Sun StorEdge™ Data Snapshot Software With Oracle Databases Usage Guide

For Sun StorEdge 6920 Systems

Sun Microsystems, Inc.
www.sun.com

Part No. 819-3326-10 (v1)
July 2005, Revision A

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

Contents

CHAPTER 1

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

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 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.

About Snapshots

A snapshot is a point-in-time copy of the source (parent) volume. A snapshot is available to servers on the storage area network (SAN) as a unique volume, but it always retains an association with the parent volume. Within the Sun StorEdge Data Snapshot software, copy-on-write (COW) technology is used to create logical copies of the parent 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 parent volume (B^1), the old data blocks are moved to a location known as the *snapshot reserve space*, as shown in FIGURE 1-1. The Sun StorEdge Data Snapshot software then combines unchanged blocks from the parent volume with the old blocks stored in the snapshot reserve space to present a unique point-in-time copy of the parent volume.

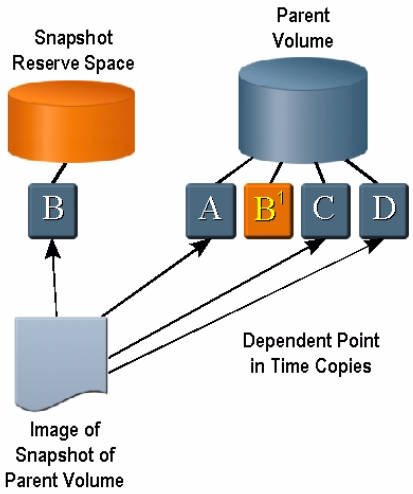

FIGURE 1-1 Old Blocks Are Moved to the Snapshot Reserve Space

When a host sends a request to the snapshot, the system determines whether the requested blocks have changed on the parent volume since the time the snapshot was created. One of the following then occurs:

- If the blocks have changed, the read request is satisfied from the data stored in the snapshot reserve space.
- If the blocks have not changed, the read request is satisfied from the parent volume.

When a write operation occurs for a data block of the parent volume in which the data has not changed since the snapshot's creation, the system does the following:

- Copies the (original) data about to be changed to the snapshot reserve space. Snapshots can accept write operations, and this data is also dependent on the snapshot reserve space.
- Writes the new data to the parent volume. Subsequent rewrites of that same data block on the parent volume do not recopy to the snapshot. In this manner, the snapshot records the data on the parent volume at the time the snapshot is taken.
- Adds a record to the snapshot that indicates the new location for the old data.

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 parent 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 Names (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 parent volume.
- **Space-saving design** – Multiple versions of data can be created without multiplying the disk space consumed. Sun StorEdge Data Snapshot software's COW technology allows a snapshot to initially consume minimal disk space. As existing data blocks on the parent volume are updated, the old blocks are moved to the snapshot reserve space, which means that the system requires only 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 they are created. This can save significant time when compared with creating a full copy of data through backup and restore, host-based copying utilities, or array-based mirroring operations.
- **One-step updates** – To support processes that are periodic in nature (such as daily backups), Sun StorEdge Data Snapshot software provides a resnap feature 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 parent volume.
- **Space consumption threshold alerts** – As snapshot space fills up, alerts at several thresholds are issued through the disk array event system. These notifications provide warning when the snapshot reserve space is 70, 90, 95, and 100 percent consumed. This enables you to take action to expand the snapshot reserve space before it is filled. In addition, you can view the snapshot reserve space in use at any time through the Sun StorEdge Configuration Service interface.
- **Selectable location for the snapshot reserve space** – During configuration, you can determine where you want to locate the snapshot reserve space. By locating it on a separate disk, you can minimize the effects of copy-on-write operations. Additionally, the snapshot reserve space can be allocated from a different storage pool than the parent volume.
- **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 rollback** – Through snapshot rollback, a parent volume can be reverted to a selected snapshot. In the event of data corruption, rollback enables quick application restart from a known snapshot without your having to wait for a tape restore.
- **Expandable snapshot reserve space** – The snapshot reserve space can be expanded as needed and located anywhere on the Sun StorEdge 6920 system. This enables you to react to unplanned changes in usage or growth for snapshots. For example, if you decide to increase snapshot usage or to store snaphots for long periods, you can easily add space on a just-in-time basis, without disrupting system activity.
- **Third-party and legacy Sun storage support** – Snapshot software capabilities can be added to supported existing Sun storage and storage from other vendors. This feature enhances existing storage and enables a single snapshot methodology for different storage array types.

About Snapshot Reserve Space

Before you can create snapshots of a parent volume, you must set up the snapshot reserve space for that volume. The snapshot reserve space can be configured either during the process of creating the parent 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 parent volume requires its own snapshot reserve space and all snapshots made from that parent volume share the same snapshot reserve space.

Once a snapshot is created, as parent volume data blocks are updated for the first time, they are copied to the snapshot reserve space. Therefore the snapshot reserve space begins to fill as the parent 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 remaining. You can also view the amount of snapshot reserve space that is available. Snapshot reserve space can be reclaimed by deleting snapshots in the order of oldest to newest.

Allocating snapshot reserve space when you create the volume helps ensure that there is be sufficient free space for the snapshots. When you allocate snapshot reserve space later, the required amount of free space may not be available. In this case, you need to delete snapshots, use the resnap operation to retake snapshots, or expand the snapshot reserve space.

Note – Once you allocate the size of the snapshot reserve space, you cannot decrease it without deleting and then re-creating it.

You must consider two important factors in determining the size of snapshot reserve space:

- The predicted quantity of snapshots.
- The frequency of write operations expected for the parent volume. Only data that is overwritten in the parent volume is copied into the snapshot reserve space.

For example, assume that a 12-Tbyte production database generates 100 Gbytes of changes on a daily basis. If the requirement is to keep daily copies online for a week (seven snapshot copies of the data), the set of copies requires approximately 700 Gbytes of snapshot reserve space.

Planning Snapshot Reserve Space

To plan snapshot reserve space, you need to know:

- **Length of snapshot activity** – Once a snapshot is taken, subsequent changes to the parent volume data are saved in the snapshot reserve space. Over time, as more blocks of the parent volume are updated, more data is copied to the snapshot reserve space.
- **Number of snapshots to be taken** – All snapshots of a parent volume share space from the same snapshot reserve space. The more snapshots are taken, the more snapshot reserve space may be required. When you configure snapshot reserve space, you specify the size as a percentage of the parent volume. Given that eight snapshots are supported per volume, you can specify up to 800 percent of the parent volume size.
- **Degree of write activity anticipated on the parent volume** – Write activity is the number of storage blocks that may change on the parent volume after each snapshot is taken and is a measure of the net change in the parent 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 percent) between snapshots.
- **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 can be modified, independent of the data that is on the parent volume. Modifications made to the snapshot are stored in the snapshot reserve space.
- **Disk space required for snapshot metadata** – Sun StorEdge Data Snapshot software 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 parent volume. When the snapshot reserve space is established, Sun StorEdge Data Snapshot software automatically configures an additional seven percent of space to the snapshot reserve space. Although you do not need to account for the metadata in you calculations, you should be aware that this additional space is allocated from the available disk space.
- **Snapshot reserve space policy –** In the unlikely event that the snapshot reserve space becomes exhausted or becomes unavailable because of a fault, there is an adjustable policy governing system behavior. The use of snapshots to perform backups is intended to maintain a high level of availability for the original volume. Therefore, if you are using snapshots to perform backups, use the favor original policy to allow for failure of the snapshots in order to have uninterrupted access to the original volume. In contrast, if you are using snapshots as fail-safe copies of the original volume or if user data is written to the snapshots, select the favor snapshot policy to preserve the snapshots.

- **Favor parent volume over snapshots (default)** – In this policy the parent volume remains online while snapshots are taken offline and invalidated. Since copy-on-write operations cannot be completed, the integrity of the snapshots can no longer be assured, and the snapshots must therefore be declared invalid. In this case, the invalid snapshots need to be manually deleted before snapshot operations are resumed for the associated parent volume.
- **Favor snapshot** – This option stops all writes to the parent volume when copy-on-write activity can no longer be completed for snapshots. Since both snapshots and the parent volume are static, the snapshots remain valid. However, the snapshots and the parent volume remain offline until the condition can be addressed.

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 a snapshot. To do this, place the Oracle database in cold backup or hot backup mode. This helps 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 enable a snapshot to be created. In preparation for the snapshot, the system briefly suspends I/O to the parent volume, establishes the necessary pointers for the snapshot, and then resumes I/O to the parent 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 Oracle operations by using snapshots.

- **Reducing database backup times** – Most data centers need to provide uninterrupted data access to users, which means there is little time to take the database offline for backups. At the same time, the amount of data that needs to

be backed up is increasing and requires more time to backup. Backing up snapshots of the database, 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 help ensure that backups of your database occur on a regular basis. Once a backup is completed, you can restore backup copies if necessary.

- **Developing and testing new applications** – Snapshots can help you develop and test new database application by enabling new applications to be tested on real, up-to-date data.
- **Analyzing up-to-date information** – 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 subsequent database corruption, application services can be restarted on snapshots, allowing a full and normal recovery to be delayed until a more convenient time. If this approach is used, you can later restore the database from these snapshots.
- **Restoring snapshots to the parent volume** – Should the need arise, you can use snapshots to restore an Oracle database. You can do this in any of the following ways:
  - By copying the snapshots of the Oracle database volumes back to the parent volumes
  - By performing a block-to-block copy from the snapshot of the Oracle database volumes to secondary volumes, and then copying the data to the parent volume
  - By using a backup application to perform a backup of the snapshots of the Oracle database volumes and then performing a restore to the parent volumes
- **Reducing the time-period for data loss** – Sun StorEdge Data Snapshot software can be used to reduce the window for data loss in the event of a database or file system corruption. For example, whereas a typical backup policy may provide for a nightly backup, creating two snapshots a day can reduce the potential window for data loss from 24 hours to 12 hours. The lighter impact of snapshots can enable more frequent copies and result in a better level of data protection and availability.

CHAPTER 2

Requirements and Installation

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

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	Remote scripting command-line interface (CLI)
For browser interface	One of the following: • 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 about the Sun StorEdge 6920 system, see the following documents, which are supplied with the software.

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

Creating 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 store 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 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 Volume

Follow these steps to create a parent volume that can store the Oracle database, as well as to set up snapshot reserve space.

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

```
# sscs create -p pool-name -s size TB|GB|MB|B
[-C snapshot-count 0|1|2|3|4|5|6|7|8] [-L snapshot-level
low|verylittle|little|average|high|full] [-P pool-name]
[-f volume|snapshot] [-t] [-d volume-description]
[-S storage-domain-name] 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`	The storage pool from which the volume is to be created.
`-s`	The size of the volume that is to be created. The volume size is configured in terabytes, gigabytes, megabytes, or bytes.
`[-C]`	The snapshot count for the snapshot reserve. This value is used in conjunction with the `-L` option to calculate the size of the snapshot reserve.
`[-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: • `low` – 10 percent • `verylittle` – 25 percent • `little` – 40 percent • `average` – 50 percent • `high` – 75 percent • `full` – 100 percent
`[-P]`	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 `-S` option.
`[-f]`	The policy setting specifying which should be favored, the snapshot or the volume.

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

Options and Arguments	Description
`[-t]`	A directive to 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 virtual disks in the pool.
`[-d]`	An alphanumeric character description of the volume of up to 64 characters.
`[-S]`	The name of the storage domain name 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 volume.

Refer to "About Snapshot Reserve Space" on page 5 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
```

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 snapshot-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 600075C000226900000000001000012F4
# sscs map initiator 4800_sc3 -P readwrite -S mr19
-l 200 -v 600075C000226900000000001000012F4
```

Installing Oracle

The following information is provided as an overview of installing the Oracle binaries, using 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 the 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.**

CHAPTER 3

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. This enables you to use a copy of the database, instead of your live database, to do testing, reporting, or analysis.

This chapter describes the tasks involved in taking snapshots of Oracle databases. It contains the following sections:

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 an application is using a snapshot volume, read / write operations occur on the parent volume. This might have an impact on performance at the parent volume.
- You can create snapshots using the Java™ Web Console browser page or the command-line interface (CLI). Use the browser 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 remote scripting CLI software. This software also enables you to perform remote configuration from a CLI, and it gives you access to storage array configuration and management services through a client on an external host. Refer to the *Sun StorEdge 6920 System Getting Started Guide* for information on installation and configuration of the remote scripting CLI.

Working With Backups and Sun StorEdge Data Snapshot Software

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 or use the snapshot for other purposes. You can use hot back or cold backup mode:

- A hot backup is created while a database is online and in use. Placing the Oracle software into hot backup mode helps ensure that the database is in a consistent state while you take the snapshot.
- Placing the Oracle software in cold backup mode takes the database offline and helps ensure that it is in a consistent state while you take snapshots.

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

▼ To Perform a Backup and Create a Snapshot

1. **Decide which volumes are going to require point-in-time 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 volume, use this command:

```
# sscs snapshot -v source-volume-name volume new-snapshot-name
```

For example, the following commands create snapshots of the data files, redo logs, and archived logs:

```
# sscs snapshot -v
/dev/dsk/c5t600015D000226900000000000000700d0s6/dbs
volume /dev/dsk/c5t600015D000226900000000010000700d0s7/dbs_snap
# sscs snapshot -v
/dev/dsk/c5t600015D000226900000000000000704d0s6/redo
volume /dev/dsk/c5t600015D000226900000000010000704d0s7/redo_snap
# sscs snapshot -v
/dev/dsk/c5t600015D000226900000000000000708d0s6/arch
volume /dev/dsk/c5t600015D000226900000000010000708d0s7/arch_snap
```

- If you did not create the snapshot reserve space when you created the volume, use this command:

```
# sscs snapshot -v source-volume-name
[-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
```

TABLE 3-1 describes the options and arguments that apply to the `snapshot` command.

TABLE 3-1 `snapshot` Command Options and Arguments

Options and Arguments	Description
`-v`	The volume for which you are taking a snapshot.
`[-C]`	The snapshot count for the snapshot reserve. This value is used in conjunction with the `-L` option to calculate the size of the snapshot reserve.
`[-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 equate to the following percentages: • `low` – 10 percent • `verylittle` – 25 percent • `little` – 40 percent • `average` – 50 percent • `high` – 75 percent • `full` – 100 percent
`volume`	The name of the new snapshot volume.

4. **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
```

The initiator could be a backup host or another host that can use the snapshot to start a new instance of the database. The new instance can then be employed to perform analysis or tests.

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.**

A message similar to the following is displayed after the snapshot is created:

```
# /opt/se6x20/cli/bin/sscs list -S volume new200
Volume:                          new200
  Description:
  Creation Date:                 Tue May 10 20:05:38 PDT 2005
  WWN:                           600075C0002269000000000000012F4
  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 snapshot is created, it is available for use.

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

▼ 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 the parent volume or to a different 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 in the original database, rename the data files in the new location (not the original database).

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

Other Uses for Snapshots

In addition to being used for database backups, the snapshots can be used to perform data analysis or to test applications with actual, current data. Use the most recent snapshots for data analysis or testing, instead of working with the online database, which might interfere with critical transactions.

You can also use the backup of a snapshot to populate a new volume with data from a particular point in time. You can then take a snapshot of the new volume, perform tests, then roll back the snapshot to the original data and perform the tests again.

Furthermore, you can use snapshots to restart the database if an application problem causes questionable data to be written to the parent volume. You can restart the database with the last known good snapshot until it is convenient to perform a full recovery.

▼ To Roll Back Snapshots

- **Issue the command to roll back a snapshot.**

```
# sscs snapshot -b -S storage-domain-name volume snapshot-volume-
name
```

Updating Snapshots

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

▼ To Resnap (Update) a Snapshot

- 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.

1. **To identify the snapshot name and see other information, run the** `sscs list volume` **command.**

```
# /opt/se6x20/cli/bin/sscs list volume new200
Volume:                           new200
  Description:
  Creation Date:                  Tue May 10 20:05:38 PDT 2005
  WWN:                            600015D000226700000000000012F4
  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 12 14:18:14 PDT 2005
Associations:
    Initiator:  4800_sc2  LUN:  200  Permission:  Write/Read  Map State:
Online
    Initiator:  4800_sc3  LUN:  200  Permission:  Write/Read  Map State:
Online
```

2. **Unmount the snapshot volume named** `vol200_snap1`**.**

3. **Run the** `resnap` **command, which resnaps** `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 10 20:05:38 PDT 2005
  WWN:                            600015D000226700000000000000012F4
  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: Mon May 16 14:14:53 PDT 2005
  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`.

Expanding Snapshot Reserve Space

You can expand snapshot reserve space in order to make room for more snapshots.

▼ To Expand Snapshot Reserve Space

1. **From the Java Web Console page, click Sun StorEdge 6920 Configuration Service > Logical Storage > Snapshots.**

   The Snapshot Summary page is displayed.

2. **Select the parent volume for which you want to expand snapshot reserve space.**

   The Volume Details page is displayed.

3. **Click Snapshot Reserve.**
4. **Increase the values of the settings, as appropriate, for the following fields**:
   - Space Allocated For
   - Expected Write Activity

   For more information on these fields, click the Help button.
5. **Click Save.**

▼ To Expand the Snapshot Reserve Space and Take a Snapshot Using the Command Line

- **Type the** `snapshot` **command on the command line.**

```
# sscs snapshot -v source-volume-name -C snapshot-count
-L snapshot-level -S storage-domain-name volume snapshot-volume-
name
```

TABLE 3-2 describes the options and arguments that apply to the `snapshot` command.

TABLE 3-2 `snapshot` Command Options and Arguments

Options and Arguments	Description
`-v`	The volume for which you are taking a snapshot.
`[-C]`	The snapshot count for the snapshot reserve. This value is used in conjunction with the `-L` option to calculate the size of the snapshot reserve.

TABLE 3-2 `snapshot` Command Options and Arguments *(Continued)*

Options and Arguments	Description
`[-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 equate to the following percentages: • `low` – 10 percent • `verylittle` – 25 percent • `little` – 40 percent • `average` – 50 percent • `high` – 75 percent • `full` – 100 percent
`[-S]`	The name of the storage domain name in which the volume is located. If nothing is specified, then the DEFAULT storage domain is used.
`volume`	The name of the new snapshot volume.

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** `delete` **command on 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) 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.

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.

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.

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.

resnap To create a snapshot again and replace the old with the new.

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.

thin-scripting client See remote scripting CLI client.