

Oracle® Identity Manager

Best Practices Guide

Release 9.0

B25937-01

May 2006

Oracle Identity Manager Best Practices Guide, Release 9.0

B25937-01

Copyright © 1991, 2006, Oracle. All rights reserved.

Primary Authors: Craig West, Deepa Aswani

Contributors: Mario Lim, Semyon Shulman

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	1-2
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 Form Versions	1-3
Export Root to Export 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	1-5
Export Entity Adapters Separately	1-5
Group Permissions	1-5
Report Permissions	1-5
Back Up the Database	1-5
Perform the Import 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)	3-2

Database Performance Monitoring.....	3-3
--------------------------------------	-----

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

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 and Upgrade Guide for JBoss*
- *Oracle Identity Manager Installation and Upgrade Guide for WebLogic*
- *Oracle Identity Manager Installation and Upgrade Guide for WebSphere*
- *Oracle Identity Manager Administrative and User Console Guide*
- *Oracle Identity Manager Design Console Guide*
- *Oracle Identity Manager Administrative and User Console Customization Guide*
- *Oracle Identity Manager Tools Reference Guide*
- *Oracle Identity Manager Audit Report Developer Guide*

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 a deployment from one server to another while minimizing the problems normally associated with such a migration. The Deployment Manager also 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: While the Deployment Manager allows multiple developers to export and import their work, be aware that the most recent import will overwrite the previous one. It is therefore important that each developer not export data that could overwrite the work of another developer.

This chapter describes how you can best use the Deployment Manager. It covers 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 Form Versions](#)
- [Export Root to Export 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](#)
- [Perform the Import 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:

- Allows individual components of the deployment to be updated in different test environments
- Identifies objects associated with components to be exported, so that those resources can be included
- Provides information on exported files and allows the exporter to add comments

The types of information managed by the deployment manager include:

- 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

Before you can fully capitalize on the Deployment Manager functionality, you must understand the following limitations:

- **Merge Utility:** The Deployment Manager is not a Merge utility and cannot handle changes done in both production and test environments. It will replace the object in target system with what is there in the XML.
- **Version Control Utility:** Deployment Manager does not keep track of versions of import, and does not provide the rollback functionality. You can only use it as a means to move data between environments.
- **Code Moving:** Deployment manager does not move the `.jar` files placed in `JavaTasks` or other. You must do this manually.

Export System Objects Only when Necessary

System objects such as Request, Xellerate User, and System Administrator should not be exported or imported unless absolutely necessary. Exporting them from testing and

staging environments into production might result in unforeseen problems. Exercise utmost caution, and if possible exclude them, while performing an export or import.

Some circumstances might require you to export and import them, however. For example, when defining Trusted Source Reconciliation on Xellerate Users resource objects.

Export Related Groups of Objects

It is much better to use the Deployment Manager to export related objects such as groups, instead of trying to export everything in the database or items one at a time. For example, an integration with one target system that includes processes, resource objects, adapters, IT resource type definitions, IT resource definitions, scheduled tasks, lookups, error codes, email definitions, and so forth should first have its related objects grouped before exporting.

Define the unit of export as a set of related items. For example, if email definitions are reused with multiple integrations, export the email definitions as one unit, and integrations as different unit. This way, changes to email definitions can be imported independent of target system integration changes. As another example, if multiple resources use the same IT resource type definition, export and import the type definition separately. A unit of export is the collection of logical items that you want to group together.

During import you can use one or more of exports together. For example, you can import the resource object definition, email definition, and IT resource type definition together.

Group Definition Data and Operational Data Separately

Definition data and Operational data should be grouped and exported separately, because definition data is defined in test and staging, while operational data is defined in production.

Definition includes resource objects, processes, rules, and so forth. Operational data includes groups and group permissions and is normally edited in the production environment. The test and staging servers normally do not include all of this data.

By grouping the different types of data according to where they are changed in the server system, you simplify your exports by knowing which ones go to test and staging, and which go directly to production.

For example, if resource objects and provisioning processes are developed in the test and staging environment, and then unchanged in the production environment, and the approval processes are changed in production, then the approval processes should be grouped and exported with the operational data.

Use Logical Naming Conventions for Form Versions

Forms often go through many revisions before they are exported. Rather than leave them with names like "v23", it is better to rename them something that is easily recognized later on, such as "Before Pre Production" or "After Production Verification". Do not use special characters, including quotes, in the version names.

Caution: The Deployment Manager version control keeps track of import components and structures, but not of completed imports. Once an import is complete, it cannot be rolled back to a previous version. Instead, a new import is required.

Export Root to Export a Complete Organizational Hierarchy

When you export a leaf or an organization within 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 certain information automatically, such as the date of the export, who performed the export, and the source database. Fill in the export description with a one-sentence description of the content of the export, for example, "resource definition after foo attributes added in reconciliation". This way the importer of the file knows what the import contains.

Check All Warnings Before Importing

When importing to production, check all the warnings before completing the import. These warning are important, and each should be treated seriously as a yes or no question.

Check Dependencies Before Exporting Data

The wizard in the top right pane shows resources that must be available in the target system. There are different possible scenarios concerning dependencies:

1. If the resources are already available in the target system, then they do not need to be exported.
2. If the resources are new, and therefore 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 them in a separate file so they can be imported if necessary.

Matching Scheduled Task Parameters

Scheduled tasks depend on certain parameters to run properly. The scheduled task parameters can be imported to the production server, however there are rules to determine how they are imported. Sometimes parameters are available for tasks that are no longer 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

Adapters are set to recompile and scheduled task statuses are disabled when imported. This is to prevent them from executing before their associated resources and settings are in place. After importing the classes and adjusting the task attributes, manually recompile the adapters and enable the scheduled tasks.

Export Entity Adapters Separately

Entity adapters are modified to bring just the entity adapter, not their usage. Each use of entity adapter with a data object must be exported separately, 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. For this reason, pay special attention when exporting data objects. For example, to export a Form, also add the data object corresponding to the form to make sure that any associated entity adapter can use the form.

Group Permissions

When exporting groups, group permissions on different data objects are also exported. However, during import any permissions on 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 the permissions requirement is met. For example, if the group has permissions for objects A, B, and C, but the target system only had objects A and B, then the permissions for object C are ignored. If object C is added at a later time, then either the group permissions for C have to 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 these reports exist in the target environment. If reports are missing, consider removing the permissions before exporting the group.

Back Up the Database

Before importing the deployment to production, backup the database. This provides you with a fail safe in case something goes wrong with the import. Backing up the database is always a good precaution before making any significant changes to the system.

Note: Importing forms and user-defined fields adds entries to the database. These database entries cannot be rolled back or deleted. Before importing, check to be sure that the correct forms version is active.

Perform the Import when the System Is Quiet

You cannot perform an entire import in a single transaction because the import includes schema changes. These changes affect currently running transactions on the system. To limit the impact of the import on work performed by the system, separate the web application from general use and perform the import when the system has the least activity, such as overnight.



JBoss Components That May be Removed

For a JBoss installation, certain JBoss components are not required by Oracle Identity Manager and may therefore be removed. These files vary for standalone and cluster installations.

The following components can be removed:

- 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 managerScheduler service
- Test queues and topics
- Mail service
- HTTP Invoker
- CORBA/IIOP
- AOP Application
- Web services support

Files and Directories to Remove

This section lists the files and directories that you can remove after installing Oracle Identity Manager.

Non-Clustered Installations

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

- *cache-invalidation-service.xml*
- *client-deployer-service.xml*
- *jboss-hibernate.deployer* directory
- *jmx-console.war* directory

-
- management directory
 - monitoring-service.xml
 - uuid-key-generator.sar directory
 - scheduler-service.xml
 - schedule-manager-service.xml
 - http-invoker.sar
 - mail-service.xml
 - jboss-aop.deployer
 - jboss-ws4ee.sar

Remove the `jbossmq-destinations-service.xml` file from the `JBOSS_HOME/server/default/deploy/jms/` directory.

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 and directories from the `JBOSS_HOME/server/all/deploy/` directory:

- client-deployer-service.xml
- jboss-hibernate.deployer directory
- jmx-console.war directory
- management directory
- monitoring-service.xml
- uuid-key-generator.sar directory
- scheduler-service.xml
- schedule-manager-service.xml
- httpa-invoker.sar
- mail-service.xml
- jboss-aop.deployer
- jboss-ws4ee.sar

Remove the `jbossmq-destinations-service.xml` file from the `JBOSS_HOME/server/all/deploy-hasingleton/jms/` directory.

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, turning the system for optimum performance depends on many unknown factors. Therefore, there is no simple procedure that will work for all systems. This section describes one sample configuration and outlines principles used to customize your configuration to best meet your needs.

Oracle Identity Manager has many configuration options. The best way to identify bottlenecks and remove them is by monitoring key database performance indicators in your production environment. Adjust the configuration from time to time in order to optimize performance. This chapter serves as a guideline to help you choose the initial baseline database configuration. It includes 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 various database objects. In order to achieve better performance, create multiple locally managed tablespaces, and use a dedicated tablespace to store each category of database objects.

Oracle Identity Manager data is classified into the following categories:

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

Place each category of data in different tablespaces in order to optimize storage for efficient data access. Furthermore, Oracle strongly recommends that you place the User Profile Audit (UPA) component tables (UPA, UPA_FIELDS, UPA_GRP_MEMBERSHIP, UPA_RESOURCE, and UPA_USR) and indexes in their own tablespaces. These tables may need to store significant amounts of historical data and might be used by historical reports.

The database schema includes a set of nine reconciliation tables (RCA, RCB, RCD, RCE, RCH, RCM, RCP, RCU, and RPC) that are used to store reconciliation data. If your environment generates a lot of reconciliation data, move these tables to a different tablespace. Use the locally managed tablespaces with automatic extent allocation.

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

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

Oracle Identity Manager uses sequence objects in order to generate unique record identifiers. Oracle Identity Manager also uses stored procedures in order to perform certain database operations. To achieve optimal performance in production deployment, pin the sequence objects and stored procedures in SGA. For this purpose, a script `create_db_trigger.sql` is shipped with the Oracle Identity Manager

installation. This script is located in the installation distribution in `\installServer\Xellerate\db\oracle`. Log in as `SYS` to run this script and perform the following steps:

1. Start Oracle SQL*Plus (the Oracle client tool), at a command prompt, by typing the following command:

```
sqlplus /nolog
```

2. 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.
```

3. The `create_db_trigger.sql` script is written for the Oracle Identity Manager database account `SYSADM`. This is a sample Oracle login account. Edit the `create_db_trigger.sql` script, and specify your actual Oracle Identity Manager database account name prior to running the script.
4. Open the `create_db_trigger.sql` script in a text editor and substitute all references to `sysadm` with the account name you actually used.
5. For example, if your Oracle Identity Manager database account name is `myschema`, edit your script so it looks like:

```
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. This way, any subsequent database bounces pin the sequences and stored procedures in SGA automatically.
7. Now, 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

Monitor the real-time performance metrics for the Oracle Identity Manager database to help identify any performance bottlenecks. Perform the following steps regularly:

- Monitor the real-time performance using performance-monitoring tools (such as Oracle `statspack`, or any other third party tool).

-
- Collect initial schema statistics upon implementation of Oracle Identity Manager. Then update schema statistics regularly, so that the Cost-Based Optimizer (CBO) has access to the most up-to-date statistics. This helps CBO determine the 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, which allows any part of the system to access information stored in this cache. The global cache uses OSCache from OpenSymphony. One of the major advantages 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 later on within a single transaction. For example, a query that is issued many times during the execution of the transaction will use data from the ThreadLocal cache. The data used for this query does not change for that transaction.

Oracle Identity Manager allows caching by category. You can specifically enable and disable caching for specific entities, and have a separately configurable expiration time.

This chapter covers the following topics of discussion:

- [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

The following is a snippet from the Cache section in the `xlconfig.xml` file:

```
<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: It is recommended that caching be disabled in development environments. The reason for this is because data in development environments changes frequently. If cached data is not refreshed in time, it might 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. The following table describes the entries in the Cache section:

Table 4–1 Cache Configuration Parameters

Property	Description
Enable	This property is used to explicitly enable components in the cache configuration for categories that are not explicitly defined in the configuration file. In other words, if the configuration file does not contain a particular category, the cache uses this entry to enable or disable such category.
ThreadLocalCacheEnabled	This property is to enable or disable ThreadLocal caching.
ExpireTime	This property is used to specify a default expiration time value for components in the cache configuration.
CacheProvider	This is the complete class path of the provider used for caching. This property should not be changed.
XLCacheProvider	This section is used to specify cache provider properties. In the example earlier in this section, the Size and a Multicast Address properties have been 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, extra items to the cache will be stored 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 a given 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 the general cache configuration, each of the components or categories is shown with their own tag name. Each tag name reflects the category name that is used within the code to store information in the cache. Each of the categories can be enabled or disabled independent of other categories. Also, the expiration time can be set for each component or category. You can use the following categories in the cache configuration file:

Table 4–2 Category Based Cache Configuration Parameters

Category Name	Description
DataObjectEventHandlers	List of event handlers to be executed 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 flush and does not fall into a particular category. This category does not have an expiration time. This information is typically populated during the initial database setup and never changes within the same installation.
MetaData	DB field metadata information. This category is enabled by default.
AdapterInformation	Adapter variables, compilation status and so on.
OrganizationName	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.

By default, all the categories are disabled in the cache configuration file unless otherwise mentioned in [Table 4–2](#).

Class Reloading

Class reloading settings are useful for scheduled tasks and adapter-related files. It is recommended that you enable reloading in development environments so that the classes are reloaded automatically without restarting the server. However, if cache reloading is disabled and any new adapters are imported, existing adapters are changed, or any `.jar` files are modified, the Oracle Identity Manager server must be restarted.

Note: It is recommended 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 different types are:
 - `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 `ClassCastExceptions`.
 - `ParentLoader` is the `ThreadContext Class Loader`.

Purging the Cache

Sometimes, you may need to purge the cache before the allocated amount of time. To do this, use the `PurgeCache` utility in the `<XL_HOME>/bin` directory. This utility purges all the elements in the cache.

In the `PurgeCache` batch file, or shell script depending on the platform, you need to change the `XEL_HOME` and `JAVA_HOME` environment variables to point to the correct places. After that, the `PurgeCache` utility can be run from the command line to purge the cache data.

The `PurgeCache` file contains the following line:

```
"%JAVA_HOME%\bin\java" -DXL.HomeDir=%XEL_HOME% com.thortech.xl.cache.PurgeCache
All
```

Note that the last word in this line is `ALL`. This refers to all the categories within Oracle Identity Manager. More specific scripts can be developed to purge specific categories instead of all the categories.

For example:

```
"%JAVA_HOME%\bin\java" -DXL.HomeDir=%XEL_HOME% com.thortech.xl.cache.PurgeCache
"FormDefinition"
```

In this example, all entries related to `FormDefinition` are purged across the entire system (even on a cluster environment).

Note: A `java.lang.NullPointerException` is thrown after running this script. This is a problem in `OSCache` that is currently under investigation. However, this problem does not prevent the data from being purged.

Optimal Cache Configuration for a Production Environment

Because post-deployment changes to the cache configuration may affect performance and usage, it is important that you configure your cache with utmost caution. To configure the Oracle Identity Manager cache for the production environment, it is recommended that you do the following:

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

The following settings are the optimal values recommended for the Oracle Identity Manager cache configuration file (`xlconfig.xml`) for use in the 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 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 -->
    <OrgnizationName>
      <Enable>true</Enable>
      <ExpireTime>14400</ExpireTime>
    </OrgnizationName>

<!-- 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>
</Cache>
```



Index

A

adapters, 1-5

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, 1-2
- 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, 3-3
- physical data placement, 3-2
- pinning sequences, 3-2
- sample instance configuration, 3-1
- stored procedures, 3-2
- System Global Area, 3-2

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, 1-5

- parameter matching, 1-4

system objects

- exporting, 1-2

T

ThreadLocal cache, 4-1

W

warnings, 1-4