



Sun StorEdge™ Data Mirror Software With Oracle Databases Usage Guide

For Sun StorEdge 6920 Systems

Sun Microsystems, Inc.
www.sun.com

Part No. 819-3327-10 (v1)
July 2005, Revision A

Submit comments about this document at: <http://www.sun.com/hwdocs/feedback>

Copyright 2005 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, California 95054, U.S.A. All rights reserved.

Sun Microsystems, Inc. has intellectual property rights relating to technology that is described in this document. In particular, and without limitation, these intellectual property rights may include one or more of the U.S. patents listed at <http://www.sun.com/patents> and one or more additional patents or pending patent applications in the U.S. and in other countries.

This document and the product to which it pertains are distributed under licenses restricting their use, copying, distribution, and decompilation. No part of the product or of this document may be reproduced in any form by any means without prior written authorization of Sun and its licensors, if any.

Third-party software, including font technology, is copyrighted and licensed from Sun suppliers.

Parts of the product may be derived from Berkeley BSD systems, licensed from the University of California. UNIX is a registered trademark in the U.S. and in other countries, exclusively licensed through X/Open Company, Ltd.

Sun, Sun Microsystems, the Sun logo, docs.sun.com, Sun StorEdge, and Solaris are trademarks or registered trademarks of Sun Microsystems, Inc. in the U.S. and in other countries.

All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the U.S. and in other countries. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc. Legato and Legato NetWorker are trademarks or registered trademarks of Legato Systems, Inc. Netscape, Netscape Navigator, and Mozilla are trademarks or registered trademarks of Netscape Communications Corporation in the United States and other countries.

The OPEN LOOK and Sun™ Graphical User Interface was developed by Sun Microsystems, Inc. for its users and licensees. Sun acknowledges the pioneering efforts of Xerox in researching and developing the concept of visual or graphical user interfaces for the computer industry. Sun holds a non-exclusive license from Xerox to the Xerox Graphical User Interface, which license also covers Sun's licensees who implement OPEN LOOK GUIs and otherwise comply with Sun's written license agreements.

Legato NetWorker is a trademark or registered trademark of Legato Systems, Inc.

Mozilla and Netscape Navigator are trademarks or registered trademarks of Netscape Communications Corporation in the United States and other countries.

U.S. Government Rights – Commercial use. Government users are subject to the Sun Microsystems, Inc. standard license agreement and applicable provisions of the FAR and its supplements.

DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID.

Copyright 2005 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, California 95054, Etats-Unis. Tous droits réservés.

Sun Microsystems, Inc. a les droits de propriété intellectuelle relatant à la technologie qui est décrit dans ce document. En particulier, et sans la limitation, ces droits de propriété intellectuelle peuvent inclure un ou plus des brevets américains énumérés à <http://www.sun.com/patents> et un ou les brevets plus supplémentaires ou les applications de brevet en attente dans les Etats-Unis et dans les autres pays.

Ce produit ou document est protégé par un copyright et distribué avec des licences qui en restreignent l'utilisation, la copie, la distribution, et la décompilation. Aucune partie de ce produit ou document ne peut être reproduite sous aucune forme, par quelque moyen que ce soit, sans l'autorisation préalable et écrite de Sun et de ses bailleurs de licence, s'il y en a.

Le logiciel détenu par des tiers, et qui comprend la technologie relative aux polices de caractères, est protégé par un copyright et licencié par des fournisseurs de Sun.

Des parties de ce produit pourront être dérivées des systèmes Berkeley BSD licenciés par l'Université de Californie. UNIX est une marque déposée aux Etats-Unis et dans d'autres pays et licenciée exclusivement par X/Open Company, Ltd.

Sun, Sun Microsystems, le logo Sun, docs.sun.com, Sun StorEdge, et Solaris sont des marques de fabrique ou des marques déposées de Sun Microsystems, Inc. aux Etats-Unis et dans d'autres pays.

Toutes les marques SPARC sont utilisées sous licence et sont des marques de fabrique ou des marques déposées de SPARC International, Inc. aux Etats-Unis et dans d'autres pays. Les produits portant les marques SPARC sont basés sur une architecture développée par Sun Microsystems, Inc. Legato et Legato NetWorker sont des marques de fabrique ou des marques déposées de Legato Systems, Inc. Netscape, Netscape Navigator, et Mozilla sont des marques de Netscape Communications Corporation aux Etats-Unis et dans d'autres pays.

L'interface d'utilisation graphique OPEN LOOK et Sun™ a été développée par Sun Microsystems, Inc. pour ses utilisateurs et licenciés. Sun reconnaît les efforts de pionniers de Xerox pour la recherche et le développement du concept des interfaces d'utilisation visuelle ou graphique pour l'industrie de l'informatique. Sun détient une licence non exclusive de Xerox sur l'interface d'utilisation graphique Xerox, cette licence couvrant également les licenciés de Sun qui mettent en place l'interface d'utilisation graphique OPEN LOOK et qui en outre se conforment aux licences écrites de Sun.

LA DOCUMENTATION EST FOURNIE "EN L'ÉTAT" ET TOUTES AUTRES CONDITIONS, DECLARATIONS ET GARANTIES EXPRESSES OU TACITES SONT FORMELLEMENT EXCLUES, DANS LA MESURE AUTORISEE PAR LA LOI APPLICABLE, Y COMPRIS NOTAMMENT TOUTE GARANTIE IMPLICITE RELATIVE A LA QUALITE MARCHANDE, A L'APTITUDE A UNE UTILISATION PARTICULIERE OU A L'ABSENCE DE CONTREFAÇON.



Please
Recycle



Adobe PostScript

Contents

- 1. Overview 1**
 - About Sun StorEdge Data Mirror Software 1
 - Sun StorEdge Data Mirror Features 2
 - How Sun StorEdge Data Mirror Software Works 3
 - Example Uses of Sun StorEdge Data Mirror Software 6
 - Using Sun StorEdge Data Snapshot and Data Mirror Software Together 7
 - Benefits 7
 - Restrictions 7

- 2. Requirements, Planning, and Installation 9**
 - System Requirements 9
 - Planning Mirror Components 10
 - Creating Mirrored Volumes 12
 - ▼ To Create a Mirrored Volume 12
 - Installing Oracle 15
 - Oracle Installation Considerations 15
 - ▼ To Prepare for Oracle Installation 15
 - ▼ To Install Oracle Binaries 16

- 3. Using Sun StorEdge Data Mirror Software With Oracle Software 17**

Protecting an Oracle Database With Mirroring and Snapshots	17
▼ To Split a Component From a Mirror	18
▼ To Rejoin a Component to the Mirror	19
Working with Backups and Sun StorEdge Data Mirror and Snapshot Software	20
About Snapshots and Oracle Databases	20
▼ To Perform a Backup and Create a Snapshot	21
▼ To Restore the Database From Tape	23
Using a Mirror of the Database for Other Purposes	24
Glossary	27

Overview

This document provides guidelines for using the Sun StorEdge™ Data Mirror software in conjunction with Oracle® databases.

This chapter describes the Sun StorEdge Data Mirror software, discusses considerations in planning for its implementation, and explores a variety of its uses. It also discusses using Sun StorEdge Data Mirror and Snapshot software together. This chapter contains the following sections:

- “About Sun StorEdge Data Mirror Software” on page 1
- “Sun StorEdge Data Mirror Features” on page 2
- “How Sun StorEdge Data Mirror Software Works” on page 3
- “Example Uses of Sun StorEdge Data Mirror Software” on page 6
- “Using Sun StorEdge Data Snapshot and Data Mirror Software Together” on page 7

About Sun StorEdge Data Mirror Software

Sun StorEdge Data Mirror software is designed to enable you to create mirrors that are full copies of application data volumes. Mirrors can be used for backup, application testing, and other business processing. Sun StorEdge Data Mirror software can be combined with Oracle database software to produce readily available copies of critical application data.

For example, if you require daily reports for accounting purposes but do not want to interfere with ongoing transaction processing, you can create mirrors of the database volumes. You can then split the mirrors (temporarily separating the mirror

components) to make mirrored components available to a report-processing server, where they become available for developing reports. When the report-processing is completed, the components can be rejoined and synchronized.

To help ensure that Oracle databases and Sun StorEdge Data Mirror software perform together as expected, Oracle provides a suite of tests, under the Oracle Storage Compatibility Program (OSCP). These tests are used to validate the compatibility of mirroring software solutions with Oracle databases. These tests were performed according to the Oracle approved scenarios with the Sun StorEdge Data Mirror software.

Sun StorEdge Data Mirror Features

The Sun StorEdge Data Mirror software has a number of features:

- **Support for two, three, or four mirror components** – Using two-way, three-way, or four-way mirrors, you can achieve multiple objectives simultaneously, such as assigning two of the mirror components for data redundancy while the third component is used for an operation like backup.
- **Read/write support for split components** – Read and write access to split components means that virtually any type of testing or analysis can be applied to split components without jeopardizing your primary production data. Sun StorEdge Data Mirror software tracks changes to both the mirror and split component and updates appropriately when a rejoin or reverse rejoin is initiated.
- **Rapid rejoins** – When a component is split, Sun StorEdge Data Mirror software keeps track of changed blocks, allowing only the changes to be updated when the split component is later rejoined. The result is a quick update process.
- **Reverse rejoins** – If the mirror is corrupted while a component is split, the split component can assume the role of the mirror, enabling an immediate restart of applications. The result is an instant restore. The corrupt mirror component is then updated to reflect the data that is on the split component. Note the difference between a rejoin and a reverse rejoin.
 - Rejoin – The split component is updated to be identical to the mirror.
 - Reverse rejoin – The mirror is updated to be identical to the split component.
- **Break operations and data migration** – Sun StorEdge Data Mirror software enables you to reduce scheduled maintenance through non-disruptive data migration. You can migrate data while it is in use by creating a mirror of an existing volume to a new location, pointing the application to the new volume, and then breaking the mirror to remove the original volume. This data migration capability can also be used to move data to a different storage tier for archiving and cost reduction, to retire aging storage without impacting operations, and to move data to higher performing storage.

- **Sun StorEdge Data Mirror software coupled with existing Sun and non Sun storage** – Sun StorEdge Data Mirror software can be used with all storage supported on the Sun StorEdge 6920 system, which includes a variety of other Sun storage products as well as select products from other vendors such as EMC and HP. This enables new value added capabilities on existing storage assets.
- **Legacy Volumes and Sun StorEdge Data Mirror software** – Sun StorEdge Data Mirror software can be used to create a mirror between an existing storage area network (SAN) volume and a Sun StorEdge 6920 system virtual volume. This enables you to perform a live migration of data from the SAN into the virtualized Sun StorEdge 6920 system environment. Establishing legacy volumes requires a short disruption of the application hosts. Therefore, the migration can be accomplished at any time you choose while the applications are in use.
- **Scripting interface** – You can create scripts to automatically create and manage mirrors using the remote scripting CLI client. For example, using the remote scripting CLI client, you can set `cron` jobs to automatically split a component, mount a backup server, and then rejoin the split component when the backup is complete.

How Sun StorEdge Data Mirror Software Works

Also called RAID level 1 and real-time copy, mirroring is a form of storage whereby two or more identical copies of data are maintained on separate media. The Sun StorEdge 6920 system enables you to create and manipulate a mirror with up to four equal mirror components that can be accessed independently. You can also update each independent copy of the mirror's data set by performing a variety of operations on any mirror component within the mirror.

A mirror is made up of equal, independent mirror components periodically synchronized through resilvering to reflect the same data set. Resilvering synchronizes the mirror so that it comprises equal, independent copies of the same data set when resilvering is completed. Resilvering occurs when you perform operations on a mirror component such as rejoin and reverse rejoin. Although the first mirrored volume (also called the original volume and the primary component) you create establishes the initial data set for the mirror that is copied to mirror components you add, the various operations that you perform on the mirror components affect the entire mirror. You can use any mirror component for any of the following: as a picture of your data at a specific time, to update the data set for the mirror, or for data-redundancy protection.

Note – The mirror assumes the World Wide Name (WWN) of the first mirrored volume you create to establish the mirror. Any existing volumes you add to the mirror preserve the WWN of the first mirrored volume you create.

A mirror consists of the following:

- **Mirror** – A volume with a unique WWN and volume that can be mapped to hosts for use.
- **Mirror components** – Individual full copies of the volume. Mirror components each have a unique WWN and volume name. Mirror components cannot be mapped to hosts for use unless they are first split.
- **Mirror log** – A disk space region that maintains status and metadata information about the mirror. The mirror log contains no actual users data but instead contains pointers and bitmaps to track create, rejoin, reverse rejoin, and other operations. Sun StorEdge Data Mirror software automatically sets up and manages its own mirror log space. Additionally, the mirror log is itself mirrored for redundancy.

Mirroring can be an important part of a database migration strategy. The Sun StorEdge 6920 system, Version 3.0, enables you to virtualize external, heterogeneous storage devices in the midrange storage category. Suppose you need to migrate data from legacy storage to the Sun StorEdge 6920 system. You can mirror the data on a legacy volume to a volume on the Sun StorEdge 6920 system, during which time the database application can continue to access the legacy volume. When mirroring is complete, the mirror components have been resilvered so that the legacy data also exists as an independent copy on the new mirrored volume. The database application can now be instructed to mount the new mirrored volume. Since you can access your data on the new mirrored volume, you can choose to delete the legacy volume and perhaps repurpose the legacy storage.

You can perform the following actions on mirror components within a mirror:

- **Split** – You can separate a mirror component from the mirror, with the intent to rejoin it later. Because a split mirror component continues to be tracked and maintains an association with other components within the mirror, it counts toward the limit of four mirror components per mirror. You can use a split mirror component to create a picture of data at the time of the split. Because mirror components are independent copies of the same data set, a split mirror component offers a picture of data that, unlike a snapshot, is not dependent upon snapshot reserve space.

During the split condition, Sun StorEdge Data Mirror software logs a record of the blocks changed on the mirror (writes by the primary application that have occurred since the split) and blocks changed on the split component (writes to the

split component by a secondary application that mounts and uses it). The result is a complete record of the differences between the mirror components and the split component.

- **Rejoin** – You can join a split mirror component back into the mirror. Resilvering immediately begins, and the previously split mirror component becomes identical to (is synchronized to match) its counterparts in the mirror.
- **Reverse rejoin** – You can perform a reverse rejoin action on a split mirror component when you need to use its data as the basis for the mirror as a whole. The result is that every mirror component in the mirror becomes identical to the previously split mirror component, and the data set represented in the mirror matches that of the previously split mirror component.
- **Break** – You can remove a mirror component from the mirror, voiding its association with the other mirror components. The broken mirror component becomes a standalone volume in the system, and synchronization with the other mirror components is not maintained. You might break a mirror component in order to create an independent copy of an existing volume while the mirror is active and save it as a picture of the data at a specific time.

Note – If you choose to add a previously broken mirror component to the mirror, you need to add it as a normal volume. This is different from the split operation, because any changes in data for a split mirror component continue to be tracked, and a relationship between the split component and other mirror components is maintained. Typically, you break a mirror component when your intent is to remove it from the mirror, and you split a mirror component when your intent is to separate it and rejoin it later.

If you break the last remaining component of a mirror, the result is a non-mirrored volume with the same name as the mirror. The name of the previous mirror component is no longer used.

- **Force break** — You can execute a force break operation to remove a mirror component that has not been 100 percent resilvered and is therefore not synchronized with the other components in the mirror.

If you know that a mirror component's data has been compromised (during a fire or other catastrophe, for instance) and data integrity is lost, you can force the break of a mirror component, regardless of the percentage of resilvering that has occurred, to protect the integrity of the mirror's data. Just as with a break operation, the mirror component for which you implement a force break becomes a standalone volume in the system. However, the condition of its data is degraded, and it is inaccessible after it is removed from the mirror during the resilvering process.

Example Uses of Sun StorEdge Data Mirror Software

Sun StorEdge Data Mirror software enables you to improve your business operations by “cloning” data sets. This enables you to run more applications in parallel, which can result in a better use of your information. Following are examples of circumstances in which Sun StorEdge Data Mirror software can be particularly useful for Oracle database applications:

- **Business Intelligence and Data Warehousing (BIDW)** – By sharing a Sun StorEdge 6920 system between transaction and analysis database systems, you can use Sun StorEdge Data Mirror software to automate the transfer of the latest data from the transaction system to the warehouse. You can do this through a periodic automated process (such as nightly) to help ensure that the latest data available is used to help make business decisions.
- **Reporting systems** – By sharing a Sun StorEdge 6920 system between the transaction and reporting database systems, you can use Sun StorEdge Data Mirror software to automatically transfer data to the reporting system on a periodic basis.
- **Restarting an application on a mirror** – Recovering an application in the event of a virus or other form of corruption can be a lengthy process. However, you can restart an application immediately by using a split mirror component. The result is an instant restore, minimizing downtime.
- **Enterprise Resource Planning (ERP) or other new application testing** – Large-scale, complex database applications like ERP often require a multiple phase implementation that adds new modules to live systems over time. Minimizing the risk of these rollouts is essential to the stability of these systems. Using split components to provide a current but isolated copy of data enables real-world testing without jeopardizing the production systems. Additionally, if a test corrupts data, split components can be rejoined to synchronize the components, then split again to accelerate preparation for a retest. Shortening test cycles can shorten the time for a production rollout or provide more testing cycles to improve application stability.

Using Sun StorEdge Data Snapshot and Data Mirror Software Together

When designing a solution using Sun StorEdge Data Snapshot and Data Mirror software together, it is important to understand how they interact. This section describes the benefits of and restrictions on using these two products together.

Benefits

You can use these two products together to accomplish the following:

- Higher levels of data protection
 - Use Sun StorEdge Data Mirror software to help ensure that two full and separate copies of the Oracle database are maintained on disk and thereby to protect data against loss from a major hardware failure.
 - Use Sun StorEdge Data Snapshot software to create frequent snapshot volumes from which you can roll back from in the event of a software or user error.
- Support of multiple uses for the same data
 - Use Sun StorEdge Data Mirror software to provide full independent copies for I/O-intensive operations like data mining. This avoids the possible performance impact that could occur through the use of snapshots.
 - Use Sun StorEdge Data Snapshot software to create frequent snapshot volumes for use in backup processes and rollback operations.
- Application testing
 - Use Sun StorEdge Data Mirror software to create a full and recent copy of production application data.
 - After breaking off a mirror component for testing, use Sun StorEdge Data Snapshot software on the new volume to run a test, roll back the data to the original, and then run the test again.

Restrictions

The following restrictions apply to the use of these two products together:

When snapshots of a mirror are created, they are associated with the mirror component, to which some restrictions apply. Split operations are not allowed for that mirror component unless the snapshots are first deleted. You must use a

different component for split and rejoin operations. Additionally, although you can create a snapshot of the mirror, you cannot create a snapshot of the split components.

If a software or user error causes data corruption, you can perform a snapshot rollback. However, to do so when also using Sun StorEdge Data Mirror software, you must break the mirror, conduct the rollback, and then re-create the mirror. Although this involves more steps than a snapshot rollback without a mirror, it is still a disk-based restore with major advantages over tape restore.

- If you want to perform a break operation, the method you use depends on how you want to treat the snapshots associated with the mirror:
 - If you want to retain the snapshots and have them continue to be associated with the mirror, you must perform the break operation on a mirror component that does not have the snapshots associated with it. After the break operation is complete, you are free to create snapshots, since the broken component is now just a normal volume. You can use this method to create a full copy of data for application testing, break the mirror, and then use snapshot on the new volume to manage versioning on the test copy of the data.
 - If you want the snapshots to go with the component that is removed from the mirror, perform the break operation on the mirror component that has snapshots associated with it. The snapshots are associated with the newly created volume and are no longer associated with the active mirror. At this stage new snapshots can be established on the mirror volume.
- If an existing volume is going to be added to an existing mirror, any snapshots on the volume to be added as a mirror component must first be deleted.

Requirements, Planning, and Installation

This chapter describes system requirements and installation procedures for Oracle software, as well as information on creating primary volumes and mirrors. It contains the following sections:

- “System Requirements” on page 9
 - “Planning Mirror Components” on page 10
 - “Creating Mirrored Volumes” on page 12
 - “Installing Oracle” on page 15
-

System Requirements

The following table lists the hardware and software requirements for running both the Sun StorEdge Data Mirror and Oracle database software.

Hardware	Sun StorEdge 6920 system
Software	Solaris™ 8 (update 4 or higher) or Solaris 9 Operating System
For CLI interface	Remote scripting command-line interface (CLI)
For browser interface	One of the following: <ul style="list-style-type: none">• Netscape Navigator™ version 7.x• Microsoft Internet Explorer version 5.x• Mozilla™ 1.2.1

For further information, refer to the *Sun StorEdge 6920 System Release Notes*.

For additional information on the Sun StorEdge 6920 system, see the following documents, which are supplied with the software.

Description	Title	Part Number
Installation information	<i>Sun StorEdge 6920 System Getting Started Guide</i>	817-5227-nn
Release information	<i>Sun StorEdge 6920 System Release Notes</i>	817-5229-nn
man pages	sscs(1M)	N/A

Planning Mirror Components

A mirror contains up to four independent, identical copies of data in separate yet equal mirror components that are accessed to manage the mirror as a whole. Mirror components are individual parts of the mirror, identical in both data and capacity. An action you perform on one mirror component can cause all other mirror components to synchronize to it, and vice versa.

For example, when you rejoin a split mirror component back into the mirror, it is synchronized to reflect the other mirror components. When the resilvering process completes, the mirror is once again an entity in which each component is separate subentity with identical data.

To create a mirror, you can create a new mirrored volume or mirror an existing volume and then add other volumes as equal mirror components. Before starting the process, you must determine the following:

- The mirror components that make up the mirror. Keep in mind the following restrictions:
 - Because mirror components must be identical in size, you must make sure each mirror component you add is the same size as the other components in the mirror.
 - You cannot mirror across storage domains.
- The capacity for the original mirrored volume (primary component). This capacity is the same for new mirror components that are added to the mirror.
- The isolation policy you require for mirror components:
 - Required – Mirror components must reside on different arrays. When mirror components exist on different arrays, redundancy protection is at its highest for data preservation.

- Optional – Mirror components can reside on the same array. If you know that a mirror component already shares an array with parts of the mirror, you must set the isolation policy to Optional. When mirror components share an array, the redundancy factor is reduced for the data set associated with each mirror component, should the array or access to the array fail.
- The resilver priority setting you require. Resilvering synchronizes each component to be identical with the others. It occurs each time you perform an operation on the mirror. This setting determines whether I/O activity or resilvering is a higher priority for the system:
 - High – The system attempts to resilver the mirror quickly, but this might noticeably slow down host I/O activities.
 - Medium – The system balances I/O activities with the resilvering process, yielding an incremental, steady pace for all jobs that are executed.
 - Low – The system takes longer to resilver the mirror, but it processes host I/O activities more quickly.
- Whether you want to use split mirror components or volume snapshots for pictures of your data at a specific time:
 - Split mirror components – When you want to take a picture of the data set represented in a mirror, you can split a mirror component from the mirror and save it as an independent copy or reference point of your data at that time. Although you cannot roll back mirrored volume snapshots (because their data is unchangeable), you can split a mirror component and later rejoin it to the mirror, and you can even perform a reverse rejoin action on the previously split component if you want the other mirror components to synchronize to its data. The split mirror component has no snapshot reserve space limitations.
 - Volume snapshots – The volume snapshot feature is most useful as a method of tracking data through precise, point-in-time pictures that reside in snapshot reserve space. You can modify the snapshot reserve space of a mirrored volume as you would any other volume. The older a volume snapshot, the more the data it represents is likely to have changed. If the volume is expected to change substantially between snapshots, using split mirror components may be more appropriate. More information on using Sun StorEdge Data Snapshot software with Oracle Databases can be found in: *Sun StorEdge Data Snapshot Software With Oracle Databases Usage Guide for Sun StorEdge 6920 Systems*.
- If the selected mirror components do not have snapshot reserve space and you need to take snapshots, how to allocate snapshot reserve space.

Creating Mirrored Volumes

You manage your system's physical disks as a pool of storage space for creating volumes. Volumes are “containers” into which applications, databases, and file systems can put data. Volumes are created from the virtual disks in storage pools.

You can create up to 32 volumes on each virtual disk. While creating a volume, it is possible to create a mirrored volume with snapshot reserve space. After you create the volumes, assign each initiator in your environment to one of the volumes.

An initiator is the protocol-specific physical port that initiates the I/O exchanges with the system. An initiator is the Fibre Channel (FC) port that is identified by a port World Wide Name (WWN). If a host is connected to the system by two host bus adapters (HBAs), the system handles the host as two different initiators. Fibre Channel array LUN masking and mapping uses initiator port identifiers to authenticate storage customers.

For additional information on creating volumes and mapping initiators, refer to the *Sun StorEdge 6920 System Getting Started Guide*, the online help, or the `sscs(1M)` man page.

▼ To Create a Mirrored Volume

Follow these steps to create a mirrored volume with snapshot reserve space.

1. Type the `create` command on the command line.

```
# sscs create -m pool1, pool2, [pool3], [pool4] -s size TB|GB|MB|B
-P pool-name [-C snapshot-count 0|1|2|3|4|5|6|7|8] [-L snapshot-
level low|verylittle|little|average|high|full] -T mirror-
component-name [-f volume|snapshot] [-i required|optional]
[-d volume-description] [-g pool-name] [-v resilver-priority
low|medium|high] [-S storage-domain-name]
volume mirrored-volume-name
```


The following options and arguments apply to the `create` command.

TABLE 2-1 `create` Command Options and Arguments

Options and Arguments	Description
<code>-m</code>	The storage pools, from two to four, from which the mirrored volume and associated mirror components will be created. The number of storage pools entered equals the number of mirror components. Storage pool names can be repeated. For example, if you want three components from a given storage pool, you would enter that storage pool's name three times.
<code>-s</code>	The size of the volume that is to be created. The volume size can be configured in terabytes, gigabytes, megabytes, or bytes.
<code>-P</code>	The storage pool from which the snapshot reserve space is to be allocated. The specified storage pool must be in the storage domain specified by the <code>sdomain</code> option.
<code>[-C]</code>	The snapshot count for the snapshot reserve. This value is used in conjunction with the <code>-L</code> option to calculate the size of the snapshot reserve.
<code>[-L]</code>	The amount of data overwriting expected between snapshots. This value is used in conjunction with the <code>-C</code> option to calculate the size of the snapshot reserve. The snapshot levels correspond to the following percentages: <ul style="list-style-type: none">• <code>low</code> - 10 percent• <code>verylittle</code> - 25 percent• <code>little</code> - 40 percent• <code>average</code> - 50 percent• <code>high</code> - 75 percent• <code>full</code> - 100 percent
<code>-T</code>	The mirror component that is to be associated with the snapshot for the mirror. For this command, the components do not exist yet, so the names of the pools used to create the component (the <code>-m</code> option) should be entered.
<code>[-f]</code>	The policy setting specifying which should be favored, the volume or the snapshot.
<code>[-i]</code>	Specifies the array level isolation policy for the mirror components. <code>required</code> specifies that all mirror components must be from different physical arrays. <code>optional</code> specifies that all mirror components may share arrays. The default is <code>required</code> .
<code>[-d]</code>	An alphanumeric character description of the volume of up to 64 characters.

TABLE 2-1 create Command Options and Arguments (Continued)

Options and Arguments	Description
[-g]	The storage pool to be used for the mirror logs. The specified pool must be in the storage domain specified by the -s option argument, and must be made up of at least two virtual disks and sufficient free space to host the two 2-MB mirror logs. If this option is not specified, the system chooses the log's storage pool.
[-v]	The resilver priority setting, which can be low, medium, or high. The default is medium.
[-S]	The name of the storage domain in which the volume is located. If nothing is specified, then the DEFAULT storage domain is used.
volume	The volume name, which can be up to 16 alphanumeric characters.

2. (Optional) Create the snapshot reserve space to be associated with this mirror.

For example:

```
# sscs create -m PoolA, PoolA, PoolA -i optional -C 5 -L average
-T PoolA -P MyPool -S MyDomain volume MyMirror
```

3. Select any initiators that you want to map to the volume and specify the LUN and permissions for each mapping.

```
# sscs map initiator initiator-name -P readwrite -S storage-
domain-name
-l LUN-ID -v volume-name
```

For example, if you want to create a volume that is connected to the Sun StorEdge 6920 system by two HBAs, then your volume would have two initiators. Each HBA must be mapped to the volume.

```
# sscs map initiator 4800_sc2 -P readwrite -S mr19
-l 200 -v MyMirror
# sscs map initiator 4800_sc3 -P readwrite -S mr19
-l 200 -v MyMirror
```

Installing Oracle

The following information, which is provided as an overview of installing the Oracle binaries, uses Oracle 9i as an example. Your installation is likely to vary. See the *Oracle Installation Guide* for further installation details.

Oracle Installation Considerations

Before you install the Oracle binaries, you should be aware of the following:

- You can use the path `/oracle/9.x.x` (where `x` is a version number) as the `ORACLE_HOME` environment variable. The software then creates various directories under `$ORACLE_HOME`. The `/bin` directory contains all binaries, and the `/rdbms/admin` directory contains all utility SQL files.
- The installation software automatically creates a default database named `starter`. You must assign a four-character `ORACLE_SID` system identifier (SID) name to the database.
- When the installation is complete, you can use the `dbassist` utility to create numerous Oracle databases on the server. The number of databases per server is restricted by resource availability.

▼ To Prepare for Oracle Installation

You must perform the following steps before you install the Oracle binaries.

1. **Add the following entries to the `/etc/system` file by using a text editor:**

```
set semsys:seminfo_semmni=100
set semsys:seminfo_semmns=256
set semsys:seminfo_semmsl=256
set shmsys:shminfo_shmmax=4294967295
set shmsys:shminfo_shmmin=1
set shmsys:shminfo_shmmni=100
set shmsys:shminfo_shmseg=10
```

2. **Create the directory `/opt/bin`.**
3. **Create the group `oinstall`.**
4. **Create the user `oracle` and attach the `oinstall` group to `oracle`.**

5. **Have a mount point ready and make `oracle` the owner.**
6. **Reboot the server.**

▼ To Install Oracle Binaries

1. **Log in to the server as user `oracle`.**
2. **Insert the Oracle 9.x Universal Install CD into the CD-ROM drive.**
3. **Accept the Oracle 9i Server Installation option and follow the instructions.**

Using Sun StorEdge Data Mirror Software With Oracle Software

Using the Sun StorEdge Data Mirror software, you can easily and quickly mirror your databases for backup purposes or for other uses. This enables you to use a copy of the database, instead of your live database, to do testing, reporting, and analysis. You can also acquire a snapshot of the Oracle database at a specific point in time whenever you choose.

This chapter describes the tasks involved in managing mirrors of Oracle databases. The task descriptions provide the steps using the command-line interface (CLI), which enables you to create scripts to automate the tasks. However, all of these tasks can also be performed through the Sun StorEdge 6920 Configuration Service software.

This chapter contains the following sections:

- “Protecting an Oracle Database With Mirroring and Snapshots” on page 17
- “Working with Backups and Sun StorEdge Data Mirror and Snapshot Software” on page 20
- “Using a Mirror of the Database for Other Purposes” on page 24

Protecting an Oracle Database With Mirroring and Snapshots

You can implement data protection for your Oracle database by combining Sun StorEdge Data Mirror and Snapshot software. The tasks involved in accomplishing this are the following:

1. Creating two or more separate copies of data to protect against loss from a major hardware failure
2. Splitting the mirror
3. Backing up the split component to tape
4. Rejoining the mirror

▼ To Split a Component From a Mirror

You can perform a split operation to temporarily separate a mirror component from a mirror in order to perform a backup or for other purposes. The system persistently records changes you make to the split component and the mirror. You can view the split mirror component as a picture of the data for the mirror at the time of the split. In this way, a split mirror component is similar to a volume snapshot. However, a split mirror component acts as an independent copy of the data set represented in the mirror, whereas a snapshot is a copy that is dependent upon snapshot reserve space.

Note – You cannot split a mirror component that has volume snapshots.

1. Issue the command to split a component from a mirror.

```
# sscs modify -t [-S storage-domain-name] volume component-volume-name
```

The following options and arguments apply to the `modify` command.

Options and Arguments	Description
<code>-t</code>	Splits one component volume from a mirrored volume.
<code>[-S]</code>	The storage domain volume operands. If no value is specified, the system uses the default storage domain.
<code>volume</code>	The name of the component volume.

For more information on additional options to the `sscs modify` command see the `sscs(1M)` man page.

2. **Map the split component to the appropriate initiator for the secondary database application host.**

```
# sscs map initiator initiator-name -P readwrite -S storage-  
domain-name -l LUN-ID -v component-volume-name
```

3. **Mount the volume on the secondary database application host.**
4. **Start the secondary database instance.**

▼ To Rejoin a Component to the Mirror

You can perform a rejoin operation on a split mirror component and bring it back into the mirror after the backup is completed. When you do so, the system immediately begins to resilver (synchronize) the split mirror component to be identical to the other mirror components in the mirror.

1. **Stop the secondary database instance.**
2. **Unmount the volume on the secondary database application host.**
3. **Map the split component to the appropriate initiator for the primary database application host.**

```
# sscs map initiator initiator-name -P readwrite -S storage-  
domain-name -l LUN-ID -v component-volume-name
```

4. **Issue the command to rejoin the split mirror component to the mirror.**

```
# sscs modify -j [-S storage-domain-name] volume component-volume-  
name
```

The following options and arguments apply to the `modify` command.

Options and Arguments	Description
<code>-j</code>	The argument that joins the specified split component volume back to the mirrored volume.
<code>[-S]</code>	The storage domain volume operand. If no value is specified, the system uses the default storage domain.
<code>volume</code>	The name of the component volume.

For more information on additional options to the `sscs modify` command see the `sscs(1M)` man page.

Working with Backups and Sun StorEdge Data Mirror and Snapshot Software

You can also use Sun StorEdge Data Snapshot software to take frequent snapshots that can also be mounted for backup purposes. This section describes the general steps you must follow to perform a hot or cold backup, create a snapshot of that backup, and then restore the database to a volume. You can use hot backup or cold backup mode:

- Putting the Oracle software into hot backup mode helps ensure that the database is in a consistent state while you take snapshots.
- Putting the Oracle software into cold backup mode takes the database offline and help ensure that it is in a consistent state while you take snapshots.

For further information regarding backups, refer to your Oracle documentation.

About Snapshots and Oracle Databases

The snapshot application works at a volume level independent of higher level applications. To maintain data integrity, you should ensure that the Oracle database is in a known state before creating a snapshot. To do this, place the Oracle database in cold backup or hot backup mode helping to ensure that if an Oracle instance is started on the snapshot, the database is in a consistent state.

The Oracle database needs to be quiesced only long enough to allow creation of a snapshot. When you take the snapshot, the system briefly suspends I/O to the primary volume, establishes the necessary pointers for the snapshot, and then resumes I/O to the primary volume. After the snapshot is created, it can be mapped to another server on the storage area network (SAN) through the unique World Wide Name (WWN) and ID that were assigned to it when it was created.

▼ To Perform a Backup and Create a Snapshot

1. Decide which mirrored volumes require snapshots.

Choose as the parent volumes the volumes where the production data files, control files, redo logs, and archived logs reside.

2. If you are performing a cold backup, shut down the database. If you are performing a hot backup, perform the following steps from the SQL*Plus prompt:

a. Put all of the tablespaces in hot backup mode (refer `v$tablespaces`).

```
SQL> alter tablespace tablespacename begin backup;
```

b. Take an SQL copy of the control file.

```
SQL> alter database backup controlfile to filename;
```

c. Archive the latest redo logs (refer `v$logfile`, `v$log`).

```
SQL> alter system switch logfile;
```

3. Issue the command to take the snapshot.

- If you created the snapshot reserve space when you created the mirror, use this command:

```
# sscs snapshot -v source-volume-name [-d volume-description]  
[-S storage-domain-name] volume new-snapshot-name
```

- If you did not create the snapshot reserve space when you created the mirror, enter the command:

```
# sscs snapshot -v source-volume-name -P pool
-C snapshot-count 0|1|2|3|4|5|6|7|8 -L snapshot-level
low|verylittle|little|average|high|full -T mirror-component
[-f volume|snapshot] [-d volume-description] [-S storage-domain-name]
volume new-snapshot-name
```

The following options and arguments apply to the snapshot command.

TABLE 3-1 snapshot Command Options and Arguments

Options and Arguments	Description
-v	The volume for which the snapshot is to be taken.
-P	The name of the snapshot reserve space for a mirrored volume that does not have snapshot reserve space configured.
-C	The number of intended snapshots for the mirrored volume. This value is used in conjunction with the -L option to calculate the size of the snapshot reserve space.
-L	The amount of data overwriting expected between snapshots. This value is used in conjunction with the -C option to calculate the size of the snapshot reserve. The snapshot levels correspond to the following percentages: <ul style="list-style-type: none"> • low - 10 percent • verylittle - 25 percent • little - 40 percent • average - 50 percent • high - 75 percent • full - 100 percent
-T	The mirror component is to be associated with the snapshot for the mirror.
[-f]	The policy setting specifying which should be favored, the volume or the snapshot.
[-d]	Specify a volume description of up to 64 alphanumeric characters.
[-S]	The name of the storage domain in which the volume is located. If nothing is specified, then the DEFAULT storage domain is used.
volume	The name of the snapshot volume. After the snapshot volume is created, it can be treated like any other volume, except that it cannot be used to create another snapshot.

4. Map the snapshot to the appropriate initiator for the backup host.

```
# sscs map initiator initiator-name -P readwrite -S storage-  
domain-name -l LUN-ID -v snapshot-volume-name
```

5. From the SQL*Plus prompt, issue the following commands for all tablespaces to end the hot backup mode (refer v\$tablespace).

```
SQL> alter tablespace tablespace-name end backup;
```

6. Check the status of the snapshot by issuing the `sscs list volumes` command.

After the snapshot is created, it is available for use.

7. Mount the snapshot volume to the backup host.
8. Using a backup program such as VERITAS NetBackup or Legato NetWorker, back up the snapshot to tape. When the backup is complete, unmount the snapshot volume and map the initiator back to the database host.

▼ To Restore the Database From Tape

1. If the database is running, stop the database.

```
SQL> shutdown immediate;
```

2. Using a backup program such as VERITAS NetBackup or Legato NetWorker, restore the database from tape to the volume.

This enables you to restore a volume to the point in time at which the database was backed up and a snapshot was created. You can restore the database to any component in the mirror.

3. Verify that the volume has been restored by reviewing the VERITAS NetBackup or Legato NetWorker logs.
4. Bring the database to the mounted state from the SQL*Plus prompt.

```
SQL> startup mount initsid.ora;
```

5. Restore the database by doing one of the following.

- If the data file locations are different than in the original database, rename the data files in the new location (not the original database). For Example:

```
SQL> cd $ORACLE_HOME/bin/sqlplus
SQL> connect sys/manager as sysdba;
SQL> startup restrict mount;
SQL> alter database rename file 'original-name' to 'new-name';
SQL> recover automatic database;
SQL> alter database open;
```

- If the data file locations are the same as in the original database, rename the data files in the new location (not the original database). For Example:

```
SQL> cd $ORACLE_HOME/bin/sqlplus
SQL> connect sys/manager as sysdba;
SQL> startup open pfile=$DBS_DIR/initSNAP1.ora
```

6. Recover the database.

```
SQL> recover automatic database;
```

7. When recovery is complete, open the database.

Using a Mirror of the Database for Other Purposes

You can use a mirror of the database to run multiple application on a single set of data. This is particularly useful for I/O-intensive applications, such as data mining, decision support, reporting, and analysis. One method of using mirrors in this way involves doing the following:

1. Splitting the mirror of the database and using the split to mount and run another instance of the database, as describe in “To Split a Component From a Mirror” on page 18.
2. Periodically rejoining the split so that it can be synchronized with current data.

For example, you could write a script that would rejoin the mirror every night to be resynchronized, and then split a component from the mirror again for use by the secondary database application, as described in “To Rejoin a Component to the Mirror” on page 19.

When you rejoin a mirror, the system immediately begins to resilver (synchronize) the split mirror component to be identical to the other mirror components in the mirror. Before rejoining the mirror, you must first stop the database instance running on the split mirror, unmount the split mirror volume, and remap the initiator to the primary database application host.

Glossary

break	To remove a mirror component from the mirror, voiding its relationship with the other mirror components. The broken mirror component becomes a standalone volume in the system, and synchronization with the other mirror components is not maintained. See also mirror component and split.
cold backup mode	A backup mode in which you take the Oracle database that you are backing up offline, at least until the backup is complete.
component	See mirror component.
copy-on-write (COW)	A technique for maintaining a point in time copy of a collection of data by copying only data which is modified after the instant of replicate initiation. The original source data is used to satisfy read requests for both the source data itself and for the unmodified portion of the point in time copy. (SNIA) See also snapshot reserve space.
force break	To remove a mirror component (mirrored volume) from the mirror before the resilvering process is complete, causing the condition of its data to be degraded. Just as with a break operation, the mirror component for which you implement a force break becomes a standalone volume in the system; however, it is inaccessible and should probably be deleted. See also break, mirror component, and resilvering.
hot backup mode	A backup mode in which you quiesce the Oracle database long enough to start the point-in-time copy or update.
initiator	A system component that initiates an I/O operation over a Fibre Channel (FC) network. If allowed by FC fabric zoning rules, each host connection within the FC network has the ability to initiate transactions with the storage array. Each host in the FC network represents a separate initiator, so if a host is connected to the system through two host bus adapters (HBAs), the system identifies two different initiators (similar to multi-homed, Ethernet-based hosts). In contrast, when multipathing is used in round-robin mode, multiple HBAs are grouped together, and the multipathing software identifies the group of HBAs as a single initiator.

legacy volume	An entire logical unit number (LUN) on an external storage array that you can use in specific ways as if it were a local volume, while preserving the user data on that external storage array. You can apply the system's data services to a legacy volume; however, you cannot extend a legacy volume.
mirror	A special type of volume in the Sun StorEdge 6920 system, consisting of up to four separate yet equal mirror components you can access independently (through the mirror) and use to track changes to and update your data set and manage your data migration strategy.
mirror component	One of up to four individual copies of the same data set that constitute a mirror in the Sun StorEdge 6920 system. When you perform an action on a mirror component, the resilvering process synchronizes the mirror so that each component is an equal yet separate copy of the same data set. You can perform a variety of operations on a mirror component. See also break, rejoin, resilvering, and split.
mirroring	A form of storage — also called RAID Level 1, independent copy, and real-time copy — whereby two or more independent, identical copies of data are maintained on separate media. Typical mirroring technologies enable the cloning of data sets to provide redundancy for a storage system. The Sun StorEdge 6920 system enables you to create and manipulate up to four equal mirror components for each mirror and to access each component individually in order to track changes to and preserve the integrity of your data. See also mirror and mirror component.
mirror log	Area of the storage pool used to track the state of mirror components (such as resilvering progress) relative to the mirror as a whole.
OSCP	Oracle Storage Compatibility Program. A program that is designed to test the compatibility of storage vendor technologies by providing a test kit with which the integrity of the Oracle database can be validated with snapshot copies.
parent volume	The volume for which a snapshot is taken. See also snapshot.
point-in-time copy	A fully usable copy of a defined collection of data that contains an image of the data as it appeared at a single point in time. The copy is considered to have logically occurred at that point in time, but implementations may perform part or all of the copy at other times (e.g., via database log replay or rollback) as long as the result is a consistent copy of the data as it appeared at that point in time. Implementations may restrict point in time copies to be read-only or may permit subsequent writes to the copy. (SNIA) See also snapshot.
rejoin	To move a split mirror component back into the mirror such that, when the resilvering process is complete, the mirror component has been made identical to all other mirror components in the mirror. See also mirror component and split.

remote scripting CLI client	A command-line interface (CLI) that enables you to manage the system from a remote management host. The client communicates with the management software through a secure out-of-band interface, HTTPS, and provides the same control and monitoring capability as the browser interface. The client must be installed on a host that has network access to the system.
resilvering	The synchronization of mirror components such that, when the process is complete, the mirror comprises equal, independent copies of the same data set. Resilvering occurs when you have performed a rejoin or reverse rejoin operation on a mirror component.
resnap	To create a snapshot again and replace the old with the new.
reverse rejoin	To move a split mirror component back into the mirror such that, when the resilvering process is complete, all mirror components in the mirror have been made identical to the previously split mirror component. See also break, mirror component, rejoin, resilvering, and split.
rollback	The process by which a volume's data is reset to become identical to a snapshot taken of that volume.
snap	The act of creating a point-in-time data image of a volume using the data snapshot software.
snapshot	An instantaneous copy of volume data at a specific point in time. Snapshots are stored in snapshot reserve space on the (parent) volume for which they are taken.
snapshot reserve space	Storage space, taken from a pool, where the system stores snapshot copies of the parent volume's original data before it is overwritten.
split	To separate a mirror component from the mirror, with the intent to rejoin it later. A split component is counted toward the limit of four mirror components per mirror and continues to be tracked by the system as part of the mirror. See also mirror component, rejoin, and snapshot.
synchronization	The act of aligning or making entries be equivalent at a specified point in time. (SNIA). In the context of mirroring, see resilvering.
thin-scripting client	See remote scripting CLI client.
volume	A logically contiguous range of storage blocks allocated from a single pool and presented by a disk array as a logical unit number (LUN). A volume can span the physical devices that constitute the array, or it can be wholly contained within a single physical disk, depending on its virtualization strategy, size, and the internal array configuration. The array controller makes these details transparent to applications running on the attached server system.

**World Wide Name
(WWN)**

A unique identifier for a port, initiator, virtual disk, or volume, assigned by the system. The WWN of an object does not change throughout its lifetime and is never reused to name another object.