

APPLICATION NOTE

March 2006

Oracle database remote replication using Sun StorEdge Data Replicator and Oracle Data Guard

ABSTRACT

For a variety of reasons, system administrators often need to back up Oracle databases remotely. This replication can be performed easily using Sun StorEdge Data Replicator and/or Oracle Data Guard. Each has specific steps, applications, and advantages, which are detailed in this paper.

- 1.0 Executive summary 2**
- 2.0 Remote replication applications 2**
 - 2.1 Business continuance, disaster recovery, and data protection2
 - 2.2 Consolidation and data migration3
 - 2.3 End-to-End information management.....3
 - 2.4 Content distribution, archiving, and backup3
- 3.0 Key planning considerations 3**
 - 3.1 Recovery point and recovery time objectives.....3
 - 3.2 Failover and failback between sites4
 - 3.3 End-to-End integration with server applications5
 - 3.4 Data and replication network security5
- 4.0 Sun StorEdge Data Replicator — the Sun difference 5**
 - 4.1 StorEdge Data replication capabilities6
- 5.0 Implementing a remote Oracle database replication..... 7**
 - 5.1 Remote replication site instantiation7
 - 5.2 Replicating Oracle databases8
 - 5.2.1 Replicating with Oracle Data Guard8
 - 5.2.2 Replicating with Sun StorEdge Data Replicator8
 - 5.2.3 Sun Data Replicator tasks9
 - 5.2.4 Re-initializing the Remote Oracle Database9
 - 5.3 Ongoing replication to the Remote Oracle Database.....9
 - 5.3.1 Ongoing replication with Oracle Data Guard9
 - 5.3.2 Ongoing replication with Sun Data Replicator10
- 6.0 Updating remote databases after initialization 10**
- 7.0 Oracle hot standby database vs. mirroring with Sun StorEdge Data Replicator . . 10**
- 8.0 Summary 11**
- 9.0 Glossary 11**

1.0 Executive summary

Oracle databases may need to be replicated remotely for a number of reasons, including business continuance, disaster recovery, data protection, data consolidation and migration, Information Management (including Information Lifecycle Management), content distribution, and archiving.

There are two excellent options for replicating Oracle databases and other database information remotely:

- You can replicate Oracle 10g, 9i, and 8i databases on Sun StorEdge 6920 systems to a remote Oracle system using Sun Data Replicator and Oracle Data Guard technology. Oracle Data Guard, which is built into the Oracle database engine, controls and manages a remote database in standby mode to maintain transactional consistency with the source database.
- Sun StorEdge Data Replicator mirrors entire Oracle and non-Oracle database volumes on a remote system.

You can employ these replication technologies individually or combine them to provide a strong joint replication solution. This paper covers planning considerations for any remote replication procedure, the capabilities of each method, the type of remote replication results achieved, actual replication tasks, and the rationale for using these technologies.

This paper is intended for Oracle DBAs, storage managers, and others with a good working knowledge of database and storage concepts.

2.0 Remote replication applications

Let's begin with an overview of why organizations might want to remotely replicate Oracle databases. There are numerous reasons, including:

- Business continuance, disaster recovery, and data protection
- Consolidation and data migration
- End-to-end Information Management
- Content distribution, archiving, and backup

The following sections provide details on each of these areas.

2.1 Business continuance, disaster recovery, and data protection

Business Continuance (BC), Disaster Recovery (DR), and Data Protection (DP) all require replication of data, such as that contained in an Oracle database, to a secondary location. This ensures that a restorable copy of the data is available that can be recovered to in the event of a disaster or unplanned outage. These three strategies represent the full spectrum of possible responses to IT breakdowns.

Business continuance includes the policies and practices that use replicated data to restore a business information infrastructure with zero data loss after a disaster or unplanned outage. Data replication is high-speed and always synchronous. BC site applications usually run in operational or standby mode so that they can be immediately instanced when an operational failover occurs.

Disaster recovery is a mid-level information protection strategy that is more time-sensitive than data protection and is less costly than BC. It requires high-speed — but not fully synchronous — replication. Recovery time after a disaster can take anywhere from a few minutes to a few hours and consists of data restoration and the time required to bring production applications on-line.

You might need to replicate a database remotely for reasons ranging from business continuance and archiving to consolidation and distribution.

Business continuance, disaster recovery, and data protection represent the full spectrum of replication needs.

DP is typically the lowest level of information protection. While it requires that a recoverable replica of data exist, there is no recovery time frame associated with the data. It can consist of low- to medium-speed replication and is usually asynchronous in nature.

2.2 Consolidation and data migration

Organizations can replicate an Oracle database instance clone at a remote location to facilitate both consolidation and data migration. For example, organizations may clone multiple Oracle databases from different storage arrays onto a single storage array to consolidate data. Also, an Oracle database can be migrated to a different storage system by cloning it to another storage system and then removing from the source storage.

2.3 End-to-End information management

Worldwide, information is being generated at a rapid pace. Oracle databases store a large percentage of this information. New compliance regulations require that much of this information be retained for an extended period of time. The problem is the operational business value of information often diminishes more rapidly than the required information retention period. Remote replication enables organizations to shift all or part the database to more economical storage tiers, such as SATA disk drive-based storage arrays, as needed.

2.4 Content distribution, archiving, and backup

A split mirror copy of an Oracle database could be replicated to multiple remote storage arrays as a means of distributing duplicate content to multiple network edge sites. For example, an organization could periodically load multiple worldwide web edge mirror sites with updated Oracle database content generated at the central source site.

Replicated Oracle database instances can also be used for archiving within different storage tiers as well as for backup to tape or virtual tape without production database disruption.

3.0 Key planning considerations

Now that we know why we may remotely replicate Oracle databases, there are several essential concerns to address when planning any database replication, including:

- Recovery point and recovery time objectives
- Failover and failback between replication sites
- End to-end integration with server applications
- Data and replication network security

The following sections describe each of these considerations.

3.1 Recovery point and recovery time objectives

Any remote Oracle database replication must take Recovery Point Objective (RPO) and Recovery Time Objective (RTO) into account, especially if it is part of a business continuance implementation.

RPO defines by how much time the remote Oracle database copy can lag. In other words, how out of synchronization with the source database instance it can be. This translates into how much data loss an organization can tolerate, because the source data that has not yet been received by and applied to the remote site database is assumed to be lost if a disaster or unplanned outage occurs at the primary site.

You can use remote replication techniques to combine storage or simply move data around your networks.

Any remote replication process requires you to decide how much data you can afford to lose and how long your systems can remain down.

An RPO of zero or zero data loss (ZDL) requires either synchronous replication or synchronous redo log shipping:

- Synchronous replication with Sun StorEdge Data Replicator ensures that any change to the source database is immediately applied and reflected by the remote database copy before any new changes to the source database can occur.
- Through real-time redo log shipping, Data Guard ships the database changes logged in the redo logs to the remote system using a TCP/IP network. When a server system at the remote site receives the redo logs, it immediately applies them to the remote database to keep it synchronized with the primary database.

For non-ZDL implementations, asynchronous replication lets the application of changes to the remote database lag source database changes by a user-specified amount, meaning that the state of the remote database is behind that of the source database. Organizations determine the RPO for asynchronous replication based on the planned amount of time by which the remote database replication will lag the change occurring on the source database.


The amount of source data change that has occurred during that time lag is considered the *tolerable data loss*. That data may be permanently lost. On the other hand, if the redo logs have been replicated but not applied to the remote database, the data can be recovered during the procedures that support the Recovery Time Objective.

The RTO is the period of time within which systems, applications, or functions must be recovered after an outage. RTOs are often used as the basis for the development of recovery strategies, and as a determinant as to whether or not to implement the recovery strategies during a disaster situation. A similar term is *Maximum Allowable Downtime*.

3.2 Failover and failback between sites

Along with RPO and RTO, any remote Oracle database replication business continuance strategy requires a plan for failover and failback procedures.

Figure 1. Failover and failback planning help ensure business continuance.



Failover to the secondary site occurs in the event of a disaster or unplanned outage at the primary site. Both Sun StorEdge Data Replicator and Oracle Data Guard support remote site failover. In the event of a disaster or an unplanned outage at the production (primary) site, the secondary site can become operational as the production site.

Failback from the secondary site back to the primary site takes place after a failover has occurred and the problem at the primary site has been corrected. Failback from the remote to the primary site to restore it as the production site is achieved by replicating changes made to the remote Oracle database back to the primary site database.

Your remote replication strategy must include plans for switching operations from the primary to the secondary sites and back again.

The need to replicate non-Oracle database information is an important consideration when planning a remote Oracle database replication. Either Sun StorEdge Data Replicator or Oracle Data Guard can replicate the database changes, but Sun StorEdge Data Replicator would be required to replicate any non-Oracle database information that has changed. Oracle Data Guard cannot replicate any other information than that contained in the source database instance.

3.3 End-to-End integration with server applications

Uninterrupted access to both the primary and secondary site servers and their applications is an essential part of Oracle database business continuance planning. Server and application continuance procedures ensure the least amount of disruption during a failover or failback event occurs and so that the overall RTO can be achieved.

Depending on implementation strategy requirements, server, critical applications, backup applications, and cluster integration may have to be highly integrated with the replicated database. Server and process virtualization technologies, such as Solaris Containers, may require additional integration to ensure the intended business continuance goals are achieved.

3.4 Data and replication network security

As with any information exchange over a communications network, remote database replication planning must take security into consideration. The most common network security methods in use are:

- Encryption, which encodes transmitted information with various levels of protection so that it is difficult if not impossible to decipher
- Key management, which manages the multiple encryption keys that may be used throughout an organization's computing infrastructure
- Network access controls, which help ensure that only authorized individuals can access and change procedures for the replication network

4.0 Sun StorEdge Data Replicator — the Sun difference

Sun takes remote Oracle database replication one step further than traditional remote mirror replication technologies by supporting management and replication of Oracle databases located on heterogeneous storage arrays. This enables a common replication infrastructure regardless of the storage platform on which the Oracle database resides.

Sun StorEdge 6920 heterogeneous array replication technology enables scalable Oracle database storage tiering. Scalable tiering ensures that administrators can seamlessly integrate the right methods of protection, data movement, archiving, and backup with the least possible complication. Oracle databases can be migrated and replicated and database capacity can be scaled within a SAN to accommodate database expansion or rehosting without the constraints of switch-based, heterogeneous replication technologies.

Sun StorEdge 6920 heterogeneous replication technology combines with snapshot, mirror, and remote mirror replication technologies to provide industry-leading capabilities. Using a common implementation methodology, StorEdge 6920 gives Oracle databases residing on multiple, heterogeneous storage platforms uniform end-to-end protection between multiple storage tiers. This significantly exceeds the capabilities of single-tier platforms that can replicate data only between similar storage platforms and require completely different replication products for enterprise and midrange storage system replication.

Sun StorEdge Data Replicator allows replication between heterogeneous storage media, potentially saving time, money, and ensuring compliance.

4.1 StorEdge Data Replication capabilities

The StorEdge Data Replicator software enables mission-critical information protection through real-time synchronous or asynchronous data replication to campus, metro, or remote data centers. The StorEdge Data Replicator remote replication topologies are shown in table 1.

Table 1. StorEdge Data Replicator remote replication topologies.

Distance	Topology	Replication Type(s)	Network Extension Equipment
0–10 km	Campus	Synchronous or Asynchronous	Multimode (500m) or Single Mode Fibre Channel
10–50 km	Extended Campus or Metro	Synchronous or Asynchronous	IP SAN Extender FC to IP Converter SONET/SDH C/DWDM
50–200 km	Extended Metro or Regional	Synchronous or Asynchronous	IP SAN Extender FC to IP Converter SONET/SDH C/DWDM
200+ km	Interregional or International	Asynchronous	IP SAN Extender FC to IP Converter SONET/SDH

Sun StorEdge Data Replicator can perform remote replication across a campus, a city, or a continent.

The StorEdge Data Replicator replication network can either be TCP/IP or Fibre Channel over public or private telecommunications infrastructures. This enables replication between sites located throughout the world and enables data to be written transparently to both primary and secondary sites either simultaneously or with a managed delay.

StorEdge Data Replicator software also includes a “fast synchronize” feature to protect against possible link failures. This feature allows data to be resynchronized quickly and can also be used to resynchronize data at predetermined intervals.

Table 2. StorEdge Data Replicator software features and benefits.

Table 2 below summarizes the benefits of StorEdge Data Replicator software features and functions.

Feature	Benefit
Synchronous and asynchronous replication modes	<ul style="list-style-type: none"> .. Synchronous provides remote replication with zero data loss .. Asynchronous enables economical and long-distance remote replication .. Supports synchronous and asynchronous modes over both IP and fibre channel links .. Replication distances can range from building-to-building within a campus or between sites over long distances
Write-order consistency across volumes	<ul style="list-style-type: none"> .. Preserves write transaction order across remote volumes .. Protects against data corruption .. Enables remote volumes to be immediately used as restartable volumes in the event of a primary site failure .. Replicated data can be used for application testing, data mining, or backup to tape
Fast start from tape	<ul style="list-style-type: none"> .. Allows initialization of volume groups at the remote site by placing data on the systems from previously shipped tape storage .. Configure large volumes at the remote site with fast, inexpensive tape instead of doing a full remote volume initialization over the replication network
Replication on heterogeneous storage volumes	<ul style="list-style-type: none"> .. Allows data from heterogeneous storage to be replicated to comparable heterogeneous storage volumes at remote sites .. Protects access to older data (as mandated by Federal compliance laws)
Role reversal	<ul style="list-style-type: none"> .. Can be used to make the secondary site the new primary site .. Allows operations be restarted and recovered quickly at the remote site with production-level data if the primary site experiences a disaster or other disruptive event
Scripting interface	<ul style="list-style-type: none"> .. Use the Remote Thin Scripting Client to automate tasks .. Can set cron jobs to replicate data between primary and secondary sites at night or reduced activity time when application use is low or can be suspended
Multi-site support	<ul style="list-style-type: none"> .. Allows up to four Ethernet or eight Fibre Channel replications to one or more StorEdge 6920 systems simultaneously

5.0 Implementing a remote Oracle database replication

So what exactly is involved in a remote Oracle database replication environment?

Some common tasks are:

- .. Remote replication site instantiation
- .. Replicating source Oracle database changes to update the remote database
- .. Periodic bulk update of the remote Oracle database

The following sections provide details on each of these tasks.

5.1 Remote replication site instantiation

The first step in any new remote Oracle database replication is that administrators must instantiate the remote site with an initial copy of the Oracle database. This can be done either with Oracle Data Guard or Sun StorEdge690 Data Replicator. Each has unique instantiation considerations.

All remote replications begin with remote database initialization. How it's done depends on your chosen replication technique.

For example, Oracle Data Guard requires that the remote database instance be loaded from a physically identical backup copy of the source site database.

Sun StorEdge 6920 Data Replicator provides two ways to instantiate the remote Oracle database instance; either through direct replication over the assigned replication network links or through the Fast Start utility:

- Direct replication requires that all the data on a volume be copied to the remote site before going live and continuing with update replication.
- Fast Start allows administrators to transport and load a tape copy of the source volumes onto the remote mirror volumes to speed initializing the mirrors to a state where updates to the mirrors can be replicated across the network. This reduces the initialization time, especially for asynchronous replication, where the network bandwidth available to replicate update information may be low.

5.2 Replicating Oracle Databases

Both Data Guard and Sun StorEdge Replicator have specific advantages and planning considerations when it comes to actual replication.

5.2.1 Replicating with Oracle Data Guard

Oracle Data Guard mirrors the online transaction redo log data from the source production database instance site to the remote site database instance, providing transactional consistency. Instead of mirroring the entire database volume content and layout, Data Guard mirrors the effect of data content changes by applying the redo data to the remote database instance in the same manner in which it's applied to the source database instance.

Oracle Data Guard transmits redo logs to replicate the database, thus reducing bandwidth requirements.

Transmitting Oracle database redo logs generally requires less network bandwidth than does a remote mirroring solution. Therefore, the time lag between the transaction commitment on the source and remote databases is and can remain small. This can be especially important when using low-speed, long distance replication networks.

Standby databases using Data Guard can also be used to perform reporting and backup duties. You can generate reports against the standby database and perform initial or incremental backups with no impact on the production Oracle database.

As with any replication technique, there are considerations when using Data Guard:

- Both the remote server and storage array resources must be used to apply the redo log data when using Oracle Data Guard to achieve the appropriate database state. This differs from Sun Data Replicator, which uses only storage array resources to create a database replica.
- Data Guard does not ensure that the remote database volumes are mirror duplicates of the source database volumes, which might lead to performance inconsistencies between the source and remote sites
- Data Guard cannot replicate any information other than that contained within the Oracle database

5.2.2 Replicating with Sun StorEdge Data Replicator

Remote replication with Sun StorEdge Data Replication requires that the database files, online logs, archive logs, and control file be mirrored. Administrators must place the database into Hot Backup mode so that all of the data can be replicated.

While mirroring with Sun Data Replicator requires more bandwidth, it reduces the workload on remote servers.

As was mentioned earlier, remote mirroring solutions normally require more network bandwidth than Data Guard-administered changes. However, the remote storage infrastructure (servers) will have less work to keep the database up to date because the redo data won't have to be applied to create the necessary remote site database change. Remote mirroring also enables organizations to replicate non-database information.

5.2.3 Sun Data Replicator tasks

There are several tasks required for remote Oracle database replication using Sun Data Replicator. While this is not an exhaustive list, you must at least:

- Allocate storage on the remote StorEdge system or heterogeneous storage system to be used as a mirror for the source StorEdge or heterogeneous storage volumes.
- Determine the data change rate on the source Oracle database production volumes and size the replication network to accommodate the peak synchronous or median asynchronous data change rate replication.
- Allocate any local mirror (snap) volumes that will be used for the Oracle database on the source system.
- Allocate any remote local mirror (snap) volumes that will be used for the Oracle database on the remote system.
- Initialize the remote replication site by either replicating the source volumes across the replication link or by using StorEdge Data Replicator Fast Start to load the database from tape.
- Run the Oracle RECOVER DATABASE command on the remote replica volumes to instantiate the remote database.

5.2.4 Re-initializing the remote Oracle database

In the event of a failover, replicated mirror duplicates of Oracle databases require re-initialization before they can be used as production databases. Re-initialization ensures the remote Oracle database tablespaces and other elements are consistent with the known state of the database stored in the database control and initialization files. You create a replicatable copy of the production database control files by issuing a DBA action command such as ALTER SYSTEM BACKUP STANDBY CONTROLFILE.

Refer to the Oracle database documentation to determine the exact command requirements to create the control file backup that is replicated to the remote site. You can replicate the control file backup using Sun Data Replicator or a file transfer utility.

5.3 Ongoing Replication to the Remote Oracle Database

After you have the remote site standby database instantiated you can put remote replication into operation. Subsequent changes to the source site database replicate to the remote site as they occur.

Once you instantiate a remote site, just a few tasks are required before ongoing replication is largely automatic.

5.3.1 Ongoing replication with Oracle Data Guard

- Use Oracle Data Guard to synchronously transmit the redo log changes over a TCP/IP Ethernet connection. A host system applies redo log changes to the remote standby database to keep the remote database synchronized with the source database.
- Periodically refresh the standby database as a mirror duplicate of the source database during off-peak periods.

5.3.2 Ongoing replication with Sun Data Replicator

- Place the source Oracle database in Hot Standby mode.
- Replicate all Oracle database volumes with Sun StorEdge Data Replicator software to the remote site.
- Remove the source Oracle database from Hot Standby mode.
- Run the Oracle RECOVER DATABASE command on the remote replica volumes to reinstantiate the remote database.

6.0 Updating remote databases after initialization

After the remote Oracle database has been initialized, you can use either Sun StorEdge Data Replicator or Oracle Data Guard to update the remote database instance with changes that occur to the source database. With Oracle Data Guard, the source databases changes are continuously sent as redo log updates and applied to the remote database by a host system at the remote site.

With Sun StorEdge Data Replicator, you can replicate either continuous or periodic database updates, depending on the acceptable data loss goals for the remote Oracle database.

Remember, along with replicating the remote Oracle database information you can use Sun StorEdge Data Replicator to replicate any non-database information residing on the database volumes or other volumes. Other volumes are included into the remote Oracle database replication set and changes replicate to the remote site as they occur.

7.0 Oracle hot standby database vs. mirroring with Sun StorEdge Data Replicator

The best approach for you depends upon the planned use for the replicated database at the remote site. If the replicated database instance will not be active, that is, is in place primarily to protect against a disaster or unplanned outage at the source site, a mirror replica of the database might be preferred. In this case, the remote database replica does not need to be continuously initialized and may only require initialization when it is instantiated during disaster recovery procedures.

The other situation in which you would prefer a remote mirror replica of the database is when server resources are not permanently attached to the storage array at the remote site. Oracle Hot Standby Database requires some number of servers be permanently attached to the storage array at the remote site to apply the redo log changes. If server resources are simply allocated on-demand during disaster recovery procedures, then a storage-array-to-storage-array remote mirror replication is preferred.

Conversely, if the remote database is to be an active database instance, such as for read operation load balancing between the production and standby databases, then you might prefer Oracle Hot Standby Database. Remote mirroring can also be used to accommodate this, but the remote database might require frequent re-initialization to be used in standby mode. Storage-array-to-storage-array remote replication will still be required to replicate non-database information to the remote site. The combination of Sun StorEdge Data Replicator and Oracle Data Guard Hot Standby Database would provide optimum information protection during business operating hours.

Mirroring with Sun StorEdge Data Replicator may be preferred if you are creating a non-active backup or do not have attached server resources.

During non-operational hours you could use Sun StorEdge Data Replicator to create a fully synchronized database mirror copy either daily, weekly, or as frequently as desired. This would ensure that the remote database copy layout exactly reflects the layout of the source production database. This mirroring is especially important when using database performance tuning to periodically optimize the production database. Replicating a mirrored copy of the source database at the remote site would ensure that any database layout changes made through performance optimization would also be reflected in the remote database instance.

8.0 Summary

The Sun StorEdge Data Replicator software and Oracle Data Guard Standby Database capability both represent highly capable means by which you can replicate a production Oracle database to a remote site for disaster protection, secondary active database access, backup, or content distribution. You can employ them individually or combine them to provide different types of remote data protection to match organizational needs.

The Sun StorEdge Data Replicator software provides an additional level of capability beyond contemporary remote mirroring technologies by supporting replication between heterogeneous storage. This capability would enable an Oracle database resident on an EMC CLARiiON CX700, for example, to be replicated to native volumes on a Sun StorEdge 6920 storage array or onto another CLARiiON system.

This flexibility provides greater ease of implementation because you use a common replication implementation interface, regardless of the platform. It also provides cost savings and a lower TCO because you can retain proven methods and don't have to license and learn different replication products to implement the Oracle database replication.

9.0 Glossary

Appliance-based data services: Data services that reside in intelligent, network-based appliances to enable replication between heterogeneous storage resources.

Array-based replication technology: Data services that reside within storage arrays to enable replication between storage resources.

Cluster integration: Provides functional capability to enable host cluster software to recognize and failover to recovery site storage resources using the capabilities provided by the storage systems.

DWDM: Dense Wavelength Division Multiplexing; aggregates up to 16 different laser light wavelengths onto a fibre optic path.

Host-based replication technology: Data services that reside in host server systems to enable replication between heterogeneous storage resources.

Hot backup mode: An Oracle term referring to a special mode into which a database is put prior to being backed up.

Hot standby mode: A secondary database with active applications and mirrored database content available for immediate business continuance support.

Hybrid network-array-based data services: Data services that reside within storage arrays to enable replication between heterogeneous storage resources.

Redo log: A set of files that record all changes made to an Oracle database. Redo logs are created when the active logs in memory are written to disk.

Rehosting: The migration of applications from one host, such as a legacy IBM mainframe, to another host, such as Sun systems.

Remote Thin Scripting Client: Also called a remote CLI client, runs the command-line interface on any qualified host in a network.

Solaris Containers: Technology that provides separate virtual instances of the Solaris Operating Environment running on a single machine.

SONET/SDH: Synchronous Optical NETWORK (SONET) and Synchronous Digital Hierarchy (SDH) refer to a group of fiber-optic transmission rates that can transport digital signals with different capacities.

Standby Database model: A secondary database with mirrored primary content for business continuance or recovery. Standby mode means that the secondary database is in a passive or read only state instead of in an active or read/write-enabled state.

Switch-based replication technology: Data services that reside in intelligent SAN switches to enable replication between heterogeneous storage resources.

Sun Microsystems, Inc. 4150 Network Circle, Santa Clara, CA 95054 USA **Phone** 1-650-960-1300 or 1-800-555-9SUN **Web** sun.com

© 2006 Sun Microsystems, Inc. All rights reserved. Sun, Sun Microsystems, the Sun logo, Sun StorEdge Data Replicator, Sun StorEdge 6920 are trademarks, registered trademarks, or service marks of Sun Microsystems, Inc. in the U.S. and other countries.

© 2006 Oracle. All rights reserved. Oracle, Oracle Data Guard are trademarks, registered trademarks, or service marks of Oracle in the U.S. and other countries.

JT 0031 A 03/06 SunWin #470625

