



Sun StorEdge™ Data Snapshot Software With Oracle Databases Usage Guide

For Sun StorEdge 6920 Systems

Sun Microsystems, Inc.
www.sun.com

Part No. 817-6985-10 (V2)
December 2004, Revision A

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

Copyright 2004 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, AnswerBook2, docs.sun.com, Sun StorEdge, Sun Enterprise, Java, 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.

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 2004 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, California 95054, Etats-Unis. Tous droits réservés.

Sun Microsystems, Inc. a les droits de propriété intellectuels relatants à la technologie qui est décrit dans ce document. En particulier, et sans la limitation, ces droits de propriété intellectuels 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, AnswerBook2, docs.sun.com, Sun StorEdge, Sun Enterprise, Java, 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.



Contents

- 1. Overview 1**
 - About the Sun StorEdge Data Snapshot Software 1
 - How Volume Snapshots Work 2
 - Sun StorEdge Data Snapshot Features and Benefits 2
 - Snapshot Reserve Space Requirements 3
 - Planning Snapshot Reserve Space 4
 - Example of Snapshot Reserve Space Configuration 5
 - About Snapshots and Oracle Databases 5
 - Example Uses of Sun StorEdge Data Snapshot Software 6

- 2. Requirements and Installation 9**
 - System Requirements 9
 - Additional Sun StorEdge Data Snapshot Software Documentation 10
 - Creating Primary Volumes 10
 - ▼ To Create a Volume 11
 - Installing Oracle 13
 - Oracle Installation Considerations 13
 - ▼ To Prepare for Oracle Installation 14
 - ▼ To Install Oracle Binaries 14

3. Using Sun StorEdge Data Snapshot Software With Oracle Software	15
Planning Snapshots	15
Working with Hot Backups and Sun StorEdge Data Snapshot Software	16
▼ To Perform a Hot Backup and Snapshot	16
▼ To Restore the Database	19
Working With Cold Backups and Sun StorEdge Data Snapshot	21
▼ To Perform a Cold Backup and Snapshot	21
▼ To Restore the Database to a Volume	23
Updating Snapshots	24
▼ To Resnap (Update) a Snapshot	25
Deleting a Snapshot	26
▼ To Delete a Snapshot	26
Glossary	27

Overview

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

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

- [“About the Sun StorEdge Data Snapshot Software” on page 1](#)
- [“Sun StorEdge Data Snapshot Features and Benefits” on page 2](#)
- [“Snapshot Reserve Space Requirements” on page 3](#)
- [“About Snapshots and Oracle Databases” on page 5](#)
- [“Example Uses of Sun StorEdge Data Snapshot Software” on page 6](#)

About the Sun StorEdge Data Snapshot Software

The Sun StorEdge Data Snapshot software is a point-in-time copy application that creates copies of storage volumes to be used for backup, application testing, and other business processing. The Sun StorEdge Data Snapshot 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 take snapshots of the database volumes. You can then make the snapshots available to a report-processing server, where they become available for developing reports. These volume snapshots reflect the data at the exact point in time at which they were taken (close-of-business,

for instance) and are available with little or no impact to ongoing database operations. You need to quiesce Oracle only long enough to issue a point-in-time copy command.

To ensure that Oracle databases and Sun StorEdge Data Snapshot 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 point-in-time copy software solutions with Oracle databases. These tests were performed according to the Oracle approved scenarios with the Sun StorEdge Data Snapshot point-in-time copy software.

How Volume Snapshots Work

Volume snapshots are point-in-time copies of the primary volume. A volume snapshot is available to servers on the storage area network (SAN) as a unique volume, but it always retains an association with the primary volume. Within the Sun StorEdge Data Snapshot software, copy-on-write (COW) technology is used to create logical copies of the primary volume without consuming equivalent disk space. Initially, when a snapshot is created, no data is copied for use by the snapshot. Instead, as blocks are updated on the primary volume, the old data blocks are moved to a location known as the *snapshot reserve space*. The Sun StorEdge Data Snapshot software then combines unchanged blocks from the primary volume with the old blocks stored in the snapshot reserve space to present a unique point-in-time copy of the primary volume.

Sun StorEdge Data Snapshot Features and Benefits

The Sun StorEdge Data Snapshot software has a number of features and benefits:

- **Support for up to eight snapshots** – You can create up to eight separate snapshots of a primary volume, which can then be accessed by multiple applications or maintained as a series of images that represent the data at different times.
- **Read/write support** – Snapshots can be mapped to other systems on the SAN through their unique addresses and world wide name (WWNs). Application servers can then use the snapshots as they would normal data volumes with read/write access. While the snapshot version of your data is being accessed by a secondary application, the original data continues to be modified by the first application on the primary volume.

- **Space-saving design** – Multiple versions of data can be created without multiplying the disk space consumed. Sun StorEdge Data Snapshot software allows a snapshot to initially consume minimal disk space. As existing data blocks on the primary volume are updated, the old blocks are moved to the snapshot reserve space, which means that Sun StorEdge Data Snapshot only requires space to hold the changed blocks. The snapshot data is consolidated so that one copy of old data can support multiple active snapshots of the same volume.
- **Rapid availability** – Snapshots become available within a few seconds after creation. This can save significant time when compared with the creation of a full copy of data through backup and restore, host-based copying utilities, or array-based mirroring operations.
- **One-step updates** – To facilitate processes that are periodic in nature (such as daily backups), Sun StorEdge Data Snapshot provides a resnap feature that allows you to update a snapshot with the latest data. The ID, WWN, and mapping of the original snapshot all remain unchanged while the snapshot data is updated from the primary volume.
- **Snapshot space threshold alerts** – As snapshot space fills up, alerts are issued through the disk array event system. These notifications provide you with timely information that enables you to take action.
- **Scripting interface** – You can create scripts to automate integration with your backup packages or applications. Scripting is conducted through a secure scripting interface installed on the application server.
- **Persistence** – Snapshots are often used for transient operations, but they are also persistent through power outages and maintenance shutdowns. Persistent and accessible snapshots alleviate the need for additional planning or recovery time and provide continuous availability to your end users.

Snapshot Reserve Space Requirements

Before you can create snapshots of a primary volume, you must set up the snapshot reserve space for that volume. The snapshot reserve space can be configured either during the creation of the primary volume or when the first snapshot is created. When you setup the snapshot reserve space, you designate how large it should be, Sun StorEdge Data Snapshot then completes the configuration and manages the snapshot reserve space. Each primary volume requires its own snapshot reserve space and all snapshots made from that primary volume share the same snapshot reserve space.

Once a snapshot has been created, as primary volume data blocks are updated for the first time, they are copied to the snapshot reserve space. Therefore the snapshot reserve space will begin to fill as the primary volume changes. When there are multiple snapshots, only one copy of each old data block needs to be maintained which results in space efficient operations.

As the snapshot reserve space is consumed, messages are generated at several thresholds to indicate the amount of space left. Snapshot reserve space can be reclaimed by deleting snapshots in the order oldest to newest.

In the event that the snapshot reserve space becomes full, the system will no longer be able to maintain old data blocks and the accuracy of the snapshot can no longer be guaranteed. If the snapshot reserve space becomes full, Sun StorEdge Data Snapshot takes the snapshots off-line. With the snapshot reserve space offline, the snapshots can be deleted, after which new snapshots can be taken.

Snapshot reserve space can be reconfigured to increase its size. However, a snapshot reserve space can not be reconfigured while snapshots are present. The snapshots must be deleted before the space can be modified. In order to reconfigure and expand the snapshot reserve space, there must be additional capacity available on the array volumes from which the primary volume is built.

Planning Snapshot Reserve Space

To plan your snapshot reserve space effectively, consider the following:

- **The length of time that snapshots will be kept in operation** – Once a snapshot is taken, subsequent changes to the primary volume's data are saved in the snapshot reserve space. Over time, as more blocks of the primary volume are updated, more data is moved to the snapshot reserve space. Snapshots that are kept longer tend to require more snapshot reserve space.
- **The number of snapshots to be taken** – All snapshots of a primary volume consume the same snapshot reserve space. The more snapshots you take, the more snapshot reserve space you will need. When setting up snapshot reserve space, you specify its size as a percentage of the primary volume. Given that eight snapshots are supported per volume, you can specify up to 800% of the master volume size.
- **The degree of write activity anticipated on the primary volume** – Write activity is the number of writes that occur on the primary volume and a measure of the net change in that primary volume over time. You can determine the snapshot reserve space you need as a function of the amount of anticipated change to the master volume (1% to 100%) between snapshots.
- **The degree of write activity anticipated on the snapshot** – When the Sun StorEdge Data Snapshot software creates a snapshot, it is creating a "volume" that can be mounted and then used by another application. You can work with a

snapshot as you would with any other volume, performing read and write transactions as necessary. If you are using a snapshot to analyze your database, it is possible that the data in the snapshot will be modified, independent of the data that is on the primary volume. Modifications made to the snapshot will be stored in the snapshot reserve space.

- **Disk space required for snapshot metadata** – Sun StorEdge Data Snapshot maintains metadata and pointers about each snapshot. This information is stored in the snapshot reserve space along with old data blocks that are moved during updates to the primary volume. When the snapshot reserve space is established, Sun StorEdge Data Snapshot will automatically configure an additional 7% space to the snapshot reserve space. While you do not need to account for the metadata in their calculations, you should be aware that this additional space will be allocated from the available disk space.

Example of Snapshot Reserve Space Configuration

As defined in the considerations above, two variables are required to actually configure the snapshot reserve space. The first is the snapshot count, which represents the estimated number of snapshots that will be used. The snapshot count does not limit how many snapshots can actually be taken, but helps you to estimate the size required for the snapshot reserve space. The second variable is the snapshot activity, which represents the estimated amount of data change between successive snapshots.

For example, you plan to create seven snapshots, one for each day of the week, and anticipate that the primary volume will change by 10% daily, between each snapshot. When configuring the snapshot reserve space use a snapshot count of 7 and an activity of 10%. StorEdge Data Snapshot then allocates enough snapshot reserve space to meet your specification. It reports the total amount of snapshot reserve space it allocated for you as a percentage of the total volume size. In this case it reports that it allocated a snapshot reserve space equal to 70% of the total size of the primary volume.

About Snapshots and Oracle Databases

The snapshot application works at a volume level independent of higher level applications. To maintain data integrity for other operations, you should ensure that the Oracle database is in a known state before creating the snapshot. To do this, place the Oracle database in “cold backup” or “hot backup” mode, ensuring that if

an Oracle instance is started on the snapshot, the database will be in a consistent state.

The Oracle database needs to be quiesced only long enough to allow creation of a snapshot. In order to 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 SAN through the unique WWN and ID that were assigned to it when it was created.

Example Uses of Sun StorEdge Data Snapshot Software

Sun StorEdge Data Snapshot software enables companies to improve their business operations by “cloning” their data sets, which enables them to run more applications in parallel and can thereby result in improved “return on information.” The following examples present possible methods of improving business operations.

- **Reducing backup times** – Backups are a critical part of any business continuity strategy. However, in order to provide uninterrupted data access to users, backup times are shrinking while the amount of information requiring protection is rapidly expanding. Backing up snapshots, rather than online data, enables critical transactions to keep running during the backup process.

Using a backup utility or scripts, and multiple snapshots, you can set up automatic backup procedures to ensure that backups of your database occur on a regular basis. Once a backup has been completed, you can restore backup copies if necessary.

- **Developing and testing new applications** – Many enterprise systems are being re-engineered to respond more efficiently to the networked business environment. Sun StorEdge Data Snapshot software can enhance these efforts by enabling business applications to be tested on real, up-to-date data. In this way, Sun StorEdge Data Snapshot software helps accelerate new application development and promote fast time-to-market.
- **Analyzing up-to-date information** – Businesses are constantly looking for ways to gain a competitive marketplace advantage. The faster they can convert data to information, the faster they can make critical decisions. With Sun StorEdge Data Snapshot software, you can create snapshots of a database that can then be used by a decision support application to analyze current information, leading to faster and better business decisions.

- **Restarting an application on a snapshot** – In the event of a user error and database corruption, application services can be restarted on snapshot volumes allowing a full and normal recovery to be delayed until a more convenient time. This can enable a fast return to operation. If this approach is used, you will then need to restore from a snapshot as discussed below.
- **Restoring snapshots to the primary volume** – Should the need arise, you can use snapshots to restore an Oracle database. This can be accomplished in several ways.
 - Copy the snapshots of the Oracle database volumes back to the primary volumes
 - Perform a block-to-block copy from the snapshot of the Oracle database volumes to secondary volumes, and then copy the data to the primary volume
 - Use a backup application to conduct a backup of the snapshots of the Oracle database volumes then conduct a restore to the primary volumes.

Requirements and Installation

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

- [“System Requirements” on page 9](#)
 - [“Additional Sun StorEdge Data Snapshot Software Documentation” on page 10](#)
 - [“Creating Primary Volumes” on page 10](#)
 - [“Installing Oracle” on page 13](#)
-

System Requirements

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

Hardware	Sun StorEdge 6920 system
Software	Solaris™ 8 (update 4 or higher) or Solaris 9 Operating System
For CLI interface	Sun StorEdge Remote Configuration CLI (sscs)
For GUI interface	One of the following: Netscape Navigator™ version 7.x or, Microsoft Internet Explorer version 5.x or, Mozilla™ 1.2.1

For further information, refer to the *Sun StorEdge 6920 System Host Installation Software Guide*.

Additional Sun StorEdge Data Snapshot Software Documentation

The following documents are supplied with the software.

Description	Title	Part Number
Installation information	<i>Sun StorEdge 6920 System Getting Started Guide</i>	817-5227-nn
	<i>Sun StorEdge 6920 System Host Installation Software Guide</i>	817-5831-nn
Reference information	<i>Sun StorEdge 6920 Technician's Reference Guide</i>	817-6499-nn
Release information	<i>Sun StorEdge 6920 System Release Notes</i>	817-5229-nn
man pages	sscs(1M)	N/A

Creating Primary 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 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 port that is identified by a port worldwide 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 Volume

These steps provide information on creating a primary volume that will store the Oracle database, as well as setting up snapshot reserve space.

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

```
# sscs create -p pool-name -s size TB|GB|MB|B [-S storage-domain-name]  
[-C snapshot-count 0|1|2|3|4|5|6|7|8] [-L snapshot-level  
low|verylittle|little|average|high|full] [-d volume-description]  
volume volume-name
```

You can use the following options and arguments with the `create` command.

TABLE 2-1 create Command Options and Arguments

Options and Arguments	Description
-p	This specifies name of the snapshot reserve space to be created.
-s	This is the size of the volume that will be created. The volume size is configured in terabytes, gigabytes, megabytes, and bytes.
[-S]	This variable specifies the storage domain name in which the volume is located. If nothing is specified, then the DEFAULT storage domain is used.
[-C]	When used in conjunction with the -L parameter, this variable represents the estimated number of snapshots that will be used, indicated by a number starting at 0 and continuing up to 8, for which space will be reserved.
[-L]	This variable is used to specify the amount of data overwrites between each snapshot. When used in conjunction with the -C parameter, this variable helps you to determine the appropriate amount of snapshot reserve space required for your system. <ul style="list-style-type: none">• low indicates that the volume will change by 10%• verylittle indicates that the volume will change by 25%• little indicates that the volume will change by 40%• average indicates that the volume will change by 50%• high indicates that the volume will change by 75%• full indicates that the volume will change by 100%
[-d]	This variable allows you to enter an up to 64 alphanumeric character description of the volume
[-t]	This variable directs the system to build the widest possible stripe in the selected pool. Use this variable in conjunction with a pool whose profile specifies a virtualization strategy of stripe to create a volume striped across all of the vdisks in the pool.

When creating the volume, you can create the snapshot reserve space that will be associated with this volume. Refer to “Snapshot Reserve Space Requirements” on page 3 for guidelines on setting snapshot reserve space.

For example:

```
# sscs create -p raid1_19pool -s 2 -S mr19 -C 8 -L full -d oraclevol
volume new200
```

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

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

For example, if you want to create a volume that is connected to the Sun StorEdge 6320 system by two host bus adapters (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 600075C00022690000000001000012F4  
# sscs map initiator 4800_sc3 -P readwrite -S mr19  
-l 200 -v 600075C00022690000000001000012F4
```

Installing Oracle

The following information, which is provided as an overview to install the Oracle binaries, uses Oracle 9i as an example. 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 create numerous Oracle databases on the server. The number of databases per server is restricted by resource availability. You can create databases using the `dbassist` utility.

▼ 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 9i Server Installation option and follow the instructions.**

Using Sun StorEdge Data Snapshot Software With Oracle Software

Using the Sun StorEdge Data Snapshot software, you can easily and quickly acquire a snapshot of the Oracle database at a specific point in time whenever you choose. This enables you to use a copy of the database, instead of your live database, to do testing, reporting and analysis.

This chapter describes the tasks needed to take snapshots of Oracle databases. It contains the following sections:

- “Planning Snapshots” on page 15
- “Working with Hot Backups and Sun StorEdge Data Snapshot Software” on page 16
- “Working With Cold Backups and Sun StorEdge Data Snapshot” on page 21
- “Updating Snapshots” on page 24
- “Deleting a Snapshot” on page 26

Planning Snapshots

Before you start creating snapshots, be aware of the following:

- The software can store up to eight snapshots.
- Snapshots are the fastest way to create a copy of the database. However, when using a snapshot volume, read/write operations occur on the primary volume. This might have an impact on performance at the primary volume.
- You can create snapshots using the Sun™ Web Console graphical user interface (GUI) or the command-line interface (CLI). Use the GUI or the CLI to manually create snapshots for one-time copies.

- To create automated scripts, you must use the CLI commands. To use the CLI, you must install the Sun StorEdge Remote Configuration CLI (*sscs*) software. This software provides the ability to perform remote configuration from a CLI.
- Also known as the *thin-scripting client software*, the software provides access to storage array configuration and management services through a client on an external host. Refer to the *Sun StorEdge 6920 System Host Installation Software Guide* for information on installation and configuration of the Remote Configuration CLI (*sscs*) software.

Working with Hot Backups and Sun StorEdge Data Snapshot Software

This section describes the general steps you must follow to perform a hot backup, create a snapshot of that backup, and then restore the database to a volume. Placing the Oracle software into “hot backup” mode ensures that the database will be in a consistent state while you take the snapshot. For further information regarding performing hot backups, refer to your Oracle documentation.

▼ To Perform a Hot Backup and Snapshot

1. Decide which volumes will require point in time snapshots.

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

2. Perform the following steps from the SQL*Plus prompt.

a. Place all 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. Unmount the volume.

4. Issue the command to take the snapshot.

- If the snapshot reserve space was created when you created the volume, enter the command:

```
# sscs snapshot -v primary-volume-name volume snapshot-name
```

- If the snapshot reserve space was not created when you created the volume, enter the command:

```
# sscs snapshot [-v source-volume-name] [-R resnap ]  
[-C snapshot-count 0|1|2|3|4|5|6|7|8] [-L snapshot-level  
low|verylittle|little|average|high|full] volume snapshot -volume-name
```

The following options and arguments apply to the snapshot command.

TABLE 3-1 snapshot command variables

Options and Arguments	Description
[-v]	This variable specifies name of the primary volume of which to take a snapshot.

TABLE 3-1 snapshot command variables (*Continued*)

Options and Arguments	Description
[-C]	When used in conjunction with the -L parameter, this variable is represents the estimated number of snapshots that will be used, indicated by a number starting at 0 and continuing up to 8, for which space will be reserved.
[-L]	This variable is used to specify the amount of data overwrites between each snapshot. When used in conjunction with the -C parameter, this variable helps you to determine the appropriate amount of snapshot reserve space required for your system. <ul style="list-style-type: none">• low indicates that the volume will change by 10%• verylittle indicates that the volume will change by 25%• little indicates that the volume will change by 40%• average indicates that the volume will change by 50%• high indicates that the volume will change by 75%• full indicates that the volume will change by 100%
volume	This specifies the snapshot volume name.

5. Map the snapshot to the appropriate initiator.

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

6. Remount the volume.

7. 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;
```

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

A message similar to the following is displayed once the snapshot has been created:

```
# /opt/se6x20/cli/bin/sscs list volume new200
Volume:                               new200
Description:
Creation Date:                         Tue May 11 20:05:38 PDT 2004
WWN:                                    600075C00022690000000000000012F4
Storage Domain:                        mire19
Pool:                                   raid1_19pool
Profile:                                new19_prof
Size:                                   2.0000 GB
State:                                  MAPPED
Condition:                              INTACT
Type:                                    Concat
Snapshot Reserve Size:                 100
Snapshot Reserve Percent Full:         5
Snapshots:
  Snapshot: vol200_snap1 Creation Time: Thu May 13 14:18:14 PDT 2004
Associations:
  Initiator: 4800_sc2 LUN: 200 Permission: Write/Read Map State:
Online
  Initiator: 4800_sc3 LUN: 200 Permission: Write/Read Map State:
Online
```

Once the snapshot has been created, it is available for use.

9. Using a backup program such as VERITAS NetBackup or Legato NetWorker, back up the snapshot to tape.

▼ To Restore the Database

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 will enable 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 the primary volume, or to a secondary volume.

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.

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

```
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 the original database, rename the data files in the new location (not the original database).

```
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.

```
SQL> alter database open resetlogs;
```

Working With Cold Backups and Sun StorEdge Data Snapshot

This section describes the general steps you must follow to perform a cold backup, create a snapshot of that backup, and then restore the database to a volume. Placing the Oracle software into cold backup mode ensures that the database will be in a consistent state while you take snapshots. For further information regarding performing cold backups, refer to your Oracle documentation.

▼ To Perform a Cold Backup and Snapshot

1. Decide which volumes will be copied with a point-in-time snapshot.

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

2. Shut down the database.

```
SQL> shutdown normal;
```

3. Unmount the volume.

4. Issue the command to take the snapshot.

- If the snapshot reserve space was created when you created the volume, type the following command:

```
# sscs snapshot -v primary-volume-name volume snapshot-name
```

- If the snapshot reserve space was not created when you created the volume, enter the command:

```
# sscs snapshot [-v source-volume-name] [-R resnap ]  
[-C snapshot-count 0|1|2|3|4|5|6|7|8] [-L snapshot-level  
low|verylittle|little|average|high|full] volume snapshot -volume-name
```

The following options and arguments apply to the snapshot command.

TABLE 3-2 snapshot command variables

Options and Arguments	Description
[-v]	This variable specifies name of the primary volume of which to take a snapshot.
[-C]	When used in conjunction with the -L parameter, this variable represents the estimated number of snapshots that will be used, indicated by a number starting at 0 and continuing up to 8, for which space will be reserved.
[-L]	This variable is used to specify the amount of data overwrites between each snapshot. When used in conjunction with the -C parameter, this variable helps you to determine the appropriate amount of snapshot reserve space required for your system. <ul style="list-style-type: none">• <code>low</code> indicates that the volume will change by 10%• <code>verylittle</code> indicates that the volume will change by 25%• <code>little</code> indicates that the volume will change by 40%• <code>average</code> indicates that the volume will change by 50%• <code>high</code> indicates that the volume will change by 75%• <code>full</code> indicates that the volume will change by 100%
volume	This specifies the snapshot volume name.

5. Map the snapshot to the appropriate initiator.

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

6. Mount the volume.

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

A message similar to the following is displayed after the snapshot has been created.

```
# /opt/se6x20/cli/bin/sscs list volume new200
Volume:                               new200
Description:
Creation Date:                         Tue May 11 20:05:38 PDT 2004
WWN:                                    600015D00022690000000000000012F4
Storage Domain:                        mire19
Pool:                                   raid1_19pool
Profile:                                new19_prof
Size:                                   2.0000 GB
State:                                  MAPPED
Condition:                              INTACT
Type:                                    Concat
Snapshot Reserve Size:                  100
Snapshot Reserve Percent Full:          5
Snapshots:
  Snapshot: vol200_snap1 Creation Time: Thu May 13 14:18:14 PDT 2004
Associations:
  Initiator: 4800_sc2 LUN: 200 Permission: Write/Read Map State:
Online
  Initiator: 4800_sc3 LUN: 200 Permission: Write/Read Map State:
Online
```

After the copy has been completed, the snapshot reserve space is available immediately as a snapshot copy for backup.

8. Using a backup program such as VERITAS NetBackup or Legato NetWorker, back up the snapshot to tape.

▼ To Restore the Database to a Volume

1. Using a backup program such as VERITAS NetBackup or Legato NetWorker, restore the snapshot from tape to the primary volume.

This will enable 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 the primary volume, or to a secondary volume.

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

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

4. Restore the database.

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

```
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 the original database, rename the data files in the new location (not the original database).

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

5. Recover the database.

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

6. When the recovery is complete, open the database.

```
SQL> alter database open resetlogs;
```

Updating Snapshots

You can update an existing snapshot using the `resnap` command. Updating a snapshot enables you to use the same snapshot volume information, while updating the point-in-time copy.

▼ To Resnap (Update) a Snapshot

1. Type the snapshot command on the command line. .

```
# sscs snapshot -v volume -R volume snapshot -volume-name
```

In the following example, the volume `new200` has an existing snapshot. To see the volume information, run the `sscs list volume` command.

```
# /opt/se6x20/cli/bin/sscs list volume new200
Volume:                               new200
Description:
Creation Date:                         Tue May 11 20:05:38 PDT 2004
WWN:                                    600015D00022670000000000000012F4
Storage Domain:                        mire19
Pool:                                    raid1_19pool
Profile:                                 new19_prof
Size:                                    2.0000 GB
State:                                   MAPPED
Condition:                              INTACT
Type:                                    Concat
Snapshot Reserve Size:                 100
Snapshot Reserve Percent Full:         45
Snapshots:
Snapshot: vol200_snap1 Creation Time: Thu May 13 14:18:14 PDT 2004
Associations:
  Initiator: 4800_sc2 LUN: 200 Permission: Write/Read Map State:
Online
  Initiator: 4800_sc3 LUN: 200 Permission: Write/Read Map State:
Online
```

2. After identifying the snapshot volume name, unmount the snapshot volume named `vol200_snap1`.
3. Run the `resnap` command, which resnaps a snapshot of `vol1200_snap1` on the volume `new200`.

```
# sscs snapshot -v new200 -R volume vol1200_snap1
```

4. Verify that the snapshot was resnapped by issuing the `sscs list volume` command.

```
# /opt/se6x20/cli/bin/sscs list volume new200
Volume:                               new200
  Description:
  Creation Date:                       Tue May 11 20:05:38 PDT 2004
  WWN:                                  600015D00022670000000000000012F4
  Storage Domain:                      mire19
  Pool:                                 raid1_19pool
  Profile:                              new19_prof
  Size:                                 2.0000 GB
  State:                                MAPPED
  Condition:                            INTACT
  Type:                                 Concat
  Snapshot Reserve Size:                100
  Snapshot Reserve Percent Full:        45
  Snapshots:
Snapshot: vol1200_snap1 Creation Time: Mon May 17 14:14:53 PDT 2004
  Associations:
    Initiator: 4800_sc2 LUN: 200 Permission: Write/Read Map State:
Online
    Initiator: 4800_sc3 LUN: 200 Permission: Write/Read Map State:
Online
```

The snapshot that has been resnapped, or updated, is now listed with the new creation time. Once you have verified that the resnap has occurred, mount `vol1200_snap1`.

Deleting a Snapshot

You can delete a snapshot to create space for new snapshots.

▼ To Delete a Snapshot

1. Unmount the snapshot.
2. Type the `snapshot` command on From the command-line,.

```
# sscs delete -s domain volume snapshot-volume-name
```

Glossary

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.
copy-on-write (COW)	The process used to preserve point-in-time data when new data is written to the disk. With each write, the system preserves the “old” data in the snapshot reserve space, so that it can re-create the volume as it existed at the time of the snapshot.
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	An initiator is the protocol-specific physical port that initiates the I/O exchanges with the system.
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.
point-in-time copy	A copy or image of volume data captured or copied at a particular point in time by the data snapshot software. Also called a <i>snapshot</i> .
primary volume	The volume containing the original data that is copied into the snapshot.
resnap	The act of updating, or “resnapping,” an existing snapshot.
snap	The act of creating a point-in-time data image of a volume using the data snapshot software.
snapshot	A copy of a primary volume that is taken to reflect an exact point in time.
snapshot reserve	When a snapshot is created, blocks are updated on the primary volume, and the old data blocks are moved to a location known as the <i>snapshot reserve space</i> .
thin-scripting client	A command-line interface that provides the ability to perform remote configuration.

