

Oracle® Identity Manager

Best Practices Guide

Release 9.0

B32139-01

October 2006

The Programs (which include both the software and documentation) contain proprietary information; they are provided under a license agreement containing restrictions on use and disclosure and are also protected by copyright, patent, and other intellectual and industrial property laws. Reverse engineering, disassembly, or decompilation of the Programs, except to the extent required to obtain interoperability with other independently created software or as specified by law, is prohibited.

The information contained in this document is subject to change without notice. If you find any problems in the documentation, please report them to us in writing. This document is not warranted to be error-free. Except as may be expressly permitted in your license agreement for these Programs, no part of these Programs may be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose.

If the Programs are delivered to the United States Government or anyone licensing or using the Programs on behalf of the United States Government, the following notice is applicable:

U.S. GOVERNMENT RIGHTS Programs, software, databases, and related documentation and technical data delivered to U.S. Government customers are "commercial computer software" or "commercial technical data" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the Programs, including documentation and technical data, shall be subject to the licensing restrictions set forth in the applicable Oracle license agreement, and, to the extent applicable, the additional rights set forth in FAR 52.227-19, Commercial Computer Software--Restricted Rights (June 1987). Oracle USA, Inc., 500 Oracle Parkway, Redwood City, CA 94065.

The Programs are not intended for use in any nuclear, aviation, mass transit, medical, or other inherently dangerous applications. It shall be the licensee's responsibility to take all appropriate fail-safe, backup, redundancy and other measures to ensure the safe use of such applications if the Programs are used for such purposes, and we disclaim liability for any damages caused by such use of the Programs.

Oracle, JD Edwards, PeopleSoft, and Siebel are registered trademarks of Oracle Corporation and/or its affiliates. Other names may be trademarks of their respective owners.

The Programs may provide links to Web sites and access to content, products, and services from third parties. Oracle is not responsible for the availability of, or any content provided on, third-party Web sites. You bear all risks associated with the use of such content. If you choose to purchase any products or services from a third party, the relationship is directly between you and the third party. Oracle is not responsible for: (a) the quality of third-party products or services; or (b) fulfilling any of the terms of the agreement with the third party, including delivery of products or services and warranty obligations related to purchased products or services. Oracle is not responsible for any loss or damage of any sort that you may incur from dealing with any third party.

Contents

Preface	v
Audience.....	v
Documentation Accessibility	v
Related Documents	vi
Documentation Updates	vi
Conventions	vi
 1 Using the Deployment Manager	
Features of the Deployment Manager.....	1-1
Limitations of the Deployment Manager	1-2
Export System Objects Only when Necessary.....	1-2
Export Related Groups of Objects	1-3
Group Definition Data and Operational Data Separately	1-3
Use Logical Naming Conventions for Versions of a Form	1-3
Export Root to Preserve a Complete Organizational Hierarchy.....	1-4
Provide Clear Export Descriptions.....	1-4
Check All Warnings Before Importing	1-4
Check Dependencies Before Exporting Data.....	1-4
Matching Scheduled Task Parameters	1-4
Compile Adapters and Enable Scheduled Tasks	1-4
Export Entity Adapters Separately.....	1-5
Group Permissions.....	1-5
Report Permissions	1-5
Back Up the Database.....	1-5
Import Data When the System Is Quiet.....	1-5
 2 JBoss Components That May be Removed	
Files and Directories to Remove.....	2-1
Non-Clustered Installations.....	2-1
Clustered Installations.....	2-2
 3 Tuning Oracle Database for Oracle Identity Manager	
Sample Instance Configuration Parameters.....	3-1
Physical Data Placement	3-2
Pinning Sequences and Stored Procedures in System Global Area (SGA)	3-3

Database Performance Monitoring.....	3-4
4 Managing the Cache	
Sample Cache Configuration	4-1
General Cache Configuration Properties	4-2
Category-Based Cache Configuration Properties	4-3
Class Reloading	4-4
Purging the Cache	4-4
Optimal Cache Configuration for a Production Environment.....	4-5
5 Securing Your Deployment	
Securing the Administrative and User Console.....	5-1
Securing the Self Registration and Change Password Pages	5-1

Index

Preface

This preface introduces you to the *Oracle Identity Manager Best Practices Guide* discussing the intended audience and conventions of this document. It also includes a list of related Oracle documents.

Note: This is a transitional release following Oracle's acquisition of Thor Technologies. Some parts of the product and documentation still refer to the original Thor company name and Xellerate product name and will be rebranded in future releases.

Audience

The *Oracle Identity Manager Best Practices Guide* is intended for Database Administrators, System Administrators, and developers who plan to use Oracle Identity Manager extensively in production environments.

Documentation Accessibility

Our goal is to make Oracle products, services, and supporting documentation accessible, with good usability, to the disabled community. To that end, our documentation includes features that make information available to users of assistive technology. This documentation is available in HTML format, and contains markup to facilitate access by the disabled community. Accessibility standards will continue to evolve over time, and Oracle is actively engaged with other market-leading technology vendors to address technical obstacles so that our documentation can be accessible to all of our customers. For more information, visit the Oracle Accessibility Program Web site at

<http://www.oracle.com/accessibility/>

Accessibility of Code Examples in Documentation

Screen readers may not always correctly read the code examples in this document. The conventions for writing code require that closing braces should appear on an otherwise empty line; however, some screen readers may not always read a line of text that consists solely of a bracket or brace.

Accessibility of Links to External Web Sites in Documentation

This documentation may contain links to Web sites of other companies or organizations that Oracle does not own or control. Oracle neither evaluates nor makes any representations regarding the accessibility of these Web sites.

TTY Access to Oracle Support Services

Oracle provides dedicated Text Telephone (TTY) access to Oracle Support Services within the United States of America 24 hours a day, seven days a week. For TTY support, call 800.446.2398.

Related Documents

This guide assumes that you have read and understood the following documents:

For more information, see the following documents in the Oracle Identity Manager documentation set:

- *Oracle Identity Manager Installation Guide for JBoss*
- *Oracle Identity Manager Installation Guide for WebLogic*
- *Oracle Identity Manager Installation Guide for WebSphere*
- *Oracle Identity Manager Globalization Guide*
- *Oracle Identity Manager Design Console Guide*
- *Oracle Identity Manager Administrative and User Console Guide*
- *Oracle Identity Manager Administrative and User Console Customization Guide*
- *Oracle Identity Manager Tools Reference Guide*
- *Oracle Identity Manager Audit Report Developer Guide*
- *Oracle Identity Manager API Usage Guide*
- *Oracle Identity Manager Glossary of Terms*

Documentation Updates

Oracle is committed to delivering the best and most recent information available. For information about updates to the Oracle Identity Manager 9.0 documentation set, visit Oracle Technology Network at

<http://www.oracle.com/technology/documentation/index.html>

Conventions

The following text conventions are used in this document:

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

Using the Deployment Manager

The Deployment Manager enables developers to move an Oracle Identity Manager deployment from one server to another while minimizing the problems that often occur during a migration. The Deployment Manager allows multiple developers to work on different parts of a deployment and upload only the part of the deployment that they have changed, rather than waiting for the entire deployment to be rebuilt.

Caution: Be aware that the most recently imported data overwrites the previous data. Developer should not export data that can overwrite the work of another developer.

This chapter discusses the following topics:

- [Features of the Deployment Manager](#)
- [Export System Objects Only when Necessary](#)
- [Export Related Groups of Objects](#)
- [Group Definition Data and Operational Data Separately](#)
- [Use Logical Naming Conventions for Versions of a Form](#)
- [Export Root to Preserve a Complete Organizational Hierarchy](#)
- [Provide Clear Export Descriptions](#)
- [Check All Warnings Before Importing](#)
- [Check Dependencies Before Exporting Data](#)
- [Matching Scheduled Task Parameters](#)
- [Compile Adapters and Enable Scheduled Tasks](#)
- [Export Entity Adapters Separately](#)
- [Group Permissions](#)
- [Back Up the Database](#)
- [Import Data When the System Is Quiet](#)

Features of the Deployment Manager

The Deployment Manager helps migrate Oracle Identity Manager deployments from one server environment to another, such as from a test environment to a staging environment, or from a staging environment to a production environment.

The Deployment Manager provides the following benefits:

- Updating individual components of a deployment in different test environments
- Identifying objects associated with components to be exported, so that those resources can be included
- Providing information on exported files
- Allowing the exporter to add comments

The Deployment Manager handles the following types of information:

- Resource objects
- Adapters
- IT resource types
- User-defined forms
- Organizations
- User-defined field definitions
- Rule definitions
- Email definitions
- System and error Codes
- Lookup definitions
- User groups
- Password policies
- Access policies
- Scheduled tasks

Limitations of the Deployment Manager

The following are important limitations of the Deployment Manager:

- **Merge Utility:** The Deployment Manager is not a merge utility.
It cannot handle changes done in both production and test environments. It will replace the object in target system with what is in the XML.
- **Version Control Utility:** Deployment Manager does not track versions of imported files, and does not provide rollback functionality.
You can only use it as a means to move data between environments.
- **Code Moving:** Deployment manager does not move `.jar` files in `JavaTasks` or other locations.
You must do this manually.
- **Custom Labels Move:** Deployment Manager does not move labels defined in the `customResources.properties` file or the property files in the `connectorResources` directory. You must do this manually.

Export System Objects Only when Necessary

Only export or import system objects, for example, Request, Xellerate User, and System Administrator, when it is absolutely necessary. Exporting them from testing

and staging environments into production can cause problems. If possible, exclude system objects when exporting or importing data.

Some circumstances might require you to export or import system objects, for example, when defining Trusted Source Reconciliation on Xellerate Users resource objects.

Export Related Groups of Objects

Oracle recommends that you use the Deployment Manager to export sets of related objects. A unit of export should be a collection of logical items that you want to group together.

Avoid exporting everything in the database in one operation, or exporting items one at a time. For example, suppose that you manage an integration between Oracle Identity Manager and a target system that includes processes, resource objects, adapters, IT resource type definitions, IT resource definitions, scheduled tasks, and so on. For this environment, you should create groups of related objects before exporting.

For example, if you use the same email definitions in multiple integrations, you should export the email definitions as one unit, and integrations as different unit. This enables you to import changes to email definitions independently of target system integration changes. Or, if multiple resources use the same IT resource type definition, you can export and import the type definition separately from other data.

You can import one or more sets of exported data at a time. For example, you can import a resource object definition, email definition, and IT resource type definition in one operation.

Group Definition Data and Operational Data Separately

You should group and export definition data and operational data separately.

You configure definition data in test and staging. Definition includes resource objects, processes, rules, and so on.

You typically configure operational data in production. Operational data includes groups and group permissions. The test and staging servers normally do not include this data.

By grouping data according to where it is changed, you know what data goes to test and staging, and what goes to production. For example, if approval processes are changed in production, you should group approval processes and export them with other operational data.

Use Logical Naming Conventions for Versions of a Form

You often revise forms multiple times before exporting them. Avoid generic names, for example, "v23" to differentiate among versions of a form. Create meaningful names, for example, "Before Pre Production" or "After Production Verification". Do not use special characters, including double quotes, in version names.

Caution: The Deployment Manager keeps track of imported components and structures, but not of completed imports. After an import is completed, you cannot roll it back to a previous version. A new import is required.

Export Root to Preserve a Complete Organizational Hierarchy

When you export a leaf or an organization in an organizational hierarchy, only one dependency level is exported. To export a complete organizational hierarchy you must export the root of the hierarchy.

Provide Clear Export Descriptions

The Deployment Manager records some information automatically, for example, the date of the export, who performed the export, and the source database. You should also provide a meaningful description of the content of the export, for example, "resource definition after xxx attributes added in reconciliation". This informs the importer of the file of the contents of the data being imported.

Check All Warnings Before Importing

When importing to production, check all the warnings before completing the import. Treat each warning seriously.

Check Dependencies Before Exporting Data

The wizard in the top right pane shows resources that must be available in the target system.

The following types of dependencies may occur:

1. If the resources are already available in the target system, they do not need to be exported.
2. If the resources are new (not in the target system), they must be exported.
3. If the target system may or may not include the resources, such as lookups, IT resource definitions, or others that are reused, then record the data and export it in a separate file so it can be imported if necessary.

Matching Scheduled Task Parameters

Scheduled tasks depend on certain parameters to run properly. You can import scheduled task parameters to the production server. [Table 1–1](#) shows the rules for determining how to import scheduled tasks. Note that parameters may be available for tasks that no longer reside on the target system.

Table 1–1 *Parameter Import Rules*

Parameter Exists in Target System	Parameter Exists in the XML File	Action Taken
Yes	No	Remove the parameter from the target system.
No	Yes	Add the parameter and current value from the XML file.
Yes	Yes	Use the more recent value of the parameter.

Compile Adapters and Enable Scheduled Tasks

After an import operation, adapters are set to recompile and scheduled tasks are disabled. This prevents these items from running prior to configuring their associated resources and settings.

After importing the classes and adjusting the task attributes, manually recompile adapters and enable scheduled tasks.

Export Entity Adapters Separately

Entity adapters are modified to bring just the entity adapter, not their usage. You must separately export each use of an entity adapter with a data object by exporting the data object. Exporting a data object exports all the adapters and event handlers attached to the object along with the permissions on the object. You need to pay special attention when exporting data objects. For example, to export a form, you should also add the data object corresponding to the form. This ensures that associated entity adapters can use the form.

Group Permissions

When exporting groups, group permissions on different data objects are also exported. However, when importing data, any permissions for missing data objects are ignored. If the group is exported as a way of exporting group permission setup, check the warnings carefully to make sure that permission requirements are met. For example, if a group has permissions for objects A, B, and C, but the target system only has objects A and B, the permissions for object C are ignored. If object C is added later, the group permissions for C must be added manually, or the group has to be imported again.

Report Permissions

When exporting groups that have permissions to view certain reports, ensure that the reports exist in the target environment. If reports are missing, consider removing the permissions before exporting the group.

Back Up the Database

Back up the database before importing data into a production system. This enables you to restore the data if anything goes wrong with the import. Backing up the database is always a good precaution before making significant changes.

Note: When you import forms and user-defined fields, you add entries to the database. These database entries cannot be rolled back or deleted. Before importing, be sure that the correct form version is active.

Import Data When the System Is Quiet

You cannot complete an import operation in a single transaction because the import includes schema changes. These changes affect currently running transactions on the system. To limit the impact of an import operation, separate the web application from general use and perform the operation when the system has the least activity, for example, overnight.

JBoss Components That May be Removed

For a JBoss installation, some JBoss components are not required by Oracle Identity Manager. These files vary for stand-alone and cluster installations.

You can remove the following components:

- Cache invalidation service (keep for a clustered installation)
- J2EE client deployer service
- Integrated HAR deployer and Hibernate session Management services
- JMX Console
- Management console
- Console/email monitor alerts
- UUID key Generation
- JBoss scheduler manager
- Scheduler service
- Test queues and topics
- Mail service
- HTTP Invoker
- CORBA/IIOP
- AOP Application
- Web services support

Files and Directories to Remove

The following sections list the files and directories that you can remove after installing Oracle Identity Manager.

Non-Clustered Installations

Remove the following files from *JBOSS_HOME*/server/default/deploy/:

- cache-invalidation-service.xml
- client-deployer-service.xml
- monitoring-service.xml
- scheduler-service.xml

-
- schedule-manager-service.xml
 - http-invoker.sar
 - mail-service.xml
 - jboss-aop.deployer
 - jboss-ws4ee.sar

Remove the following directories from *JBOSS_HOME*/server/default/deploy/:

- jboss-hibernate.deployer
- jmx-console.war
- management
- uuid-key-generator.sar

Remove the following file:

JBOSS_HOME/server/default/deploy/jms/jbossmq-destinations-service.xml.

Open the *JBOSS_HOME*/server/default/conf/jboss-service.xml file and remove the following attribute:

```
<attribute name="RMI_IIOPService">jboss:service=CorbaORB</attribute>
```

Clustered Installations

Remove the following files from *JBOSS_HOME*/server/all/deploy/:

- client-deployer-service.xml
- monitoring-service.xml
- scheduler-service.xml
- schedule-manager-service.xml
- httpa-invoker.sar
- mail-service.xml
- jboss-aop.deployer
- jboss-ws4ee.sar

Remove the following directories from the *JBOSS_HOME*/server/all/deploy/:

- jboss-hibernate.deployer
- jmx-console.war
- management
- uuid-key-generator.sar

Remove the following file:

JBOSS_HOME/server/all/deploy-hasingleton/jms/jbossmq-destinations-service.xml

Open the *JBOSS_HOME*/server/all/conf/jboss-service.xml file and remove the following attribute:

```
<attribute name="RMI_IIOPService">jboss:service=CorbaORB</attribute>
```

Tuning Oracle Database for Oracle Identity Manager

As with any enterprise class business application, there is no simple procedure for tuning that works for all systems. This section describes one sample configuration and outlines principles for tuning.

Oracle Identity Manager has many configuration options. The best way to identify bottlenecks and optimize performance is to monitor key database performance indicators in your production environment and adjust the configuration from time to time. This chapter serves as a guideline to help you choose the initial baseline database configuration.

This chapter discusses the following topics:

- [Sample Instance Configuration Parameters](#)
- [Physical Data Placement](#)
- [Pinning Sequences and Stored Procedures in System Global Area \(SGA\)](#)
- [Database Performance Monitoring](#)

Sample Instance Configuration Parameters

The following parameter settings are based on a server with 4 CPUs and 8GB RAM

Table 3–1 Sample Configuration Parameters

Parameter	Recommended Initial Settings
compatible	9.2.0.0.0
cursor_sharing	SIMILAR
db_block_size	8192
db_cache_size	3200M
db_keep_cache_size	800M
log_buffer	262144
shared_pool_size	500M
sga_max_size	5500M
db_file_multiblock_read_count	8
db_writer_processes	2
hash_join_enabled	true

Table 3–1 (Continued) Sample Configuration Parameters

Parameter	Recommended Initial Settings
java_pool_size	150M
open_cursors	600
optimizer_feature_enable	9.2.0
optimizer_index_cost_adj	Between 0 and 20
pga_aggregate_target	1000M
workarea_size_policy	auto
processes	Set the processes parameter to the maximum number of concurrent users + the number of background processes + the number of SQL*PLUS and other user processes + 10 (as an extra cushion)
query_rewrite_enabled	true
query_rewrite_integrity	trusted
session_cached_cursors	300

Physical Data Placement

The out-of-the-box installation of Oracle Identity Manager uses only one physical tablespace to store database objects. Oracle Identity Manager database objects belong to one of the following categories:

- Physical tables
- Index
- Large objects (LOBS/CLOBS)

For better performance, create multiple locally managed tablespaces and store each category of database object in a dedicated tablespace. This optimizes storage for efficient data access. Oracle recommends that you place the following User Profile Audit (UPA) component tables and indexes in their own tablespaces:

- UPA
- UPA_FIELDS
- UPA_GRP_MEMBERSHIP
- UPA_RESOURCE
- UPA_USR

These tables can store significant amounts of historical data and can be used by historical reports.

The database schema includes the following tables for reconciliation data:

- RCA
- RCB
- RCD
- RCE
- RCH
- RCM

- RCP
- RCU
- RPC

If your environment generates a large amount of reconciliation data, move these tables to a new tablespace. Use the locally managed tablespaces with automatic extent allocation.

You can use performance metrics to identify tables that are accessed frequently, or *hot* tables. To reduce I/O contention, move hot tables to dedicated tablespaces. See ["Database Performance Monitoring"](#) on page 4 for more on performance metrics.

Pinning Sequences and Stored Procedures in System Global Area (SGA)

Oracle Identity Manager uses sequence objects to generate unique record identifiers. Oracle Identity Manager also uses stored procedures to perform specific database operations. To optimize performance in production, pin the sequence objects and stored procedures in SGA. A script named `create_db_trigger.sql` is shipped with the Oracle Identity Manager installation for this purpose. The `create_db_trigger.sql` script is written for the Oracle Identity Manager database account `SYSADM`. It is a sample Oracle login account.

This script is located in the following installation directory:

```
\installServer\Xellerate\db\oracle
```

To pin the sequence objects and stored procedures:

1. Log in as `SYS`.
2. Start Oracle SQL*Plus (the Oracle client tool), at a command prompt, by typing the following command:

```
sqlplus /nolog
```

3. Connect to the Oracle instance as `SYS` user with `SYSDBA` role.

For example, type the following command:

```
CONNECT SYS/sys_password@db_instance AS SYSDBA
```

Where `sys_password` is the password for the `SYS` user account, and `db_instance` is the Net 8 service name for connecting to the database instance.

For example:

```
CONNECT SYS/sys@xeltest AS SYSDBA
Connected.
```

4. Edit the `create_db_trigger.sql` script in a text editor, and specify your actual Oracle Identity Manager database account name.
5. In `create_db_trigger.sql`, substitute all references to `sysadm` with the account name you actually used.

For example, if your Oracle Identity Manager database account name is `myschema`, edit your script as follows:

```
create or replace trigger cache_seq after startup on database begin
myschema.pin_obj;
-- pin all sysadm's sequences in shared_pool
myschema.pin_sp;
```

```
-- pin all commonly executed XELL stored procedures/functions
end;
/
```

6. Run the `create_db_trigger.sql` script.

This script creates a database trigger that is fired every time the database starts up. Any subsequent database bounces automatically pin the sequences and stored procedures in SGA.

7. While connected to Oracle as the `SYS` user, run the following commands:

```
EXEC database_user.PIN_OBJ;
EXEC database_user.PIN_SP;
```

Where `database_user` is the database account.

Run these commands only once during initial schema creation. Bouncing the Oracle server is not required.

Database Performance Monitoring

To identify performance bottlenecks, you can monitor real-time performance metrics for the Oracle Identity Manager database.

Perform the following at regular intervals:

- Monitor the real-time performance using performance-monitoring tools, for example, Oracle `statspack`, or another tool.
- Collect initial schema statistics upon implementation of Oracle Identity Manager.

Update schema statistics regularly, so that the Cost-Based Optimizer (CBO) has access to the most up-to-date statistics.

This helps the CBO determine a efficient query execution plan that is based on the current state of data.

- Look for relevant recommendations provided in advisory sections in the Oracle `statspack` report, and adjust the instance configuration parameters according to the advised settings.

Managing the Cache

Oracle Identity Manager uses two types of caching: **global** and **ThreadLocal**.

The global cache stores information globally. Any part of the system can access information that is stored in this cache. The global cache uses OSCache from OpenSymphony. One advantage of using OSCache is its support for cluster environments. Database queries are usually stored in the global cache so that repeated queries are not run against the database again.

The ThreadLocal cache stores information that is used multiple times in a single transaction. For example, a query that is issued many times during a transaction uses data from the ThreadLocal cache. The data used for this query does not change for the transaction.

Oracle Identity Manager allows caching by category. You can enable and disable caching for specific entities and configure separate expiration times.

This chapter discusses the following topics:

- [Sample Cache Configuration](#)
- [General Cache Configuration Properties](#)
- [Category-Based Cache Configuration Properties](#)
- [Class Reloading](#)
- [Purging the Cache](#)
- [Optimal Cache Configuration for a Production Environment](#)

Sample Cache Configuration

[Example 4-1](#) is a snippet from the Cache section in the xlconfig.xml file:

Example 4-1 *xlconfig.xml Snippet*

```
<Cache>
  <Enable>false</Enable>
  <ThreadLocalCacheEnabled>false</ThreadLocalCacheEnabled>
  <ExpireTime>14400</ExpireTime>

  <CacheProvider>com.thortech.xl.cache.OSCacheProvider</CacheProvider>
  <XLCacheProvider>
    <Size>5000</Size>
    <MultiCastAddress>231.121.212.133</MultiCastAddress>
  </XLCacheProvider>

  <!-- Individual cache categories -->
```

```

<!-- Adapters and event handlers to be executed on update/insert/delete -->
<DataObjectEventHandlers>
  <Enable>false</Enable>
  <ExpireTime>14400</ExpireTime>
</DataObjectEventHandlers>

...
...
...
</Cache>

```

Note: Oracle recommends that you disable caching in development environments. Data in development environments changes frequently. If cached data is not refreshed in time, it can cause problems for developers working with the product.

General Cache Configuration Properties

The Cache tag refers to the cache configuration and what is contained between the beginning and the end Cache tag. [Table 4–1](#) describes the entries in the Cache section:

Table 4–1 Cache Configuration Parameters

Property	Description
Enable	This property enables components in the cache configuration for categories that are not explicitly defined in the configuration file. If the configuration file does not contain a particular category, the cache uses this entry to enable or disable the category.
ThreadLocalCacheEnabled	This property enables or disables ThreadLocal caching.
ExpireTime	This property specifies a default expiration time for components in the cache configuration.
CacheProvider	This is the complete class path of the provider used for caching. Do not change this property.
XLCacheProvider	This section specifies cache provider properties. In Example 4–1 , the Size and Multicast Address properties are specified.
XLCacheProvider - Size	This property specifies the size of the cache. This number reflects the number of items that the cache stores. If the size is reached, new items are stored in the cache while the least used are pushed out of the cache.
XLCacheProvider - MultiCastAddress	The IP address used for multicast communication among all of Oracle Identity Manager components.

Note: The same MultiCast Address must be used for all Oracle Identity Manager installations in an environment, for example, for all the nodes in a cluster. Cache flushes are propagated to all installations using MultiCast IP. If multicasting is disabled, cache flush is not possible.

Category-Based Cache Configuration Properties

After you preform general cache configuration, each component or category is shown with its own tag name. The tag name reflects a category name that is used in the code to store information in the cache. You can enable or disable each category independently of other categories, and you can set the expiration time for each component or category.

Table 4–2 lists the categories in the cache configuration file. By default, all the categories are disabled in the cache configuration file unless otherwise mentioned in Table 4–2.

Table 4–2 Category Based Cache Configuration Parameters

Category Name	Description
DataObjectEventHandlers	List of event handlers to be run when data object changes occur. This is the location where custom event handler and entity adapters are attached to a data object.
ProcessDefinition	Process definition information, for example process attributes, tasks, task mappings, and so on.
RuleDefinition	Rule definition information.
FormDefinition	Form definition information.
ColumnMap	DB column name from a column code. This is enabled by default.
UserDefinedColumns	User Defined Form and column definitions
ObjectDefinition	Object definition information.
StoredProcAPI	Used to stored total counts when calling APIs with paging capability. Because information changes frequently, the default expiration time for this category is 10 minutes.
NoNeedToFlush	This category defines data that does not need to be flushed and does not fall into a particular category. This category does not have an expiration time. This information is typically populated during initial database setup and never changes in an installation.
MetaData	DB field metadata information. This is category is enabled by default.
AdapterInformation	Adapter variables, compilation status and so on.
OrgnizationName	Cache organization names.
Reconciliation	Reconciliation rules.
SystemProperties	Caches system properties.
LookupDefinition	Caches the conversions between lookup names and fields.
UserGroups	Caches user groups.
LookupValues	Caches the lookup values for a given lookup name.
ITResourceKey	IT Resources DB key cache.
ServerProperties	Caches what data is to be encrypted along with System Properties
ColumnMetaData	Database metadata information for common queries.
CustomResourceBundle	Caches custom resource bundle.
CustomDefaultBundle	Caches custom default bundle.

Table 4–2 (Continued)Category Based Cache Configuration Parameters

Category Name	Description
ConnectorResourceBundle	Caches connector resource bundles

Class Reloading

Class reloading refers to automatically reloading classes without restarting the server. Class reloading settings are useful for scheduled tasks and adapter-related files. Oracle recommends that you enable reloading in development environments. You must restart the Oracle Identity Manager server if cache reloading is disabled and any new adapters are imported, existing adapters are changed, or any .jar files are modified.

Note: Oracle recommends that you disable class reloading in production environments to improve performance.

The class reloading configuration information is included in the xlconfig.xml file as follows:

```
<ClassLoading>
  <ReloadEnabled>true</ReloadEnabled>
  <ReloadInterval>15</ReloadInterval>
  <LoadingStyle>ParentFirst</LoadingStyle>
</ClassLoading>
```

- `ReloadEnabled` enables class reloading on regular basis.
- `ReloadInterval` specifies the time to reload (in seconds).
- `LoadingStyle` specifies the type of loading used.

The following are the different types of loading:

- `ParentFirst` looks for the classes in the parent before loading them from the jar files in `ADPClassLoader` classpath.
- `ParentLast` overrides the classes from the parent. Using `ParentLast` may cause `ClassCastException`s.
- `ParentLoader` is the `ThreadContext Class Loader`.

Purging the Cache

If you want to purge the cache before the allocated amount of time, use the `PurgeCache` utility in the `XL_HOME/bin` directory. This utility purges all elements in the cache.

Depending on the platform, the `PurgeCache` utility is a batch file or a shell script. After you edit the `XEL_HOME` and `JAVA_HOME` environment variables to point to the correct location, you can run the `PurgeCache` from the command line.

To use the `PurgeCache` utility, run `PurgeCache.bat category name` on Windows systems or `PurgeCache.sh category name` on UNIX/Linux systems. The category name argument represents the name of the category that needs to be purged. For example, the following commands purge all `FormDefinition` entries from a system and its clusters:

```
PurgeCache.bat FormDefinition
PurgeCache.sh FormDefinition
```

To purge all Oracle Identity Manager categories, pass a value of "ALL" to the PurgeCache utility.

Note:

- The category name argument of the PurgeCache utility is case sensitive.
 - A `java.lang.NullPointerException` is thrown after running this script. However, this exception does not prevent data from being purged.
-

Optimal Cache Configuration for a Production Environment

Post-deployment changes to the cache configuration may affect performance and usage. Configure your cache using utmost caution.

The following are guidelines for configuring the Oracle Identity Manager cache for a production environment:

- Set all properties to true, except for the `<StoredProcAPI>` setting.
- Increase the `<XLCacheProvider>` size to 15000 (default value is 5000).

[Example 4-2](#) shows the recommended values for the Oracle Identity Manager cache configuration file (`xlconfig.xml`) in a production environment.

Example 4-2 Recommended Cache Values for `xlconfig.xml` in a Production Environment

```
<Cache>
  <Enable>true</Enable>
  <ThreadLocalCacheEnabled>true</ThreadLocalCacheEnabled>
  <ExpireTime>14400</ExpireTime>
  <CacheProvider>com.thortech.xl.cache.OSCacheProvider</CacheProvider>
  <XLCacheProvider>
    <Size>15000</Size>
    <MultiCastAddress>231.172.169.176</MultiCastAddress>
  </XLCacheProvider>

  <!-- Individual cache categories -->
  <!-- Adapters and event handlers to be executed on update/insert/delete -->
  <DataObjectEventHandlers>
    <Enable>true</Enable>
    <ExpireTime>14400</ExpireTime>
  </DataObjectEventHandlers>
  <ProcessDefinition>
    <Enable>true</Enable>
    <ExpireTime>14400</ExpireTime>
  </ProcessDefinition>
  <RuleDefinition>
    <Enable>true</Enable>
    <ExpireTime>14400</ExpireTime>
  </RuleDefinition>
  <FormDefinition>
    <Enable>true</Enable>
    <ExpireTime>14400</ExpireTime>
  </FormDefinition>
  <ColumnMap>
    <Enable>true</Enable>
```

```

        <ExpireTime>14400</ExpireTime>
    </ColumnMap>
    <UserDefinedColumns>
        <Enable>true</Enable>
        <ExpireTime>14400</ExpireTime>
    </UserDefinedColumns>
    <ObjectDefinition>
        <Enable>true</Enable>
        <ExpireTime>14400</ExpireTime>
    </ObjectDefinition>
    <StoredProcAPI>
        <Enable>false</Enable>
        <ExpireTime>600</ExpireTime>
    </StoredProcAPI>

<!-- This information never needs to be flushed out. For example, key for requests
organization and so on. -->

    <NoNeedToFlush>
        <Enable>true</Enable>
        <ExpireTime>-1</ExpireTime>
    </NoNeedToFlush>

<!-- Metadata Information -->
    <MetaData>
        <Enable>true</Enable>
        <ExpireTime>14400</ExpireTime>
    </MetaData>

<!-- Adapter Mapping Information -->
    <AdapterInformation>
        <Enable>true</Enable>
        <ExpireTime>14400</ExpireTime>
    </AdapterInformation>

<!-- Name of the organization for a given key and vice versa -->
    <OrganizationName>
        <Enable>true</Enable>
        <ExpireTime>14400</ExpireTime>
    </OrganizationName>

<!-- Reconciliation rules -->
    <Reconciliation>
        <Enable>true</Enable>
        <ExpireTime>14400</ExpireTime>
    </Reconciliation>

<!-- System Properties -->
    <SystemProperties>
        <Enable>true</Enable>
        <ExpireTime>14400</ExpireTime>
    </SystemProperties>
    <LookupDefinition>
        <Enable>true</Enable>
        <ExpireTime>14400</ExpireTime>
    </LookupDefinition>
    <UserGroups>
        <Enable>true</Enable>
        <ExpireTime>14400</ExpireTime>
    </UserGroups>

```

```
<LookupValues>
  <Enable>true</Enable>
  <ExpireTime>14400</ExpireTime>
</LookupValues>
<ITResourceKey>
  <Enable>true</Enable>
  <ExpireTime>14400</ExpireTime>
</ITResourceKey>
<RecordExists>
  <Enable>true</Enable>
  <ExpireTime>14400</ExpireTime>
</RecordExists>
<ServerProperties>
  <Enable>true</Enable>
  <ExpireTime>14400</ExpireTime>
</ServerProperties>

<!-- Column Meta Data -->
<ColumnMetaData>
  <Enable>true</Enable>
  <ExpireTime>14400</ExpireTime>
</ColumnMetaData>
<CustomResourceBundle>
  <Enable>true</Enable>
  <ExpireTime>-1</ExpireTime>
</CustomResourceBundle>
<CustomDefaultBundle>
  <Enable>true</Enable>
  <ExpireTime>-1</ExpireTime>
</CustomDefaultBundle>
<ConnectorResourceBundle>
  <Enable>true</Enable>
  <ExpireTime>-1</ExpireTime>
</ConnectorResourceBundle>
</Cache>
```

Securing Your Deployment

This chapter describes how to use Oracle Application Server Single Sign-On to secure your Oracle Identity Manager deployment.

This chapter discusses the following topics:

- [Securing the Administrative and User Console](#)
- [Securing the Self Registration and Change Password Pages](#)

See Also: See the *Oracle Application Server Single Sign-On Administrator's Guide* for information on how to protect URLs.

Securing the Administrative and User Console

To secure the Administrative and User Console, use Oracle Application Server Single Sign-On to protect the following URLs:

```
http://hostname:port/xlWebApp  
http://hostname:port/Nexaweb
```

Securing the Self Registration and Change Password Pages

To secure the Self Registration and Change Password pages, use Oracle Application Server Single Sign-On to protect the following URL:

```
http://hostname:port/xlWebApp/Logon.do
```

After using Oracle Application Server Single Sign-On to protect the Self Registration and Change Password pages, you can use the following URLs to directly access the pages:

```
http://hostname:port/xlWebApp/selfRegister.do?method=New%20Registration  
http://hostname:port/xlWebApp/forgetPassword.do?method=displayVerifyUserId
```

Index

A

adapters, 1-4

C

cache configuration
 category-based properties, 4-3
 class reloading, 4-4
 general properties, 4-2
 optimal, 4-5
 production environment, for, 4-5
 purging, 4-4
 sample, 4-1
cache management
 best practices, 4-1
class reloading, 4-4

D

database back up, 1-5
definition data, 1-3
Deployment, 1-1
Deployment Manager, 1-1
 best practices, 1-1
 exporting system objects, 1-2
 features, 1-1
 limitations, 1-2

E

entity adapters, 1-5
export descriptions, 1-4
exporting data
 dependencies, 1-4

G

global cache, 4-1
group permissions, 1-5

I

importing data, 1-5

J

JBoss components, 2-1
JBoss installation
 removing files and directories, 2-1
Jboss installation
 clustered, 2-2
 non-clustered, 2-1

N

naming conventions, 1-3

O

operational data, 1-3
Oracle Database
 performance monitoring, 3-4
 physical data placement, 3-2
 pinning sequences, 3-3
 sample instance configuration, 3-1
 stored procedures, 3-3
 System Global Area, 3-3
Oracle Identity Manager
 JBoss installation, 2-1
 tuning Oracle Database, 3-1
organizational hierarchy
 exporting, 1-4

P

purging, 4-4

R

related groups of objects
 exporting, 1-3
report permissions, 1-5

S

scheduled tasks, 1-4
 parameter matching, 1-4
system objects
 exporting, 1-2

T

ThreadLocal cache, 4-1

W

warnings, 1-4