



Sun™ Cluster 3.0/3.1 and Sun StorEdge™ Availability Suite 3.2 Software Integration Guide

Sun Microsystems, Inc.
www.sun.com

Part No. 817-4224-10
December 2003, [Revision 52](#)

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

Copyright© 2003 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 this product. 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, 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 Adobe® logo is a registered trademark of Adobe Systems, Incorporated.

U.S. Government Rights—Commercial use. Government users are subject to the Sun Microsystems, Inc. standard license agreement and applicable provisions of the FAR and its supplements.

Products covered by and information contained in this service manual are controlled by U.S. Export Control laws and may be subject to the export or import laws in other countries. Nuclear, missile, chemical biological weapons or nuclear maritime end uses or end users, whether direct or indirect, are strictly prohibited. Export or reexport to countries subject to U.S. embargo or to entities identified on U.S. export exclusion lists, including, but not limited to, the denied persons and specially designated nationals list is strictly prohibited.

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© 2003 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, Californie 95054, Etats-Unis. Tous droits réservés.

Sun Microsystems, Inc. a les droits de propriété intellectuelle relatants à la technologie qui est décrit dans ce document. En particulier, et sans la limitation, ces droits de propriété intellectuelle 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, 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.

Ce produit est soumis à la législation américaine en matière de contrôle des exportations et peut être soumis à la réglementation en vigueur dans d'autres pays dans le domaine des exportations et importations. Les utilisations, ou utilisateurs finaux, pour des armes nucléaires, des missiles, des armes biologiques et chimiques ou du nucléaire maritime, directement ou indirectement, sont strictement interdites. Les exportations ou réexportations vers les pays sous embargo américain, ou vers des entités figurant sur les listes d'exclusion d'exportation américaines, y compris, mais de manière non exhaustive, la liste de personnes qui font objet d'un ordre de ne pas participer, d'une façon directe ou indirecte, aux exportations des produits ou des services qui sont régis par la législation américaine sur le contrôle des exportations et la liste de ressortissants spécifiquement désignés sont rigoureusement interdites.

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 QUALITÉ MARCHANDE, A L'APTITUDE A UNE UTILISATION PARTICULIÈRE OU A L'ABSENCE DE CONTREFAÇON.



Contents

| | |
|---|----------|
| Preface | v |
| 1. Overview | 1 |
| Terminology Used in This Guide | 2 |
| Supported Software and Hardware | 3 |
| Using the Sun StorEdge Availability Suite Software in a Sun Cluster Environment | 4 |
| Global and Local Use of the Sun StorEdge Availability Suite Software | 5 |
| Switching Over Global Devices Only | 5 |
| Volumes Eligible for Use | 6 |
| VTOC Information | 7 |
| 2. Installing and Configuring the Sun StorEdge Availability Suite Software | 9 |
| Shutting Down Nodes | 10 |
| Overview of Installation Tasks | 10 |
| Choosing the Configuration Location | 11 |
| Installing the Software | 13 |
| ▼ To Install the Software. | 14 |
| Editing the Bitmap Parameter Files | 16 |
| Setting the Bitmap Operation Mode | 16 |
| The <code>/usr/kernel/drv/ii.conf</code> File | 17 |

| | | |
|-----------|---|-----------|
| ▼ | To Edit the <code>ii.conf</code> File | 17 |
| | Shutting Down and Restarting Nodes | 18 |
| ▼ | To Shut Down and Restart a Node | 18 |
| | Supported Configurations for The Remote Mirror Software | 19 |
| | Adding Host Names | 19 |
| ▼ | Edit the <code>/etc/hosts</code> File | 19 |
| | Using Autosynchronization | 19 |
| | Rules For the Remote Mirror Software | 20 |
| | Remote Mirror Primary Host Is On a Cluster Node | 21 |
| | Remote Mirror Secondary Host On a Cluster Node | 21 |
| | Remote Mirror Primary and Secondary Hosts On a Cluster Node | 22 |
| | Supported Configurations for the Point-in-time Copy Software | 22 |
| | Rules For the Point-in-Time Copy Software | 22 |
| | Disk Device Groups and the Sun StorEdge Availability Suite Software | 23 |
| | Configuring the Sun Cluster Environment | 23 |
| ▼ | Configuring Sun Cluster for HAStorage or HAStoragePlus | 24 |
| | Configuring the HAStoragePlus Resource Types with Volume Sets | 28 |
| 3. | Using the Sun StorEdge Availability Suite <code>iiadm</code> and <code>sndradm</code> Commands | 29 |
| | Mounting and Replicating Global Volume File Systems | 30 |
| | Global Device Command Syntax | 31 |
| | Remote Mirror Example | 32 |
| | Point-in-Time Copy Example | 32 |
| | Local Device Command Syntax | 33 |
| | Point-in-Time Copy Example | 33 |
| | Which Host Do I Issue Commands From? | 35 |
| | Putting All Cluster Volume Sets in an I/O Group | 37 |
| | Preserving Point-in-time Copy Volume Data | 39 |

Preface

The *Sun™ Cluster 3.0/3.1 and Sun StorEdge™ Availability Suite 3.2 Software Integration Guide* describes how to integrate the Sun StorEdge™ Availability Suite 3.2 Remote Mirror and Point-in-Time Copy software products in Sun™ Cluster 3.0 Update3 and Sun Cluster 3.1 environments.

Note – The Sun StorEdge Availability Suite 3.2 Remote Mirror and Point-in-Time Copy software products are supported only in the Sun Cluster 3.0 Update3 and Sun Cluster 3.1 initial release environments.

This guide is intended for system administrators who have experience with the Solaris™ operating environment, Sun Cluster software, and related disk storage systems.

Before You Read This Book

Note – Before you install the Sun StorEdge Availability Suite software as described in the installation and release documentation in “[Related Documentation](#)” on page viii, see [Chapter 2](#).

To fully use the information in this document, you must have thorough knowledge of the topics discussed in the books in “[Related Documentation](#)” on page viii.

How This Book Is Organized

[Chapter 1](#) is an overview of the Sun Cluster and Sun StorEdge Availability Suite software integration.

[Chapter 2](#) describes installing and configuring the Sun StorEdge Availability Suite software for use in a Sun Cluster environment.

[Chapter 3](#) describes using the Sun StorEdge Availability Suite software commands in a Sun Cluster environment.

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 the following for this information:

- Software documentation that you received with your system
- Solaris™ operating environment documentation, which is at

<http://docs.sun.com>

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 | # |

Typographic Conventions

| Typeface ¹ | 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> <code>Password:</code> |
| <i>AaBbCc123</i> | Book titles, new words or terms, words to be emphasized. Replace command-line variables with real names or values. | 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> . |

¹ The settings on your browser might differ from these settings.

Related Documentation

| Application | Title | Part Number |
|-----------------------|---|--------------------|
| Hardware | <i>Sun Cluster 3.0 U1 Hardware Guide</i> | 806-7070 |
| | <i>Sun Cluster 3.0 12/01 Hardware Guide</i> | 816-2023 |
| Software Installation | <i>Sun Cluster 3.0 U1 Installation Guide</i> | 806-7069 |
| | <i>Sun Cluster 3.0 12/01 Software Installation Guide</i> | 816-2022 |
| | <i>Sun StorEdge Availability Suite Software Installation Guide</i> | 817-2783 |
| Data Services | <i>Sun Cluster 3.0 U1 Data Services Installation and Configuration Guide</i> | 806-7071 |
| | <i>Sun Cluster 3.0 12/01 Data Services Installation and Configuration Guide</i> | 816-2024 |
| Concepts | <i>Sun Cluster 3.0 U1 Concepts</i> | 806-7074 |
| | <i>Sun Cluster 3.0 12/01 Concepts</i> | 816-2027 |
| Error Messages | <i>Sun Cluster 3.0 U1 Error Messages Manual</i> | 806-7076 |
| | <i>Sun Cluster 3.0 12/01 Error Messages Manual</i> | 816-2028 |
| Release Notes | <i>Sun Cluster 3.0 U1 Release Notes</i> | 806-7078 |
| | <i>Sun Cluster 3.0 12/01 Release Notes</i> | 816-2029 |
| | <i>Sun Cluster 3.0 U1 Release Notes Supplement</i> | 806-7079 |
| | <i>Sun Cluster 3.0 12/01 Release Notes Supplement</i> | 816-3753 |
| | <i>Sun StorEdge Availability Suite 3.2 Software Release Notes</i> | 817-2782 |
| | <i>Sun" Cluster 3.0/3.1 and Sun StorEdge" Availability Suite 3.2 Software Release Note Supplement</i> | 817-4225 |
| System Administration | <i>Sun Cluster 3.0 U1 System Administration Guide</i> | 806-7073 |

| Application | Title | Part Number |
|-------------|--|-------------|
| | <i>Sun Cluster 3.0 12/01 System Administration Guide</i> | 816-2026 |
| | <i>Sun StorEdge Availability Suite 3.2 Remote Mirror Software Administration and Operations Guide</i> | 817-2784 |
| | <i>Sun StorEdge Availability Suite 3.2 Point-in-time Copy Software Administration and Operations Guide</i> | 817-2781 |

Accessing Sun Documentation

You can view, print, or purchase a broad selection of Sun documentation, including localized versions, at:

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

Contacting Sun Technical Support

If you have technical questions about this product that are not answered in this document, go to:

<http://www.sun.com/service/contacting>

Sun Welcomes Your Comments

Sun is interested in improving its documentation and welcomes your comments and suggestions. You can submit your comments by going to:

<http://www.sun.com/hwdocs/feedback>

Include the title and part number of this document with your feedback: *Sun Cluster 3.0/3.1 and Sun StorEdge Availability Suite 3.2 Software Integration Guide*, 817-4224-10

Overview

This guide assumes that you have already installed the volume manager software and the Sun Cluster software on each node in your cluster.

Note – The Sun StorEdge Availability Suite 3.2 Remote Mirror and Point-in-Time Copy software products are supported only in the Sun Cluster 3.0 Update3 and Sun Cluster 3.1 environments.

The Sun Cluster and Sun StorEdge Availability Suite 3.2 software combine to provide a highly available environment for cluster storage. The Remote Mirror software is a data replication application that provides access to data as part of business continuance and disaster recovery plans. The Point-in-Time Copy software is a point-in-time snapshot copy application that enables you to create copies of application or test data.

The topics in this chapter include:

- [“Terminology Used in This Guide” on page 2](#)
- [“Supported Software and Hardware” on page 3](#)
- [“Using the Sun StorEdge Availability Suite Software in a Sun Cluster Environment” on page 4](#)
- [“VTOC Information” on page 7](#)

Terminology Used in This Guide

Data service

Highly Available (HA) applications within the Sun Cluster environment are also known as data services. The term *data service* is used to describe a third-party application that has been configured to run on a cluster rather than on a single server. A data service includes the application software and Sun Cluster software that starts, stops, and monitors the application.

Primary and secondary hosts and nodes

In this guide and the Remote Mirror software documentation, the terms *primary host* and *secondary host* are used as follows.

The primary and secondary hosts are physically-separate servers running the Remote Mirror software. The primary host contains the primary volume and bitmap volume to be initially replicated to a remote server called a secondary host. The secondary hosts contains the secondary volume and bitmap volume

The terms *primary node* and *secondary node* refers to cluster nodes with respect to device group mastering in a cluster.

Supported Software and Hardware

TABLE 1-1 Supported Software and Hardware

| | |
|--|---|
| Operating Environment Software | Solaris™ 8 and Solaris 9 Update 3 and higher; all releases that are supported by the Sun Cluster 3.0 Update 3 software |
| Sun Cluster Software | Sun Cluster 3.0 05/02 software (also known as Update 3) |
| Volume Manager Software | Solstice DiskSuite, Solaris Volume Manager VERITAS Volume Manager (VxVM) The Sun StorEdge software does not support metatrans (metapartition) devices created by using the Sun Solstice DiskSuite and Solaris Volume Manager. |
| Sun StorEdge Software | Sun StorEdge Availability Suite 3.2 Remote Mirror and Point-in-Time Copy software. |
| Supported Cluster Configuration | The Sun Cluster 3.0 Update 3 release, the Sun Cluster 3.1 initial release, and the Sun StorEdge Availability 3.2 software are supported in a two-node cluster environment only. |
| Hardware | If you plan to install the software from the product CD, a CD-ROM drive connected to the host server where the software is to be installed. |
| Disk space requirements | Disk space requirements: 15 Mbytes <ul style="list-style-type: none">• The Remote Mirror software requires approximately 1.7 Mbytes• The Point-in-Time Copy software requires approximately 1.9 Mbyte• The Sun StorEdge configuration location requires 5.5 Mbytes• Supporting Sun StorEdge core packages require approximately 5.4 Mbytes |

Using the Sun StorEdge Availability Suite Software in a Sun Cluster Environment

To use cluster failover features with the Sun StorEdge Availability Suite 3.2 software, your software environment requires the Sun Cluster 3.0 Update 3 software or the Sun Cluster 3.1 initial release software.

In this environment, the Sun StorEdge Availability Suite software is *cluster aware*. See TABLE 1-2.

The `sndradm` and `iiadm` commands are used to control the Sun StorEdge Availability Suite software. You can use the command options `C tag` and `-C tag` in a cluster environment only. If you accidentally use these options in a noncluster environment, the specified operation does not execute. See [Chapter 3](#) in this guide for more information.

TABLE 1-2 Cluster Terminology and Status

| Term | Definition | Sun StorEdge Availability Suite Status |
|--------------------------------|---|--|
| Cluster aware | A software product is Sun Cluster aware if it can coexist with the Sun Cluster environment and fails over and fails back as the logical host containing the software product fails over and fails back. A Sun Cluster aware product can then be made highly available by utilizing the High Availability framework that Sun Cluster provides. | The Sun StorEdge Availability Suite 3.2 software is cluster-aware in a two-node, Sun Cluster 3.0 Update 3 software environment or a Sun Cluster 3.1 initial release environment. |
| Cluster tolerant or coexistent | A software product is Sun Cluster tolerant if it can coexist with the Sun Cluster environment and does not interfere with the Sun Cluster software and applications running in this environment. A product that is cluster tolerant is not expected to fail over or fail back when a Sun Cluster logical host fails over and fails back. | The Sun StorEdge Availability Suite 3.2 software <i>is not cluster-tolerant</i> in the initial release of the Sun Cluster 3.0 software. |

Global and Local Use of the Sun StorEdge Availability Suite Software

Note – See [“Rules For the Remote Mirror Software”](#) on page 20 and [“Rules For the Point-in-Time Copy Software”](#) on page 22

The Sun StorEdge Availability Suite software can use volumes that are local or global devices. Global devices are those Sun StorEdge Availability Suite software or other volumes accessible from any cluster node and which will fail over under the control of the Sun Cluster framework. Local devices are volumes that are local to the individual node (host machine), not defined in a disk device or resource group, and not managed within a cluster file system. Local devices do not fail over and switch back.

To access local devices, use the `C local` or `-C local` options as part of the `sndradm` commands, or the `-C local` option with `iiadm` commands. To access global devices, use the command options `C tag` and `-C tag`. (Typically, you do not need to specify the `-C tag` option as `iiadm` and `sndradm` automatically detect the disk device group.)

See [Chapter 3](#) in this guide and the Sun StorEdge Availability Suite administration and operations guides listed in [“Related Documentation”](#) on page viii.

Switching Over Global Devices Only

The `scswitch(1M)` command enables you to change all resource groups and device groups manually from the primary mastering node to the next preferred node. The Sun Cluster documentation describes how to perform these tasks.

Local devices do not fail over and switch back so do not configure them as part of your cluster. A file system mounted on a volume and designated as a local device must not be configured as a device to fail over and switch back in the Sun Cluster environment.

Volumes Eligible for Use

Note – When creating shadow volume sets, do not create shadow or bitmap volumes using partitions that include cylinder 0 because data loss might occur. See [“VTOC Information” on page 7](#).

You can replicate the following critical volumes using the Remote Mirror software:

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

You can exclude volumes from replication 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 volume set (including the configuration location), 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 Point-in-Time Copy software supports all Sun-supported storage. It works independently of the underlying data reliability software (for example, RAID-1, RAID-5, or volume manager). Additionally, you can use it as a tool when migrating data to and from differing storage types.

Typical uses for the Point-in-Time Copy software include:

- Backup of live application data
- Load data warehouses and fast resynchronization of data warehouses at predefined intervals
- Application development and test on a point-in-time snapshot of live data
- Migrate data across different types of storage platforms and volumes
- Hot back up of application data from frequent point-in-time snapshots

VTOC Information

The Solaris system administrator must be knowledgeable about the virtual table of contents (VTOC) that is created on raw devices by Solaris.

The creation and updating of a physical disk's VTOC is a standard function of Solaris. Software applications like AV Suite, the growth of storage virtualization, and the appearance of SAN-based controllers have made it easy for an uninformed Solaris system administrator to inadvertently allow a VTOC to become altered. Altering the VTOC increases the possibility of data loss.

Remember these points about the VTOC:

- A VTOC is a software generated virtual table of contents based on the geometry of a device and written to the first cylinder of that device by the Solaris `format(1M)` utility.
- Various software components such as `dd(1M)`, backup utilities, Point-in-Time Copy software, and Remote Mirror software can copy the VTOC of one volume to another volume if that volume includes cylinder 0 in its mapping.
- If the VTOC of the source and destination volumes are not 100% identical, then there is a possibility of some type of data loss occurring.

This data loss may not be initially detectable, but can be detected later when other utilities are used, like `fsck(1M)`.

When first configuring and validating volume replication, save copies of all affected device's VTOCs using the `prtvtoc(1M)` utility. The `fmthard(1M)` utility can be used to restore them later, if necessary.

- When using volume managers like VxVM and SVM, copying between individual volumes created under these volume managers is safe. VTOC issues are avoided because the VTOC is excluded from volumes created by these volume managers.
- When formatting individual partitions on a raw device, for all partitions except the backup partition, make sure they do not map cylinder 0, which contains the VTOC. When using raw partitions as volumes, you are the volume manager and you need to exclude the VTOC from partitions that you configure.
- When formatting the backup partition of a raw device, make sure that the physical geometries of the source and destination devices are identical. (Partition 2, by default, maps all cylinders under the backup partition.) If identical device sizing is not possible, make sure that the source backup partition is smaller than the destination partition, and that the destination partition does not map cylinder 0.

Installing and Configuring the Sun StorEdge Availability Suite Software

Note – This guide assumes that you have already installed the volume manager software and the Sun Cluster software on each node in your cluster.



Caution – Do not install the Sun StorEdge Availability Suite 3.2 software on a system running the initial release of the Sun Cluster 3.0 software.

The *Sun StorEdge Availability Suite 3.2 Software Installation Guide* listed in “[Related Documentation](#)” on [page viii](#) describe how to install the Sun StorEdge Availability Suite software in a nonclustered environment. The installation steps to install this software in a Sun Cluster environment are generally the same as described in the installation guides. This chapter describes the differences when you install the software in a Sun Cluster environment.

The topics in this chapter include:

- “[Shutting Down Nodes](#)” on [page 10](#)
- “[Overview of Installation Tasks](#)” on [page 10](#)
- “[Disk Device Groups and the Sun StorEdge Availability Suite Software](#)” on [page 23](#)
- “[Choosing the Configuration Location](#)” on [page 11](#)
- “[Installing the Software](#)” on [page 13](#)
- “[Editing the Bitmap Parameter Files](#)” on [page 16](#)
- “[Shutting Down and Restarting Nodes](#)” on [page 18](#)
- “[Supported Configurations for The Remote Mirror Software](#)” on [page 19](#)
- “[Supported Configurations for the Point-in-time Copy Software](#)” on [page 22](#)
- “[Configuring the Sun Cluster Environment](#)” on [page 23](#)

Shutting Down Nodes

Because the installation process requires you to shut down and restart each node in the cluster, make sure that you install the Sun StorEdge Availability Suite 3.2 software and related patches during your normal maintenance window.

As a result of this shutdown and restart, you might experience a panic condition on the node you are restarting. The node panic is expected behavior in the cluster and is part of the cluster software's *failfast mechanism*. The *Sun Cluster 3.0 Concepts* manual describes this mechanism and the Cluster Membership Monitor (CMM). See [“Shutting Down and Restarting Nodes” on page 18](#).

Overview of Installation Tasks

For each node, use the following order of installation order:

1. Install the volume manager software.
2. Install the Sun Cluster software.
3. Install the Sun StorEdge Availability Suite software as shown in TABLE 2-1.

TABLE 2-1 Installation and Configuration Steps for the Sun StorEdge Availability Suite 3.2 Software

| Installation Steps | For More Information, See... |
|--|---|
| 1. Select a configuration location. | “Choosing the Configuration Location” on page 11 |
| 2. Install the Sun StorEdge Availability Suite core, Remote Mirror, and Point-in-Time Copy software on a cluster node. | Sun StorEdge Availability Suite installation guides listed in “Related Documentation” on page viii . “Supported Software and Hardware” on page 3 . |
| 3. Edit the <code>/usr/kernel/drv/rdc.conf</code> or <code>/usr/kernel/drv/ii.conf</code> files, if necessary. | “Editing the Bitmap Parameter Files” on page 16 |
| 4. Shut down and restart the node. | “Shutting Down and Restarting Nodes” on page 18 |
| 5. Repeat Step 2 through Step 4 for each additional cluster node. | |
| 6. Configure the Sun Cluster software for use with the Sun StorEdge Availability Suite software. | “Supported Configurations for The Remote Mirror Software” on page 19 “Configuring the Sun Cluster Environment” on page 23 |

Choosing the Configuration Location

Place the configuration database on a slice of the cluster quorum device.

Note – Ensure that slice 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. See [“VTOC Information” on page 7](#).

When you install the Sun StorEdge Availability Suite software on the first cluster node, the installation process asks you to specify a raw slice on a `did` device for the single configuration location used by all Sun StorEdge Availability Suite software you plan to install.

The configuration location must be available to all nodes running the Sun StorEdge Availability Suite software.

See TABLE 2-2 for the requirements for this configuration location.

The `scdidadm -L` command shows the local and shared disks by device ID.

TABLE 2-2 Configuration Location Requirements and Considerations

| Item | Requirement or Consideration |
|--------------|--|
| Location | A raw device that is cluster-addressable. For example: <code>/dev/did/rdisk/d0s7</code> . The slice used for the configuration database must reside on the quorum device. |
| Availability | <ul style="list-style-type: none">• The raw device must be accessible by both nodes of the cluster.• The location must be writable by the superuser user.• The location is available or persistent at system startup and reboots.• The slice used for the configuration database cannot be used by any other application (for example a file system or a database). |
| Disk space | The configuration location requires 5.5 Mbytes of disk space. If you specify a file for the configuration location during the installation, the file of the appropriate size is automatically created. Note: If you specify a volume or a slice for the configuration location, only 5.5 Mbytes of the space is used, the remainder is unused. |
| Mirroring | Consider configuring RAID (such as mirrored partitions) for the location and ensure that you mirror the location to another disk in the array. The location cannot be stored on the same disk as the replicated volumes. |

Installing the Software

Install the Remote Mirror software on the primary and secondary host machines. This process also installs the Sun StorEdge Availability Suite core and Point-in-Time Copy software.

Note – Install the software on the primary hosts first.

You can install all Sun StorEdge Availability Suite software or an individual product. Each option also installs the core software, required for all products. The script checks whether the core software is already installed. If it is not, the script installs it.

The `install.sh` installation script on the product CD has the following syntax.

```
install.sh [-j] {-a | -p | -r}
```

where:

-
- | | |
|-----------------|---|
| <code>-j</code> | Installs the packages where the root installation path is a path other than the standard root slice (/). For example, use this option when root is located on a remotely mounted device and you want to install the packages on a remotely mounted device. |
| <code>-a</code> | Installs the core, remote mirror, and point-in-time copy software. Use the following order: <ol style="list-style-type: none">1. The remote mirror software on the primary host machine2. The remote mirror software on the secondary host machine.3. The point-in-time copy software on the primary machine. |
| <code>-p</code> | Installs the core and the point-in-time software. |
| <code>-r</code> | Installs the core and the remote mirror software. Use the following order: <ol style="list-style-type: none">1. The remote mirror software on the primary host machine2. The remote mirror software on the secondary host machine. |
-

▼ To Install the Software.

1. Log in as superuser in single-user mode on the primary host machine.
2. Insert the CD into the CD-ROM drive that is connected to your system.
3. If the Volume Manager daemon `vold(1M)` is not started, use the following command to start it. This allows the CD to automount the `/cdrom` directory.

```
# /etc/init.d/volmgt start
```

Start the Volume Manager daemon only once. Do not start the daemon again.

4. Install the Sun StorEdge core, point-in-time copy, and remote mirror software.

For example, enter the following:

```
# cd /cdrom/cdrom0
# ./install.sh -a
```

You see the following system message:

```
System is ready for Sun StorEdge Availability Suite 3.2 installation.
```

The core software package installation starts and displays the following message:

```
-----ENTER DATABASE CONFIGURATION LOCATION-----
Note: Please ensure this location meets all requirements specified
in the Availability Suite 3.2 Installation Guide.

Enter location:
```

5. Type a raw device for the single configuration location used by all Sun StorEdge software you plan to install.

For example: `/dev/did/rdisk/d0s7`

For configuration location requirements, see [“Choosing the Configuration Location” on page 6](#). For example, `/dev/rdisk/clt1d0s7` or `/config` are typical names. When you enter the location, you see the following message:

```
NOTE: Adding entry to root crontab file. This entry will
automatically back-up the Data Services Configuration Database
daily at 1am to /etc/opt/SUNWesm/dscfg.bak.current
```


NOTE: Effective with the 3.2 version of Availability Suite: Read caching of data volumes is no longer supported, but read caching of bitmap volumes is supported.

When the software installation finishes, the script displays an “installation complete” message.

6. Eject the CD.

```
# cd /  
# eject cdrom
```

7. Perform any-post installation steps for the software as described in “[Editing the Bitmap Parameter Files](#)” on page 16 and the Sun StorEdge Availability Suite installation guides listed in “[Related Documentation](#)” on page viii.

Note – Ensure that you place the names and IP addresses of all machines you plan to use with the Remote Mirror software in the `/etc/hosts` file. Make sure you include the logical host names and IP addresses of the logical hosts you plan to use with the Remote Mirror software in the `/etc/hosts` file.

Edit this file on each machine where you are installing and running the Remote Mirror software.

8. Shut down and restart this node.

See “[Shutting Down Nodes](#)” on page 10 and “[Shutting Down and Restarting Nodes](#)” on page 18.

9. Log on as the root user at the next cluster node where you are installing the software and repeat these installation steps.

10. See “[Mounting and Replicating Global Volume File Systems](#)” on page 30 for information about global file systems.

Editing the Bitmap Parameter Files

Bitmap volumes are used by the Remote Mirror and Point-in-Time Copy software to track differences between volumes and provide information for volume updates. The Sun StorEdge software documentation listed in [“Related Documentation” on page viii](#) describes the bitmap size and other requirements.

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

The Remote Mirror and Point-in-Time Copy software include two configuration files that determine how bitmap volumes are written to and saved:

- remote mirror - /usr/kernel/drv/rdc.conf
- point-in-time copy - /usr/kernel/drv/ii.conf



Caution – The Sun StorEdge Availability Suite 3.2 Remote Mirror and Point-in-Time Copy software do not support bitmap files. The software uses regular raw devices to store bitmaps. These raw devices must be located on a disk separate from the disk that contains your data.

Setting the Bitmap Operation Mode

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`. The default setting is 0. Set the bitmap mode to 1, as in the following example:

- **Edit the `rdc.conf` file and locate the following section. Edit the value for the bitmap mode, save the file, and close it.:**

```
# 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=1;
```

The /usr/kernel/drv/ii.conf File

The /usr/kernel/drv/ii.conf file contains one setting that sets the point-in-time copy bitmap save mode:

- `ii_bitmap` - modify to change how the bitmap volume is saved during a shut down or system crash. In a Sun Cluster environment, set this to 1.

A bitmap maintained on disk can persist across a system crash when this field is set to 1.

▼ To Edit the ii.conf File

1. Open the /usr/kernel/drv/ii.conf file using a text editor such as `vi(1)`.
2. In a Sun Cluster environment, set the bitmap mode to 1.

For example:

```
# bitmap volume storage strategy:
# 0 indicates kernel memory loaded from bitmap volume when shadow is resumed
#   and saved to bitmap volume when shadow is suspended.
# 1 indicates permanent SDBC storage, bitmap volume is updated directly as
#   bits are changed.
# 2 indicates that if FWC is present strategy 1 is used, otherwise strategy 0.
ii_bitmap=1;
```

3. Save and exit the file.
4. Shut down and restart your server as described in [“Shutting Down and Restarting Nodes” on page 18](#).

Shutting Down and Restarting Nodes



Caution – After a shutdown and restart, you might experience a panic condition on the node you are restarting. The node panic is expected behavior in the cluster and is part of the cluster software’s *failfast mechanism*. The *Sun Cluster 3.0 Concepts* manual describes this mechanism and the Cluster Membership Monitor (CMM).

After performing the steps listed in [“Overview of Installation Tasks” on page 10](#), shutdown and restart each node.

Note – The `shutdown(1M)` command shuts down a single node or machine; the `scshutdown(1M)` command shuts down all nodes in a cluster. To shut down a single node, use the `scswitch(1M)` command as describes in the Sun Cluster documentation.

▼ To Shut Down and Restart a Node

- Shut down and restart your node as follows:

```
# scswitch -S -h nodelist
# /etc/shutdown -y -g0 -i 6
```

| | |
|-----------------------------|--|
| <code>-S</code> | Evacuates all device and resource groups from the node. |
| <code>-h node1,node2</code> | Specifies the cluster nodes that can master this resource group. If you do not specify these nodes, it defaults to all the nodes in the cluster. |

Supported Configurations for The Remote Mirror Software

Adding Host Names

This step ensures that the host names in the `/etc/hosts` file are read and known by machines running the version 3.2 software. Place the names and IP addresses of all machines you plan to use with the Remote Mirror software in the `/etc/hosts` file. Make sure you include the logical host names and IP addresses of the logical hosts you plan to use with the Remote Mirror software in the `/etc/hosts` file. Edit this file on each machine where you are installing and running the Remote Mirror software.

▼ Edit the `/etc/hosts` File

- **Add the names and IP addresses of all machines you plan to use with the remote mirror software to the `/etc/hosts` file.**

Edit this file on each machine where you are installing and running the remote mirror software.

Using Autosynchronization

Consider the following when using autosynchronization with Sun Cluster:

- If you want automatic resynchronization to occur in the event of a cluster failover, turn on the autosync feature. With this feature enabled, any cluster failover will automatically put the remote mirror volume sets back into replication mode after an update occurs.
- If you want to manually force clusters to failover, you must be careful to allow all remote mirror components to fully recognize the condition, including the remote host that is not in the cluster. In practice, this means that you should not attempt an immediate update sync after the failover. You should wait at least thirty seconds after the completion of the `scswitch` command and before starting an update sync to allow time for `suncluster` to complete its logical host interface reconfiguration.

Rules For the Remote Mirror Software

- **The primary volume and its bitmap volume and possible disk queue volume or the secondary volume and its bitmap volume in a remote mirror volume set must reside in the same disk device group per node.** (A remote mirror volume set also includes information about primary and secondary hosts and operating mode.)

For example, you cannot have a primary volume with a disk device group name of `sndrdg` and a primary bitmap volume with a disk device group name of `sndrdg2` in the same remote mirror volume set.

- **With the Remote Mirror software, you can use more than one disk device group for cluster switchover and failover,** but each primary or secondary disk device component in the cluster node's volume set must reside in the same disk device group.
- **The Remote Mirror software also requires a resource group containing the disk device group and logical failover host.** The disk device group is used to create a lightweight resource group containing the disk and a logical failover host.

The Remote Mirror software requires that the `SUNW.HAStorage` or `SUNW.HAStoragePlus` resource is configured in the same resource group as the logical host, as described in the procedures in [“Configuring Sun Cluster for HAStorage or HAStoragePlus” on page 24](#).

The resource group name you specify consists of the disk device group name appended with `-stor-rg`. For example, if the group name is `sndrdg`, then the resource group name would be:

```
sndrdg-stor-rg
```

- **Remote mirror replication within the cluster is not supported.** An example is when the primary host is cluster node 1 and the secondary host is cluster node 2 in the cluster and the primary, secondary, and bitmap volumes in a volume set reside in the same disk device group.

Typically, the remote mirror primary host is part of one cluster configuration, while the replicating secondary host might or might not be part of a different cluster.

Three configurations for the Remote Mirror software are supported:

- [“Remote Mirror Primary Host Is On a Cluster Node” on page 21](#)
- [“Remote Mirror Secondary Host On a Cluster Node” on page 21](#)
- [“Remote Mirror Primary and Secondary Hosts On a Cluster Node” on page 22](#)

Remote Mirror Primary Host Is On a Cluster Node

In this configuration, the remote mirror primary host is the logical host you created in the remote mirror resource group for the remote mirror disk group using the `scrgadm` command; for example, see [“Configuring Sun Cluster for HAStorage or HAStoragePlus” on page 24](#).

If you have configured the remote mirror autosynchronization feature on the primary host, the Remote Mirror software starts an update resynchronization from the primary host for all affected remote mirror volume sets following a switchover or failover event (if the autosynchronization feature is enabled for those volume sets). This operation is performed after the resource group and network switchover operation is complete. See the `sndradm` man page and the *Sun StorEdge Availability Suite 3.2 Remote Mirror Software Administration and Operations Guide* for a description of the `sndradm -a` command to set the autosynchronization feature.

Remote Mirror Secondary Host On a Cluster Node

In this configuration, the remote mirror secondary host is the logical host you created in the remote mirror resource group for the remote mirror disk group using the `scrgadm` command; for example, see [“Configuring Sun Cluster for HAStorage or HAStoragePlus” on page 24](#).

Operations such as update resynchronizations occur and are issued from the primary host machine. Following a switchover (or failover) event, the Remote Mirror software attempts to start an update resynchronization for all affected remote mirror volume sets (if the autosynchronization feature is enabled for those volume sets). However, the remote mirror secondary host in a remote mirror volume set cannot initiate an update resynchronization.

This operation is performed after the resource group and network switchover operation is complete. In this case, the remote mirror secondary host switchover appears to be a short network outage to the remote mirror primary host.

If you have configured the remote mirror autosynchronization feature on the primary host, the `sndrsyncd` synchronization daemon attempts to resynchronize the volume sets if the system reboots or link failures occur. See the `sndradm` man page and the *Sun StorEdge Availability Suite 3.2 Remote Mirror Software Administration and Operations Guide* for a description of the `sndradm -a` command to set the autosynchronization feature.

If this feature is disabled (its default setting) and volume sets are logging but not replicating, perform the updates manually using the `sndradm` command.

Remote Mirror Primary and Secondary Hosts On a Cluster Node

Remote mirror replication within the cluster is not supported; that is, when the primary and secondary hosts reside in the same cluster and the primary, secondary, and bitmap volumes in a volume set reside in the same disk device group.

However, if the remote mirror primary and secondary hosts are configured in different clusters, see [“Remote Mirror Primary Host Is On a Cluster Node” on page 21](#) and [“Remote Mirror Secondary Host On a Cluster Node” on page 21](#) for operating considerations.

Supported Configurations for the Point-in-time Copy Software

Rules For the Point-in-Time Copy Software

- **All Point-in-Time Copy volume set components must reside in the same disk device group.** (A Point-in-Time Copy volume set includes the master, shadow, bitmap, and optional overflow volumes.)
- **With the Point-in-Time Copy software, you can use more than one disk device group for cluster switchover and failover,** but each component in the volume set must reside in the same disk device group.

For example, you cannot have a master volume with a disk device group name of `ii-group` and a shadow volume with a disk device group name of `ii-group2` in the same volume set.

- **If a Solaris operating environment failure or Sun Cluster failover occurs during a point-in-time copy or update operation to the master volume,** specifically where the shadow volume is copying (`iiadm -c m`) or updating (`iiadm -u m`) data to the master volume, the master volume might be in an inconsistent state (that is, the copy or update operation might be incomplete). [“Preserving Point-in-time Copy Volume Data” on page 39](#) describes how to avoid this situation.

Disk Device Groups and the Sun StorEdge Availability Suite Software

The Solstice Disk Suite (SDS) and VERITAS Volume Manager (VxVM) can arrange disk devices into a group to be mastered by a cluster node. You can then configure these disk device groups to fail over to another cluster node, as described in [“Configuring the Sun Cluster Environment” on page 23](#).

The SDS and VxVM device paths contain the disk device group. When operating in a Sun Cluster environment, the Sun StorEdge Availability Suite commands `sndradm` and `iiadm` automatically detect and use the disk device group as configured in [“Configuring the Sun Cluster Environment” on page 23](#).

You can also use the `sndradm` and `iiadm` commands to select specified disk device groups or to operate on a volume set as a local node-only configuration entry. See [“Using the Sun StorEdge Availability Suite `iiadm` and `sndradm` Commands” on page 29](#)

Configuring the Sun Cluster Environment

Note – The Sun StorEdge Availability Suite software is supported *only* in a two-node Sun Cluster 3.0 Update 3 or Sun Cluster 3.1 initial release environment.

The procedures in this section describe how to configure the Sun Cluster software for use with the Remote Mirror and Point-in-Time Copy software. The *Sun Cluster 3.0 Data Installation and Configuration Guide* contains more information about configuring and administering Sun Cluster data services. See the `scrgadm(1M)` and `scswitch(1M)` man pages for more information.

The general configuration steps are:

TABLE 2-3

-
1. Log on to any node in the cluster.
 2. Configure a disk device group using your volume manager.
 3. Register the `SUNW.HAStorage` or `SUNW.HAStoragePlus` resource type.
-

TABLE 2-3

-
4. Create a resource group.
 5. Add `SUNW.HAStorage` or `SUNW.HAStoragePlus` to the disk device group.
 6. (Remote Mirror step only) Add a logical failover host to the resource group.
 7. Enable the resource group and bring it online.
-

See [“Configuring Sun Cluster for HAStorage or HAStoragePlus” on page 24](#).

When you complete the selected procedure, the resource group is configured and ready to use.

▼ Configuring Sun Cluster for HAStorage or HAStoragePlus



Caution – You must adhere to the naming conventions and configuration rules specified in this procedure. If you do not, the resulting configuration is unsupported and might lead to cluster hangs and panics. The naming convention for device groups is to use the suffix `-stor-rg`.

1. **Log on as the root user on any node in the cluster.**
2. **Configure a disk device group using your volume manager software.**
See the documentation that came with your volume manager software. Also you might check the currently configured groups before configuring a new disk device group. For example, use the `metaset(1M)`, `vxdg`, or `vxprint` commands, depending on your volume manager software.
3. **Register `SUNW.HAStorage` or `SUNW.HAStoragePlus` as a resource type.**

```
# scrgadm -a -t SUNW.HAStorage
```

```
# scrgadm -a -t SUNW.HAStoragePlus
```

4. **Create a resource group for the *devicegroup*.**

```
# scrgadm -a -g devicegroup-stor-rg -h node1,node2
```

devicegroup is the required disk device group name.

-h node1,node2 specifies the cluster nodes that can master this resource group. If you do not specify these nodes, it defaults to all the nodes in the cluster.



Caution – Do not add resources other than HAStorage or HAStoragePlus, and a logical host, to this lightweight resource group. Failure to follow this rule might cause the Sun StorEdge Availability Suite software to not fail over or switch over properly.

5. For a SUNW.HAStorage resource, use the following command to add the resource to the resource group:

```
# scrgadm -a -j devicegroup-stor -g devicegroup-stor-rg \  
-t SUNW.HAStorage \  
-x ServicePaths=devicegroup -x AffinityOn=True
```

devicegroup Disk device group name.

-x ServicePaths= Specifies the extension property that the Sun StorEdge Availability Suite software relies on. In this case, use the disk device *devicegroup*.

-x AffinityOn=True Specifies that the SUNW.HAStorage resource needs to perform an affinity switchover for the global devices and cluster file systems defined in *-x ServicePaths*. It also enforces co-location of resource groups and disk device groups on the same node, thus enhancing the performance of disk-intensive data services. If the device group is switched to another node while the SUNW.HAStorage resource is online, *AffinityOn* has no effect and the resource group does not migrate along with the device group. On the other hand, if the resource group is switched to another node, *AffinityOn* being set to *True* causes the device group to follow the resource group to the new node.

For a a SUNW.HAStoragePlus resource, use the following command to add the resource to the resource group.

```
# scrgadm -a -j devicegroup-stor -g devicegroup-stor-rg \  
-t SUNW.HAStoragePlus \  
-x GlobalDevicePaths=devicegroup -x AffinityOn=True
```

- x GlobalDevicePaths= specifies the extension property that the Sun StorEdge Availability Suite software relies on. In this case, use the disk device *devicegroup*.
- x AffinityOn=True specifies that the SUNW.HAStoragePlus resource needs to perform an affinity switchover for the global devices and cluster file systems defined in -x GlobalDevicePaths. It also enforces co-location of resource groups and disk device groups on the same node, thus enhancing the performance of disk-intensive data services. If the device group is switched to another node while the SUNW.HAStoragePlus resource is online, AffinityOn has no effect and the resource group does not migrate along with the device group. On the other hand, if the resource group is switched to another node, AffinityOn being set to True causes the device group to follow the resource group to the new node.

6. Add a logical hostname resource to the resource group.

Note – Perform this step for the remote mirror volumes only. This step is not needed for point-in-time copy volumes.

```
# scrgadm -a -L [-j lhost-stor] -g devicegroup-stor-rg \  
-l lhost1,lhost2,...lhostN  
-n nafo0@node,nafo0@node
```

- j *lhost-stor* Optional resource *lhost-stor*. If you do not specify this option and resource, the name defaults to the first logical hostname specified in the -l option.
- l *lhost1,lhost2,...lhostN* Specifies a comma-separated list of UNIX hostnames (logical hostnames) by which clients communicate with the Sun StorEdge Availability Suite software in the resource group.
- n *nafo0@node,nafo0@node* Specifies the comma-separated list of Network Adapter Failover (NAFO) groups on each node.
- node* can be a node name or ID. You can display the node ID using the `scconf -p` command.

7. Enable the resources in the resource group, manage the resource group, and bring the resource group online.

```
# scswitch -Z -g devicegroup-stor-rg
```

8. Verify that the resource is online.
 - a. Run the following command on any cluster node.

```
# scstat -g
```

- b. Look for the resource group state field to determine if the resource group is online on the nodes specified in the node list.
9. For the HAStoragePlus resource, verify that the resource group can be failed between nodes.

```
# scswitch -z -g <dg>-stor-rg -h <fail-to node> (fails resource group to specified node)
```

Or:

```
# scswitch -s -h <fail-from node> (fails ALL resources from specified node)
```

Configuring the HAStoragePlus Resource Types with Volume Sets

This example shows how to configure a resource group on a locally-mounted Sun Cluster global device partition.

You can configure the HAStoragePlus resource to fail over resource groups as well as individual volume sets to another node in the cluster. When configuring a resource type with volume sets, consider the following:

- When you add a new volume set to the Sun StorEdge Availability Suite software, you must disable the configured resource group and place it offline.
- You must specify each volume in the set. For example, the following command shows how to define a volume set to an existing resource group using the HAStoragePlus resource:

```
# scrgadm -a -j iidg-rs -g iidg -t SUNW.HAStoragePlus \  
-x GlobalDevicePaths=/dev/vx/rdisk/iidg/ii01,/dev/vx/rdisk/ii02, \  
/dev/vx/rdisk/iidg/ii11,/dev/vx/rdisk/iidg/ii12,/dev/vx/rdisk/iidg/iibitmap1, \  
/dev/vx/rdisk/iidg/iibitmap2
```

where:

| | |
|-----------------------|---|
| -j iidg-rs | is the resource name. |
| -g iidg | is the resource group name. |
| -x GlobalDevicePaths= | specifies the extension property <code>GlobalDevicePath</code> and raw device volume names for the point-in-time copy volume set. |

Using the Sun StorEdge Availability Suite `iiadm` and `sndradm` Commands

This chapter describes using the Sun StorEdge Availability Suite commands `iiadm` and `sndradm` in a Sun Cluster environment. The Sun StorEdge Availability Suite administrator guides listed in [“Related Documentation” on page viii](#) describe the full command syntax and options for `iiadm` and `sndradm`.

The Sun StorEdge Availability Suite software can use volumes that are global or local devices.

- *Global devices* are Sun StorEdge Availability Suite or other volumes accessible from any cluster node and which fail over and switch back under the control of the Sun Cluster framework.
- *Local devices* are Sun StorEdge Availability Suite software volumes that are local to the individual node (host machine), not defined in a disk or resource group, and not managed within a cluster file system. Local devices do not fail over and switch back.

The topics in this chapter include:

- [“Mounting and Replicating Global Volume File Systems” on page 30](#)
- [“Global Device Command Syntax” on page 31](#)
- [“Local Device Command Syntax” on page 33](#)
- [“Putting All Cluster Volume Sets in an I/O Group” on page 37](#)
- [“Preserving Point-in-time Copy Volume Data” on page 39](#)

Mounting and Replicating Global Volume File Systems

If a volume contains a file system and you wish to replicate the file system using the Sun StorEdge Availability Suite software, you must create and mount a related global file system on all cluster nodes. These steps ensure that file system is available to all nodes and hosts when you copy or update the volume sets.

Note – See the Sun Cluster documentation for information about administering cluster file systems, including creating and mounting global file systems. See also the `mount(1M)` and `mount_ufs(1M)` commands.

For example:

1. **Create the file systems on the appropriate diskset metadevices or disk group volumes.**

```
# newfs raw-disk-device
```

For example, using the VERITAS Volume Manager, you might specify *raw-disk-device* as `/dev/vx/rdisk/sndrdg/vol01`.

2. **On each node, create a mount point directory for the file system.**

```
# mkdir -p /global/device-group/mount-point
```

- *device-group* is the name of the directory that corresponds to the name of the device group that contains the device.
- *mount-point* is the name of the directory on which to mount the file system.

3. **On each node, add an entry to the `/etc/vfstab` file for the mount point and use the `global` mount option.**
4. **On a cluster node, use `sccheck(1M)` to verify the mount points and other entries.**
5. **From any node in the cluster, mount the file system.**

```
# mount /global/device-group/mount-point
```

6. **Verify that the file system is mounted using the `mount` command with no options.**

Global Device Command Syntax

Note – During the initial enable of the remote mirror or point-in-time copy volume sets, you can optionally specify the global device disk group with the `-C tag` cluster option when you use the `iiadm` or `sndradm` commands. As this section shows, however, you do not have to use the `-C tag` cluster option. Also see [“The C tag and -C tag Options” on page 31](#).

The Sun StorEdge Availability Suite software automatically derives the disk device group name from the volume path when you first enable volume sets. During this initial enable operation, the Remote Mirror and Point-in-Time Copy software creates a configuration entry for each volume set. Part of the entry is the disk device group name for use in a cluster.

The Remote Mirror software shows this name as `C tag`, where `tag` is the disk device group name. point-in-time copy shows this name as `Cluster tag: tag`.

The C tag and -C tag Options

`C tag` is displayed as part of a volume set’s configuration information as shown in [“Global Device Command Syntax” on page 31](#).

Typically, the Sun StorEdge Availability Suite software derives the disk device group name from the volume path and does not require the `-C tag` option.

Use the `-C tag` option and `C tag` volume set option to execute the `iiadm` and `sndradm` commands on the enabled volume sets in the disk device group name `tag`, when the disk device group name is not indicated by the volume path. The commands are not executed on any other volume sets in your configuration; `-C tag` excludes those volume sets not contained in the `tag` disk device group from the specified operation.

For example, the following command makes a point-in-time copy volume set in the `iigrp2` disk device group wait for all copy or update operations to finish before you can issue other point-in-time copy commands.

```
# iiadm -w /dev/vx/rdisk/iigrp2/nfsvol-shadow -C iigrp2
```

Remote Mirror Example

When you enable this remote mirror volume set where `host1` is a logical failover host name:

```
# sndradm -e host1 /dev/vx/rdsk/sndrdg/datavol /dev/vx/rdsk/sndrdg/datavolbml \  
host2 /dev/rdsk/clt3d0s0 /dev/rdsk/clt2d0s4 ip sync
```

the corresponding configuration information as shown by the `sndradm -i` command is:

```
# sndradm -i  
  
host1 /dev/vx/rdsk/sndrdg/datavol /dev/vx/rdsk/sndrdg/datavolbml \  
host2 /dev/rdsk/clt3d0s0 /dev/rdsk/clt2d0s4 ip sync \  
C sndrdg
```

The `C` portion of the entry shows a disk device group name `sndrdg`.

Point-in-Time Copy Example

When you enable a point-in-time copy volume set on a cluster node (logical failover host):

```
# iiadm -e ind /dev/vx/rdsk/iidg/clt3d0s0 /dev/vx/rdsk/iidg/clt3d0s4 \  
/dev/vx/rdsk/iidg/clt2d0s5
```

the corresponding configuration as shown by `iiadm -i` command is:

```
# iiadm -i  
  
/dev/vx/rdsk/iidg/clt3s0d0: (master volume)  
/dev/vx/rdsk/iidg/clt3d0s4: (shadow volume)  
/dev/vx/rdsk/iidg/clt2d0s5: (bitmap volume)  
Cluster tag: iidg  
Independent copy  
Volume size: 208278  
Percent of bitmap set: 0
```

The `Cluster tag` entry shows the derived disk device group name `iidg`.

Local Device Command Syntax

Note – Enabling a local disk device group named `local` prevents you from configuring a cluster disk device group named `local`.

- When you enable a point-in-time copy volume set, use the `-C local` option to specify that the volume set's disk device group name is `local`:

```
iiadm -C local -e {dep | ind} master shadow bitmap
```
- When you enable a remote mirror volume set, use the `C local` option as part of the `vol-set` volume set definition:

```
sndradm -e vol-set
```

where `vol-set` is:

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

The `local` disk device group is local to the individual cluster node and is not defined in a cluster disk or resource group. Local devices do not fail over and switch back. This initial configuration is similar to using the Sun StorEdge Availability Suite software in a nonclustered environment.

When you enable a volume set with the `local` disk device group, its configuration entry includes the name of its host machine.



Caution – Volumes and bitmaps used in a local remote mirror volume set cannot reside in a shared disk device group or metaset.

Point-in-Time Copy Example

When you enable this point-in-time copy volume set where `local` indicates a disk device group:

```
# iiadm -C local -e ind /dev/rdisk/c1t90d0s5 /dev/rdisk/c1t90d0s6 \  
/dev/rdisk/c1t90d0s7
```

the corresponding configuration as shown by `iiadm -i` command is:

```
# iiadm -i

/dev/rdsk/iidg/clt90d0s5: (master volume)
/dev/rdsk/iidg/clt90d0s6: (shadow volume)
/dev/rdsk/iidg/clt90d0s7: (bitmap volume)
Cluster tag: localhost (local)
Independent copy
Volume size: 208278
Percent of bitmap set: 0
```

where *localhost* is the local host name as returned by the `hostname(1)` command.

The corresponding configuration information as shown by the `dscfg -l` command is:

```
# dscfg -l | grep /dev/rdsk/clt3d0s0

ii: /dev/rdsk/clt90d0s5 /dev/rdsk/clt90d0s6 /dev/rdsk/clt90d0s7 I - l.localhost - -
```

Which Host Do I Issue Commands From?

The Sun StorEdge Availability Suite software requires that you issue the `iiadm` or `sndradm` commands from the node that is the current primary host for the disk device group that the command applies to.

In a clustered environment, you can issue the command from the node mastering the disk device group you specified in [Step 2](#) in “[Configuring Sun Cluster for HAStorage or HAStoragePlus](#)” on page 24.

When you enable the Remote Mirror software for the first time, issue the `sndradm enable` command from the primary and secondary hosts. See TABLE 3-1.

TABLE 3-1 Which Host to Issue Remote Mirror 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 Remote Mirror 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 Remote Mirror software. | Primary and secondary host | Perform this operation on both hosts if you are deleting a volume set. When enabling the Remote Mirror software for the first time, issue the command from both hosts. |
| Full forward or reverse synchronization (copy). | Primary host | Ensure that both hosts are enabled. |
| Forward or reverse synchronization (update). | Primary host | Ensure that both hosts are enabled. |

TABLE 3-1 Which Host to Issue Remote Mirror Commands From *(Continued)*

| Task | Where Command Is Issued | Comments |
|---------------------------------------|--------------------------------|---|
| Log. | Primary host | Perform on the primary host only if a synchronization is in progress. |
| | | Perform on the secondary host if the primary host failed. |
| | Primary or secondary host | Perform on either host if no synchronization is in progress. |
| Toggle the autosynchronization state. | Primary host | |
| Update an I/O group. | Primary and secondary hosts | |

Putting All Cluster Volume Sets in an I/O Group

Note – Placing volume sets in an I/O group does not affect the cluster operations of all volume sets configured in disk device and resource groups.



Caution – 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.

The Remote Mirror and Point-in-Time Copy software enables you to assign volume sets to I/O groups. Instead of issuing one command for each volume set, you can:

- Assign specific volume sets to an I/O group
- Issue one command specifying the I/O group
- Perform operations on those volume sets only

Like the `-c tag` and `c tag` options, the I/O group name excludes all other enabled volume sets from operations you specify.

In a clustered environment, you can assign some or all volume sets in a specific disk device group to an I/O group when you enable each volume set.

Example

1. Enable three point-in-time copy volume sets and place them in an I/O group named `cluster1`.

```
# iiadm -g cluster1 -e ind /dev/rdisk/iigrp2/clt3d0s0 \  
/dev/rdisk/iigrp2/clt3d0s4 /dev/rdisk/iigrp2/clt2d0s5  
  
# iiadm -g cluster1 -e dep /dev/rdisk/iigrp2/clt4d0s0 \  
/dev/rdisk/iigrp2/clt4d0s4 /dev/rdisk/iigrp2/clt3d0s5  
  
# iiadm -g cluster1 -e ind /dev/rdisk/iigrp2/clt5d0s0 \  
/dev/rdisk/iigrp2/clt5d0s4 /dev/rdisk/iigrp2/clt4d0s5
```

2. Wait for any disk write operations to complete before issuing another command.

```
# iiadm -g cluster1 -w
```

3. Allow your applications to write to the master volumes.
4. Update the shadow volumes.

```
# iiadm -g cluster1 -u s
```

Preserving Point-in-time Copy Volume Data

If a Solaris operating environment system failure or Sun Cluster failover occurs during a point-in-time copy or update operation to the master volume, specifically where the shadow volume is copying (`iiadm -c m`) or updating (`iiadm -u m`) data to the master volume, the master volume might be in an inconsistent state (that is, the copy or update operation might be incomplete).

To avoid or reduce the risk of inconsistent data if a system failover occurs during such a copy or update operation, perform the following before performing the shadow volume-to-master volume copy or update operation:

1. Create a second independent shadow volume copy of the master volume by issuing an `iiadm -e ind` command.

This operation results in a full shadow volume copy of the master volume data.

2. Ensure that all copy or update operations to this second shadow volume are finished by issuing a wait command (`iiadm -w shadowvol`) after issuing the `iiadm -e ind` command.

You can now perform the copy or update operation from the original shadow volume to the master volume. If a system failure or failover occurs during this operation, you at least have a known good copy of your original master volume data. When this operation is complete, you can keep the second shadow volume under point-in-time copy control or return it to your storage pool.

Index

SYMBOLS

`/etc/hosts` 19
`/usr/kernel/drv/rdc.conf` 16

A

a 2

B

bitmap volume
 parameter files 16
bitmap volumes
 ii.conf file 17

C

-C local 33
commands
 iiadm and sndradm 29
configuration
 location 11
 procedures 23
 supported for point-in-time copy software 22
 supported for remote mirror software 19
configuration location
 disk space required for 12
 requirements 12
 specifying 14
configuring the Sun StorEdge software 9, 23

D

data service
 defined 2
disk device groups 23

F

files
 `/etc/hosts` 19
 `/usr/kernel/drv/rdc.conf` 16

G

global devices 5
 command syntax 31
grouping volume sets 37

H

High Availability applications
 see data services 2

I

I/O group
 grouping volume sets in a cluster 37
ii.conf file 17
installation
 remote mirror software 14

L

local devices 5
command syntax 33

P

point-in-time copy
configuration rules 22
ii.conf file 17
point-in-time copy software
data in a system failover 39
I/O groups 37
iiadm command 4, 29
supported configurations 22

R

reboot
shutdown and restart node 18
Remote mirror software
supported configurations 19
remote mirror software
configuration rules 20
I/O groups 37
sndradm command 4, 29
requirements
bitmap volumes 16

S

shutdown and restart node 18
software
installing 14
Sun StorEdge
installation 13
Sun StorEdge software
bitmap volumes 16
configuration
location 11
configuring 9
iiadm command 5
installing 9
reboot node 18
sndradm command 5
SUNWnvm 4

supported configurations 20, 22

syntax

global device command 29
iiadm and sndradm commands 29
local device command 33

V

volume sets
I/O group 37