



Sun StorEdge™ Availability Suite 3.1 Remote Mirror Software Administration and Operations Guide

Sun Microsystems, Inc.
4150 Network Circle
Santa Clara, CA 95054 U.S.A.
650-960-1300

Part No. 816-4415-10
June 2002, [Revision A](#)

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

Copyright 2002 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, California 95054, U.S.A. All rights reserved.

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

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

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

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

Sun, Sun Microsystems, the Sun logo, AnswerBook2, docs.sun.com, Sun StorEdge, SunATM, SunSolve, Sun Fire, Ultra, Netra, Java, and Solaris are trademarks or registered trademarks of Sun Microsystems, Inc. in the U.S. and in other countries.

All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the U.S. and in other countries. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc.

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.

Use, duplication, or disclosure by the U.S. Government is subject to restrictions set forth in the Sun Microsystems, Inc. license agreements and as provided in DFARS 227.7202-1(a) and 227.7202-3(a) (1995), DFARS 252.227-7013(c)(1)(ii) (Oct. 1998), FAR 12.212(a) (1995), FAR 52.227-19, or FAR 52.227-14 (ALT III), as applicable.

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 2002 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, California 95054, Etats-Unis. Tous droits réservés.

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

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

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

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

Sun, Sun Microsystems, le logo Sun, AnswerBook2, docs.sun.com, Sun StorEdge, SunATM, SunSolve, Sun Fire, Ultra, Netra, Java, et Solaris sont des marques de fabrique ou des marques déposées de Sun Microsystems, Inc. aux Etats-Unis et dans d'autres pays.

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

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 licences de Sun qui mettent en place l'interface d'utilisation graphique OPEN LOOK et qui en outre se conforment aux licences écrites de Sun.

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



Contents

Preface xiii

1. Introduction to the Remote Mirror Software 1

Overview 2

Sun StorEdge Software Architecture 3

About Replicating File Systems 4

Working with Volume Sets 5

 About Volume Sizes 5

 About Bitmap Volume Size 5

 What Kind of Volumes Can I Use? 6

 Using Raw Partitions and Volumes 6

 How Many Volumes Can I Configure? 6

 Using I/O Groups 7

 Grouping Volume Sets in a Cluster Environment 8

Using Volume Manager Software 8

Network Protocols and TCP/IP Connection 9

Using the Software in a Firewall Environment 9

One-to-Many, Many-to-One, and Multihop Volume Sets 10

Order-Dependent Writes and Volume Set Grouping 11

Mutual Backup 12

Features 13

2. Replication and Synchronization Modes 15

Replication Modes 16

Synchronous Replication 16

Asynchronous Replication 17

Synchronization Modes 18

Full Synchronization 19

Update Resynchronization 21

Reverse Full Synchronization 23

Reverse Update Synchronization 24

Logging 26

Choosing the Resynchronization Type: Update or Full 27

Choosing Automatic or Manual Resynchronization 28

Using the Remote Mirror Software With the Point-in-Time Copy Software 29

Autosynchronization With the Point-in-Time Copy Software 30

Example Replication Scenarios 31

One-to-Many Replication 32

Many-to-One Replication 33

Multihop Replication 35

3. Preparing to Use the Software 37

Increasing the Default Number of Volumes Allowed 38

Using More Than 64 Volume Sets 39

Increasing the Storage Volume Device Limit 40

▼ To Increase the Storage Volume Limit 40

Bitmap Volumes 41

Bitmap Size Requirements 41

Customizing Volume Sets	42
Restricted Access When Creating and Configuring Sun StorEdge Volume Sets	42
Using a Volume Set File	43
Commands and I/O Group Operations	45
Failed Operations in an I/O Group	45
Commands and Sun Cluster Operations	45
4. Using the Remote Mirror Software	47
Example Host and Volume Names for This Chapter	48
Using the Storage Volume Interface	49
Monitoring Remote Mirror Software Operations	49
Initial Steps for Using the Software	50
Step 1: Enabling the Volume Sets	50
▼ To Enable the Volume Sets	50
Step 2: Establishing Volume Copies for the First Time	51
▼ To Fully Synchronize the Remote Volumes While Applying New Updates	52
Step 3: Updating the Secondary Volume	53
▼ To Resynchronize the Primary and Secondary Volumes	53
What Happens If a Network Link Fails?	54
When Not to Resynchronize Volumes	54
Autosynchronization	55
Manually Resynchronizing Volumes After a Network Link Failure or Other Interruption	55
Rehearsing Disaster Recovery	56
▼ To Rehearse a Primary Volume or Site Failure	56
▼ To Rehearse a Secondary Volume or Site Failure	58
About Primary Volume Failures	60

Recovering From a Primary Site Disaster	61
Restoring a Primary Site From the Secondary Site	61
Disabling Remote Replication	62
Swapping the Remote Mirror Hosts	63
Sample <code>rdc.cf</code> File Used in This Example	64
▼ To Disable the Software at Site-A	65
▼ To Change the Site-B Secondary Host to the Primary Host	66
5. The <code>sndradm</code> Command	67
Log File	68
<code>sndradm</code> Syntax Summary	69
Command Switches	70
Command Options	74
Which Host Do I Issue Commands From?	76
Enabling and Disabling Volume Sets	77
Enabling Volume Sets	78
After You Enable Volume Sets	78
Using a Volume Set Name	78
Disabling Volume Sets	80
Adding and Deleting Point-in-Time Copy Software Volumes	81
Stopping Replication and Starting Logging	83
Synchronizing and Updating Volume Sets	84
If an Initial Synchronization is Interrupted	85
Performing a Reverse Copy or Update After a Primary Volume or Network Link Failure	85
Starting a Full Synchronization (Copy Between Volumes)	86
Starting an Update Resynchronization (Update Between Volumes)	88
Waiting for Synchronization to Finish Executing	90

Reconfiguring Or Modifying a Volume Set	91
Changing the Bitmap Volume in a Volume Set	92
Moving a Volume Set to a Different I/O Group	93
▼ To Remove Volume Sets From an I/O Group	94
▼ To Move Volume Sets to Another I/O Group	95
Changing the Replication Mode of a Volume Set	96
▼ To Remove Volume Sets from an I/O Group and Change the Replication Mode	97
Updating the Volume Set Information	98
Updating the Disk Cluster Tag Name	99
Displaying Volume Set Information	100
Displaying Volume Set and I/O Group Status	101
Displaying a List of Current Volume Sets and I/O Group Names	102
Displaying Link Status	103
Enabling or Disabling the Autosynchronization Feature	104
Setting the Asynchronous Queue	105
Tuning the Asynchronous Queue	106
Asynchronous Mode and the Queue	107
Displaying the Queue Size	108
▼ To Display the Current Queue Size	108
Sample <code>kstat</code> Output For a Correctly-Sized Queue	109
Sample <code>kstat</code> Output for an Incorrectly-Sized Queue	110
A. Sun StorEdge Software and Storage Cache Statistics	111
Remote Mirror and Point-in-Time Copy Software Statistics: The <code>dsstat</code> Command	112
More Remote Mirror Software and Storage Cache Statistics: The <code>scmadm</code> Command	119
Example <code>scmadm -S -M</code> Screen Display	122

Example `scmadm -S` Screen Display 124

Glossary 127

Figures

FIGURE 1-1	Remote Mirror Software Architecture	4
FIGURE 1-2	Mutual Backup	12
FIGURE 2-1	Full Synchronization (Volume-to-Volume Copy)	20
FIGURE 2-2	Update Resynchronization of a Secondary Volume Set	22
FIGURE 2-3	Reverse Full Synchronization	23
FIGURE 2-4	Reverse Update Synchronization	25
FIGURE 2-5	One-to-Many Volume Set Replication	33
FIGURE 2-6	Many-to-One Volume Set Replication	34
FIGURE 2-7	Multihop Volume Set Replication	35

Tables

TABLE 1-1	Remote Mirror Software Features	13
TABLE 3-1	Volume Set File Format Fields	44
TABLE 4-1	Example Names and Devices Used in This Chapter	48
TABLE 4-2	Example Names and Devices Used in This Procedure	64
TABLE 5-1	Command Switches Sorted By Function	70
TABLE 5-2	Command Switches in Alphabetical Order	72
TABLE 5-3	Command Options	74
TABLE 5-4	Which Host to Issue Commands From	76
TABLE A-1	<code>dsstat</code> Options	113
TABLE A-2	<code>dsstat</code> Display Fields	117
TABLE A-3	Keys Used in <code>scmadm</code> Display	120
TABLE A-4	<code>scmadm -S</code> Options	121
TABLE A-5	Fields for Remote Mirror Software Screen Display	122
TABLE A-6	<code>scmadm -S</code> Fields	124

Preface

The *Sun StorEdge Availability Suite 3.1 Remote Mirror Software Administration and Operations Guide* describes the Sun StorEdge™ remote mirror software and its use. This guide is intended for system administrators who have experience with the Solaris™ operating environment and related disk storage systems.

Note – This software was previously known as the Sun StorEdge Network Data Replicator (Sun SNDR) software.

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 Availability Suite 3.1 Remote Mirror Software Installation Guide*
- *Sun StorEdge Availability Suite 3.1 Remote Mirror Software Release Notes*

To operate correctly, the remote mirror software must be installed as described in the installation and release documentation.

How This Book Is Organized

[Chapter 1](#) provides general information about the remote mirror software.

[Chapter 2](#) describes the remote mirror software architecture and replication modes.

[Chapter 3](#) describes how to configure and prepare to use the software.

[Chapter 4](#) explains the remote mirror software operations including startup, synchronization, and recovery rehearsals.

[Chapter 5](#) describes the remote mirror software `sndradm` command.

[Appendix A](#) describes the `dsstat` and `scmadm` reporting and monitoring commands.

Using UNIX Commands

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

See one or more of the following for this information:

- *Solaris Handbook for Sun Peripherals*
- AnswerBook2[™] online documentation for the Solaris operating environment
- Other software documentation that you received with your system

Typographic 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. <code>% You have mail.</code>
AaBbCc123	What you type, when contrasted with on-screen computer output.	<code>% su</code> 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 superuser to do this. To delete a file, type <code>rm filename</code> .
[]	In syntax, brackets indicate that an argument is optional.	<code>scmadm [-d sec] [-r n[:n],[n]...] [-z]</code>
{ arg arg }	In syntax, braces and pipes indicate that one of the arguments must be specified.	<code>sndradm -R b {p s}</code>

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

Application	Title	Part Number
Man pages	sndradm iiadm dsstat kstat scmadm svadm	N/A
Latest release information	<i>Sun StorEdge Availability Suite 3.1 Remote Mirror Software Release Notes</i>	816-4414
	<i>Sun Cluster 3.0 and Sun StorEdge Software Release Note Supplement</i>	816-5128
	<i>Sun StorEdge Availability Suite 3.1 Point-In-Time Copy Software Release Notes</i>	816-4314
Installation and User	<i>Sun StorEdge Availability Suite 3.1 Remote Mirror Software Installation Guide</i>	816-4413
	<i>Sun StorEdge Availability Suite 3.1 Point-In-Time Copy Software Installation Guide</i>	816-4312
	<i>SunATM 3.0 Installation and User's Guide</i>	805-0331
	<i>SunATM 4.0 Installation and User's Guide</i>	805-6552
	<i>Sun Gigabit Ethernet FC-AL/P Combination Adapter Installation Guide</i>	806-2385
	<i>Sun Gigabit Ethernet/S 2.0 Adapter Installation and User's Guide</i>	805-2784
	<i>Sun Gigabit Ethernet/P 2.0 Adapter Installation and User's Guide</i>	805-2785
	<i>Sun Enterprise 10000 InterDomain Networks User Guide</i>	806-4131
System administration	<i>Sun StorEdge Availability Suite 3.1 Point-In-Time Copy Software Administration and Operations Guide</i>	816-4314
	<i>TCP/IP and Data Communications Administration Guide</i>	805-4003
Cluster	<i>Sun Cluster 3.0 and Sun StorEdge Software Integration Guide</i>	816-5127
Configuration	<i>Sun Enterprise 10000 InterDomain Network Configuration Guide</i>	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/>

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 number (816-4415-10) of your document in the subject line of your email.

Introduction to the Remote Mirror Software

Note – This software was previously known as the Sun StorEdge Network Data Replicator (Sun SNDR) software.

This chapter describes the following:

- [“Overview” on page 2](#)
- [“Sun StorEdge Software Architecture” on page 3](#)
- [“About Replicating File Systems” on page 4](#)
- [“Working with Volume Sets” on page 5](#)
- [“Using Volume Manager Software” on page 8](#)
- [“Network Protocols and TCP/IP Connection” on page 9](#)
- [“Using the Software in a Firewall Environment” on page 9](#)
- [“One-to-Many, Many-to-One, and Multihop Volume Sets” on page 10](#)
- [“Order-Dependent Writes and Volume Set Grouping” on page 11](#)
- [“Mutual Backup” on page 12](#)
- [“Features” on page 13](#)

Overview

The Sun StorEdge Availability Suite 3.1 remote mirror software is a volume-level replication facility for the Sun Solaris™ operating environment. You can use this software to replicate disk volumes between physically separate primary and secondary hosts in real time.

As part of a disaster recovery and business continuance plan, the remote mirror software enables you to keep up-to-date copies of critical data at remote sites. You can also rehearse your data recovery strategy to failover data to remote sites. Later, you can write any data changes that occurred to the original disk that “failed”. To transport data, the remote mirror software uses any Sun network adapter that supports TCP/IP.

The remote mirror software is active while your applications are accessing the data volumes and continually replicates the data to remote sites. The software operates at the volume level on storage devices attached to one or more hosts.

You can also manually update the data on the secondary site volume by issuing a command to *synchronize* the primary and secondary site volumes. You can restore data from the secondary volume to the primary volume by issuing a command to *reverse synchronize* the volumes. (Reverse synchronization is also known as performing a *reverse update*. Updates from the primary site to the secondary site are also known as *forward resynchronizations*.)

Sun StorEdge Software Architecture


The core remote mirror software code is a kernel module that interfaces to the network storage control module (`nsctl`) framework. The software is configurable on any device that is accessible through the `nsctl` framework. The `sndradm` CLI is the external user interface used to manage the remote mirror software.

[FIGURE 1-1](#) shows the relationship between the Storage Volume driver, the remote mirror software, and the rest of the `nsctl` framework.

The I/O commands targeted to remote mirror volumes are intercepted by the SV driver and routed through the Sun StorEdge I/O stack before being passed on to the storage device driver or the volume manager. The SV driver is a very thin layer in the I/O stack and operates by interposing onto the entry points to the underlying device driver. The I/O commands originating in user space are intercepted at the top of the Sun StorEdge I/O stack. The SV driver routes them through the stack and feeds them back to the storage device driver or volume manager at the bottom of the stack. Data also flows in the opposite direction, from the storage back to user space.

Because the remote mirror software is located at the top of the stack before the point-in-time copy software, the remote mirror volume sets must be placed in logging mode (`sndradm -l`) before performing point-in-time enable, copy, update or reset operations on any volume that is also part of a remote mirror volume set.

If the volume set is not in logging mode, the point-in-time copy operation fails and the remote mirror software reports that the operation is denied. See also [“Using the Remote Mirror Software With the Point-in-Time Copy Software”](#) on page 29.



1. I/O commands and data enter and exit the remote mirror software through the Sun StorEdge Storage Volume (SV) driver software.
2. Mediated by `nsctl`, the data flows through remote mirror software (and optionally the point-in-time copy software) and the Storage Device Block Cache (SDBC) drivers to its destination on the storage array or in user space.

FIGURE 1-1 Remote Mirror Software Architecture

About Replicating File Systems

The remote mirror software is not a file or file system replicator. It is a volume replicator. When you replicate a volume on the primary Site-A to a volume on the secondary Site-B, Site-B receives an exact bit-for-bit copy. Make sure that any Site-A file systems you want to replicate are mounted.

When replicating, the primary host file system is mounted. Do not mount the file system on Site-B until you are ready to fail over or write to that site's volume. Changes appear on a replicated file system volume only after mounting the volume.

Working with Volume Sets

Note – In this document, a volume is a raw disk partition or a volume created by a volume manager.

The remote mirror software replicates data to volume sets that you define. A volume set consists of a primary volume residing on a local (primary) host and secondary volume residing on a remote (secondary) host. The volume set also includes a bitmap volume on each host to track write operations and differences between the volumes. See “[Customizing Volume Sets](#)” on page 42.

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 remote mirror 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.

About Volume Sizes

The size of the secondary site volume must be equal to or greater than the corresponding primary site volume. If you initiate a resynchronization on a volume set where the secondary volume is smaller than the primary, the software fails with an error.

About Bitmap Volume Size

Each volume to be replicated requires a bitmap volume in the volume set. 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.

What Kind of Volumes Can I Use?

Note – See [“Using Volume Manager Software”](#) on page 8.

Volumes are also defined here as logical volumes that can be linear, striped, or RAID volumes. You can create the logical volumes by using the Sun Solstice DiskSuite™, Solaris Volume Manager (available in Solaris 9), or VERITAS Volume Manager software.

You can use RAID volumes as part of your remote mirror software strategy. Volumes can be any RAID level. The RAID levels of volumes in a volume set do not have to match.



Caution – You cannot use the remote mirror software to replicate the boot device or any system level partitions like `/var` and `/usr`.

Using Raw Partitions and Volumes

When selecting a volume to be used in a volume set (including the [configuration location](#)), ensure that volume does not contain disk label private regions (for example, slice 2 on a Solaris operating environment-formatted volume). The disk label region is contained in the first sectors of a disk. The safest method is to ensure that cylinder 0 is not part of any logical volume that is replicated.



Caution – When the volume in a volume set is a raw partition, the partition *must not* include the cylinder that contains the label for the disks. On Sun disks, this is cylinder 0.

How Many Volumes Can I Configure?

By default, the remote mirror and point-in-time copy software support the configuration of 1024 volumes and 64 Mbytes for caching. You can increase both amounts if system resources allow. The number of volumes allowed is divided between both software products. For example, if you use the point-in-time software only, you can have 341 volume sets, each consisting of master, shadow, and bitmap volumes.

For more information, see [“Increasing the Default Number of Volumes Allowed”](#) on page 38.



Caution – In a clustered environment, only one system administrator or root user at a time is allowed to create and configure Sun StorEdge volume sets. This restriction helps avoid creating an inconsistent Sun StorEdge Availability Suite volume set configuration.

Two administrators should not be writing to the Sun StorEdge Availability Suite software configuration at the same time. The operations that access the configuration include but are not limited to:

- Creating and deleting volume sets
- Adding and removing volume sets from I/O groups
- Assigning new bitmap volumes to a volume set
- Updating the disk device group or resource name
- Any operation that changes the Sun StorEdge Availability Suite software and related volume set configuration

Using I/O Groups

The software enables you to group volume sets in *I/O groups*. You can assign specific volume sets to an I/O 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.

An I/O group is a collection of 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.

By using an I/O group, you can issue a remote mirror software command that is executed on every member of the group, enabling volume sets to be controlled as a single unit.

I/O group operations are atomic. The change from replicating mode to logging mode is guaranteed to occur on every set in an I/O group and to fail on all the sets if it fails on a single set in the group.

The 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. See [“Order-Dependent Writes and Volume Set Grouping”](#) on page 11.

Note – The I/O group concept does not matter for synchronous replication; that is, write-ordering is preserved among those volume sets configured as `sync`.

The auto-resynchronization feature supports the I/O grouping concept. It allows the feature to be enabled or disabled on a per-group basis and controls the resynchronization operation atomically on the group.

I/O grouping has an adverse affect on the remote mirror software asynchronous operation, as I/O flushing is reduced to a single thread. In this case, consider the size of the data to be transferred since all I/O will be routed through a single queue.

Grouping Volume Sets in a Cluster Environment

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

- **Not cluster-tolerant** with the Sun Cluster 3.0 environment
- **Cluster-aware** in the Sun Cluster 3.0 Update 1 and Update 2 environments and provides high-availability for the Sun StorEdge software

See the *Sun Cluster 3.0 and Sun StorEdge Software Integration Guide* for more information about configuring the Sun StorEdge Availability Suite software in a Sun Cluster environment.

Using Volume Manager Software

The remote mirror and point-in-time copy software does not support metatrans devices (also known as trans metadevices) created by the Sun Solstice DiskSuite or Solaris Volume Manager software.

Metatrans devices are intended for use with UNIX file systems (`ufs`) without using any other layered services. Use the `ufs logging mount` option as an alternative to the use of metatrans devices. The Sun StorEdge Availability Suite software supports `ufs logging` which should be used when available instead of metatrans devices.

Network Protocols and TCP/IP Connection

The remote mirror software requires a TCP/IP connection between the primary and secondary hosts. A dedicated TCP/IP link is not required.

Although the remote mirror software is most likely to be used with SunATM link-level interfaces, the remote mirror software design is link-neutral, meaning that it can use any Sun network adapter that supports TCP/IP.

Each host must have the proper ATM or Ethernet hardware installed to support the TCP/IP link. The remote mirror software operates over any TCP/IP networking technology but has been qualified only on 10, 100, and 1000 Mbit Ethernet and SunATM155 and SunATM622 technologies.

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*.

Using the Software in a Firewall Environment

You can use the software with any enhanced network product that Sun offers. Port 121 is the default port for use by the remote mirror `sndrd` daemon. To change the port number, edit the `/etc/services` file using a text editor. See the *Sun StorEdge Availability Suite 3.1 Remote Mirror Software Installation Guide* for more information.

If you change the port number, you must change it on all hosts (that is, primary and secondary hosts, all hosts in one-to-many, many-to-one, and multihop configurations).

You can use IP masquerading. The software uses connection-oriented service (COTS) based remote procedure calls (RPC; that is, TCP). The RPC part of the packet does not contain any address information.

One-to-Many, Many-to-One, and Multihop Volume Sets

Note – See [“One-to-Many Replication” on page 32](#), [“Many-to-One Replication” on page 33](#) and [“Multihop Replication” on page 35](#) for more information about these scenarios.

The remote mirror software enables you to create one-to-many, many-to-one, and multihop volume sets.

One-to-many replication enables you to replicate data from one primary volume to many secondary volumes residing on one or more hosts. One primary and each secondary site volume is a single volume set. For example, with one primary and three secondary host volumes, you need to configure three volume sets: primary A and secondary B1, primary A and secondary B2, and primary A and secondary B3.

Many-to-one replication enables you to replicate volumes across more than two hosts through more than one network connection. The software supports the replication of volumes located on many different hosts to volumes on a single host. The terminology differs from the one-to-many configuration terminology, where the *one* and the *many* referred to are volumes.

Multihop replication indicates that the secondary host volume of one volume set acts as the “primary” host volume of another volume set (it is still the secondary volume of the first volume set). In the case of one primary host volume A and one secondary host volume B, the secondary host volume B appears as primary host volume A1 to the secondary host volume B1.

Order-Dependent Writes and Volume Set Grouping

Write ordering is maintained for *groups* of asynchronously replicating volume sets. (The general definition of write ordering is that write operations directed to the target volume occur in the same order as write operations to the source.) The group of target volumes is a copy of the group of source volumes.

This feature can be valuable in those cases where you can avoid application requirements that limit operations. For example, a database application might limit partition sizes to no greater than 2 Gbytes. In this case, you might group volume sets to create a virtual large “volume” that preserves write operations. Otherwise, you might risk having inconsistent data by trying to update volume sets individually instead of as a group.

When an application has multiple logical volumes assigned, application data integrity can be maintained by one of the following:

- Specifying that all remote mirror software volumes associated with that application are in `sync` mode
- Using the Sun StorEdge Availability Suite 3.1 point-in-time copy software to make periodic recoverable point-in-time copies

If you use point-in-time copy software, the remote point-in-time snapshot is taken while the application is in the recoverable state. For example, most database applications allow for a hot backup. If a remote point-in-time copy were made of the entire replicated database while the primary was in hot backup mode, then a consistent remote database is available by using the point-in-time copy and the log files taken while the database was in hot backup mode.

Note – You must place the remote mirror volume set in logging mode for the point-in-time copy software to successfully perform an enable, copy, update, or reset operation on a remote mirror volume. If the volume set is not in logging mode, the point-in-time copy operation fails and the remote mirror software reports that the operation is denied.

Mutual Backup

Typically, volumes are replicated from local Site-A to remote Site-B. However, as applications are geographically distributed, a storage system at remote Site-B can be a remote volume backup to local Site-A and also be a direct storage resource for applications on Host-B. Under these circumstances, you can replicate Host-B volumes to Site-A.

This reciprocal backup arrangement supported by the remote mirror software is known as *mutual backup* or *dual backup*.

With mutual backup, the remote mirror 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.

In this case, you need to configure two *unique* volume sets for each site. For example,

- Volume set 1 contains primary Site-A volumes replicated to secondary Site-B
- Volume set 2 contains primary Site-B volumes replicated to secondary Site-A

See [FIGURE 1-2](#).




FIGURE 1-2 Mutual Backup

Features

[TABLE 1-1](#) describes the remote mirror software features.

TABLE 1-1 Remote Mirror Software Features

This Feature. . .	Enables You To. . .
Active logging	Continue logging operations whenever the remote mirror software is disabled or interrupted.
Backward compatible <code>rdcadm</code> command	Maintain the integrity of scripts you have written and want to continue using. The version 3.1 <code>sndradm</code> command is linked to the <code>rdcadm</code> command and is generally compatible with remote mirror software version 2.0 <code>rdcadm</code> command.
Disaster rehearsal rollback	Simulate a disaster without committing data changes to volumes.
Groups of remote mirror volume sets	<p>Control the remote mirror software volume sets simultaneously by grouping them. This feature is essential in installations requiring you to maintain consistent contents of a group of volumes.</p> <p>An I/O group is a collection of remote mirror 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.</p> <p>The remote mirror 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.</p>
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, Many-to-One, and Multihop Volume Sets” on page 10.
Mutual backup	Concurrently transmit to and receive data copies from a remote volume. Also known as a <i>bilateral relationship</i> .
One-to-many sets	Replicate data from one primary volume to many secondary volumes residing on one or more sites. 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, Many-to-One, and Multihop Volume Sets” on page 10.

TABLE 1-1 Remote Mirror 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 since the last synchronization
RAID support	Use RAID volumes as part of your remote mirror software strategy. Volumes can be any RAID level.
Security and Internet Protocols	Use version 3.1 securely and efficiently in the Sun Solaris 8 and 9 operating environments, which support Internet Security Protocol (IPsec). The Solaris 8 and 9 operating environments also support Internet Protocol Version 6 (IPv6). (The Solaris 7 operating environment does not support IPv6). Version 3.1 does not require the use of an <code>.rhosts</code> file. Place the host names of each host running the remote mirror Version 3.1 software in the <code>/etc/hosts</code> file on each machine.
Stop and start control of remote replication	Have fine control over replication at the logical volume level. You can include volumes containing critical data by specifying these volumes as part of a remote mirror software volume set. You can also exclude volume with noncritical data from remote mirror software options.
Sun StorEdge point-in-time copy 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 remote mirror 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 remote mirror software is configured).

Replication and Synchronization Modes

This chapter describes the following:

- [“Replication Modes” on page 16](#)
- [“Synchronization Modes” on page 18](#)
- [“Example Replication Scenarios” on page 31](#)

Replication Modes

The remote mirror software supports two modes of data replication:

- Synchronous replication
- Asynchronous replication

The replication mode is a user-selectable parameter for each remote mirror volume set. (Use the `sndradm enable` command and select the volume set's `sync` or `async` parameter. Use the `sndradm -R m` command to change the replication mode thereafter.) The 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 software replicates changes to the secondary volume, keeping the two volumes consistent.

In the event of planned or unplanned outages, the software maintains per-device bitmap volumes that are marked to indicate changed blocks with a granularity of 32 Kbytes per segment. This technique optimizes resynchronization by allowing the software to resynchronize only the blocks that have changed since the last synchronization.

Synchronous Replication

In synchronous mode, a write operation is not confirmed as complete until the remote volume has been updated. Synchronous mirroring forces the remote mirror software to wait until an acknowledgement of the receipt of the data is received from the secondary by the primary before returning to the application. The application is not acknowledged until the write at the secondary site is complete.

The advantages of synchronous replication are that it is more reliable and can help reduce the risk of data loss. One disadvantage might be an increase in write response time, especially for large data sets or long-distance replication (where write operations can incur additional latency because of the time required to transfer data and return acknowledgements).

Asynchronous Replication

In asynchronous mode, a write operation is confirmed as complete before the remote volume has been updated. Asynchronous replication allows the remote mirror software to return to the host as soon as the write has been completed on the primary volume and been placed on a per-volume queue for the secondary site. Subsequently, the secondary site receives the queued requests in the order that they were queued. Once the I/O has been completed at the secondary site and the bitmap has been updated to reflect the volume state, notification is sent to the primary.

The advantages of asynchronous replication are that it provides fast response and has the least impact on the response time of the primary application. A disadvantage is that data loss might occur at the secondary site after primary site or network failures.

For more information about asynchronous replication and adjusting the asynchronous queue, see [“Setting the Asynchronous Queue” on page 105](#) and [“Tuning the Asynchronous Queue” on page 106](#).

Synchronization Modes

When a volume set is enabled using the `sndradm -e` command, you must initially synchronize the primary and secondary volumes in the set. (Use the `sndradm -E` command if the volumes are already identical). Using one of the remote mirror synchronization modes ensures that the primary and secondary volumes contain the same data and that they are identical at a clearly defined time. Synchronization is driven by the software through the `sndradm` command and progresses to completion.

Once a volume set has been synchronized, the software ensures that the primary and secondary volumes contain the same data through *replication*. Replication is driven by user-layer application write operations. Remote mirror replication is an ongoing process.

The software synchronizes data in a *forward* (from the primary to the secondary) or *reverse* (from the secondary to the primary) direction. The software synchronizes data in five modes:

- [Full synchronization](#)
- [Update resynchronization](#)
- [Reverse full synchronization](#)
- [Reverse update synchronization](#)
- [Logging](#)

This section also describes the following topics:

- [“Choosing the Resynchronization Type: Update or Full” on page 27](#)
- [“Choosing Automatic or Manual Resynchronization” on page 28](#)
- [“Using the Remote Mirror Software With the Point-in-Time Copy Software” on page 29](#)
- [“Autosynchronization With the Point-in-Time Copy Software” on page 30](#)

Full Synchronization

Full synchronization starts 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. After the operation is complete, the remote mirror software maintains the normal replicating mode for the volume: either synchronous or asynchronous replication.

Note – The volumes can be made identical using other methods, not just a full synchronization. When network latencies justify it, you can perform the initial synchronization of a volume set by backing up a source or primary volume on magnetic tape on one site, then restoring the volume from the tape on the other site. During the intervening period (that is, the period between when the backup is completed and the restore is started), place the source or primary volume in logging mode.

Make sure that the backup copy is a physical copy (for example, by using the `dd(1M)` command) and not a logical copy (for example, one made using the `tar(1M)` or `cpio(1M)` commands). The copies must have identical blocks, not just identical files.

FIGURE 2-1 shows the full forward synchronization process.

1. The remote mirror 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 software transmits the disk blocks, with destaging instructions, over the connection to a cache region on the secondary system.
3. The software on the secondary system updates its remote volume and acknowledges the update to the primary system.




FIGURE 2-1 Full Synchronization (Volume-to-Volume Copy)

Update Resynchronization

During update resynchronization mirroring, the remote mirror software initiates replication of only the changed primary site volume data to the secondary site, based on the bitmap. Only the blocks marked dirty in the bitmap are copied to the target volume.

After the update is complete, the software maintains the normal replicating mode. The software can also be placed in logging mode. See [“Logging” on page 26](#).

The software resynchronizes the secondary volume from the primary volume. It updates the secondary volume according to the changes based on 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.

If a remote copy interruption lasts numerous hours and the updates are widespread, logging and update resynchronization provide diminishing returns. As time passes, the proportion of bits set to true in the bitmap volumes of a volume set might reach 100 percent. The overhead of logging, coupled with update resynchronization, must then be balanced against that of a full synchronization without intervening periods of logging.

Logging and update resynchronization serve as a built-in safety net should one of your replication processes be interrupted. The software monitors the network connections between the primary and secondary hosts. Link failures and remote system failures are detected by the transport interface and are passed to the remote mirror software.

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

1. The remote mirror software on `host1` examines a bitmap from the primary and secondary hosts for the remote mirror software-managed volumes affected by the interruption.
2. The 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 software on `host1` transmits the update blocks 3R to `host2` remote mirror software.
4. The software on `host2` refreshes its stale replicated image with the updated blocks and acknowledges the action to `host1`.
5. The software revises the bitmap to track the remote update.

All steps repeat until the remote replicated image is updated.




FIGURE 2-2 Update Resynchronization of a Secondary Volume Set

Reverse Full Synchronization

During reverse full synchronization, the remote mirror software replicates the volume data from the secondary site to the primary site. After you issue the `sndradm -m -r` command, the software starts a full reverse copy operation from the secondary volume to the primary volume. It also enables concurrent replication from the primary volume to the secondary volume; any new writes to the primary volume are also replicated to the secondary volume. This remote mirror feature can help reduce application downtime.

You can use the primary volume during reverse synchronization. The primary volume shows a consistent volume image of the latest data as soon as the reverse synchronization starts. If your application had been writing to the secondary volume as part of a failure or disaster rehearsal, you can move the application back to the primary volume when the reverse synchronization starts.

FIGURE 2-3 shows the full reverse synchronization process.

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




FIGURE 2-3 Reverse Full Synchronization

Reverse Update Synchronization

During reverse update synchronization mirroring, the remote mirror software compares the bitmaps between the primary and secondary sites and replicates only the changed blocks from the secondary site to the primary site.

It resynchronizes the primary volume from the secondary volume. It updates the primary volume according to the changes based on 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. This remote mirror feature can help reduce application downtime.

You can use the primary volume during reverse update synchronization. The primary volume shows a consistent volume image of the latest data as soon as the reverse update synchronization starts. If your application had been writing to the secondary volume as part of a failure or disaster rehearsal, you can move the application back to the primary volume when the reverse update synchronization starts.

FIGURE 2-4 shows a reverse update resynchronization from the secondary system to its primary system.

1. The remote mirror software on `host1` retrieves the secondary bitmap 1R from `host2` for one of the remote mirror software-managed volumes affected by the interruption.
2. The 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 software on `host2` transmits the updated blocks 3R to `host1` remote mirror software region of cache using the intersite link.
4. The software on `host1` refreshes its stale image with the updated blocks.
5. The software on `host1` revises the bitmap to track the remote update.

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

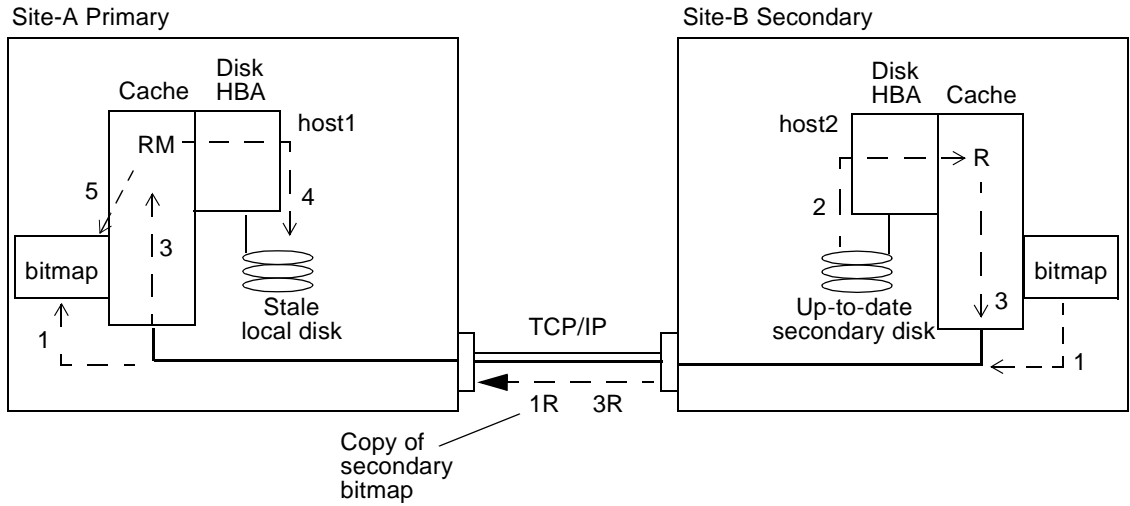


FIGURE 2-4 Reverse Update Synchronization

Logging

Note – With synchronous and asynchronous replication, the remote mirror software automatically switches to logging mode if a break in the network occurs or if the primary site is down.

During logging, the remote mirror software only updates the bitmaps at the primary site; no replication occurs. At a later time, the bitmaps at the primary and secondary sites are compared and the changed blocks in the primary site volume are mirrored by update resynchronization to the secondary site.

If all volume sets in an I/O group are replicating (that is, 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 I/O group will enter logging mode automatically. This scheme ensures that the secondary volumes will contain a valid copy of the data.

You can use logging to save on telecommunications or connection costs. The risk, however, is the costs incurred by increased data loss if the primary site is lost. If you lose the primary site, you do not have the data at the secondary site that was written to the primary site during logging.

You can also perform logging on the secondary site before a failover. You can then update the primary site using a reverse synchronization or reverse update synchronization command.

Note – To resume remote mirror software operations after using the `sndradm -l` logging 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.

Choosing the Resynchronization Type: Update or Full

Tip – 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.

An update resynchronization applies changes that occurred at the primary site to the secondary site while replication was suspended. The primary site can also be updated from the secondary, if desired.

A full synchronization is a complete disk-to-disk copy. This operation is the most time consuming of the synchronization operations. A full synchronization is typically only performed when the remote mirror software volume set is:

- Initially established
- Damaged due to a disaster
- Includes data integrity that is in question

If the integrity of the volume data is questioned, then the only way to get the volume to a synchronized set is to perform a full volume copy.

The volume set data might be in question if, for example, a double disk failure on a RAID 5 set occurred or if the remote mirror software is shut down manually and writes occurred to either the primary or secondary volumes without logging being active. The full copy can be performed from the primary to the secondary or, if appropriate, from the secondary to the primary.

Choosing Automatic or Manual Resynchronization

Note – [Chapter 5](#) describes the remote mirror software commands.

The remote mirror software provides two methods for synchronization after a scheduled or unscheduled link failure:

- Automatic (also known as autosynchronization), where synchronization occurs automatically when the link is reestablished (you need to set this feature for each volume set using the `sndradm -a` command)
- Manual, where synchronization requires that you perform an update or full resynchronization by issuing one of the `sndradm` commands

Autosynchronization is discouraged if the interruption is the warning of a larger rolling disaster. 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. The autosynchronization option is disabled by default for this reason. For example:

1. The primary and secondary volumes in a volume set are synchronized and then in replicating mode when a link failure is detected.
2. Both the primary and secondary volumes revert to logging mode and bits in the bitmap volumes are marked as dirty.
3. The link returns online and a reverse synchronization is started: the bitmap volumes are logically OR'ed.
4. The secondary volume fails.

The primary volume is now in an inconsistent state because it now is a mix of data written between the two volumes before the second (volume) failure and data which has not yet been synchronized. In this instance, the fact that one of the volumes is now unrecoverable makes it difficult to return both volumes to full integrity.

See also [“Using the Remote Mirror Software With the Point-in-Time Copy Software” on page 29](#). Before you start a resynchronization operation, ensure that you have an appropriate point-in-time copy of the target volume.

Using the Remote Mirror Software With the Point-in-Time Copy Software

Note – You must place the related remote mirror volume set in logging mode for the point-in-time copy software to successfully perform an enable, copy, update, or reset operation on a remote mirror volume. If the volume set is not in logging mode, the point-in-time copy operation fails and the remote mirror software reports that the operation is denied

To help ensure a high level of data integrity on both sites during normal operations or during data recovery, use the remote mirror software with the Sun StorEdge Availability Suite 3.1 point-in-time copy software.

Use the point-in-time copy software just before you perform a resynchronization, to help ensure that a consistent copy of data exists. If a failure occurs during the resynchronization, you have the point-in-time data copy to roll back to and you can start again when its safe to do so.

During the resynchronization process of updating the local and remote sites, the data on a secondary remote mirror 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 help ensure data integrity, use point-in-time copy software regularly to create a point-in-time copy of data at both sites. See the point-in-time copy documentation listed in [“Related Documentation” on page xvi](#).

Autosynchronization With the Point-in-Time Copy Software

Note – See the command description for “[Adding and Deleting Point-in-Time Copy Software Volumes](#)” on page 81.

The `/usr/opt/SUNWrdc/lib/sndrsyncd` daemon automates update resynchronization after a network link or machine failure. If the point-in-time copy software is also installed and you have added point-in-time copy volume groups, the daemon 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 remote mirror software becomes unavailable, the daemon attempts to execute the remote mirror software update commands to resynchronize all volume sets that have autosynchronization enabled and are using the network link.

Use the `sndradm -I` command to create 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 remote mirror software synchronization daemon on the target system whenever a synchronization is started and waits for `sndrsyncd` to perform any necessary PIT copies before allowing the synchronization to proceed.

The daemon is also notified when any remote mirror software resynchronization starts or finishes executing. The daemon performs 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.

Example Replication Scenarios

This section describes three example scenarios:

- [“One-to-Many Replication” on page 32](#)
- [“Many-to-One Replication” on page 33](#)
- [“Multihop Replication” on page 35](#)

One-to-Many Replication

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 site volume is a single volume set. Each volume set requires its own unique bitmap volume.

With one primary and three secondary host volumes, you need to configure three volume sets: primary A and secondary B1, primary A and secondary B2, and primary A and secondary B3. [FIGURE 2-5](#) shows one primary and three secondary host volumes and therefore three volume sets: A and B1, A and B2, and A and B3.

Consider the following:

- Each volume set on each host requires its own unique bitmap volume. For example, the volume set configurations shown in [FIGURE 2-5](#) require three unique bitmap volumes at each host: one for each volume set (A+B1, A+B2, and A+B3)
- To forward or reverse synchronize one volume set or all volume sets, issue a separate command for each set. You update the primary volume by using a specific secondary volume.

Note – Do not reverse synchronize the primary volume from more than one secondary volume or host at a time.

- You can group one-to-many sets that share a common primary volume into a single I/O group to forward synchronize all sets simultaneously instead of issuing a separate command for each set. You cannot use this technique to reverse synchronize volume sets, however. In this case, you must issue a separate command for each set and reverse update the primary volume by using a specific secondary volume

When one-to-many replication is performed in synchronous mode, I/O from the primary is sent to the first secondary in the configuration A+B1. The software does not wait for any I/O acknowledgment before starting to send the I/O to the second secondary volume in the configuration (B2). Writes are queued and processed in parallel. This pattern is repeated until I/O is acknowledged on all secondary volumes in the one-to-many configuration (B3).

In a synchronous one-to-many configuration, the latency at the primary host is the combined I/O latency for each connection to a secondary host and for each disk access on the secondary hosts.

When one-to-many replication is performed in asynchronous mode, I/O is queued at the primary host for later transmission and acknowledgment for every secondary host. This scheme allows replication to proceed in parallel during one-to-many asynchronous replications.

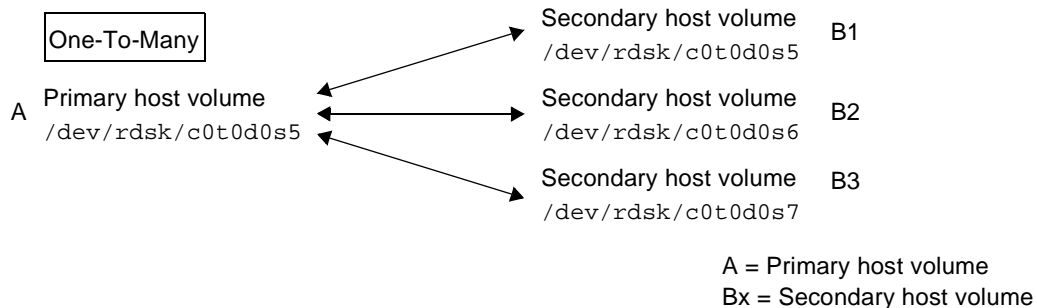


FIGURE 2-5 One-to-Many Volume Set Replication

Many-to-One Replication

The remote mirror software also supports the replication of volumes located on many different hosts to volumes on a single host. The terminology differs from the one-to-many configuration terminology, where the one and the many referred to are volumes. Many-to-one configuration refers to the ability to replicate volumes across more than two hosts through more than one network connection. An example of a many-to-one configuration is shown in [FIGURE 2-6](#).

[FIGURE 2-6](#) shows a simple use of the many-to-one configuration. Host A backs up volumes on both Host B and Host C. As the remote mirror software does not place restrictions on many-to-one configurations, Host A can be configured as the primary host for some replicated volumes and as the secondary host for others.

Many-To-One

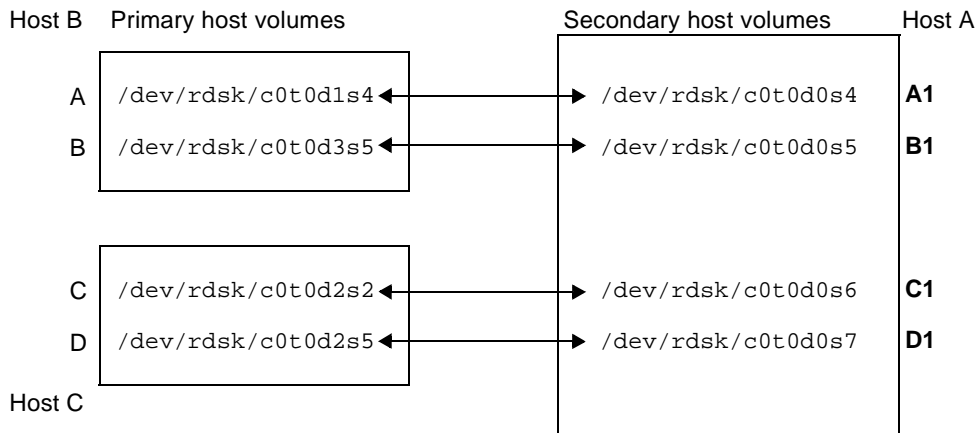


FIGURE 2-6 Many-to-One Volume Set Replication

Multihop Replication

In a multihop set, the secondary host volume of one volume set can act as the “primary” host volume of another volume set (it is still the secondary volume of the first volume set). In the case of one primary A and one secondary host volume B, the secondary host volume B appears as primary host volume A1 to the secondary host volume B1.

FIGURE 2-7 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.

Multi-hop configurations can become complex and the use and administration of multihop sets must be carefully considered.

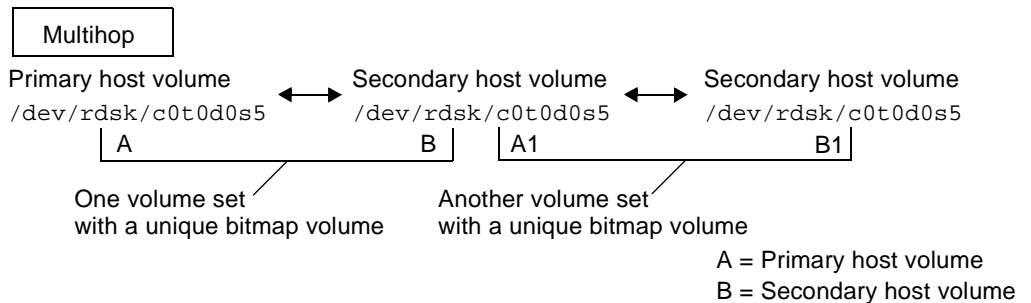


FIGURE 2-7 Multihop Volume Set Replication

For example, consider what happens if resynchronization operations for every volume set in the multihop chain are performed in synchronous mode. The I/O proceeds along each link of the chain and the I/O acknowledgment will not be confirmed until the last link is reached, at which point the process will complete.

If both sets were configured to replicate synchronously in the example in FIGURE 2-7:

- The I/O acknowledgment from B1 would be received at A1
- The I/O acknowledgment at B would be received at A

In a multihop configuration where every set in the chain is configured to replicate synchronously, the I/O latency at the primary node (assuming a forward replication) is the combined latency of every link and disk access along the chain.

Conversely, when volume sets are part of an multihop configuration where all sets replicate asynchronously, the contents of any given non-primary volume is unpredictable with respect to its neighbor until the resynchronization completes on all nodes.

These examples are for illustration only. The remote mirror software does not place any restrictions on the configurations between sets along the chain. A mix of synchronous and asynchronous sets is most useful.

As another example, configure the A+B volume set as a synchronous set running over a dark fiber in the same room (that is, ensure a consistent copy of the volume without adversely affecting performance on the primary site). Make the A1+B1 volume set an asynchronous set, running across a network to a remote location (that is, replicate the volume to a remote location at a comparatively fast rate by performing the replication asynchronously because of the high network latencies).

Multihop configurations can be expanded and the performance of these configurations improved when the point-in-time copy software and the remote mirror software are used together.

Preparing to Use the Software

This chapter provides information about the software before you first use it. The topics include the following:

- [“Increasing the Default Number of Volumes Allowed” on page 38](#)
- [“Bitmap Volumes” on page 41](#)
- [“Customizing Volume Sets” on page 42](#)
- [“Commands and I/O Group Operations” on page 45](#)
- [“Commands and Sun Cluster Operations” on page 45](#)

Increasing the Default Number of Volumes Allowed

The following sections describe how to change the default number of volumes you can use with the software.

- [“Using More Than 64 Volume Sets” on page 39](#)

The default number of *remote mirror volume sets* you can enable is 64. Follow these procedures to increase this number.

- [“Increasing the Storage Volume Device Limit” on page 40](#)

The default number of Storage Volume (SV) driver devices you can configure is 1024. This number of devices are divided for use between the remote mirror and point-in-time copy software. Follow these procedures to increase this number.

Note – After editing the files in this section, shut down and restart your server using the `shutdown` command for changes to take effect. Also, if you edit the `rdc.conf` file to use more than 64 volume sets, ensure that you have enough system resources (such as a large swap space).

Using More Than 64 Volume Sets

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

For example, to use 128 sets, change the file as follows; note the semicolon character (;) 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;
```

Note – The default number of Storage Volume (SV) driver devices you can configure is 1024, as set by the `nsc_max_devices` setting in the `nsctl.conf` file. See [“Increasing the Storage Volume Device Limit” on page 40](#) to change this setting.

Increasing the Storage Volume Device Limit

The default number of SV driver devices (that is, volumes) you can configure is 1024, as set by the `nsc_max_devices` setting in the `nsctl.conf` file. The number of volumes allowed is divided for use between the remote mirror and point-in-time copy software.

For example, if you use the point-in-time copy software only, you can have 341 volume sets, each consisting of master, shadow, and bitmap volumes. Also, If you use the remote mirror and point-in-time copy software products together, the number of volume sets are divided between the two products.

The following procedure describes how to increase this default limit.

▼ To Increase the Storage Volume Limit



Caution – Increasing this limit causes more memory to be consumed. You might have to adjust the default `nsc_global_pages` value of 2 in the `/usr/kernel/drv/mc_rms.conf` file. Only an experienced system administrator should make these changes.

1. **Log on as superuser.**
2. **Open the `/usr/kernel/drv/nsctl.conf` file using a text editor.**
3. **Search for the `nsc_max_devices` field.**
4. **Edit the number in this field to increase your volume limit.**
The default number is 1024.
5. **Save and exit the file.**
6. **Reboot your server using the `shutdown` command.**

Bitmap Volumes

The remote mirror version 3.1 software does not support bitmap files. The remote mirror software uses regular raw devices to store bitmaps.

These raw devices must be stored on a disk separate from the disk that contains the data from the replicated volumes. Configure RAID (such as mirrored partitions) for these bitmap devices and ensure that you mirror the bitmap to another disk in a different array. *The bitmap must not be stored on the same disk as the replicated volumes.*

If the bitmap and the replicated volumes reside on the same disk or array, then a single point of failure exists. In case of a disk or array failure, a greater chance of data loss exists. The bitmap might become corrupted.

In a clustered environment, a bitmap must reside only on a volume and cannot be a file. 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.

Bitmap Size Requirements

The bitmap size can be calculated using the following formula:

- 1 Kbyte + 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.)

Customizing Volume Sets

Before you start creating volume sets, see the following topics:

- [“Restricted Access When Creating and Configuring Sun StorEdge Volume Sets” on page 42](#)
- [“Using a Volume Set File” on page 43](#)

See also [“Reconfiguring Or Modifying a Volume Set” on page 91](#).

Restricted Access When Creating and Configuring Sun StorEdge Volume Sets



Caution – In a clustered environment, only one system administrator or root user at a time is allowed to create and configure Sun StorEdge volume sets. This restriction helps avoid creating an inconsistent Sun StorEdge Availability Suite volume set configuration.

Two administrators should not be writing to the Sun StorEdge Availability Suite software configuration at the same time. The operations that access the configuration include but are not limited to:

- Creating and deleting volume sets
- Adding and removing volume sets from I/O groups
- Assigning new bitmap volumes to a volume set
- Updating the disk device group or resource name
- Any operation that changes the Sun StorEdge Availability Suite software and related volume set configuration

Using a Volume Set File

When you enable the remote mirror software, you can specify an optional *volume set file* containing information about the volume set: volumes, primary and secondary hosts, bitmaps, operating mode, and so on. Use the `sndradm -f volset-file` option when you use a volume set file.

You can also enter information about each volume set from the command line, but it might be more convenient to put this information in a file when you have multiple volume sets.

One advantage when using one or more volume set files is that you can operate on specific volume sets and exclude other sets from the operation. Unlike adding the volume sets to an I/O group, you can mix replication modes in a volume set file.

The fields for the volume set file specified using the `-f volset-file` option are:

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

An example file entry is as follows:

```
atm10 /dev/vx/rdisk/oracle816/oratest /dev/vx/rdisk/oracle816/oratest_bm \  
atm20 /dev/vx/rdisk/oracle816/oratest /dev/vx/rdisk/oracle816/oratest_bm \  
ip sync g oragroup
```

See [TABLE 3-1](#). See the `rdc.cf` man page for more information about the volume set file format.

TABLE 3-1 Volume Set File Format Fields

Field	Meaning	Description
<i>phost</i>	Primary host	Server on which the primary volume resides.
<i>pdev</i>	Primary device	Primary volume partition. Specify full path names only (for example, /dev/rdisk/c0t1d0s4).
<i>pbitmap</i>	Primary bitmap	Volume partition in which the bitmap of the primary partition is stored. Specify full path names only.
<i>shost</i>	Secondary host	Server on which the secondary volume resides.
<i>sdev</i>	Secondary device	Secondary volume partition. Specify full path names only.
<i>sbitmap</i>	Secondary bitmap	Volume partition in which the bitmap of the secondary partition is stored. Specify full path names only.
<i>ip</i>	Network transfer protocol	Specify <i>ip</i> .
<i>sync</i> <i>async</i>	Operating mode	<i>sync</i> is the mode where the I/O operation is confirmed as complete only when the remote volume has been updated. <i>async</i> is the mode where the primary host I/O operation is confirmed as complete before updating the remote volume.
<i>g io-groupname</i>	I/O group name	An I/O group name can be specified using the <i>g</i> character. In this example, it is <i>oragroup</i> .

Commands and I/O Group Operations

Adding the remote mirror 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, excluding all other volume 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.

Failed Operations in an I/O Group

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

1. Correct any known problems with the failing sets
2. Reissue the command on the I/O group

Commands and Sun Cluster Operations

You use the `C tag` and `-C tag` options described in [Chapter 5](#) in Sun Cluster 3.0 Update 1 and Update 2 environments only. If you accidentally use these options in a noncluster environment, the remote mirror operation does not execute.

Using the Remote Mirror Software

This chapter includes examples of how to use the remote mirror software command `sndradm`. The topics in this chapter include:

- [“Example Host and Volume Names for This Chapter” on page 48](#)
- [“Initial Steps for Using the Software” on page 50](#)
- [“What Happens If a Network Link Fails?” on page 54](#)
- [“Rehearsing Disaster Recovery” on page 56](#)
- [“About Primary Volume Failures” on page 60](#)
- [“Disabling Remote Replication” on page 62](#)
- [“Swapping the Remote Mirror Hosts” on page 63](#)

Example Host and Volume Names for This Chapter

In the examples in this chapter, the following example names are used:

TABLE 4-1 Example Names and Devices Used in This Chapter

Primary Host Name	rmshost1
Primary Volume	/dev/vx/rdisk/c0t117d0s3
Primary Bitmap	/dev/vx/rdisk/bmap/bm1
Secondary Host Name	rmshost2
Secondary Volume	/dev/vx/rdisk/c0t117d0s5
Secondary Bitmap	/dev/vx/rdisk/bmap/bm2
Set Name (assigned by the software)	rmshost2:/dev/vx/rdisk/c0t117d0s5

Depending on the example, either site can be the primary or secondary of the remote copy operation. As shown in [TABLE 5-4](#), you must perform all synchronization operations from the primary host session.

Using the Storage Volume Interface

When you enable volume sets using the `/usr/opt/SUNWesm/sbin/sndradm` command, the specified volumes are automatically placed under control of the Storage Volume (SV) driver. See [“Enabling and Disabling Volume Sets” on page 77](#).

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 volset-file
```

where *volset-file* is an ASCII file listing the volumes. See [“Using a Volume Set File” on page 43](#). You can also add individual volumes from the command line. See the `svadm` man page.

Monitoring Remote Mirror Software Operations

- Use the `/usr/opt/SUNWscm/sbin/dsstat` command described in [Appendix A](#).
- Use the `/usr/opt/SUNWscm/sbin/scmadm -S -M` command described in [Appendix A](#) to monitor the progress of the remote mirror software operations.

Initial Steps for Using the Software

The following sections describe the initial steps for using the remote mirror software:

- “Step 1: Enabling the Volume Sets” on page 50
- “Step 2: Establishing Volume Copies for the First Time” on page 51
- “Step 3: Updating the Secondary Volume” on page 53

Step 1: Enabling the Volume Sets

The first step to using the remote mirror software is to enable the software on volume sets. Make sure you perform this step on both the primary and secondary hosts. A common user error is to enable volume sets on only one host. In this example, enable the hosts in asynchronous mode. The example also shows that the volumes need to be synchronized and do not match (in this case, use the `sndradm -e` command).

▼ To Enable the Volume Sets

1. **Log into the primary host `rmshost1` as the superuser.**
2. **Enable the volume sets.**

```
rmshost1# sndradm -e rmshost1 /dev/vx/rdisk/c0t117d0s3 /dev/vx/rdisk/bmap/bm1 \  
rmshost2 /dev/vx/rdisk/c0t117d0s5 /dev/vx/rdisk/bmap/bm2 ip async
```

3. **Log into the secondary host `rmshost2` as the superuser.**
4. **Enable the volume sets.**

```
rmshost2# sndradm -e rmshost1 /dev/vx/rdisk/c0t117d0s3 /dev/vx/rdisk/bmap/bm1 \  
rmshost2 /dev/vx/rdisk/c0t117d0s5 /dev/vx/rdisk/bmap/bm2 ip async
```

The following events have now occurred:

- The software has defined the volume set name as `rmshost2:/dev/vx/rdisk/c0t117d0s5`, meaning that you can use this name with the `sndradm` command, instead of specifying each item in the volume set each time.
- The bitmaps are now set to 1, indicating that a full synchronization between volumes is required.

Step 2: Establishing Volume Copies for the First Time

The next step is to perform a full forward synchronization to copy the primary volume contents to the secondary volume. The remote mirror software performs this initial primary-volume-to-secondary-volume copy while also forwarding any new primary volume updates to the secondary volume.

Note – The volumes can be made identical using other methods, not just a full synchronization. When network latencies justify it, you can perform the initial synchronization of a volume set by backing up a source or primary volume on magnetic tape on one site, then restoring the volume from the tape on the other site. While the volumes are out of sync (that is, the period between when the backup is finished and the restore is started), place the source or primary volume in logging mode.

Make sure that the backup copy is a physical copy (for example, by using the `dd(1M)` command) and not a logical copy (for example, one made using the `tar(1M)` or `cpio(1M)` commands). The copies must have identical blocks, not just identical files.

In this case, use the `sndradm -E` command instead of `sndradm -e` to enable the volume sets.

▼ To Fully Synchronize the Remote Volumes While Applying New Updates

1. **Log into the primary host** `rmshost1` **as superuser.**

2. **Unmount the secondary volume.**

You can keep the primary volume mounted.



Caution – Do not mount the secondary volume until the synchronization is complete. See [“About Replicating File Systems” on page 4](#).

3. **Synchronize the volumes.**

```
rmshost1# sndradm -m rmshost2:/dev/vx/rdisk/c0t117d0s5
```

4. **Check the synchronization progress.**

```
rmshost1# scmadm -S -M
```

After the synchronization is complete, the remote mirror software continues to replicate any primary volume changes to the secondary volume. Keep the secondary volume unmounted during replication or until you are ready to permit your application to write to it. When you are ready to permit write operations to the secondary volume, place the volume sets into logging mode and mount the volume. The software continues tracking changes through the bitmap until you are ready to update or resynchronize the volumes.

Step 3: Updating the Secondary Volume

This section describes the commands to use when you are ready to update the secondary volume to resynchronize volumes.

Optionally, you can also use the Sun StorEdge Availability Suite 3.1 point-in-time copy software to ensure data consistency at the primary and secondary volumes. Using this software helps guarantee a known good copy of your data in case of a network link failure during the synchronization. See the *Sun StorEdge Availability Suite 3.1 Point-in-Time Copy Software Administration and Operations Guide* for details about the `iiadm` command.

Note – You must place the related remote mirror volume set in logging mode for the point-in-time copy software to successfully perform an enable, copy, update, or reset operation on a remote mirror volume. If the volume set is not in logging mode, the point-in-time copy operation fails and the remote mirror software reports that the operation is denied

▼ To Resynchronize the Primary and Secondary Volumes

1. **Log into the primary host `rmshost1` as the superuser.**
2. **Quiesce any applications writing to the primary volume.**

Note – You are not necessarily required to quiesce your applications but doing so helps guarantee a consistent copy of your data. Quiescing applications also helps guarantee a consistent copy for the point-in-time copy software. If you choose not to quiesce your applications and to not use the point-in-time copy software, the remote mirror software still forwards any data updates to the secondary volume during replication.

3. **(Optional) Take a point-in-time copy snapshot of the primary volume using the `iiadm` command.**
Ensure that the primary volume is in logging mode (`sndradm -l`) first.
4. **(Optional) Take a point-in-time copy snapshot of the secondary volume using the `iiadm` command.**
Ensure that the secondary volume is in logging mode first.

5. Copy only the changed data from the primary volume to the secondary volume.

```
rmshost1# sndradm -u rmshost2:/dev/vx/rdisk/c0t117d0s5
```

6. Check the synchronization progress.

```
rmshost1# scmadm -S -M
```

When the update synchronization completes, the secondary volume is a block-for-block copy of the primary volume and the bitmaps are cleared to 0. See [“What Happens If a Network Link Fails?”](#) on page 54.

What Happens If a Network Link Fails?

The remote mirror software uses a periodic signal to monitor the health of primary and secondary systems. If the software cannot sense a health monitor signal, it assumes an interruption in the remote mirror software service.

The remote mirror software then places all volume sets into logging mode. During logging mode, the software only updates the primary volume bitmap. (The software assumes that the secondary volume is not mounted and not being written to.) See [“Logging”](#) on page 26 and [“Stopping Replication and Starting Logging”](#) on page 83.

Note – You can introduce interruptions intentionally to exercise remote failure strategies—for example, during the disaster recovery rehearsals described in [“Rehearsing Disaster Recovery”](#) on page 56.

When Not to Resynchronize Volumes

Resynchronization is discouraged if the interruption is the warning of a larger rolling disaster. 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. The autosynchronization option is disabled by default for this reason. See [“Choosing Automatic or Manual Resynchronization”](#) on page 28.

Autosynchronization

Note – If you enable autosynchronization and later issue a synchronization command like `sndradm -u`, autosynchronization remains enabled.

After a network link failure is restored, the software attempts to synchronize the primary and secondary volumes if you have enabled autosynchronization. To enable autosynchronization, use the `sndradm -a` command described in [“Enabling or Disabling the Autosynchronization Feature” on page 104](#).

See also [“Choosing Automatic or Manual Resynchronization” on page 28](#) and [“Autosynchronization With the Point-in-Time Copy Software” on page 30](#).

Manually Resynchronizing Volumes After a Network Link Failure or Other Interruption



Caution – While resynchronization is occurring, the secondary volume data is temporarily inconsistent and cannot be relied on for recovery. Consistency is restored when the resynchronization finishes executing. To ensure data integrity, use the point-in-time copy software regularly to create a snapshot of data at both sites.

Typically, interruptions in the remote mirror software services are infrequent.

If the secondary volume state is unknown because of system or disk failure, make full volume copies to reestablish matching remote mirror software volume sets. In this case, use the `sndradm -m` command to fully update the secondary volume set.

- **Perform the procedures in [“To Fully Synchronize the Remote Volumes While Applying New Updates” on page 52](#).**

Rehearsing Disaster Recovery

The remote mirror software enables you to perform disaster rehearsals, encouraging verification of your disaster plans. Perform rehearsals regularly and refine them whenever a significant change is made to the primary or secondary host environments.

When rehearsing disasters or in case of an actual disk disaster or failure, keep the failed volumes under control of the remote mirror software. Do not disable the software. The remote mirror software marks the device as failed when it is unable to read from or write to it. In the case of a primary failure, for example, the remote mirror software continues to provide read and write services to the host application using the secondary volume at the remote site.

▼ To Rehearse a Primary Volume or Site Failure

- 1. Simulate a primary volume or site disaster by using one of the following methods:**
 - Break the network link between the primary and secondary sites by disconnecting the network cable or otherwise shutting down the network link. In this case, the remote mirror software places the volumes into logging mode after detecting the break.
 - Place the primary and secondary volumes into logging mode. See [“Stopping Replication and Starting Logging”](#) on page 83.
- 2. Ensure that any cached data at the secondary site is destaged by using the `scmadm -S` command. [Appendix A](#) describes this command.**
- 3. When the data is destaged, mount the secondary volume in read-write mode so your application can write to it.**
- 4. Configure your application to read and write to the secondary volume.**

The secondary bitmap volume tracks the volume changes.
- 5. Fix the “failure” at the primary volume by using one of the following methods:**
 - Reconnect the network.
 - Disable logging mode by issuing a resynchronization command. Quiesce your application and unmount the secondary volume if you choose this method.

Note – If the autosynchronization feature is enabled, the remote mirror software resynchronizes the primary volume from the secondary volume when the link is re-established. If you also install and configure the point-in-time copy software, this software takes a snapshot copy of your secondary volume data before performing the reverse update synchronization. Consider if this approach is appropriate for your disaster recovery plan.

You can now choose to resynchronize your volumes:

- 6. Ensure that any cached data at the secondary site is purged or synchronized by using the `scmadm -p` or `scmadm -s` command. [Appendix A](#) describes this command.**

If you do not perform this step, when you attempt to perform a reverse copy or update operation, the software reports that the update failed and that the volume sets are in logging mode. See [“Stopping Replication and Starting Logging” on page 83](#).

- 7. Ensure that any cached data at the secondary site is destaged by using the `scmadm -S` command.**

The software automatically destages it.

Note – If you check the volume status at this point using the `scmadm -S` command, the volume status might report VOF (volume failed) or RSN (reverse synchronization needed). When you perform the reverse copy or update, this volume status is cleared.

- 8. Perform the update by choosing one of the following methods:**

- Roll back the secondary volume to its original pre-disaster state to match the primary volume. Use the `sndradm -m copy` or `sndradm -u update` commands.
- Keep the changes from the updated secondary volume and resynchronize so that both volumes match. Use the `sndradm -m r reverse copy` or `sndradm -u r reverse update` commands.

▼ To Rehearse a Secondary Volume or Site Failure

1. **Simulate a secondary volume or site disaster by using one of the following methods:**
 - Break the network link between the primary and secondary sites by disconnecting the network cable or otherwise shutting down the network link. In this case, the remote mirror software places the volumes into logging mode after detecting the break.
 - Place the primary and secondary volumes into logging mode. See [“Stopping Replication and Starting Logging” on page 83](#).
2. **Fix the “failure” at the primary volume by using one of the following methods:**
 - Reconnect the network.
 - Disable logging mode by issuing a resynchronization command. Unmount the secondary volume if you choose this method.

Note – If the autosynchronization feature is enabled, the remote mirror software resynchronizes the secondary volume from the primary volume when the link is re-established. If you also install and configure the point-in-time copy software, this software takes a snapshot copy of your secondary volume data before performing the reverse update synchronization. Consider if this approach is appropriate for your disaster recovery plan.

You can now choose to resynchronize your volumes:

3. **Ensure that any cached data at the secondary site is purged or synchronized by using the `scmadm -p` or `scmadm -s` command.**
See [Appendix A](#).
If you do not perform this step, when you attempt to perform a reverse copy or update operation, the software reports that the update failed and that the volume sets are in logging mode. See [“Stopping Replication and Starting Logging” on page 83](#).
4. **Ensure that any cached data at the secondary site is destaged by using the `scmadm -S` command. See [Appendix A](#).**

Note – If you check the volume status at this point using the `scmadm -S` command, the volume status might report VOF (volume failed) or RSN (reverse synchronization needed). When you perform the reverse copy or update, this volume status is cleared.

5. Perform the update by choosing one of the following methods:

- Update the secondary volume to match the primary volume. Use the `sndradm -u` update command.
- If you are unsure about the secondary volume contents, perform a full forward synchronization to match the primary volume. Use the `sndradm -m` synchronization command.

About Primary Volume Failures

Note – Keep the failed volumes under control of the remote mirror software. Do not disable the software. The remote mirror software marks the device as failed when it is unable to read to or write from it. The remote mirror software continues to provide read and write services to the host application using the secondary volume at the remote site.

The remote mirror software provides continuous data access during primary volume failures. The remote mirror 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 remote mirror software's 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 remote mirror 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 remote mirror 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 remote mirror software remote volume access only when two or more of the disks in the RAID 5 stripe fail.

Recovering From a Primary Site Disaster

The remote mirror software minimizes the effects of a disaster at the primary site by enabling you to keep the secondary storage images updated. Although the secondary remote mirror 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 remote mirror software service, the remote mirror software automatically destages the secondary remote mirror software cache to its corresponding secondary volumes.

After all the secondary volumes have been updated with the latest remote mirror software cache images, the secondary volumes can be accessed by the secondary hosts. The `scmadm -S` command shows information confirming 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 remote mirror software enabled at the secondary site to track disk areas that are being modified.

Restoring a Primary Site From the Secondary Site

If the primary host becomes inoperative and primary data on the primary disks is lost, update logs at the secondary systems have little value. You must flush the cache and perform a full reverse 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 remote mirror software-managed volumes. This reverse synchronization process ensures that only the latest data is deposited on the primary disks. See [“Rehearsing Disaster Recovery” on page 56](#).

Disabling Remote Replication



Caution – Disable remote replication *only* when the primary and secondary volumes will no longer be associated.

Disabling the remote mirror software breaks the connection between primary and secondary volumes, discards any bitmap information, and removes the host and volume information from the Sun StorEdge configuration. Following the remote mirror software disable, enable and full synchronization (full volume copy) operations are necessary to re-establish the remote mirror software relationship and ensure that the contents of each volume match. See [“Enabling and Disabling Volume Sets” on page 77](#) and [“Step 2: Establishing Volume Copies for the First Time” on page 51](#).

Swapping the Remote Mirror Hosts



Caution – Before performing this procedure, ensure that I/O operations are not occurring to the volumes at the primary and secondary hosts. Data corruption will result if I/O operations continue.

In case of disaster recovery or link failure situations, you can also swap the remote mirror host roles to provide access to your critical data. That is, the primary host can become the secondary host and the secondary host can become the primary host. This alternate scheme enables you to recover the old primary host and, if you choose, switch back to the original roles.

The basic steps to swap host roles are as follows. Before performing these steps, quiesce the application writing to the remote mirror volumes and then unmount those volumes.

1. Disable the remote mirror software at the primary host (named Site-A). This step also discards the primary bitmap volume at Site-A.
2. Create or edit a volume set file to change the host information. See [“Sample rdc.cf File Used in This Example” on page 64](#).
3. At Site-A, unmount the volumes, if possible.
4. At the secondary host (named Site-B), disable the remote mirror software.
5. At Site-B, edit the `/etc/opt/SUNWrdc/rdc.cf` file to change the host information to match Site-A.
6. At Site-B (which is now the primary host), enable the remote mirror software.
7. At Site-B, synchronize the volumes from Site-B to Site-A.
8. Perform any modifications or recovery procedures required by your application.

Note – For example, if you are using a database application, you might have to copy data and control files to the new secondary host after the synchronization.

Sample `rdc.cf` File Used in This Example

The following is an example `/etc/opt/SUNWrdc/rdc.cf` volume set file. See also [“Using a Volume Set File” on page 43](#).

```
rmshost1 /dev/vx/rdisk/c0t117d0s3 /dev/vx/rdisk/bmap/bm1 \  
rmshost2 /dev/vx/rdisk/c0t117d0s5 /dev/vx/rdisk/bmap/bm2 ip sync
```

Note – Your actual volume set file can be any name. The `rdc.cf` file name is used as an example here.

TABLE 4-2 Example Names and Devices Used in This Procedure

Primary Host Name (Site-A)	<code>rmshost1</code>
Primary Volume	<code>/dev/vx/rdisk/c0t117d0s3</code>
Primary Bitmap	<code>/dev/vx/rdisk/bmap/bm1</code>
Secondary Host (Site-B)	<code>rmshost2</code>
Secondary Volume	<code>/dev/vx/rdisk/c0t117d0s5</code>
Secondary Bitmap	<code>/dev/vx/rdisk/bmap/bm2</code>
Transmission Protocol	<code>ip</code>
Replication Mode	<code>sync</code>
Set Name (assigned by the software)	<code>rmshost2:/dev/vx/rdisk/c0t117d0s5</code>

▼ To Disable the Software at Site-A



Caution – Before performing this procedure, ensure that I/O operations are not occurring to the volumes at the primary and secondary hosts. Data corruption will result if I/O operations continue to occur. Before performing this procedure, quiesce the application writing to the remote mirror volumes and then unmount those volumes.

Note – The following example assumes that the `/etc/opt/SUNWrdc/rdc.cf` volume set file is already created and the volumes specified in it are enabled.

1. At Site-A, disable the remote mirror software and discard the remote mirror scoreboard bitmap.

```
rmshost1# sndradm -dn -f /etc/opt/SUNWrdc/rdc.cf
```

2. Edit the `rdc.cf` file to swap the Site-A primary host information and Site-B secondary host information.

For example, in the example entry shown in [“Sample rdc.cf File Used in This Example” on page 64](#), change `rmshost1` to `rmshost2` and `rmshost2` to `rmshost1`.

3. If possible, unmount the remote mirror volumes.

```
rmshost1# umount mount-point
```

▼ To Change the Site-B Secondary Host to the Primary Host

1. At Site-B, disable the remote mirror software and discard the remote mirror scoreboard bitmap.

```
rmshost2# sndradm -dn -f /etc/opt/SUNWrdc/rdc.cf
```

2. Edit the `rdc.cf` file to swap the Site-A primary host information and Site-B secondary host information.

For example, in the example entry shown in [“Sample rdc.cf File Used in This Example” on page 64](#), change `rmshost1` to `rmshost2` and `rmshost2` to `rmshost1`.

3. Enable the remote mirror software *at both hosts*.

```
rmshost1# sndradm -En -f /etc/opt/SUNWrdc/rdc.cf
rmshost2# sndradm -En -f /etc/opt/SUNWrdc/rdc.cf
```

Note – Use the `-E` enable option to ensure that the bitmap contents are clear (0), indicating that a synchronization is not required.

4. If desired, at Site-A, perform a full synchronization from Site-B to Site-A.

```
rmshost1# sndradm -mn -f /etc/opt/SUNWrdc/rdc.cf
```

5. Perform any modifications or recovery procedures required by your application.

The `sndradm` Command

Note – To use the remote mirror software, you must be the root user.

This chapter describes the remote mirror software
`/usr/opt/SUNWesm/sbin/sndradm` command and options.

The topics in this chapter include:

- “Log File” on page 68
- “`sndradm` Syntax Summary” on page 69
- “Which Host Do I Issue Commands From?” on page 76
- “Enabling and Disabling Volume Sets” on page 77
- “Synchronizing and Updating Volume Sets” on page 84
- “Reconfiguring Or Modifying a Volume Set” on page 91
- “Displaying Volume Set Information” on page 100
- “Enabling or Disabling the Autosynchronization Feature” on page 104
- “Setting the Asynchronous Queue” on page 105
- “Tuning the Asynchronous Queue” on page 106

Log File

The `/var/opt/SUNWesm/ds.log` file contains operation log messages for the remote mirror commands. For example, these log messages show a volume set enabling and updating:

```
Oct 17 13:48:10 sndr: sndradm -e atm10 /dev/vx/rdisk/oracle816/oratest
/dev/vx/rdisk/oracle816/oratest_bm atm20 /dev/vx/rdisk/oracle816/oratest
/dev/vx/rdisk/oracle816/oratest_bm ip sync g oragroup
Successful
Oct 17 13:48:10 sv: enabled /dev/vx/rdisk/oracle816/oratest
Oct 17 13:48:10 sv: enabled /dev/vx/rdisk/oracle816/oratest_bm
Oct 18 10:23:54 sndr: sndradm -u atm10 /dev/vx/rdisk/oracle816/oratest
/dev/vx/rdisk/oracle816/oratest_bm atm20 /dev/vx/rdisk/oracle816/oratest
/dev/vx/rdisk/oracle816/oratest_bm ip sync g oragroup
Starting
```

sndradm Syntax Summary

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

The `/usr/opt/SUNWesm/sbin/sndradm` command uses the command switches in [TABLE 5-1](#). ([TABLE 5-2](#) shows the commands sorted in alphabetical order) and the options in [TABLE 5-3](#). The command syntax is as follows:

Syntax

```
sndradm switches [options]
```

For example, to update the secondary volume from the primary volume:

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

Command Switches

TABLE 5-1 Command Switches Sorted By Function

Function	Command and Switch	See This Section
Enable and Disable Volume Sets		
<ul style="list-style-type: none"> • Enable the software for the specified set (sync not required). 	sndradm -E	“Enabling Volume Sets” on page 78
<ul style="list-style-type: none"> • Enable the software for the specified set (sync required). 	sndradm -e	“Enabling Volume Sets” on page 78
<ul style="list-style-type: none"> • Disable the software for the specified set. 	sndradm -d	“Disabling Volume Sets” on page 80
<ul style="list-style-type: none"> • Add or delete a point-in-time copy software volume group. 	sndradm -I	“Adding and Deleting Point-in-Time Copy Software Volumes” on page 81
<ul style="list-style-type: none"> • Stop the software replication, enter logging mode, and track changes using bitmaps. 	sndradm -l	“Stopping Replication and Starting Logging” on page 83
Synchronize or Update Volume Sets		
<ul style="list-style-type: none"> • Copy the entire contents of the primary volume to the secondary volume; also known as <i>full synchronization</i>. 	sndradm -m	“Starting a Full Synchronization (Copy Between Volumes)” on page 86
<ul style="list-style-type: none"> • Copy the entire contents of the secondary volume to the primary volume; also known as <i>full reverse synchronization</i>. 	sndradm -m -r	“Starting a Full Synchronization (Copy Between Volumes)” on page 86
<ul style="list-style-type: none"> • Update only the changed data of the primary volume to the secondary volume; also known as <i>update synchronization</i> or <i>resynchronization</i>. 	sndradm -u	“Starting an Update Resynchronization (Update Between Volumes)” on page 88
<ul style="list-style-type: none"> • Update only the changed data of the secondary volume to the primary volume; also known as <i>reverse synchronization</i> or <i>reverse update</i>. 	sndradm -u -r	“Starting an Update Resynchronization (Update Between Volumes)” on page 88
<ul style="list-style-type: none"> • Wait for a synchronization operation to finish executing. 	sndradm -w	“Waiting for Synchronization to Finish Executing” on page 90
Reconfigure or Change Volume Sets		
<ul style="list-style-type: none"> • Change the bitmap of a volume set. 	sndradm -R b	“Reconfiguring Or Modifying a Volume Set” on page 91
<ul style="list-style-type: none"> • Update or reconfigure the disk group name or cluster resource tag. 	sndradm -R C	“Updating the Disk Cluster Tag Name” on page 99
<ul style="list-style-type: none"> • Update or reconfigure the software volume sets. 	sndradm -R -f <i>volset-file</i>	“Updating the Volume Set Information” on page 98

TABLE 5-1 Command Switches Sorted By Function (*Continued*)

Function	Command and Switch	See This Section
<ul style="list-style-type: none"> • Move all the specified remote mirror sets to a different group. 	<code>sndradm -R g</code>	“Moving a Volume Set to a Different I/O Group” on page 93
<ul style="list-style-type: none"> • Reset the replication mode of a volume set. 	<code>sndradm -R m {sync async}</code>	“Changing the Replication Mode of a Volume Set” on page 96
Display Information		
<ul style="list-style-type: none"> • Display the command syntax. 	<code>sndradm -h</code>	See the on-screen display results from entering this command at the command prompt
<ul style="list-style-type: none"> • Display the version information. 	<code>sndradm -v</code>	See the on-screen display results from entering this command at the command prompt
<ul style="list-style-type: none"> • Display the software status. 	<code>sndradm -p</code> <code>sndradm -P</code>	“Displaying Volume Set and I/O Group Status” on page 101
<ul style="list-style-type: none"> • Display the software volume set and I/O group names. 	<code>sndradm -i</code>	“Displaying a List of Current Volume Sets and I/O Group Names” on page 102
<ul style="list-style-type: none"> • Display the status of the link connecting systems running the software. 	<code>sndradm -H</code>	“Displaying Link Status” on page 103
Miscellaneous		
<ul style="list-style-type: none"> • Set and tune the asynchronous queue. 	<code>sndradm -W</code> <code>sndradm -F</code>	“Setting the Asynchronous Queue” on page 105 “Tuning the Asynchronous Queue” on page 106
<ul style="list-style-type: none"> • Toggle the autosynchronization state. 	<code>sndradm -a {on off}</code>	“Enabling or Disabling the Autosynchronization Feature” on page 104

TABLE 5-2 Command Switches in Alphabetical Order

Command and Switch	Description	See This Section
<code>sndradm -a {on off}</code>	Autosynchronization mode; turn on or off.	“Enabling or Disabling the Autosynchronization Feature” on page 104
<code>sndradm -e</code>	Enable the software for the specified set (sync required).	“Enabling Volume Sets” on page 78
<code>sndradm -E</code>	Enable the software for the specified set (sync not required).	“Enabling Volume Sets” on page 78
<code>sndradm -d</code>	Disable the software.	“Disabling Volume Sets” on page 80
<code>sndradm -h</code>	Display the command syntax.	See the on-screen display results from entering this command at the command prompt
<code>sndradm -H</code>	Display the status of the link connecting systems running the software.	“Displaying Link Status” on page 103
<code>sndradm -i</code>	Display the software volume set and I/O group names.	“Displaying Volume Set and I/O Group Status” on page 101
<code>sndradm -I</code>	Add or delete a Sun StorEdge point-in-time copy software volume group for use with the software.	“Adding and Deleting Point-in-Time Copy Software Volumes” on page 81
<code>sndradm -l</code>	Logging mode; place the volume sets into logging mode.	“Stopping Replication and Starting Logging” on page 83
<code>sndradm -m</code>	Full forward synchronization. Copy the entire contents of the primary volume to the secondary volume.	“Starting a Full Synchronization (Copy Between Volumes)” on page 86
<code>sndradm -m -r</code>	Full reverse synchronization. Copy the entire contents of the secondary volume to the primary volume.	“Starting a Full Synchronization (Copy Between Volumes)” on page 86
<code>sndradm -p</code>	Display software status.	“Displaying a List of Current Volume Sets and I/O Group Names” on page 102
<code>sndradm -P</code>	Display software status.	“Displaying a List of Current Volume Sets and I/O Group Names” on page 102

TABLE 5-2 Command Switches in Alphabetical Order (*Continued*)

Command and Switch	Description	See This Section
<code>sndradm -R b</code>	Reconfigure the volume set by changing the bitmap volume.	“Changing the Bitmap Volume in a Volume Set” on page 92
<code>sndradm -R C</code>	Reconfigure the disk group name or cluster resource tag.	“Updating the Disk Cluster Tag Name” on page 99
<code>sndradm -R g</code>	Reconfigure the specified volume sets by moving them to a different group.	“Moving a Volume Set to a Different I/O Group” on page 93
<code>sndradm -R -f <i>volset-file</i></code>	Reconfigure the volume sets specified in the <i>volset-file</i> .	“Updating the Volume Set Information” on page 98
<code>sndradm -R m {sync async}</code>	Reconfigure the replication mode of a volume set.	“Changing the Replication Mode of a Volume Set” on page 96
<code>sndradm -u</code>	Update only the changed data of the primary volume to the secondary volume; also known as <i>update synchronization</i> or <i>resynchronization</i> .	“Starting an Update Resynchronization (Update Between Volumes)” on page 88
<code>sndradm -u -r</code>	Update only the changed data of the secondary volume to the primary volume; also known as <i>reverse synchronization</i> or <i>reverse update</i> .	“Starting an Update Resynchronization (Update Between Volumes)” on page 88
<code>sndradm -v</code>	Displays the version information.	See the on-screen display results from entering this command at the command prompt
<code>sndradm -w</code>	Wait for a synchronization operation to finish executing.	“Waiting for Synchronization to Finish Executing” on page 90
<code>sndradm -W</code> <code>sndradm -F</code>	Set and tune the asynchronous queue.	“Setting the Asynchronous Queue” on page 105 “Tuning the Asynchronous Queue” on page 106

Command Options

TABLE 5-3 Command Options

Option	Definition
-n	<p>Does not prompt the user after issuing a <code>sndradm</code> command. The default is to prompt and request a response. For example, after starting a full synchronization from the primary to secondary volumes, the remote mirror software prompts <code>Overwrite secondary with primary? (Y/N) [N]</code>.</p> <p>This option is useful when you include the commands in a script file.</p>
-f <i>volset-file</i>	<p>Specifies a volume set file defining the remote mirror software volume sets. If you do not specify a <i>volset-file</i>, <i>volset</i>, or <i>set-name</i>, the command acts on all volume sets.</p>
<i>volset</i>	<p>Specifies the complete remote mirror software set configuration information. If you do not specify a <i>volset-file</i>, <i>volset</i>, or <i>set-name</i>, the command acts on all volume sets. The <i>volset</i> format is as follows:</p> <pre><i>phost pdev pbitmap shost sdev sbitmap</i> ip {sync async} [g <i>io-groupname</i>] [C <i>tag</i>]</pre> <p><i>phost</i> - Server on which the primary volume resides.</p> <p><i>pdev</i> - Primary volume partition to be copied. Specify full path names only; for example, <code>/dev/rdisk/c0t1d0s4</code>.</p> <p><i>pbitmap</i> - Volume partition where the bitmap of the primary partition is stored. Specify full path names only; for example, <code>/dev/rdisk/c0t1d0s6</code>.</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; for example, <code>/dev/rdisk/c0t1d0s7</code>.</p> <p><i>sbitmap</i> - Volume partition in which the bitmap of the secondary partition is stored. Specify full path names only; for example, <code>/dev/rdisk/c0t1d0s8</code>.</p> <p><i>ip</i> - Specifies the network protocol.</p> <p><i>sync</i> <i>async</i> - remote mirror software operating mode. <i>sync</i> is the remote mirror software mode where the I/O operation is not confirmed as complete until the remote volume has been updated. <i>async</i> is the remote mirror software mode where the primary host I/O operation is confirmed as complete before updating the remote volume.</p> <p><i>g io-groupname</i> - Specifies which I/O group the volume set belongs to.</p> <p><i>C tag</i> - 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, <code>/dev/rdisk/md/dg/vol</code> and <code>/dev/vx/rdisk/dg/vol</code> both indicate a disk group name of <i>dg</i>.) It is the responsibility of the user to ensure that the cluster tag specified to the remote mirror software matches the appropriate cluster resource group tag and to keep all installed data services updated in the event of cluster resource group reconfigurations.</p> <p>You can use the <i>C tag</i> and <i>-C tag</i> options in a Sun Cluster 3.0 Update 1 or Update 2 environment only. If you accidentally use these options in a non-Sun Cluster 3.0 environment, the remote mirror operation does not execute.</p>

TABLE 5-3 Command Options (*Continued*)

Option	Definition
<i>set-name</i>	Name of the remote mirror software volume set as assigned by the remote mirror software. The remote mirror software assigns a default volume set name of <i>shost:sdev</i> , where <i>shost</i> is the secondary host name and <i>sdev</i> is the secondary volume partition name, separated by a colon (:).
-g <i>io-groupname</i>	Name of the I/O group containing the collection of remote mirror software volume sets.
	Specifying -g <i>io-groupname</i> limits any operation performed to the sets in the named <i>io-groupname</i> .
-C <i>tag</i>	You can use the C <i>tag</i> and -C <i>tag</i> options in a Sun Cluster 3.0 Update 1 or Update 2 environment only. If you accidentally use these options in a noncluster environment, the remote mirror operation does not execute. See the C <i>tag</i> description in this table for more information.
	Specifying -C <i>tag</i> limits any operation performed to the sets in the named <i>tag</i> cluster resource tag.

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 5-4](#).

TABLE 5-4 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 software.	Primary or secondary host	You can disable on one host, leave the other host enabled, and then re-enable the disabled host.
Enable the software.	Primary and secondary hosts	Perform this operation on both hosts if you are deleting a volume set.
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.
Logging mode.	Primary host	Perform on the primary host only if a synchronization is in progress.
	Secondary host	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 host	
Update an I/O group.	Primary and secondary hosts	

Enabling and Disabling Volume Sets

Use the following commands to enable and disable volume sets:

Function	Command	See This Section
Enable the software for the specified set (sync not required).	<code>sndradm -E</code>	“Enabling Volume Sets” on page 78
Enable the software for the specified set (sync required).	<code>sndradm -e</code>	“Enabling Volume Sets” on page 78
Disable the software for the specified set.	<code>sndradm -d</code>	“Disabling Volume Sets” on page 80
Add or delete a point-in-time copy software volume group.	<code>sndradm -I</code>	“Adding and Deleting Point-in-Time Copy Software Volumes” on page 81
Stop the software replication, enter logging mode, and track changes using bitmaps.	<code>sndradm -l</code>	“Stopping Replication and Starting Logging” on page 83

Enabling Volume Sets

Note – When you first enable the remote mirror software, you must enable it on the primary and secondary hosts. Always maintain the same [volume set field order](#) on primary and secondary hosts.

Use these commands to enable the remote mirror volume set (*volset*) and start logging changes to the volumes. Also use these commands to create one-to-many, many-to-one, and multihop volume sets. See [“One-to-Many, Many-to-One, and Multihop Volume Sets”](#) on page 10.

After You Enable Volume Sets

After you enable a volume set, you can synchronize or update the volumes. See:

- [“Starting a Full Synchronization \(Copy Between Volumes\)”](#) on page 86
- [“Starting an Update Resynchronization \(Update Between Volumes\)”](#) on page 88

Using a Volume Set Name

When you first enable a set, the 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 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 to indicate that a full synchronization between volumes is required and it enables logging. It also adds the local volume set (including the local bitmap volume) to the Storage Volume (SV) driver control.

Syntax

```
sndradm -e [-n] {-f volset-file | volset}
```

`sndradm -E`



Caution – Before using `sndradm -E`, ensure that the volumes are already synchronized by means other than the remote mirror 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 logging to indicate that the specified volumes are fully synchronized and enables logging. It also adds the local volume set (including the local bitmap volume) to the Storage Volume (SV) driver control.

Syntax

```
sndradm -E [-n] {-f volset-file | volset}
```

Disabling Volume Sets

Note – This command also removes the information for the specified set from the Sun StorEdge software configuration.

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

This command discards any active logging information in the bitmap volumes. See [“Enabling and Disabling Volume Sets” on page 77](#) to re-enable software replication and [“Starting a Full Synchronization \(Copy Between Volumes\)” on page 86](#) to re-establish identical replicated sets.

`sndradm -d`

When you issue this command, the 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 bitmaps that track temporary differences between primary and secondary volumes.

Syntax

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

Adding and Deleting Point-in-Time Copy Software Volumes

Note – Use the `sndradm -e` command to enable software volume sets first before using this command. [“Enabling and Disabling Volume Sets” on page 77](#) describes this command.

Use the `sndradm -I` command to add or delete point-in-time copy software volumes for use with enabled 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 or manual resynchronization.

The `sndradm -I a` command creates 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 remote mirror software synchronization daemon whenever a synchronization is started and waits for `sndrsyncd` to perform any necessary PIT copies before allowing the synchronization to proceed.

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

Consider the following when configuring this volume set:

- The master volume is the remote mirror secondary volume.
- The shadow and bitmap volumes cannot be a volume currently in use by the PIT or remote mirror software. If the volumes are in use, data loss can occur as the result of this misconfiguration.
- Before the reverse synchronization occurs, the point-in time copy software enables the volume set as a dependent set. Should the synchronization fail for any reason and the secondary volume data is lost or corrupted, you can restore the remote mirror secondary volume (the master volume in this set) by updating it with the shadow volume contents.

See also:

- [“Choosing Automatic or Manual Resynchronization” on page 28](#)
- [“Using the Remote Mirror Software With the Point-in-Time Copy Software” on page 29](#)
- [“Autosynchronization With the Point-in-Time Copy Software” on page 30.](#)

`sndradm -I a`

This command adds a point-in-time copy volume group entry to the Sun StorEdge configuration.

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 point-in-time copy volumes. In this case, the *master-vol* is the remote mirror secondary volume.



Caution – The shadow and bitmap volumes cannot be a volume currently in use by the point-in-time copy or remote mirror software. If the volumes are in use, data loss can occur as the result of this misconfiguration.

See the *Sun StorEdge Availability Suite 3.1 Point-in-Time Copy Software Administration and Operations Guide* for more information about these volume sets.

Note – Before the `/usr/opt/SUNWrdc/lib/sndrsyncd` daemon performs an update synchronization after a network link or machine failure, it creates a matching point-in-time copy volume group entry. After the update finishes executing, you can choose to delete this entry with the `/usr/opt/SUNWscm/sbin/iadm -d` command. However, if you do not delete this entry, the daemon can use it in subsequent resynchronizations involving the point-in-time copy software.

`sndradm -I d`

This command deletes a point-in-time copy volume group entry from the Sun StorEdge configuration.

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 point-in-time copy volumes.

Stopping Replication and Starting Logging

Note – To resume remote mirror software operations after using the `sndradm -l` command, use the `sndradm -u` command to perform an update resynchronization or the `sndradm -m` command to perform a full 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.

Use the `sndradm -l` command to start bitmap logging on enabled volume sets and I/O groups. Issue the `sndradm -l` command as follows:

- **From the primary host if a synchronization is in progress.**
When you place the primary host into logging mode, both the primary and the secondary hosts begin logging.
- **From the secondary host if the primary host or volume has failed.**
When you place the secondary host into logging mode, the primary host continues in replication mode (that is, the primary host is not automatically placed into logging mode).
- **From either host if there is no replication occurring.**

Tip – Place the secondary volume into logging mode and then mount any file systems in read-write mode to enable your application to write to the secondary volume. The remote mirror software logs the changes in the bitmap volume.

`sndradm -l`

This command enables you to stop replication between primary and secondary volumes and to start logging data changes on these volumes. The remote mirror software continues logging even 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 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 copy.

Syntax

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

Synchronizing and Updating Volume Sets

The following commands enable you to synchronize and update volume sets.

See also [“If an Initial Synchronization is Interrupted”](#) on page 85 and [“Performing a Reverse Copy or Update After a Primary Volume or Network Link Failure”](#) on page 85.

Function	Command	See This Section
Copy the entire contents of the primary volume to the secondary volume; also known as <i>full synchronization</i> .	<code>sndradm -m</code>	“Starting a Full Synchronization (Copy Between Volumes)” on page 86
Copy the entire contents of the secondary volume to the primary volume; also known as <i>full reverse synchronization</i> .	<code>sndradm -m -r</code>	“Starting a Full Synchronization (Copy Between Volumes)” on page 86
Update only the changed data of the primary volume to the secondary volume; also known as <i>update synchronization</i> or <i>resynchronization</i> .	<code>sndradm -u</code>	“Starting an Update Resynchronization (Update Between Volumes)” on page 88
Update only the changed data of the secondary volume to the primary volume; also known as reverse synchronization or reverse update.	<code>sndradm -u -r</code>	“Starting an Update Resynchronization (Update Between Volumes)” on page 88
Wait for a synchronization operation to finish executing before being able to execute other commands.	<code>sndradm -w</code>	“Waiting for Synchronization to Finish Executing” on page 90

If an Initial Synchronization is Interrupted

Consider the following if, after enabling the volume sets, the initial synchronization operation using the `sndradm -m` or `sndradm -u` commands is interrupted

If You Enabled the Volume Sets With This Command	And Used This Command to Synchronize Volume Sets	The <code>sndradm -u</code> Command Results In
<code>sndradm -e</code>	<code>sndradm -m</code> <code>sndradm -m -r</code>	Another full synchronization operation starting from the beginning.
<code>sndradm -E</code>	<code>sndradm -m</code> <code>sndradm -m -r</code>	The synchronization restarts from the point of interruption (that is, not a full synchronization starting from the beginning).

Performing a Reverse Copy or Update After a Primary Volume or Network Link Failure

As specified in [“Rehearsing Disaster Recovery” on page 56](#), perform the following commands after the network link or disk failure is repaired:

1. Ensure that any cached data at the secondary site is purged or synchronized by using the `scmadm -p` or `scmadm -s` command. See [Appendix A](#).
If you do not perform this step, when you attempt to perform a reverse copy or update operation, the software reports that the update failed and that the volume sets are in logging mode.
2. Place the primary and secondary host volumes into logging mode if they are not already in logging mode.
See [“Stopping Replication and Starting Logging” on page 83](#).

Note – If you check the volume status at this point using the `scmadm -S` command, the volume status might report VOF (volume failed) or RSN (reverse synchronization needed). When you perform the reverse copy or update, this volume status is cleared.

3. Perform the reverse copy or update.

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 finish executing the synchronization. Use `/usr/opt/SUNWscm/sbin/scmadm -S` to view the synchronization progress. You can also use the `sndradm -p`, `sndradm -P`, and `dsstat` commands to view volume set status. See also “[If an Initial Synchronization is Interrupted](#)” on page 85 and “[Performing a Reverse Copy or Update After a Primary Volume or Network Link Failure](#)” on page 85.

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

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

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

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

Note – You can use the primary volume during reverse synchronization (`sndradm -m -r`). The primary volume shows a consistent volume image of the latest data as soon as the reverse synchronization starts. If your application had been writing to the secondary volume as part of a failure or disaster rehearsal, you can move the application back to the primary volume when the reverse synchronization starts. See also “[Rehearsing Disaster Recovery](#)” on page 56.

`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 volset-file | volset | 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 volset-file | volset | set-name]
```

Starting an Update 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. You can also use the `sndradm -p`, `sndradm -P`, and `dsstat` commands to view volume set status.

See also “[If an Initial Synchronization is Interrupted](#)” on page 85 and “[Performing a Reverse Copy or Update After a Primary Volume or Network Link Failure](#)” on page 85.

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

- The primary volume has the desired contents.
- Replication was stopped by using the `sndradm -l` command or the network link was interrupted and bitmap logging was activated automatically.
- You want to overwrite those segments of the secondary volume that are different from the primary as defined in the bitmaps.
- 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:

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

Note – You can use the primary volume during reverse synchronization (`sndradm -u -r`). The primary volume shows a consistent volume image of the latest data as soon as the reverse synchronization starts. If your application had been writing to the secondary volume as part of a failure or disaster rehearsal, you can move the application back to the primary volume when the reverse synchronization starts. See also “[Rehearsing Disaster Recovery](#)” on page 56.

sndradm -u

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

Syntax

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

sndradm -u -r

Tip – You can also use this command to roll back volume changes on the primary volume from a point-in-time copy software 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 bitmaps maintained while replication was stopped. It also enables concurrent replication between the primary volume and secondary volumes; any new write operations to the primary volume are also replicated to the secondary volume.

Syntax

```
sndradm -u -r [-g io-groupname] [-C tag] [-n] [-f volset-file | volset | 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 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 commands

`sndradm -w`

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

Syntax

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

Reconfiguring Or Modifying a Volume Set

Note – Before reconfiguring or changing volume sets (except for changing the volume set bitmap), place the volume sets in logging mode by using the `sndradm -l` command.

The following commands enable you to reconfigure or modify a volume set.

Reconfigure or Change Volume Sets	Command	See This Section
Change the bitmap of a volume set.	<code>sndradm -R b</code>	“Changing the Bitmap Volume in a Volume Set” on page 92
Move all the specified remote mirror sets to a different group.	<code>sndradm -R g</code>	“Moving a Volume Set to a Different I/O Group” on page 93
Change the replication mode of a volume set.	<code>sndradm -R m {sync async}</code>	“Changing the Replication Mode of a Volume Set” on page 96
Update or reconfigure the software volume sets.	<code>sndradm -R -f volset-file</code>	“Updating the Volume Set Information” on page 98
Update or reconfigure the disk group name or cluster resource tag.	<code>sndradm -R C</code>	“Updating the Disk Cluster Tag Name” on page 99

Changing the Bitmap Volume in a Volume Set

Use the `sndradm -R b` command when you need to assign a new 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 remote mirror volume set. This command copies any data from the bitmap file to the bitmap volume. The *Sun StorEdge Availability Suite 3.1 Remote Mirror Software Installation Guide* describes this procedure.

`sndradm -R b`

This command enables you to assign a new 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 [-C tag] [-n] [-f volset-file | volset | 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.

Moving a Volume Set to a Different I/O Group

Note – Before reconfiguring or changing volume sets (except for changing the volume set bitmap), place the volume sets in logging mode by using the `sndradm -l` command. After you finish changing the sets, perform an update synchronization using the `sndradm -u` command.

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-names* on the command line.

See the following procedures:

- “To Remove Volume Sets From an I/O Group” on page 94
- “To Move Volume Sets to Another I/O Group” on page 95
- “To Remove Volume Sets from an I/O Group and Change the Replication Mode” on page 97

`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*. Issue this command from the primary and secondary hosts.

Syntax

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

▼ To Remove Volume Sets From an I/O Group

1. Place the I/O group volume sets into logging mode.

```
# sndradm -g io-groupname -l
```

2. Remove volume sets:

- To remove a volume set from an I/O group, use double quotes to specify a null group as follows:

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

- To 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 volset-file` option.

```
# sndradm -R g "" -g io-groupname -f volset-file
```

3. After you finish changing the sets, perform an update synchronization using the `sndradm -u` command.

▼ To Move Volume Sets to Another I/O Group

1. Place the I/O group volume sets into logging mode.

```
# sndradm -g io-groupname -l
```

2. Move volume sets:

- To move all volume sets in one I/O group to another I/O group:

```
# sndradm -g io-groupname -R g new-io-groupname
```

- To move selected sets to an I/O group, enter the volume set information in a file and use the `-f volset-file` option with these commands.

```
# sndradm -R g new-io-groupname -f volset-file
```

3. After you finish changing the sets, perform an update synchronization using the `sndradm -u` command.

Changing the Replication Mode of a Volume Set

Note – Before reconfiguring or changing volume sets (except for changing the volume set bitmap), place the volume sets in logging mode by using the `sndradm -l` command. After you finish changing the sets, perform an update synchronization using the `sndradm -u` command.

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

Note – You can remove one volume set, selected volume sets, or all volume sets from a group using the `sndradm -R g` command. See “[To Remove Volume Sets from an I/O Group and Change the Replication Mode](#)” on page 97 and “[Moving a Volume Set to a Different I/O Group](#)” on page 93.

`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 volset-file | volset | 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 Volume Sets from an I/O Group and Change the Replication Mode

1. Place the I/O group volume sets into logging mode.

```
# sndradm -g io-groupname -l
```

2. Choose one of the following commands:

- Remove the volume set from its group.

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

- To remove all volume sets from an I/O group:

```
# sndradm -R g "" -g io-groupname
```

3. Change the volume set replication mode of each volume set.

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

4. If desired, add the volume sets you have changed to a different or the previous I/O group.

See also [“Moving a Volume Set to a Different I/O Group”](#) on page 93.

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

5. After you finish changing the sets, perform an update synchronization using the `sndradm -u` command.

Updating the Volume Set Information

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

The fields for the volume set file specified using the `-f` option are similar to those in the *volset*, described in [TABLE 5-3](#):

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

See [“Using a Volume Set File” on page 43](#) and the `rdc.cf` man page for more information.

```
sndradm -R -f
```

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

Syntax

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

Updating the Disk Cluster Tag Name



Caution – Do not install or use the Sun StorEdge version 3.1 software on servers in a Sun Cluster 3.0 initial release environment. **The Version 3.1 software is not coexistent with the Sun Cluster 3.0 environment.** The Version 3.1 software is cluster-aware in the Sun Cluster 3.0 Update 1 and Update 2 environments 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 volset-file | volset | set-name]
```

Displaying Volume Set Information

Use the following commands to display information about volume sets.

Function	Command	See This Section
Display the command syntax information from the command prompt.	<code>sndradm -h</code>	See the on-screen display results from entering this command at the command prompt
Display the software version information from the command prompt.	<code>sndradm -v</code>	See the on-screen display results from entering this command at the command prompt
Display the software status.	<code>sndradm -p</code> <code>sndradm -P</code>	“Displaying Volume Set and I/O Group Status” on page 101
Display the software volume set and I/O group names.	<code>sndradm -i</code>	“Displaying a List of Current Volume Sets and I/O Group Names” on page 102
Display the status of the link connecting systems running the software.	<code>sndradm -H</code>	“Displaying Link Status” on page 103

Displaying Volume Set and I/O Group Status

Use the `sndradm -p` and `sndradm -P` commands to display status information. Use the `volset | set-name` options to view the status for a single volume set. 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`.

Note – Issue this command from the primary host to display information correctly.

`sndradm -p`

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

```
# sndradm -p
/dev/rdisk/c4t96d0s1    -> fast7:/dev/rdisk/c2t0d0s1
/dev/rdisk/c4t97d0s1    -> fast7:/dev/rdisk/c2t0d0s1
/dev/rdisk/c4t98d0s1    -> fast7:/dev/rdisk/c2t0d0s1
```

Syntax

```
sndradm -p [volset | set-name] [-g io-groupname] [-C tag]
```

`sndradm -P`

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

```
# sndradm -P

/dev/rdisk/c4t96d0s1    -> fast7:/dev/rdisk/c2t0d0s1
autosync: off, max q writes: 4194304, max q fbas: 16384, mode: sync

/dev/rdisk/c4t98d0s1    -> fast7:/dev/rdisk/c2t0d0s1
autosync: off, max q writes: 4194304, max q fbas: 16384, mode: sync
```

Syntax

```
sndradm -P [volset | set-name] [-g io-groupname] [-C tag]
```

Displaying a List of Current Volume Sets and I/O Group Names

Use the `sndradm -i` command to display volume set and I/O group names. This command displays volume set information in a volume set file format. See [Section , “Using a Volume Set File” on page 3-43](#).

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

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

```
# sndradm -i

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

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

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

Syntax

`sndradm -i [volset | set-name] [-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/rdisk/freddg/sndr_vol01
Report SNDR link health? (Y/N) [N]: y

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

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

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 remote mirror software assigns a default volume set name of *shost:sdev*.)

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

Enabling or Disabling the Autosynchronization Feature

Use the `sndradm -a` command to enable or disable autosynchronization. See [“Choosing Automatic or Manual Resynchronization” on page 28](#) for more information about autosynchronization.

Note – If you enable autosynchronization and later issue a synchronization command like `sndradm -u`, autosynchronization remains enabled.

`sndradm -a`

This command enables or disables the remote mirror software autosynchronization. If remote mirror software autosynchronization is enabled on the primary host, the synchronization daemon attempts to resynchronize volume sets 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 volset-file | volset | set-name]
```

Setting the Asynchronous Queue

Note – Ensure that you do not configure the queue size to be larger than available cache. As a rule of thumb, reserve 5 percent of the total cache.

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

See “[Tuning the Asynchronous Queue](#)” on page 106.

`sndradm -W`

This command enables you to set the maximum number of write operations to queue for asynchronous mode volume sets. The default *value* number is 4194304. For example, set this value to 1000 to ensure that the secondary volume is never more than 1000 write operations behind the primary volume.

Syntax

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

`sndradm -F`

This command enables you to set the maximum number of 512-byte blocks in the asynchronous queue. The default number is 16384, allowing about 8 Mbytes of data per queue.

Syntax

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

Tuning the Asynchronous Queue

Note – Ensure that you do not configure the queue size to be larger than available cache. As a rule of thumb, reserve 5 percent of the total cache.

The queues containing the remote mirror software asynchronous write operations are stored in system kernel memory. You can tune or adjust these queues so that they are correctly sized for your remote mirror software use and volume data set size. Left untuned, the queues can appear to slow performance and consume more system memory than is desirable.

Consider the following, in order, when tuning the asynchronous queues:

- Number of volume sets and whether a group contains a large number of volume sets in asynchronous mode (this group will share a single queue and might affect performance)
- Physical system memory
- Network latency

The number of volume sets is most important. Large queue settings matched with a small number of volume sets helps ensure better performance than large queue settings with a large number of volume sets. However, try not to size queues so large that you exhaust kernel memory and degrade your system's performance.

This section includes the following topics:

- [“Asynchronous Mode and the Queue” on page 107](#)
- [“Displaying the Queue Size” on page 108](#)
- [“Sample kstat Output For a Correctly-Sized Queue” on page 109](#)
- [“Sample kstat Output for an Incorrectly-Sized Queue” on page 110](#)

Asynchronous Mode and the Queue

The remote mirror software asynchronous mode helps when your network has sufficient bandwidth to handle many write transactions or when your replicating application (such as a database) performs burst-writes at a rate higher than the network limits allow.

Generally, you use asynchronous mode for a volume set when the I/O write rate is less than the throughput of the network you are using for replication. This mode provides low latency rates to the software, resulting in faster response at the host (see [“Asynchronous Replication” on page 17](#)).

As the number of write operations in the queue increases, they can exceed the ability of the queue and network to process them (depending on network bandwidth and latency). The remote mirror software appears slow and reports less than optimal write transaction performance. Unless you tune (or size) the queue correctly, asynchronous mode performance might appear to be the same as or worse than synchronous mode, depending on the volume data set size.

Consider changing the software’s default queue size if:

- The amount of system memory cannot support the number of asynchronous mode volume sets in the queue.
- The replicating application writes more data than the queue can store, impacting performance.
- You do not want the application’s write operations to become “out of sync” with the primary host, where pending write operations try to queue faster than the queue or network can handle them.
- You do not care how “out of sync” the secondary host is from the primary host, you only care about response time on the primary host.

Displaying the Queue Size

You can change the asynchronous queue size by using the `sndradm -W` and `sndradm -F` commands. The remote mirror software version 3.1 has set the following default sizes for the queue:

Default maximum number of write operations in the queue (default for the <code>sndradm -W</code> command)	4194304
Default maximum number of 512-byte data blocks (default for the <code>sndradm -F</code> command)	16384

▼ To Display the Current Queue Size

1. Type the following to display the current queue size:

```
# sndradm -P  
  
/dev/vx/rdisk/rootdg/ds4-clone -> nws:/dev/vx/rdisk/rootdg/ds4-sndr-s  
autosync: on, max q writes: 4194304, max q fbas: 16384, mode: async
```

The `kstat` command also displays queue information:

2. Perform any of the following `kstat(1M)` commands to show more information.

- To show all sets, type:

```
# kstat sndr::maxqfbas
```

- To show the first instance (0) in the queue, type:

```
# kstat sndr:0::maxqfbas
```

- To show more information:

```
# kstat sndr::maxqitems  
# kstat sndr::async_throttle_delay
```

Sample `kstat` Output For a Correctly-Sized Queue

Note – This example shows only a portion of the command output needed for this section; the `kstat` command actually displays more information.

The following `kstat(1M)` kernel statistics output shows information about the asynchronous queue. In this example, the queue is sized correctly.

```
# kstat sndr:0:setinfo
module: sndr                               instance: 0
name:   setinfo                             class:   storedge
        async_block_hwm                     878
        async_item_hwm                      483
        async_throttle_delay                0
        maxqfbas                            16384
        maxqitems                           4194304
        primary_host                        regina
        primary_vol                         rootdg/ds4-clone
```

This example shows the following settings and statistics:

<code>maxqfbas 16384</code>	Default maximum number of 512-byte blocks in the asynchronous queue. Set this using by <code>sndradm -F</code>
	The default setting allows about 8 Mbytes of data per queue.
<code>maxqitems 4194304</code>	Default maximum number of write operations to queue for asynchronous mode volume sets. Set this by using <code>sndradm -W</code> .
	According to this setting, each write consumes 2 Kbytes for the 8 Mbytes of data.
<code>async_block_hwm 878</code>	¹ A total of 878 512-bytes blocks (approximately 439 Kbytes) have been put into the queue.

<code>async_item_hwm</code>	483	¹ A total of 483 write transactions have been put into the queue.
<code>async_throttle_delay</code>	0	No delay in the queue; that is, the queue is not yet full.

¹ `async_block_hwm` and `async_item_hwm` show the maximum number that have been put into the queue since replication started. They do not show the current number that are in the queue.

Sample `kstat` Output for an Incorrectly-Sized Queue

Note – This example shows only a portion of the command output needed for this section; the `kstat` command actually displays more information.

The following `kstat(1M)` kernel statistics output shows information about the asynchronous queue, which is incorrectly sized.

```
# kstat sndr:4:setinfo
module: sndr                               instance: 4
name:   setinfo                             class:   storedge
        async_block_hwm                     16380
        async_item_hwm                      2045
        async_throttle_delay                16497
        maxqfbas                           16384
        maxqitems                          4194304
        primary_host                        andrea
        primary_vol                         rootdg/ds-forall
```

This example shows the default queue settings but the application is writing more data than the queue can handle. The `async_block_hwm` value of 16380 indicates that the application is approaching the maximum allowed limit of 512-bytes blocks. It is possible that the next few I/O operations are not going to be placed into the queue.

The value of `async_throttle_delay` indicates that the application has been delayed 16497 x 2 milliseconds. In this example, if the system has sufficient memory, increase the size of `maxqfbas` using the `sndradm -F` command.

Sun StorEdge Software and Storage Cache Statistics

This section describes the following topics to enable you to gather statistics about the Sun StorEdge software and cache:

- [“Remote Mirror and Point-in-Time Copy Software Statistics: The dsstat Command” on page 112](#)
- [“More Remote Mirror Software and Storage Cache Statistics: The scmadm Command” on page 119](#)

Remote Mirror and Point-in-Time Copy Software Statistics: The `dsstat` Command

The `/usr/opt/SUNWscm/sbin/dsstat` command displays I/O statistics for the remote mirror, point-in-time copy, and cache software. See the `dsstat` man page for more information.

Unless you specify otherwise, all fields displayed show per-second averages based on the data collected during the interval. For example, if you specify 5-second intervals, the data displayed is a per-second average of the data collected during the last 5-second interval.

Syntax

```
dsstat -m mode [-r report-options] [-d display-options] [-s volume-sets] [-f | -F] [-z] [interval [count]]
```

See [TABLE A-1](#) for the command options.

TABLE A-1 dsstat Options

Option	Description
<code>-m mode</code>	<p>Specifies the statistics reporting mode, where <i>mode</i> is:</p> <ul style="list-style-type: none"> • <code>cache</code> • <code>ii</code> • <code>sndr</code> <p>This option determines the options you can use with the <code>-r report-option</code>. You cannot use the <code>-r report-option</code> if you specify <code>-m cache</code>.</p> <p>If <code>-m</code> is not specified, the default is for the statistics for all modes to be displayed. If more than one mode is displayed, field headers are output for each report, overriding the <code>[-f -F]</code> option.</p>
<code>-r report-options</code>	<p>No <i>report-options</i> are available for the <code>-m cache</code> mode.</p> <ul style="list-style-type: none"> • If you specify <code>-m ii</code>, you can use this option as follows. If you do not specify <code>-r report-options</code>, the default display is <code>-r msbo</code>. <ul style="list-style-type: none"> <code>-r m</code> show master volume information <code>-r s</code> show shadow volume information <code>-r b</code> show bitmap volume information <code>-r o</code> show overflow volume information (if this volume type is attached) • If you specify <code>-m sndr</code>, you can use this option as follows. If you do not specify <code>-r report-options</code>, the default display is <code>-r bn</code>. <ul style="list-style-type: none"> <code>-r b</code> show bitmap volume information (this command shows only the bitmap of the host you are issuing the command from) <code>-r n</code> show network volume information
<code>-d display-options</code>	<p>Specifies the types of information to show. You can specify one or more of the following options on the command line. (If you specify <code>-d t</code>, then the <code>r</code> and <code>w</code> options are ignored.)</p> <ul style="list-style-type: none"> • If you specify <code>-m cache</code>, the available options are as follows. If you do not specify this option, the default display is <code>-d sf</code>. <ul style="list-style-type: none"> <code>-d r</code> detailed read information <code>-d w</code> detailed write information <code>-d d</code> destaged data information <code>-d c</code> cancelled write information <code>-d s</code> summary of information <code>-d f</code> cache behavior flags

TABLE A-1 dsstat Options (Continued)

	<ul style="list-style-type: none">• If you specify <code>-m sndr</code> or <code>-m ii</code>, the available options are as follows. If you do not specify this option, the default display for <code>-m ii</code> is <code>-d sf</code> and for <code>-m sndr</code> is <code>-d spf</code>.
	<code>-d r</code> detailed read information
	<code>-d w</code> detailed write information
	<code>-d t</code> timing information
	<code>-d s</code> summary of information
	<code>-d f</code> volume type and status flags
	<code>-d p</code> (<code>-m sndr</code> only) percentage of volume requiring a synchronization
	<code>-d a</code> (<code>-m sndr</code> only) shows the secondary volume names
<code>-s volume-set</code>	Shows information about the specified <i>volume-set</i> only. <i>volume-set</i> can also be a comma-delimited list of volume names.
<code>-f</code>	Displays a field header for every report cycle.
<code>-F</code>	Displays a field header once, when reporting begins.
<code>-z</code>	Do not report or display lines that have a value of zero (indicates no activity)
<code>interval [count]</code>	<i>interval</i> specifies how often the report is updated. The default interval is one second. <i>count</i> specifies how many reports to display. If <i>count</i> is not specified, report output continues at the specified <i>interval</i> time until you interrupt it by a break (^C).
	If you do not specify <i>interval</i> and <i>count</i> , the command shows a single report for a one-second interval.
	Unless you specify otherwise, all fields displayed show per-second averages based on the data collected during the interval. For example, if you specify 5-second intervals, the data displayed is a per-second average of the data collected during the last 5-second interval.

Examples

TABLE A-2 describes the fields displayed.

- Report cache statistics (`dsstat -m cache`)
- Show detailed information for read and write statistics (`-d rw`)
- Show information for the volume `c1t35d0s6` (`-s /dev/rdisk/c1t35d0s6`)
- Generate reports at five-second intervals (5)

Note – Only the trailing sixteen characters of a volume name appear in a report. For example: `/dev/rdisk/c1t1d0s0` is displayed as `ev/rdisk/c1t1d0s0`.

```
# dsstat -m cache -d rw -s /dev/rdisk/c1t1d0s0 5
```

volume	- read -			- write -		
	ckps	dkps	hit	ckps	dkps	hit
ev/rdisk/c1t1d0s0	0	0	0.00	0	0	0.00
ev/rdisk/c1t1d0s0	3	2396	0.13	983	763	100.00
ev/rdisk/c1t1d0s0	2399	799	75.00	2815	2686	100.00
ev/rdisk/c1t1d0s0	3200	800	80.00	2755	2908	100.00
ev/rdisk/c1t1d0s0	3999	799	83.33	2809	2868	100.00
ev/rdisk/c1t1d0s0	4800	800	85.71	2867	2931	100.00

- Report master and shadow volume point-in-time copy software statistics (`dsstat -m ii -r ms`)
- Generate reports at five-second intervals (5).

```
# dsstat -m ii -r ms 5
```

set name	t	s	- master -			- shadow -		
			kps	tps	svt	kps	tps	svt
ev/rdisk/c0t1d0s0	I	-	0	0	0	0	0	0
ev/rdisk/c0t1d0s0	I	-	9047	219	3	9040	219	7
ev/rdisk/c0t1d0s0	I	-	13548	317	3	9760	243	6
ev/rdisk/c0t1d0s0	I	-	5946	155	3	9684	227	8
ev/rdisk/c0t1d0s0	I	-	16539	417	2	9242	225	7

- Reports network statistics for the remote mirror software (-m sndr, -r n)
- Shows information for the secondary volume c1t35d0s6 (-s /dev/rdisk/c1t35d0s6)
- Shows detailed information for read and write statistics, timing, and volume type or status flags (-d rwtpf)
- Generates reports at five-second intervals (5)

```
# dsstat -m sndr -r n -d rwtpf -s /dev/rdisk/c1t1d0s0 5
```

- network -									
set name	t	s	sn	rkps	rtps	wkps	wtps	svt	
ev/rdisk/c1t1d0s0	P	L	77.27	0	0	0	0	0	
ev/rdisk/c1t1d0s0	P	SY	75.07	0	0	2317	72	11	
ev/rdisk/c1t1d0s0	P	SY	71.67	0	0	3443	108	9	
ev/rdisk/c1t1d0s0	P	SY	69.37	0	0	2426	76	12	
ev/rdisk/c1t1d0s0	P	SY	66.68	0	0	2765	86	11	

TABLE A-2 dsstat Display Fields

Fields	Description
set name	Volume set specified by <code>-s</code> option
t	Volume type for remote mirror volumes: P primary host volume S secondary host volume Volume type for the point-in-time copy volumes: I independent shadow volume set D dependent shadow volume set Valid cache modes C cache reads and writes D disk reads and writes
s	Volume status for the remote mirror software: L volume is in logging mode; changes are being logged SY forward synchronization is in progress RS reverse synchronization is in progress SN forward synchronization needed RN reverse synchronization needed VF volume failed BF bitmap failed FF network link failed Volume status for the point-in-time copy volumes: C copy in progress - no copy in progress
sn	Percentage of volume requiring sync
rtps	Number of reads
rkps	Kilobytes read
wtps	Number of writes
wkps	Kilobytes written
svt	Service time per operation
tps	Sum of rtps and wtps
kps	Sum of rkps and wkps
r	Cache read behavior

TABLE A-2 dsstat Display Fields

Fields	Description
w	Cache write behavior
ckps	Kilobytes read from cache
dkps	Kilobytes read from disk
hit	Read hits during interval
ds/s	Kilobytes destaged from cache
cn/s	Number of write cancellations

More Remote Mirror Software and Storage Cache Statistics: The `scmadm` Command

This command displays storage cache information and can be used to view remote mirror software statistics. See the `scmadm` man page for more information about this command.

Purging or Synchronizing the Cache

Use one of the following options with the `/usr/opt/SUNWscm/sbin/scmadm` command to purge or synchronize the cache.

```
scmadm -p { device-name | all }
```

This command discards the bad blocks for the specified *device-name* or for `all` disk devices on the system.

```
scmadm -s { device-name | all }
```

This command restores data on the specified *device-name* or for `all` disk devices on the system.

scmadm -S

After you type `/usr/opt/SUNWscm/sbin/scmadm -S` with the desired options, you can use the keys described in [TABLE A-3](#).

TABLE A-3 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 remote mirror software screens.
T	Display the cumulative key cache statistics generated since the last time the counters were reset (the <code>-z</code> option resets the counters).

During synchronization operations, the remote mirror 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 A-4 describes the `scmadm -S` options.

TABLE A-4 `scmadm -S` Options

Option	Description
-M	Displays the remote mirror software statistics. If <code>scmadm -S</code> is invoked without the -M option, it displays statistics related to the storage device cache.
-d <i>delay-time</i>	Sets the display update time to <i>delay-time</i> seconds.
-l <i>logfile</i>	Writes all screen outputs to the specified <i>logfile</i> .
-r [<i>range</i>]	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: <i>n[:n,n]</i> 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): -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 -M` Screen Display

The statistics display includes the fields described in [TABLE A-5](#).

Type `scmadm -S -M` to display the remote mirror software-specific statistics. Press the M key to toggle between the remote mirror software screens and regular cache statistics after typing the command.

TABLE A-5 Fields for Remote Mirror Software Screen Display

Field	Description
<code>primary</code>	Specifies the primary partition.
<code>link status</code>	Indicates whether replication is enabled or disabled, or indicates the synchronization direction if a synchronization is in progress.
<code>secondary</code>	Specifies the secondary partition.
<code>dual copy status</code>	Indicates current status using the following keywords: <code>logging</code> - Data is not being replicated and changes are being logged to the scoreboard. <code>need rev sync</code> - Logging, but an earlier reverse synchronization did not finish executing. The primary volume might be in an inconsistent state. The <code>need rev sync</code> state is cleared following a successful reverse synchronization, or by a point-in-time copy software update to the primary volume (which restores an earlier valid snapshot and removes the need to repeat the reverse synchronization to restore consistency). <code>need sync</code> - Logging, but an earlier synchronization did not finish executing. The secondary volume might be in an inconsistent state. The <code>need sync</code> state is cleared following a successful synchronization, or by a point-in-time copy software update to the secondary volume (which restores an earlier valid snapshot and removes the need to repeat the synchronization to restore consistency). <code>replicating</code> - Data written to the primary is being replicated to the secondary. Formerly called <code>enabled</code> . <code>rev sync</code> - A synchronization from secondary to primary is under way. <code>sync</code> - A synchronization from primary to secondary is under way. <code>volume down</code> - The volume has failed.
<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 [CODE EXAMPLE A-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 A-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/clt6d0s3    sync
ma:...dsk/c0t115d0s3    *>*           atmsi:...dsk/clt9d0s3    sync
  ma:...dsk/c0t116d0s3    *>*           atmsi:...dsk/clt6d0s3    sync
  ma:...dsk/c0t96d0s3    *=*           atmsi:...rdsk/clt0d0s3    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 A-6](#). Use the `M` key to toggle between these statistics and the remote mirror 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 A-2](#) and [CODE EXAMPLE A-3](#) show the cache statistics and cumulative cache statistics displays.

TABLE A-6 `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.
<code>write-blocks</code>	<code>dirty</code> —the number of dirty cache blocks that have not been queued for destaging. <code>todisk</code> —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</code> (writes).
<code>accesses/s</code>	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.
<code>read/s</code>	Displays the number of read requests (hits) serviced per second. <code>misses/s</code> —Displays the number of read misses per second.
<code>write/s</code>	Displays the number of write requests serviced per second. <code>misses/s</code> —Displays the number of write misses per second.
<code>%readh</code>	Displays the percent of reads that are cache hits.
<code>%writeh</code>	Displays the percent of writes that are cache hits.
<code>cachesize</code>	Displays the cache size.
<code>blocksize</code>	Displays the block size.
<code>Write blocks available</code>	Displays the available write cache in blocks. <code>Net n</code> in the field represents each network.

TABLE A-6 `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 A-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 writes  dirty 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 A-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

```


Glossary

Asynchronous replication	Asynchronous replication confirms to the originating host that the primary I/O transaction is complete before updating the remote image. That is, completion of the I/O transaction is acknowledged to the host when the local write operation is finished and the remote write operation has been queued. Deferring the secondary copy removes the long distance propagation delays from the I/O response time.
Configuration location	Location where the Sun StorEdge Availability Suite software stores configuration information about all enabled volumes used by the software.
Forward resynchronization	See Update synchronization .
Full synchronization	Full synchronization performs 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 disk might require reverse synchronization, using the surviving remote mirror as the source.
Logging	Mode where a bitmap tracks writes to a disk, rather than a running log of each I/O event. This method tracks 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 volume. The software uses this log to re-establish a remote mirror through an optimized update synchronization rather than a complete volume-to-volume copy.
Primary or local: host or volume	The system or volume on which the host application is principally dependent. For example, this is where the production database is being accessed. This data is to be replicated to the secondary by the software.

Replication Once a volume set has been initially synchronized, the software ensures that the primary and secondary volumes contain the same data on an ongoing basis. Replication is driven by user-layer application write operations; replication is an ongoing process.

Reverse synchronization An operation used during recovery rehearsals. 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.

Secondary or remote: host or volume The 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 might 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 the software mirroring.

Synchronous replication Synchronous replication is limited to short distances (tens of kilometers) because of the detrimental effect of propagation delay on I/O response times.

Update synchronization Update synchronization copies only those disk blocks identified by logging, reducing the time to restore remotely mirrored sets.

Volume set file A text file containing information about specific volume sets (see [“Using a Volume Set File” on page 43](#)). This text file is not the same as the [configuration location](#), which contains information about all configured volume sets used by the remote mirror and point-in-time copy software.

Index

A

- architecture, 3
- assign a new bitmap, 92
- asynchronous
 - queue, 105
 - replication, 17
- autosynchronization, 28, 55
 - state, 104
 - with the point-in-time copy software, 30

B

- bitmap
 - requirements, 41
 - size requirement, 5
 - volumes, 41

C

- cluster operations, 45
- commands
 - assign a new bitmap, 92
 - copy primary to secondary, 87
 - copy secondary to primary, 87
 - disable, 80
 - display
 - link status, 103
 - status, brief, 101
 - status, detailed, 101
 - volume and group name list, 102
 - dsstat, 112

- enable, 79
- refresh primary from secondary, 89
- refresh secondary from primary, 89
- reset the replication mode of a volume set, 96
- scmadm -S, 119
- set the asynchronous queue, 105
- stop remote mirror operations and start logging, 83
- syntax summary, 69
- toggle the autosynchronization state, 104
- update
 - configuration, 98
 - disk cluster tag name, 99
 - I/O group, 93
 - wait for synchronization to complete, 90
- configuration files
 - format, 43
- copy
 - initial volume copies, 51
 - primary to secondary, 86, 87
 - secondary to primary, 87
 - volume to volume, 20

D

- disable
 - volume sets, 80
- disaster recovery rehearsals, 56
- disk cluster tag name, 99
- displaying volume and group names, 102
- dsstat command, 112

- E**
 - enable
 - for the first time, 50
 - volume set, 77
 - examples
 - scmadm -S, 122, 124

- F**
 - failures
 - primary site, 60
 - file system replication, 4
 - firewall, 9
 - changing the port number of, 9
 - forward synchronization, 87
 - full synchronization, 19, 20

- G**
 - grouping volume sets, 7, 45

- H**
 - hosts
 - secondary, 63
 - swapping, 63
 - swapping the primary and secondary, 63

- I**
 - I/O groups, 7, 45
 - interruptions
 - resynchronizing after, 55

- L**
 - link failure, 54
 - logging
 - mode, 26
 - start, 83

- M**
 - many-to-one volume sets, 10
 - monitoring operations, 49
 - monitoring the software operations, 119
 - multihop volume sets, 10
 - mutual backup, 12

- N**
 - network link failure, 54
 - network protocols
 - supported, 9

- O**
 - one-to-many volume sets, 10
 - order-dependent writes, 11

- P**
 - point-in-time copy software
 - adding volumes for remote mirror use, 81
 - deleting volumes from remote mirror software use, 81
 - with the remote mirror software, 29
 - primary
 - site disaster, recovery, 61
 - site failures, 60

- R**
 - recovery, primary site disaster, 61
 - refreshing
 - primary from secondary, 89
 - secondary from primary, 89
 - rehearsals
 - disaster recovery, 56
 - remote mirror software
 - architecture, 3
 - replicating file systems, 4
 - replication
 - asynchronous, 17
 - examples, 32, 33, 35

- modes, 16
- synchronous, 16
- reset the volume set replication mode, 96
- restoring
 - primary from secondary, 88
 - primary from secondary, complete, 61
- resynchronization
 - choosing the type, 27, 28
 - forward and reverse update, 89
 - manual, 28, 55
 - when not to, 54
- reverse synchronization, 23
- reverse update synchronization, 24

S

- scmadm -S
 - example, 122, 124
 - options, 121
- secondary
 - restoring primary from, 61
- start logging, 83
- status
 - displaying, 101
- stop remote mirror operations and start logging, 83
- supported network protocols, 9
- synchronization
 - full, 19, 20, 52, 87
 - modes, 18
 - reverse, 23
 - reverse update, 24
 - update, 21, 22, 53, 89
- synchronous replication, 16

T

- toggle the autosynchronization state, 104

U

- update
 - bitmap in a volume set, 92
 - configuration, 98
 - disk cluster tag name, 99

- I/O group, 93
- update synchronization, 21, 22, 53

V

- volume set, 5
 - disable, 80
 - enable, 77, 78
 - enabling for the first time, 51
 - grouping, 11, 45
 - how many in configuration, 6
 - initial copies, 51
 - many-to-one, 10
 - multihop, 10
 - one-to-many, 10
 - raw partitions, 6
 - reset the replication mode, 96
 - size requirement, 5
- volume-to-volume copy, 20

W

- wait for synchronization to complete command, 90
- writes
 - order dependency, 11

