Sun StorEdge™ Network Data Replicator 3.0 System Administrator's Guide

Sun Microsystems, Inc.
901 San Antonio Road
Palo Alto, CA 94303-4900 U.S.A.
650-960-1300

Part No. 806-7512-10
June 2001, Revision A

Send comments about this document to: docfeedback@sun.com

Please
Recycle

Adobe PostScript

Contents

Figures

Code Samples

Tables

Preface

The *Sun StorEdge Network Data Replicator 3.0 System Administrator's Guide* describes the Sun StorEdge™ Network Data Replicator (Sun SNDR) software and its use. This guide is intended for system administrators who have experience with the Solaris™ operating environment and related disk storage systems.

Before You Read This Book

To fully use the information in this document, you must have thorough knowledge of the topics discussed in these books:

- *Sun StorEdge Network Data Replicator 3.0 Installation Guide*
- *Sun StorEdge Network Data Replicator 3.0 Release Notes*

The Sun SNDR software must be installed as described in the installation and release documentation.

How This Book Is Organized

Chapter 1 provides general information about the Sun SNDR software and describes the Sun SNDR software terminology.

Chapter 2 describes the Sun SNDR software configuration and bitmapvolumes.

Chapter 3 describes the Sun SNDR software commands.

Chapter 4 explains the Sun SNDR software operations including synchronization, recovery rehearsals, and mirrored partitions.

The Glossary contains definitions of terms used in this document.

Using UNIX Commands

This document may not contain information on basic UNIX® commands and procedures such as shutting down the system, booting the system, and configuring devices.

Refer to the software documentation that you received with your system.

Documentation Conventions

Typeface or Symbol	Meaning	Examples
`AaBbCc123`	The names of commands, files, and directories; on-screen computer output.	Edit your `.login` file. Use `ls -a` to list all files. `% You have mail.`
`AaBbCc123`	What you type, when contrasted with on-screen computer output.	`% su` `Password:`
AaBbCc123	Book titles, new words or terms, words to be emphasized. Command-line variable; replace with a real name or value.	Read Chapter 6 in the *User's Guide*. These are called *class* options. You *must* be `root` to do this. To delete a file, type `rm` *filename*.
[ ]	In syntax, brackets indicate that an argument is optional.	`scmadm [-d` *sec*`] [-r` *n*[:*n*][,*n*]...`] [-z]`
{ *arg* | *arg*}	In syntax, braces and pipes indicate that one of the arguments must be specified.	`sndradm -R b {p|s}`

Shell Prompts

Shell	Prompt
C shell	*machine_name*%
C shell superuser	*machine_name*#
Bourne shell and Korn shell	$
Bourne shell and Korn shell superuser	#

Related Documentation

Application	Title	Part Number

man pages	`sndradm`
	`scmadm`
	`svadm`
	`iiadm`
		N/A

Latest information	*Sun StorEdge Network Data Replicator 3.0 Release Notes*	806-7513
	Sun StorEdge Instant Image 3.0 Release Notes	806-7678
Installation and User	*Sun StorEdge Network Data Replicator 3.0 Installation Guide*	806-7514
	Sun StorEdge Instant Image 3.0 Installation Guide	806-7675
	SunATM 3.0 Installation and User's Guide *Sun ATM 4.0 Installation and User's Guide*	805-0331 805-6552
	Sun Gigabit Ethernet FC-AL/P Combination Adapter Installation Guide	806-2385
	Sun Gigabit Ethernet/S 2.0 Adapter Installation and User's Guide *Sun Gigabit Ethernet/P 2.0 Adapter Installation and User's Guide*	805-2784 805-2785
	Sun Enterprise 10000 InterDomain Networks User Guide	806-4131
System administration	*Sun StorEdge Instant Image 3.0 System Administrator's Guide*	806-7677
	TCP/IP and Data Communications Administration Guide	805-4003
Configuration	*Sun StorEdge Network Data Replicator 3.0 Configuration Guide*	806-7550
	Sun Enterprise 10000 InterDomain Network Configuration Guide	806-5230

Accessing Sun Documentation Online

A broad selection of Sun system documentation is located at:

`http://www.sun.com/products-n-solutions/hardware/docs`

A complete set of Solaris documentation and many other titles are located at:

`http://docs.sun.com`

For the latest version of released storage software documentation, go to:
`http://www.sun.com/products-n-solutions/hardware/docs/Software/`

Ordering Sun Documentation

Fatbrain.com, an Internet professional bookstore, stocks select product documentation from Sun Microsystems, Inc.

For a list of documents and how to order them, visit the Sun Documentation Center on Fatbrain.com at:

`http://www.fatbrain.com/documentation/sun`

Sun Welcomes Your Comments

Sun is interested in improving its documentation and welcomes your comments and suggestions. You can email your comments to Sun at:

`docfeedback@sun.com`

Please include the Part No. 806-7512-10 in the subject line of your email.

CHAPTER 1

Overview

Caution – Do not install or use the Sun StorEdge Version 3.0 Core and data services software on servers in a Sun Cluster 3.0 environment. **The Version 3.0 software is not co-existent with the Sun Cluster 3.0 environment.** The Version 3.0 software **is** co-existent in the Sun Cluster 2.2 environment, where it will not interfere with failover. The Version 3.0 core and data services software is cluster aware in the Sun Cluster 3.0 Update 1 environment and provides high availability for the Sun StorEdge software.

This chapter describes the Sun StorEdge Network Data Replicator (Sun SNDR) software and its capabilities. The topics in this chapter include:

Sun SNDR Software Description

The Sun SNDR software is a remote replication facility for the Solaris™ operating environment. It is intended for use as part of a disaster recovery and business continuance plan to provide redundant storage of critical information across physically separate sites.

The Sun SNDR software enables you to replicate disk volumes between physically separate primary and secondary hosts in real time. To transport data, the Sun SNDR software uses any Sun network adapter that supports TCP/IP.

A Sun SNDR software volume set consists of a primary volume residing on a local host and secondary volume residing on a remote host. The volume set also includes a bitmap volume on each host to track write operations and differences between the volumes.

The Sun SNDR software enables you to group volume sets. You can assign specific volume sets to a group to perform replication on these volume sets and not on others you have configured. Grouping volume sets also guarantees write ordering: write operations to the secondary volume occur in the same order as the write operations to the primary volume.

You can also group volume sets according to their cluster or resource tag to perform replication in a clustered environment. The Sun SNDR Version 3.0 software is:

- **Co-existent** in the Sun Cluster 2.2 environment. It will not interfere with failover
- **Not co-existent** with the Sun Cluster 3.0 environment
- **Cluster aware** in the Sun Cluster 3.0 Update 1 environment and provides high availability for the Sun StorEdge software

See “Using the Sun StorEdge Data Services Software in a Sun Cluster Environment” on page 7.

You can update the data on the secondary volume by issuing a command to *resynchronize* the volumes. You can also restore data from the secondary volume to the primary volume by issuing a command to *reverse resynchronize* the volumes.

The secondary volumes can be updated synchronously in real time or asynchronously using a store-and-forward technique. Typically, a primary volume is first explicitly copied to a designated secondary volume to establish matching contents. As applications write to the primary volume, the Sun SNDR software replicates changes to the secondary volume, keeping the two images consistent.

In synchronous mode, a write operation is not confirmed as complete until the remote volume has been updated. In asynchronous mode, a write operation is confirmed as complete before the remote volume has been updated.

Features

TABLE 1-1 describes the Sun SNDR software features.

TABLE 1-1 Sun SNDR Software Features

This Feature. . .	Enables You To. . .
Active logging	Continue scoreboard logging operations whenever the Sun SNDR software is disabled or interrupted.
Backward compatible `rdcadm` command	Maintain the integrity of scripts you have written and want to continue using. The Version 3.0 `sndradm` command is linked to the `rdcadm` command and is generally compatible with Sun SNDR software Version 2.0 `rdcadm` command. See TABLE 3-1 on page 21.
Disaster rehearsal rollback	Simulate a disaster without committing data changes to volumes.
Granularly configurable	Configure the Sun SNDR software to operate on a logical volume; you can exclude volumes containing noncritical data from the Sun SNDR software operations.
Groups of Sun SNDR volume sets	Control the Sun SNDR software volume sets simultaneously by grouping them. This feature is essential in installations requiring you to maintain consistent contents of a group of volumes. A group is a collection of Sun SNDR software sets that have the same group name, primary and secondary interfaces, and mirroring mode. Mixed groups (those where mirroring modes are asynchronous for one set and synchronous for another set) are not allowed. The Sun SNDR software maintains write ordering for volumes in a group to ensure that the data on the secondary volumes is a consistent copy of the corresponding primary volumes.
Multihop sets	Replicate data from one primary volume to a secondary volume; the secondary volume then replicates the data again to another secondary volume, and so on, in a daisy-chain fashion. See “One-to-Many and Multihop Volume Sets” on page 6.
Mutual backup	Concurrently transmit and receive remote dual copies to and from its remote counterpart. Also known as a *bilateral relationship*.
One-to-many sets	Replicate data from one primary volume to many secondary volumes residing on one or more hosts. When you perform a forward resynchronization, you can synchronize one volume set or all volume sets. Issue a separate command for each set. You can also update the primary volume using a specific secondary volume. See “One-to-Many and Multihop Volume Sets” on page 6.

TABLE 1-1 Sun SNDR Software Features *(Continued)*

This Feature. . .	Enables You To. . .
Optimized resynchronization	Resynchronize volumes following disk, link, system, and storage platform outages; you only replicate those blocks that were modified
RAID support	Use RAID volumes as part of your Sun SNDR software strategy. Volumes can be any RAID level.
Security and Internet Protocols, Version 3.0	Use Version 3.0 securely and efficiently in the Sun Solaris 7 and 8 operating environments, which support Internet Security Protocol (IPsec). The Solaris 8 operating environment also supports Internet Protocol Version 6 (IPv6). (The Solaris 7 operating environment does not support IPv6). Version 3.0 does not require the use of an `.rhosts` file. You place the hosts to be used in the `/etc/hosts` file of each host running the Sun SNDR Version 3.0 software in your configuration.
Stop and start control of remote replication on a logical volume basis	Have fine control over replication at the logical volume level.
Sun StorEdge Instant Image software integration	Make point-in-time volume copies of your data for additional point-in-time coordination and recovery.
Synchronous and asynchronous volume replication between physically separate sites	Plan for disaster recovery and business continuance using physically separate primary and secondary sites. The Sun SNDR software design is link-neutral, meaning that it can use any Sun network adapter that supports TCP/IP. Continue remote data access despite local disk failure (depending on how the Sun SNDR software is configured).
Cluster capability	See “Using the Sun StorEdge Data Services Software in a Sun Cluster Environment” on page 7.

Volume Size Requirements

Note – Volume size checking is performed only when you initiate a synchronization or resynchronization.

The primary and secondary volumes must be the same size. If you initiate a resynchronization on a Sun SNDR software volume set where the secondary is larger than the primary, a warning message appears but the initiation is allowed to continue. If you initiate a resynchronization on a Sun SNDR software volume set where the secondary volume is smaller than the primary, the Sun SNDR software fails with an error.

One-to-Many and Multihop Volume Sets

The Sun SNDR software enables you to create one-to-many and multihop volume sets.

In a one-to-many volume set, you can replicate data from one primary volume to many secondary volumes residing on one or more hosts. One primary and each secondary host volume is a single volume set (each secondary volume requires its own unique secondary bitmap volume). When you perform a forward resynchronization, you can synchronize one volume set or all volume sets; in this case, issue a separate command for each set. You can also update the primary volume by using a specific secondary volume. FIGURE 1-1 shows one primary and three secondary host volumes and therefore three volume sets: A and B1, A and B2, and A and B3.

Tip – You can group one-to-many sets that share a common primary volume in a single I/O group to operate on all sets simultaneously instead of issuing a separate command for each set. See "Sun SNDR Commands and I/O Group Operations" on page 22.

In a multihop set, the secondary host volume of one volume set can be the primary host volume of another volume set. FIGURE 1-1 shows one primary and one secondary host volume; the secondary host volume B becomes the primary host volume A1 to the secondary host volume B1.

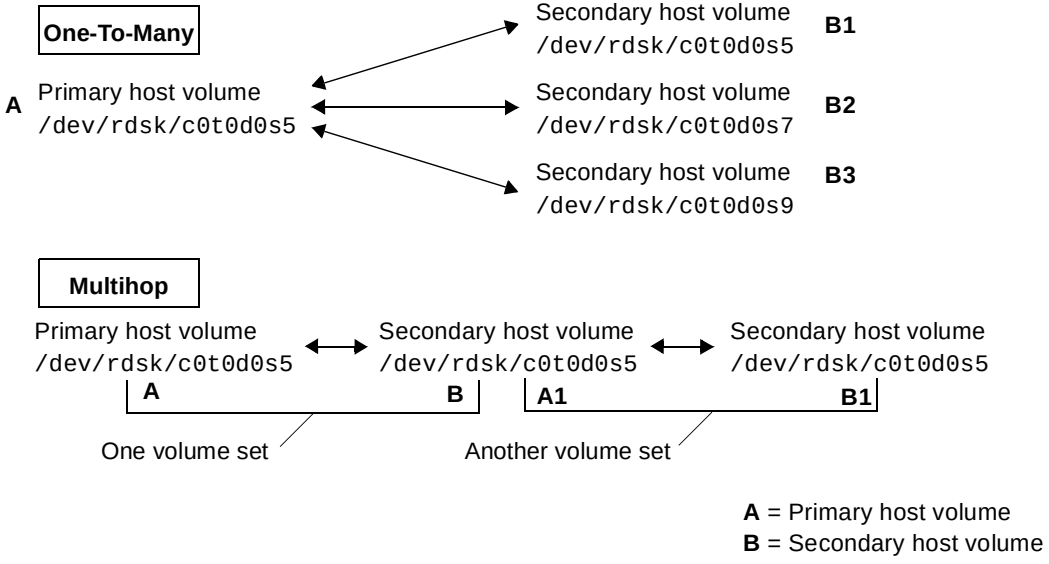

FIGURE 1-1 One-to-Many and Multihop Volume Sets

Command Compatibility With the Sun SNDR Software Version 2.0

The Sun SNDR Version 3.0 software command-line interface (CLI) `/usr/opt/SUNWesm/sbin/sndradm` is generally compatible with the Sun SNDR software Version 2.0 `rdcadm` CLI. The Version 3.0 `sndradm` command is also linked to the `rdcadm` command, so that if your existing scripts reference the `rdcadm` command, the Version 3.0 software will use the `sndradm` command to perform the specified operations.

See TABLE 3-1 on page 21.

There are two ways to use the CLI to execute commands:

- As part of an existing or new script file
- From the Solaris operating environment root user prompt

Using the Sun StorEdge Data Services Software in a Sun Cluster Environment

Caution – Do not install or use the Sun StorEdge Version 3.0 Core and data services software on servers in a Sun Cluster 3.0 environment. **The Version 3.0 software is not co-existent with the Sun Cluster 3.0 environment.** The Version 3.0 software **is** co-existent in the Sun Cluster 2.2 environment, where it will not interfere with failover. The Version 3.0 core and data services software is cluster aware in the Sun Cluster 3.0 Update 1 environment and provides high availability for the Sun StorEdge software.

You can use the `C` *tag* and `-C` *tag* `sndradm` command options described in Chapter 3 in a Sun Cluster 3.0 Update 1 environment only. If you accidentally use these options in a noncluster environment, the specified Sun SNDR operation does not execute.

You cannot use the Sun StorEdge Fast Write Cache (FWC) product, including the `SUNWnvm` Version 3.0 software, in a Sun Cluster environment because cached data is inaccessible from other machines in a cluster. To compensate, you can use a caching array such as the Sun StorEdge A3500 disk array.

Note – The Sun StorEdge data services software volumes, such as those used with the Sun SNDR and Instant Image Version 3.0 Software, do not fail over in a Sun Cluster 2.2 environment.

In a Sun Cluster 2.2 environment, the Sun StorEdge Version 3.0 data services software requires that the following patches be installed:

Solaris 7 operating environment	Patch number 109209-06
Solaris 8 operating environment	Patch number 109210-05

Before starting a Sun Cluster, ensure that the Sun SNDR volume sets do not have a copy or update operation in progress.

During a failover or takeover operation on a Sun Cluster logical host, the list of all volumes being deported is checked to determine whether they are part of one or more Sun StorEdge data services software volume sets. For each volume that is part of such a volume set, the following occurs:

- All processes accessing these volumes or raw devices will be politely killed
- All mounted file systems (`nfs`, `ufs`, or `VxFS`) will be unshared or unmounted
- The Sun StorEdge data services software volume sets will be disabled and then removed from the data services and Storage Volume (SV) driver configuration.

After a Sun Cluster failover or takeover operation, the volume sets remain disabled and require manual reconfiguration using the Sun SNDR software before you can use them.

Sun StorEdge Instant Image Software

The Sun StorEdge Instant Image Version 3.0 software is an optional software component. You can combine the Sun SNDR software with volumes from Instant Image software to create multiple copies of each device. See “Adding and Deleting Instant Image Software Volumes” on page 32.

The Instant Image software is a point-in-time volume copy data service for the Solaris operating environment. With Instant Image software, you create a volume pair: a point-in-time logical volume copy (shadow) from the original logical volume (master), which you specify. Once the shadow is established, you can read from and write to this shadow volume and the master volume.

The Instant Image software is best used just before you perform a resynchronization, to ensure that a consistent copy of data exists. In this case, if the resynchronization is interrupted, you at least have a copy of known good data that is usable.

During the resynchronization process of updating the local and remote sites, the data on a secondary Sun SNDR volume is temporarily inconsistent with the primary volume. The secondary volume cannot be relied on for data recovery. Consistency is restored when the resynchronization is complete. To ensure data integrity, use Instant Image software regularly to create a point-in-time copy of data at both sites. See the Instant Image documentation listed in “Related Documentation” on page xviii.

Network Protocols

Although the Sun SNDR software is most likely to be used with SunATM link-level interfaces, the Sun SNDR software can be used with any Sun-supported link-level interface that is TCP/IP-capable, such as Gigabit Ethernet, Gigabit Ethernet Fibre Channel, and others.

When using ATM, ensure that the configuration supports TCP/IP by using either Classical IP or LAN Emulation. For more information on configuring the SunATM interface for these protocols, refer to the *SunATM Installation and User’s Guide*.

See the network protocol manuals listed in “Related Documentation” on page xviii for more information about other protocols.

What Happens When Volume Sets Do Not Match

Physical separation provides highly available storage configurations but requires special attention to logistics. Intersite link outages, component failures, and other situations temporarily cause the primary and secondary volume to be out of sync. The Sun SNDR software uses several logging and resynchronization techniques to reestablish mirror copies:

- Full synchronization (that is, a complete volume-to-volume copy) is used when the remote volume specified for redundancy is newly formatted or its contents are unknown.
- Update resynchronization provides a streamlined way to make remote sets match if their changes have been logged earlier by the Sun SNDR software.
- Rollback resynchronization is a form of update synchronization used for disaster recovery rehearsals and other applications where test updates must be reversed. Chapter 4 describes the mechanism to use in a given circumstance. When multiple volumes require synchronization, the Sun SNDR software can perform several of these techniques concurrently.

Getting Started: Task Summary

TABLE 1-2 lists the high-level steps and tasks for getting started with the Sun SNDR software.

TABLE 1-2 Getting Started Task Summary

Task	See This Section or Chapter
1. Review the Sun SNDR software commands and supporting file requirements.	Chapter 2 and Chapter 3
2. Define the primary and secondary sites.	
3. Select the volumes for the Sun SNDR software operations.	"Volume Size Requirements" on page 5 "One-to-Many and Multihop Volume Sets" on page 6
4. Enable the Sun SNDR software for the selected volumes.	"Enabling the Sun SNDR Software" on page 28
5. Perform the Sun SNDR software copy, update, and other operations on volumes.	"Command Option and Parameter Summary" on page 20 "Full Forward Synchronization: Establishing Replicated Images for the First Time" on page 64 Chapter 4 for recovery scenarios
6. Check volume status.	"Displaying Volume Set and I/O Group Status" on page 41 "Displaying a List of Current Volume Sets and Group Names" on page 42 "Getting Sun SNDR Software and Storage Cache Statistics: The scmadm Command" on page 51

CHAPTER 2

Configuration Files and Bitmaps

This chapter describes specifying a Sun SNDR software Version 3.0 configuration file and converting configuration files from the Sun SNDR software Version 2.0. It also describes bitmap considerations and how to create bitmaps.

The topics in this chapter include:

Using a Configuration File

When you enable the Sun SNDR software using the `/usr/opt/SUNWesm/sbin/sndradm` command, you can specify an optional configuration file containing information about the volume set: volumes, primary and secondary hosts, bitmaps, Sun SNDR operating mode, and so on. You can also enter this information from the command line.

The fields for the configuration file specified using the `-f` *config-file* option are similar to those in the *SNDR-set*, described in "Command Option and Parameter Summary" on page 20:

phost pdev pbitmap shost sdev sbitmap `ip` {`sync` | `async`} [`g` *io-groupname*] [`C` *tag*]

See TABLE 2-1. See the `rdc.cf` man page for more information about this configuration file format.

TABLE 2-1 Optional Configuration File Format Fields

Primary Host	Server on which the primary volume resides.
Primary Device	Primary volume partition to be copied. Specify full path names only (for example, `/dev/dsk/c0t1d02s4`)
Primary Bitmap	Volume partition in which the bitmap (scoreboard logs) of the primary partition is stored. Specify full path names only
Secondary Host	Server on which the secondary volume resides.
Secondary Device	Primary volume partition to be copied. Specify full path names only.
Secondary Bitmap	Volume partition in which the bitmap (scoreboard logs) of the secondary partition is stored. Specify full path names only
Protocol	Network transfer protocol. Specify `ip`
Mode	{`sync` | `async`} Sun SNDR software operating mode. sync is the mode where the I/O operation is confirmed as complete only when the remote volume has been updated. async is the mode where the primary host I/O operation is confirmed as complete before updating the remote volume.
Options	[`g` *io-groupname*] [`C` *tag*] An I/O group name can be specified using the `g` character. When running in a Sun Cluster Version 3.0 Update 1 environment, a cluster resource group tag can be specified using the `C` character.

Configuration Files from Version 2.0

If you upgrade from Sun SNDR software Version 2.0 to Version 3.0, the Sun SNDR software converts your Version 2.0 configuration files to the latest format during installation. The Sun SNDR Version 2.0 software uses three configuration files:

- The default configuration file named `/etc/opt/SUNWrdc/rdc.cf` is used to specify all volumes under the Sun SNDR software control. You could also create a customized configuration file, depending on your server connection and disaster recovery plans in the Sun SNDR software Version 2.0; if this customized configuration file is named `/etc/opt/SUNWrdc/rdc.cf`, the Sun SNDR 3.0 installation process will use it. (If it is not named `rdc.cf`, include this information in the `rdc.cf` file so that you can use it in Version 3.0.)
- The `/etc/opt/SUNWrdc/rdc_ii.cf` configuration file is used to list all secondary volumes on which Sun StorEdge Instant Image software was enabled by the `rdc_ii_enable` script.
- The `/etc/opt/SUNWspsv/sv.cf` Storage Volume (SV) driver interface file is used to place Sun SNDR software volumes under SV control.

Sun SNDR software Version 3.0 enables you to keep using the same volumes that you used with the Sun SNDR software Version 2.0 as specified in the `rdc.cf`, `rdc_ii.cf`, and `sv.cf` configuration files.

When you remove the Version 2.0 software, the removal process using `pkgrm`(1M) preserves the `rdc.cf`, `rdc_ii.cf`, and `sv.cf` configuration files in their original locations. When the Sun SNDR software Version 3.0 installation process finds them in their original locations, it converts them for use with Version 3.0.

Bitmaps

Note – After editing the `/usr/kernel/drv/rdc.conf` file, reboot your server.

The Sun SNDR software uses regular raw devices to store bitmaps. These raw devices should be stored on a disk separate from the disk that contains the data. Configure RAID (such as mirrored partitions) for these bitmap devices and ensure that the mirrored members are not stored on the same disk as the data.

In a clustered environment, a bitmap must reside only on a volume. The bitmap volume in this case must be part of the same disk group or cluster resource group as the corresponding primary or secondary data volume.

A bitmap maintained on disk might persist across a system crash, depending on the setting of `rdc_bitmap_mode` in `/usr/kernel/drv/rdc.conf`. For example:

```
# rdc_bitmap_mode
# - Sets the mode of the RDC bitmap operation, acceptable values are:
#   0 - autodetect bitmap mode depending on the state of SDBC (default).
#   1 - force bitmap writes for every write operation, so an update resync
#       can be performed after a crash or reboot.
#   2 - only write the bitmap on shutdown, so a full resync is
#       required after a crash, but an update resync is required after
#       a reboot.
#
rdc_bitmap_mode=0;
```

If your server is configured in a clustered environment, set the bitmap mode to 1. If your server is not configured in a clustered environment, you can also choose the bitmap mode setting of 1 to improve error recovery.

Bitmap Size Requirements

The bitmap size can be calculated using the following formula:

- 1 Kbytes + 4 Kbytes per Gbyte of device storage space

For example, a 2-Gbyte data device requires a bitmap size of 9 Kbytes. (You can create bitmaps that are larger than the calculated size. Also, round up any partial Gbyte counts to the next Gbyte.)

Using More Than 64 Sun SNDR Software Volume Sets

Note – After editing the `/usr/kernel/drv/rdc.conf` file, reboot your server.

If you configure more than 64 Sun SNDR software volume sets, you must edit the `rdc_max_sets` field in the `/usr/kernel/drv/rdc.conf` file on each machine running the Sun SNDR software. The default number of configured volume sets is 64.

For example, to use 128 sets, change the file as follows; note the semicolon (;) at the end of the `rdc_max_sets` field:

```
#
# rdc_max_sets
# - Configure the maximum number of RDC sets that can be enabled on
# this host. The actual maximum number of sets that can be
# enabled will be the minimum of this value and nsc_max_devices
#(see nsctl.conf) at the time the rdc kernel module is loaded.
#
rdc_max_sets=128;
```

CHAPTER 3

Sun SNDR Software Commands

Caution – Do not install or use the Sun StorEdge Version 3.0 Core and data services software on servers in a Sun Cluster 3.0 environment. **The Version 3.0 software is not co-existent with the Sun Cluster 3.0 environment.** The Version 3.0 software **is** co-existent in the Sun Cluster 2.2 environment, where it will not interfere with failover. The Version 3.0 core and data services software is cluster aware in the Sun Cluster 3.0 Update 1 environment and provides high availability for the Sun StorEdge software.

Note – To use the Sun SNDR software, you must be the root user.

This chapter describes the Sun SNDR commands available from the command line using `/usr/opt/SUNWesm/sbin/sndradm`. Chapter 4 describes how to use the command in various data recovery scenarios.

The topics in this chapter include:

Command Option and Parameter Summary

The `/usr/opt/SUNWesm/sbin/sndradm` command uses the options in TABLE 3-2 and the parameters in TABLE 3-3.

Note – If you use a command without specifying any volume set arguments, the command operates on all Sun SNDR volume sets.

Log File

The `/var/opt/SUNWesm/ds.log` file contains operation log messages for the Sun SNDR commands.

Command Compatibility with Version 2.0

The Sun SNDR software Version 3.0 `sndradm` command is generally compatible with the Sun SNDR software Version 2.0 `rdcadm` command. The Version 3.0 `sndradm` command is also linked to the `rdcadm` command, so that if your existing scripts reference the `rdcadm` command, the Version 3.0 software will use the `sndradm` command to perform the specified operations.

TABLE 3-1 shows the instances where the two command versions differ. The Version 3.0 command provides more options and capabilities, such as the ability to operate on groups of volume sets.

TABLE 3-1 `rdcadm` and `sndradm` Command Differences Between Sun SNDR Software Version 2.0 and Version 3.0

Sun SNDR Software Version 2.0	Sun SNDR Software Version 3.0	Difference
`rdcadm`, all options	`sndradm` and `rdcadm`, all options	When you use any option and you do not specify a *SNDR-set*: Version 2.0 operates on all *SNDR-sets* in the `rdc.cf` file or those as specified by `-f` *config-file* Version 3.0 operates on all enabled *SNDR-sets* or those as specified by `-f` *config-file*
`rdcadm -a` *value*	`sndradm -a` {`on` | `off`}	Version 2.0: *value* is 0 (off) or 1 (on) Version 3.0: specify `on` or `off` instead of 0 or 1
`rdcadm -A`	`sndradm -P`	Version 2.0: shows the autosynchronization state Version 3.0: shows detailed information about all *SNDR-sets*

Sun SNDR Commands and I/O Group Operations

Adding the Sun SNDR software volume sets to an I/O group enables you to issue a single command that operates on all volume sets in the specified I/O group or groups, excluding all other volumes sets from the operations. Most commands allow for group operations and perform them when you include `-g` *io-groupname* in the command syntax.

The operations performed are independent of each other; that is, operations performed on I/O group A, volume set 1 are independent of operations performed on I/O group A, volume set 2.

If an operation fails on one or more volume sets in an I/O group, the state of the data on the volumes in the I/O group's volume sets is unknown. To correct this:

- Correct any known problems with the failing sets.
- Reissue the command on the entire I/O group or individual volume sets that failed.

If the operation is successful, the I/O group and volume set state returns to a normal condition; that is, the data stored on all replicated volumes is consistent.

Sun SNDR Commands and Sun Cluster Operations

Caution – Do not install or use the Sun StorEdge Version 3.0 Core and data services software on servers in a Sun Cluster 3.0 environment. **The Version 3.0 software is not co-existent with the Sun Cluster 3.0 environment.** The Version 3.0 software **is** co-existent in the Sun Cluster 2.2 environment, where it will not interfere with failover. The Version 3.0 core and data services software is cluster aware in the Sun Cluster 3.0 Update 1 environment and provides high availability for the Sun StorEdge software.

You can use the `C` *tag* and `-C` *tag* options described in TABLE 3-2 in a Sun Cluster 3.0 Update 1 environment only. If you accidentally use these options in a noncluster environment, the Sun SNDR operation does not execute.

Configuration File Format Using `-f` *config-file*

The fields for the configuration file specified using the `-f` option are similar to those in the *SNDR-set*, described in “Command Options” on page 23:

phost pdev pbitmap shost sdev sbitmap `ip` {`sync` | `async`} [`g` *io-groupname*] [`C` *tag*]

See “Using a Configuration File” on page 14 and the `rdc.cf` man page for more information.

TABLE 3-2 Command Options

Option	Definition
`-n`	Does not prompt the user after starting a Sun SNDR software operation using `sndradm`. The default is to prompt and request a response. For example, after starting a full synchronization from the primary to secondary volumes, the Sun SNDR software prompts `Overwrite secondary with primary? (Y/N) [N]`.
`-f` *config-file*	Specifies a configuration file defining the Sun SNDR software volume sets. If you do not specify a *config-file*, *SNDR-set*, or *set-name*, the command acts on all volume sets.
SNDR-set	Specifies the complete Sun SNDR software set configuration information. If *SNDR-set* is not specified on the command line, the information is retrieved from a configuration file you specify using the `-f` *config-file* option. The *SNDR-set* format is as follows:

phost pdev pbitmap shost sdev sbitmap `ip {sync | async} [g` *io-groupname*`] [C` *tag*`]`

phost - Server on which the primary volume resides.

pdev - Primary volume partition to be copied. Specify full path names only; for example, `/dev/rdsk/c0t1d0s2`.

pbitmap - Volume partition or file in which the bitmap (scoreboard log) of the primary partition is stored. Specify full path names only; for example, `/dev/rdsk/c0t1d0s2`.

shost - Server on which the secondary volume resides.

sdev - Secondary volume partition. Specify full path names only; for example, `/dev/rdsk/c0t1d0s2`.

sbitmap - Volume partition or file in which the bitmap (scoreboard log) of the secondary partition is stored. Specify full path names only; for example, `/dev/rdsk/c0t1d0s2`.

`ip` - Specifies the network protocol

TABLE 3-2 Command Options *(Continued)*

Option	Definition
	`sync	async` - Sun SNDR software operating mode. `sync` is the Sun SNDR software mode where the I/O operation is not confirmed as complete until the remote volume has been updated. `async` is the Sun SNDR software mode where the primary host I/O operation is confirmed as complete before updating the remote volume.
	`g` *io-groupname* - Specifies which I/O group the volume set belongs to.
	You can use the C tag and -C tag options in a Sun Cluster 3.0 Update 1 environment only. If you accidentally use these options in a non-Sun Cluster 3.0 environment, the Sun SNDR operation does not execute.
	`C` *tag* - For operation within a Sun Cluster 3.0 Update 1 environment only.
	`C` *tag* specifies the disk group name or resource tag of the local data and bitmap volumes in cases where this is not implied by the name of the volume. (For example, `/dev/rdsk/md/dg/vol` and `/dev/vx/rdsk/dg/vol` both indicate a disk group name of `dg.`) It is the responsibility of the user to ensure that the cluster tag specified to the Sun SNDR software matches the appropriate cluster resource group tag and to keep all installed data services updated in the event of cluster resource group reconfigurations.
set-name	Name of the Sun SNDR software volume set as assigned by the Sun SNDR software. The Sun SNDR software assigns a default volume set name of *shost*:*sdev*, where *shost* is the secondary host name and *sdev* is the secondary volume partition name, separated by a colon (:).
`-g` *io-groupname*	Name of the I/O group containing the collection of Sun SNDR software volume sets.
	Specifying `-g` *io-groupname* limits any operation performed to the sets in the named *io-groupname*.
`-C` *tag*	You can use the `C` *tag* and `-C` *tag* options in a Sun Cluster 3.0 Update 1 environment only. If you accidentally use these options in a noncluster environment, the Sun SNDR operation does not execute.
	Specifying `-C` *tag* limits any operation performed to the sets in the named *tag* cluster resource tag.

TABLE 3-3 Sun SNDR Software Tasks and Related Command Parameters

Operation	Command	For Detailed Information, See:
Enable the Sun SNDR software for the specified set (sync not required).	`sndradm -E`	"Enabling the Sun SNDR Software" on page 28
Enable the Sun SNDR software for the specified set (sync required).	`sndradm -e`	"Enabling the Sun SNDR Software" on page 28
Disable the Sun SNDR software for the specified set.	`sndradm -d`	"Disabling the Sun SNDR Software" on page 30
Add or delete a Sun StorEdge Instant Image software volume group for use with the Sun SNDR software.	`sndradm -I`	"Adding and Deleting Instant Image Software Volumes" on page 32
Copy the entire contents of the primary volume to the secondary volume; also known as *full synchronization*.	`sndradm -m`	"Starting a Full Synchronization (Copy Between Volumes)" on page 34
Copy the entire contents of the secondary volume to the primary volume; also known as *full reverse synchronization*.	`sndradm -m -r`	"Starting a Full Synchronization (Copy Between Volumes)" on page 34
Update only the changed data of the primary volume to the secondary volume; also known as *update synchronization* or *resynchronization*.	`sndradm -u`	"Starting a Resynchronization (Update Between Volumes)" on page 36
Update only the changed data of the secondary volume to the primary volume; also known as reverse synchronization or reverse update.	`sndradm -u -r`	"Starting a Resynchronization (Update Between Volumes)" on page 36
Wait for a synchronization operation to finish executing.	`sndradm -w`	"Waiting for Synchronization to Finish Executing" on page 38
Display the Sun SNDR command syntax and version information from the command prompt.	`sndradm -h` `sndradm -v`	Results from entering this command at the command prompt
Display the Sun SNDR software status.	`sndradm -p` `sndradm -P`	"Displaying Volume Set and I/O Group Status" on page 41
Display the Sun SNDR software volume set and I/O group names.	`sndradm -i`	"Displaying a List of Current Volume Sets and Group Names" on page 42
Display the status of the link connecting systems running the Sun SNDR software.	`sndradm -H`	"Displaying Link Status" on page 43

TABLE 3-3 Sun SNDR Software Tasks and Related Command Parameters *(Continued)*

Operation	Command	For Detailed Information, See:
Change the scoreboard bitmap of a volume set.	`sndradm -R b`	"Assigning a New Scoreboard Bitmap to a Volume Set" on page 40
Reset the replication mode of a volume set.	`sndradm -R m {sync|async}`	"Resetting the Replication Mode of a Volume Set" on page 47
Set the asynchronous queue.	`sndradm -W` `sndradm -F`	"Setting the Asynchronous Queue" on page 50
Stop the Sun SNDR software replication and log scoreboards.	`sndradm -l`	"Stopping Sun SNDR Software Operations and Starting Logging" on page 39
Toggle the autosynchronization state.	`sndradm -a {on|off}`	"Toggling the Sun SNDR Software Autosynchronization State" on page 49
Move all the specified SNDR sets to a different group.	`sndradm -R g`	"Moving a Volume Set to a Different I/O Group" on page 44
Update or reconfigure the Sun SNDR software volume sets.	`sndradm -R -f` *config-file*	"Updating the Volume Set Configuration" on page 45
Update or reconfigure the disk group name or cluster resource tag.	`sndradm -R C`	"Updating the Disk Cluster Tag Name" on page 46

Which Host Do I Issue Commands From?

Certain commands and synchronization situations require you to issue a command from a primary host, secondary host, or both. See TABLE 3-4.

TABLE 3-4 Which Host to Issue Commands From

Task	Where Command Is Issued	Comments
Assign a new bitmap to a volume set.	Primary and secondary host	Perform this command first on the host where the new bitmap resides and is being assigned, and then perform it on the other host.
Disable the Sun SNDR software.	Primary or secondary host	You can disable on one host, leave the other host enabled, and then re-enable the disabled host.
	Primary and secondary host	Perform this operation on both hosts if you are deleting a volume set.
Enable the Sun SNDR software.	Primary and secondary host	When enabling the Sun SNDR software for the first time, issue the command from both hosts.
Full forward or reverse synchronization (copy).	Primary host	Ensure that both hosts are enabled.
Forward or reverse synchronization (update).	Primary host	Ensure that both hosts are enabled.
Log.	Primary host	Perform on primary only if a synchronization is in progress.
		Perform on the secondary host if the primary host failed.
	Primary or secondary host	Perform on either host if no synchronization is in progress
Toggle the autosynchronization state.	Primary	
Update an I/O group.	Primary and secondary	

Enabling the Sun SNDR Software

Note – When you first enable the SNDR software, you must enable it on the primary and secondary hosts.

Use the commands to enable the SNDR volume set and start scoreboard logging. Also use these commands to create one-to-many and multihop sets. See "One-to-Many and Multihop Volume Sets" on page 6.

When you enable the Sun SNDR software using `sndradm -e` or `sndradm -E`, the SNDR software is enabled for the specified *SNDR-set* and optional I/O group. These commands also enable you to add volume sets to I/O groups when you specify the `-g` *io-groupname* as part of the *SNDR-set* options.

Using a Volume Set Name

When you first enable a set, the Sun SNDR software assigns a default volume set name of *shost*:*sdev*, where *shost* is the secondary host name and *sdev* is the secondary volume partition name, separated by a colon (:). The volume set name is referred to in this document as *set-name*.

After executing these commands, you can use the *shost*:*sdev* name for a volume set each time you issue a Sun SNDR command, instead of specifying the complete primary and secondary host, volume, and bitmap information for a volume set.

`sndradm -e`

This command sets the bitmap scoreboard to indicate that a full resynchronization between volumes is required and it enables scoreboard logging. It also adds the local volume to the Storage Volume (SV) driver control; the bitmap scoreboard volume is also added to the SV driver.

Syntax

`sndradm -e [-g` *io-groupname*`] [-C` *tag*`] [-n] {-f` *config-file* | *SNDR-set*`}`

sndradm -E

Caution – Before using `sndradm -E`, ensure that the volumes are already synchronized by means other than the Sun SNDR software (for example, restoring the volumes from tape or other media). If the volumes are not synchronized before this command is used, the secondary or target volume will contain inconsistent data.

This command clears the bitmap scoreboard logging to indicate that the specified volumes are fully synchronized and enables scoreboard logging. It also adds the local volume set to the Storage Volume (SV) driver control; the bitmap scoreboard volume is also added to the SV driver.

Syntax

`sndradm -E` [`-g` *io-groupname*] [`-C` *tag*] [`-n`] {`-f` *config-file* | *SNDR-set*}

Disabling the Sun SNDR Software

Note – See also “To Remove the Volumes From SV Control and the Sun StorEdge Configuration” on page 31.

Use the `sndradm -d` command when the primary and secondary volumes no longer need to be associated with each other as Sun SNDR software volume sets.

This command discards any active Sun SNDR software scoreboard logs. See “Enabling the Sun SNDR Software” on page 28 to re-enable Sun SNDR software replication and “Starting a Full Synchronization (Copy Between Volumes)” on page 34 to reestablish identical replicated sets.

`sndradm -d`

When you issue this command, the Sun SNDR software terminates all replication services between the specified primary and secondary volumes and discontinues the relationship between these volume sets. It also discards any active scoreboard logs that track temporary differences between primary and secondary volumes.

Note – This command also removes the information for the specified set from the Sun StorEdge data services.

Syntax

`sndradm -d` [`-g` *io-groupname*] [`-C` *tag*] [`-n`] [`-f` *config-file* | *SNDR-set* | *set-name*]

▼ To Remove the Volumes From SV Control and the Sun StorEdge Configuration

The Sun SNDR `sndradm -e` and `sndradm -E` commands automatically add configured volumes to the SV driver's configuration. During the disabling of the Sun SNDR volumes, the volumes are not automatically removed from the Sun StorEdge configuration. Also, you might accidentally remove enabled volumes using the Sun StorEdge SV driver `svadm -d` command.

The `svadm` command currently has no checking in place to prevent a configured volume from being removed from an active Sun SNDR volume set. To remove a volume, ensure that it is not configured by using the `sndradm -i` command options to display volume information. If the volumes are not listed, they can be safely removed.

1. **Disable the volumes in the Sun StorEdge software by using the `sndradm -d` command.**
2. **Remove the volumes from SV control using the `svadm -d` command.**

Adding and Deleting Instant Image Software Volumes

Use the `sndradm -I` command to add or delete Instant Image software volumes for use with enabled Sun SNDR software volume sets. You can use this command on the secondary and primary hosts to offer data protection during reverse and forward synchronization update operations. You also use this command to offer data protection during an automatic resynchronization; see "Autosynchronization With Instant Image Software Configured" on page 69.

Note – Use the `sndradm -e` command to enable Sun SNDR software volume sets first before using this command. "Enabling the Sun SNDR Software" on page 28 describes this command.

`sndradm -I a`

This command adds an Instant Image volume group entry to the data services in a format similar to that of the `rdc_ii.cf` file used with the Sun SNDR Version 2.0 software. (See "Configuration Files from Version 2.0" on page 15.)

Syntax

`sndradm -I a` *`master-vol shadow-vol bitmap-vol`*

where *`master-vol`*, *`shadow-vol`*, and *`bitmap-vol`* are the full path to the raw device nodes for configured Instant Image volumes. See the *Sun StorEdge Instant Image Software 3.0 System Administrator's Guide* for more information.

Note – Before the `/usr/opt/SUNWrdc/lib/sndrsyncd` daemon performs an update synchronization after a network link or machine failure, it creates a matching Instant Image volume group entry. After the update finishes executing, you may choose to delete this entry with the `/usr/opt/SUNWscm/sbin/iiadm -d` command. However, do not delete this entry so that the daemon can use it in subsequent resynchronizations involving the Instant Image software. See "Autosynchronization With Instant Image Software Configured" on page 69.

---

sndradm -I d

This command deletes an Instant Image volume group entry from the data services.

Syntax

`sndradm -I d` *master-vol shadow-vol bitmap-vol*

where *master-vol*, *shadow-vol*, and *bitmap-vol* are the full path to the raw device nodes for configured Instant Image volumes. See the *Sun StorEdge Instant Image 3.0 System Administrator's Guide* for more information.

Starting a Full Synchronization (Copy Between Volumes)

Note – If a synchronization initiated by the `sndradm -m` or `sndradm -m -r` commands is interrupted, use the `sndradm -u` update command to complete the synchronization. Use `/usr/opt/SUNWscm/sbin/scmadm -S` to view the synchronization progress.

Use the `sndradm -m` command when all of the following cases exist:

- The contents of the primary and secondary volumes may be inconsistent
- No scoreboard logging information exists to incrementally resynchronize the volumes
- The primary volume has the desired contents
- You wish to completely overwrite the contents of the secondary volume

Use the `sndradm -m -r` command when all of the following cases exist:

- The contents of the primary and secondary volumes may be inconsistent
- No scoreboard logging information exists to incrementally resynchronize the volumes
- The secondary volume has the desired contents
- You wish to completely overwrite the contents of the primary volume

`sndradm -m`

This command enables you to start a full copy operation from the primary volume to the secondary volume. It also enables replication concurrently from the primary volume to the secondary volume; any new writes to the primary volume are also replicated to the secondary volume.

Syntax

`sndradm -m [-g` *io-groupname*`] [-C` *tag*`] [-n] [-f` *config-file* | *SNDR-set* | *set-name*`]`

sndradm -m -r

This command enables you to start a full reverse copy operation from the secondary volume to the primary volume. It also enables replication from the primary volume to the secondary volume; any new writes to the primary volume are also replicated to the secondary volume.

Syntax

`sndradm -m -r` [`-g` *io-groupname*] [`-C` *tag*] [`-n`] [`-f` *config-file* | *SNDR-set* | *set-name*]

Note – Do not start the primary application (such as a database application) that writes to the volumes until the full reverse copy operation finishes executing.

Starting a Resynchronization (Update Between Volumes)

Note – If a synchronization initiated by the `sndradm -m` or `sndradm -m -r` commands is interrupted, use the `sndradm -u` update command to finish executing the synchronization. Use `/usr/opt/SUNWscm/sbin/scmadm -S` to view the synchronization progress.

Use the `sndradm -u` command when all of the following cases exist:

- Replication was stopped by using the `sndradm -l` command or the link was interrupted, and scoreboard logging was active.
- The primary volume has the desired contents.
- You want to overwrite those segments of the secondary volume that are different from the primary as defined in the scoreboard logs.
- You intend to start replication from the primary to the secondary volume.

Use the `sndradm -u -r` command when all of the following cases exist:

- Replication was stopped using the `sndradm -l` command or the link was interrupted, and scoreboard logging was active.
- The secondary volume has the desired contents.
- You want to overwrite those segments of the primary volume that are different from the secondary as defined in the scoreboard logs.
- You intend to start replication from the primary to the secondary volume.

`sndradm -u`

This command resynchronizes the secondary volume from the primary volume. It updates the secondary volume according to the changes based on scoreboard logs maintained while replication was stopped. It also enables concurrent replication between the primary and secondary volumes; any new write operations to the primary volumes are also replicated to the secondary volumes.

Syntax

`sndradm -u` [`-g` *io-groupname*] [`-C` *tag*] [`-n`] [`-f` *config-file* | *SNDR-set* | *set-name*]

sndradm -u -r

Tip – You can use this command to roll back volume changes on the primary volume to a point-in-time image captured on the stopped secondary volume.

This command resynchronizes the primary volume from the secondary volume. It updates the primary volume according to the changes based on scoreboard logs maintained while replication was stopped. It also enables concurrent replication between the primary volume and secondary volumes; any new write operations to the primary are also replicated to the secondary volumes.

Syntax

`sndradm -u -r` [`-g` *io-groupname*] [`-C` *tag*] [`-n`] [`-f` *config-file* | *SNDR-set* | *set-name*]

Waiting for Synchronization to Finish Executing

Use the `sndradm -w` command in one of the following situations:

- When you need to be sure that the copy or update operation is finished executing before you use a primary application (like a database application) or another Sun SNDR software command on the specified volume set or I/O group
- Before disabling a volume set, so that any write or copy operations can finish executing
- As part of a script executing Sun SNDR software commands

`sndradm -w`

This command enables you to make the Sun SNDR software wait until an in-progress copy or update synchronization operation finishes executing. It prevents another Sun SNDR software command from executing on the specified volume set or I/O group.

Syntax

`sndradm -w [-g` *io-groupname*`] [-C` *tag*`] [-n] [-f` *config-file* | *SNDR-set* | *set-name*`]`

Stopping Sun SNDR Software Operations and Starting Logging

Use the `sndradm -l` command to start status and scoreboard logging on enabled volume sets and I/O groups.

Note – To resume Sun SNDR software operations after using the `sndradm -l` command, use the `sndradm -m` command to perform a full resynchronization or the `sndradm -u` command to perform an update resynchronization. Note also that, when issued from the secondary host, the `sndradm -l` command does not work on the secondary volume for any volume that is currently synchronizing.

`sndradm -l`

This command enables you to stop replication between primary and secondary volumes and to start scoreboard logging on these volumes. The Sun SNDR software will continue logging if the operations of a volume set or I/O group are interrupted.

If all volume sets in an I/O group are replicating (meaning that the secondary volumes contain a valid point-in-time copy of the corresponding primary volumes), when one volume set enters logging mode, all other sets in the group will enter logging mode automatically. This scheme ensures that the secondary volumes will contain a valid point-in-time copy.

Syntax

`sndradm -l` [`-g` *io-groupname*] [`-C` *tag*] [`-n`] [`-f` *config-file* | *SNDR-set* | *set-name*]

Assigning a New Scoreboard Bitmap to a Volume Set

Use the `sndradm -R b` command when you need to assign a new scoreboard bitmap to a volume set.

Tip – If you upgraded from the Sun SNDR Version 2.0 software and used bitmap files instead of bitmap volumes in Version 2.0, use this command to assign new primary and secondary host bitmap *volumes* to the Sun SNDR volume set. This command copies any data from the bitmap file to the bitmap volume. The *Sun StorEdge Network Data Replicator 3.0 Installation Guide* describes this procedure.

`sndradm -R b`

This command enables you to assign a new scoreboard bitmap to an existing volume set. You can change the bitmaps on the primary or secondary host. This command copies any data from the old bitmap to the new bitmap.

Syntax

`sndradm -R b {p|s}` *new-bitmap-name* `[-g` *io-groupname*`] [-C` *tag*`] [-n] [-f` *config-file* `|` *SNDR-set* `|` *set-name*`]`

- To change a primary host bitmap, issue the command from the primary host.
- To change a secondary host bitmap, issue the command from the secondary host.
- Optionally, to ensure consistency for status reporting, issue the command from both hosts. For example, to change a secondary host bitmap, issue the command from the secondary host first and then issue it from the primary host.

Displaying Volume Set and I/O Group Status

Use the `sndradm -p` and `sndradm -P` commands to display status information. Use the `-g` *io-groupname* and `-C` *tag* options to filter the output so that it includes only those volumes that are part of *io-groupname* or *tag*.

sndradm -p

This command displays a brief version of the Sun SNDR software status for volume sets and I/O groups. The following is an example of output of this command on the primary host:

```
# sndradm -p
/dev/rdsk/c4t96d0s1     <-      fast7:/dev/rdsk/c2t0d0s1
/dev/rdsk/c4t97d0s1     <-      fast7:/dev/rdsk/c2t0d0s1
/dev/rdsk/c4t98d0s1     <-      fast7:/dev/rdsk/c2t0d0s1
```

Syntax

`sndradm -p [-g` *io-groupname*`] [-C` *tag*`]`

sndradm -P

This command displays a detailed version of the Sun SNDR software status for volume sets and I/O groups. The following is an example of output of this command on the primary host:

```
# sndradm -P
/dev/rdsk/c4t96d0s1     <-      fast7:/dev/rdsk/c2t0d0s1
autosync: off, max q writes: 60, max q fbas: 800, mode: sync

/dev/rdsk/c4t97d0s1     <-      fast7:/dev/rdsk/c2t0d0s1
autosync: off, max q writes: 60, max q fbas: 800, mode: sync

/dev/rdsk/c4t98d0s1     <-      fast7:/dev/rdsk/c2t0d0s1
autosync: off, max q writes: 60, max q fbas: 800, mode: async
```

Syntax

`sndradm -P [-g` *io-groupname*`] [-C` *tag*`]`

Displaying a List of Current Volume Sets and Group Names

Use the `sndradm -i` command to display volume set and I/O group names.

Tip – You can also use the output of this command to create a configuration file suitable for editing and using with the `sndradm -R -f` *config-file* command.

`sndradm -i`

This command enables you to display a list of the current volume set and I/O group names in use. The following is an example of the output:

```
# sndradm -i

fast7 /dev/rdsk/c2t0d0s1 /dev/rdsk/c2t1d0s0 fast8 /dev/rdsk/c4t96d0s1
/dev/rdsk/c6t0d0s4 ip sync

fast7 /dev/rdsk/c2t0d0s1 /dev/rdsk/c2t1d0s3 fast8 /dev/rdsk/c4t97d0s1
/dev/rdsk/c6t0d0s6 ip sync

fast7 /dev/rdsk/c2t0d0s1 /dev/rdsk/c2t1d0s4 fast8 /dev/rdsk/c4t98d0s1
/dev/rdsk/c6t0d0s8 ip async
```

Note that the output is formatted as it would be entered in a *config-file* :
phost pdev pbitmap shost sdev sbitmap `ip` {`sync`|`async`}.

Syntax

`sndradm -i` [`-g` *io-groupname*] [`-C` *tag*]

Use the `-g` *io-groupname* and `-C` *tag* options to filter the output so that it includes only those volumes that are part of *io-groupname* or *tag.*

Displaying Link Status

Use the `sndradm -H` command to check the status of the link connecting primary and secondary hosts.

Note – Use the `ping`(1M) command to supplement the link status shown by `sndradm -H`.

`sndradm -H`

This command displays the currently configured primary and secondary host machine names, volumes, and bitmap volumes for the selected sets. It also shows whether the link between machines is active or inactive. For example:

```
# sndradm -H atm-fred:/dev/vx/rdsk/freddg/sndr_vol01
Report SNDR link health? (Y/N) [N]: y

SNDR: atm-ethel /dev/rdsk/c3t9d0s3 /dev/rdsk/c6t0d0s4
atm-fred /dev/vx/rdsk/freddg/sndr_vol01 /dev/rdsk/c6t0d0s6
Inactive
```

`Active` indicates that replication or synchronization operations are occurring on the link used by the specified volume set. `Inactive` indicates that the link might be down or that all sets are in logging mode.

Syntax

`sndradm -H` *shost*:*sdev*

where *shost* is the secondary host name and *sdev* is the secondary volume partition name, separated by a colon (:). (When you first enable a set, the Sun SNDR software assigns a default volume set name of *shost*:*sdev*.)

If you specify no arguments, the Sun SNDR software displays all configured volume sets; however, for best display results, specify a volume set using *shost:sdev*.

Moving a Volume Set to a Different I/O Group

Use the `sndradm -R g` command to move volume sets to a different I/O group. To move multiple sets to a single *io-groupname*, keep the *io-groupname* the same and specify different *set-name*s on the command line.

Tip – You can also move all volume sets in one I/O group to another I/O group. Use the command as follows: `sndradm -g` *io-groupname* `-R g` *new-io-groupname*

To move selected sets from an I/O group, enter the volume set information in a file and use the `-f` *config-file* option.

Tip – To remove a volume set from an I/O group, use double quotes to specify a null group used with the `-g` command option as follows:
`sndradm -R g "" {`*SNDR-set* `|` *set-name*`}`

To remove all volume sets from an I/O group:
`sndradm -R g "" -g` *io-groupname*

To remove selected sets from an I/O group, enter the volume set information in a file and use the `-f` *config-file* option.

`sndradm -R g`

This command enables you to move volume sets to a different I/O group, thereby updating an existing I/O group. You must specify at least one *set-name*. Enter this command on the primary and secondary hosts.

Syntax

`sndradm -R g` *io-groupname* [`-g` *io-groupname*] [`-C` *tag*] [`-n`] [`-f` *config-file* | *SNDR-set* | *set-name*]

Updating the Volume Set Configuration

Use the `sndradm -R -f` command to change the current volume sets being operated on by the Sun SNDR software. This command is useful when you use a configuration file and have made changes to it.

The fields for the configuration file specified using the `-f` option are similar to those in the *SNDR-set*, described in TABLE 3-2:

phost pdev pbitmap shost sdev sbitmap `ip` {`sync` | `async`} [`g` *io-groupname*] [`C` *tag*]

See “Using a Configuration File” on page 14 and the `rdc.cf` man page for more information.

`sndradm -R -f`

This command enables you to update or reconfigure the current volume set configuration from a specified configuration file.

Syntax

`sndradm -R -f` *config-file* [`-g` *io-groupname*] [`-C` *tag*] [`-n`]

Updating the Disk Cluster Tag Name

Caution – Do not install or use the Sun StorEdge Version 3.0 Core and data services software on servers in a Sun Cluster 3.0 environment. **The Version 3.0 software is not co-existent with the Sun Cluster 3.0 environment.** The Version 3.0 software **is** co-existent in the Sun Cluster 2.2 environment, where it will not interfere with failover. The Version 3.0 core and data services software is cluster aware in the Sun Cluster 3.0 Update 1 environment and provides high availability for the Sun StorEdge software.

Use the `sndradm -R C` command in cases where the disk group name or cluster resource tag are not indicated by the volume path name. This command does not affect remote volumes and cannot be used in noncluster environments.

Tip – You can also move all volume sets in one cluster tag to another cluster tag. Use the command as follows:

`sndradm -C` *tag* `-R C` *new-tag*

`sndradm -R C`

This command enables you to update or reconfigure the current disk group name or cluster resource tag of the local volume in a volume set. *tag* here is defined as the disk group name or cluster resource tag.

Syntax

`sndradm -R C` *tag* [`-g` *io-groupname*] [`-C` *tag*] [`-n`] [`-f` *config-file* | *SNDR-set* | *set-name*]

Resetting the Replication Mode of a Volume Set

Use the `sndradm -R m` command to change the replication mode of a volume set.

Note – If the volume set is part of an I/O group, remove the volume set from the group and change its mode. You can remove one volume set, selected volume sets, or all volume sets from a group using the `sndradm -R g` command. See "Moving a Volume Set to a Different I/O Group" on page 44 and the procedure described in "To Remove a Volume Set from a Group and Change Its Replication Mode" on page 48.

`sndradm -R m`

This command enables you to reconfigure the replication or mirroring mode of a volume set.

Syntax

`sndradm -R m` {`sync` | `async`} [`-C` *tag*] [`-n`] [`-f` *config-file* | *SNDR-set* | *set-name*]

Considerations

- **Mixed mode I/O group** – Creating a mixed I/O group, where mirroring modes are asynchronous for one set and synchronous for another set, is not allowed.
- **Volume set requirements** – All volume sets in the I/O group must have the same primary host, secondary host, and mirroring mode.

▼ To Remove a Volume Set from a Group and Change Its Replication Mode

1. **Remove the volume set from its group.**

```
# sndradm -R g ""{SNDR-set | set-name}
```

2. **Change the volume set replication mode.**

See also “Syntax” on page 47.

```
# sndradm -R m {sync|async} set-name
```

3. **If desired, add the volume sets you have changed to an I/O group.**

See also “Moving a Volume Set to a Different I/O Group” on page 44.

```
# sndradm -R g io-groupname [-f config-file |SNDR-set | set-name]
```

Toggling the Sun SNDR Software Autosynchronization State

Use the `sndradm -a` command to enable or disable autosynchronization.

`sndradm -a`

This command enables or disables the Sun SNDR software autosynchronization. If Sun SNDR software autosynchronization is enabled on the primary host, the synchronization daemon attempts to resynchronize if the system reboots or link failures occur.

The default state is disabled.

Syntax

`sndradm -a {on|off} [-g` *io-groupname*`] [-C` *tag*`] [-n] [-f` *config-file* `|` *SNDR-set* `|` *set-name*`]`

Setting the Asynchronous Queue

Use the following commands to set the maximum number of writes or 512 Kbyte blocks that can be queued in asynchronous mode.

sndradm -W

This command enables you to set the maximum number of write operations queued to an asynchronous-mode volume set. The default *value* number is 60. For example, set this value to 1 to ensure that the secondary volume is never more than one write operation behind the primary volume.

Syntax

`sndradm -W` *value* [`-g` *io-groupname*] [`-C` *tag*] [`-n`] [`-f` *config-file* | *SNDR-set* | *set-name*]

sndradm -F

This command enables you to set the maximum number of 512 Kbyte blocks in the asynchronous queue. The default number is 800.

Syntax

`sndradm -F` *value* [`-g` *io-groupname*] [`-C` *tag*] [`-n`] [`-f` *config-file* | *SNDR-set* | *set-name*]

Getting Sun SNDR Software and Storage Cache Statistics: The scmadm Command

This command displays storage cache information and can be used to view Sun SNDR software statistics. Refer to the scmadm man page for more information about this command.

scmadm -S

After you type /usr/opt/SUNWscm/sbin/scmadm -S with the desired options, you can use the keys described in TABLE 3-5.

TABLE 3-5 Keys Used in scmadm Display

Press This Key. . .	To . . .
F	Scroll forward through the display.
B	Scroll backward through the display.
M	Toggle between regular cache statistics and the Sun SNDR software screens.
T	Display the cumulative key cache statistics generated since the last time the counters were reset (the -z option resets the counters).

During synchronization operations, the Sun SNDR software screen displays:

- Percentage of the synchronization completed for each volume
- Arrows indicating the direction of the synchronization
- Status of each volume

Syntax

/usr/opt/SUNWscm/sbin/scmadm -S [-M] [–d *delay-time*] [–l *logfile*] [–r [*range*]] [–z]

TABLE 3-6 describes the scmadm -S options.

TABLE 3-6 `scmadm -S` Options

Option	Description
`-M`	Displays the Sun SNDR software statistics. If `scmadm -S` is invoked without the `-M` option, it displays statistics related to the storage device cache.
`-d` *delay-time*	Sets the display update time to *delay-time* seconds.
`-l` *logfile*	Writes all screen outputs to the specified *logfile*.
`-r` [*range]*	Specifies the cache descriptor for one volume or a combination of a single volume, an inclusive range of volumes, and multiple volumes. The default is that all volumes are displayed. Specify the range in this format: *n*[`:`*n*`,`*n*] Where *n* is the numbers of the specified decimal volumes. A colon (`:`) is a separator specifying an inclusive range of volumes. A comma (`,`) is a separator specifying another volume. The following two examples specify the same volumes (3, 6, 7, 8, 9, 10, 11, 12, 14, and 15): **`-r 3,6,7, 8,9:12,14,15`** **`-r 3,6:12,14,15`**
`-z`	Clears the statistics before displaying current statistics.

Example `scmadm -S` Screen Display

The statistics display includes the fields described in TABLE 3-7. Type `scmadm -S -M` to display the Sun SNDR software-specific statistics; press the M key to toggle between the Sun SNDR software screens and regular cache statistics after typing the command.

TABLE 3-7 Fields for Sun SNDR Software Screen Display

Field	Description
`primary`	Specifies the primary partition.
`link status`	Indicates whether replication is enabled or disabled, or indicates the synchronization direction if a synchronization is in progress.
`secondary`	Specifies the secondary partition.
`dual copy status`	Indicates current status using the following keywords:
	`logging` - Data is not being replicated and changes are being logged to the scoreboard.
	`need rev sync` - Logging, but an earlier reverse synchronization did not finish executing. The primary volume might be in an inconsistent state. The `need rev sync` state is cleared following a successful reverse synchronization, or by an Instant Image software update to the primary volume (which restores an earlier valid snapshot and removes the need to repeat the reverse synchronization to restore consistency).
	`need sync` - Logging, but an earlier synchronization did not finish executing. The secondary volume might be in an inconsistent state. The `need sync` state is cleared following a successful synchronization, or by an Instant Image software update to the secondary volume (which restores an earlier valid snapshot and removes the need to repeat the synchronization to restore consistency).
	`replicating` - Data written to the primary is being replicated to the secondary. Formerly called `enabled`.
	`rev sync` - A synchronization from secondary to primary is under way.
	`sync` - A synchronization from primary to secondary is under way.
	`volume down` - The volume has failed.
`recovery needed`	Displays the percentage of segments of the partition that require recovery.
`recovery completed`	Displays the percentage of the partition that has been resynchronized.

See CODE EXAMPLE 3-1 for a sample output.

- The equals sign (=) in the link status field indicates that the link is up and replication is enabled
- The asterisks (*) in the recovery fields are graphical representations of the percentage of recovery needed and completed
- The arrows (>) indicate the synchronization direction

CODE EXAMPLE 3-1 `scmadm -S` Link Status Display

```
-------------------------------------------------------------------
SAMPLE 109056   ****** Dual Copy Statistics ******    17:40:10

primary               link status     secondary            dual copy status

ma:...dsk/c0t117d0s3      *>*        atmsi:...rdsk/c1t6d0s3   sync
ma:...dsk/c0t115d0s3      *>*        atmsi:...dsk/c1t9d0s3    sync
 ma:...dsk/c0t116d0s3     *>*        atmsi:...dsk/c1t6d0s3    sync
 ma:...dsk/c0t96d0s3      *=*        atmsi:...rdsk/c1t0d0s3   replicating
-------------------------------------------------------------------
partition         recovery needed                recovery completed

/dev/rdsk/c0t117d0s3 [*****             ]  29.53%  [**************     ] 70.47%
/dev/rdsk/c0t115d0s3 [*****             ]  29.16%  [**************     ] 70.84%
/dev/rdsk/c0t116d0s3 [*****             ]  29.38%  [**************     ] 70.62%
-------------------------------------------------------------------
```

Example scmadm -S Screen Display

The scmadm -S statistics include the fields described in TABLE 3-8. Use the M key to toggle between these statistics and the Sun SNDR software statistics. Use the T key to display the cumulative cache statistics generated since the last time the counters were reset. Type scmadm -S -z to reset the counters.

The sample displays in CODE EXAMPLE 3-2 and CODE EXAMPLE 3-3 contain the cache statistics and cumulative cache statistics displays.

TABLE 3-8 scmadm -S Fields

Field	Description
cd	Specifies the cache descriptor number.
cached-partition	Specifies the disk partition being monitored.
disk-io	Displays the number of Kbytes per second read from or written to physical disks.
cache	Displays the number of Kbytes per second read from or written to data cache.
write-blocks	dirty—the number of dirty cache blocks that have not been queued for destaging. todisk—the number of blocks that have been queued for destaging and are waiting to be written to disks. The blocks displayed here are eventually moved to disk-io (writes).
accesses/s	Displays the number of I/O operations (reads per second + writes per second) serviced by the cache per second. The accesses/s total is the sum of the hits and misses.
read/s	Displays the number of read requests (hits) serviced per second. misses/s—Displays the number of read misses per second.
write/s	Displays the number of write requests serviced per second. misses/s—Displays the number of write misses per second.
%readh	Displays the percent of reads that are cache hits.
%writeh	Displays the percent of writes that are cache hits.
cachesize	Displays the cache size.
blocksize	Displays the blk-size.
Write blocks available	Displays the available write cache in blocks. Net *n* in the field represents each network.

TABLE 3-8 `scmadm -S` Fields *(Continued)*

Field	Description
`LRU stats`	Displays the least recently used (LRU) algorithm statistics: `Blocks`—the number of cache blocks in the LRU (always the total number of blocks in the cache). This figure stays constant. Requeued—unused. `Optimized`—cache blocks not requeued to the tail of the LRU.
`Total Cache Memory Usage`	Displays the size of system memory that is used by the cache when cache starts.
`Total Stats Memory Usage`	Displays the size of system memory that keeps the cache statistics data.

CODE EXAMPLE 3-2 Sample `scmadm -S` Display Cache Statistics

```
SAMPLE 9            ******      Storage Cache     ******            14:43:44

                          disk_io       cache          write_blocks
  cd cached_partition  reads writes  reads writesdirty todisk failed
   0 ...k/c0t115d0s3    2559      0      0      0      0      0      0
   1 ...k/c0t116d0s3    3071      0      0      0      0      0      0
   2 ...k/c0t117d0s3    2559      0      0      0      0      0      0
   3 ...sk/c0t98d0s3       0      0      0      0      0      0      0
   4 ...sk/c0t98d0s1       0      0      0      0      0      0      0
   5 ...sk/c0t98d0s5       0      0      0      0      0      0      0
   6 ...sk/c0t96d0s3       0      0      0      0      0      0      0
                      ------ ------ ------ ------
       Kbytes/s total:  8189      0      0      0

 accesses/s    read/s    write/s   %readh   %writeh
             (misses/s) (misses/s)
     16.00       0.00       0.00      0.0       0.0
             (  16.00 ) (   0.00 )
```

CODE EXAMPLE 3-3 Sample `scmadm -S` Display Cumulative Cache Statistics

```
SAMPLE 10     ******     Storage Cache (Cumulative)      ******      14:43:44

                              disk_io                 cache
  cd cached_partition      reads     writes      reads     writes
   0 ...k/c0t115d0s3     2034688          0          0          0
   1 ...k/c0t116d0s3     2031104          0          0          0
```

```
  2 ...k/c0t117d0s3     2028032          0          0          0
  3 ...sk/c0t98d0s3           9         10          4          0
  4 ...sk/c0t98d0s1           9         10          4          0
  5 ...sk/c0t98d0s5           9         10          4          0
  6 ...sk/c0t96d0s3           0          0          0          0
                     ---------- ---------- ---------- ----------
        Kbytes total:   6093851         30         12          0
 accesses        read        write    %readh  %writeh
          (    misses) (    misses)
    11950          27            0       0.2      0.0
          (     11908) (        15)

cachesize  blocksize
   65536K       8192

Write blocks available:
Net 0:   4096   Net 1:      0    Net 2:      0    Net 3:      0

LRU stats:  Blocks     Requeued    Optimized
              8192            0            0

Total Cache Memory Usage: 2320 Kbytes
```

CHAPTER 4

Using The Sun SNDR Software to Restore Volumes and Sites

You must enable the Sun SNDR software at the primary and secondary sites before attempting the tasks in this chapter. See Chapter 3 for information on the command syntax.

The topics in this chapter include:

Preparing to Use the Sun SNDR Software

This section describes several issues you need to consider before you start using the Sun SNDR software. These topics are as follows:

Definition of Primary and Secondary Sites and Network Link Type

In the examples in this chapter, Site A, `host1` is the local primary host site and machine name, and Site B, `host2` is the remote secondary host site and machine name, respectively. Depending on the example, either site can be the primary or secondary of the remote copy operation. As shown in TABLE 3-4, you must perform all synchronization operations from the primary host session.

Note – The diagrams in this section show the network link type as ATM. Although the Sun SNDR software is most likely to be used with SunATM link-level interfaces, the Sun SNDR software can be used with any Sun-supported link-level interface that is TCP/IP-capable, such as Gigabit Ethernet, Gigabit Ethernet Fibre Channel, and others.

Selecting Volumes

The first planning step is to determine which volumes to include in the remote replication operations. The secondary volume must be as large as the primary volume.

It is important to balance remote accessibility and recoverability against capacity usage and I/O response time. Generally, include the following critical volumes in the remote replication configuration:

- Databases and database management system (DBMS) logs
- Access control files

You can enable these volumes for remote copy individually as part of a volume set or collectively as part of an I/O group as described in Chapter 3.

You can exclude volumes from the Sun SNDR software configuration if they can be reconstructed at the recovery site or if they seldom change:

- Temporary volumes (such as those used in sort operations)
- Spool files
- Paging volumes

Using the Storage Volume Interface

When you enable volume sets using the `/usr/opt/SUNWesm/sbin/sndradm` command, you place the specified volumes under control of the Storage Volume (SV) driver. See “Enabling the Sun SNDR Software” on page 28.

You can also manually place volumes under SV control through a command-line interface named `svadm`. For example, you can add the volumes as follows:

```
# /usr/opt/SUNWesm/sbin/svadm -e -f config-file
```

where *config-file* is an ASCII file listing the volume (the full path to the raw device node). You can also add individual volumes from the command line.

See the `svadm` man page; it also describes the format of a configuration file.

Monitoring Sun SNDR Software Operations

Use the `/usr/opt/SUNWscm/sbin/scmadm -S` command described in the section “Getting Sun SNDR Software and Storage Cache Statistics: The scmadm Command” on page 51 to monitor the progress of the Sun SNDR software operations.

Mutual Backup Considerations

At times the distinction between primary and secondary sites can become blurred. As applications are geographically distributed, a storage system at Site B can function as a remote volume backup to Site A and as a direct storage resource for applications on Host B. Under these circumstances, keep replicated volume copies of Host B volumes at Site A. This reciprocal backup arrangement supported by the Sun SNDR software is known as *mutual backup*.

With mutual backup, the Sun SNDR software volumes considered primary by Site B are administered from the Site B session. Site B replicated-volume devices are considered secondary volumes at Site A.

Using the Sun SNDR Software With Sun StorEdge Instant Image Software

The Sun StorEdge Instant Image software is an optional software component. You can combine the Sun SNDR software with volumes from Instant Image software to create multiple copies of each device. See “Adding and Deleting Instant Image Software Volumes” on page 32.

The Instant Image software is a point-in-time volume copy data service for the Solaris operating environment. With Instant Image software, you create a volume pair: a point-in-time logical volume copy (shadow) from the original logical volume (master), which you specify. Once the shadow is established, you can read from and write to this shadow volume and the master volume.

The Instant Image software is best used just before you perform a resynchronization, to ensure that a consistent copy of data exists. In this case, if the resynchronization fails, you at least have a copy of known good data that is usable.

During the resynchronization process of updating the local and remote sites, the data on a secondary Sun SNDR volume is temporarily inconsistent with the primary volume. The secondary volume cannot be relied on for data recovery. Consistency is restored when the resynchronization is complete. To ensure data integrity, use Instant Image software regularly to create a point-in-time copy of data at both sites. See the Instant Image documentation listed in “Related Documentation” on page xviii.

Full Forward Synchronization: Establishing Replicated Images for the First Time

Before using data on the Sun SNDR software-replicated volumes, ensure that the contents of the local and remote volume sets match. All Sun SNDR software synchronization and update operations in this case start from the primary system, because you issue the commands from the primary system.

Existing Primary and Newly Formatted Secondary Volumes

If the primary volume has already been updated locally (without the Sun SNDR software replication) and contains important information, you can copy its image to the newly formatted secondary volume. The Sun SNDR software performs this initial primary-volume-to-secondary-volume copy while also forwarding new online updates to the secondary volume.

▼ To Fully Synchronize Specific Remote Volumes While Applying New Updates

- **Use this command**:

```
host1# sndradm -m [-g io-groupname] [-C tag] [-n] [-f config-file | SNDR-set |
set-name]
```

FIGURE 4-1 shows the full resynchronization process using an ATM link. This process also enables simultaneous write updates from the local to remote volumes.

Note – The illustrations in this section show the network link type as ATM. Although the Sun SNDR software is most likely to be used with SunATM link-level interfaces, the Sun SNDR software can be used with any Sun-supported link-level interface that is TCP/IP-capable., such as Gigabit Ethernet, Gigabit Ethernet Fibre Channel, and others.

Data Flow

FIGURE 4-1 shows the full resynchronization process using an ATM link.

1. The Sun SNDR software on the primary system (`host1`) requests disk blocks from the active primary volume. The data might already be resident in the primary system data cache, or might require a local disk access.
2. The Sun SNDR software transmits the disk blocks, with destaging instructions, over the ATM connection to a cache region on the secondary system.
3. The Sun SNDR software on the secondary system updates its remote volume and acknowledges the update to the primary system.

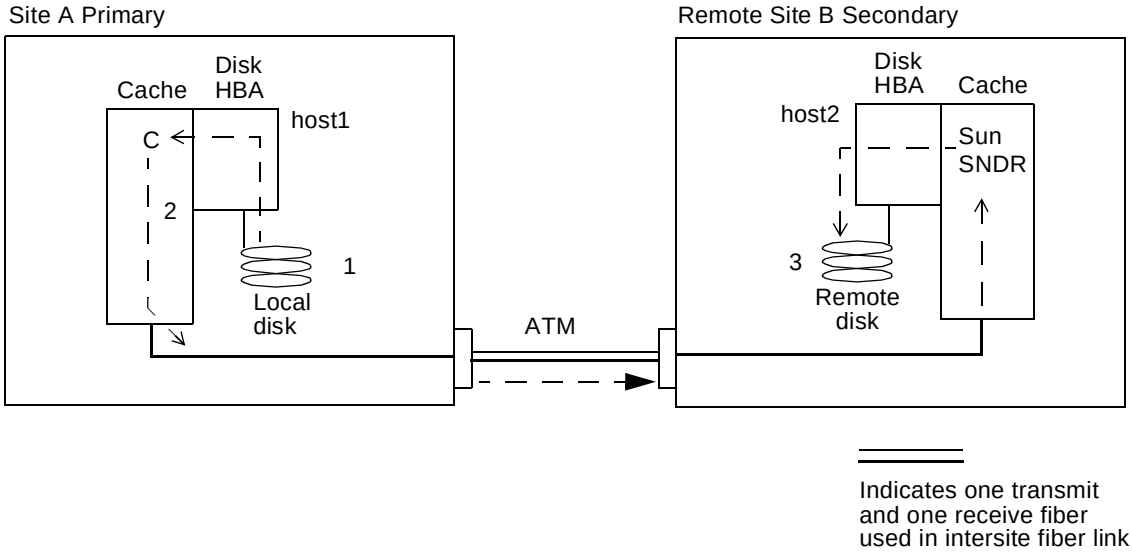

FIGURE 4-1 Full Synchronization (Volume-to-Volume Copy) Using ATM

Handling Interruptions

The Sun SNDR software uses a periodic intersite signal to monitor the health of primary and secondary systems. The absence of a health monitor signal indicates an interruption in the Sun SNDR software service. A failure or impairment in the intersite link or an outage at the remote site can cause interruptions. You can introduce interruptions intentionally to exercise remote failure strategies—for example, during the disaster recovery rehearsals described in “Rehearsing Disaster Recovery” on page 74.

Primary Logging During Interruptions

During interruptions, the Sun SNDR software tracks primary volume areas that are being updated locally but have not yet been copied to the secondary site. Under several circumstances, when the Sun SNDR software service is restored, update resynchronization can be requested automatically or manually, using the scoreboard logging information to update the remote site.

Forward Resynchronization Update: Resynchronizing Volumes After an Interruption

Caution – Ensure that you understand the resynchronization process of updating both sites before invoking it. While resynchronization is occurring, the secondary volumes are temporarily inconsistent and cannot be relied on for recovery. Consistency is restored when the resynchronization finishes executing. To ensure data integrity, use the Instant Image software regularly to create a point-in-time copy of data at both sites.

Typically, interruptions in the Sun SNDR software services are infrequent. The type of interruption (an isolated incident or part of a larger disaster) determines the action you should take.

If the Sun SNDR software interruption is a symptom of a larger rolling disaster, avoid resynchronization of the sites. Maintain the secondary site in a dated but consistent state, rather than risk a disastrous interruption that leaves the secondary site inconsistent and difficult to recover from.

▼ To Enable Update Synchronization

- **When you determine that the Sun SNDR software interruption is an isolated incident and the condition has been repaired, enable update synchronization using the following Sun SNDR software command:**

```
host1# sndradm -u [-g io-groupname] [-C tag] [-n] [-f config-file | SNDR-set |
set-name]
```

Note – If the secondary volume state is unknown because of system or disk failure, you might need to make full volume copies to reestablish matching Sun SNDR software volume sets. In this case, use the `sndradm -m` command to fully update the secondary volume set.

Data Flow

FIGURE 4-2 shows an update resynchronization using an ATM link from the primary system to its secondary system, when the secondary volumes are stale from the interruption.

1. The Sun SNDR software on `host1` examines a scoreboard from the primary and secondary hosts for the Sun SNDR software-managed volumes affected by the interruption.
2. The Sun SNDR software on `host1` requests the blocks that were updated during the interruption from the up-to-date volume. The data might already reside in the `host1` data cache or on the local disk.
3. The Sun SNDR software on `host1` transmits the update blocks 3R to `host2` Sun SNDR software using the SunATM connection.
4. The Sun SNDR software on `host2` refreshes its stale replicated image with the updated blocks and acknowledges the action to `host1`.
5. The Sun SNDR software revises the scoreboard to track the remote update.

All steps repeat until the remote replicated image is up-to-date. Use `scmadm -S -M` to monitor the resynchronization progress.

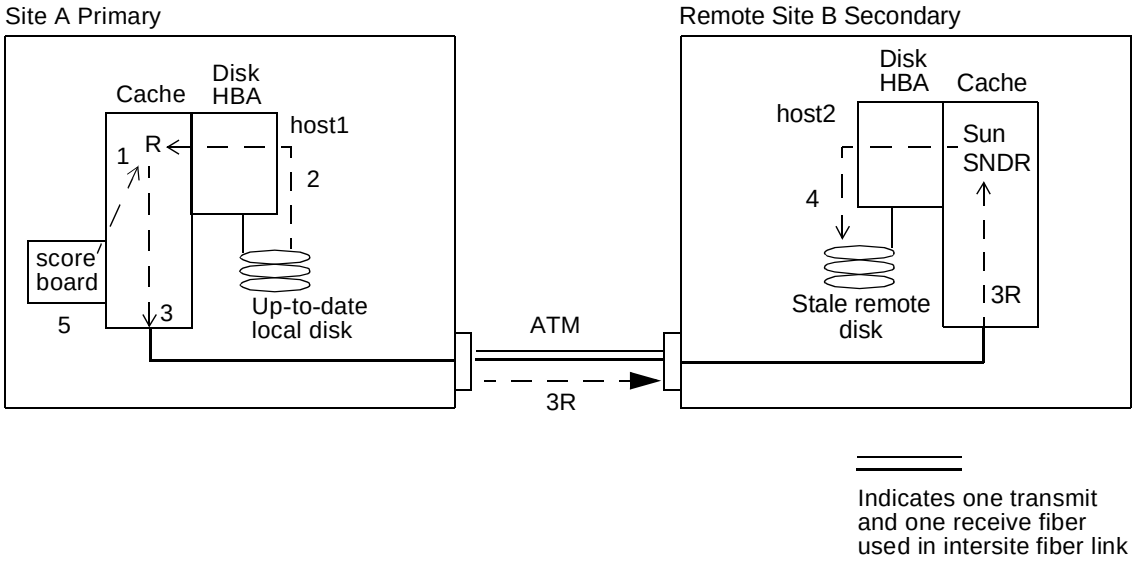

FIGURE 4-2 Update Synchronization of a Secondary Volume Set

Autosynchronization After an Interruption

Note – See the command descriptions for “Toggling the Sun SNDR Software Autosynchronization State” on page 49 and “Setting the Asynchronous Queue” on page 50.

The Sun SNDR software synchronization daemon `/usr/opt/SUNWrdc/lib/sndrsyncd` starts when the Sun SNDR software starts (by default, autosynchronization is disabled. You enable it on a volume set basis using the `sndradm -a` command). The daemon monitors the existing Sun SNDR software configuration and attempts to resynchronize the specified volumes if the secondary system reboots or a link failure occurs. You can initiate automatic resynchronizations only on the primary system.

If the secondary host is enabled, the update synchronization operations are started on the primary system to update all the specified secondary hosts.

Autosynchronization With Instant Image Software Configured

Note – See the command description for “Adding and Deleting Instant Image Software Volumes” on page 32.

The `/usr/opt/SUNWrdc/lib/sndrsyncd` daemon automates update resynchronization after a network link or machine failure; if the Sun StorEdge Instant Image software is also installed and you have added Instant Image software volume groups to the data service, it invokes point-in-time copies when necessary to protect the data volumes being updated during a resynchronization.

When a network link being used by the Sun SNDR software becomes unavailable, the daemon attempts to invoke the Sun SNDR software update commands to resynchronize all volume sets that have autosynchronization enabled and are using the network link.

The daemon is also notified when any Sun SNDR software resynchronization starts or finishes executing. The daemon performs Instant Image software point-in-time copy operations on the secondary or target host, if configured.

On a secondary host, the daemon checks whether a file system is currently mounted on the secondary volume and informs the kernel not to allow the synchronization to start if the file system is currently mounted.

Use the `sndradm -I` command to create data service configuration entries marked with the `ndr_ii` key. The `ndr_ii` entries contain an additional state field that the kernel uses to determine when point-in-time (PIT) copies must be made. The kernel notifies the Sun SNDR software synchronization daemon on the target system whenever a sync is started and waits for `sndrsyncd` to perform any necessary PIT copies before allowing the synchronization to proceed.

Restoring Volumes After a Secondary Site Failure

The effort required to restore a secondary site to its remote copy state depends on the hardware and software that is replaced or repaired. This section describes restoring disk, system, and intersite links. If a volume is not recoverable after a disk failure, follow the steps described in this section.

Restoring a Failed Secondary Volume

FIGURE 4-3 shows restoration of a failed secondary disk. The process follows the same steps used to establish an initial Sun SNDR software volume set; see “Full Forward Synchronization: Establishing Replicated Images for the First Time” on page 64. The concurrent write updates have been omitted from the figure.

▼ To Restore a Volume Set at a Remote Site

- **Use this command:**

```
host1# sndradm -m [-g io-groupname] [-C tag] [-n] [-f config-file |SNDR-set |
set-name]
```

Data Flow

FIGURE 4-3 shows restoration of a failed secondary disk.

1. The Sun SNDR software on `host1` requests disk blocks from the active primary volume. The data might already be resident in `host1` data cache or might require a local disk access.
2. The Sun SNDR software transmits the data blocks 2R, with destaging instructions, over the ATM connection to a Sun SNDR software region on remote `host2`.
3. The Sun SNDR software on `host2` updates its remote volume and acknowledges the action to `host1`.

All steps repeat until the entire volume is copied. Use `scmadm -S -M` to monitor the restoration process.

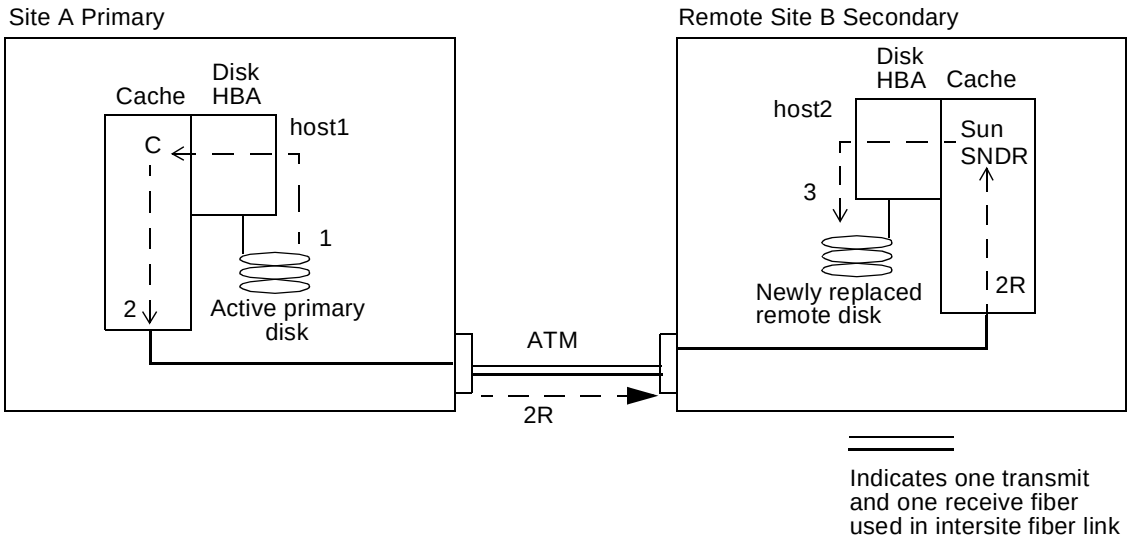

FIGURE 4-3 Restoring a Failed Secondary Volume

Restoring a Failed Secondary Server

Restoring a failed secondary server might involve both updating and fully resynchronizing the secondary disks, depending on the severity and duration of the failure. The update operation updates any secondary volumes whose contents were maintained during the failure. The full synchronization operation writes a complete copy of the primary volumes to volumes residing on secondary disks that were replaced or whose state is unknown.

▼ To Perform an Update Operation

- **Use this command:**

```
host1# sndradm -u [-g io-groupname] [-C tag] [-n] [-f config-file |SNDR-set
| set-name]
```

▼ To Perform a Full Synchronization

- **Use this command:**

```
host1# sndradm -m [-g io-groupname] [-C tag] [-n] [-f config-file |SNDR-set |
set-name]
```

Switching to an Alternate Fiber (ATM Only)

Most intersite link impairments are temporary but a fiber cable can be accidentally cut. Fiber plant suppliers support separately routed links between the two sites for added protection. They might use an automatic switch or might need to manually switch fiber cable pairs at a patch panel to redirect the signal to the alternate route.

Rehearsing Disaster Recovery

The Sun SNDR software enables you to perform disaster rehearsals easily, which encourages frequent verification of disaster preparedness.

Validating contingency procedures through rehearsals is an important aspect of any disaster recovery plan. Perform rehearsals regularly and refine them whenever a significant change is made to the primary or secondary processing environments.

Rehearsing Recovery From Secondary Volumes

You can simulate a disaster by temporarily interrupting the Sun SNDR software network connection and practicing takeover on the secondary site. To do this with the Sun SNDR software, break the network connection between the primary and secondary sites.

Complete these steps before proceeding with recovery:

- Use `scmadm -S -M` to verify that all Sun SNDR software secondary volume updates have been destaged to disk. Destaging might be in progress if the Sun SNDR software replication activity was high prior to the rehearsal.
- Declare the secondary volumes available for access from the recovery hosts.
- Start application-level recovery procedures to ensure a consistent starting point for future transactions.

During a read operation from the secondary recovery copy:

- The secondary remote site host issues a read request on `host2` against the recovery copy.
- As the read request returns a cache miss, the appropriate blocks are read from the replicated copy volume into `host2` data cache.
- The read operation results are returned to the secondary remote site host.

Subsequent host reads and writes populate the recovery system data cache and normal local access is restored.

Secondary Updates During a Takeover Rehearsal

During the rehearsal, you can apply test updates from a secondary host to the secondary volumes to evaluate the recovery mechanisms. If these test updates are not part of the permanent business record, they must be undone when the recovery copy is returned to its secondary Sun SNDR software role. See “To Perform a Secondary Rollback” on page 77” for the appropriate procedures.

Scoreboard logs track the test updates occurring on the secondary volumes during recovery rehearsals.

Data Flow

FIGURE 4-4 shows secondary write logging during rehearsals:, described as follows.

1. The remote site host issues a write request to its local volumes.
2. The write data D is placed into `host2` data cache.
3. The Sun SNDR software keeps a scoreboard for each Sun SNDR software-managed volume to indicate changes from the last known image synchronized with `host1`.

   These scoreboard logs might be used later to overwrite the changed segments with the most recent data from the primary volumes.
4. Software on `host2` destages the write data to the local volume.

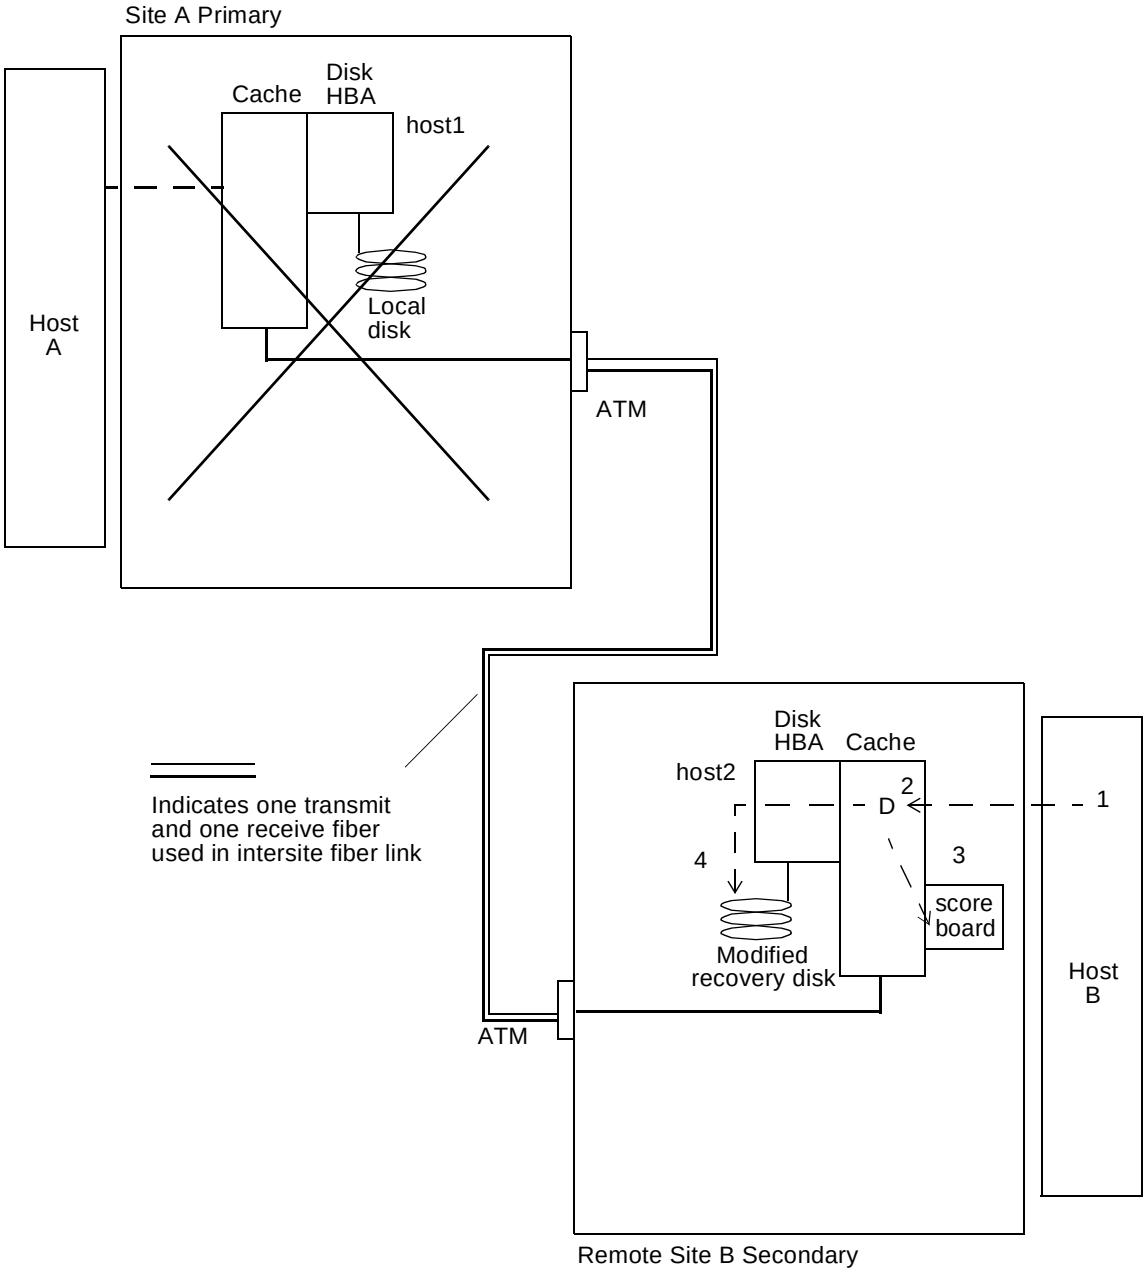

FIGURE 4-4 Secondary Scoreboard Logging

Rolling Back Updates After a Takeover Rehearsal

When the rehearsal completes, the Sun SNDR software uses the secondary scoreboard logs to update the appropriate blocks with the current information from the primary volumes; the update procedure rolls back the test changes.

A more complete recovery rehearsal requires that the rehearsal updates be captured and later reflected on the primary server before the workload is switched back. See “Reverse Synchronization: Updating the Primary Site From the Secondary Site” on page 83” for more detail.

▼ To Perform a Secondary Rollback

- **On the primary host, use this command:**

```
host1# sndradm -u [-g io-groupname] [-C tag] [-n] [-f config-file |SNDR-set |
set-name]
```

Primary Site Failures

This section contains information needed for handling primary site failures.

Failed Primary Disk

The Sun SNDR software provides continuous data access during primary volume failures. The Sun SNDR software high-availability features are a superset of RAID 1 and RAID 5 storage protection that can be optionally configured for the primary volumes. The Sun SNDR software remote volume access features start only after the disk protection schemes on the primary system are unable to provide data access to the local devices.

In the linear and striped (RAID 0) cases, failure of a single disk storing the primary volume triggers the Sun SNDR software to transparently redirect disk reads and writes to the remote storage system.

If the primary logical volume is locally mirrored (RAID 1) across two physical disks on the same system, a single disk failure results in its local mirror disk handling all requests for cache staging on a read miss and cache destaging. The Sun SNDR software relies on the remote site secondary devices only if both local mirrors fail.

If the primary volume is RAID 5 protected, its contents are striped across several physical disks. The local system considers the primary volume inaccessible and yields to the Sun SNDR software remote volume access only when two or more of the disks in the RAID 5 stripe fail.

Data Flow

FIGURE 4-5 shows the Sun SNDR software transparent pass-through of read requests to the remote system, bypassing the local disk failure. Writes continue to be handled as with a working local disk, though no destaging occurs on the primary system.

1. Host A issues a read request R for failed local disk on `host1`.
2. If the disk block is resident in `host1` cache, it is immediately returned. However, if the read results in a cache miss, the attempt to access the local disk fails and the Sun SNDR software forwards the read request to `host2`.
3. The Sun SNDR software on `host2` acts on the remote request by reading its remote mirror disk.
4. The Sun SNDR software on `host2` responds to `host1` with the requested disk blocks. The read data is used to refresh `host1` cache.
5. The read data is returned to Host A from `host1` cache.

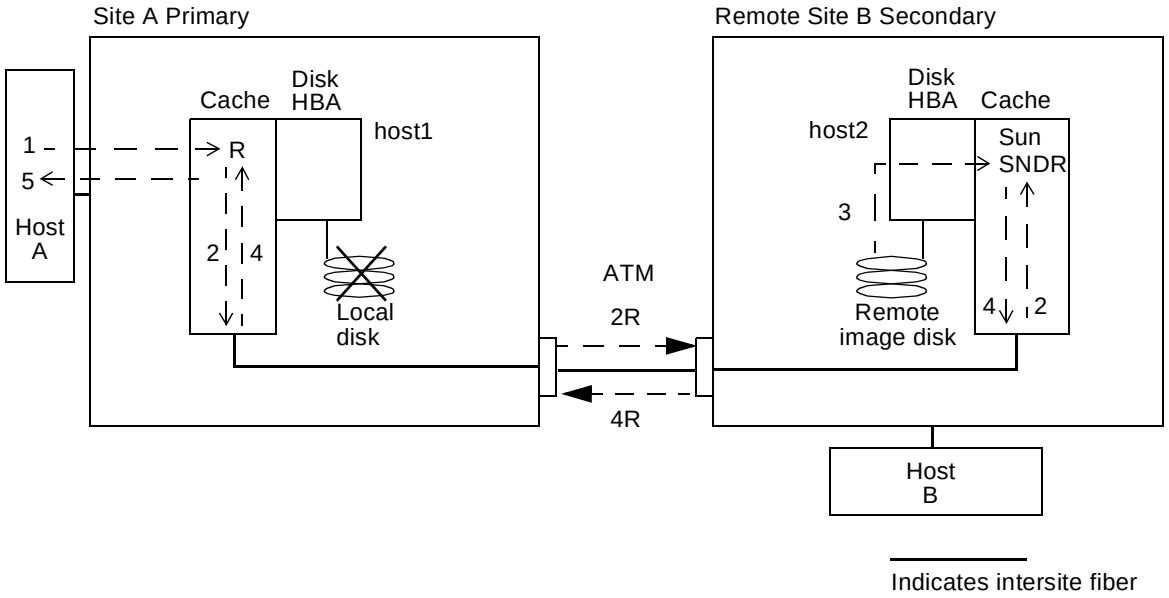

FIGURE 4-5 Pass-Through Reads to Bypass Local Disk Failure

▼ Example: To Restore a Failed Primary Volume

Caution – Disabling the Sun SNDR software on the failed device causes application access to the device to fail and disassociates the failed primary volume from its secondary copy.

1. **Keep the failed volume enabled under Sun SNDR software control.**

   The Sun SNDR software marks the device as failed when it is unable to read or write from it. The Sun SNDR software continues to provide read and write services to the host application using the secondary volume at the remote site. The remote volume can be used to satisfy I/O requests only while in replicating mode. As soon as logging mode is entered, the I/O fails.

   To determine if the problem is a hardware failure requiring part replacement or a transient error, an experienced system administrator should carefully examine the `/var/adm/messages` file.

2. **If it is a hardware problem, contact Sun Enterprise Services to arrange for replacement of the failed device and restoration of the associated logical volumes.**

3. **Quiesce the application using the failed primary volume.**

4. **If autosynchronization is off and you have installed Sun StorEdge Instant Image, perform any Instant Image operations necessary to preserve a valid point-in-time copy of the data on the secondary system before starting a Sun SNDR software synchronization to update the secondary.**

   You can configure Sun StorEdge Instant Image volume groups for use with the Sun SNDR software. See "Adding and Deleting Instant Image Software Volumes" on page 32 and "Using the Sun SNDR Software With Sun StorEdge Instant Image Software" on page 63.

5. **If autosynchronization is on, if the volume specified as the Sun StorEdge Instant Image shadow of the Sun SNDR software secondary volume is mounted, the Sun SNDR software secondary volume data will be in an unknown state when the Sun SNDR software starts a synchronization. If this state is not appropriate for a Sun StorEdge Instant Image snapshot, do not put the Sun SNDR software into autosynchronization mode.**

6. **After the disk failure is corrected and any related Volume Manager volumes are recreated, use this command to perform the following reverse synchronization while the primary affected volumes are unmounted:**

```
host1# sndradm -m -r [-g io-groupname] [-C tag] [-n] [-f config-file |SNDR-set |
set-name]
```

This command starts reverse full synchronization, as the secondary volume on `host2` is resynchronizing the new primary volume on `host1`. FIGURE 4-6 shows the full reverse synchronization process.

1. The data might already be resident in `host1` data cache, or it might require a secondary disk access. If so, the Sun SNDR software on `host1` requests blocks from the up-to-date secondary volume on `host2`.
2. The Sun SNDR software on `host2` transmits the cache blocks 2R over the intersite fiber link to a Sun SNDR software region on `host1` with destaging instructions.
3. The Sun SNDR software on `host1` updates its disk.

7. **Restart the application on the primary.**

All steps repeat until the entire volume is copied. Subsequent reads from Host A are serviced locally from the newly updated disk. Reverse synchronization occurs concurrently with pass-through reads.

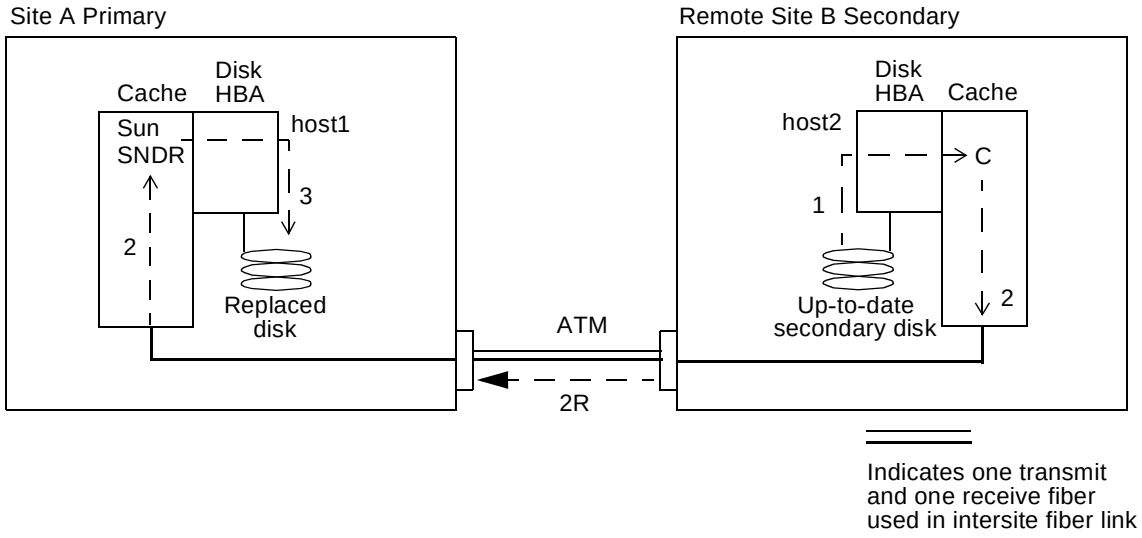

FIGURE 4-6 Reverse Full Synchronization

Recovering From a Primary Site Disaster

The Sun SNDR software minimizes the effects of a disaster at the primary site by enabling you to keep the secondary storage images updated. Although the secondary Sun SNDR software cache contains the latest writes issued on the primary before the disaster, that data might not have been destaged to the secondary disks yet. After detecting an interruption in the Sun SNDR software service, the Sun SNDR software automatically destages the secondary Sun SNDR software cache to its corresponding secondary volumes.

After all the secondary volumes have been updated with the latest Sun SNDR software cache images, the secondary volumes can be accessed by the secondary hosts. The `scmadm -S` screen displays confirm that destaging is complete. Run application-level recovery procedures to ensure a well-known state at the secondary site. The workload can then be switched to the secondary hosts for continued business operation.

Until the extent of the primary failure is understood, keep the Sun SNDR software enabled at the secondary site to track disk areas that are being modified. Under some confined disasters, the secondary update logs speed primary restoration, as described in “Full Reverse Synchronization: Completely Restoring a Primary Site From the Secondary Site” on page 85.

Reverse Synchronization: Updating the Primary Site From the Secondary Site

The primary volume state can be frozen as host workloads are switched to the secondary site. This is true in more realistic disaster rehearsals and might be true in confined disasters. Under these conditions, the secondary site activates its Sun SNDR software scoreboard logs to track changes not yet seen at the primary. When it is time to bring the primary back into service, the scoreboards can be used to refresh the primary contents from the current secondary images.

▼ To Start the Reverse Synchronization From the Secondary to the Primary Volumes

- **On the primary host, use this command:**

```
host1# sndradm -u -r [-g io-groupname] [-C tag] [-n] [-f config-file |SNDR-set
| set-name]
```

Data Flow

FIGURE 4-7 shows a reverse update resynchronization from the secondary system to its primary system.

1. The Sun SNDR software on `host1` retrieves the secondary scoreboard 1R from `host2` for one of the Sun SNDR software-managed volumes affected by the interruption.
2. The Sun SNDR software on `host1` requests the blocks updated during the interruption from the up-to-date secondary volume of `host2`. The data might already be resident in `host2`'s data cache, or it might require secondary disk access.
3. The Sun SNDR software on `host2` transmits the updated blocks 3R to `host1` Sun SNDR software region of cache using the intersite link.
4. The Sun SNDR software on `host1` refreshes its stale image with the updated blocks.
5. The Sun SNDR software on `host1` revises the scoreboard to track the remote update.

All steps repeat until the primary volume is up-to-date.

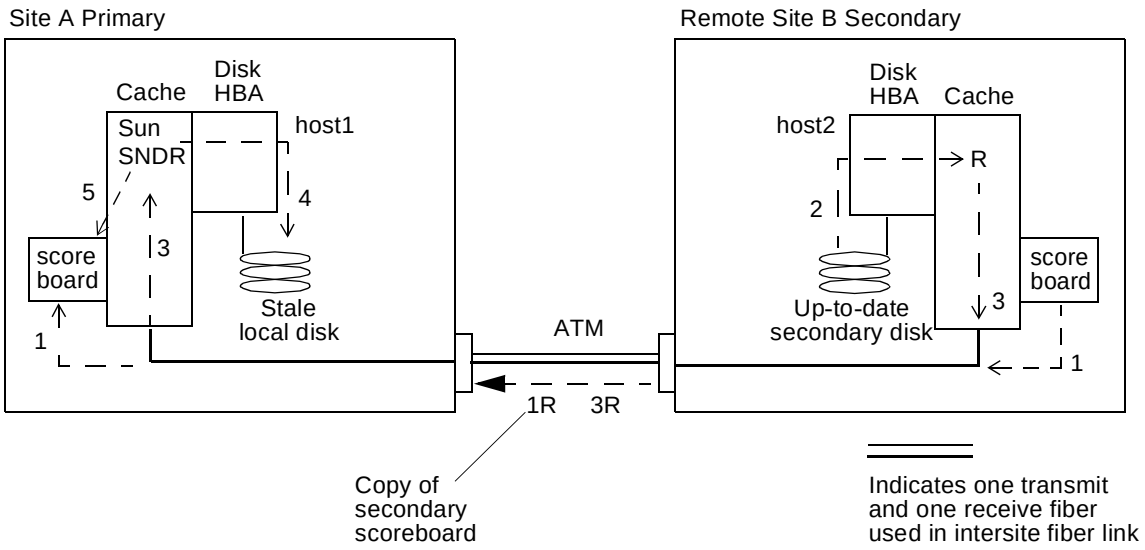

FIGURE 4-7 Reverse Update Synchronization

Full Reverse Synchronization: Completely Restoring a Primary Site From the Secondary Site

If the primary server becomes inoperative and primary data on the primary disks is lost, update logs at the secondary systems have little value. You must perform a reverse full synchronization on the repaired or replaced primary host. In other words, volume-to-volume copies from the secondary to the primary are required for all Sun SNDR software-managed volumes. This reverse synchronization process ensures that only the latest data is deposited on the primary disks.

▼ To Perform a Full Reverse Synchronization

- **On the primary host, use this command:**

```
host1# sndradm -m -r [-g io-groupname] [-C tag] [-n] [-f config-file |SNDR-set
| set-name]
```

FIGURE 4-8 shows this restoration process for the primary storage platform. The procedure is similar to that of restoring a single primary disk (see “Example: To Restore a Failed Primary Volume” on page 80), but a larger number of devices is included in the reverse resynchronization request.

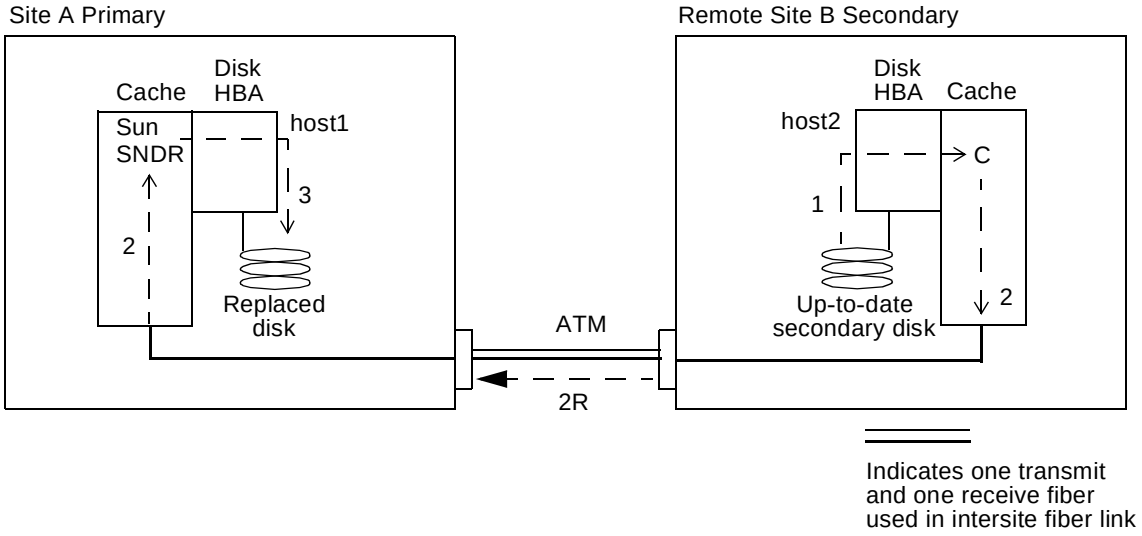

FIGURE 4-8 Restoring the Primary Site from the Secondary Site

Disabling Remote Replication

Caution – Disable remote replication *only* when it is not necessary that the primary and secondary volumes be associated with each other any longer.

Disabling the Sun SNDR software breaks the connection between primary and secondary volumes, discards any scoreboards, and removes the host and volume information from the data services configuration. Following the Sun SNDR software disable, enable and full synchronization (full volume copy) operations will be necessary to re-establish the Sun SNDR software relationship and ensure that the volumes' contents match. See "Enabling the Sun SNDR Software" on page 28 and "Full Forward Synchronization: Establishing Replicated Images for the First Time" on page 64.

▼ To Permanently Terminate the Sun SNDR Software Remote Replication and Resynchronization Services

- **From both hosts, use this command:**

```
host1# sndradm -d [-g io-groupname] [-C tag] [-n] [-f config-file | SNDR-set |
set-name]
```

Glossary

asynchronous replication Asynchronous replication confirms to the originating host that the primary I/O transaction is complete before updating the remote image. Asynchronous replication is used when the distance between primary and secondary sites might introduce prohibitive latency times to the synchronous operations.

data service One of the Sun StorEdge software components: Sun StorEdge Network Data Replicator, Fast Write Cache, Instant Image, or Target Emulation.

disk A host view of the logical storage device. Also referred to as volumes. Several host volumes can be physically stored in a single head disk assembly (HDA). Alternatively, a single host volume can be spread (striped) across multiple HDAs. The Sun SNDR software options are performed on discrete logical volumes rather than on physical HDAs.

full synchronization A complete volume-to-volume copy, which is the most time-consuming of the synchronization operations. In most cases, a secondary volume is synchronized from its source primary volume. However, restoration of a failed primary volume might require reverse synchronization, using the surviving remote mirror as the source.

health monitor Signal sent by each system to its remote set confirming that it is operational. The absence of a health monitor signal is the first indication that the intersite link or remote system is impaired.

HDA Head disk assembly.

I/O group A collection of Sun SNDR software volume sets that have the same group name, primary and secondary hosts, and mirroring mode.

local mirrors A mirrored partition or RAID 1 volume, which are established in a single system and do not rely on the Sun SNDR software or any hardware to maintain local synchronization. See *mirrored partitions*.

logging A method for tracking volume updates that have not been remotely copied while the remote service is interrupted or impaired. The blocks that no longer match their remote sets are identified for each volume. The Sun SNDR software uses these logs (scoreboards) to re-establish matching sets through an optimized update synchronization rather than through a complete volume-to-volume copy.

mirrored partitions Multiple copies of disk partitions that provide RAID 1 protection against possible data loss from disk failure. Data can be read back from any of the members in the set; in the event of a failure, a read is attempted on each set member until a successful read is accomplished.

mirroring The process of maintaining two or more identical images of a designated disk volume.

mutual backup A process where each server can concurrently transmit and receive copies to and from its remote counterpart. Each system contains primary disks in a Sun SNDR software volume set that are accessible by local hosts, as well as remote mirrors secondary to remote hosts. Mutual backup might be used where critical applications and storage are split across sites, and both sites require remote redundant copies.

partitions UNIX device names of the form `/dev/rdsk/c`*n*`t`*n*`d`*n*`s`*n* associated with the host disks on a Solaris system.

primary The system or volume being accessed directly by the host application. System or volume on which the host application is principally dependent.

quiesce To stop operations of a database application momentarily so that Sun SNDR software update and copy operations will not encounter data errors. After the update or copy command is issued, the database application can be restarted.

RAID Redundant Arrays of Independent Disks.

rollback synchronization A resynchronization operation that discards any blocks modified during recovery rehearsals. Scoreboard logging keeps track of test updates applied to the secondary system during the rehearsal. When the primary is restored, the test updates are overwritten with the blocks from the primary image, restoring matching remote sets.

rolling disaster A collection of damaging events spread over a period of hours that might impair multiple components in a Sun SNDR software configuration.

scoreboard log A special bitmap that tracks writes to a volume, rather than providing a running log of each I/O event. Scoreboard logging is a method for tracking volume updates that have not been remotely copied while the remote service is interrupted or impaired. The blocks that no longer match their remote sets are identified for each source volume. The Sun SNDR software uses this scoreboard log to re-establish a remote mirror through an optimized update synchronization rather than through a complete volume-to-volume copy.

secondary The remote counterpart of a primary system or volume where copies are destined Remote copies are transmitted without host intervention between peer servers. A server might act as primary storage for some volumes and secondary (remote) storage for others.

source device A device used as the originating device in the remote copy.

Sun SNDR software Sun StorEdge Network Data Replicator software.

Sun SNDR software log See *scoreboard log*.

synchronization The process of establishing an identical copy of a source volume onto a target volume as a precondition to Sun SNDR software mirroring.

synchronous mirroring The process of synchronous mirroring is limited to short distances (tens of kilometers) because of the detrimental effect of propagation delay on I/O response times.

takeover rehearsal Simulation of a switchover to the secondary system after a staged failure of the primary system.

target device Device receiving the remote copy.

update resynchronization The resynchronization operation that copies only the blocks modified since a Sun SNDR software interruption, as identified by scoreboard logging, reducing the time to restore remotely mirrored sets.

volume See *disk*.

volume set The specification that specifies the complete Sun SNDR software volume configuration information: primary and secondary hosts, volume partitions, scoreboard log bitmaps, and replication mode (asynchronous or synchronous).

Index

E

F

H

I

L

M

O

P

R

S

T

U

V

W