

Best Practices for Migrating to
Oracle® Exadata Storage Server

*Oracle Maximum Availability Architecture White Paper
September 2008*

Maximum
Availability
Architecture

Oracle Best Practices for High Availability

Best Practices for Migrating to Oracle Exadata Storage Server

Executive Summary	2
Prerequisites and Scope.....	3
Pre Migration Best Practices.....	3
Perform Capacity Planning.....	3
Employ Exadata Configuration Automation.....	4
Ensure Infiniband Network Is Used for Oracle Clusterware and RAC Communication.....	5
Ensure Proper ASM Software and Disk Group Attributes.....	6
Ensure Proper Database Software and Compatibility.....	6
Enable DB_BLOCK_CHECKSUM.....	6
Migration Best Practices	7
Configure the ASM Allocation Unit Size to 4MB	7
Configure Optimal Database Extent Sizes	7
Choose the Right Migration Strategy.....	7
Logical Migrations - Database Extent Size Change	9
Creating New Tablespaces on Oracle Exadata Storage Server... 10	
Migration Using Logical Standby Databases	10
Migration Using Oracle Streams and Data Pump.....	11
Migration Using Oracle Data Pump	11
Physical Migrations – No Database Extent Size Change.....	12
Migration Using ASM Redundancy	12
Migration Using a Physical Standby Database	13
Migration Using Recovery Manager (RMAN).....	14
Migration Using Transportable Database	14
Migration Using Transportable Tablespaces	14
Post Migration Best Practices.....	16
Check Disk Groups for Rebalance	16
Assess Index Requirements.....	16
References	17

Best Practices for Migrating to Oracle Exadata Storage Server

EXECUTIVE SUMMARY

Oracle Exadata Storage Server is a revolutionary advancement in database technology where database intelligence has been implemented at the storage level resulting in significant application performance gains. For a technical overview of Oracle Exadata Storage Server, see the standard user documentation.

This paper defines best practices for pre-migration, migration, and post migration from legacy storage to Oracle Exadata Storage Server. The best practices aid in determining the most appropriate migration strategy given application service levels and attributes.

PREREQUISITES AND SCOPE

This paper discusses Oracle Exadata Storage Server concepts and migration strategies at a high level. This document assumes that you are already familiar with Exadata Storage Server and with storage system concepts in general. More detailed information about specific Oracle Exadata Storage Server concepts, functionality, syntax, and examples, is in the standard user documentation. For more detailed information about specific migration concepts and strategies, see the documents in the [References](#) section.

PRE MIGRATION BEST PRACTICES

Implement the following best practices before beginning the migration:

- [Perform Capacity Planning](#)
- [Verify Exadata Hardware](#)
- [Employ Exadata Configuration Automation](#)
- [Ensure Infiniband Network Is Used for Oracle Clusterware and RAC Communication](#)
- [Ensure Proper ASM Software and Disk Group Attributes](#)
- [Ensure Proper Database Software and Compatibility](#)
- [Ensure DB_BLOCK_CHECKSUM Is Enabled](#)

Perform Capacity Planning

Follow these steps to perform capacity planning:

1. Understand the capacity of your current environment.

Before migrating to Oracle Exadata Storage Server, you must understand the I/O capacity and characteristics of your current environment. Collect the usable size¹, IOPS (I/Os per second) and MBPS (Megabytes per second) for your current storage system. Use the system statistic `physical IO disk bytes` to help in understanding current MBPS in the source system. Use the system statistics `physical reads` and `physical writes` to help in understanding current IOPS in the source system. These stats are best collected and viewed via an AWR (Automatic Workload Repository) report.

2. Define the proper number of Oracle Exadata Storage Server cells.

Once you understand the capacity of your current system, you can configure the correct number of Oracle Exadata Storage Server cells. It is important to size both for performance and capacity. One HP DL180 G5

¹ Typical ASM disk groups using Oracle Exadata Storage Server are normal redundancy (as opposed to high redundancy) so usable space is half of total capacity.

Exadata cell configured with 12 SAS drives provides 1.78 TB of usable storage, 1 GB/second throughput for large I/Os, and 2700 IOPS for small I/Os².

3. Ensure that post-failure I/O capacity is sufficient to tolerate failures.

It is also important to consider failures when planning capacity. Oracle Exadata Storage Server cell and disk failures are transparently tolerated using Automatic Storage Management (ASM) redundancy. However, it is a best practice to ensure that post-failure I/O capacity is sufficient to meet the redundancy requirements and performance service levels:

- From a redundancy perspective, if a cell or disk fails, ASM will automatically redistribute the cell or disk contents across the remaining disks in the disk group, as long as there is enough space to hold the data.
- From a performance perspective, if a cell or disk fails, the remaining disks should be able to generate the MBPS and IOPS necessary to sustain the performance service level agreement.

For an existing disk group using ASM redundancy, the columns `USABLE_FILE_MB` and `REQUIRED_FREE_MIRROR_MB` in the `V$ASM_DISKGROUP` view give the amount of usable space and space for redundancy, respectively.

Verify Exadata Hardware

Follow these steps to verify Exadata hardware:

1. Verify cable connections via the following steps:
 - Visually inspect all cables for proper connectivity.
 - Run the command, `cat /sys/class/net/ib?/carrier` on all nodes and ensure the output is '1'.
 - Run the command, `ls -l /sys/class/infiniband/*/ports/*/*errors*` and check for errors.
2. Verify the cell provides expected MBPS and IOPS by running the `CALIBRATE` command in `cellcli`.
3. Verify RDS-level network connectivity and bandwidth by running the `rds-stress` command between all cells and database server nodes.
4. Verify the proper cells and database servers are in the Infiniband network topology by running the `ibnetdiscover` command from all cells and database server nodes.

Employ Exadata Configuration Automation

Take advantage of the following Oracle Exadata Storage Server tools and features to automate and simplify configuration tasks:

- `CELLCLI` commands:

² This configuration may change in a future release.

- `CREATE CELLDISK ALL`—This `CELLCLI` command automatically creates celldisks on all available logical unit numbers (LUNs).
- `CREATE GRIDDISK ALL`—This `CELLCLI` command automatically creates griddisks on all available celldisks.
- ASM automated failure group detection:

When creating ASM failure groups on Oracle Exadata Storage Server, grid disks in the same Oracle Exadata Storage Server cell are automatically placed in the same ASM failure group. There is no need to specify the failure group during disk group creation, which greatly simplifies the `CREATE DISKGROUP` syntax.
- `Dcli`

Oracle Exadata Storage Server includes the tool `Dcli` on each cell that you can use to execute commands or scripts across a defined set of cells in parallel. The `Dcli` tool greatly simplifies any operations that must be run across a subset or all cells. Configuration of SSH user equivalency across all cells is an important prerequisite to optimize the use of `Dcli`. `Dcli` provides the “-k” option to automate the distribution of SSH private keys into the `authorized_keys` file.

Ensure Infiniband Network Is Used for Oracle Clusterware and RAC Communication

If using Oracle Clusterware and Oracle RAC, it is a best practice to use the Infiniband Network for all network and storage communication. Infiniband configuration is preconfigured on the cells but the following should be verified on the database servers:

1. Ensure the private network used for Oracle Clusterware communication is using the Infiniband interface. This can be verified with the command, `oifcfg getif -type cluster_interconnect`
2. Ensure the private network used for Oracle RAC communication is running RDS over the Infiniband network by checking the alert log for the following message:

```
Cluster communication is configured to use the
following interface(s) for this instance
 192.168.20.21
cluster interconnect IPC version:Oracle RDS/IP
(generic)
```

If it is not running RDS, relink the oracle binary via the following:

- Shutdown any process using the Oracle binary
- `cd $ORACLE_HOME/rdbms/lib`
- `make -f ins_rdbms.mk ipc_rds ioracle`

Ensure Proper ASM Software and Disk Group Attributes

Verify the ASM software and compatible attribute settings are set as follows:

1. The Oracle ASM software must be running release 11.1.0.7.
2. For the ASM disk group containing Oracle Exadata Storage Server grid disks, specify the attribute settings as follows:
 - Set the `COMPATIBLE . ASM` attribute to 11.1.0.7.
 - Set the `COMPATIBLE . RDBMS` attribute to 11.1.0.7.
 - Set the `CELL . SMART_SCAN_CAPABLE` attribute to `TRUE`.

Ensure Proper Database Software and Compatibility

To access Oracle Exadata Storage Server grid disks, set the Oracle Database software version and the `COMPATIBLE` parameter as follows:

1. Oracle Database software version is running release 11.1.0.7.
2. Set the database `COMPATIBLE` initialization parameter to 11.1.0.7.

Enable DB_BLOCK_CHECKSUM

Oracle Exadata Storage Server is `HARD`³ (Hardware Assisted Resilient Data) compliant. To fully utilize the `HARD` feature, set the `DB_BLOCK_CHECKSUM` initialization parameter to `TYPICAL` or `FULL` to ensure that database checksums are computed and stored in the blocks.

³ For more information about `HARD` compliancy, see <http://www.oracle.com/technology/deploy/availability/htdocs/HARD.html>

MIGRATION BEST PRACTICES

The following sections describe the best practices for migrating an existing database to Oracle Exadata Storage Server:

- [Configure the ASM Allocation Unit Size to 4MB](#)
- [Configure Optimal Database Extent Sizes](#)
- [Choose the Right Migration Strategy](#)

Configure the ASM Allocation Unit Size to 4MB

Oracle Exadata Storage Server performs best when scanning at least 4 MB contiguous chunks. To ensure that this occurs at an ASM level, the disk group's allocation unit size should be set to 4 MB. This parameter is set at disk group creation time and cannot be changed thereafter. Use the following query to determine the allocation unit size for existing disk groups:

```
select NAME,ALLOCATION_UNIT_SIZE from
v$asm_diskgroup;
```

Configure Optimal Database Extent Sizes

As mentioned in the “[Configure Optimal ASM Allocation Unit Size](#)” section, Oracle Exadata Storage Server performs best when scanning at least 4 MB contiguous chunks. To ensure that this occurs at the database level, the database extent sizes should be at least 4 MB. The simplest way to ensure this is to create the tablespaces as locally managed with at least 4 MB uniform extents sizes. If this strategy causes too much space wastage because of the 4 MB database extents, control the extent space usage at a segment level via the INITIAL and NEXT attributes. The `cell_partition_large_extents` is also available to set the minimum extent size for all segments in a partitioned object to 8 MB for auto-allocate tablespaces. Use the following query to check for any segments that have extent sizes smaller than 4MB:

```
select segment_name,partition_name from dba_extents where bytes <
4194304 and owner = '<schema name>';
```

Choose the Right Migration Strategy

Migration to Oracle Exadata Storage Server can be done logically or physically. If the source database does not contain the optimal extent sizes previously discussed, use [Table 1](#) to determine the best logical migration strategy. If the source database already contains the optimal extent sizes, use [Table 2](#) to determine the best physical migration strategy.

Table 1: Logical Migration Strategies: Change The Database Extent Size

Migration Strategy	Source database running 11.1.0.7?	Adjust ASM AU size to 4MB	Support heterogeneous platforms	Time and Operational Investment	Migration Outage Time
Creating New Tablespaces on Oracle Exadata Storage Server <i>Generally best for Data Warehouse customers.</i>	Yes	Yes	No	Lowest	No outage
Oracle Logical Standby <i>Generally best for minimal downtime and the same source / target platform.</i>	No	Yes	No	Medium	Low ⁴
Oracle Streams and Oracle Data Pump <i>Generally best for minimal downtime and different source / target platform.</i>	No	Yes	Yes	High	Low ⁵
Oracle Data Pump <i>Generally best when you have a planned maintenance window to perform the migration.</i>	No	Yes	Yes	Low	High

⁴ When using RMAN backups, there is no effect on the source database during instantiation.

⁵ The source database is affected during instantiation. The effect is controllable by how many parallel processes are used.

Table 2: Physical Migration Strategies: No Change To Database Extent Size

Migration Strategy	Source database running 11.1.0.7?	Adjust ASM AU size to 4MB	Support heterogeneous platforms	Time and Operational Investment	Migration Outage Time
ASM Redundancy <i>Generally best when already using ASM redundancy and the source / target version and platform are the same.</i>	Yes	No ⁶	No	Lowest	No outage
Physical Standby Database <i>Generally best for same release and with supported cross platforms</i>	Yes	Yes	Limited ⁷	Low	Low
Recovery Manager <i>Generally best when migrating at a tablespace level with the same source / target version and platform</i>	Yes	Yes	No	Low	Medium
Transportable Database <i>Generally best migrating to a different platform with the same endian format</i>	Yes	Yes	Yes (same endian)	Low	Medium
Transportable Tablespace <i>Generally best when migrating to a different platform, different endian format, and different release.</i>	No	Yes	Yes	Medium	Medium

Logical Migrations - Database Extent Size Change

Logical migration consists of data extraction from the source database to a platform-independent format followed by data insertion/load into the target database. Logical migration consists of the following strategies:

- [Creating new tablespaces on Oracle Exadata Storage Server](#)
- [Migration using logical standby databases](#)
- [Migration using Oracle Streams and Oracle Data Pump](#)
- [Migration using Oracle Data Pump](#)

⁶ Not changing the ASM AU size is acceptable if you are already using a 4MB AU size.

⁷ See [OracleMetalink](#) Note 413484.1 for information about heterogeneous support.

Creating New Tablespaces on Oracle Exadata Storage Server

If your application inherently supports the logical roll-in of new data and age-out of old data at the tablespace level, a strategy commonly used with the Oracle Partitioning option, then migrate to Oracle Exadata Storage Server by creating all new tablespaces in a new ASM disk group containing Oracle Exadata Storage Server grid disks.

This strategy has the following pros and cons:

- Incurs zero application downtime.
- Experiences less risk due to the phased-in approach.
- Requires that the database that houses both legacy storage and Oracle Exadata Storage Server storage must:
 - Run Oracle Database release 11.1.0.7 software.
 - Have the COMPATIBLE initialization parameter set to 11.1.0.7 before the migration process begins.

Migration Using Logical Standby Databases

If your application service level agreements permit little-to-no downtime, then use an Oracle Data Guard logical standby database to move the data, and track and merge the changes while the application continues to run. At a high-level, this solution first instantiates a physical standby database, then converts it to a logical standby database, and finally, if required, converts the database extent sizes to the optimal size on the logical standby database⁸.

This strategy has the following pros and cons:

- You can eliminate the impact on the source database during instantiation by using RMAN backups to instantiate the physical standby database.
- The data is moved twice, once for the initial instantiation, and then a second time during the planned maintenance on the logical standby database.
- You do not need to upgrade the source database.
- Platform migration is not supported
- This is a more complex solution to architect because the data is changing during the migration but it comes with the benefit of very little downtime.

For more information on using Oracle Logical Standby database for migration, see [OracleMetalink](#) Note 737460.1.

⁸ Other changes could be made during this planned maintenance as well, but the larger extent sizes are the most important for Oracle Exadata Storage Server scan performance

Migration Using Oracle Streams and Data Pump

If your application service level agreements permit little-to-no downtime, then use Oracle Streams with Oracle Data Pump to move the data, and to track and merge the changes while the application continues to run.

This strategy has the following pros and cons:

- The source database is potentially impacted during the initial instantiation. Limit the number of parallel processes performing the instantiation to control the effect on the source database.
- The data is moved once during the initial instantiation.
- You do not need to upgrade the source database.
- Platform migration is supported.
- This is a more complex solution to architect because the data is changing during the migration but it comes with the benefit of very little downtime.

For more information about the Oracle Streams with Oracle Data Pump solution, see Appendix C in [Oracle Streams Concepts and Administration](#) [11].

Migration Using Oracle Data Pump

If your application service level agreements permit enough downtime to move all of the data during a planned maintenance window, then use Oracle Data Pump to move the data in bulk from the legacy storage to Oracle Exadata Storage Server.

This strategy has the following pros and cons:

- This is a simpler solution because data is not changing during migration but it incurs more downtime.
- You do not need to upgrade the source database.
- Platform migration is supported.
- Oracle Data Pump provides the network import feature that eliminates the need to stage the data outside of the database. When using this feature with large partitioned tables, investigate using the QUERY parameter to manually separate the data being imported across several jobs so that all available computing capacity can be leveraged.

For more information about this solution, see the Oracle Data Pump chapter in [Oracle Database Utilities 11g Release 1 \(11.1\)](#) [8].

Physical Migrations – No Database Extent Size Change

Physical migration consists of platform-dependent data extraction from the source followed by transfer to the target. In some cases, you may need to run a conversion routine after the data files are moved to the target⁹.

Because the migration is done at the physical level, the source database might not contain the optimal database extent sizes mentioned in the [Configure Optimal Database Extent Size](#)” section. If performance is the top priority, use the techniques described in the [Logical Migrations to Change the Database Extent Size](#)” section.

All of the physical migration strategies below, except for [Migration Using ASM Redundancy](#),” deal with database or tablespace migration. With a few exceptions¹⁰, a database or tablespace migration to Oracle Exadata Storage Server is the same as a database or tablespace migration to Oracle ASM. Therefore, you can take advantage of the best practices in existing MAA white papers for physical database or tablespace migration. The appropriate paper is referenced in each of the following sections that describe the various migration strategies:

- [Migration Using ASM Redundancy](#)
- [Migration Using a Physical Standby Database](#)
- [Migration Using Recovery Manager \(RMAN\)](#)
- [Migration Using Transportable Database](#)
- [Migration Using Transportable Tablespace](#)

Migration Using ASM Redundancy

If using ASM normal or high redundancy for the legacy storage, you can migrate the entire disk group with no downtime by adding the Oracle Exadata Storage Server storage to the disk group, removing the legacy storage from the disk group, and setting the disk group attribute `CELL.SMART_SCAN_CAPABLE` to `TRUE`. To optimize the operation, remove and add storage as one command in the ASM instance.

An example flow for this strategy would be as follows:

1. Start with existing legacy disks in a normal or high redundancy disk group with a 4MB AU size.
2. Ensure that the RDBMS software and database supports the compatibility requirements in the [Database Compatibility](#) section.

⁹ For example, a conversion routine may be required if this is a change in endian.

¹⁰ You must configure the server-side `OSSINIT.ORA/CELLIP.ORA` file, and the `ASM_DISKSTRING` parameter requires that you include the correct path to the Oracle Exadata Storage Server grid disks (either by taking the default or by specifying the correct path).

3. Ensure that the ASM software and disk groups support the compatibility requirements in the [ASM Compatibility](#) section with the exception of setting `CELL.SMART_SCAN_CAPABLE=TRUE`.
`CELL.SMART_SCAN_CAPABLE` should be set to `FALSE` at this point because legacy disks still exist in the disk group and the Exadata grid disks cannot be added if set to `TRUE`.
4. Add the new Exadata grid disks into the disk group and remove the legacy disks from the disk group with a single `ASM ALTER DISKGROUP` command. Executing both the add disk and drop disk with one command makes the operation more efficient and only one rebalance is required. Ensure that the `POWER` limit used during the rebalance does not impact the application SLAs.
5. After the rebalance is completed, confirm that it did not encounter any errors and that all legacy disks have been removed from the disk group. This can be confirmed by checking the ASM alert log for messages about the rebalance completing successfully, the cache closing the disks, and the disk headers being erased.
6. Finally, now that only Exadata grid disks exist in the disk group, set the `CELL.SMART_SCAN_CAPABLE` to `TRUE`.

This strategy has the following pros and cons:

- Incurs zero downtime.
- Requires that ASM release 11.1.0.7
- Platform migration is not supported
- Application impact can be controlled with the rebalance `POWER` setting

Migration Using a Physical Standby Database

The high-level strategy for this option is to create a physical standby database in Oracle Exadata Storage Server and perform a Data Guard switchover to migrate.

This strategy has the following pros and cons:

- Incurs 1 minute or less downtime.
- Requires that you upgrade the source database to release 11.1.0.7.
- Limited support for platform migration¹¹
- Requires that you specify a new `DB_UNIQUE_NAME`.

More details about this migration strategy are available in the MAA white paper, [Best Practices for Minimal Downtime Migration to ASM](#) [5].

¹¹ See [OracleMetalink](#) Note 413484.1 for information about heterogeneous support.

Migration Using Recovery Manager (RMAN)

The high-level strategy for this option is to create the initial RMAN full backups and subsequent incremental backups on Oracle Exadata Storage Server and use the RMAN switch-to-copy feature to migrate.

This strategy has the following pros and cons:

- The original DB_UNIQUE_NAME is retained; a new DB_UNIQUE_NAME is not required.
- Optional piecewise tablespace-level migration.
- Requires that you upgrade the source database to release 11.1.0.7.
- Platform migration is not supported
- Incurs more downtime than the strategy described in the [Migration Using a Physical Standby Database](#) section, especially if many updates are occurring during the last incremental backup prior to the switch to copy.

More details about this migration strategy are available in the MAA white paper [Oracle Database 10g Best Practices: Migration to Automatic Storage Management \(ASM\)](#) [4].

Migration Using Transportable Database

The high level strategy for this option is to use the Transportable Database feature to migrate the entire database to a different platform with the same endian format.

This strategy has the following pros and cons:

- Provides a simple operational solution.
- Requires the same endian platform on source and target databases.
- Requires that you upgrade the source database to release 11.1.0.7.
- Requires downtime to make data available to the target platform.

More details about this migration strategy are available in the MAA white paper [Platform Migration Using Transportable Database](#) [6].

Migration Using Transportable Tablespaces

The high level strategy for this option is to use the Transportable Tablespace feature to migrate tablespaces.

This strategy has the following pros and cons:

- Permits different endian platform on source and target
- Source database upgrade not required
- Platform migration is supported

Maximum Availability Architecture

- Allows partial database migrations at the tablespace level

More details about this migration strategy are available in the MAA white paper, [Platform Migration using Transportable Tablespaces](#) [7].

POST MIGRATION BEST PRACTICES

Perform the following tasks after completing the migration:

- [Check disk groups for rebalance](#)
- [Assess Index requirements](#)

Check Disk Groups for Rebalance

One of ASM's responsibilities is to ensure that database data is evenly distributed across all disks in the disk group. This happens automatically after any allocation or deallocation. Occasionally, a disk group may become imbalanced due to uncommon errors such as a failed rebalance. It is therefore an operational best practice to check the balance on disk groups on a regular basis and run a manual rebalance, if necessary. A script is available to check disk group balance in [OracleMetalink](#) Note 367445.1. Also, Enterprise Manager Grid Control offers an alert if disk group balance goes below a defined threshold. You can customize the threshold.

Assess Index Requirements

The superior scan rates available from Oracle Exadata Storage Server make it possible that indexes, previously required for good performance on legacy storage, are no longer required. You should assess application execution plans that use indexes to see if they would run faster with Oracle Exadata Storage Server scans. Table scans and index fast full scans are offloaded to Oracle Exadata Storage Server so you can use the `FULL` and `INDEX_FFS` hints to force scans and assess performance.

REFERENCES

1. Oracle Maximum Availability Architecture Web site
<http://www.oracle.com/technology/deploy/availability/htdocs/maa.htm>
2. *Oracle Database High Availability Overview (Part #B14210)*
<http://otn.oracle.com/pls/db111/db111.toc?partno=b28281>
3. *Oracle Database High Availability Best Practices (Part B25159)*
<http://otn.oracle.com/pls/db111/db111.toc?partno=b28282>
4. *Oracle Database 10g Migration to Automatic Storage Management*
http://www.oracle.com/technology/deploy/availability/pdf/MAA_WP_10gASMMigration.pdf
5. *Best Practices for Minimal Downtime Migration to ASM Oracle 10g Release 2*
http://www.oracle.com/technology/deploy/availability/pdf/maa_wp_10gr2_asmmigration_withdg.pdf
6. *Platform Migration using Transportable Database*
http://www.oracle.com/technology/deploy/availability/pdf/MAA_WP_10gR2_PlatformMigrationTDB.pdf
7. *Platform Migration using Transportable Tablespaces*
http://www.oracle.com/technology/deploy/availability/pdf/MAA_WP_10gR2_PlatformMigrationTTS.pdf
8. *Oracle Database Utilities 11g Release 1 (11.1)*
<http://otn.oracle.com/pls/db111/db111.toc?partno=b28319>
9. *Oracle Database Storage Administrator's Guide 11g Release 1 (11.1)*
<http://otn.oracle.com/pls/db111/db111.toc?partno=b31107>
10. *Oracle Streams Concepts and Administration*
<http://otn.oracle.com/pls/db111/db111.toc?partno=b28321>



Best Practices for Migrating to Oracle Exadata Storage Server

September 2008

Author: Michael Nowak

Contributors: Andrew Babb and Douglas Utzig

Oracle USA, Inc.

World Headquarters

500 Oracle Parkway

Redwood Shores, CA 94065

U.S.A.

Worldwide Inquiries:

Phone: +1.650.506.7000

Fax: +1.650.506.7200

oracle.com

Copyright © 2008, Oracle. All rights reserved.

This document is provided for information purposes only and the contents hereof are subject to change without notice.

This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document and no contractual obligations are formed either directly or indirectly by this document. This document may not be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without our prior written permission. 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.