

Repository File

by Thomas Mahler, Daren Drummond, Brian McCallister, Armin Waibel, Thomas Dudziak

Table of contents

1 Introduction - repository syntax.....	3
2 descriptor-repository.....	3
2.1 Elements.....	3
2.2 Attributes.....	4
2.2.1 version.....	4
2.2.2 isolation-level.....	4
2.2.3 proxy-prefetching-limit.....	4
3 jdbc-connection-descriptor.....	5
3.1 Elements.....	5
3.2 Attributes.....	5
3.2.1 jdbcAlias.....	6
3.2.2 default-connection.....	6
3.2.3 platform.....	6
3.2.4 jdbc-level.....	6
3.2.5 eager-release.....	7
3.2.6 batch-mode.....	7
3.2.7 useAutoCommit.....	7
3.2.8 ignoreAutoCommitExceptions.....	7
3.2.9 jndi-datasource-name.....	7
3.2.10 username.....	7
4 connection-pool.....	8
5 sequence-manager.....	9

6 object-cache.....	10
7 custom attribute.....	11
8 class-descriptor.....	12
9 extent-class.....	14
10 field-descriptor.....	14
11 reference-descriptor.....	17
12 foreignkey.....	18
13 collection-descriptor.....	19
14 inverse-foreignkey.....	21
15 fk-pointing-to-this-class.....	22
16 fk-pointing-to-element-class.....	22
17 query-customizer.....	22
18 index-descriptor.....	22
19 index-column.....	23
20 Stored Procedure Support.....	23
20.1 insert-procedure.....	23
20.2 update-procedure.....	24
20.3 delete-procedure.....	24
20.4 runtime-argument.....	25
20.5 constant-argument.....	26

1. Introduction - repository syntax

The syntax of the OJB repository xml files is defined by the *repository.dtd*.

The [repository.dtd can be found here](#) (../repository.dtd.txt) .

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](#) (../repository.xml.txt) . Main file for metadata declaration. This file is split into several sub files using xml-Entity references.
2. [the repository_database.xml](#) (../repository_database.xml.txt) . This file contains the mapping information for database/connection handling.
3. [the repository_internal.xml](#) (../repository_internal.xml.txt) . 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](#) (../repository_user.xml.txt) . This file contains mappings for the tutorial applications and may be used to hold further user defined class mappings.
5. [the repository_junit.xml](#) (../repository_junit.xml.txt) . 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.

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](#) (../docu/guides/metadata.html#without-repository) .

2.1. 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* )
>
```

2.2. Attributes

```
<!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"
>
```

2.2.1. version

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.

2.2.2. isolation-level

The *isolation-level* attribute defines the default locking isolation level used by OJB's [pessimistic locking](http://../docu/guides/lockmanager.html#pessimistic-locking) (../docu/guides/lockmanager.html#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.

2.2.3. proxy-prefetching-limit

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 |
```

Repository File

```
        read-committed  
        repeatable-read  
        serializable  
        optimistic ) "read-uncommitted"  
    proxy-prefetching-limit CDATA "50"  
>
```

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](#) ([../docu/guides/connection.html](#)).

3.1. Elements

```
<!ELEMENT jdbc-connection-descriptor (documentation?, attribute*,  
    object-cache?, connection-pool?, sequence-manager?)>
```

The *object-cache* element specifies the [object-cache implementation](#) ([../docu/guides/objectcache.html](#)) 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? )  
>
```

3.2. Attributes

The *jdbc-connection-descriptor* element contains a bunch of required and implied attributes:

```
<!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
>

```

3.2.1. jdbcAlias

The *jdbcAlias* attribute is a shortcut name for the defined connection descriptor. OJB uses the *jdbc* alias as key for the defined connections.

3.2.2. default-connection

The *default-connection* attribute used to define if this connection should be used as default connection with OJB. You could define only one connection as 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* was not set at runtime, it is mandatory that *username and password* is set in repository file.

3.2.3. platform

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 was *Hsqldb*.

3.2.4. jdbc-level

Repository File

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 was *1.0*.

3.2.5. eager-release

The *eager-release* attribute was adopted to solve a problem occurred 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.

3.2.6. batch-mode

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 was enabled. On *PB.close()* OJB switches off batch mode, thus you have to do '*...setBatchMode(true)*' on each obtained PB instance again.

3.2.7. useAutoCommit

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 - Set autoCommit explicitly to *true* when a connection was created and temporarily set to *false* when necessary (default mode).
- 2 - Set autoCommit explicitly to *false* when a connection was created.

3.2.8. ignoreAutoCommitExceptions

If the *ignoreAutoCommitExceptions* attribute is set to *true*, all exceptions caused by setting autoCommit state, will be ignored. Default mode is *false*.

3.2.9. jndi-datasource-name

If a *jndi-datasource-name* for JNDI based lookup of Jdbc connections is specified, the four attributes *driver*, *protocol*, *subprotocol*, and *dbalias* used for Jdbc DriverManager based construction of Jdbc Connections must not be declared.

3.2.10. username

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](#) ([../../docu/faq.html](#)) .

4. connection-pool

The *connection-pool* element specifies the connection pooling parameter. More info about the [connection handling can be found here](#) ([../../docu/guides/connection.html](#)) .

```
<!ELEMENT connection-pool (documentation? )
>
```

Valid attributes for the *connection-pool* element are:

```
<!ATTLIST connection-pool
  maxActive           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
  logAbandoned        ( true | false ) #IMPLIED
  removeAbandoned    ( true | false ) #IMPLIED
  removeAbandonedTimeout CDATA #IMPLIED
>
```

maxActive is the maximum number of connections that can be borrowed from the pool at one time. When non-positive, there is no limit.

maxIdle controls the maximum number of connections that can sit idle in the pool at any time. When non-positive, there is no limit

maxWait - the maximum time block to get connection instance from pool, after that exception is thrown. When non-positive, block till last judgement

whenExhaustedAction

- 0 - fail when pool is exhausted
- 1 - block when pool is exhausted
- 2 - grow when pool is exhausted

Repository File

testOnBorrow when *true* the pool will attempt to validate each object before it is returned from the pool.

testOnReturn set to *true* will force the pool to attempt to validate each object before it is returned to the pool.

testWhileIdle indicates whether or not idle objects should be validated. Objects that fail to validate will be dropped from the pool.

timeBetweenEvictionRunsMillis indicates how long the eviction thread should sleep before "runs" of examining idle objects. When non-positive, no eviction thread will be launched.

minEvictableIdleTimeMillis specifies the minimum amount of time that a connection may sit idle in the pool before it is eligible for eviction due to idle time. When non-positive, no connection will be dropped from the pool due to idle time alone (depends on *timeBetweenEvictionRunsMillis* > 0)

numTestsPerEvictionRun - the number of connections to examine during each run of the idle object evictor thread (if any)

validationQuery allows to specify a validation query used by the *ConnectionFactory* implementations using connection pooling, to test a requested connection (e.g. "select 1 from dual") before leave the pool (used by *ConnectionFactoryDBCPIImpl* and *ConnectionFactoryPooledImpl*).

If not set, only *connection.isClosed()* will have been called before the connection was delivered.

logAbandoned is only supported when using *org.apache.obj.broker.accesslayer.ConnectionFactoryDBCPIImpl* *ConnectionFactory* implementation. Then it is a flag to log stack traces for application code which abandoned a *Statement* or *Connection*. Defaults to false. Logging of abandoned *Statements* and *Connections* adds overhead for every *Connection* open or new *Statement* because a stack trace has to be generated.
DEPRECATED attribute!

removeAbandoned and *removeAbandonedTimeout* When using *org.apache.obj.broker.accesslayer.ConnectionFactoryDBCPIImpl* *ConnectionFactory* implementation, the *removeAbandoned* flag controls the removal of abandoned connections if they exceed the *removeAbandonedTimeout*. Set to true or false, default false. 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.
DEPRECATED attributes!

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](#) ([../docu/guides/sequencemanager.html](#)) .

Use the [custom-attribute](#) element to pass implementation specific properties.

```
<!ELEMENT sequence-manager (
  documentation?,
  attribute* )
>
```

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 *SequenceManagerXXXImpl*

More info about the usage of the Sequence Manager implementations [can be found here](#). ([../docu/guides/sequencemanager.html](#))

```
<!ATTLIST sequence-manager
  className CDATA #REQUIRED
>
```

6. object-cache

The *object-cache* element can be used to specify the ObjectCache implementation used by OJB. There are three levels of declaration:

- in [OJB.properties](#) ([../OJB.properties.txt](#)) file, to declare the standard (default) ObjectCache implementation
- on jdbc-connection-descriptor level, to declare ObjectCache implementation on a per connection/user level
- on class-descriptor level, to declare ObjectCache implementation on a per class level

Note:

The priority of the declared object-cache elements are:
per class > per jdbc descriptor > standard

E.g. if you declare ObjectCache implementation 'my.cacheDef' as standard, set ObjectCache implementation 'my.cacheA' in class-descriptor for class A and class B does not declare an

Repository File

object-cache element. Then OJB use 'my.cacheA' as ObjectCache for class A and 'my.cacheDef' for class B.

```
<!ELEMENT object-cache (documentation?, attribute*)>
```

Use the [custom-attribute](#) element to pass implementation specific properties.

```
<!ATTLIST object-cache
  class CDATA #REQUIRED
>
```

Attribute 'class' specifies the full qualified class name of the used ObjectCache implementation.

7. custom attribute

An *attribute* element allows arbitrary name/value pairs to be represented in the repository. See the [repository.dtd](#) (../repository.dtd.txt) 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](http://api/org/apache/obj/broker/metadata/ClassDescriptor.html) ([../api/org/apache/obj/broker/metadata/ClassDescriptor.html](http://api/org/apache/obj/broker/metadata/ClassDescriptor.html)) . To read the attribute at runtime do:

```

// get the ClassDescriptor
ClassDescriptor cld = broker.getClassDescriptor(TestClass.class);
String value = cld.getAttribute("myAttribute");

```

8. class-descriptor

A *class-descriptor* and the associated java class [ClassDescriptor](http://api/org/apache/obj/broker/metadata/ClassDescriptor.html) ([../api/org/apache/obj/broker/metadata/ClassDescriptor.html](http://api/org/apache/obj/broker/metadata/ClassDescriptor.html)) encapsulate metadata information of an interface, abstract or concrete class.

For *interfaces* or *abstract* classes a *class-descriptor* holds a sequence of *extent-class* elements which specify the [types extending this class](http://docu/guides/advanced-technique.html#extents) ([../docu/guides/advanced-technique.html#extents](http://docu/guides/advanced-technique.html#extents)) .

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?,

```

Repository File

```
        update-procedure? ,  
        delete-procedure? )  
    )  
>
```

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](#) (../docu/guides/lockmanager.html#pessimistic-locking) api).

Note: This does NOT touch the jdbc-level of the connection.

Note:

The *isolation-level* does not touch the jdbc-connection isolation level. It's completely independent from the database connection setting.

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](#) (../docu/guides/basic-technique.html#using-proxy) .

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](#) (../docu/guides/advanced-technique.html#using-rowreader) implementation used to materialize instances of the persistent class.

The *extends* attribute *****TODO: description*****

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"
>
```

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
>
```

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

Repository File

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](#) ([../docu/guides/advanced-technique.html#persistent-field](#)) .

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](#) ([../docu/guides/jdbc-types.html](#)) .

The *primarykey* specifies if the column is a primary key column, default value is *false*.

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](#) ([../docu/guides/sequencemanager.html](#)) .

The *sequence-name* attribute can be used to state explicitly a sequence name used by the sequence manager implementations. Check the [javadocs](#) ([../api/index.html](#)) 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 was not set. More info about [sequence key generation here](#) ([../docu/guides/sequencemanager.html](#)) .

The *locking* attribute is set to *true* if the persistent attribute is used for *optimistic locking*. More about [optimistic locking](#) ([../docu/faq.html#optimisticLocking](#)) . The default value is *false*.

The *update-lock* 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.obj.accesslayer.conversions.FieldConversion`. A `FieldConversion` can be used to implement conversions between Java- attributes and database columns. More about [field conversion](#) (`../../docu/guides/jdbc-types.html#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 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](#) (`../../docu/guides/advanced-technique.html#anonymous-keys`) .

```
<!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"
>
```

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](#) (../docu/guides/basic-technique.html#one-to-one) .

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](#) (../docu/guides/advanced-technique.html#persistent-field) 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 was loaded from cache. If *true* OJB try to retrieve the reference (dependent on the [auto-xxx settings](#) (../docu/guides/basic-technique.html#cascading)) 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](#) (../docu/guides/objectcache.html) 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](#) (../docu/guides/objectcache.html) (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](#) (../docu/guides/basic-technique.html#auto-retrieve) .

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](#) ([../../docu/guides/basic-technique.html#cascading](#)) .

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](#) ([../../docu/guides/basic-technique.html#cascading](#)) .

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"
>
```

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

Repository File

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
>
```

13. collection-descriptor

A collection-descriptor contains mapping info for a `liCollection-` or `Array-`attribute of a persistent class that contains persistent entity Objects. See more about [1:n](#) ([../docu/guides/basic-technique.html#one-to-n](#)) and [m:n](#) ([../docu/guides/basic-technique.html#m-to-n](#)) references.

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](#) ([../docu/guides/advanced-technique.html#persistent-field](#)) .

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](#) ([../docu/guides/advanced-technique.html#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.

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](#) ([../docu/guides/basic-technique.html#using-proxy](#)) .

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 was loaded from cache. If *true* OJB try to retrieve the reference (dependent on the [auto-xxx settings](#) ([../docu/guides/basic-technique.html#cascading](#))) 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](#) ([../docu/guides/objectcache.html](#)) 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](#) ([../docu/guides/objectcache.html](#)) (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](#) ([../docu/guides/basic-technique.html#auto-retrieve](#)) .

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](#) ([../docu/guides/basic-technique.html#cascading](#)) .

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).

Repository File

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](#) ([../../docu/guides/basic-technique.html#cascading](#)) .

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"
>
```

14. 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
```

>

15. 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
>
```

16. 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
>
```

17. 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](#) ([../docu/guides/advanced-technique.html#query-customizer](#)).

Use the [custom-attribute](#) element to pass implementation specific properties.

```
<!ELEMENT query-customizer (
  documentation?,
  attribute*)>

<!ATTLIST query-customizer
  class CDATA #REQUIRED
>
```

18. index-descriptor

Repository File

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">
```

19. 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>
```

20. Stored Procedure Support

OJB supports stored procedures for insert, update and delete operations. [How to use stored procedures within OJB can be found here](#)

(../docu/howtos/howto-work-with-stored-procedures.html) .

20.1. 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"
>
```

20.2. 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"
>
```

20.3. delete-procedure

Identifies the procedure/function that should be used to handle deletions for a specific

Repository File

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"
>
```

20.4. 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"
>
```

20.5. 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
>
```