



Sun StorEdge™ Network Data Replicator 2.0 System Administrator's Guide

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

Part No. 806-4402-10
April 2000, Revision A

[Send comments about this document to: docfeedback@sun.com](mailto:docfeedback@sun.com)

Copyright 2000 Sun Microsystems, Inc., 901 San Antonio Road • Palo Alto, CA 94303-4900 USA. All rights reserved.

This product or document is protected by copyright and distributed under licenses restricting its use, copying, distribution, and decompilation. No part of this product or document may be reproduced in any form by any means without prior written authorization of Sun and its licensors, if any. Third-party software, including font technology, is copyrighted and licensed from Sun suppliers.

Parts of the product may be derived from Berkeley BSD systems, licensed from the University of California. UNIX is a registered trademark in the U.S. and other countries, exclusively licensed through X/Open Company, Ltd. For Netscape Communicator™, the following notice applies: Copyright 1995 Netscape Communications Corporation. All rights reserved.

Sun, Sun Microsystems, the Sun logo, AnswerBook2, docs.sun.com, Sun StorEdge, and Solaris are trademarks, registered trademarks, or service marks of Sun Microsystems, Inc. in the U.S. and other countries. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the U.S. and other countries. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc.

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

RESTRICTED RIGHTS: Use, duplication, or disclosure by the U.S. Government is subject to restrictions of FAR 52.227-14(g)(2)(6/87) and FAR 52.227-19(6/87), or DFAR 252.227-7015(b)(6/95) and DFAR 227.7202-3(a).

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

Copyright 2000 Sun Microsystems, Inc., 901 San Antonio Road • Palo Alto, CA 94303-4900 Etats-Unis. Tous droits réservés.

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

Des parties de ce produit pourront être dérivées des systèmes Berkeley BSD licenciés par l'Université de Californie. UNIX est une marque déposée aux Etats-Unis et dans d'autres pays et licenciée exclusivement par X/Open Company, Ltd. La notice suivante est applicable à Netscape Communicator™: Copyright 1995 Netscape Communications Corporation. Tous droits réservés.

Sun, Sun Microsystems, the Sun logo, AnswerBook2, docs.sun.com, Sun StorEdge, et Solaris sont des marques de fabrique ou des marques déposées, ou marques de service, de Sun Microsystems, Inc. aux Etats-Unis et dans d'autres pays. Toutes les marques SPARC sont utilisées sous licence et sont des marques de fabrique ou des marques déposées de SPARC International, Inc. aux Etats-Unis et dans d'autres pays. Les produits portant les marques SPARC sont basés sur une architecture développée par Sun Microsystems, Inc.

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

CETTE PUBLICATION EST FOURNIE "EN L'ETAT" ET AUCUNE GARANTIE, EXPRESSE OU IMPLICITE, N'EST ACCORDEE, Y COMPRIS DES GARANTIES CONCERNANT LA VALEUR MARCHANDE, L'APTITUDE DE LA PUBLICATION A REpondre A UNE UTILISATION PARTICULIERE, OU LE FAIT QU'ELLE NE SOIT PAS CONTREFAISANTE DE PRODUIT DE TIERS. CE DENI DE GARANTIE NE S'APPLIQUERAIT PAS, DANS LA MESURE OU IL SERAIT TENU JURIDIQUEMENT NUL ET NON AVENU.



Contents

Preface xi

What is Network Data Replicator? xi

How This Book Is Organized xi

Using UNIX Commands xii

Documentation Conventions xii

Shell Prompts xiii

Related Documentation xiii

Accessing Sun Documentation Online xiii

Ordering Sun Documentation xiv

Sun Welcomes Your Comments xiv

1. General Information About SNDR 1-1

Overview 1-2

SNDR Terminology 1-3

Volume Requirements 1-4

Performance 1-4

Bitmaps 1-5

What Happens When Sets Do Not Match? 1-5

2. SNDR Files and Commands 2-1

Contents	2-1
Configuration Files	2-2
/etc/opt/SUNWrdc/rdc.cf File	2-2
/etc/opt/SUNWrdc/rdc_ii.cf File	2-3
Summary of SNDR Commands	2-5
Copy Primary Volume to Secondary Volume	2-7
When To Use This Command	2-8
Stop Replication	2-8
When To Use This Command	2-9
Start Replication	2-9
When To Use This Command	2-10
Restore Primary Volume From Secondary	2-10
When To Use This Command	2-11
Refresh Primary Volume From Secondary Volume	2-12
When To Use This Command	2-13
Disable Replication	2-14
Description	2-14
When To Use This Command	2-14
Enable SNDR	2-15
Description	2-15
When To Use This Command	2-15
Toggle the SNDR Auto-Synchronization Value	2-16
Description	2-16
When To Use This Command	2-16
Display the SNDR Auto-Synchronization Value	2-16
Description	2-16
When To Use This Command	2-17
Display a List of Configured SNDR Sets	2-17

Description	2-17
When To Use This Command	2-17
Getting SNDR and Storage Cache Statistics: the <code>scmadm</code> Command	2-18
Example <code>scmadm -S</code> SNDR Screen Display	2-19
Example <code>scmadm -S</code> Screen Display	2-21
3. SNDR Operations	3-1
Before You Begin Using SNDR	3-2
Establishing Replicated Images for the First Time	3-4
Existing Primary and Newly-Formatted Secondary	3-5
Handling SNDR Interruptions	3-7
Primary Logging During SNDR Interruptions	3-7
Re-synchronizing Volumes After an SNDR Interruption	3-7
Auto-Resynchronization After an SNDR Interruption	3-10
Restoring Volumes After Secondary Site Failures	3-11
Restoring a Failed Secondary Disk	3-11
Restoring a Failed Secondary Server	3-13
Switching to an Alternate Fiber (ATM Only)	3-13
Disaster Recovery Rehearsals	3-14
Rehearsing Recovery From Secondary Volumes	3-14
Secondary Updates During a Takeover Rehearsal	3-15
Rolling Back Updates After a Takeover Rehearsal	3-17
Primary Site Failures	3-17
Failed Primary Disk	3-17
▼ To Restore a Failed Primary Volume (Example)	3-20
Recovering From a Primary Site Disaster	3-23
Updating the Primary Site From the Secondary Site	3-23
Completely Restoring a Primary Server from the Secondary	3-26

Temporarily Stopping Remote Replication 3-28

Disabling Remote Replication 3-29

Mirrored Partitions for Multiple Copies 3-29

Glossary GL-1

Index IN-1

Figures

- FIGURE 2-1 `rdcadm -m` Command Result 2-8
- FIGURE 2-2 `rdcadm -u` Command Result 2-10
- FIGURE 2-3 `rdcadm -m -r` Command Result 2-11
- FIGURE 2-4 `rdcadm -u -r` Command Result 2-12
- FIGURE 2-5 `scmadm -S` Link Status Display 2-21
- FIGURE 2-6 Sample `scmadm -S` Display -Cache Statistics 2-23
- FIGURE 2-7 Sample `scmadm -S` Display -Cumulative Cache Statistics 2-24
- FIGURE 3-1 Full Synchronization (Volume-to-Volume Copy) Using ATM 3-6
- FIGURE 3-2 Update Synchronization of a Secondary, Using an ATM Link 3-9
- FIGURE 3-3 Restoring a Failed Secondary Disk, Using an ATM Link 3-12
- FIGURE 3-4 Secondary Scoreboard Logging, Using an ATM Link 3-16
- FIGURE 3-5 Pass-Through Reads to Bypass Local Disk Failure, Using an ATM Link 3-19
- FIGURE 3-6 Reverse Full Synchronization, Using an ATM Link 3-22
- FIGURE 3-7 Reverse Update Synchronization, Using an ATM Link 3-25
- FIGURE 3-8 Restoring Primary from Secondary 3-27

Tables

TABLE P-1	Documentation Conventions	xii
TABLE P-2	Shell Prompts	xiii
TABLE P-3	Related Documentation	xiii
TABLE 2-1	<code>/etc/opt/SUNWrdc/rdc.cf</code> File Fields	2-3
TABLE 2-2	<code>/etc/opt/SUNWrdc/rdc_ii.cf</code> File Fields	2-4
TABLE 2-3	SNDR Command Options	2-5
TABLE 2-4	SNDR Command Parameter Summary	2-6
TABLE 2-5	<code>scmadm -S</code> Options	2-18
TABLE 2-6	Fields for SNDR Screen Display	2-19
TABLE 2-7	<code>scmadm -S</code> Fields	2-21

Preface

The *Sun StorEdge™ Network Data Replicator (SNDR) System Administrator's Guide* describes the SNDR software and its use.

What is Network Data Replicator?

SNDR software lets you replicate disks between different, physically-separate servers in real time. SNDR is intended to be used as part of Disaster Recovery and Business Continuity plans, which often rely on redundant storage of critical information across physically separate sites.

To transport data, SNDR uses any Sun™ network adaptor which supports TCP/IP.

How This Book Is Organized

Chapter 1 describes general information about SNDR and describes SNDR terminology.

Chapter 2 describes SNDR configuration files and commands.

Chapter 3 describes SNDR operations including synchronization, administrative utilities, 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

TABLE P-1 Documentation Conventions

Typeface or Symbol	Meaning	Examples
AaBbCc123	The names of commands, files, and directories; on-screen computer output.	Edit your <code>.login</code> file. Use <code>ls -a</code> to list all files. % You have mail.
AaBbCc123	What you type, when contrasted with on-screen computer output.	% su Password:
<i>AaBbCc123</i>	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 <i>User's Guide</i> . These are called <i>class</i> options. You <i>must</i> be root to do this. To delete a file, type <code>rm filename</code> .
[]	In syntax, brackets indicate that an argument is optional.	<code>sd_stats [-d sec] [-r n[:n][,n]...] [-z]</code>

Shell Prompts

TABLE P-2 Shell Prompts

Shell	Prompt
C shell	<i>machine_name%</i>
C shell superuser	<i>machine_name#</i>
Bourne shell and Korn shell	\$
Bourne shell and Korn shell superuser	#

Related Documentation

TABLE P-3 Related Documentation

Application	Title	Part Number
Release	<i>Sun StorEdge™ Network Data Replicator (SNDR) 2.0 Release Notes</i>	806-4404
	<i>Sun StorEdge™ Instant Image 2.0 Release Notes</i>	806-0231
Installation	<i>Sun StorEdge™ Network Data Replicator 2.0 Installation Guide</i>	806-4403
	<i>Sun StorEdge Instant Image™ 2.0 Installation Guide</i>	806-4004
System Administration	<i>SunATM 3.0 Installation and User's Guide</i>	805-0331
	<i>Sun StorEdge Instant Image™ 2.0 System Administrator's Guide</i>	806-0230

Accessing Sun Documentation Online

The `docs.sun.comsm` web site enables you to access Sun technical documentation on the Web. You can browse the `docs.sun.com` archive or search for a specific book title or subject at:

<http://docs.sun.com>

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://www1.fatbrain.com/documentation/sun>

Sun Welcomes Your Comments

We are interested in improving our documentation and welcome your comments and suggestions. You can email your comments to us at:

docfeedback@sun.com

Please include the part number (806-4402-10) of your document in the subject line of your email.

General Information About SNDR

This chapter describes SNDR and its capabilities:

- Overview
- SNDR Terminology
- Volume Requirements
- Performance
- Bitmaps
- What Happens When Sets Do Not Match?

Overview

SNDR features include:

- Synchronous and asynchronous volume replication between physically separate primary and secondary sites
- The SNDR link-neutral design can use any Sun network adaptor which supports TCP/IP
- Granularly configurable per logical volume (noncritical volumes may be excluded)
- Mutual backup (also known as a bilateral relationship)
- Continued remote data access despite local disk failure (depending on how SNDR is configured)
- Stop and start control of remote replication on a logical volume basis
- Active logging when SNDR is stopped or interrupted
- Optimized resynchronization following disk, link, system, and storage platform outages
- Disaster rehearsal rollback provisions
- Locally mirrored (RAID 1), RAID 5 protected, striped (RAID 0), and linear source and target volumes
- Integrated with Sun StorEdge Instant Image™ for additional point-in-time coordination and recovery

SNDR Terminology

This section describes important SNDR terms that are used throughout this manual. The Glossary contains more general terminology.

Term	Description
Asynchronous replication	<p>In contrast to synchronous mirroring, asynchronous replication confirms to the originating host that the primary I/O transaction is complete before updating the remote image. Deferring the secondary copy removes the long distance propagation delays from the I/O response time (there is uncertainty, as the replication operation is not confirmed to the originating application in real time).</p> <p>Asynchronous replication is used when the distance between primary and secondary sites may introduce prohibitive latency times to the synchronous operations.</p>
Full synchronization	<p>Complete disk-to-disk copy, which is the most time-consuming of the synchronization operations. In most cases, a secondary disk is synchronized from its source primary disk. However, restoration of a failed primary disk may require reverse synchronization, using the surviving remote mirror as the source.</p>
Mutual backup	<p>Each server can concurrently transmit and receive remote dual copies to and from its remote counterpart. Each system contains primary disks in an SNDR set that are accessible by local hosts, as well as remote mirrors secondary to remote hosts. Mutual backup may be used where critical applications and storage are split across sites, and both sites require remote redundant copies.</p>
Primary	<p>System or volume on which the host application is principally dependent.</p>
Rollback synchronization	<p>Operation used 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.</p>
Rolling disaster	<p>A collection of damaging events spread over a period of hours that may potentially impair multiple components in an SNDR configuration.</p>

Term	Description
Scoreboard logging	A special bitmap that tracks writes to a disk, rather than a running log of each I/O event. This is a method for tracking disk 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 disk. SNDR uses this scoreboard log to re-establish a remote mirror through an optimized update synchronization rather than a complete disk-to-disk copy.
Secondary	Remote counterpart of the primary where data copies are written to and read from. Remote copies are transmitted without host intervention between peer servers. A server may act as primary storage for some volumes and secondary (remote) storage for others.
Synchronization	The process of establishing an identical copy of a source disk onto a target disk as a precondition to SNDR mirroring.
Synchronous mirroring	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 server after a staged failure of the primary server.
Update synchronization	Copies only those disk blocks identified by scoreboard logging, reducing the time to restore remotely mirrored sets.

Volume Requirements

The primary and secondary volumes should be the same size. If you initiate an SNDR set where the secondary is larger than the primary, you will get a WARNING but will be allowed to continue. If you initiate an SNDR set where the secondary is smaller than the primary, `rdcadm` fails with an ERROR.

Performance

For optimum performance, use SNDR with Sun StorEdge™ Fast Write Cache and Instant Image.

Bitmaps

SNDR uses a bitmap to track writes. The bitmap is maintained on disk and may persist across a system crash, depending on the setting of `rdc_bitmap_mode` in `/usr/kernel/drv/rdc.conf`. For instance:

```
#
# 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;
```

SNDR can use regular files or raw devices to store bitmaps. If raw devices are used, they should be on a disk separate from the disk that has the data.

What Happens When 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 images to be unequal, or out of sync. SNDR uses several logging and resynchronization techniques to re-establish 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 SNDR.

Rollback resynchronization is a form of update synchronization used for disaster recovery rehearsals and other applications where test updates must be reversed. Chapter 3 describes the mechanism to use in a given circumstance. When multiple volumes require synchronization, SNDR performs several of these techniques concurrently.

SNDR Files and Commands

This chapter describes the SNDR configuration files and commands available from the command line. Chapter 3 describes the use of these commands for SNDR operations.

This chapter describes the most-commonly used commands first, followed by the least-used commands.

Note – Only the `rdcadmin -e`, `-d`, `-l`, and `-p` options can be issued from the primary or the secondary server. All other options must be issued from the primary server only.

Contents

- Configuration Files
- Summary of SNDR Commands
- Copy Primary Volume to Secondary Volume
- Stop Replication
- Start Replication
- Restore Primary Volume From Secondary
- Refresh Primary Volume From Secondary Volume
- Disable Replication
- Enable SNDR
- Toggle the SNDR Auto-Synchronization Value
- Display the SNDR Auto-Synchronization Value

- Display a List of Configured SNDR Sets
- Getting SNDR and Storage Cache Statistics: the `scmadm` Command

Configuration Files

Configuration files contain information for the dual copy sets. The SNDR commands rely on a default configuration file: `/etc/opt/SUNWrdc/rdc.cf`. You can create a customized configuration file depending on your server connection and disaster recovery plans.

`/etc/opt/SUNWrdc/rdc.cf` File

The default SNDR configuration file is `/etc/opt/SUNWrdc/rdc.cf`. This file defines the SNDR volume sets on which SNDR operates.

You can specify a configuration file other than the default `/etc/opt/SUNWrdc/rdc.cf` file with the `-f` option on the SNDR command line. For example, use a configuration file to specify multiple dual copy volume sets with a single SNDR command instead of entering a command for each set.

If you use the `-f` option to specify a configuration file that is not in the default location, ensure that the startup and shutdown scripts contains the `-f` option.

To read the manual page for `/etc/opt/SUNWrdc/rdc.cf` in a system window, type:

```
# man rdc.cf
```

Note – Volumes in one SNDR volume set cannot appear as volumes in another SNDR volume set.

TABLE 2-1 /etc/opt/SUNWrdc/rdc.cf File Fields

Fields	Description
Primary Host	Server on which the primary volume resides.
Primary Device	Primary volume partition to be copied. Specify full path names only (/dev/rdisk/c0t1d0s2).
Primary Bitmap	Volume partition or file in which the bitmap of the primary partition is stored. Specify full path names only (/dev/rdisk/c0t1d0s3).
Secondary Host	Server on which the secondary volume resides.
Secondary Device	Secondary volume partition. Specify full path names only (/dev/rdisk/c0t1d0s4).
Secondary Bitmap	Volume partition or file in which the bitmap of the secondary partition is stored. Specify full path names only (/dev/rdisk/c0t1d0s5).
Protocol	Network transfer protocol. Specify ip.
Mode	SNDR operating mode. sync is the SNDR mode where the I/O operation is not confirmed as complete until the remote volume has been updated. async is the other SNDR mode where the primary host I/O operation is confirmed as complete before updating the remote volume.

Caution – You must use consistent volume naming practices when configuring volumes in the rdc.cf configuration file, especially with regard to Veritas volumes created in the rootdg disk group. You must either use 'rootdg' in all configuration file pathnames, or use it in none; either: /dev/vx/rdisk/rootdg/vol01 or /dev/vx/rdisk/vol01

/etc/opt/SUNWrdc/rdc_ii.cf File

The /etc/opt/SUNWrdc/rdc_ii.cf configuration file lists all the secondary volumes on which Instant Image™ will be enabled by rdc_ii_enable.

To read the manual page for `/etc/opt/SUNWrdc/rdc_ii.cf` in a system window, type:

```
# man rdc_ii.cf
```

TABLE 2-2 `/etc/opt/SUNWrdc/rdc_ii.cf` File Fields

Fields	Description
Secondary Volume	Secondary volume on which Instant Image will be enabled.
Shadow Volume	Associated Instant Image shadow volume.
Bitmap Volume	Associated Instant Image bitmap volume.

The following is an example of an `/etc/opt/SUNWrdc/rdc_ii.cf` entry.

```
/dev/rdisk/rdc_secondary /dev/rdisk/ii_shadow /dev/rdisk/ii_bitmap
```


Summary of SNDR Commands

The `rdcadm` command can use the options in TABLE 2-3. If you do not specify a *config_file* when using the `-f` option, and you do not specify the *rdc_set*, the command uses the default configuration file `/etc/opt/SUNWrdc/rdc.cf`.

TABLE 2-4 briefly describes the SNDR commands, which have the following syntax:

```
rdcadm -a value
rdcadm -A
rdcadm [options] -d [rdc_set]
rdcadm [options] -e [rdc_set]
rdcadm [options] -l [rdc_set]
rdcadm [options] -m [rdc_set]
rdcadm [options] -m -r [rdc_set]
rdcadm [options] -u [rdc_set]
rdcadm [options] -u -r [rdc_set]
rdcadm -p
scmadm -S
```

TABLE 2-3 SNDR Command Options

Parameter	Definition
<code>-f</code> [<i>config_file</i>]	Specifies the configuration file defining the volume sets. If you do not specify a <i>config_file</i> , and you do not specify the <i>rdc_set</i> , the command uses the default configuration file <code>/etc/opt/SUNWrdc/rdc.cf</code> . If you use a configuration file that is not in the default location, modify the <code>/etc/init.d/rdc</code> file to reflect the change.
<code>-n</code>	Does not prompt the user after starting an SNDR operation using <code>rdcadm</code> . The default is to prompt and request a response. For example, after starting a full synchronization from the primary to secondary volumes, SNDR prompts Overwrite secondary with primary? (Y/N) [N].
<i>rdc_set</i>	Specifies the SNDR set configuration information. If <i>rdc_set</i> is not specified on the command line, the information is retrieved from the configuration file. The <i>rdc_set</i> format is as follows: <i>phost pdev pbitmap shost sdev sbitmap ip {sync async}</i> <i>phost</i> - Server on which the primary volume resides. <i>pdev</i> - Primary volume partition to be copied. Specify full path names only (<code>/dev/rdsk/c0t1d0s2</code>).

TABLE 2-3 SNDR Command Options

Parameter	Definition
	<p><i>pbitmap</i> - Volume partition or file in which the bitmap (scoreboard log) of the primary partition is stored. Specify full path names only (/dev/rdisk/c0t1d0s3).</p> <p><i>shost</i> - Server on which the secondary volume resides.</p> <p><i>sdev</i> - Secondary volume partition. Specify full path names only (/dev/rdisk/c0t1d0s4).</p> <p><i>sbitmap</i> - Volume partition or file in which the bitmap (scoreboard log) of the secondary partition is stored. Specify full path names only (/dev/rdisk/c0t1d0s5).</p> <p><i>ip</i> - Network transfer protocol.</p> <p><i>sync</i> <i>async</i> - SNDR operating mode. <i>sync</i> is the SNDR mode where the I/O operation is not confirmed as complete until the remote volume has been updated. <i>async</i> is the SNDR mode where the primary host I/O operation is confirmed as complete before updating the remote volume.</p>

Note – Only the `rdcadm -e`, `-d`, `-l`, and `-p` parameters can be issued from the primary or the secondary server. All other options must be issued from the primary server only.

TABLE 2-4 SNDR Command Parameter Summary

Command	Description
<code>rdcadm -a value</code>	Specifies a zero or nonzero value to set <code>rdc_auto_sync</code> to 0 or 1. If the <code>-a</code> parameter is specified, any other parameters specified are ignored.
<code>rdcadm -A</code>	Displays the state of <code>rdc_auto_sync</code> . The <code>-A</code> option causes <code>rdcadm</code> to exit with current value of <code>rdc_auto_sync</code> . If the <code>-A</code> parameter is specified, any other parameters specified are ignored.
<code>rdcadm -d</code>	Terminates all SNDR replication services between the specified primary and secondary volumes and discontinues the SNDR relationship between these volume sets. Note: You cannot disable on the secondary if the primary is running.
<code>rdcadm -e</code>	Enables SNDR for the volume set, determines if scoreboard logging is necessary, and enables scoreboard logging if it is necessary.

TABLE 2-4 SNDR Command Parameter Summary (*Continued*)

Command	Description
<code>rdcadm -l</code>	Stops SNDR replication and copy operations between primary and secondary volumes and starts independent SNDR scoreboard logging on these volumes. Note: This option does not work on the secondary for any volumes that are currently synchronizing.
<code>rdcadm -m</code>	Starts copying the current primary volume contents to the secondary volume. It also enables concurrent SNDR replication between the primary and secondary volumes.
<code>rdcadm -p</code>	Displays a list of configured SNDR volumes or sets.
<code>rdcadm -u</code>	Starts SNDR replication from primary volumes to secondary volumes currently stopped by SNDR. It also updates the corresponding secondary volumes to match the primary volumes, based on SNDR scoreboards maintained while replication was stopped.
<code>rdcadm -m -r</code>	Starts copying the current secondary volume contents to the primary volume. It also starts concurrent SNDR replication from the primary volumes to the secondary volumes.
<code>rdcadm -u -r</code>	Resynchronizes the primary volume from the secondary volume based on SNDR logs maintained while SNDR replication was stopped. It also enables SNDR replication between the primary and secondary volumes.
<code>scmadm -S</code>	Displays SNDR and cache statistics for the system on which it is invoked.

Copy Primary Volume to Secondary Volume

```
rdcadm -m
rdcadm -m -f config_file
rdcadm -m phost pdev pbitmap shost sdev sbitmap ip {sync|async}
```

After you type the command, the response is displayed:

```
Overwrite secondary with primary? (Y/N) [N]
```

The default response is no (N). If you type Y (yes), the `rdcadm -m` command starts a full copy of the primary volume to the secondary volume. It also enables concurrent SNDR replication between the primary and secondary volumes; any new writes to the primary are also replicated to the secondary volume. See FIGURE 2-1. Use `scmadm -S` to monitor the operation.

When To Use This Command

Use this command when all of the following cases exist:

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




FIGURE 2-1 `rdcadm -m` Command Result

Stop Replication

```
rdcadm -l
rdcadm -l -f config_file
rdcadm -l phost pdev pbitmap shost sdev sbitmap ip {sync|async}
```

After you type the command, the response is displayed:

```
# rdcadm -l
RDC: auto sync is now OFF
rdcadm: warning: Using /etc/opt/SUNWrdc/rdc.cf
Suspend RDC replication and log? (Y/N) [N]
```

The default response is **N** (no). If you type **Y** (yes), the `rdcadm -l` command temporarily stops SNDR replication, copy, refresh, and restore operations between the primary and secondary volumes. It also starts independent SNDR scoreboard logging on these volumes.

SNDR scoreboard logging identifies the parts of each volume that are updated while SNDR is stopped. SNDR later uses these scoreboard logs to re-synchronize the volumes, when you issue the `rdcadm -u` command.

When To Use This Command

Use this command to temporarily stop SNDR volume updates while maintaining the SNDR scoreboard logs or to start an SNDR set for the first time (forcing the set into scoreboard logging mode, leaving the system no opportunity to misjudge the situation and start up in replicating mode).

Start Replication

Note – Before using this command, quiesce the workload to the affected volumes—stop the host application from writing to the volumes until the updates and copies are complete and the volume sets are fully synchronized. This ensures that the primary and secondary volumes match before replication of new updates starts. Use `scmadm -S` to view the synchronization progress.

```
rdcadm -u
rdcadm -u -f config_file
rdcadm -u phost pdev phitmap shost sdev sbitmap ip {sync|async}
```

The `rdcadm -u` command starts SNDR replication from the primary volume to the secondary volume. It also uses the SNDR scoreboard logs to start any copy, refresh, or restore operations that were underway when SNDR was stopped. Any new writes to the primary are also replicated to the secondary volumes. See FIGURE 2-2. Use `scmadm -S` to monitor the operation.

When To Use This Command

Use this command when replication on the specified primary volume has been stopped and active SNDR scoreboard logs exist.

Tip – You can use this command after a system failure or an intersite link failure.




FIGURE 2-2 `rdcadm -u` Command Result

Restore Primary Volume From Secondary

Note – Before using this command, quiesce the workload to the affected volumes—stop the host application from writing to the volumes until the updates and copies are complete and the volume sets are fully synchronized. This ensures that the primary and secondary volumes match before replication of new updates starts. Use `scmadm -S` to view the synchronization progress.

```
rdcadm -m -r
rdcadm -m -r -f config_file
rdcadm -m -r phost pdev pbitmap shost sdev sbitmap ip {sync|async}
```

After you type the command, the response is displayed:

```
Overwrite primary with secondary? (Y/N) [N]
```

The default response is no (N). If you type Y (yes), the `rdcadm -m -r` command starts a full copy of the secondary volume to the primary volume. It also enables concurrent SNDR replication between the primary and secondary volumes; any new writes to the primary are also replicated to the secondary volumes. See FIGURE 2-3. Use `scmadm -S` to monitor the operation.




FIGURE 2-3 `rdcadm -m -r` Command Result

When To Use This Command

Use this command when all of the following cases exist:

- The contents of the primary and secondary volumes may be different
- The secondary volume has the desired contents
- You wish to completely overwrite the contents of the primary volume
- No scoreboard logging information exists to incrementally re-synchronize the volumes
- You intend to start replication from the primary to the secondary volume

Refresh Primary Volume From Secondary Volume

Note – Before using this command, quiesce the workload to the affected volumes—stop the host application from writing to the volumes until the refresh is complete and the volume sets are fully synchronized. This ensures that the primary and secondary volumes match before replication of new updates starts. Use `scmadm -S` to view the synchronization progress.

```
rdcadm -u -r
rdcadm -u -r -f config_file
rdcadm -u -r phost pdev pbitmap shost sdev sbitmap ip {sync|async}
```

After you type the command, the response is displayed:

```
Refresh primary with secondary? (Y/N) [N]
```

The default response is no (N). If you type Y (yes), the `rdcadm -u -r` command re-synchronizes the primary volume from the secondary volume, based on SNDR scoreboard logs maintained while replication was stopped. It also enables concurrent SNDR replication between the primary and secondary volumes; any new writes to the primary are also replicated to the secondary volumes. See FIGURE 2-4. Use `scmadm -S` to monitor the operation.

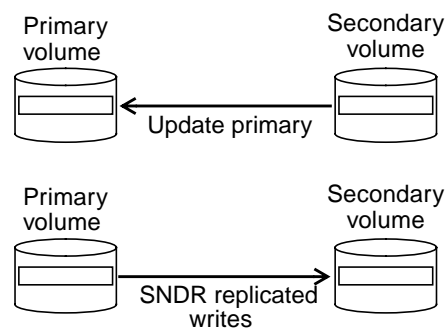


FIGURE 2-4 `rdcadm -u -r` Command Result

When To Use This Command

Use this command when all of the following cases exist:

- SNDR replication was stopped and SNDR scoreboard logging was active
- The secondary volume has the desired contents
- You wish to overwrite those segments of the primary volume that are different from the secondary as defined in the SNDR scoreboard logs
- You intend to start replication from the primary to the secondary volume

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

Disable Replication

```
rdcadm -d
rdcadm -d -f config_file
rdcadm -d phost pdev pbitmap shost sdev sbitmap ip {sync|async}
```

Description

After you type the command, the response is displayed:

```
Discard scoreboards and terminate SNDR association? (Y/N) [N]
```

The default response is no (N). If you type Y (yes), the `rdcadm -d` command terminates all SNDR replication services between the specified primary and secondary volumes and discontinues the SNDR relationship between these volume sets. It also discards any active SNDR scoreboard logs that track temporary differences between primary and secondary volumes.

When To Use This Command

Use this command when the primary and secondary volumes no longer need to be associated with each other as SNDR volume sets.

As `rdcadm -d` discards any active SNDR scoreboard logs, use one of the following commands to re-establish identical replicated sets and re-enable SNDR replication:

- `rdcadm -m`
- `rdcadm -m -r`

Note – Use `rdcadm -l` instead of `rdcadm -d` to temporarily stop replication and maintain the SNDR volume set association and related SNDR scoreboard logs.

Enable SNDR

```
rdcadm -e  
rdcadm -e -f config_file  
rdcadm -e phost pdev pbitmap shost sdev sbitmap ip {sync|async}
```

Description

After you type the command, the response is displayed:

```
Enable RDC? (Y/N)[N]:
```

The default response is no (N). If you type Y (yes), the `rdcadm -e` command enables scoreboard logging if needed. If scoreboard logging is not needed, the `rdcadm -e` command enables the SNDR sets so that they can continue in replicating mode.

When To Use This Command

Use this command to enable the SNDR volume set and to determine if scoreboard logging is necessary.

Toggle the SNDR Auto-Synchronization Value

```
rdcadm -a value
```

Description

The `rdcadm -a value` command enables or disables SNDR auto-synchronization. If *value* is 0, SNDR auto-synchronization is disabled. If *value* is 1, SNDR auto-synchronization is enabled.

When To Use This Command

Use this command when you want to enable or disable auto-synchronization. If SNDR auto-synchronization is enabled on both the primary and the secondary hosts, the SNDR synchronization daemon attempts to re-synchronize if the system reboots or there are link failures.

Display the SNDR Auto-Synchronization Value

```
rdcadm -A
```

Description

The `rdcadm -A` command displays the status of SNDR auto-synchronization. If SNDR auto-synchronization is enabled, the following response is displayed:

```
RDC: auto sync is now ON
```

If SNDR auto-synchronization is disabled, the following response is displayed:

```
RDC: auto sync is now OFF
```

When To Use This Command

Use this command when determining whether SNDR auto-synchronization is enabled or disabled.

Display a List of Configured SNDR Sets

```
rdcadm -p
```

Description

The `rdcadm -p` command displays a list of the configured SNDR sets. The following is an example of output of this command on the Primary:

```
RDC: auto sync is now OFF
/dev/rdisk/c0t10d0s3      atm2:/dev/rdisk/c1t0d0s3
/dev/rdisk/c0t16d0s4     atm2:/dev/rdisk/c1t6d0s4
```

When To Use This Command

Use this command to display the configured SNDR sets.

Getting SNDR and Storage Cache Statistics: the `scmadm` Command

`scmadm -S` displays storage cache information and can be used to view SNDR statistics. Refer to the `scmadm(1M)` man page for more detail about this command. The basic command syntax for SNDR is:

```
/usr/opt/SUNWscm/sbin/scmadm -S -M [-d delay_time] [-l logfile]  
[-r [range]] [-z]
```

Note – Brackets [] indicate optional arguments.

TABLE 2-5 describes the `scmadm -S` options.

After you type `scmadm -S` with the desired options, you can use the following keys:

- `f`—scroll forward through the display
- `b`—scroll backward through the display
- `m`—toggle between regular cache statistics and SNDR 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 SNDR screen displays:

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

TABLE 2-5 `scmadm -S` Options

Option	Description
<code>-M</code>	Displays SNDR statistics.
<code>-d <i>delay_time</i></code>	Sets the display update time to <i>delay_time</i> seconds.

TABLE 2-5 `scmadm -S` Options (Continued)

Option	Description
<code>-l logfile</code>	Writes all screen outputs to the specified <i>logfile</i> .
<code>-r [range]</code>	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 all volumes are displayed. Specify the range in this format: <code>n[:n,n]</code> Where <i>n</i> 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): <code>-r 3,6:7,8,9:12,14:15</code> <code>-r 3,6:12,14,15</code>
<code>-z</code>	Clears the statistics before displaying current statistics.

Example `scmadm -S` SNDR Screen Display

The SNDR statistics display includes the fields in TABLE 2-6. Type `scmadm -S -M` to display SNDR-specific statistics; press the `m` key to toggle between SNDR screens and regular cache statistics after typing the command.

TABLE 2-6 Fields for SNDR Screen Display

Field	Description
<code>primary</code>	Specifies the primary dual copy partition.
<code>link status</code>	Indicates whether replication is enabled, disabled, or synchronization direction if a synchronization is in progress.
<code>secondary</code>	Specifies the secondary dual copy partition.
<code>dual copy status</code>	Indicates current status of the dual copy using the following keywords: <code>logging</code> - Data is not being replicated and changes are being logged.

TABLE 2-6 Fields for SNDR Screen Display (*Continued*)

Field	Description
	<p><code>need rev sync</code> - Logging, but an earlier reverse synchronization did not complete. The primary volume may be in an inconsistent state. The <code>need rev sync</code> state is cleared following a successful reverse synchronization, or by an Instant Image™ update to the primary volume (which restores an earlier valid snapshot and removes the need to repeat the reverse synchronization to restore consistency).</p> <p><code>need sync</code> - Logging, but an earlier synchronization did not complete. The secondary volume may be in an inconsistent state. The <code>need sync</code> state is cleared following a successful synchronization, or by an Instant Image update to the secondary volume (which restores an earlier valid snapshot and removes the need to repeat the synchronization to restore consistency).</p> <p><code>replicating</code> - Data written to the primary is being replicated to the secondary. Formerly called <code>enabled</code>.</p> <p><code>rev sync</code> - A synchronization from secondary to primary is under way.</p> <p><code>sync</code> - A synchronization from primary to secondary is under way.</p> <p><code>volume down</code> - The volume has failed.</p>
<code>recovery needed</code>	Displays the percentage of segments of the partition that require recovery.
<code>recovery completed</code>	Displays the percentage of the partition that has been resynchronized.

See FIGURE 2-5. The = 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.

```

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

primary           link status      secondary       dual copy status

  ma:...sk/c0t117d0s3 *>*      atmsi:...rdsk/c1t6d0s3   sync
  ma:...sk/c0t115d0s3 *>*      atmsi:...dsk/c1t9d0s3   sync
  ma:...sk/c0t116d0s3 *>*      atmsi:...dsk/c1t6d0s3   sync
  ma:...dsk/c0t96d0s3 *=*      atmsi:...rdsk/c1t0d0s3   replicating

-----
partition         recovery needed          recovery completed

...k/c0t117d0s3 [*****          ] 29.53% [*****          ] 70.47%
...k/c0t115d0s3 [*****          ] 29.16% [*****          ] 70.84%
...k/c0t116d0s3 [*****          ] 29.38% [*****          ] 70.62%
-----

```

FIGURE 2-5 `scmadm -S` Link Status Display

Example `scmadm -S` Screen Display

The `scmadm -S` statistics include the fields in TABLE 2-7. Use the `m` key to toggle between these statistics and the SNDR 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 display in FIGURE 2-6 contains the cache statistics and cumulative cache statistics displays.

TABLE 2-7 `scmadm -S` Fields

Field	Description
<code>cd</code>	Specifies the cache descriptor number.
<code>cached_partition</code>	Specifies the disk partition being monitored.
<code>disk_io</code>	Displays the number of Kbytes per second read from or written to physical disks.
<code>cache</code>	Displays the number of Kbytes per second read from or written to data cache.

TABLE 2-7 `scmadm -S` Fields (Continued)

Field	Description
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 <code>disk_io(writes)</code> .
accesses/s	Displays the number of I/O operations (reads per second + writes per second) serviced by the cache per second. The <code>accesses/s</code> total is the sum of the hits and misses.
read/s	Displays the number of read requests (hits) serviced per second.
	<code>misses/s</code> —Displays the number of read misses per second.
write/s	Displays the number of write requests serviced per second.
	<code>misses/s</code> —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, as specified in the <code>/etc/sd.cf</code> file. Cache size is calculated as:
blocksize	Displays the <code>blk_size</code> as specified in the <code>/etc/sd.cf</code> file.
Write blocks available	Displays the available write cache in blocks. <code>Net n</code> in the field represents each network. A network is added for each board that is added.
LRU stat	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.

```

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 )

```

FIGURE 2-6 Sample `scmadm -S` Display -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
```

FIGURE 2-7 Sample `scmadm -S Display` -Cumulative Cache Statistics

SNDR Operations

SNDR operations range from establishing initial replicated sets to recovering from a failed primary storage platform. This chapter describes each operation.

- Before You Begin Using SNDR
 - Define primary and secondary sites
 - Monitoring SNDR operations
 - Mutual backup considerations
 - Selecting volumes for SNDR
 - Adding volumes to the storage volume driver configuration file
 - Sun StorEdge Instant Image
- Establishing Replicated Images for the First Time
- Handling SNDR Interruptions
- Restoring Volumes After Secondary Site Failures
- Disaster Recovery Rehearsals
- Primary Site Failures
- Recovering From a Primary Site Disaster
- Updating the Primary Site From the Secondary Site
- Completely Restoring a Primary Server from the Secondary
- Temporarily Stopping Remote Replication
- Disabling Remote Replication
- Mirrored Partitions for Multiple Copies

You can combine SNDR with volumes from Sun StorEdge Instant Image™ software to create multiple copies of each device.

Before You Begin Using SNDR

This section describes things to consider before you start using the full capabilities of SNDR.

- Define primary and secondary sites

In the examples in this chapter, Site A is the local primary site and Site B is the remote secondary site. Depending on the example, either site can be the primary or secondary of the remote dual copy operation.

- Monitoring SNDR operations

Use `scmadm -S` to monitor the progress of SNDR operations. Use a separate `scmadm -S` window for each system. Chapter 2 describes this command. You may monitor a subset of SNDR statistics on the secondary system from the secondary system.

- Mutual backup considerations

Sometimes the distinction between primary and secondary sites is blurred. As applications are geographically distributed, a storage system at Site B may function as a remote disk backup to Site A and as a direct storage resource for applications on Host B. Under these circumstances, keep remote-dual-copies of Host B volumes at Site A. This reciprocal backup arrangement supported by SNDR is called *mutual backup*.

With mutual backup, the SNDR volumes considered primary by Site B are administered from the Site B console. Site B remote-dual-copied devices are considered secondary volumes at Site A.

- Selecting volumes for SNDR

One initial SNDR planning step is to determine the volumes to include in the remote replication operations. Balance remote accessibility and recoverability against capacity usage and I/O response time.

Generally, you will 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 on the SNDR command line or collectively using a configuration file, as described in Chapter 2.

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

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

- Adding volumes to the storage volume driver configuration file

Before you can use SNDR, you must add the SNDR volumes to the Storage Volume (SV) driver configuration file and reconfigure the driver.

You cannot enable and use the SNDR software on volumes that are not in the SV driver configuration file and enabled by the SV driver. The SV driver provides a simple layer of volume management for the SNDR software.

- Add Volumes to the Configuration File

Tip – If you want to add devices later, repeat Step i to Step iv.

i. Decide which volumes you are going to use with SNDR.

Caution – You must use consistent volume naming practices when configuring volumes in the `sv.cf` configuration file, especially with regard to Veritas volumes created in the `rootdg` disk group. You must either use 'rootdg' in *all* configuration file pathnames, or use it in none; either: `/dev/vx/rdisk/rootdg/vol01` or `/dev/vx/rdisk/vol01`

ii. Using a text editor, add the volumes to the Storage Volume (SV) driver configuration file `/etc/opt/SUNWspcv/sv.cf`.

This file includes two fields per line for each volume: the resource name field and access mode field. Separate the fields in the file with spaces.

The resource name field contains the name of a raw device for which the SV driver is to be enabled. *It must be the full path to the raw device node and cannot be the block device node.*

The access mode field can be `cache` or `raw`. This field specifies the way in which SNDR accesses the volume: through the Storage Cache Management module `cache` or directly through the `raw` device. The `cache` mode is preferable.

The following example shows four primary and four secondary devices. Two of each type will be accessed directly and two of each type will be cached.

Note – The example file shows volume names in order: primary followed by secondary. This is not required; you can put the volume names in any order.

```
# devices clt3d0s0 to clt3d0s3 are primary volumes
/dev/rdisk/clt3d0s0 cache
/dev/rdisk/clt3d0s1 cache
/dev/rdisk/clt3d0s2 raw
/dev/rdisk/clt3d0s3 raw
#
# devices clt3d0s4 to clt3d0s7 are secondary volumes
/dev/rdisk/clt3d0s4 cache
/dev/rdisk/clt3d0s5 cache
/dev/rdisk/clt3d0s6 raw
/dev/rdisk/clt3d0s7 raw
```

iii. **Save and exit** `/etc/opt/SUNWspsv/sv.cf`.

iv. **Use the `svadm(1SV)` utility to reconfigure the SV driver to add the new volumes in the configuration file to the SV layer.**

```
# /usr/opt/SUNWesm/SUNWspsv/sbin/svadm -r
```

■ **Sun StorEdge Instant Image**

Instant Image is a prerequisite for SNDR; it is a point-in-time volume copy Data Service for the Solaris operating environment. With Instant Image, 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.

Establishing Replicated Images for the First Time

Before using SNDR replicated updates, make sure that the contents of the local and remote volume sets match. In the examples, Site A is the local primary site and Site B is the remote secondary site. Depending on the example, either site can be the primary or secondary of the remote dual copy operation. All SNDR operations start from the primary system, as you issue the commands from the primary system.

Existing Primary and Newly-Formatted Secondary

If the primary volume has already been updated locally (without SNDR replication) and contains important information, you must first copy its image to the newly formatted secondary volume. Perform this full synchronization for the specified primary and secondary with the `rdcadm -m` command. SNDR can perform this initial volume-to-secondary-volume copy while concurrently forwarding new online updates to the secondary.

To fully synchronize specific remote volumes while applying new updates, type:

```
host1# rdcadm -m
```

FIGURE 3-1 shows the full resynchronization process using an ATM link. This process also enables concurrent write updates from the local to remote volumes.

The data flow is:

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

These steps repeat until the entire volume is copied. Use `scmadm -S` to monitor the progress. As these re-synchronizations finish, others begin until all the requested volume sets have been synchronized.

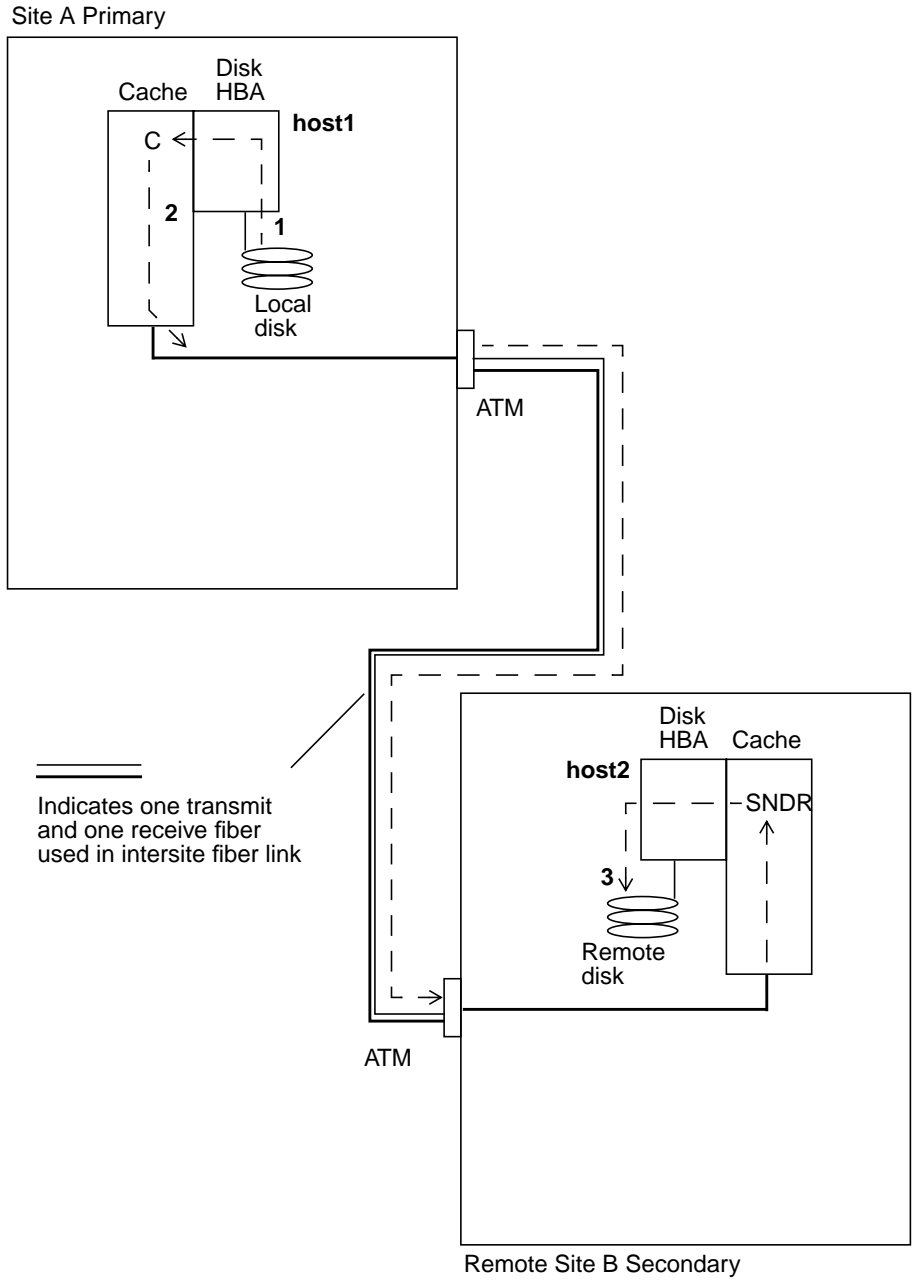


FIGURE 3-1 Full Synchronization (Volume-to-Volume Copy) Using ATM

Handling SNDR Interruptions

SNDR 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 SNDR service. Interruptions can arise from a failure or impairment in the intersite link or an outage at the remote site. In some cases, interruptions are introduced intentionally to exercise failure contingencies, as during the disaster recovery rehearsals described later in this chapter.

Primary Logging During SNDR Interruptions

During interruptions, 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 SNDR service is restored, update resynchronization may be requested using the scoreboard logging information to update the remote site.

Re-synchronizing Volumes After an SNDR Interruption

Typically, interruptions in SNDR services are infrequent. Whether the event is an isolated incident or part of a larger disaster determines the action you should take.



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

If the SNDR 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 inconsistent and difficult to recover from. When you determine that the SNDR interruption is an isolated incident and the condition has been repaired, enable update synchronization using the following SNDR command:

```
host1# rdcadm -u
```

Note – If the secondary volume state is unknown because of system or disk failure, you may need to make full volume copies to re-establish matching SNDR sets. In this case, use the `rdcadm -m` command.

FIGURE 3-2 shows an update resynchronization using an ATM link from the primary system to its secondary system, when the secondary volumes have been left stale by the interruption.

1. SNDR on `host1` examines a scoreboard for one of the SNDR-managed volumes affected by the interruption.
2. SNDR on `host1` requests the blocks updated during the interruption from the up-to-date volume. The data may already be resident in `host1` data cache, or it may require a local disk access.
3. SNDR on `host1` transmits the update blocks to `host2` SNDR using the SunATM connection.
4. SNDR on `host2` refreshes its stale replicated image with the updated blocks and acknowledges the action to `host1`.
5. SNDR revises the scoreboard to track the remote update.

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

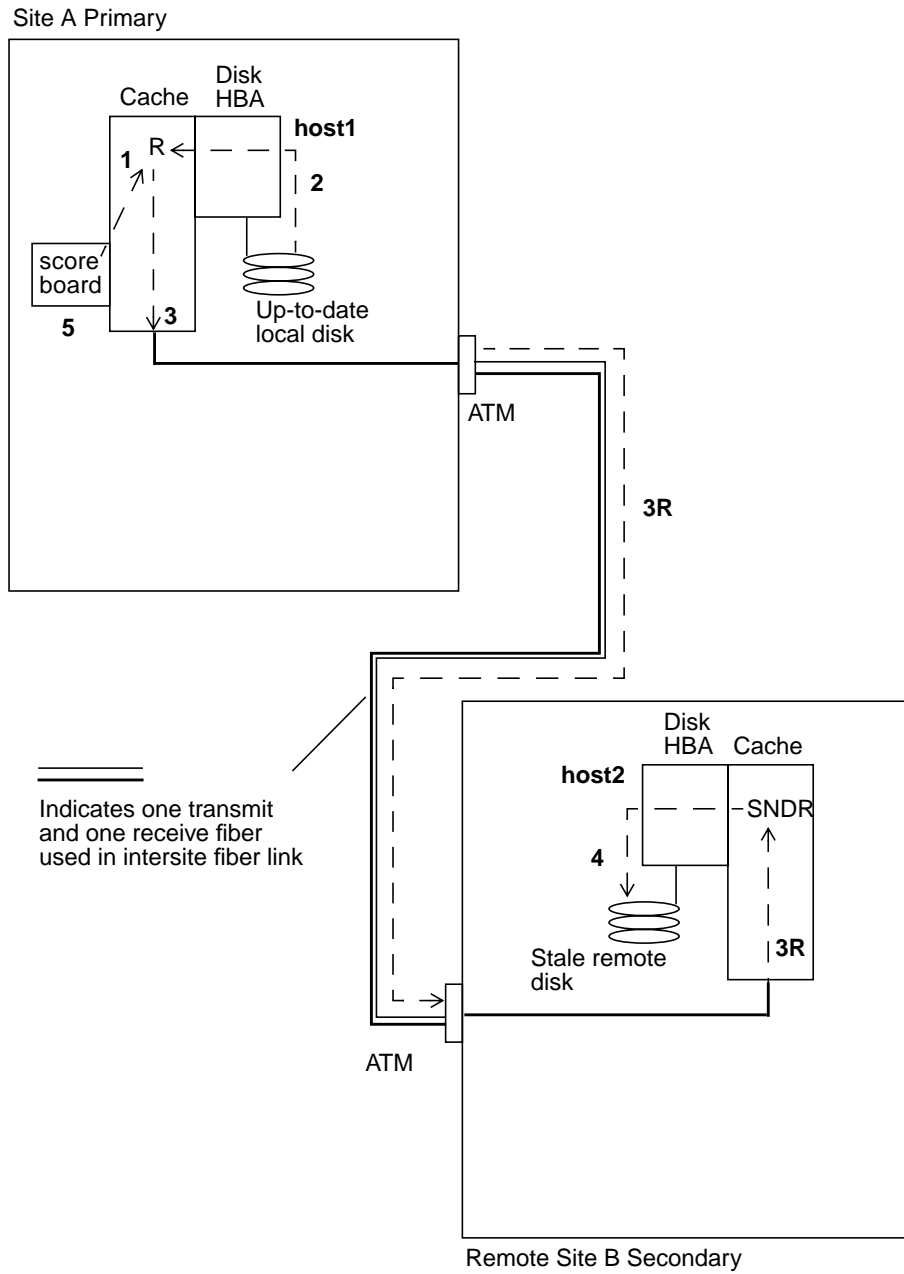


FIGURE 3-2 Update Synchronization of a Secondary, Using an ATM Link

Auto-Resynchronization After an SNDR Interruption

The SNDR synchronization daemon (`rdcsyncd`) starts when SNDR does. It monitors the existing SNDR configuration and attempts to re-synchronize if the system reboots or the link is dropped.

Automatic re-synchronizations are initiated only on the primary system, and only if `rdc_auto_sync` is set to *on* (1) for both the primary and secondary systems. By default, `rdc_auto_sync` is set to *off* (0) in `/usr/kernel/drv/rdc.conf`. You can change this to 1 when the SNDR configuration is set up.

An auto-synchronization is needed to enable a secondary system if any of the volumes for that secondary system are logging and none are *Syncing*, *Reverse Syncing*, or *Need Reverse Sync*.

`rdcsyncd` performs the auto-resynchronization using the `rdc_link_up` shell script. `rdc_link_up` uses the `rdc_ii_enable` and `rdc_ii_disable` scripts. `rdcsyncd` and these shell scripts are in `/usr/opt/SUNWrdc/lib`.

`rdc_link_up` first runs `rdc_ii_enable` on the secondary system. This enables Instant Image for all volumes listed in `/etc/opt/SUNWrdc/rdc_ii.cf`. The `rdc_ii.cf` file contains the following type of entry:

```
/dev/rdisk/rdc_secondary /dev/rdisk/ii_shadow /dev/rdisk/ii_bitmap
```

where `/dev/rdisk/rdc_secondary` is the SNDR secondary volume name. `/dev/rdisk/ii_shadow`, and `/dev/rdisk/ii_bitmap` are the associated Instant Image shadow and bitmap volumes. The entire operation aborts if any of the enables fail.

No errors occur if any of the secondary devices are not in the `rdc_ii.cf` file (that is, they are not Instant Image devices) or if Instant Image for a secondary device is already enabled.

If the secondary is enabled, `rdcadm -u` (update sync) operations are started on the primary system to update all the secondaries. When the synchronizations complete, `rdc_ii_disable` runs on the secondary system to disable the Instant Images that were enabled.

Using `rdcadm -l` or starting a reverse synchronization turns off `rdc_auto_sync` because you have taken manual control of SNDR.

The `rdcsyncd` may also turn off `rdc_auto_sync` if it detects problems with synchronizations it has started.

Restoring Volumes After Secondary Site Failures

The effort required to restore a secondary site to its remote dual copy state depends on the hardware and software that is replaced or repaired. This section describes restoring disk, system, and intersite links.

Restoring a Failed Secondary Disk

If a volume is not recoverable after a disk failure, follow the steps in this section.

FIGURE 3-3 shows restoration of a failed secondary disk. The process follows the same steps used to establish an initial SNDR set. The concurrent write updates have been omitted from the figure.

Type on one line:

```
host1# rdcadm -m host1 /dev/rdisk/abc /dev/rdisk/pbitmap host2  
/dev/rdisk/xyz /dev/rdisk/sbitmap ip [sync|async]
```

where */dev/rdisk/abc* is the active primary volume on *host1*, */dev/rdisk/pbitmap* is the scoreboard log for the primary partition, */dev/rdisk/xyz* is the newly formatted secondary volume on *host2*, and */dev/rdisk/sbitmap* is the scoreboard log for the secondary partition. *abc* and *xyz* are device names in *cntndnsn* format. Brackets indicate optional arguments; choose *sync*, *async*, or neither.

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

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

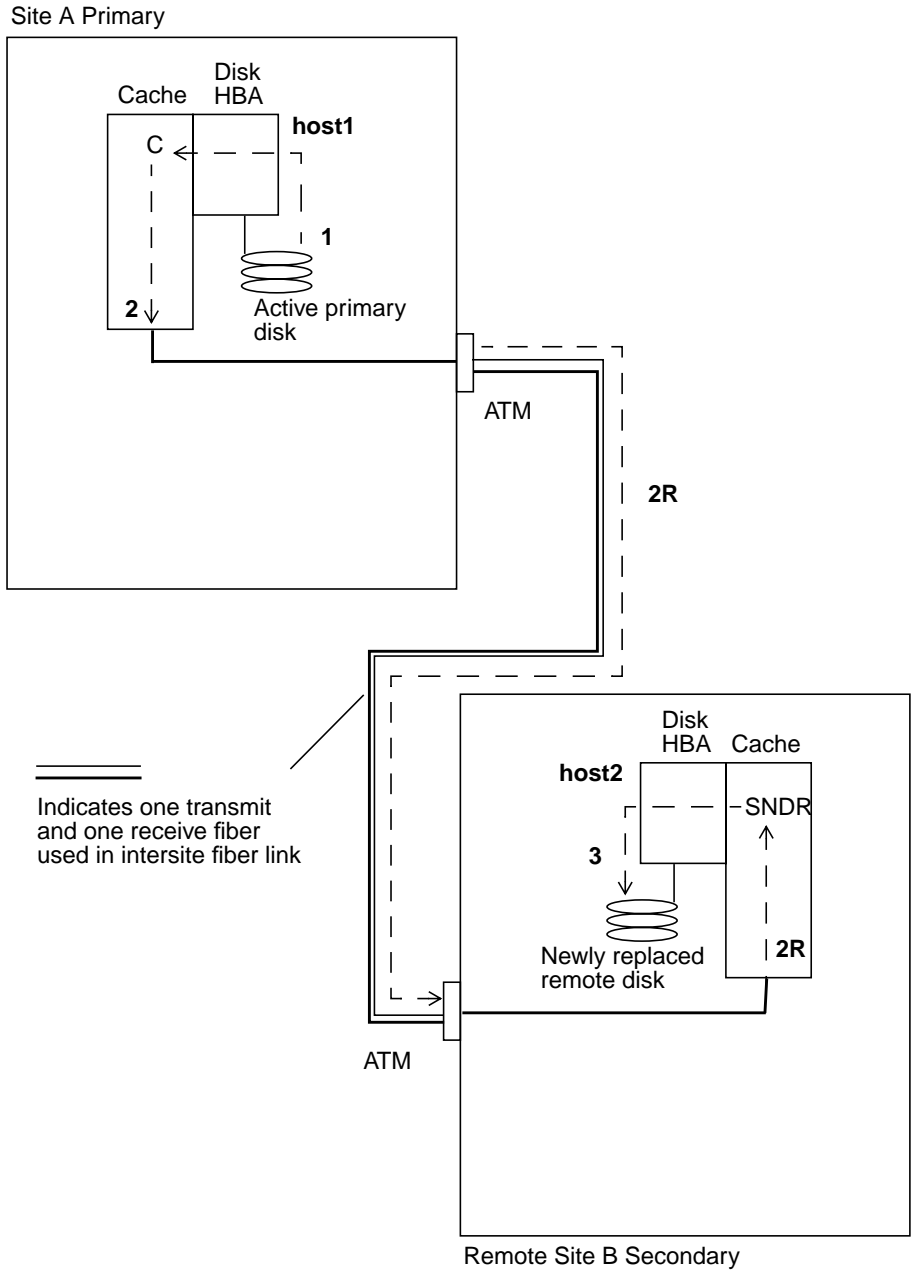


FIGURE 3-3 Restoring a Failed Secondary Disk, Using an ATM Link

Restoring a Failed Secondary Server

Restoring a failed secondary server may involve both updating and fully re-synchronizing the secondary disks, depending on the severity and duration of the failure.

The `start` command updates any secondary volumes whose contents were unchanged by the failure, and whose primary members have seen little activity during the outage.

```
host1# rdcadm -u -f config_file
```

where the *config_file* lists volumes to be updated.

Fully synchronize volumes residing on secondary disks that were replaced or whose state is unknown:

```
host1# rdcadm -m -f config_file
```

where the *config_file* lists volumes requiring full volume-to-volume copies.

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 may use an automatic switch or may need to manually switch fiber cable pairs at a patch panel to redirect the signal to the alternate route.

Disaster Recovery Rehearsals

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.

SNDR lets you perform disaster rehearsals easily, encouraging frequent verification of disaster preparedness.

Rehearsing Recovery From Secondary Volumes

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

Complete these steps before proceeding with recovery:

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

During a read operation from the secondary recovery copy:

1. The secondary remote site host issues a read request on `host2` against the recovery copy.
2. Step 1 being a cache miss, the appropriate blocks are read from the recovery copy disk into `host2` data cache.
3. 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 SNDR role. See “Rolling Back Updates After a Takeover Rehearsal” on page 17” for the appropriate procedures.

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

FIGURE 3-4 shows secondary write logging during rehearsals:

1. The remote site host issues a write request to its directly-attached server.
2. The write data *D* is placed into *host2* data cache.
3. SNDR keeps a scoreboard for each SNDR-managed volume to indicate changes from the last known image synchronized with *host1*.
These scoreboard logs may 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 disk.

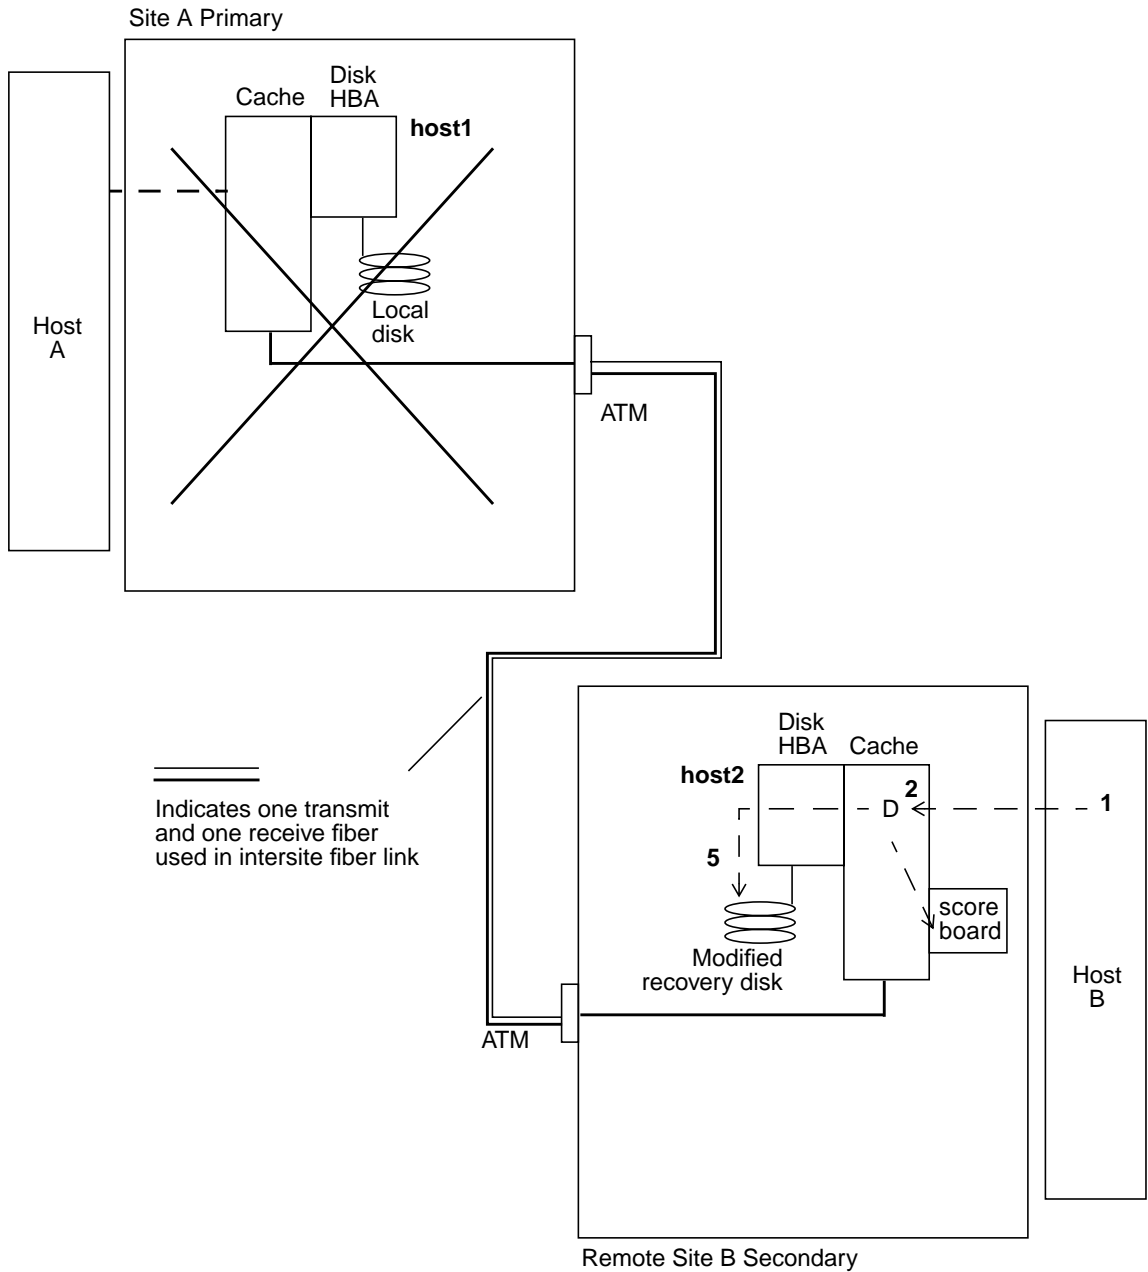


FIGURE 3-4 Secondary Scoreboard Logging, Using an ATM Link

Rolling Back Updates After a Takeover Rehearsal

When the rehearsal completes, SNDR uses the secondary scoreboard logs to update the appropriate blocks with the current information from the primary devices; the update procedure rolls back the test changes. At the primary console, enable SNDR secondary rollback by typing:

```
host1# rdcadm -u
```

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 “Updating the Primary Site From the Secondary Site” on page 23” for more detail.

Primary Site Failures

This section contains information needed for handling primary site failures.

Failed Primary Disk

SNDR provides continuous data access during primary disk failures. The SNDR high-availability features are a superset of RAID 1 and RAID 5 storage protection that can be optionally configured for the primary volumes. The SNDR remote disk 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 SNDR 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. SNDR 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 SNDR remote disk access only when two or more of the disks in the RAID 5 stripe fail.

FIGURE 3-5 shows the SNDR 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 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 SNDR forwards the read request to `host2`.
3. SNDR on `host2` acts on the remote request by reading its remote mirror disk.
4. SNDR 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.

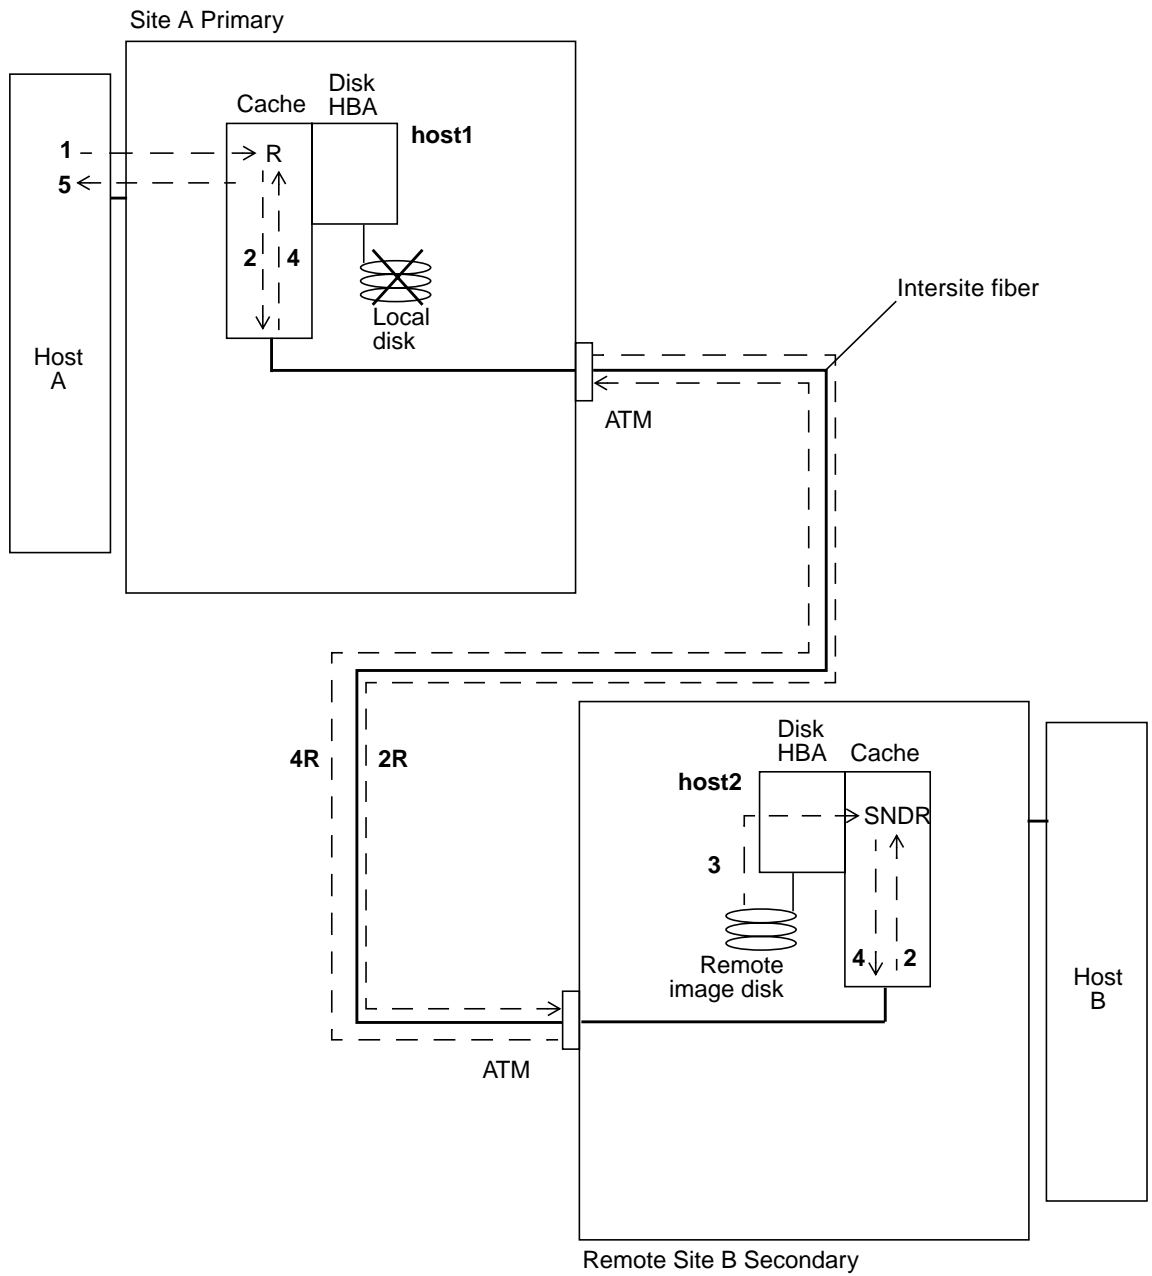


FIGURE 3-5 Pass-Through Reads to Bypass Local Disk Failure, Using an ATM Link

▼ To Restore a Failed Primary Volume (Example)

1. Keep the failed volume enabled under SNDR.

SNDR marks the device as failed when it is unable to read or write from it. SNDR 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.



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

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. Use the `scmadm -S` command to see that the cached partition has pinned data.

To clear the pinned data on the `cached_partition` relating to a disk to be replaced, use `scmadm -p diskname` to purge the pinned data. If a corrected transient error was at fault, use `scmadm -s diskname` to sync pinned data.

5. If auto-synchronization is off, perform any Instant Image operations necessary to preserve a valid point-in-time copy of the data on the secondary system before starting an SNDR sync to update the secondary.

6. If auto-synchronization is on, if the volume specified as the Instant Image shadow of the SNDR secondary volume is mounted, the SNDR secondary will be in an unknown state when SNDR starts a sync. If this state is not appropriate for an Instant Image snapshot, then SNDR should not be put into auto-synchronization mode.

7. After the disk failure is corrected and any related Volume Manager volumes are recreated (refer to the Volume Manager documentation for details), perform the following reverse sync while the primary affected volumes are unmounted:

```
host1# rdcadm -m -r atm1 /dev/rdisk/c0t0d0s3 /dev/rdisk/c1516d0s4 amt2 \  
      /dev/rdisk/c2t2d0s3 /dev/rdisk/c3t0d0s4 ip [sync|async]
```


This is a reverse full synchronization, as the secondary volume on `host2` is re-synchronizing the new primary volume on `host1`. FIGURE 3-6 shows the full reverse synchronization process.

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

8. 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.

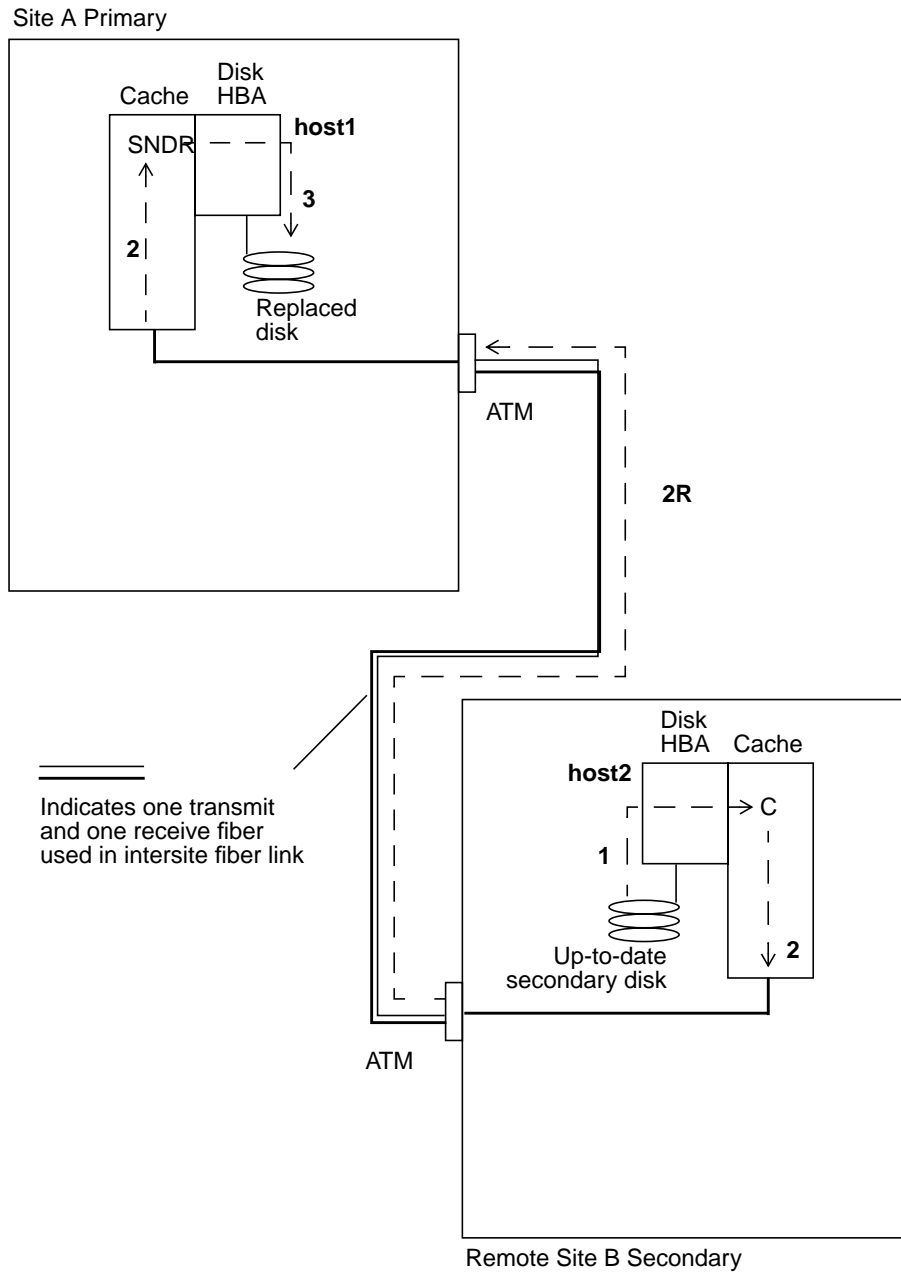


FIGURE 3-6 Reverse Full Synchronization, Using an ATM Link

Recovering From a Primary Site Disaster

SNDR minimizes the effects of a disaster at the primary site by keeping the secondary storage images up-to-date. Although the secondary SNDR cache contains the latest writes issued on the primary before the disaster, that data may not have been destaged to the secondary disks yet. After detecting an interruption in the SNDR service, SNDR automatically destages the secondary SNDR cache to its corresponding secondary volumes.

After all the secondary volumes have been updated with the latest SNDR cache images, the secondary volumes may 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 SNDR 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 “Completely Restoring a Primary Server from the Secondary” on page 26.

Updating the Primary Site From the Secondary Site

The primary volume state may be frozen as host workloads are switched to the secondary site. This is true in more realistic disaster rehearsals and may be true in confined disasters. Under these conditions, the secondary site activates its SNDR logs (scoreboards) 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. The following SNDR command, typed at the primary system, starts the reverse update resynchronization from the secondary to the primary:

```
host1# rdcadm -u -r
```

FIGURE 3-7 shows a reverse update resynchronization from the secondary system to its primary system.

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

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

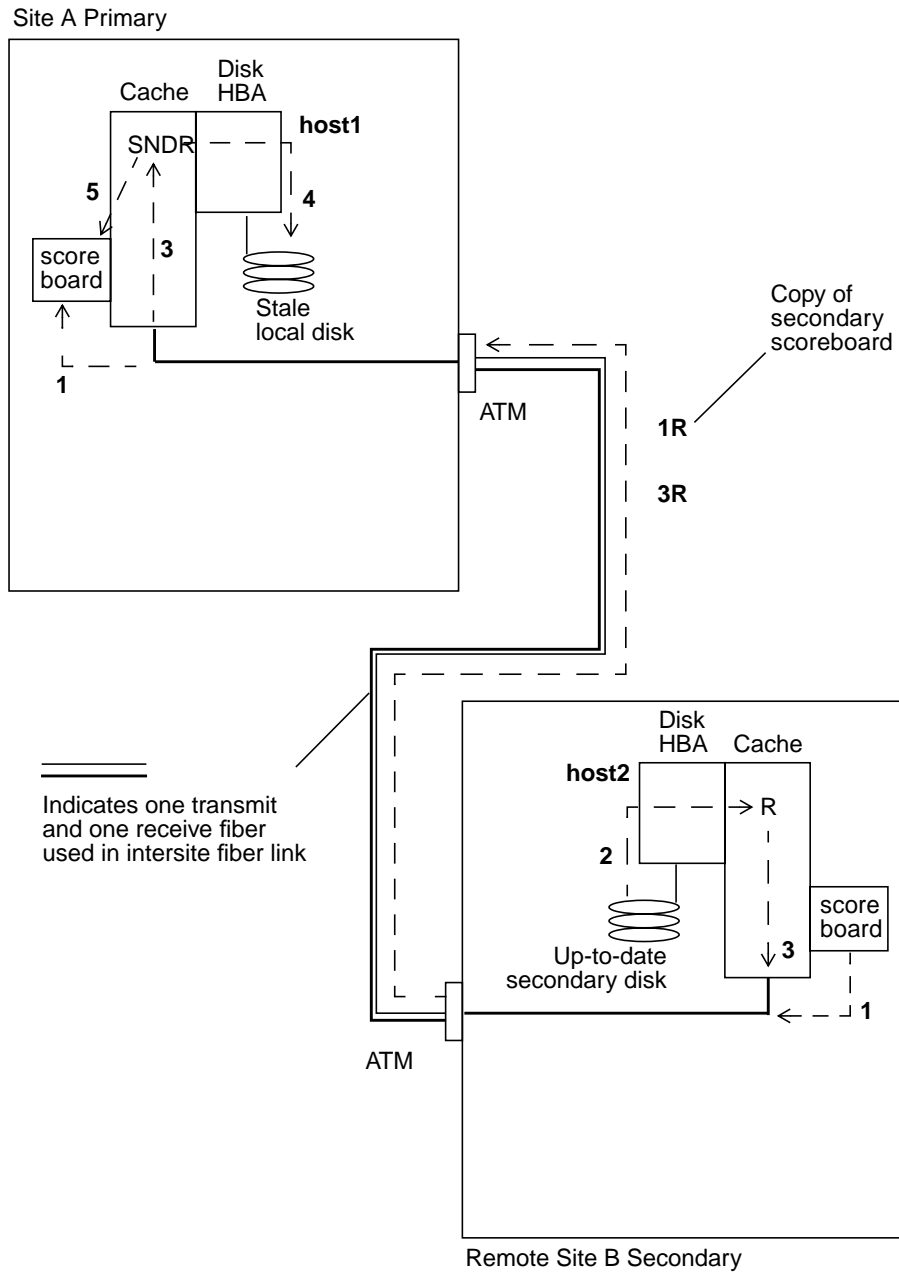


FIGURE 3-7 Reverse Update Synchronization, Using an ATM Link

Completely Restoring a Primary Server from the Secondary

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 server. In other words, volume-to-volume copies from the secondary to the primary are required for all SNDR-managed volumes. This reverse synchronization process ensures that only the latest data is deposited on the primary disks. On the primary server, type:

```
host1# rdcadm -m -r
```

FIGURE 3-8 shows this restoration process for the primary storage platform. The procedure is similar to that of restoring a single primary disk, but a larger number of devices is included in the reverse resynchronization request.

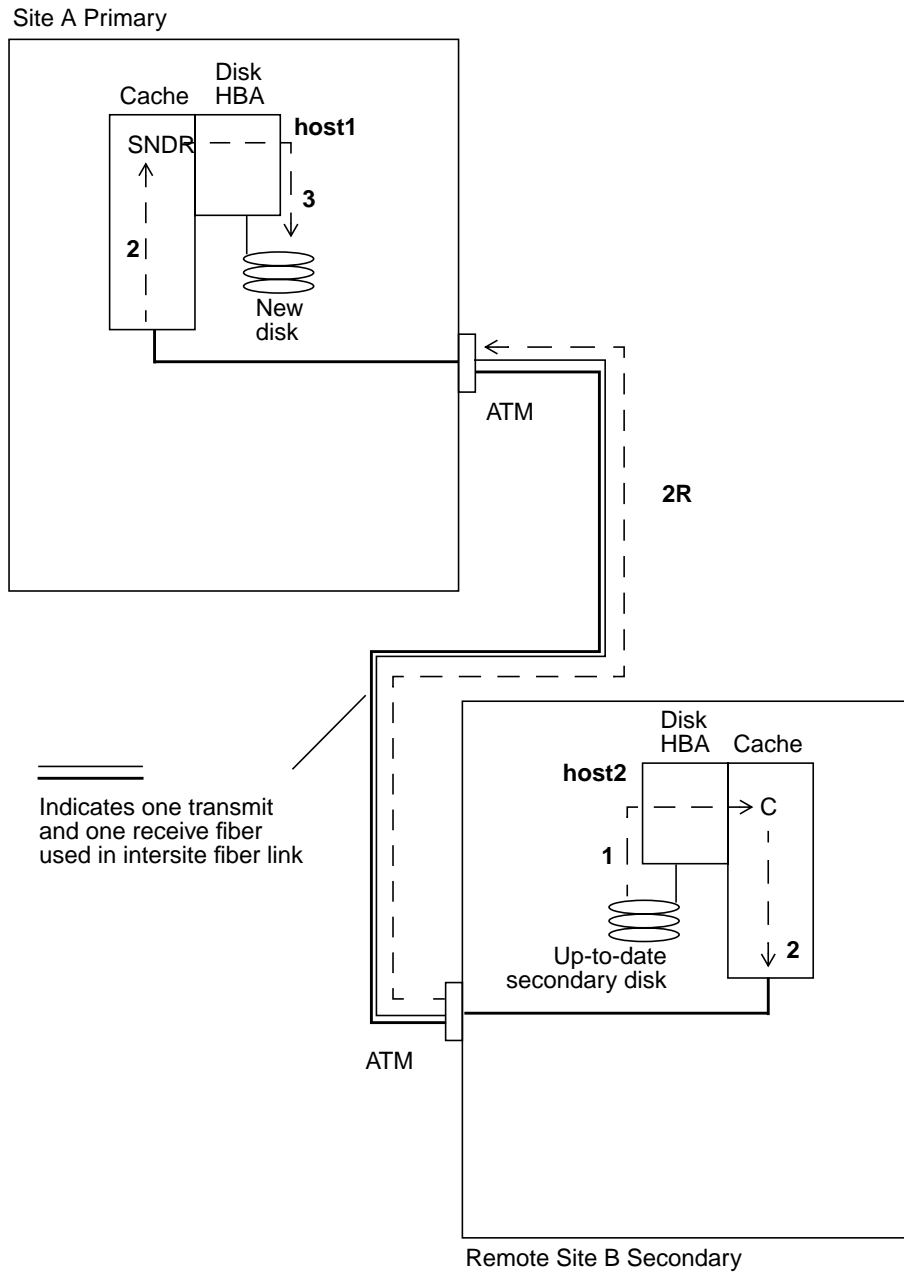


FIGURE 3-8 Restoring Primary from Secondary

Temporarily Stopping Remote Replication

Occasionally it is convenient to temporarily stop remote replication between primary and secondary volumes and still keep track of any changes that occur during the suspension. Establishing point-in-time copies for backups or report generation are likely cases. The `rdcadm -l` command provides such control at the logical volume level.

While stopped, SNDR tracks new writes to both the primary and secondary volumes in their respective scoreboards. You can later perform an update resynchronization to restore remotely replicated images and start active dual copy operations.

The direction of resynchronization determines which updates are applied and which updates are overwritten. Typically, the latest contents are taken from the primary site volume, used to incrementally update the secondary volume, and replace any secondary writes that occurred during the stopped period. A reverse resynchronization treats the secondary volume as having the desired contents and updates the primary volume accordingly.

- To stop remote replication and activate scoreboard logging, type:

```
host1# rdcadm -l
```

- To start remote replication and re-synchronize the secondary volumes on `host2` from the primary volumes on `host1`, type:

```
host1# rdcadm -u
```

- To reverse the resynchronization, using the secondary volume on `host2` as the source to overwrite the primary volume, type:

```
host1# rdcadm -u -r
```

- You can also start remote replication by requesting a full volume copy against the stopped devices. To fully resynchronize the secondary volumes on `host2` using `host1` volumes as the primary, ignoring any incremental scoreboards, type:

```
host1# rdcadm -m -f config_file
```


- To re-synchronize the primary volumes on `host1` using `host2` volumes as the source, ignoring any incremental scoreboards, type:

```
host1# rdcadm -m -r -f config_file
```

Disabling Remote Replication

Unlike stopping SNDR replication with the `rdcadm -l` command, disabling SNDR breaks the connection between primary and secondary volumes and discards any scoreboards.



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

Following the SNDR disable, a full synchronization (full volume copy) will be necessary to re-establish the SNDR relationship and ensure that the volumes' contents match. See “Establishing Replicated Images for the First Time” on page 4.

To permanently terminate SNDR remote replication and SNDR resynchronization services, type:

```
host1# rdcadm -d
```

Note – Use the `rdcadm -l` command instead of the `rdcadm -d` command when you want to only temporarily stop SNDR services.

Mirrored Partitions for Multiple Copies

You can combine SNDR with volumes from Sun StorEdge Instant Image™ software to create multiple copies of each device.

Glossary

- disk** Host view of the logical storage device. Disks are 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. SNDR options are performed on discrete logical volumes rather than on physical HDAs.
- health monitor** Health monitoring 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.
- local mirrors** Also known as mirrored partitions or RAID 1, which are established in a single system and do not rely on SNDR software or hardware to maintain local synchronization. See *mirrored partitions*.
- logging** Method for tracking disk 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. SNDR uses these logs (scoreboards) to re-establish matching sets through an optimized update synchronization rather than 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 and, in the event of a failure, a read is attempted on each set member until a successful read is accomplished.
- mirroring** Process of maintaining two or more identical images of a designated disk volume.
- partitions** UNIX device names of the form `/dev/rdisk/cnt.ndnsn` associated with the host disks on a Solaris system.
- primary** System or volume being accessed directly by the host application.
- RAID** Redundant Arrays of Independent Disks.

rollback	
synchronization	Resynchronization operation that discards any blocks modified during recovery rehearsals.
scoreboard log	Special bitmap that tracks hits to a disk area rather than a running log of each I/O event. It is a method for tracking disk 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. SNDR uses these scoreboard logs to re-establish matching sets through an optimized update synchronization rather than a complete volume-to-volume copy.
secondary	Remote counterpart of a primary system or volume where copies are destined.
SNDR	Network Data Replicator.
SNDR log	See <i>scoreboard</i> .
source device	Device used as the originating device in the remote copy.
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	Resynchronization operation that copies only the blocks modified since an SNDR interruption.
volume	See <i>disk</i> .

Index

A

- alternate fiber, 3-13
- asynchronous mirroring, 1-3
- auto sychronization, 2-16
- auto sychronization value, 2-16

C

- commands, 2-5
 - copy primary to secondary, 2-8
 - disable replication, 2-14
 - display auto sychronization value, 2-16
 - display list of configured sets, 2-17
 - enable RDC, 2-15
 - enable scoreboard logging, 2-15
 - refresh primary from secondary, 2-12
 - restore primary from secondary, 2-11
 - resume replication, 2-10
 - suspend replication, 2-9
 - toggle auto sychronization, 2-16
- configuration files
 - /etc/opt/SUNWrdc/rdc.cf, 2-2
- copy
 - primary to secondary, 2-7, 2-8
 - volume to volume, 3-6

D

- devices
 - source, GL-2
 - target, GL-2

- disable replication, 2-14
- disaster recovery rehearsals, 3-14
- disks, GL-1
 - restoring, 3-11
 - failed secondary, 3-12
 - restoring failed primary, 3-17
- display auto sychronization value, 2-16

E

- enable
 - RDC, 2-15
 - scoreboard logging, 2-15
- examples
 - scmadm -S, 2-19, 2-21

F

- failed primary disk
 - restoring, 3-17
- failed secondary disk, 3-12
- failures
 - primary site, 3-17
 - secondary site, 3-11
- fiber
 - alternate, 3-13
- full synchronization, 1-3, 3-6

H

HDA, GL-1
health monitor, GL-1

I

Instant Image
Storage Volume driver (sv), 3-3
svadm, 3-3
interruptions, 3-7
primary logging, 3-7
resynchronizing after, 3-7

L

local
mirrors, GL-1
log, scoreboard, GL-2
logging, 1-4, GL-1
secondary, 3-16

M

mirrored partitions, GL-1
mirroring, GL-1
asynchronous, 1-3
establishing images, 3-4
local, GL-1
synchronous, 1-4
monitoring operations, 3-2
mutual backup, 1-3, 3-2

O

overview, 1-2

P

partitions, GL-1
mirrored, GL-1
pass-through reads, 3-19
primary, 1-3, GL-1

site disaster, recovery, 3-23
site failures, 3-17
updating site, 3-23

R

RAID, GL-1
RDC terminology, 1-3
rdc.cf, 2-2
reads, pass-through, 3-19
recovery, primary site disaster, 3-23
refreshing
primary from secondary, 2-12
rehearsals
disaster recovery, 3-14
recovery from secondary volumes, 3-14
rolling back updates, 3-17
secondary updates, 3-15
takeover, 1-4, GL-2
replication
suspend, 2-9
restoring
failed primary disk, 3-17
failed secondary disk, 3-12
primary from secondary, 2-10, 2-11
primary from secondary, complete, 3-26
secondary disk, 3-11
secondary disk, restoring failed, 3-12
restoring primary from secondary, 3-27
resume replication, 2-10
resuming mirroring, 2-9
reverse
full synchronization, 3-22
update, 3-25
update synchronization, 3-25
rollback synchronization, 1-3, GL-2
rolling back updates, 3-17
rolling disaster, 1-3

S

scmadm -S, 2-7
example, 2-19, 2-21
options, 2-18
scoreboard log, GL-2

- scoreboard logging, 2-15
- secondary, 1-4, GL-2
 - logging, 3-16
 - restoring primary from, 3-27
 - site failures, 3-11
 - updates during rehearsal, 3-15
- sites
 - primary, 3-2
 - secondary, 3-2
- source device, GL-2
- Storage Volume driver
 - configuration file sv.cf, 3-3
- suspend replication, 2-9
- suspending mirroring, 2-8
- synchronization operations, 1-4, 3-25
 - full, 1-3, 3-6
 - reverse full, 3-22
 - rollback, 1-3, GL-2
 - update, 1-4, 3-9, GL-2
- synchronous mirroring, 1-4

T

- takeover rehearsal, 1-4, GL-2
- target device, GL-2
- terminology
 - asynchronous mirroring, 1-3
 - full synchronization, 1-3
 - logging, 1-4
 - mutual backup, 1-3
 - primary, 1-3
 - RDC, 1-3
 - rollback synchronization, 1-3
 - rolling disaster, 1-3
 - secondary, 1-4
 - synchronization, 1-4
 - synchronous mirroring, 1-4
 - update synchronization, 1-4

U

- update synchronization, 1-4, 3-9, GL-2
- updating primary site, 3-23

V

- VOLSER, GL-2
- volumes
 - see disk, GL-2
 - selecting, 3-2
- volume-to-volume copy, 3-6

