



Information Lifecycle Management with Oracle[®] Database 10g[™] Release 2 and NetApp SnapLock[®]

Regulatory Compliance for Enterprise Data with Oracle Database 10g
Release 2 and NetApp SnapLock

Arvind Shrivastava, Blaine McFadden, and Brajesh Goyal, Network Appliance, Inc.
Lilian Hobbs, Oracle Corporation
January 2007 | TR-3534

Abstract

The number of digital assets that organizations must retain for active reference is subject to regulatory compliance and is ever increasing. Regulatory scrutiny has become increasingly aggressive, with an estimated 10,000 global compliance regulations. Information lifecycle management (ILM) includes technologies and processes to effectively manage enterprise data. Oracle Database 10g Release 2 offers numerous technologies for information lifecycle management of enterprise data for regulatory compliance. Network Appliance offers storage technologies for retaining regulatory-compliant data. This paper discusses how you can utilize Oracle Database 10g Release 2 with NetApp data migration and regulatory compliance solutions for lifecycle management of your critical enterprise data. It also discusses technologies that you can use to migrate and retain Oracle database data in regulatory-compliant storage, and provides guidelines on how to best perform these tasks.

Table of Contents

1. Introduction	3
2. Background	3
2.1 Oracle Database 10g Release 2 ILM Technologies.....	3
2.2 NetApp SnapLock Compliance and SnapLock Enterprise Software	3
3. How to use Oracle Database 10g ILM with NetApp SnapLock	4
3.1 Create Storage Tiers	4
3.2 Create the Data Classes for the Oracle Database Data (Using Partitioning)	6
3.3 Create and Define Data Migration Policies	8
4. Summary	12
5. Revision History	13

1. Introduction

The number of digital assets that organizations must retain for active reference is subject to regulatory compliance and is ever increasing. Regulatory scrutiny has become increasingly aggressive; it is estimated that more than 10,000 compliance regulations exist globally. These regulations describe the process by which records must be created, stored, accessed, retained, and destroyed over increasingly long periods of time. Information lifecycle management (ILM) includes technologies and processes that can be used to effectively manage, organize, and migrate data during its entire lifecycle. ILM is increasingly used to develop regulatory-compliance solutions.

A big percentage of critical enterprise data is organized in databases. Oracle Database 10g Release 2 provides an ideal platform to implement an ILM solution by providing a number of unique features to simplify implementation of the solution. From a storage perspective, Network Appliance offers unique technologies for ILM. In particular, it offers SnapLock technology, which delivers high-performance and high-security data permanence to disk-based near-line and primary NetApp storage. It also provides technologies to simplify migration of data to regulatory-compliant storage.

This paper discusses the various technologies offered by Oracle and NetApp for ILM and how these technologies can be used together. It also provides examples of how customers can implement an ILM solution using these technologies.

2. Background

This section discusses the technologies offered by Oracle and NetApp for ILM.

2.1 Oracle Database 10g Release 2 ILM Technologies

Oracle Database 10g Release 2 is ready today for business ILM. It is capable of storing many different types of data. Storing all of your data in an Oracle database means that it is much easier to manage, because the data is all in one place, instead of being stored using many different formats. The Oracle database is the ideal platform to implement an information lifecycle management policy, because it has a number of unique features that make it very easy to implement an ILM solution:

- Fine-grained: Data is managed at the individual row level
- Application transparency: Data classification is transparent
- Low cost: Uses low-cost storage to reduce costs
- Enforceable: Compliance policies are enforceable

An Oracle database is a structured collection of data. *Data* refers to the characteristics of people, things, and events. Oracle stores each data item in its own field, and stores records relating to each other in a table. For example, Oracle would store all the records for employees of a company in one table, the employee table.

Partitioning allows data stored in a table to be physically segmented according to a data value. Partitioning data is completely transparent to the application, providing an easy way to distribute data across appropriate storage tiers depending on its value, while keeping the data online and stored on the most cost-effective device. For example, all of last year's orders can be stored in one or more partitions on a low-cost storage tier, while all of the current year's orders are stored on the high-performance storage tier.

2.2 NetApp SnapLock Compliance and SnapLock Enterprise Software

The NetApp SnapLock software product family delivers high-performance and high-security data permanence to disk-based near-line and primary NetApp storage. An integrated element of the proven NetApp Data ONTAP® operating system, SnapLock software helps ensure the permanence, accuracy,

integrity, and security of data by allowing business records to be both unalterable and rapidly accessible online for long periods of time.

SnapLock is available in two versions:

- **SnapLock Compliance:** Enables organizations to satisfy strict records-retention regulations such as SEC Rule 17a-4 (broker-dealers), HIPAA (healthcare), Sarbanes-Oxley (public companies), 21CFR Part 11 (life sciences), and DOD 5015.2 (government). Only an act of willful destruction, such as physically removing disks from a NetApp system, can result in record deletion or alteration prior to the specified retention date.
- **SnapLock Enterprise:** Enables adherence to rigorous organizational best practices through functionality similar to that of SnapLock Compliance, but allows administrators to delete entire SnapLock Enterprise volumes. Under no circumstances is it possible for any SnapLock Enterprise user or administrator to delete or modify individual SnapLock Enterprise write once, read many (WORM) records or to undermine SnapLock Compliance WORM volumes.

Retention period support is how long a file must be kept in read-only mode in a SnapLock volume. The retention period for data can vary from 3 to 6 years, according to SEC 240.17a-4. Setting the retention period on a file is based on the last accessed time. A compliance clock prohibits the deletion of content before it expires.

3. How to use Oracle Database 10g ILM with NetApp SnapLock

The Oracle database provides technologies to organize data in the database and to segregate data at the logical level using the partitioning technology. When used in conjunction with the NetApp storage system, data can be physically stored on the various storage tiers and then migrated from one storage tier to another. This section describes the process of implementing ILM using the Oracle database and NetApp SnapLock.

Here are the steps to implement ILM:

1. Create storage tiers (including a WORM storage tier for regulatory-compliant read-only data).
2. Create the data classes for the Oracle database data (using partitioning).
3. Create and define data migration policies.

3.1 Create Storage Tiers

Administrators can create storage tiers based on the various classes of data that they need to support. These can be based on the operational characteristics required for the data, such as performance, availability, etc. For regulatory compliance, data must be archived on regulatory-compliant WORM storage. NetApp SnapLock technology enables administrators to create such a storage tier.

Here are the steps to create a WORM storage tier using SnapLock:

1. Create a SnapLock aggregate.
2. Create a flexible volume within a SnapLock aggregate.
3. Set up a retention policy.
4. Set up an NFS export on the flexible volume.

Create a SnapLock Aggregate

SnapLock aggregates must be created from either the NetApp storage system console or a telnet session to the storage appliance. After a session is established, use the following command:

```
filer> aggr create sle_ilmaggr -L 14
```

Using the `-L` switch in this command creates a SnapLock aggregate called `sle_ilmaggr` with a default RAID-DP™ group and RAID size depending on the appliance, using 14 disk drives.

Next, ensure that the SnapLock aggregate is available and online:

```
filer> aggr status
      Aggr State      Status      Options
      vol0 online     raid4, trad root
sle_ilmaggr online   raid_dp, aggr  snaplock_enterprise
```

Note: The FilerView® GUI cannot currently be used to create a SnapLock aggregate.

Create a Flexible Volume within an Aggregate

Once the SnapLock aggregate has been created, the next step in the process is creating a FlexVol® volume on top of the aggregate. An aggregate can store numerous flexible volumes, and the amount of storage allocated to each FlexVol volume can vary. Each FlexVol volume created on a SnapLock aggregate takes on the SnapLock properties of the underlying aggregate, including either the SnapLock Compliance or SnapLock Enterprise functionality of the aggregate.

Creating the FlexVol volume on a SnapLock aggregate works exactly the same as a FlexVol volume does on a regular read-write aggregate. The FlexVol `create` command syntax is:

```
filer> vol create sle_ilmvol sle_ilmaggr 2g
Creation of volume 'sle_ilmvol' with size 2g on containing aggregate
'sle_ilmaggr' has completed.
```

Where `sle_ilmvol` is the name of the FlexVol volume, and `sle_ilmaggr` is the underlying aggregate from the previous section. The trailing `2g` denotes that the size of the FlexVol volume is 2GB of storage.

Next, verify that the FlexVol volume is available and online:

```
filer> vol status
      Volume State      Status      Options
      vol0 online     raid4, trad root, maxdirsize=10240
sle_ilmvol online   raid_dp, flex  no_atime_update=on,
                    snaplock_enterprise
```

Set Up a Retention Period

Each SnapLock traditional volume or FlexVol volume can have unique maximum, minimum, and default retention periods, which are set using the `vol options` command. The purpose of each setting is as follows:

- Maximum retention period is the longest retention period that Data ONTAP allows to be set on a file, regardless of what the application attempts to overload the last access time attribute for a file with. This protects against situations such as an application being set with a longer retention period than the business requires.
- Minimum retention period is the shortest retention period that Data ONTAP allows to be set on a file, regardless of what the application attempts to overload the last access time attribute for a file with. This protects against files having too short a retention period to meet regulatory compliance.

- Default retention period is the retention period for a record or file if the last access time attribute was not overloaded prior to putting a file into a SnapLock WORM state. Default retention dates can be any value between the minimum and maximum retention period settings.

By default, the maximum retention period is 30 years, the minimum retention period is zero years, and the default is set to equal the minimum retention period. To change any of the retention values, use the `vol options` command:

```
filer> vol options volume_name snaplock_minimum_period 3y
filer> vol options volume_name snaplock_default_period min
filer> vol options volume_name snaplock_maximum_period 10y
```

Set Up an NFS Export within the Flexible Volume

The last step required is to set up a NFS export on the FlexVol volume, then mount it from the appropriate servers. This can be done by creating a directory within the flexible volume and then mounting this directory as an NFS export from the host servers.

For more details on SnapLock configurations, recommendations, and best practices, refer to [TR3342](#).

In the example case study, the following mount points have been set up:

- `/2006data`: Mount point created on a normal read-write volume. This volume is not a SnapMirror or SnapLock volume.
- `/2005data`: In the following example, this data is mounted from a normal NetApp volume. Later this data is copied to a SnapLock volume and then mounted.

3.2 Create the Data Classes for the Oracle Database Data (Using Partitioning)

DBAs can set up database tables with the partitions based on how they want to distribute the data on various storage tiers. Each partition can be stored on a different tablespace. The tablespace is in turn stored on an appropriate storage tier. In addition, administrators can create local partitioned indexes for that table. The partition for the index can be stored on the same tablespace as the partition for the table. This ensures that when they are moved, both the table data and the corresponding index data are moved to the correct storage tier. Figure 1 illustrates an example of a partitioned table and Figure 2 on the next page illustrates the mount points and flexible volumes in utilized in the test environment.

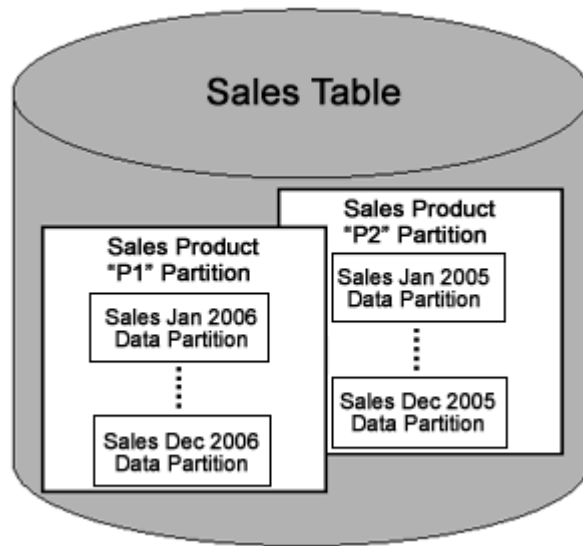


Figure 1) Example of a partitioned table

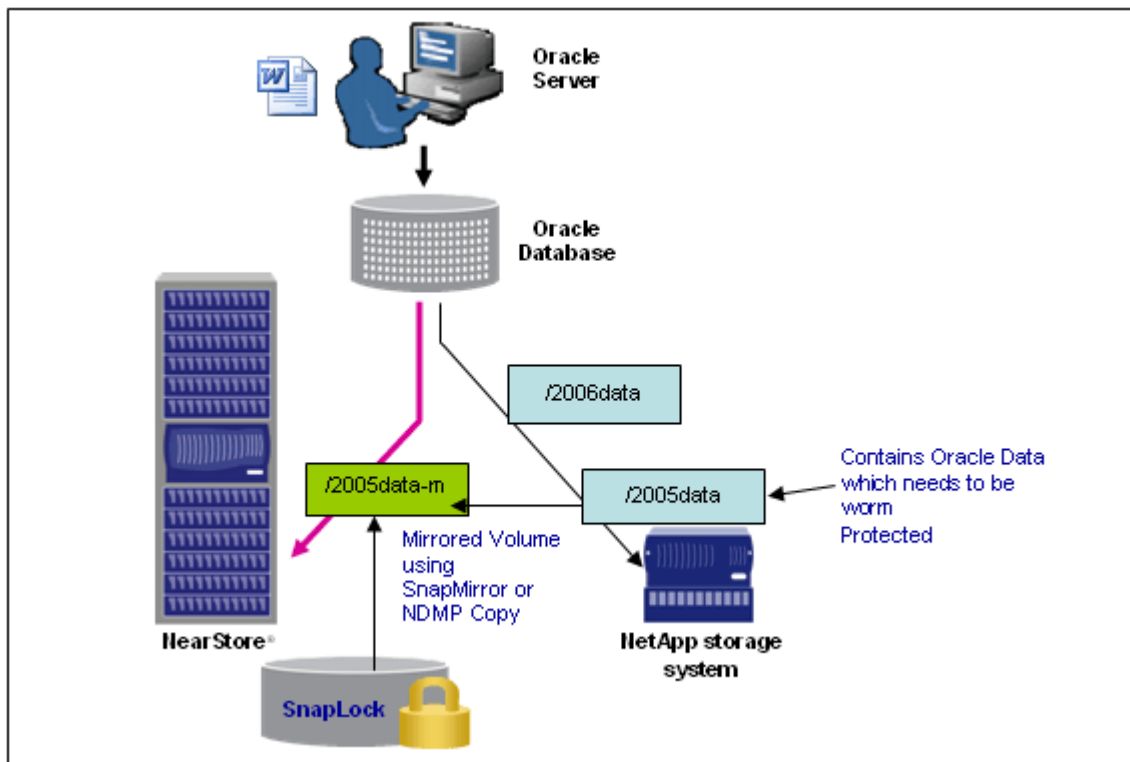


Figure 2) Testing environment for SnapLock with Oracle

The following steps are required to create a partitioned table with partitions on appropriate storage tiers:

1. Create the tablespaces on the appropriate storage tiers.
2. Create a partitioned table with partitions on the appropriate tablespace.

Create Tablespaces

You can create the tablespace using the data files that are located on the appropriate storage tier.

```
SQL> create tablespace Sales_2006 datafile '/2006data/2006_data.dbf' size 10M
SQL> create tablespace Sales_2005 datafile '/2005data/2005_data.dbf' size 10M
```

The `Sales_2006` tablespace is created on a read-write storage tier. This example changes the `/2005data` mount point to point to a WORM storage tier.

Create Partitioned Table

You can create a partitioned table with partitions on the relevant tablespace:

```
CREATE TABLE Sales
(Id NUMBER(3),
 Name VARCHAR2(25),
 Amount NUMBER(7,2),
 Year NUMBER(4),
 Month NUMBER(2),
 Day NUMBER(2))
PARTITION BY RANGE (Year,Month,Day)
(PARTITION P1 VALUES LESS THAN (2006, 12, 31) TABLESPACE Sales_2006,
 PARTITION P2 VALUES LESS THAN (2005, 12, 31) TABLESPACE Sales_2005,
 PARTITION P3 VALUES LESS THAN (MAXVALUE, MAXVALUE, MAXVALUE) TABLESPACE
Sales_Current);
```

This SQL creates a table `Sales` with partitions `p1` on tablespace `Sales_2006` and `p2` on tablespace `Sales_2005`.

3.3 Create and Define Data Migration Policies

Based on the enterprise requirements that specifies what data needs to be archived and how often, enterprises can define automatic policies to migrate the data from a read-write storage tier to a WORM storage tier. Here are the steps that are needed to do so:

1. Migrate data to the WORM storage tier.
2. Unmount the read-write storage tier and mount the WORM storage tier.

Migrate Data to the WORM Storage Tier

You can use many techniques to migrate data to the WORM storage tier. Two of the NetApp recommended techniques are:

- Using SnapMirror®.
- Using NDMPCopy.

Customers choose one of the techniques based on their requirements.

SnapMirror: SnapMirror can be used to mirror the data from the read-write volume to the WORM volume. You can mirror the entire volume or a qtree within a volume. After mirroring the data, the mirror can be broken and the data on the source deleted. Both SnapMirror and NDMPCopy can either be executed from the same NetApp storage system or executed between different NetApp storage systems. The following example is executed from the same host system using a different volume on that system. A flexible volume is stored in an aggregate; for this purpose, a SnapLock volume needs to be stored in an aggregate that is configured as a SnapLock volume.

Command Syntax:

```
filer> snapmirror initialize -S(indicates source vol) <source filer  
ip/hostname>:<source_vol> -w <dest filer hostname only>:<dest_vol>
```

Example:

```
filer> snapmirror initialize -S filer:/vol/dbvol -w filer:/vol/sle_ilmvol
```

If a SnapLock volume is stored on a different NetApp storage system, then the host name changes, as shown in the following example.

```
Example: filer> snapmirror initialize -S filerA:/vol/dbvol -w  
filerB:/vol/sle_ilmvol
```

NDMPCopy: The `ndmpcopy` command enables a storage system administrator to transfer file system data between storage systems that support NDMP v3 or v4 and the UNIX® file system (UFS) dump format.

Using the `ndmpcopy` command, you can carry out both full and incremental data transfers; however, incremental transfers are limited to a maximum of two levels (one full and up to two incremental). You can transfer full or partial volumes, qtrees, or directories, but not individual files.

NDMPCopy can be executed from the same NetApp storage system or from different NetApp storage systems. The following example is executed from the same host system using a different volume on that system. With the use of flexible volumes, a separate aggregate must be created for non-SnapLock volumes.

Command Syntax:

```
filer> ndmpcopy -sa <username>:<password> -da <username>:<password> <source  
filerip/hostname>:<source_vol> -w <dest filer hostname only>:<dest_vol>
```

Example:

```
filer> ndmpcopy -sa root:netapp -da root:netapp filer:/vol/dbvol -w  
filer:/vol/sle_ilmvol
```

If a SnapLock volume will be stored on two different NetApp storage systems, then the host name of the system on which the database files are located changes. However, on one NetApp storage system both SATA and FC disks are supported. The following example uses different systems.

Example:

```
filer> ndmpcopy -sa root:netapp -da root:netapp filerA:/vol/dbvol -w  
filerB:/vol/sle_ilmvol
```

Table 1) Advantages and disadvantages with NDMPCopy vs. SnapMirror

Advantages of SnapMirror	Disadvantages of SnapMirror
<ul style="list-style-type: none">Consistency points and snapshots for file system consistencyNo performance penalty for write acknowledgeNo distance-dependent latencyNetApp storage controller connects right into the networkNo need for expensive ESCON extenders or convertersNo need for expensive service contractsNo need for identical hardware and software between sitesTrue mirror pair, not three or more copies of dataGrow file system dynamically	<ul style="list-style-type: none">There is a small window of potential data loss
Advantages of NDMPCopy	Disadvantages of NDMPCopy
<ul style="list-style-type: none">Auto-detects the NDMP versionSimpler and smaller syntaxThe data is moved on the filer and not across the networkUsing the <code>ndmpcopy</code> command, you can carry out both full and incremental data transfersYou can transfer full or partial volumes, qtrees, or directories	<ul style="list-style-type: none">Works only with NetApp storage controllersNo support for other NDMP serversOnly two levels of incremental backups (one full and up to two incremental)Supports only version 3 and laterMaximum number of NDMP copies limited by source or destination filerNo performance advantageIndividual files cannot be transferred

Unmount the Read-Write Storage Tier and Mount the WORM Storage Tier

The following steps are required to change the tablespace and the storage tier from read-write to WORM.

1. Make the tablespace read-only and shut down the database.

```
SQL> alter tablespace Sales_2005 read only; (to make the tablespace residing in
the SnapLock volume read only)
SQL> alter database open;
SQL> shutdown immediate;
```

2. Change the permission bits on the READ ONLY tablespace (datafile) to READ ONLY. When you change the permission bits on the file, it is converted to a READ ONLY file governed by the retention policy set.

```
# chmod 444 /2005data/2005_data.dbf
```

3. Start up the newly mounted Oracle database on the WORM storage.

```
oracle $ sqlplus '/as sysdba'

SQL> startup;
```

4. By storing the data on a SnapLock-compliant volume, the storage type is transparent to Oracle. The Oracle database can read the tables stored in the SnapLock volume with database modification.

Database Query 1:

Action: Issue a select statement on partitioned tablespace.

Result: Select statement gives output.

```
SQL> select * from sales;
```

ID NAME	AMOUNT	YEAR	MONTH	DAY
-----	-----	-----	-----	-----
		2005	1	1
		2006	2	17
		2006	2	33
		2006	2	29

Database Insert 1 Where Inserts Are Rejected:

Action: Issue an insert statement with a data clause to update the data file in the SnapLock volume on a read-only tablespace. The statement should fail.

Result:- Insert statement throws an error because the table is in read-only mode.

```
SQL> INSERT INTO sales(Year, Month, Day) VALUES (2005, 02, 28);
INSERT INTO sales(Year, Month, Day) VALUES (2005, 02, 28)
*
ERROR at line 1:
ORA-00372: file 18 cannot be modified at this time
ORA-01110: data file 18: '/2005data/2005_data.dbf'
```

Database Insert 2:

Action: Issue an insert statement with a data clause to update the data file in a read-write volume. The statement should pass.

Result: Insert statement inserts the data.

```
SQL> INSERT INTO sales(Year, Month, Day) VALUES (2006, 01, 02);  
  
1 row created.
```

The results in these examples show how seamless running an Oracle database is with SnapLock. Oracle treats the data in the tablespace stored on the SnapLock volume the same as the data in the tablespaces stored on a non-SnapLock volume. These SnapLock volumes can be stored on any NetApp device that uses Fibre Channel or SATA drives.

Change the Oracle Tablespace Back to Read-Write Mode:

Action: Within Oracle SQL*Plus, issue the SQL statement to convert the tablespace back to read-write mode. The statement should fail because it's now on a SnapLock volume.

Result: Modification of the data file on the SnapLock volume is denied.

```
SQL> ALTER TABLESPACE SALES_2005 READ WRITE;  
ALTER TABLESPACE SALES_2005 READ WRITE  
*  
ERROR at line 1:  
ORA-01114: IO error writing block to file 23 (block # 1)  
ORA-01110: data file 23: '/2005data/2005_data.dbf'  
ORA-27091: unable to queue I/O  
ORA-27072: File I/O error  
Linux Error: 30: REad-only file system  
Additional information: 4  
Additional information: 1  
Additional information: -1
```

Remove the Oracle Data Files on the SnapLock Device:

Action: Issue a UNIX or Linux[®] `remove` command to delete the data file after the database is shut down. The command should fail because the volume where the data file resides is on a SnapLock-compliant volume.

Result: Deletion of the data file is denied:

```
# rm -r /2005data/2005_data.dbf  
rm: remove write-protected regular file `/2005data/2005_data.dbf'? y  
rm: cannot remove `/2005data/2005_data.dbf': Read-only file system
```

4. Summary

Regulatory scrutiny has become increasingly aggressive, with an estimated 10,000 global compliance regulations. Information lifecycle management (ILM) includes technologies and processes to effectively manage enterprise data. Oracle Database 10g Release 2 offers numerous technologies for information lifecycle management of enterprise data for regulatory compliance. Network Appliance offers storage technologies for retaining regulatory-compliant data. Combining NetApp and Oracle technologies simplifies and provides enterprises with the means to implement their ILM solutions for regulatory compliance. This paper discussed how you can utilize Oracle Database 10g Release 2 with NetApp data migration and regulatory-compliance solutions for lifecycle management of your critical enterprise data.

5. Revision History

Date	Name	Description
01/02/07	Blaine McFadden	Update
11/06/06	Blaine McFadden/Lilian Hobbs	Update
5/31/2006	Brajesh Goyal	Update
5/20/2006	Arvind Shrivastava	Creation



www.netapp.com

© 2007 Network Appliance, Inc. All rights reserved. Specifications subject to change without notice. NetApp, the Network Appliance logo, Data ONTAP, FilerView, NearStore, SnapLock, FlexVol and SnapMirror are registered trademarks and Network Appliance and RAID-DP are trademarks of Network Appliance, Inc. in the U.S. and other countries. Linux is a registered trademark of Linus Torvalds. Oracle is a registered trademark and Oracle 10g is a trademark of Oracle Corporation. UNIX is a registered trademark of The Open Group. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such.