372036854775808).
NUMERIC	20,0

For other types, if only the precision is specified, the scale defaults to 0. If only scale is specified, precision defaults to 1.

Other attributes supported in the **obj.field** tag that have the same meaning as in the [repository descriptor](#) (and partly in the torque table schema) are:

- **default-fetch**
- **indexed**
- **locking**
- **nullable**
- **primarykey**
- **sequence-name**
- **update-lock**

*Examples:*

```
/**
 * @obj.field column="Auslaufartikel"
 *          jdbc-type="INTEGER"
```

```

*
conversion="org.apache.ojb.broker.accesslayer.conversions.Boolean2IntFieldConversion"
*      column-documentation="Some documentation on the column"
*      id="10"
*      attributes="color=green,size=small"
*/
protected boolean isSelloutArticle;

```

will result in the following field descriptor:

```

<field-descriptor
  name="isSelloutArticle"
  column="Auslaufartikel"
  jdbc-type="INTEGER"
conversion="org.apache.ojb.broker.accesslayer.conversions.Boolean2IntFieldConversion"
>
  <attribute attribute-name="color" attribute-value="green"/>
  <attribute attribute-name="size" attribute-value="small"/>
</field-descriptor>

```

The column descriptor looks like:

```

<table name="Artikel">
  ...
  <column name="Auslaufartikel"
    javaName="isSelloutArticle"
    type="INTEGER"
    description="Some documentation on the column"
  />
  ...
</table>

```

An anonymous field is declared like this:

```

/**
 * @ojb.class table="TABLE_F"
 *      include-inherited="false"
 * @ojb.field name="eID"
 *      column="E_ID"
 *      jdbc-type="INTEGER"
 * @ojb.reference class-ref="org.apache.ojb.broker.E"
 *      auto-retrieve="true"
 *      auto-update="object"
 *      auto-delete="object"
 *      foreignkey="eID"
 */
public class F extends E implements Serializable
...

```

In this case an anonymous field is declared and also used as the foreignkey of an anonymous reference. The corresponding class descriptor looks like:

```

<class-descriptor
  class="org.apache.ojb.broker.F"
  table="TABLE_F"
>
  <field-descriptor
    name="eID"
    column="E_ID"
    jdbc-type="INTEGER"
    access="anonymous"
  >
</field-descriptor>
  ...
  <reference-descriptor
    name="super"
    class-ref="org.apache.ojb.broker.E"
    auto-retrieve="true"
    auto-update="object"
    auto-delete="object"
  >

```

```
>
  <foreignkey field-ref="eID"/>
</reference-descriptor>
</class-descriptor>
```

Here the anonymous field and reference (which implicitly refers to `super`) are used to establish the super-subtype relationship between E and F on the database level. For details on this see the [advanced technique section](#).

## 5.5.18.6. References

### **obj.reference**

Similar to fields, references (java fields or accessor methods) are marked with the **obj.reference** tag. We have a reference when the type of the java field is itself a persistent class (has an [obj.class](#) tag) and therefore the java field represents an association. This means that the referenced type of an association (or the one specified by the **class-ref** attribute, see below) is required to be present in the repository descriptor (it has the [obj.class](#) tag).

Foreign keys of references are also declared in the torque table schema (see example below).

OJB currently requires that the referenced type has at least one field used to implement the reference, usually some id of an integer type.

A reference can be stated in the JavaDoc comment of the class (anonymous reference), but in this case it silently refer to `super` (see the example of [obj.field](#)) which can be used to establish an inheritance relationship. Note that anonymous references are not inherited (in contrast to anonymous fields and normal references).

*Attributes:*

#### **attributes**

Optionally contains attributes of the reference as a comma-separated list of name-value pairs.

#### **class-ref**

Allows to explicitly specify the referenced type. Normally the XDoclet OJB module searches the type of the field and its sub types for the nearest type with the [obj.class](#) tag. If the type is specified explicitly, then this type is used instead. For anonymous references, the **class-ref** has to specified as there is no field to determine the type from.

Note that the type is required to have the [obj.class](#) tag.

#### **database-foreignkey : true (default) | false**

Specifies whether a database foreignkey shall be generated for the reference. Note that this attribute is only evaluated if the XDoclet module has determined that a database foreignkey could be generated. You cannot force the generation with this attribute, and the value of the attribute is not considered when checking if database foreignkeys can be generated in case the referencing class has subtypes (in which case database foreignkeys can only be generated if all subtypes map to the same table or don't map to a table or the inheritance is mapped via a super-reference).

#### **documentation**

Optionally contains documentation on the reference.

#### **foreignkey**

Contains one or more foreign key fields separated by commas (required). The foreign key fields are fields with the [obj.field](#) tag in the same class as the reference, which implement the association, i.e. contains the values of the primarykeys of the referenced object.

Other supported attributes (see [repository.dtd](#) for their meaning) written directly to the repository

descriptor file:

- **auto-delete**
- **auto-retrieve**
- **auto-update**
- **otm-dependent**
- **proxy**
- **proxy-prefetching-limit**
- **refresh**

*Example:*

```
public abstract class AbstractArticle implements InterfaceArticle,
java.io.Serializable
{
    protected InterfaceProductGroup productGroup;

    /**
     * @ojb.reference class-ref="org.apache.ojb.broker.ProductGroup"
     *                foreignkey="productGroupId"
     *                documentation="this is the reference to an articles
productgroup"
     *                attributes="color=red,size=tiny"
     */
    protected InterfaceProductGroup productGroup;
    /**
     * @ojb.field
     */
    protected int productGroupId;
    ...
}
```

Here the java type is InterfaceProductGroup although the repository reference uses the sub type ProductGroup. The generated reference descriptor looks like:

```
<field-descriptor
  name="productGroupId"
  column="Kategorie_Nr"
  jdbc-type="INTEGER"
>
</field-descriptor>
<reference-descriptor
  name="productGroup"
  class-ref="org.apache.ojb.broker.ProductGroup"
>
  <documentation>this is the reference to an articles
productgroup</documentation>
  <attribute attribute-name="color" attribute-value="red"/>
  <attribute attribute-name="size" attribute-value="tiny"/>
  <foreignkey field-ref="productGroupId"/>
</reference-descriptor>
```

In the torque table schema for the Article class, the foreign key for the product group is explicitly declared:

```
<table name="Artikel">
  ...
  <column name="Kategorie_Nr"
    javaName="productGroupId"
    type="INTEGER"
  />
  ...
  <foreign-key foreignTable="Kategorien">
    <reference local="Kategorie_Nr" foreign="Kategorie_Nr"/>
  </foreign-key>
</table>
```

For an example of an anonymous reference, see the examples of [ojb.field](#).

### 5.5.18.7. Collections

#### ojb.collection

Persistent collections which implement 1:n or m:n associations are denoted by the **ojb.collection** tag. If the collection is an array, then the XDoclet OJB module can determine the element type automatically (analogous to references). Otherwise the type must be specified using the **element-class-ref** attribute. m:n associations are also supported (collections on both sides) via the **indirection-table**, **foreignkey** and **remote-foreignkey** attributes.

*Attributes:*

#### **attributes**

Optionally contains attributes of the collection as a comma-separated list of name-value pairs.

#### **collection-class**

Specifies the class that implements the collection. This attribute is usually only required if the actual type of the collection object shall be different from the variable type, e.g. if an interface like `java.util.Collection` is used as the declared type. Collections that use `java.util.Collection`, `java.util.List` or `java.util.Set` can be handled by OJB as-is so specifying **collection-class** is not necessary. For the types that do not, the XDoclet OJB module checks whether the declared collection field type implements the `org.apache.ojb.broker.ManageableCollection` interface, and if so, generates the **collection-class** attribute automatically. Otherwise, you have to specify it.

#### **database-foreignkey : true (default) | false**

Specifies whether a database foreignkey shall be generated for the collection. Note that this attribute is only evaluated if the XDoclet module has determined that a database foreignkey could be generated. You cannot force the generation with this attribute, and the value of the attribute is not considered when checking if database foreignkeys can be generated in the case of subtypes of the element type or the type having the collection (if m:n collection). For 1:n collections, database foreignkeys can only be generated if all subtypes of the element type map to the same table or don't map to a table or the inheritance is mapped via a super-reference. For m:n collections, the same applies to the class owning the collection.

#### **documentation**

Optionally contains documentation on the collection.

#### **element-class-ref**

Allows to explicitly specify the type of the collection elements. Note that the type is required to have the [ojb.class](#) tag.

#### **foreignkey**

Contains one or more foreign key field or columns separated by commas (required). If the collection implements an 1:n association, then this attribute specifies the fields in the element type that implement the association on the element side, i.e. they refer to the primary keys of the class containing this collection. Note that these fields are required to have the [ojb.field](#) tag.

When the collection is one part of an m:n association (e.g. with an indirection table), this attribute specifies the columns in the indirection table that point to the type owning this collection. This attribute is required of both collections. If the other type has no explicit collection, use the **remote-foreignkey** attribute to specify the foreign keys for

the other type.

### **foreignkey-documentation**

Optionally contains documentation for the columns in the indirection table in the database schema that point to this class.

### **indirection-table**

Gives the name of the indirection table used for m:n associations. The XDoclet OJB module will create an appropriate torque table schema. The specified foreign keys are taken from the **foreignkey** attribute in this class and the corresponding collection in the element class, or if the element class has no collection, from the **remote-foreignkey** attribute of this collection. The XDoclet OJB module associates the foreignkeys (in the order they are stated in the **foreignkey/ remote-foreignkey** attributes) to the ordered primarykey fields (for the ordering rules see the [obj.field](#) tag), and use their jdbc type (and length setting if necessary) of these primary keys for the columns.

### **indirection-table-documentation**

Optionally contains documentation for the indirection table in the database schema.

### **indirection-table-primarykeys : true | false (default)**

Specifies that the columns in the indirection table that point to this type, are primary keys of the table. If the element type has no corresponding collection, then this setting is also applied to the columns pointing to the element type.

### **orderby**

Contains the fields used for sorting the collection and, optionally, the sorting order (either `ASC` or `DESC` for ascending or descending, respectively) as a comma-separated list of name-value pairs. For instance, `field1=DESC, field2, field3=ASC` specifies three fields after which to sort, the first one in descending order and the other two in ascending order (which is the default and can be omitted).

### **query-customizer**

Specifies a query customizer for the collection. The type is required to implement `org.apache.ojb.broker.accesslayer.QueryCustomizer`.

### **query-customizer-attributes**

Specifies attributes for the query customizer. This attribute is ignored if no query customizer is specified for this collection.

### **remote-foreignkey**

Contains one or more foreign key columns (separated by commas) in the indirection table pointing to the elements. Note that this field should only be used if the other type does not have a collection itself which the XDoclet OJB module can use to retrieve the foreign keys. This attribute is ignored if used with 1:n collections (no indirection table specified).

### **remote-foreignkey-documentation**

Optionally contains documentation for the columns in the indirection table in the database schema that point to the element type. This value can be used when the element type has no corresponding collection (i.e. `remote-foreignkey` is specified) or if the corresponding collection does not specify the `foreignkey-documentation` attribute.

The same attributes as for references are written directly to the repository descriptor file (see [repository.dtd](#)) :

- **auto-delete**
- **auto-retrieve**
- **auto-update**
- **otm-dependent**
- **proxy**
- **proxy-prefetching-limit**

- refresh

Examples:

```
/**
 * @ojb.collection element-class-ref="org.apache.ojb.broker.Article"
 *                 foreignkey="productGroupId"
 *                 auto-retrieve="true"
 *                 auto-update="link"
 *                 auto-delete="object"
 *                 orderby="productGroupId=DESC"
 *
 * query-customizer="org.apache.ojb.broker.accesslayer.QueryCustomizerDefaultImpl"
 *                 query-customizer-attributes="attr1=value1"
 */
private ArticleCollection allArticlesInGroup;
```

The corresponding collection descriptor is:

```
<collection-descriptor
  name="allArticlesInGroup"
  element-class-ref="org.apache.ojb.broker.Article"
  collection-class="org.apache.ojb.broker.ArticleCollection"
  auto-retrieve="true"
  auto-update="link"
  auto-delete="object"
>
  <orderby name="productGroupId" sort="DESC"/>
  <inverse-foreignkey field-ref="productGroupId"/>
  <query-customizer
class="org.apache.ojb.broker.accesslayer.QueryCustomizerDefaultImpl">
    <attribute attribute-name="attr1" attribute-value="value1"/>
  </query-customizer>
</collection-descriptor>
```

An m:n collection is defined using the **indirection-table** attribute:

```
/**
 * @ojb.class generate-table-info="false"
 */
public abstract class BaseContentImpl implements Content
{
  /**
   * @ojb.collection element-class-ref="org.apache.ojb.broker.Qualifier"
   *                 auto-retrieve="true"
   *                 auto-update="link"
   *                 auto-delete="none"
   *                 indirection-table="CONTENT_QUALIFIER"
   *                 foreignkey="CONTENT_ID"
   *                 remote-foreignkey="QUALIFIER_ID"
   */
  private List qualifiers;
  ...
}

/**
 * @ojb.class table="NEWS"
 */
public class News extends BaseContentImpl
{
  ...
}

/**
 * @ojb.class generate-table-info="false"
 */
public interface Qualifier extends Serializable
{
  ...
}
```

The `BaseContentImpl` has a m:n association to the `Qualifier` interface. For the `BaseContentImpl` class, this association is implemented via the `CONTENT_ID` column (specified by the **foreignkey**) in the indirection table `CONTENT_QUALIFIER`. Usually, both ends of an m:n association have a collection implementing the association, and for both ends the **foreignkey** specifies the indirection table column pointing to the class at this end. The `Qualifier` interface however does not contain a collection which could be used to determine the indirection table column that implements the association from its side. So, this column is also specified in the `BaseContentImpl` class using the **remote-foreignkey** attribute. The class descriptors are:

```
<class-descriptor
  class="org.apache.ojb.broker.BaseContentImpl"
>
  <extent-class class-ref="org.apache.ojb.broker.News" />
</class-descriptor>

<class-descriptor
  class="org.apache.ojb.broker.News"
  table="NEWS"
>
  ...
  <collection-descriptor
    name="qualifiers"
    element-class-ref="org.apache.ojb.broker.Qualifier"
    indirection-table="CONTENT_QUALIFIER"
    auto-retrieve="true"
    auto-update="link"
    auto-delete="none"
  >
    <fk-pointing-to-this-class column="CONTENT_ID" />
    <fk-pointing-to-element-class column="QUALIFIER_ID" />
  </collection-descriptor>
</class-descriptor>

<class-descriptor
  class="org.apache.ojb.broker.Qualifier"
>
  <extent-class class-ref="org.apache.ojb.broker.BaseQualifierImpl" />
</class-descriptor>
```

As can be seen, the collection definition is inherited in the `News` class and the two indirection table columns pointing to the ends of the m:n association are correctly specified.

### 5.5.18.8. Nested objects

#### `ojb.nested`

The features of a class can be included in another class by declaring a field of that type and using this tag. The XDoclet OJB module will then add every tagged feature (i.e. fields/bean properties with `ojb.field`, `ojb.reference` or `ojb.collection` tag, or even with **ojb.nested**) from the type of the field to the current class descriptor. It is not required that the field's type has the `ojb.class` tag, though.

All attributes of the features are copied (even **primarykey**) and modified if necessary (e.g. the **foreignkey** of a reference is adjusted accordingly). For changing an attribute use the `ojb.modify-nested` tag.

For an example of nesting, see the example of `ojb.modify-nested`.

#### `ojb.modify-nested`



Similar to [ojb.modify-inherited](#), this tag allows to modify attributes of a nested feature.

*Attributes:* All of [ojb.field](#), [ojb.reference](#), and [ojb.collection](#) with the exception of the attributes related to indirection tables ( **indirection-table**, **remote-foreignkey**, **indirection-table-primarykeys**, **indirection-table-documentation**, **foreignkey-documentation**, **remote-foreignkey-documentation**), and also:

**ignore : true | false (default)**

Specifies that this feature will not be nested.

**name**

The name of the field/reference/collection to modify (required). Use here the name of the feature in the nested type.

*Example:*

The two classes:

```
public class NestedObject implements java.io.Serializable
{
    /** @ojb.field primaryKey="true" */
    protected int id;

    /** @ojb.field */
    protected boolean hasValue;

    /** @ojb.field */
    protected int containerId;

    /**
     * @ojb.reference foreignkey="containerId"
     */
    protected ContainerObject container;

    ...
}

/** @ojb.class */
public class ContainerObject implements java.io.Serializable
{
    /**
     * @ojb.field primaryKey="true"
     *             autoincrement="ojb"
     *             id="1"
     */
    protected int id;

    /** @ojb.field id="2" */
    protected String name;

    /**
     * @ojb.nested
     * @ojb.modify-nested name="hasValue"
     *                   jdbc-type="INTEGER"
     *
     * conversion="org.apache.ojb.broker.accesslayer.conversions.Boolean2IntFieldConversion"
     *             id="3"
     * @ojb.modify-nested name="id"
     *                   primaryKey=" "
     */
    protected NestedObject nestedObj;

    ...
}
```

result in the one class descriptor

```
<class-descriptor
```

```

class="ContainerObject"
table="ContainerObject"
>
<field-descriptor
  name="id"
  column="id"
  jdbc-type="INTEGER"
  primaryKey="true"
  autoincrement="true"
/>
<field-descriptor
  name="name"
  column="name"
  jdbc-type="VARCHAR"
  length="24"
/>
<field-descriptor
  name="nestedObj::hasValue"
  column="nestedObj_hasValue"
  jdbc-type="INTEGER"
conversion="org.apache.ojb.broker.accesslayer.conversions.Boolean2IntFieldConversion"
/>
<field-descriptor
  name="nestedObj::id"
  column="nestedObj_id"
  jdbc-type="INTEGER"
/>
<field-descriptor
  name="nestedObj::containerId"
  column="nestedObj_containerId"
  jdbc-type="INTEGER"
/>
<reference-descriptor
  name="nestedObj::container"
  class-ref="ContainerObject"
>
  <foreignkey field-ref="nestedObj::containerId"/>
</reference-descriptor>
...
</class-descriptor>

```

and the table descriptor

```

<table name="ContainerObject">
  <column name="id"
    javaName="id"
    type="INTEGER"
    primaryKey="true"
    required="true"
  />
  <column name="name"
    javaName="name"
    type="VARCHAR"
    size="24"
  />
  <column name="nestedObj_hasValue"
    type="INTEGER"
  />
  <column name="nestedObj_id"
    type="INTEGER"
  />
  <column name="nestedObj_containerId"
    type="INTEGER"
  />
  <foreign-key foreignTable="ContainerObject">\n"+
    <reference local="nestedObj_containerId" foreign="id"/>\n"+
  </foreign-key>\n"+
  ...
</table>

```

Note how one **ojb.modify-nested** tag changes the type of the nested hasValue field, add a

conversion and specifies the position for it. The other modification tag removes the primarykey status of the nested id field.

## 5.5.19. OJB Performance

### 5.5.19.1. Introduction

*"There is no such thing as a free lunch."*  
(North American proverb)

Object/relational mapping tools hide the details of relational databases from the application developer. The developer can concentrate on implementing business logic and is liberated from caring about RDBMS related coding with JDBC and SQL.

O/R mapping tools allow to separate business logic from RDBMS access by forming an additional software layer between business logic and RDBMS. Introducing new software layers always eats up additional computing resources.

In short: the price for using O/R tools is performance. But on the other hand the biggest performance consumption is the database access itself (database performance, network performance, jdbc driver, driver settings, ...). So the percentage of O/R tool performance consumption isn't big.

Software architects have to take in account this tradeoff between programming comfort and performance to decide if it is appropriate to use an O/R tool for a specific software system.

This document describes the *OJB Performance Test Suite* which was created to lighten the decision between native JDBC, OJB (the different OJB API's) and other O/R mapper.

### 5.5.19.2. The Performance Test Suite

The *OJB Performance Test Suite* allows to compare all supported OJB API's against [native single-threaded JDBC programming](#) against your RDBMS of choice and run OJB API's in a [virtual multithreaded environment](#). Further on it is possible to [compare OJB against any O/R mapping tool](#) using a simple *performance application*.

All tests are integrated in the OJB build script, you only need to perform the according ant target:

```
ant target
```

The following 'targets' exist:

- `perf-test` multi-threaded performance/stress test of PB/OTM/ODMG api against native JDBC
- `performance` single-threaded test, OJB API implementations (PB, ODMG) against native JDBC

By changing the `JdbcConnectionDescriptor` in the configuration files you can point to your specific RDBMS. Please refer to this [document for details](#).

### Interpreting test results

Interpreting the result of these benchmarks carefully will help to decide whether using OJB is viable for specific application scenarios or if native JDBC programming should be used for performance reasons.

Take care of compareable configuration properties when run performance tests with different O/R tools.

If the decision made to use an O/R mapping tool the comparison with other tools helps to find the best one for the thought scenario. But performance shouldn't be the only reason to take a specific O/R tool. There are many other points to consider:

- Usability of the supported API's
- Flexibility of the framework
- Scalability of the framework
- Community support
- The different licences of Open Source projects
- etcetera ...

### How OJB compares to native JDBC programming - single-threaded

OJB is shipped with tests compares native JDBC with ODMG and PB-API implementation. This part of the test suite is integrated into the OJB build mechanism.

A single client test you can invoke it by typing `ant performance` or `ant performance`.

If running OJB out of the box the tests will be performed against the Hypersonic SQL (in-memory mode) shipped with OJB. A typical console output looks like follows:

```
performance:
  [obj] .[performance] INFO: Test for PB-api
  [obj] [performance] INFO:
  [obj] [performance] INFO: inserting 1500 Objects: 188 msec
  [obj] [performance] INFO: updating 1500 Objects: 265 msec
  [obj] [performance] INFO: querying 1500 Objects: 0 msec
  [obj] [performance] INFO: querying 1500 Objects: 0 msec
  [obj] [performance] INFO: fetching 1500 Objects: 16 msec
  [obj] [performance] INFO: deleting 1500 Objects: 63 msec
  ....
  [obj] Time: 5,672
  [obj] OK (1 test)

  [jdbc] .[performance] INFO: Test for JDBC
  [jdbc] [performance] INFO:
  [jdbc] [performance] INFO: inserting 1500 Objects: 157 msec
  [jdbc] [performance] INFO: updating 1500 Objects: 187 msec
  [jdbc] [performance] INFO: querying 1500 Objects: 94 msec
  [jdbc] [performance] INFO: querying 1500 Objects: 94 msec
  [jdbc] [performance] INFO: fetching 1500 Objects: 16 msec
  [jdbc] [performance] INFO: deleting 1500 Objects: 62 msec
  ....
  [jdbc] Time: 8,75
  [jdbc] OK (1 test)

  [odmg] .[performance] INFO: Test for ODMG-api
  [odmg] [performance] INFO:
  [odmg] [performance] INFO: inserting 1500 Objects: 266 msec
  [odmg] [performance] INFO: updating 1500 Objects: 359 msec
  [odmg] [performance] INFO: querying 1500 Objects: 531 msec
  [odmg] [performance] INFO: querying 1500 Objects: 531 msec
  [odmg] [performance] INFO: fetching 1500 Objects: 47 msec
  [odmg] [performance] INFO: deleting 1500 Objects: 125 msec
  ....
  [odmg] Time: 13,75
  [odmg] OK (1 test)
```

Some notes on these test results:

- You see a consistently better performance in the second and third run. this is caused by warming up effects of JVM and OJB.
- PB and native JDBC need about the same time for the three runs although JDBC performance is better for most operations. This is caused by the second run of the querying operations. In the second run OJB can load all objects from the cache, thus the time is **much** shorter. Hence the

interesting result: if you have an application that has a lot of lookups, OJB can be faster than a native JDBC application (without caching extensions)!

- ODMG is much slower than PB or JDBC. This is due to the complex object level transaction management it is doing and the fact that ODMG doesn't have a specific method to lookup objects by it's identity. The second reason is responsible for slow *querying* results, because in test always a complex query is done for each object. It is possible to use the PB-api within ODMG, then the query by identity will be as fast as in PB-api.
- You can see that for HSQLDB operations like insert and update are faster with JDBC than with PB. This performance difference strongly depends on the used [cache implementation](#) and can rise up to 50% when the cache operate on object copies. This ratio is so high, because HSQLDB was running with *in memory mode* and is much faster than ordinary database servers. If you work against Oracle or DB2 the percentual OJB overhead is going down a lot (10 - 15 %), as the database latency is much longer than the OJB overhead.

It's easy to change target database. Please refer to this [document for details](#).

Also it's possible to change the number of test objects by editing the ant-target in build.xml.

Another test compares PB-api,ODMG-api and native JDBC you can find [next section](#).

### OJB performance in multi-threaded environments

This test was created to check the performance and stability of the supported API's (PB-api, ODMG-api and future API's) in a single/multithreaded environment and to compare the different api's against native JDBC calls.

Running this test out of the box (a virgin OJB version against hsql) shouldn't cause any problems.

Per default OJB use a in-memory hsql database, by changing the [JdbcConnectionDescriptor](#) in the *repository.xml* file or modify the *build.properties* file when running OJB out of the box you can point to your specific RDBMS. Please refer to this [document for details](#).

To run the multithreaded performance test call

```
ant perf-test
```

A typical output of this test, using OJB against in-memory hsql looks like this

```
[obj]
-----
[obj]          OJB PERFORMANCE TEST SUMMARY, Thu Dec 29 23:42:20 CET 2005
[obj]
-----
[obj]  12 concurrent threads, handle 500 objects per thread
[obj]  500 INSERT operations per test instance
[obj]  FETCH collection of 500 objects per test instance
[obj]  Repeat FETCH collection of 500 objects per test instance
[obj]  125 get by Identity calls per test instance
[obj]  500 UPDATE operations per test instance
[obj]  500 DELETE operations per test instance
[obj]  - performance mode - results per test instance (average)
[obj]
-----
[obj] API          Total      Insert      Fetch      Fetch 2      by Id
[obj] Update      Delete
[obj]          [%]          [msec]      [msec]      [msec]      [msec]
[obj] [msec]      [msec]
-----
[obj] JDBC          100          475(100%)   26(100%)   23(100%)   209(836%)
[obj] 477(100%)   197(100%)
[obj] PB           203          1341(282%)  153(588%)  151(656%)  25(100%)
[obj] 648(135%)   239(121%)
[obj] ODMG          250          1469(309%)  104(400%)  105(456%)  527(2108%)
```

```
800(167%)    569(288%)
[obj]
=====
[obj] PerfTest takes 70 [sec]
```

This test run shows the overhead caused by the O/R layer compared to handcoded sql statements. Most overheads results in populate the two-level cache which is useless when using a in-memory database.

Below you can see the same test against *MaxDB* running on the same machine

```
....
[obj]
=====
[obj] API          Total      Insert      Fetch      Fetch 2      by Id
Update      Delete
[obj]          [%]          [msec]      [msec]      [msec]      [msec]
[obj]          [msec]
=====
[obj] JDBC          100          5855(100%)  55(100%)    38(100%)    1628(5087%)
5588(184%)  4084(136%)
[obj] ODMG          117          12043(205%) 180(327%)   294(773%)   754(2356%)
3027(100%)  2988(100%)
[obj] PB           123          11577(197%) 94(170%)    84(221%)    32(100%)
4240(140%)  3193(106%)
[obj]
=====
[obj] PerfTest takes 440 [sec]
```

**Note:**

The performance test output is written to console and in a file called *OJB-Performance-Result.txt*.

To change the test properties go to target `perf-test` in the `build.xml` file and change the program parameter.

The test needs five parameter:

- A comma separated list of the test implementation classes (no blanks!)
- The number of test loops
- The number of concurrent threads
- The number of managed objects per thread
- The desired test mode. `false` means run in performance mode, `true` means run in stress mode (useful only for developer to check stability).

```
<target name="perf-test" depends="prepare-testdb"
  description="Simple performance benchmark and stress test for PB- and
ODMG-api">
  <java fork="yes" classname="org.apache.ojb.performance.PerfMain"
    dir="${build.test}/obj" taskname="obj" failonerror="true" >
    <classpath refid="runtime-classpath"/>
    <!-- comma separated list of the PerfTest implementations -->
    <arg value=
      "org.apache.ojb.compare.OJBPerfTest$JdbcPerfTest,
      org.apache.ojb.compare.OJBPerfTest$PBPerfTest,
      org.apache.ojb.compare.OJBPerfTest$ODMGPerfTest"
    />
    <!-- test loops, default was 6 -->
    <arg value="6"/>
    <!-- performed threads, default was 12 -->
    <arg value="12"/>
    <!-- number of managed objects per thread, default was 500 -->
    <arg value="500"/>
    <!-- if 'false' we use performance mode, 'true' we do run in stress mode
-->
    <arg value="false"/>
```

```

    <!-- if 'true' all log messages will be print -->
    <arg value="true"/>
    <jvmarg value="-Xms128m"/>
    <jvmarg value="-Xmx256m"/>
  </java>
  <!-- do some cleanup -->
  <ant target="copy-testdb"/>
</target>

```

### 5.5.19.3. How OJB compares to other O/R mapping tools?

Many user ask this question and there is more than one answer. But OJB was shipped with a simple *performance application* (independent from OJB) which allows a rudimentarily comparison of OJB with other (java-based) O/R mapping tools.

All involved classes can be found in directory `[db-ojb]/src/test` in package `org.apache.ojb.performance`.

Call `ant perf-test-jar` to build the jar file contain all necessary classes to set up a test with an arbitrary O/R mapper. After the build, the `db-ojb-XXX-performance.jar` can be found in `[db-ojb]/dist` directory.

#### Steps to set up the test for other O/R frameworks:

- Implement a class derived from [PerfTest](#)
- If persistent objects (used within your mapping tool) must be derived from a specific base class or must be implement a specific interface write your own persistent object class by implementing [PerfArticle](#) interface and **override method** `#newPerfArticle()` in your `PerfTest` implementation class. Otherwise a default implementation of interface [PerfArticle](#) was used. The default implementation class is `org.apache.ojb.performance.PerfArticleImpl`.

That's it!

You can find a example implementation called `org.apache.ojb.compare.OJBPerfTest` in the test-sources directory under `[db-ojb]/src/test` (when using source-distribution).

This implementation class is used to compare the performance of the PB-API, ODMG-API, OTM-api and native JDBC (to bunch all API's, for each API a static inner implementation class of [PerfTest](#) was used). See more [about multi-threaded performance](#).

#### Run the test

You have two possibilities to run the test:

a) Integration in the OJB build script

Add the full qualified class name of your `PerfTest` implementation class to the `perf-test` target of the OJB `build.xml` file, add all necessary jar files to `[db-ojb]/lib`. The working directory of the test is `[db-ojb]/target/test/ojb`.

b) Run `PerfMain`

It's possible to run the test using `org.apache.ojb.performance.PerfMain`.

```

java -classpath CLASSPATH org.apache.ojb.performance.PerfMain
[comma separated list of PerfTest implementation classes, no blanks!]
[number of test loops]
[number of threads]
[number of insert/fetch/delete loops per thread]
[optional boolean - run in stress mode if set true,
run in performance mode if set false, default false]
[optional boolean - if 'true' all log messages will be print, else only a test
summary, default 'true']

```

For example:

```
java -classpath CLASSPATH my.A_PerfTest,my.B_PerfTest 3 10 2000
```

This will use `A_PerfTest` and `B_PerfTest` and perform three loops each loop run 10 threads and each thread operate on 2000 objects. The test run in *normal* mode and log all messages.

Take care of compareable configuration properties when run performance tests with different O/R tools (caching, locking, sequence key generation, connection pooling, ...).

**Note:**

Please, don't start flame wars by posting performance results to mailing lists made with this simple test. This test was created for OJB QA and to give a clue how good or bad OJB performs, NOT to start discussion like *XY is 12% faster than XZ!!*.

#### 5.5.19.4. What are the best settings for maximal performance?

We don't know, that depends from the environment OJB runs (hardware, database, driver, application server, ...). But there are some settings which affect the performance:

- The API you use, e.g. PB-api is much faster then the ODMG-api. See [which API](#) for more information.
- The used [cache implementation](#).
- ConnectionFactory implementation / Connection pooling. See [connection pooling](#) for more information.
- The *autocommit* setting of used connections. For best performance it's recommended to use *autocommit* 'false' setting in the *jdbc-driver* to avoid `Connection.setAutoCommit(...)` calls by OJB.
- [PersistentField class implementation](#). See [OJB.properties](#) section '[PersistentFieldClass](#)' to change the implementation.
- Used sequence manager implementation. See [sequence manager](#) for more information.
- Use of *batch mode* (when supported by the DB). See [jdbc-connection-descriptor](#) *batch-mode* attribute for more information.
- PersistenceBroker pool size. See [OJB.properties](#) for more information.
- The *JDBC driver* settings (e.g. statement caching on/off).

To test the different settings use the tests of the [performance test suite](#).

## 5.6. Howto's

### 5.6.1. Howto's Summary

Summary of all HOWTO documentation with practical hands-on information, submitted both by OJB users and developers.

- [How to build large metadata mappings](#)
- [Using anonymous keys for cleaner objects](#)
- [Using native database sequences](#)
- [Using OJB in a clustered environment](#)
- [Using OJB with stored procedures](#)
- [Using Oracle LOB's](#)
- [Using Spring with OJB](#)

### 5.6.2. HOWTO - Build O/R Mapping Files



### 5.6.2.1. How to build O/R mapping files

Writing the repository.xml file for only a few classes can easily be done manually with the text or xml editor of your choice.

But keeping the repository in sync with the java codebase and the database gets more difficult if several hundred classes and large developer teams are involved.

This page contains tips how to integrate mapping tools and code-generators into your build process.

### 5.6.2.2. classification of O/R related transformations

Let's start with a classification of the source transformation problems that developers have to face in an O/R environment.

Typical development environments contain some or all of the following artefacts:

- A UML model containing at least class diagrams of the persistent classes. All modern UML tools can export to the XMI standard format.
- Other tools, such as Torque, also use a model based approach but use different model file formats (typically XML based)
- Java source code for the persistent classes. The Java source code can possibly be enhanced with xdoclet tags.
- The OJB repository.xml file. This file contains all the class-descriptors for the persistent classes.
- The database. This could be an online DB or a DDL script representing the database tables. The database contains all tables used to store instances of the persistent classes.

The technique you will use depends a lot on the problem you have to solve, on the methodology and the tool chain you have in use, which of transformations between those artefacts fit to your development process.

1. **Forward engineering from XMI:** A UML model in XMI format with class diagrams of your persistent classes exists and is used as the master source (model driven approach). Java code, repository.xml and DDL for the database tables have to be generated from this model.
2. **Forward engineering from Torque:** A model of the persistent entity classes exists in form of a Torque.XML file. Java code, repository.xml and DDL for the database tables have to be generated from this model.
3. **Forward engineering from the repository.xml:** The OJB repository.xml file is used a model format. Java code and DDL for the database tables have to be generated from this model.
4. **Xdoclet transformation from Java code:** Java code for the persistent classes exists and contains special comment tags in the Xdoclet ojb-module format. Repository.xml and DDL for the database tables have to be generated from the java files via Xdoclet transformation.
5. **Reverse engineering from database:** There is a database with existing tables or a DDL script. Java code and repository.xml have to be generated from the database.

These transformations are depicted in the following graphics. The numbers close to the arrows correspond to the numbers in the above enumeration. All related transformations have the same colour.

mapping tools image

In the following sections we will have a closer look at each of these transformations and discuss tools that provide support each approach.

### 5.6.2.3. Forward engineering from XMI

This approach is recommended if you start from scratch with a new project and have to deal with a

large number of persistent classes. This approach works best when there are no restrictions regarding the database, like integration of legacy tables.

Forward engineering from XMI fits perfectly into a model driven architecture (MDA) software development process.

### Tool support

- **AXGen**

AXgen is a code generator using XMI as input and Velocity templates for transformation.

The power of AXgen is in its simplicity. You don't have to understand complicated structure of your XMI file to write an XSLT stylesheet for transformation. AXgen uses nsuml to deal with the xmi file, which gives access to the Metamodel in an objectoriented way.

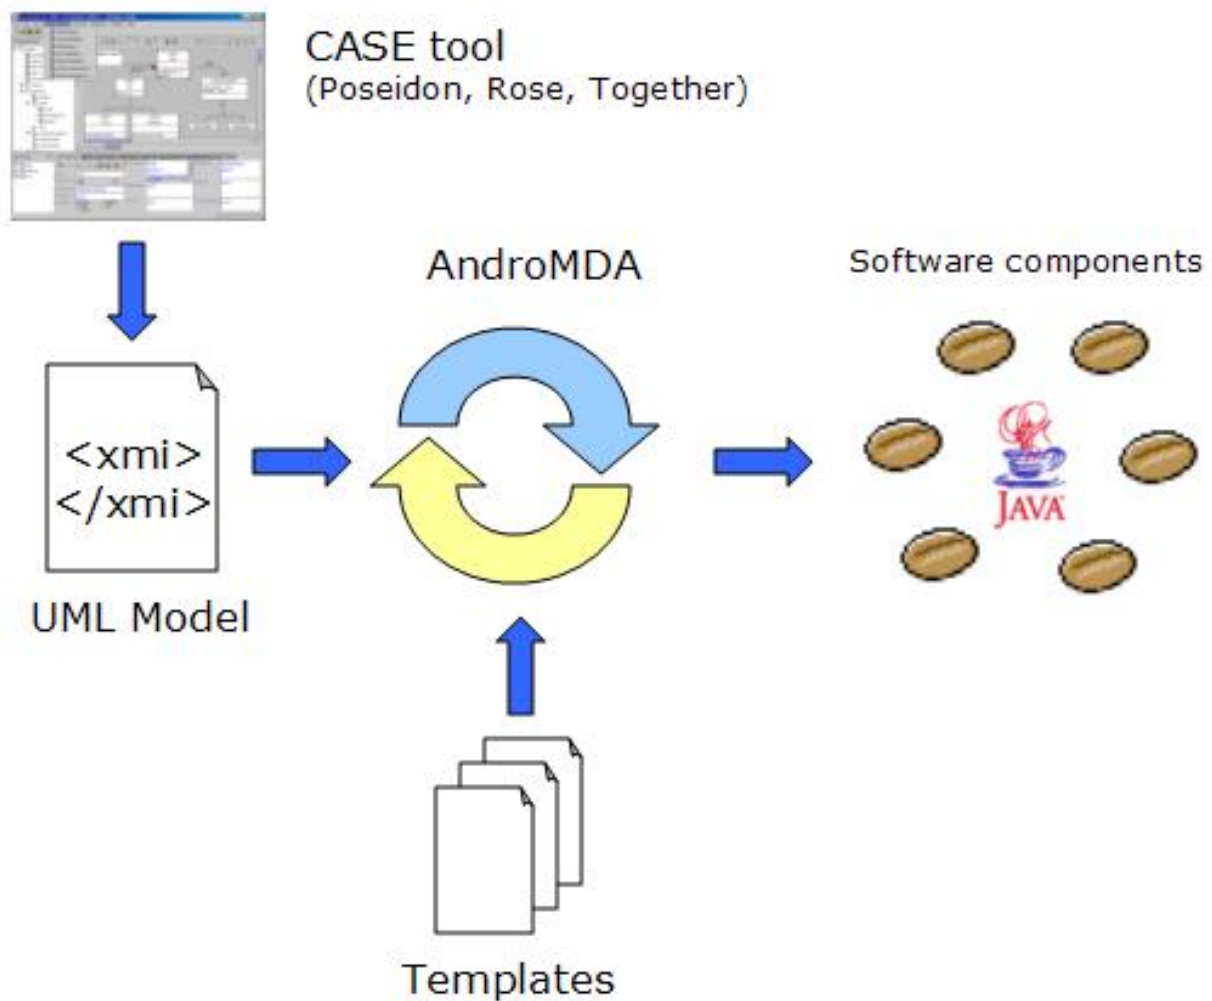
Further AXgen makes use of Jakartas Velocity. Velocity is a very sophisticated Java-based template engine. This means that inside your templates you can call Java methods. Feel free to write templates that generate anything you want.

Our motive for AXgen is to generate Java Classes for use in a O/R Mapping tool that allows transparent persistence for Java Objects against relational databases. Therefore AXgen comes with a bundle of ready to use templates for generating ObJectRelationalBridge (OJB) specific stuff like:

- Entity Classes
- XML Repository
- SQL script to build the DB scheme

- **AndroMDA**

AndroMDA is a code generator framework - it takes a Unified Modeling Language (UML) model from a CASE-tool in XMI format and generates custom components. It comes with a set of sample templates that generate classes attributed with XDoclet tags. One build step later, the XDoclet tool generates full-blown components that can readily be deployed in the JBoss application server (and the other servers that XDoclet can feed).



andromeda image

Currently AndroMDA provides no special support for OJB. But by tagging classes with tags of the [XDoclet OJB module](#) it is possible to use it as a full forward engineering engine.

- Searching the Sourceforge project list for "XMI" returns a long list of projects dealing with code generation. It may be a good idea to check if you find a tool that matches your requirements.

#### 5.6.2.4. Forward engineering from Torque

##### Torque

Torque is a persistence layer. Torque includes a generator to generate all the database resources required by your application and includes a runtime environment to run the generated classes.

Torque was developed as part of the Turbine Framework. It is now decoupled and can be used by itself. Starting with version 2.2 Turbine uses the decoupled Torque.

Torque uses a single XML database schema to generate the SQL for your target database and Torque's Peer-based object relation model representing your XML database schema.

You can use [devaki-nextobjects](#) to create the model for your application.

OJB uses Torque's generator engine to setup the testbed database and feed it with initial data.

Besides the SQL generation facilities Torque also provides special support for OJB related transformations. It provides the following two ant targets:

- **ojb-model**  
generates a simple object model for ojb
- **ojb-repository**  
generates the repository for ojb

A complete list of all available Torque targets can be found at the [Torque Generator site](#).

#### 5.6.2.5. Forward engineering from repository.xml

There is currently no tool available that directly supports this model. It is not difficult to implement an XSLT stylesheet that transforms the OJB repository.xml directly into DDL Statements.

An even simpler approach could be to transform the repository.xml file into a Torque xml file. DDL can then be generated by the Torque engine.

If you write such an XSLT file please tell us about it!

#### 5.6.2.6. XDoclet transformation from Java code

##### XDoclet

XDoclet is a code generation engine. It enables Attribute-Oriented Programming for java. In short, this means that you can add more significance to your code by adding meta data (attributes) to your java sources. This is done in special JavaDoc tags.

OJB was shipped with its own [xdoclet-module](#).

XDoclet will parse your source files and generate many artifacts such as XML descriptors and/or source code from it. These files are generated from templates that use the information provided in the source code and its JavaDoc tags.

XDoclet lets you apply Continuous Integration in component-oriented development. Developers should concentrate their editing work on only one Java source file per component.

XDoclet originated as a tool for creating EJBs, it has evolved into a general-purpose code generation engine. XDoclet consists of a core and a constantly growing number of modules.

#### 5.6.2.7. Reverse engineering from database

- **Druid**  
Druid is a tool that allows users to create databases in a graphical way. The user can add or import tables, fields, folders to group tables and can modify most of the database options that follow the SQL-92 standard. In addition to sql options, the user can document each table and each field with HTML information. It is distributed with modules for generating Java classes, OJB metadata, and JDO metadata.
- **Impart Eclipse Plugin for OJB**  
The Impart Eclipse plugin is based on the OJB ReverseDB Tool and provides the same functionality (and also some additional goodies). It ships as a plugin to the Eclipse IDE. It provides a very convenient GUI that integrates smoothly into the Eclipse platform.
- **RDBS2J**  
RDBS2J is a GUI based mapping tool from relational database scheme to persistent java classes which use JDO as persistence mechanism. The mapping can be modified by the GUI. The current version is designed to create code for OJB. The ODMG and the JDO interface are supported. RDBS2J creates the \*.jdo files and the repository\_user.xml, which are needed by OJB.
- **The OJB ReverseDB tool**  
OJB ships with a simple reverse engineering tool that allows to connect to a RDBMS via JDBC and to take the tables from the database catalog as input.

This tool provides a nice GUI to generate Java classes and the matching repository.xml file. You can invoke the ReverseDB tool with the ANT target `reverse-db`.

**Note:**

The ReverseDB tool is not up to date - any help is welcome.

### 5.6.3. HOWTO - Use Anonymous Keys

#### 5.6.3.1. Why Do We Need Anonymous Keys?

The core difference between referential integrity in Java and in an RDBMS lies in where the specific referential information is maintained. Java, and most modern OO languages, maintain referential integrity information in the runtime environment. Actual object relationships are maintained by the virtual machine so that the symbolic variable used in the application is dereferenced it will in fact provide access to the object instance which it is expected to provide access to. There is no need for a manual lookup or search across the heap for the correct object instance. Entity reference integrity is maintained and handled for the programmer by the environment.

Relational databases, on the other, purposefully place the referential integrity and lookups into the problem domain - that is the problem they are designed to solve. An RDBMS presumes you can design something more efficient for your specific circumstances than the JVM does (you trust its ability to do object lookups in the heap is sufficiently efficient). As an RDBMS has a much larger heap equivalent it is designed to not operate under that assumption (mostly). So, in an RDBMS the concept of specific foreign keys exists to maintain the referential integrity.

In crossing the object to relational entity barrier there is a mismatch between the referential integrity implementations. Java programmers do not want to have to maintain both object referential integrity and key referential integrity analogous to

```
{
    Foo child = new SomeOtherFooType();
    Foo parent = new SomeFooType();
    child.setParent(parent);
    child.setParentId(parent.getId());
}
```

This is double the work required - you set up the object relationship, then set up the key relationship. OJB knows about the relationship of the objects, thus it is only needed to do

```
{
    Foo child = new Foo();
    Foo parent = new Foo();
    child.setParent(parent);
}
```

OJB can provide transparent key relationship maintenance behind the scenes for [1:1 relations](#) via **anonymous access fields**. As object relationships change, the relationships will be propagated into the key values without the *Java object ever being aware of a relational key* being in use. This means that the java object does not need to specify a FK field for the reference.

Without use of *anonymous keys* class Foo have to look like:

```
class Foo
{
    Integer id;
    Integer fkParentFoo;
    Foo parent;
```

```
// optional getter/setter
....
{
```

When using *anonymous keys* the FK field will become obsolete:

```
class Foo
{
    Integer id;
    Foo parent;

    // optional getter/setter
    ....
}
```

**Note:**

Under specific conditions it's also possible to use anonymous keys for other relations or primary keys. More info [in advanced-technique section](#).

### 5.6.3.2. How it works

To play for safety it is mandatory to understand how this feature is working. More information how it works [please see here](#).

### 5.6.3.3. Using Anonymous Keys

Now we can start using of the *anonymous key* feature. In this section the using is detailed described on the basis of an example.

#### The Code

Take the following classes designed to model a particular problem domain. They may do it reasonably well, or may not. Presume they model it perfectly well for the problem being solved.

```
public class Desk
{
    private Integer id;
    private Finish finish;
    /** Contains Drawer instances */
    private List drawers;
    private int numberOfLegs;

    public Desk()
    {
        this.drawers = new ArrayList();
    }
    ....
    // getter/setter
}

public class Drawer
{
    private Integer id;
    /** Contains Thing instances */
    private List stuffInDrawer;

    public Drawer()
    {
        this.stuffInDrawer = new ArrayList();
    }
    ....
    // getter/setter
}
```

```

public class Finish
{
    private Integer id;
    private String wood;
    private String color;

    ....
    // getter/setter
}

public class Thing
{
    private Integer id;
    private String name;

    ....
    // getter/setter
}

```

A Desk will typically reference multiple drawers and one finish.

### The Database

When we need to store our instances in a database we use a fairly typical table per class persistence model.

```

CREATE TABLE finish
(
    id            INTEGER PRIMARY KEY,
    wood          VARCHAR(255),
    color         VARCHAR(255)
);

CREATE TABLE desk
(
    id            INTEGER PRIMARY KEY,
    num_legs     INTEGER,
    finish_id    INTEGER,
    FOREIGN KEY (finish_id) REFERENCES finish(id)
);

CREATE TABLE drawer
(
    id            INTEGER PRIMARY KEY,
    desk_id     INTEGER,
    FOREIGN KEY (desk_id) REFERENCES desk(id)
);

CREATE TABLE thing
(
    id            INTEGER PRIMARY KEY,
    name         VARCHAR(255),
    drawer_id    INTEGER,
    FOREIGN KEY (drawer_id) REFERENCES drawer(id)
);

```

At the database level the possible relationships need to be explicitly defined by the foreign key constraints. These model all the possible object relationships according to the domain model (until generics enter the Java language for the collections API, this is technically untrue for the classes used here).

### The Repository Configuration

When we go to map the classes to the database, it is almost a one-to-one property to field mapping. The exception here is the primary key on each entity. This is meaningless information in Java, so we would like to keep it out of the object model. Anonymous access keys allow us to do that.

The [repository.xml](#) must know about the database columns used for referential integrity, but OJB can maintain the foreign key relationships behind the scenes - freeing the developer to focus on more accurate modeling of her objects to the problem, instead of the the persistence mechanism. Doing this is also very simple - in the *repository.xml* file mark the [field descriptors](#) with a `access="anonymous"` attribute.

```
<class-descriptor
  class="Desk"
  table="desk">

  <field-descriptor
    name="id"
    column="id"
    jdbc-type="INTEGER"
    primarykey="true"
    autoincrement="true"
  />
  <field-descriptor
    name="numberOfLegs"
    column="num_legs"
    jdbc-type="INTEGER"
  />
  <field-descriptor
    name="finishId"
    column="finish_id"
    jdbc-type="INTEGER"
    access="anonymous" />

  <reference-descriptor
    name="finish"
    class-ref="Finish">
    <foreignkey field-ref="finishId"/>
  </reference-descriptor>

  <collection-descriptor
    name="drawers"
    element-class-ref="Drawer"
  >
    <inverse-foreignkey field-ref="deskId"/>
  </collection-descriptor>
</class-descriptor>

<class-descriptor
  class="Finish"
  table="finish">

  <field-descriptor
    name="id"
    column="id"
    jdbc-type="INTEGER"
    primarykey="true"
    autoincrement="true"
  />
  <field-descriptor
    name="wood"
    column="wood"
    jdbc-type="VARCHAR"
    size="255"
  />
  <field-descriptor
    name="color"
    column="color"
    jdbc-type="VARCHAR"
    size="255"
  />
</class-descriptor>

<class-descriptor
```



```

class="Drawer"
table="drawer">

  <field-descriptor
    name="id"
    column="id"
    jdbc-type="INTEGER"
    primarykey="true"
    autoincrement="true"
  />
  <field-descriptor
    name="deskId"
    column="desk_id"
    jdbc-type="INTEGER"
    access="anonymous"
  />
  <collection-descriptor
    name="stuffInDrawer"
    element-class-ref="Thing"
  >
    <inverse-foreignkey field-ref="drawerId"/>
  </collection-descriptor>
</class-descriptor>

<class-descriptor
class="Thing"
table="thing">

  <field-descriptor
    name="id"
    column="id"
    jdbc-type="INTEGER"
    primarykey="true"
    autoincrement="true"
  />
  <field-descriptor
    name="name"
    column="name"
    jdbc-type="VARCHAR"
    size="255"
  />
  <field-descriptor
    name="drawerId"
    column="drawer_id"
    jdbc-type="INTEGER"
    access="anonymous"
  />
</class-descriptor>

```

Look first at the class descriptor for the Thing class. Notice the [field-descriptor](#) with the name attribute "drawerId". This field is labeled as anonymous access. Because it is anonymous access OJB will not attempt to assign the value here to a "drawerId" field or property on the Thing class. Normally the name attribute is used as the Java name for the attribute, in this case it is not. The name is still required because it is used as an indicated for references to this anonymous field.

In the field descriptor for Drawer, look at the collection descriptor with the name *stuffInDrawer*. This collection descriptor references a foreign key with the `field-ref="drawerId"`. This reference is to the anonymous field descriptor in the Thing descriptor. The field-ref matches to the name in the descriptor whether or not the name also maps to the Java attribute name. This dual use of name can be confusing - be careful.

The same type mapping that is used for the collection descriptor in the Drawer descriptor is also used for the 1:1 [reference descriptor](#) in the Desk descriptor.

The primary keys are populated into the objects as it is generally a good practice to not implement primary keys as anonymous access fields. Primary keys may be anonymous-access but references will get lost if the cache is cleared or the persistent object is serialized.

### 5.6.3.4. Benefits and Drawbacks

There are both benefits and drawbacks to using anonymous field references for maintaining referential integrity between Java objects and database relations. The most immediate benefit is avoiding semantic code duplication. The second major benefit is avoiding cluttering class definitions with persistence mechanism artifacts. In a well layered application, the persistence mechanism will not really need to be so obvious in the object model implementation. Anonymous fields help to achieve this - thereby making persistence mechanisms more flexible. Moving to a different one becomes easier.

## 5.6.4. HOWTO - Use DB Sequences

### 5.6.4.1. Introduction

It is easy to use OJB with with database generated sequences. Typically a table using database generated sequences will autogenerate a unique id for a field as the default value for that field. This can be particularly useful if multiple applications access the same database. Not every application will be using OJB and find it convenient to pull unique values from a high/low table. Using a database managed sequence can help to enforce unique id's across applications all adding to the same database. All of that said, care needs to be taken as using database generated sequences imposes some portability problems.

OJB includes a sequence manager implementation that is aware of database sequences and how to use them. It is known to work against Oracle, SAP DB, and PostgreSQL. MySQL has its own sequence manager implementation because it is special. This tutorial will build against PostgreSQL, but working against Oracle or SAP will work the same way.

Additional information on sequence managers is available in the [Sequence Manager documentation](#).

### 5.6.4.2. The Sample Database

Before we can work with OJB against a database with a sequence, we need the database. We will create a simple table that pulls its primary key from a sequence named 'UniqueIdentifier'.

```
CREATE TABLE thingie
(
  name VARCHAR(50),
  id INTEGER DEFAULT NEXTVAL('UniqueIdentifier')
)
```

We must also define the sequence from which it is drawing values:

```
CREATE SEQUENCE UniqueIdentifier;
```

So that we have the following table:

Column	Type	Modifiers
name	character varying(50)	
id	integer	default nextval('UniqueIdentifier'::text)

If we manually insert some entries into this table they will have their id field set automatically.

```
INSERT INTO thingie (name) VALUES ('Fred');
INSERT INTO thingie (name) VALUES ('Wilma');
SELECT name, id FROM thingie;
```

```

name | id
-----+-----
Fred | 0
Wilma | 1
(2 rows)

```

### 5.6.4.3. Using OJB

#### The Database Repository Descriptor

The next step is to configure OJB to access our thingie table. We need to configure the correct sequence manager in the `repository-database.xml`.

The default `repository-database.xml` uses the High/Low Sequence manager. We will delete or comment out that entry, and replace it with the `org.apache.ojb.broker.util.sequence.SequenceManagerNextValImpl` manager. This manager will pull the next value from a named sequence and use it. The entry for our sequence manager in the repository is:

```

<sequence-manager
  className="org.apache.ojb.broker.util.sequence.SequenceManagerNextValImpl"
/>

```

This needs to be declared within the JDBC Connection descriptor, so an entire `repository-database.xml` might look like:

```

<jdbc-connection-descriptor
  jcd-alias="default"
  default-connection="true"
  platform="PostgreSQL"
  jdbc-level="2.0"
  driver="org.postgresql.Driver"
  protocol="jdbc"
  subprotocol="postgresql"
  dbalias="test"
  username="tester"
  password=""
  eager-release="false"
  batch-mode="false"
  useAutoCommit="1"
  ignoreAutoCommitExceptions="false"
>

  <connection-pool
    maxActive="21"
    validationQuery="" />

    <sequence-manager
      className="org.apache.ojb.broker.util.sequence.SequenceManagerNextValImpl" />
</jdbc-connection-descriptor>

```

#### Defining a Thingie Class

For the sake of simplicity we will make a very basic Java Thingie:

```

public class Thingie
{
    /** thingie(name) */
    private String name;

    /** thingie(id) */
    private int id;

    public String getName()      { return this.name; }
    public void setName(String name) { this.name = name; }
}

```

```

    public int getId()                { return this.id; }
}

```

We also need a class descriptor in `repository-user.xml` that appears as follows:

```

<class-descriptor
  class="Thingie"
  table="THINGIE"
  >
  <field-descriptor
    name="id"
    column="ID"
    jdbc-type="INTEGER"
    primarykey="true"
    autoincrement="true"
    sequence-name="UniqueIdentifier"
  />
  <field-descriptor
    name="name"
    column="NAME"
    jdbc-type="VARCHAR"
  />
</class-descriptor>

```

Look over the `id` field descriptor carefully. The `autoincrement` and `sequence-name` attributes are important for getting our desired behavior. These tell OJB to use the sequence manager we defined to auto-increment the value in `id`, and they also tell the sequence manager which database sequence to use - in this case `UniqueIdentifier`

We could allow OJB to create an extent-aware sequence and use it, however as we are working against a table that defaults to a specific named sequence, we want to make sure to pull values from that same sequence. Information on allowing OJB to create its own sequences is available in the [Sequence Manager](#) documentation.

### Using Thingie

Just to demonstrate that this all works, here is a simple application that uses our `Thingie`.

```

import org.apache.ojb.broker.PersistenceBroker;
import org.apache.ojb.broker.PersistenceBrokerFactory;

public class ThingieDriver
{
    public static void main(String [] args)
    {
        PersistenceBroker broker =
PersistenceBrokerFactory.defaultPersistenceBroker();

        Thingie thing = new Thingie();
        Thingie otherThing = new Thingie();

        thing.setName("Mabob");
        otherThing.setName("Majig");

        broker.beginTransaction();
        broker.store(thing);
        broker.store(otherThing);
        broker.commitTransaction();

        System.out.println(thing.getName() + " : " + thing.getId());
        System.out.println(otherThing.getName() + " : " + otherThing.getId());
        broker.close();
    }
}

```

When it is run, it will create two `Thingie` instances, store them in the database, and report on their

assigned id values.

```
java -cp [...] ThingieDriver
Mabob : 2
Majig : 3
```

## 5.6.5. HOWTO - Work with LOB Data Types

### 5.6.5.1. Using Oracle LOB Data Types with OJB

#### Introduction

In a lot of applications there is a need to store large text or binary objects into the database. The definition of large usually means that the object's size is beyond the maximum length of a character field. In Oracle this means the objects to be stored can grow to > 4 KB each.

Depending on the application you are developing your "large objects" may either be in the range of some Kilobytes (for example when storing text-only E-Mails or regular XML documents), but they may also extend to several Megabytes (thinking of any binary data such as larger graphics, PDFs, videos, etc.).

In practice, the interface between your application and the database used for fetching and storing of your "large objects" needs to be different depending on the expected size. While it is probably perfectly acceptable to handle XML documents or E-Mails in memory as a string and always completely retrieve or store them in one chunk this will hardly be a good idea for video files for example.

This HOWTO will explain:

1. Why you would want to store large objects in the database
2. Oracle LARGE versus LOB data types
3. LOB handling in OJB using JDBC LOB types

This tutorial presumes you are familiar with the basics of OJB.

#### 5.6.5.2. Backgrounder: Large objects in databases

This section is meant to fill in non-DBA people on some of the topics you need to understand when discussing large objects in databases.

#### Your database: The most expensive file system?

Depending on background some people tend to store anything in a database while others are biased against that. As databases use a file system for physical storage anyway, why would it make sense to store pictures, videos and the like as a large object in a database rather than just create a set of folders and store them right into the database.

When listening to Oracle's marketing campaigns one might get the impression that there is no need to have plain filesystems anymore and that they all will vanish and be replaced by Oracle database servers. If that happened this would definitely boost Oracle's revenues, but at the same time make IT cost in companies explode.

But there are applications where it in fact makes sense to have the database take care of unstructured data that you would otherwise just store in a file. The most common criteria for storing non-relational data in the database instead of storing it directly into the file system is whenever there is a strong link between this non-relational and some relational data.

Typical examples for that would be:

1. Pictures or videos of houses in a real estate agent's offer database
2. E-Mails related to a customer's order

If you are not storing these objects into the database you would need to create a link between the relational and the non-relational data by saving filenames in the database. This means that your application is responsible for managing this weak link in any respect. In order to make sure your application will be robust you need to make sure in your own code that

1. When creating a new record you create a valid and unique filename for storing the object.
2. When deleting a record you delete the corresponding file as well
3. When accessing the file referred to in the record you double-check the file is there and not locked

(There might be other, more subtle implications.)

All this is done for you by the database in case you choose to store your objects there. In addition to that, when discussing text data, a database might come with an option to automatically index the stored text documents for easy retrieval. This would allow you to perform an SQL search such as "give me all customers that ever referred to the project foo in an e-mail". (In Oracle you need to install the InterMedia option, aka Oracle Text in order to get this extra functionality. Several vendors have also worked on technologies that allowed to search rich content such as PDFs files, pictures or even sound or video stored in a database from SQL.)

#### Oracle LARGE versus LOB datatypes

Some people are worried about the efficiency of storing large objects in databases and the implications on performance. They are not necessarily entirely wrong in fearing that storing large objects in databases might be problematic the best or might require a lot of tweaks to parameters in order to be able to handle the large objects. It all depends on how a database implements storing large objects.

Oracle comes with two completely different mechanisms for that:

1. LARGE objects
2. LOB objects

When comparing the LARGE datatypes such as (\*fixme\*) to the LOB datatypes such as CLOB, BLOB, NCLOB (\*fixme\*) they don't read that different at first. But there is a huge difference in how they are handled both internally inside the database as well when storing and retrieving from the database client.

LARGE fields are embedded directly into the table row. This has some consequences you should be aware of:

1. If your record is made up of 5 VARCHAR fields with a maximum length of 40 bytes each and one LONGVARCHAR and you store 10 MB into the LONGVARCHAR column, your database row will extent to 10.000.200 bytes or roughly 10 MB.
2. The database always reads or writes the entire row. So if you do a SELECT on the VARCHAR fields in order to display their content in a user interface as a basis for the user to decide if he or she will need the content of the LONGVARCHAR at all the database will already have fetched all the 10 MB. If you SELECT and display 25 records each with a 10 MB object in it this will mean about 250 MB of I/O.
3. When storing or fetching such a row you need to make sure your fetch buffer is sized appropriately.

In practice this cannot be efficient. It might work as long as you stay in the KB range, but you will most likely run into trouble as soon as it gets into the MBs per record. Additionally, there are more limitations to the concept of LONG datatypes such as limiting the number of them you can have in

one row and how you can index them. This is probably why Oracle decided to deprecate LONG datatypes in favor of LOB columns.

A lot of non-Oracle-DBA people believe that LOB means "large Object" because some other vendors have used the term BLOB for "Binary Large Object" in their products. This is not only wrong but - even worse - misleading, because people are asking: "What's the difference between large and long?" (Bear with all non native English speakers here, please!)

Instead, LOB stands for Locator Object which exactly describes what it is. It is a pointer to the place where the actual data itself is stored. This locator will need only occupy some bytes in the row thus not harming row size at all. So all the issues discussed above vanish immediately. For the sake of simplicity think of a LOB as a pointer to a file on the databases internal file system that stores the actual content of that field. (Oracle might use plain files or different mechanisms in their implementation, we don't have to care.)

But as there is always a trade-off while LOBs are extremely handy inside a row, they are more complex to store and retrieve. As opposed to all other column types their actual content stays where it is even if you transfer the row from the database to the client. All that goes over the wire in that case will be a token representing the actual LOB column content.

In order to read the content or to write LOB content it needs to open a separate stream connection over the network that can be read from or written to similar to a file on a network file system. JDBC (starting at version \*fixme\*) comes with special objects such as `java.sql.Blob` and `java.sql.Clob` to access the content of LOBs that do not represent character arrays or strings but streams!

### 5.6.5.3. Large Objects in OJB

After having skipped the above Backgrounder (in case you do Oracle administration for a living) of having read and understood it (hopefully applies to the rest of us) now that you've most likely decided to go for LOBs and forget about LONGs how is this handled with OJB?

#### Strategy 1: Using streams for LOB I/O

##### to be written #####

#### Strategy 2: Embedding OJB content in Java objects

##### to be written #####

#### Querying CLOB content

##### to be written #####

## 5.6.6. HOWTO - Use OJB in clustered environments

### 5.6.6.1. How to use OJB in clustered environments

Object/Relational Bridge will work fine in environments that require robustness features such as load-balancing, failover, and clustering. However, there are some important steps that you need to take in order for your data to be secure, and to prevent isolation failures. These steps are outlined below.

I have tested this in a number of environments, and with Servlet engines and J2EE environments. If you run into problems, please post questions to the OJB users mail list.

This outline for clustering is based on an email from the OJB Users Mail List: [This is that mail.](#)

### 5.6.6.2. Three steps to clustering your OJB application

A lot of people keep asking for robust ways to run their OJB engines in multi-VM or clustered environments. This email covers how to do that safely and securely using Open Symphony's OSCache caching product.

OSCache is a high-performance, robust caching product that supports clustering. I've been using it for a while in production and it is excellent.

Back to the Topic: There are three main things that you should do in order to make your OJB application ready for using a cache in a multi-VM or distributed environment.

#### First: Take care of the sequence manager

that you define within *[jdbc-connection-descriptor](#)* element in your repository.xml file. If none was set OJB use per default the `SequenceManagerHighLowImpl` sequence manager implementation.

#### Note:

As of Release Candidate 5 (rc5), you can use `SequenceManagerHighLowImpl` in distributed (non-managed) environments. The `SequenceManagerHighLowImpl` now supports its own optimistic locking that makes the implementation cluster aware by versioning an entry in the `OJB_HL_SEQ` table.

However, the `SequenceManagerHighLowImpl` has not been heavily tested in clustered environments, so if you want absolute security use an sequence manager implementation which delegates key generation to database.

If your database supports database based sequence key generation (like PostgreSQL, Oracle, ...) it's recommended to use `SequenceManagerNextValImpl` (supports database based sequence keys). Using this sequence manager will prevent conflicts in multi-vm or clustered environments. More about [sequence manager here](#).

#### Handling sequence names

If you are using `SequenceManagerNextValImpl` you have two possibilities:

- Do it by your own:
  - Create a sequence object in your database.
    - An Oracle sequence lookslike: "create sequence ackSequence increment by 1 start with 1;"
    - A Postgres sequence looks like: "CREATE SEQUENCE ackSequence START 1";
  - For other databases you're on your own.
  - To tell OJB to use that sequence for your table add in your repository.xml the sequence name to the field-descriptor for your table's primary key field:

```
<field-descriptor
name="ackID"
column="ACKID"
jdbc-type="INTEGER"
primarykey="true"
autoincrement="true"
sequence-name="ackSequence"
/>
```

- Let OJB do that job for you:  
The `SequenceManagerNextValImpl` implementation create the sequence in database



automatic if none was found. If you don't want to declare a `sequence-name` attribute in your *field-descriptor*, you can enable an automatic sequence name building by setting a specific *custom-attribute*, then `SequenceManagerNextValImpl` build an internal sequence name if none was found.

```
<sequence-manager
  className="org.apache.ojb.broker.util.sequence.SequenceManagerNextValImpl">
  <attribute attribute-name="autoNaming" attribute-value="true"/>
</sequence-manager>
```

More about [sequence manager here](#).

## Second: Enable optimistic locking

You need to secure the data at the database. Thomas Mahler (lead OJB developer and considerable ORM guru) recommended in one email that you use the Optimistic Locking feature that is provided by OJB when using the **PB API** in a clustered environment. Sounds good to me. To do this you need to do three small steps:

### Note:

When using one of the top-level API in most cases Pessimistic (Object) Locking is supported. In that case it is recommended to use a *distributed lock management* instead of optimistic locking. More information about [ODMG API and Locking here](#).

- Add a database column to your table that is either an INTEGER or a TIMESTAMP
- Add the field to your java class, and getter/setter methods (depends on the used [PersistentField implementation](#)):

```
private Integer ackOptLock;

public Integer getAckOptLock()
{
  return ackOptLock;
}

public void setAckOptLock(Integer ackOptLock)
{
  this.ackOptLock = ackOptLock;
}
```

- Add the column to your table in your repository:

```
<field-descriptor
  name="ackOptLock"
  column="ACKOPTLOCK"
  jdbc-type="INTEGER"
  locking="true"/>
```

Now OJB will handle the locking for you. No explicit transactional code necessary!

## Do The Cache

The detailed steps to setup the *OSCache* can be found in [caching document](#)

**You're ready to go!** Now just create two instances of your application and see how they communicate at the cache level. Very cool.

### 5.6.6.3. Notes

- For J2EE/Servlet users: I have tested this on a number of different application servers. If you're having problems with your engine, post an email to the OJB Users mail list.
- OSCache also supports JMS for clustering here, which I haven't covered. If you either don't have access to a multicast network, or just plain like JMS, please refer to the OSCache

documentation for help with that, see [OSCache Clustering with JMS](#)).

- I have also tested this with Tangosol Coherence. Please refer to this Blog entry for that setup: [Coherence Setup](#)
- OJB also has ships with JCS. Feel free to try that one out on your own.

## 5.6.7. HOWTO - Stored Procedure Support

### 5.6.7.1. Introduction

OJB supports the use of stored procedures to handle the basic DML operations (INSERT, UPDATE, and DELETE). This document will describe the entries that you'll need to add to your repository in order to get OJB to utilize stored procedures instead of 'traditional' INSERT, UPDATE or DELETE statements.

Please note that there will be references to 'stored procedures' throughout this document. However, this is just a simplification for the purposes of this document. Any place you see a reference to 'stored procedure', you can assume that either a stored procedure or function can be used.

Information presented in this document includes the following:

- [Basic repository entries](#)
- [Common attributes for all procedure descriptors](#)
- An overview of the [insert procedure](#), [update procedure](#) and [delete procedure](#) descriptors.
- Information about the [argument descriptors](#) that are supported for all procedure
- A [simple example](#) and a [more complex example](#)

### 5.6.7.2. Repository entries

For any persistable class (i.e. "com.myproject.Customer") where you want to utilize stored procedures to handle persistence operations instead of traditional DML statements (i.e. INSERT, UPDATE or DELETE), you will need to include one or more of the following descriptors within the corresponding class-descriptor for the persistable class:

- `insert-procedure` - identifies the stored procedure that is to be used whenever a class needs to be inserted into the database.
- `update-procedure` - identifies the stored procedure that is to be used whenever a class needs to be updated in the database.
- `delete-procedure` - identifies the stored procedure that is to be used whenever a class needs to be removed from the database.

All of these descriptors must be nested within the class-descriptor that they apply to. Here is an example of a simple class-descriptor that includes each of the procedure descriptors listed above:

```
<class-descriptor class="com.myproject.Customer" table="CUSTOMER">
<field-descriptor column="ID" jdbc-type="DECIMAL" name="id" primarykey="true"/>
<field-descriptor column="NAME" jdbc-type="VARCHAR" name="name"/>
<insert-procedure name="CUSTOMER_PKG.ADD">
  <runtime-argument field-ref="id" return="true"/>
  <runtime-argument field-ref="name"/>
</insert-procedure>
<update-procedure name="CUSTOMER_PKG.CHG">
  <runtime-argument field-ref="id"/>
  <runtime-argument field-ref="name"/>
</update-procedure>
<delete-procedure name="CUSTOMER_PKG.CHG">
  <runtime-argument field-ref="id"/>
</delete-procedure>
</class-descriptor>
```

### 5.6.7.3. Common attributes

All three procedure descriptors have the following attributes in common:

- `name` - This is the name of the stored procedure that is to be used to handle the specific persistence operation.
- `return-field-ref` - This identifies the field in the class where the return value from the stored procedure will be stored. If this attribute is blank or not specified, then OJB will assume that the stored procedure does not return a value and will format the SQL command accordingly.

The basic syntax that is used to call a procedure that has a return value looks something like this:

```
{?= call <procedure-name>[<arg1>,<arg2>, ...]}
```

The basic syntax that is used to call a procedure that **does not** include a return value looks something like this:

```
{call <procedure-name>[<arg1>,<arg2>, ...]}
```

When OJB assembles the SQL to call a stored procedure, it will use the value of the 'name' attribute in place of 'procedure-name' in these two examples.

In addition, if the procedure descriptor includes a value in the 'return-field-ref' attribute that is 'valid', then the syntax that OJB builds will include the placeholder for the result parameter.

The previous section referred to the idea of a 'valid' value in the 'return-field-ref' attribute. A value is considered to be 'valid' if it meets the following criteria:

- The value is not blank
- There is a field-descriptor with a 'name' that matches the value in the 'return-field-ref' attribute.

If the 'return-field-ref' attribute is not 'valid', then the placeholder for the result parameter will not be included in the SQL that OJB assembles.

### 5.6.7.4. insert-procedure

The insert-procedure descriptor identifies the stored procedure that should be used whenever a class needs to be inserted into the database. In addition to the [common attributes](#) listed earlier, the insert-procedure includes the following attribute:

- `include-all-fields`

This attribute provides an efficient mechanism for passing all attributes of a persistable class to a stored procedure. If this attribute is set to true, then OJB will ignore any nested [argument descriptors](#). Instead, OJB will assume that the argument list for the stored procedure includes arguments for all attributes of the persistable class and that those arguments appear in the same order as the field-descriptors for the persistable class.

The default value for this attribute is 'false'.

#### Note:

If the field-descriptors in your repository do not 'align' exactly with the argument list for the stored procedure, or you want to maintain explicit control over the values that are passed to the stored procedure, then either set the 'include-all-fields' attribute to 'false' or leave it off the insert-procedure descriptor.

### 5.6.7.5. update-procedure

The update-procedure descriptor identifies the stored procedure that should be used whenever a

class needs to be updated in the database. In addition to the [common attributes](#) listed earlier, the update-procedure includes the following attribute:

- [include-all-fields](#)

This attribute provides the same capabilities and has the same caveats as the [include-all-fields](#) attribute on the [insert-procedure](#) descriptor.

#### 5.6.7.6. delete-procedure

The delete-procedure descriptor identifies the stored procedure that should be used whenever a class needs to be deleted from the database. In addition to the [common attributes](#) listed earlier, the delete-procedure includes the following attribute:

- [include-pk-only](#)  
This attribute provides an efficient mechanism for passing all of the attributes that make up the primary key for a persistable class to the specified stored procedure. If this attribute is set to true, then OJB will ignore any nested [argument descriptors](#). Instead, OJB will assume that the argument list for the stored procedure includes arguments for all attributes that make up the primary key for the persistable class (i.e. those field-descriptors where the 'primary-key' attribute is set to 'true'). OJB will also assume that those arguments appear in the same order as the corresponding field-descriptors for the persistable class. The default value for this attribute is 'false'.

#### Note:

If the field-descriptors in your repository that make up the primary key for a persistable class do not 'align' exactly with the argument list for the stored procedure, or you want to maintain explicit control over the values that are passed to the stored procedure, then either set the 'include-pk-only' attribute to 'false' or leave it off the delete-procedure descriptor.

#### 5.6.7.7. Argument descriptors

Argument descriptors are the mechanism that you will use to tell OJB two things:

1. How many placeholders should be included in the argument list for a stored procedure?
2. What value should be passed for each of those arguments?

There are two types of argument descriptors that can be defined in the repository:

- [runtime arguments](#) used to set a stored procedure argument equal to a value that is only known at runtime.
- [constant arguments](#) used to set a stored procedure argument equal to constant value.

You may notice that there is no argument descriptor specifically designed to pass a *null* value to the procedure. This capability is provided by the [runtime argument](#) descriptor.

The argument descriptors are essentially the 'mappings' between stored procedure arguments and their runtime values. Each procedure descriptor can include 0 or more argument descriptors in its definition.

After reading that last comment, you may wonder why OJB allows you to configure a procedure descriptor with no argument descriptors since the primary focus of OJB is to handle object persistence. How could OJB perform any sort persistence operation using a stored procedure that did not involve the passage of at least one value to the stored procedure? To be honest, it is extremely unlikely that you would ever set up a procedure descriptor with no argument descriptors. However, since there is no minimum number of arguments required for a stored procedure, we did not want to implement within OJB a requirement on the number of arguments that was more restrictive than the limits imposed by most/all database systems.

### runtime-argument descriptors

A runtime-argument descriptor is used to set a stored procedure argument equal to a value that is only known at runtime.

Two attributes can be specified for each runtime-argument descriptor:

- **field-ref**  
The 'field-ref' attribute identifies the specific field descriptor that will provide the argument's value. If this attribute is not specified or does not resolve to a valid field-descriptor, then a null value will be passed to the stored procedure.
- **return**  
The 'return' attribute is used to determine if the argument is used by the stored procedure as an 'output' argument. If this attribute is set to true, then the corresponding argument will be registered as an output parameter. After execution of the stored procedure, the value of the argument will be 'harvested' from the CallableStatement and stored in the attribute identified by the field-ref attribute. If this attribute is not specified or set to false, then OJB assumes that the argument is simply an 'input' argument, and it will do nothing special to the argument.

### constant-argument descriptors

A constant-argument descriptor is used to set a stored procedure argument equal to constant value.

There is one attribute that can be specified for the constant-argument descriptor:

- **value**  
The 'value' attribute identifies the value for the argument.

### 5.6.7.8. A simple example

This section provides background information and a simple example that illustrates how OJB's support for stored procedures can be utilized.

The background information covers the following topics:

- [The basic requirements](#)
- [The database objects](#) including the [table](#) that will be manipulated, the [sequence](#) that will be used by the stored procedures to assign primary key values, the [insert and update triggers](#) that maintain the four 'audit' columns and the [package](#) that provides the stored procedures that will handle the persistence operations.

Click [here](#) to skip the background information and go straight to the implementation.

### The basic requirements

These are the requirements that must be satisfied by our example

1. All insert, update and delete operations are to be performed by stored procedures.
2. All primary key values are to be by the stored procedure that handles the insert operation. The value that is assigned should be reflected in the object that 'triggered' the insert operation.
3. For auditing purposes, all tables will include the following set of columns:
  - USER\_CREATED - This will contain the id of the user who created the record
  - DATE\_CREATED - The date on which the record was created
  - USER\_UPDATED - The id of the user who last modified the record
  - USER\_UPDATED - The date on which the record was last modified

In addition to the inclusion of these columns on each table, the following requirements related to these columns had to be supported:

1. The values of the two date-related audit columns were to be maintained at the database level via insert and update triggers.
  - The insert trigger will set both DATE\_CREATED and DATE\_UPDATED to the current system date.
  - The update trigger will set DATE\_UPDATED to the current system date. The update trigger will also ensure that the original value of DATE\_CREATED is never modified.
2. The values of the two user-related audit columns are to be maintained at the database level via insert and update triggers.
  - The insert and update triggers will ensure that USER\_CREATED and USER\_UPDATED are appropriately populated.
  - The update trigger will ensure that the original value of USER\_CREATED is never modified.
3. Any changes that are made by the insert or update triggers to any of the four 'audit' columns had to be reflected in the object that caused the insert or update operation to occur.

### The database objects

The database objects that are described in this section utilize Oracle specific syntax. However, you should not infer from this that the stored procedure support provided by OJB can only be used to access data that is stored in an Oracle database. In reality, stored procedures can be used for persistence operations in any database that supports stored procedures.

- The [table](#) that will be manipulated,
- The [sequence](#) that will be used by the stored procedures to assign primary key values
- The [insert and update triggers](#) that maintain the four 'audit' columns
- The [package](#) that provides the stored procedures that will handle the persistence operations.

Click [here](#) to skip the information about the database objects and go straight to the implementation.

### The CUSTOMER table

This example will deal exclusively with persistence operations related to the a table named 'CUSTOMER' that is built using the following DDL:

```
CREATE TABLE CUSTOMER
( ID NUMBER(18) NOT NULL
, NAME VARCHAR2(50) NOT NULL
, USER_CREATED VARCHAR2(30)
, DATE_CREATED DATE
, USER_UPDATED VARCHAR2(30)
, DATE_UPDATED DATE
, CONSTRAINT PK_CUSTOMER PRIMARY KEY (ID)
);
```

### The sequence

This sequence will be used to assign unique values to CUSTOMER.ID.

```
CREATE SEQUENCE CUSTOMER_SEQ;
```

### The insert and update triggers

These two triggers will implement all of the requirements listed above that are related to the four audit columns:

```
CREATE OR REPLACE TRIGGER CUSTOMER_ITR
```

```

BEFORE INSERT ON CUSTOMER
FOR EACH ROW
BEGIN
--
-- Populate the audit dates
--
SELECT SYSDATE, SYSDATE
INTO :NEW.DATE_CREATED, :NEW.DATE_UPDATED
FROM DUAL;
--
-- Make sure the user created column is populated.
--
IF :NEW.USER_CREATED IS NULL
THEN
SELECT SYS_CONTEXT('USERENV','TERMINAL')
INTO :NEW.USER_CREATED
FROM DUAL;
END IF;
--
-- Make sure the user updated column is populated.
--
IF :NEW.USER_UPDATED IS NULL
THEN
SELECT SYS_CONTEXT('USERENV','TERMINAL')
INTO :NEW.USER_UPDATED
FROM DUAL;
END IF;
END;
/

CREATE OR REPLACE TRIGGER CUSTOMER_UTR
BEFORE UPDATE ON CUSTOMER
FOR EACH ROW
BEGIN
--
-- Populate the date updated
--
SELECT SYSDATE
INTO :NEW.DATE_UPDATED
FROM DUAL;
--
-- Make sure the user updated column is populated.
--
IF :NEW.USER_UPDATED IS NULL
THEN
SELECT SYS_CONTEXT('USERENV','TERMINAL')
INTO :NEW.USER_UPDATED
FROM DUAL;
END IF;
--
-- Make sure the date/user created are never changed
--
SELECT :OLD.DATE_CREATED, :OLD.USER_CREATED
INTO :NEW.DATE_CREATED, :NEW.USER_CREATED
FROM DUAL;
END;
/

```

### The package

This Oracle package will handle all INSERT, UPDATE and DELETE operations involving the CUSTOMER table.

```

CREATE OR REPLACE PACKAGE CUSTOMER_PKG AS
--
-- This procedure should be used to add a record to the CUSTOMER table.
--
PROCEDURE ADD ( AID          IN OUT CUSTOMER.ID%TYPE
               , ANAME       IN      CUSTOMER.NAME%TYPE
               , AUSER_CREATED IN OUT CUSTOMER.USER_CREATED%TYPE

```

```

        , ADATE_CREATED IN OUT CUSTOMER.DATE_CREATED%TYPE
        , AUSER_UPDATED IN OUT CUSTOMER.USER_UPDATED%TYPE
        , ADATE_UPDATED IN OUT CUSTOMER.DATE_UPDATED%TYPE );
--
-- This procedure should be used to change a record on the CUSTOMER table.
--
PROCEDURE CHANGE ( AID          IN      CUSTOMER.ID%TYPE
                  , ANAME       IN      CUSTOMER.NAME%TYPE
                  , AUSER_CREATED IN OUT CUSTOMER.USER_CREATED%TYPE
                  , ADATE_CREATED IN OUT CUSTOMER.DATE_CREATED%TYPE
                  , AUSER_UPDATED IN OUT CUSTOMER.USER_UPDATED%TYPE
                  , ADATE_UPDATED IN OUT CUSTOMER.DATE_UPDATED%TYPE );
--
-- This procedure should be used to delete a record from the CUSTOMER table.
--
PROCEDURE DELETE ( AID IN CUSTOMER.ID%TYPE );
END CUSTOMER_PKG;
/
CREATE OR REPLACE PACKAGE BODY CUSTOMER_PKG AS
--
-- This procedure should be used to add a record to the CUSTOMER table.
--
PROCEDURE ADD ( AID          IN OUT CUSTOMER.ID%TYPE
              , ANAME       IN      CUSTOMER.NAME%TYPE
              , AUSER_CREATED IN OUT CUSTOMER.USER_CREATED%TYPE
              , ADATE_CREATED IN OUT CUSTOMER.DATE_CREATED%TYPE
              , AUSER_UPDATED IN OUT CUSTOMER.USER_UPDATED%TYPE
              , ADATE_UPDATED IN OUT CUSTOMER.DATE_UPDATED%TYPE )
IS
    NEW_SEQUENCE_1 CUSTOMER.ID%TYPE;
BEGIN
    SELECT CUSTOMER_SEQ.NEXTVAL
        INTO NEW_SEQUENCE_1
        FROM DUAL;
    INSERT INTO CUSTOMER ( ID, NAME, USER_CREATED, USER_UPDATED )
        VALUES ( NEW_SEQUENCE_1, ANAME, AUSER_CREATED, AUSER_UPDATED )
        RETURNING ID, USER_CREATED, DATE_CREATED, USER_UPDATED, DATE_UPDATED
            INTO AID, AUSER_CREATED, ADATE_CREATED, AUSER_UPDATED, ADATE_UPDATED;
END ADD;
--
-- This procedure should be used to change a record on the CUSTOMER table.
--
PROCEDURE CHANGE ( AID          IN      CUSTOMER.ID%TYPE
                  , ANAME       IN      CUSTOMER.NAME%TYPE
                  , AUSER_CREATED IN OUT CUSTOMER.USER_CREATED%TYPE
                  , ADATE_CREATED IN OUT CUSTOMER.DATE_CREATED%TYPE
                  , AUSER_UPDATED IN OUT CUSTOMER.USER_UPDATED%TYPE
                  , ADATE_UPDATED IN OUT CUSTOMER.DATE_UPDATED%TYPE )
IS
BEGIN
    UPDATE CUSTOMER
        SET NAME          = ANAME
          , USER_CREATED = USER_CREATED
          , USER_UPDATED = AUSER_UPDATED
        WHERE ID          = AID
        RETURNING USER_CREATED, DATE_CREATED, USER_UPDATED, DATE_UPDATED
            INTO AUSER_CREATED, ADATE_CREATED, AUSER_UPDATED, ADATE_UPDATED;
END CHANGE;
--
-- This procedure should be used to delete a record from the CUSTOMER table.
--
PROCEDURE DELETE ( AID IN CUSTOMER.ID%TYPE )
IS
BEGIN
    DELETE
        FROM CUSTOMER
        WHERE ID          = AID;
END DELETE;
END CUSTOMER_PKG;
/

```



Please note the following about the structure of the CUSTOMER\_PKG package:

- The AID argument that is passed to the the ADD procedure is defined as IN OUT. This allows the procedure to return the newly assigned ID to the caller.
- In the ADD and CHANGE procedures, the arguments that correspond to the four 'audit' columns are defined as IN OUT. This allows the procedure to return the current value of these columns to the 'caller'.

### The implementation

Getting OJB to utilize the stored procedures described earlier in this document is as simple as adding a few descriptors to the repository. Here is a class-descriptor related to the CUSTOMER table that includes all of the necessary descriptors.

```
<class-descriptor class="com.myproject.Customer" table="CUSTOMER">
  <field-descriptor column="ID" jdbc-type="DECIMAL" name="id"
primarykey="true"/>
  <field-descriptor column="NAME" jdbc-type="VARCHAR" name="name"/>
  <field-descriptor column="USER_CREATED" jdbc-type="VARCHAR"
name="userCreated"/>
  <field-descriptor column="DATE_CREATED" jdbc-type="TIMESTAMP"
name="dateCreated"/>
  <field-descriptor column="USER_UPDATED" jdbc-type="VARCHAR"
name="userUpdated"/>
  <field-descriptor column="DATE_UPDATED" jdbc-type="TIMESTAMP"
name="dateUpdated"/>
  <insert-procedure name="CUSTOMER_PKG.ADD">
    <runtime-argument field-ref="id" return="true"/>
    <runtime-argument field-ref="name"/>
    <runtime-argument field-ref="userCreated" return="true"/>
    <runtime-argument field-ref="dateCreated" return="true"/>
    <runtime-argument field-ref="userUpdated" return="true"/>
    <runtime-argument field-ref="dateUpdated" return="true"/>
  </insert-procedure>
  <update-procedure name="CUSTOMER_PKG.CHG">
    <runtime-argument field-ref="id"/>
    <runtime-argument field-ref="name"/>
    <runtime-argument field-ref="userCreated" return="true"/>
    <runtime-argument field-ref="dateCreated" return="true"/>
    <runtime-argument field-ref="userUpdated" return="true"/>
    <runtime-argument field-ref="dateUpdated" return="true"/>
  </update-procedure>
  <delete-procedure name="CUSTOMER_PKG.CHG">
    <runtime-argument field-ref="id"/>
  </delete-procedure>
</class-descriptor>
```

Some things to note about this class-descriptor:

1. In the insert-procedure descriptor, the first runtime-argument descriptor corresponds to the "AID" argument that is passed to the CUSTOMER\_PKG.ADD routine. The "return" attribute on this runtime-argument is set to "true". With this configuration, OJB will 'harvest' the value that is returned by the CUSTOMER\_PKG.ADD stored procedure and store the value in the "id" attribute on the com.myproject.Customer class.
2. In both the insert-procedure and update-procedure descriptors, the runtime-argument descriptors that correspond to the four 'audit' columns all have the "return" argument set to "true". This allows any updates that are made by the procedure or the insert/update triggers to be reflected in the "Customer" object that caused the insert/update operation to occur.

### 5.6.7.9. A complex example

This example builds upon the [simple example](#) that was presented earlier by introducing some additional requirements beyond those that were specified in the simple example. Some of these

additional requirements may seem a little contrived. To be honest, they are. The only purpose of these additional requirements is to create situations that illustrate how the additional capabilities provided by OJB's support for stored procedures can be utilized.

The additional requirements for this example include the following:

- All procedures will include two additional arguments. These two new arguments will be added to the end of the argument list for all existing procedures.
  - ASOURCE\_SYSTEM - identifies the system that initiated the persistence operation. This will provide a higher level of audit tracking capability. In our example, this will always be "SAMPLE".
  - ACOST\_CENTER - identifies the 'cost center' that should be charged for the persistence operation. In our example, this argument will always be null.
- For all "ADD" and "CHG" stored procedures, the value that was assigned to the "DATE\_UPDATED" column will no longer be returned to the caller via an "IN OUT" argument. Instead, it will be returned to the caller via the procedure's return value.

Based on these new requirements, the class-descriptor for the "com.myproject.Customer" class will look like this. The specific changes are detailed below.

```
<class-descriptor class="com.myproject.Customer" table="CUSTOMER">
  <field-descriptor column="ID" jdbc-type="DECIMAL" name="id"
primaryKey="true"/>
  <field-descriptor column="NAME" jdbc-type="VARCHAR" name="name"/>
  <field-descriptor column="USER_CREATED" jdbc-type="VARCHAR"
name="userCreated"/>
  <field-descriptor column="DATE_CREATED" jdbc-type="TIMESTAMP"
name="dateCreated"/>
  <field-descriptor column="USER_UPDATED" jdbc-type="VARCHAR"
name="userUpdated"/>
  <field-descriptor column="DATE_UPDATED" jdbc-type="TIMESTAMP"
name="dateUpdated"/>
  <insert-procedure name="CUSTOMER_PKG.ADD"
  return-field-ref="dateUpdated"> <!-- See note 1 -->
    <runtime-argument field-ref="id" return="true"/>
    <runtime-argument field-ref="name"/>
    <runtime-argument field-ref="userCreated" return="true"/>
    <runtime-argument field-ref="dateCreated" return="true"/>
    <runtime-argument field-ref="userUpdated" return="true"/>
    <runtime-argument field-ref="dateUpdated"/> <!-- See note 2 -->
    <constant-argument value="SAMPLE"/> <!-- See note 3 -->
    <runtime-argument/> <!-- See note 4 -->
  </insert-procedure>
  <update-procedure name="CUSTOMER_PKG.CHG"
  return-field-ref="dateUpdated"> <!-- See note 1 -->
    <runtime-argument field-ref="id"/>
    <runtime-argument field-ref="name"/>
    <runtime-argument field-ref="userCreated" return="true"/>
    <runtime-argument field-ref="dateCreated" return="true"/>
    <runtime-argument field-ref="userUpdated" return="true"/>
    <runtime-argument field-ref="dateUpdated"/> <!-- See note 2 -->
    <constant-argument value="SAMPLE"/> <!-- See note 3 -->
    <runtime-argument/> <!-- See note 4 -->
  </update-procedure>
  <delete-procedure name="CUSTOMER_PKG.CHG">
    <runtime-argument field-ref="id"/>
    <constant-argument value="SAMPLE"/> <!-- See note 3 -->
    <runtime-argument/> <!-- See note 4 -->
  </delete-procedure>
</class-descriptor>
```

Here are an explanation of each modification:

- **Note 1:** The value that is returned by the "ADD" and "CHG" stored procedures will now be stored in the "dateUpdated" attribute on the "com.myproject.Customer" class.

- **Note 2:** Since the `ADATE_UPDATED` argument is no longer defined as an "IN OUT" argument, we have removed the "return" attribute from the corresponding runtime-argument descriptor.
- **Note 3:** This is the first of two new arguments that were added to the argument list of each procedure. This argument represents the 'source system', the system that initiated the persistence operation. In our example, we will always pass a value of 'SAMPLE'.
- **Note 4:** This is the second of two new arguments that were added to the argument list of each procedure. This argument represents the 'cost center' that should be charged for the persistence operation. In our example, we have no cost center, so we need to pass a null value. This is accomplished by including a 'runtime-argument' descriptor that has no 'field-ref' specified.

## 5.6.8. HOWTO - Spring with OJB

### 5.6.8.1. Spring with OJB

This Howto refers to external articles describe the usage of [Spring](#) with OJB's [PB-api](#).

Included are three articles that are in-depth How-to's for using Spring Framework's *declarative transaction engine* with OJB's [Persistence Broker](#) implementation. For each article, I have included a complete sample application with an embedded database that you can run.

They are simple to setup and require minimal configuration.

#### Spring, OJB, and Struts, getting started

The first article covers a basic configuration and setup for integrating Spring with OJB. It covers:

- Connecting to the database using Spring's local datasources (*LocalOjbConfigurer*)
- Register the appropriate transactionManager for use with OJB (*PersistenceBrokerTransactionManager*)
- Creating beans and interfaces for use with the database
- Having Spring handle transactions for the beans in declarative fashion

The article link is: [Spring, OJB, and Struts, getting started](#)

#### Spring, OJB, and Struts. Version 2 with DAO.

The second article covers more advanced Spring usage, separating out the database implementation code from your Data Access Objects. This is an important abstraction layer for controlling data access ubiquitously.

The article link is: [Spring, OJB, and Struts. Version 2 with DAO.](#)

#### Spring, OJB, and Struts. Version 3: JNDI Datasources and Caching

The Third article covers using OJB's datasourcing and caching strategies.

The first two articles require local datasources and [Per Broker caching](#) because of the way that the *OjbConfigurer* works.

Now that we've completed basics of integrating OJB in Spring, this last entry shows how to use JNDI datasourcing with OJB and Spring, which in turn allows users to get back to other OJB caching strategies.

The article link is: [Spring, OJB, and Struts. Version 3: JNDI Datasources and Caching](#)

## 5.7. Testing

## 5.7.1. Testing Summary

Summary of documentation about OJB Unit testing.

OJB has an extensive [JUnit](#)-based test suite with over 800 individual tests.

- [The OJB test suite](#)  
Information about how to run the test suite against your RDBMS.
- [How to write tests](#)  
Explains how to add your own tests to the OJB test suite.
- [The OJB performance tests](#)  
Link to the reference guide about OJB performance comparison and the OJB performance test-suite.

## 5.7.2. JUnit Test Suite

### 5.7.2.1. Introduction

Building an Object/Relational mapping tool with support for multiple API's is really error prone. To create a solid and stable software, the most awful thing in programmers life has to be done - Testing.

Quality assurance taken seriously! OJB and provide specific tests for each supported API. Currently more than 800 test cases for regression tests exist. As testing framework [JUnit](#) was used.

#### Where can I find the test sources?

The test sources of the *OJB Test-Suite* can be find under **[db-ojb]/src/test/org/apache/ojb**.

It's also possible to browse the test sources online using the [OJB SVN view](#). The test directory can be found here: [\[db-ojb\]/src/test/org/apache/ojb](#).

### 5.7.2.2. How to run the Test Suite

If the [platform depended settings](#) are done, the test suite can be started with the ant target:

```
ant junit
```

If compiling of the sources should be skipped use

```
ant junit-no-compile
```

If you did not manage to set up the target database with the ant `prepare-testdb` you can use

```
ant junit-no-compile-no-prepare
```

to run the testsuite without generation of the test database (and without compiling).

After running the regression tests you should see a console output similar to this:

```
junit-no-compile-no-prepare:
[junit] Running org.apache.ojb.broker.AllTests
[junit] Tests run: 620, Failures: 0, Errors: 0, Time elapsed: 81,75 sec
[junit] Running org.apache.ojb.odmg.AllTests
[junit] Tests run: 183, Failures: 0, Errors: 0, Time elapsed: 21,719 sec
[junit] Running org.apache.ojb.soda.AllTests
[junit] Tests run: 3, Failures: 0, Errors: 0, Time elapsed: 7,641 sec
[junit] Running org.apache.ojb.otm.AllTests
[junit] Tests run: 79, Failures: 0, Errors: 0, Time elapsed: 28,266 sec
```

```
junit-no-compile-no-prepare-selected:
junit-no-compile:
junit:
BUILD SUCCESSFUL
Total time: 3 minutes 26 seconds
```

We aim at shipping that releases have no failures and errors in the regression tests! If the JUnit tests report errors or failures something does not work properly! There may be several reasons:

- You made a mistake in configuration (OJB was shipped with settings pass all tests). See [platform](#), [OJB.properties](#), [repository file](#), .
- Your database doesn't support specific features used by the tests
- Evil hex
- Bug in OJB

JUnit writes a *log-file* for each tested API. You can find the logs under `[db-ojb]/target/test`. The log files named like `tests-XXX.txt`. The test logs show in detail what's going wrong.

In such a case please check again if you followed all the above steps. If you still have problems you might post a request to the OJB user mailinglist.

#### How to run the test-suite with a different database than OJB default DB

Basically all you have to do is:

- Get source version of OJB or fetch OJB from CVS (take care of branches, branch OJB\_1\_0\_RELEASE represents OJB 1.0.x).
- Adapt the profile file of your database under `[db-ojb]/profile/yourDB.profile` and set user, password, ...
- In `[db-ojb]/build.properties` file comment the "profile=hsqldb" line and uncomment the "#profile=yourDB" line.
- Drop `jdo.jar` and your database driver into `[db-ojb]/lib` directory.
- Drop `junit.jar` into your `...ant/lib` folder.
- Make sure that your database allows at least 20 concurrent connections.

Then follow the steps described [above](#).

#### 5.7.2.3. What about known issues?

All major known issues are listed in the [release-notes](#) file.

The tests reproduce open bugs will be skipped on released OJB versions. It is possible to enable these tests to see all failing test cases of the shipped version by changing a flag in `[db-ojb]/build.properties` file:

```
###
# If 'true', junit tests marked as known issue in the junit-test
# source code (see OJBTestCase class for more detailed info) will be
# skipped. Default value is 'true'. For development 'false' is recommended,
# because this will show unsolved problems.
OJB.skip.issues=true
```

#### 5.7.2.4. Donate own tests for OJB Test Suite

Details about [donate own test to OJB you can find here](#).

### 5.7.3. Write Tests

#### 5.7.3.1. Introduction

As described in the [test suite](#) document, OJB emphasizes on quality assurance and provide a huge test suite. But of course it is impossible to cover all parts of OJB with unit tests and OJB will never be perfect (although we would like to think it's s nearly perfect ;-)), thus if you are missing a testcase or think you found an bug -- don't hesitate to write your own test and send it to the [developer list](#) or, if you have an existing issue report, attach it in the [issue tracker](#).

### 5.7.3.2. How to write a new Test

Before starting to write your own test case, please pay attention to these rules.

#### The Test Class

All test classes have to inherit from `org.apache.ojb.junit.OJBTestCase` and have to provide a static main method to start the Junit test:

```
public class MyTest extends OJBTestCase
{
    public static void main(String[] args)
    {
        String[] arr = {MyTest.class.getName()};
        junit.textui.TestRunner.main(arr);
    }

    public void testMyFirstOne()
    {
        ....
    }
}
```

You will find some test classes for specific scenarios in the `org.apache.ojb.junit` package:

- `org.apache.ojb.junit.PBTestCase` - provides a `org.apache.ojb.broker.PersistenceBroker` instance for tests.
- `org.apache.ojb.junit.ODMGTestCase` - provides `org.odmg.Implementation` and `org.odmg.Database` instances for tests.
- `org.apache.ojb.junit.JUnitExtensions` - servers as a base class when writing multi-threaded test classes.

For more info, see the JavaDoc of the class.

A test case for the PB-API may look like:

```
public class ReferenceRuntimeSettingTest extends PBTestCase
{
    public static void main(String[] args)
    {
        String[] arr = {ReferenceRuntimeSettingTest.class.getName()};
        junit.textui.TestRunner.main(arr);
    }

    public void testChangeReferenceSetting()
    {
        ClassDescriptor cld = broker.getClassDescriptor(MainObject.class);
        // and so on
        ....
    }
}
```

#### Note:

The `PersistenceBroker` cleanup is done by `PBTestCase`.

#### Persistent Objects used by Test

We recommend to introduce separate persistent objects for each `TestCase` class. In the [test suite](#)

two concepts are used:

- Include your persistent objects as *public static classes* in your test class.
- Separate your test class in an independent package and include the test case and all persistent object classes in this new package.

### Test Class Metadata

Currently all test specific object metadata (class-descriptor used for tests) are shared among several xml files. The naming convention is `repository_junit_XXX.xml`. Thus metadata for new tests should be included in one of the existing junit repository (sub) files or written in an new separate one and included in the main repository file.

```
<!DOCTYPE descriptor-repository PUBLIC
    "-//Apache Software Foundation//DTD OJB Repository//EN"
    "repository.dtd"
[
<!ENTITY database SYSTEM "repository_database.xml">
<!ENTITY internal SYSTEM "repository_internal.xml">
<!ENTITY user SYSTEM "repository_user.xml">

<!-- Start of JUnit included files -->
<!ENTITY junit SYSTEM "repository_junit.xml">
<!ENTITY junit_odmg SYSTEM "repository_junit_odmg.xml">
<!ENTITY junit_otm SYSTEM "repository_junit_otm.xml">
<!ENTITY junit_ref SYSTEM "repository_junit_reference.xml">
<!ENTITY junit_meta_seq SYSTEM "repository_junit_meta_seq.xml">
<!ENTITY junit_model SYSTEM "repository_junit_model.xml">
<!ENTITY junit_cloneable SYSTEM "repository_junit_cloneable.xml">

<!-- Your entity here: -->
<!ENTITY junit_myfirstttest SYSTEM "repository_junit_myfirstttest.xml">
]>
<descriptor-repository version="1.0" isolation-level="read-uncommitted"
    proxy-prefetching-limit="50">

    <!-- include all used database connections -->
    &database;

    <!-- include ojb internal mappings here -->
    &internal;

    <!-- include user defined mappings here -->
    &user;

    <!-- include mappings for JUnit tests -->
    &junit;
    &junit_odmg;
    &junit_otm;
    &junit_ref;
    &junit_meta_seq;
    &junit_model;
    &junit_cloneable;

    &junit_myfirstttest;
```