



Sun StorEdge™ Network Data Replicator 3.0 Configuration Guide

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

Part No. 806-7550-10
June 2001, Revision A

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

Copyright 2001 Sun Microsystems, Inc., 901 San Antonio Road • Palo Alto, CA 94303 USA. All rights reserved.

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

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

Sun, Sun Microsystems, the Sun logo, Sun StorEdge, Sun Enterprise, Solstice DiskSuite, Sun Trunking, SunScreen, AnswerBook, and Solaris are trademarks, registered trademarks, or service marks of Sun Microsystems, Inc. in the U.S. and other countries.

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

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

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

Copyright 2001 Sun Microsystems, Inc., 901 San Antonio Road • Palo Alto, CA 94303 Etats-Unis. Tous droits réservés.

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

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

Sun, Sun Microsystems, le logo Sun, Sun StorEdge, Sun Enterprise, Solstice DiskSuite, Sun Trunking, SunScreen, AnswerBook, et Solaris sont des marques de fabrique ou des marques déposées, ou marques de service, de Sun Microsystems, Inc. aux Etats-Unis et dans d'autres pays. Toutes les marques SPARC sont utilisées sous licence et sont des marques de fabrique ou des marques déposées de SPARC International, Inc. aux Etats-Unis et dans d'autres pays. Les produits portant les marques SPARC sont basés sur une architecture développée par Sun Microsystems, Inc.

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

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



Contents

Preface v

1. Overview 1

Sun SDR Software Description 1

Supported Hardware and Software 3

TCP/IP Connection Hardware Components 3

Network Multipathing 4

Applications 5

Architecture 5

2. Configuration Considerations 7

Volumes Eligible for Replication 7

Choosing Volume Level Protection 8

Bitmap Volumes for Scoreboard Logs 8

Choosing a Connection Medium to Link My Sites 9

ATM Link Advantages 9

Configuring Redundant Links Between Sites 10

Link Security 10

Choosing Between Synchronous and Asynchronous Replication 11

When To Suspend Replication to the Secondary Site 12

Configuring The Sun SNDR Software for Mutual Replication	12
Order-Dependent Writes and Volume Set Grouping	13
Recovery Considerations	14
Failing Over to the Secondary Site	14
Recovering the Primary or Secondary Site	14
Using The Sun SNDR and Sun StorEdge Instant Image Software	16
Sun SNDR Software and Sun StorEdge Fast Write Cache Software	17
One-to-Many and Multihop Volume Sets	18
Using The Sun SNDR Software in a Firewall Environment	19

Preface

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

Before You Read This Book

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

- *Sun StorEdge Network Data Replicator 3.0 System Administrator's Guide*
- *Sun StorEdge Network Data Replicator 3.0 Installation Guide*
- *Sun StorEdge Network Data Replicator 3.0 Release Notes*

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

Using UNIX Commands

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

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

Documentation Conventions

Typeface or Symbol	Meaning	Examples
AaBbCc123	The names of commands, files, and directories; on-screen computer output.	Edit your <code>.login</code> file. Use <code>ls -a</code> to list all files. % You have mail.
AaBbCc123	What you type, when contrasted with on-screen computer output.	% su Password:
<i>AaBbCc123</i>	Book titles, new words or terms, words to be emphasized. Command-line variable; replace with a real name or value.	Read Chapter 6 in the <i>User's Guide</i> . These are called <i>class</i> options. You <i>must</i> be <code>root</code> 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 -b {phost shost}</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 scmadm svadm iiadm	N/A
Latest information	<i>Sun StorEdge Network Data Replicator 3.0 Software Notes</i>	806-7513
	<i>Sun StorEdge Instant Image 3.0 Release Notes</i>	806-7678
Installation and User	<i>Sun StorEdge Network Data Replicator 3.0 Installation Guide</i>	806-7514
	<i>Sun StorEdge Instant Image 3.0 Installation Guide</i>	806-7675
	<i>SunATM 3.0 Installation and User's Guide</i>	805-0331
	<i>Sun ATM 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 Trunking 1.2.1 Installation and User's Guide</i>	806-4207

Application	Title	Part Number
	<i>Platform Notes: The Sun Quad FastEthernet Device Driver</i>	806-3989
	<i>Sun Enterprise 10000 InterDomain Networks User Guide</i>	806-4131
System administration	<i>Sun StorEdge Instant Image 3.0 System Administrator's Guide</i>	806-7677
	<i>Sun StorEdge Network Data Replicator 3.0 System Administrator's Guide</i>	806-7512
	<i>TCP/IP and Data Communications Administration Guide</i>	805-4003
	<i>SunScreen 3.1 Administration Guide</i>	806-4127
	<i>Solaris Bandwidth Manager 1.6 System Administration Guide</i>	806-4704
Configuration	<i>Sun Enterprise 10000 InterDomain Network Configuration Guide</i>	806-5230
	<i>Sun StorEdge Instant Image 3.0 Configuration Guide</i>	806-7676

Accessing Sun Documentation Online

A broad selection of Sun system documentation is located at:

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

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

<http://docs.sun.com>

For the latest version of released storage software documentation, go to:

<http://www.sun.com/products-n-solutions/hardware/docs/Software/>

Ordering Sun Documentation

For the latest version of storage software documentation, go to:

<http://www.sun.com/products-n-solutions/hardware/docs/Software/>

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

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

<http://www.fatbrain.com/documentation/sun>

Sun Welcomes Your Comments

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

docfeedback@sun.com

Please include Part No. 806-7550-10 in the subject line of your email.

Overview



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

Sun SNDR Software Description

The Sun SNDR software is a remote replication facility for the Solaris™ operating environment. It is intended for use as part of a disaster recovery and business continuance plan to provide redundant storage of critical information across physically separate sites. It is designed to be active during normal application access to the data volumes, and it will continually replicate the data to the remote site.

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

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

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

The replication mode can be synchronous or asynchronous based as a user-selectable parameter for each Sun SNDR volume set. (Use the `sndradm enable` command and select the `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 Sun SNDR software replicates changes to the secondary volume, keeping the two volumes consistent.

- In synchronous mode, a write operation is not confirmed as complete until the remote volume has been updated. Synchronous mirroring forces the Sun SNDR 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.
- In asynchronous mode, a write operation is confirmed as complete before the remote volume has been updated. Asynchronous mirroring allows the Sun SNDR 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, notification is sent to the primary.

In the event of planned or unplanned outages, the Sun SNDR software maintains per-device scoreboard bitmap volumes that are marked to indicate changed blocks with a granularity of 32 Kbytes per segment. This technique allows for optimized resynchronization by allowing the Sun SNDR software to resynchronize only the blocks that have changed.

Other features included as part of the Sun SNDR software are:

- **Grouping volume sets** — The Sun SNDR software enables you to group volume sets. Grouping volume sets guarantees write ordering. For example, in a synchronization from the primary host to the secondary host, write operations to the secondary volumes in the group occur in the same order as the write operations to the primary volumes in the group.

You can also assign specific volume sets to a group to perform `sndradm` commands on these volume sets and not on others you have configured. See “Order-Dependent Writes and Volume Set Grouping” on page 13.

Note – You can also group volume sets according to their cluster or resource tag to perform replication in a Sun Cluster 2.2 or 3.0 Update 1 clustered environment only.

- **One-to-many and multihop volume sets** — The Sun SNDR software enables you to create one-to-many and multihop volume sets. In a one-to-many volume set, you can replicate data from one primary volume to many secondary volumes residing on one or more hosts. In a multihop set, the secondary host volume of one volume set can be the primary host volume of another volume set. See “One-to-Many and Multihop Volume Sets” on page 18.

Supported Hardware and Software

Operating Environment	Solaris™ 7 or 8 operating environment or subsequent compatible version
Supporting Software	Network transport software such as SunATM™ or Gigabit Ethernet transports
Optional Software	Sun StorEdge Instant Image software Version 3.0; install this package for additional point-in-time capability
Supported Volume Manager Software	Sun Logical Volume Manager Sun Solstice DiskSuite™ Veritas Volume Manager
Host Types	Server hosts using the Solaris operating environment and any Sun-supported network interface card. Hosts include: <ul style="list-style-type: none"> • Sun Enterprise™ Server models 2x0 through 4x0 • Sun Enterprise Server models 3x00 through 10000
Supported Attached Storage	The Sun SNDR software is storage-hardware independent

TCP/IP Connection Hardware Components

The Sun SNDR software requires a TCP/IP connection between the primary and secondary server. Each server must have the proper ATM or Ethernet hardware installed to support the TCP/IP link.

The Sun SNDR software operates over any TCP/IP networking technology but has been qualified only on 10, 100, and 1000 Mbit Ethernet and ATM 166 and 622 technologies.

A dedicated TCP/IP link is not required.

Network Multipathing

Network multipathing with failover is supported on hosts running the Solaris 8, Update 3 operating environment. This feature is provided as part of the standard Solaris operating environment and is transparent to the Sun SNDR software.

The Sun SNDR software does not natively support the functionality listed in TABLE 1. However, other Sun products can be used with the Sun SNDR software to provide these facilities.

TABLE 1 Sun SNDR Software Related Functionality

Functionality	Sun Product
Network interface aggregation and network alternate pathing	Sun Trunking™ software and the native facilities provided in a Solaris operating environment update.
Network encryption	SunScreen™ Software
Network bandwidth allocation	Solaris Bandwidth Manager

Applications

The Sun SNDR software provides protection of critical enterprise information. The following information types are candidates for remote replication:

- Database partitions and logs
- Data that is important for business continuance in a disaster situation

The Sun SNDR software can be leveraged for additional applications, such as:

- Data migration
- Remote backup
- Remote data analysis
- Test and development

Architecture

The Sun SNDR software operates on logical volumes that can be linear, striped, or RAID devices. You can create the logical volumes by using the Sun Solstice DiskSuite™ or Veritas Volume Manager software.

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

FIGURE 1 shows the relationship between SNDR and the rest of the NSCTL data service framework.

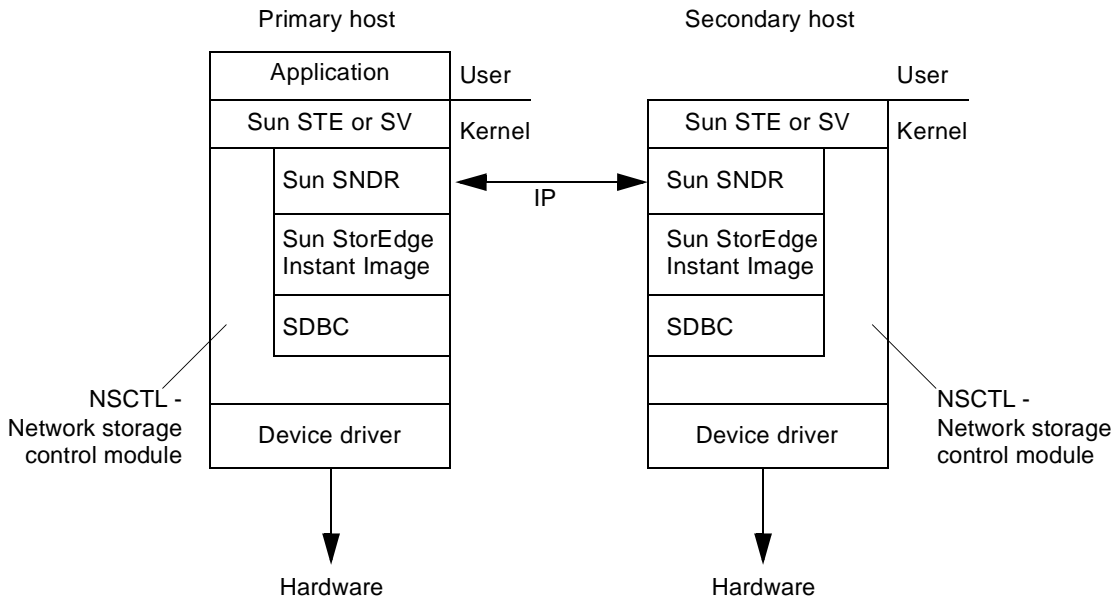


FIGURE 1 The Sun SNDR Software Architecture

Configuration Considerations

Volumes Eligible for Replication

Generally, include the following critical volumes in the Sun SNDR configuration:

- Database and database management system (DBMS) logs (the total database or online DBMS log)
- Access control files

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

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

When selecting a volume to be used in the Sun SNDR software volume set, ensure that volume does not contain disk label private areas (for example, slice 2 on a Solaris operating environment-formatted volume). The disk label region is contained in the first sectors of cylinder 0 of a disk. The safest method is to ensure that cylinder 0 is not part of any logical volume that is replicated (except for volumes under Veritas Volume Manager control, where cylinder 0 can be part of a logical volume that is replicated).

Choosing Volume Level Protection

In general, customers should protect their data with some level of RAID. With the Sun SNDR software, the primary volumes can be protected by any software or hardware RAID level desired. The protection level of the secondary volumes does not have to match that of the primary volumes. For example, if the primary volumes are RAID1, the secondary volumes could be RAID5.

Bitmap Volumes for Scoreboard Logs

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

The Sun SNDR Version 3.0 software does not support bitmap files. The Sun SNDR software uses regular raw devices to store bitmaps, used to maintain the Sun SNDR scoreboard logs. For example, you can specify a bitmap volume as `/dev/rdsk/c6t0d0s4`.

These raw devices should be stored on a disk separate from the disk that contains the data. Configure RAID (such as mirrored partitions) for these bitmap devices and ensure that the mirrored members are not stored on the same disk as the data.

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

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

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

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

Choosing a Connection Medium to Link My Sites

The primary and secondary sites must be connected by a standard medium that supports TCP/IP. For example, you can use any medium such as ATM, ISDN, Ethernet, Gigabit-Ethernet, or T1 and T3 lines to replicate data. Consider cost, bandwidth, and latency when choosing a connection medium.

The Sun SNDR software uses the maximum network bandwidth available to it. If the Sun SNDR software must share the network connection with other applications, you can use a bandwidth allocator to restrict the bandwidth available to an application.

ATM Link Advantages

ATM is a standards-based high-performance WAN protocol that establishes routes at circuit initialization. Its performance, combined with SONET/SDH's ring-healing resiliency, have made it the choice of carriers in their core backbone networks.

ATM Quality of Service parameters allow customers to choose the circuit characteristics that best suit their needs and budget. They define how different WAN ports competing for bandwidth should be prioritized relative to each other, each having "virtual paths" through the network fabric. Voice, video, and LAN traffic can move over the same connection as critical remote mirroring connections between primary and recovery sites. ATM switches make these varied protocol adaptations possible.

The Sun SNDR software will operate over any TCP/IP networking technology but has only been qualified on 10, 100, and 1000 Mbit Ethernet and ATM technologies.

Configuring Redundant Links Between Sites

The Solaris 8 operating environment offers IP multipathing. This feature permits a Solaris host to group multiple physical network interfaces into a single logical interface. Redundancy between paths permits transparent operation in the event of a failure anywhere on the path including failures of network interface cards, links, switches and routers.

This feature also provides the Sun SNDR software with the transparent means for surviving link failures. It provides the ability to aggregate bandwidth across the physical interfaces permitting a pair of ATM 622 Mbit links to operate as one logical 1244 Mbit link with resiliency.

Link Security

Solaris 8 operating environment offers IPsec, the industry standard form of integrating encryption and authentication mechanisms into standard IP internetworking. A Solaris 8 host can authenticate and encrypt traffic on a conversation-by-conversation basis.

Additionally, Sun offers SunScreen™ software technology that permits both encryption and authentication of IP traffic. Coupled with the Sun SNDR software, SunScreen software can provide an additional level of security to the data being transmitted between sites. SunScreen software uses public key encryption technology to guarantee that traffic is strongly encrypted. The public key mechanism is also used to authenticate that the sender of the traffic is the correct entity. SunScreen software is normally configured at the edge of the data center network. Of course, cooperating SunScreen software is required at each data center.

Choosing Between Synchronous and Asynchronous Replication

In synchronous mode, a write operation is not confirmed as complete until the remote volume has been updated. Generally, synchronous mirroring is limited to relatively short distances (that is, tens of kilometers) because of the detrimental effect of round-trip propagation delay on I/O response times. However, there are applications such as web server replication that can tolerate relatively slow remote updates and operate in a synchronous mode over long stretches.

With synchronous mirroring, the remote copy always matches the local host's view of committed writes. Should the local primary site be rendered inoperative, the remote secondary copy can be used to continue operations after the user community and the applications are switched to the alternate site.

What if the link goes down while running in synchronous mode? Should the local I/O remain incomplete until the link comes back up? Some applications require that the primary and secondary sites be exactly in sync and therefore would desire that the I/O not complete. However, many mission-critical customers decide not to block local access to the data upon encountering intersite transmission errors. Instead, software keeps track of these writes at the primary site and then updates the recovery site when the remote service is reliably restored. If configured this way, disaster recovery procedures must take into account loss of *in-flight* and logged data. That is, data might have been committed to the local application but never had a chance to arrive at the backup site.

Asynchronous mirroring affirms primary I/O completion to the originating host before updating the remote image. This type of mirroring is often used when the distance and relatively low bandwidth telecommunications lines of 45 Mbps or less between primary and secondary sites would introduce prohibitive latencies if performed synchronously. Here, the long-distance pipe becomes the bottleneck, forcing local writes to be queued for later transmission. Consequently, there is a higher possibility of losing buffered and in-flight data if the primary system fails.

When To Suspend Replication to the Secondary Site

For many customers, the Sun SNDR software is used to mirror the primary site's information onto the alternate location as quickly as possible under the constraints of distance and technology. Operator errors, software bugs, and hardware data corruption are solid arguments for taking more deliberate, albeit less real-time, approaches to remote copies.

A deferred-write policy under such conditions delays remote updates until the local data has been confirmed to be good. The data integrity checks are performed periodically, followed by the remote copy. Some IT organizations reason that it is better to have slightly stale, good data at the remote site than to risk corrupting the remote copy with the lock-step updates. Consciously scrutinized database redo logs facilitate the recovery and update process, while minimizing the consequences of the errors.

Another method is to ship the data remotely in real time, take a snapshot of the data at the remote site with Sun StorEdge Instant Image software at a known point in time, and verify its integrity in the background before setting it aside for recovery.

Telco costs might similarly dictate a deferred copy policy. Some customers must timeshare narrow pipes between sites during the peak processing hours, leaving only off-hours for bulk intersite copies. Such prospects look for incremental update capabilities to reduce the transmission load and copy time. In other words, they prefer to send only what has changed since the previous remote image was generated.

Configuring The Sun SNDR Software for Mutual Replication

Sometimes the distinction between primary (A) and secondary (B) sites is blurred. As applications are geographically distributed, a secondary storage system at Site B might function as a remote volume replication to Site A and as a direct storage resource for applications on Host B. Under these circumstances, you might also elect to have critical applications running on Host B to have their volumes replicated to Site A. This reciprocal replication arrangement supported by the Sun SNDR software is known as mutual replication.

Each server can concurrently transmit and receive writes to and from its remote counterpart. Each system contains primary disks in a Sun SNDR software set that are accessible by local hosts, as well as remote mirrors secondary to remote hosts. Mutual replication might be used where critical applications and storage are split across sites, and both sites require remote replicated volumes.

Order-Dependent Writes and Volume Set Grouping

For the Sun SNDR Version 2.0 software, write ordering is only maintained by the Sun SNDR software within a logical volume. In Version 3.0, write ordering is also maintained for *groups* of asynchronously replicating volume sets. (The general definition of write ordering here 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 point-in-time copy of the group of source volumes.

This feature is especially 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 all Sun SNDR software volumes associated with that application as SYNC
- Using Sun StorEdge Instant Image software to take periodic recoverable point-in-time copies

If you use Instant Image software, the remote point-in-time 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.

Recovery Considerations

The *Sun StorEdge Network Data Replicator 3.0 System Administrator's Guide* contains information about various recovery scenarios. This section describes the following topics:

- Failing Over to the Secondary Site
- Recovering the Primary or Secondary Site
- Using The Sun SNDR and Sun StorEdge Instant Image Software

Failing Over to the Secondary Site

Should the primary site fail, the secondary copy can be used to continue operations after the user community and the applications are switched to the alternate site. You should implement application level recovery procedures at the secondary site to help ensure application recovery to a well-known state.

Recovering the Primary or Secondary Site

Occasionally, remote mirroring operations are interrupted either intentionally or by unplanned outages. Normally, these interruptions are handled by the Sun SNDR software because whenever a member of a Sun SNDR software volume pair becomes unavailable, the write activity is scoreboarded in the volume's bitmap. When the service is restored, a fast update or resynchronization is performed.

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 -m -u update` command to complete the synchronization. Use `/usr/opt/SUNWscm/sbin/scmadm -S` to view the synchronization progress.

An update resynchronization applies changes that occurred at the primary data to the secondary site during the time 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 only done when the Sun SNDR software volume set is:

- Initially established
- Damaged due to a disaster
- Data integrity is questioned

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

The Sun SNDR software volume set data might be in question if, for example, there was a double disk failure on a RAID 5 set or if the Sun SNDR software is shutdown manually and writes occurred to either the primary or secondary volumes without scoreboarding being active. The full copy can either be performed from the primary to the secondary or if appropriate, from the secondary to the primary.

Choosing Automatic or Manual Resynchronization

The Sun SNDR software provides a choice of automatic or manual resynchronization after a disaster. Automatic resynchronization is discouraged if the interruption is the warning of a larger rolling disaster. It is best to 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 also “Using The Sun SNDR and Sun StorEdge Instant Image Software” on page 16. Before you start a resynchronization operation, ensure that you have an appropriate Instant Image software point-in-time copy of the Sun SNDR software target volume.

Using The Sun SNDR and Sun StorEdge Instant Image Software

To help ensure the highest level of data integrity on both sites during normal operations or during data recovery, use the Sun SNDR software with the Sun StorEdge Instant Image software.

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

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

When to Use The Sun StorEdge Instant Image Software at the Remote Site

In most cases, customers should also install Instant Image software at the secondary site to protect against a dual failure. Before you resynchronize the primary site with the secondary, take a point-in-time snapshot copy of the data. If 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.

Whenever a fast resynchronization is being performed, the I/O transfer is in block order, not in I/O order. Therefore the volume being resynchronized will not be consistent until the fast resynchronization has been completed and the volume is back in replicating mode. If a point-in-time copy is not available at the remote site, then the remote site recovery plan should cover the exposure that will exist during the time required to perform the fast resynchronization.

Additionally, a point-in-time snapshot at the secondary site can be used for applications such as remote backup, remote data analysis, or other functions.

Sun SNDR Software and Sun StorEdge Fast Write Cache Software

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

The Version 2.0 and Version 3.0 Sun StorEdge data services are binary incompatible. If your system includes Version 2.0 of the Sun StorEdge Instant Image (including Instant Image 2.0.1 with STE 1.2), Sun SNDR software, or Sun StorEdge Fast Write Cache, you must remove them before installation. For example, you cannot use Sun StorEdge Instant Image software Version 2.0 with Sun SNDR software Version 3.0. When you plan to install or upgrade to a Version 3.0 data service, you must remove all Version 2.0 and 2.01 data services.

However, the Sun StorEdge Core Services Version 3.0 Software CD contains the Sun StorEdge SUNWnvm Version 3.0 software package for those users whose systems include Version 2.0 of the Sun FWC hardware and software product and who wish to continue using the Sun FWC product.

For optimum performance in a noncluster environment, you can use the Sun SNDR software with the Sun StorEdge Fast Write Cache SUNWnvm Version 3.0 product on the primary and secondary sites for the volumes containing bitmaps and any other volumes for which write performance is critical (for example, database log files). Using the FWC product allows the scoreboard to be preserved across a system crash without adversely affecting performance during replication.

One-to-Many and Multihop Volume Sets

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

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

Note – You can group one-to-many sets that share a common primary volume in a single I/O group to operate on all sets simultaneously instead of issuing a separate command for each set. See also “Order-Dependent Writes and Volume Set Grouping” on page 13.

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

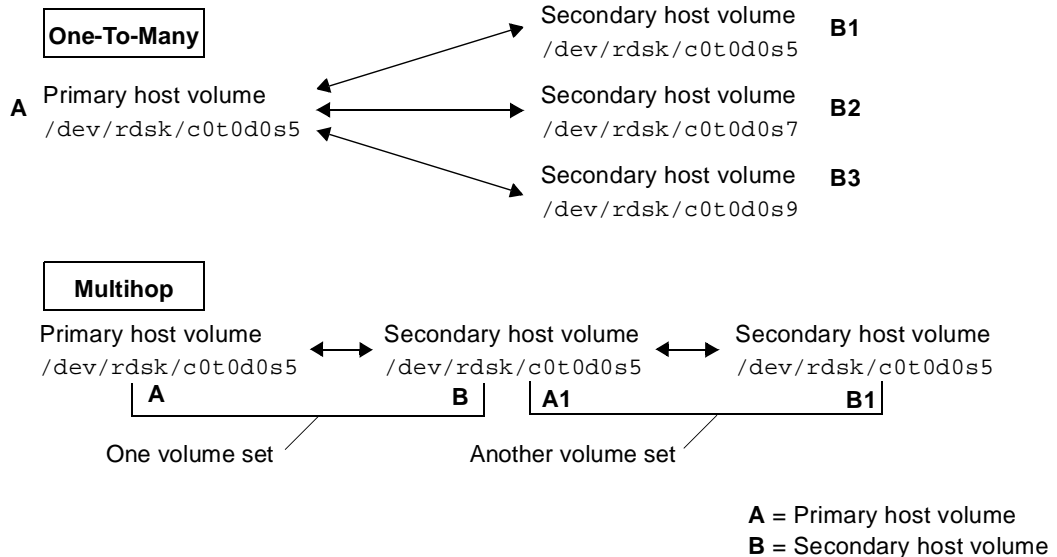


FIGURE 1 One-to-Many and Multihop Volume Sets

Using The Sun SNDR Software in a Firewall Environment

You can use the Sun SNDR software with any enhanced network product that Sun offers. Port 121 must be available for use by the Sun SNDR `sndrd` daemon.

You can use IP masquerading. SNDR 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.

