



# Sun™ Cluster 3.0 U1 and Sun StorEdge™ Software 3.0 Integration Guide

---

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

Part No. 816-1544-10  
September 2001, [Revision A](#)

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

Copyright 2001 Sun Microsystems, Inc., 901 San Antonio Road, Palo Alto, CA 94303-4900 U.S.A. All rights reserved.

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

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

Federal Acquisitions: Commercial Software—Government Users Subject to Standard License Terms and Conditions.

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 94303-4900 Etats-Unis. Tous droits réservés.

Ce produit ou document est 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 Management Center, Sun StorEdge, AnswerBook2, docs.sun.com, 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.

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



# Contents

---

<b>Preface</b>	<b>vii</b>
<b>1. Overview</b>	<b>1</b>
Terminology Used in This Guide	2
Supported Software and Hardware	3
Using the Sun StorEdge Services 3.0 Software in a Sun Cluster Environment	5
Global and Local Use of the Sun StorEdge Services Software	6
Volumes Eligible for Use with the Sun SNDR and Instant Image Software	7
Sun StorEdge Fast Write Cache	8
Sun StorEdge Services Software Version Compatibility	8
<b>2. Installing and Configuring The Sun StorEdge Services Software</b>	<b>9</b>
Shutting Down Nodes After Installing the Sun StorEdge Services Software	10
Task Summary: Steps to Configure the Sun Cluster 3.0 U1 and Install the Sun StorEdge Services Software	10
Disk Device Groups and the Sun StorEdge Services Software	12
Specifying the Configuration Location During the Sun StorEdge Installation Process	13
Requirements for the Configuration Location	13
▼ To Specify the Configuration Location	14
Editing the Sun StorEdge Bitmap Parameter Files	16

	The <code>/usr/kernel/drv/rdc.conf</code> File	17
▼	To Edit the <code>rdc.conf</code> File	17
	The <code>/usr/kernel/drv/ii.conf</code> File	18
▼	To Edit the <code>ii.conf</code> File	18
	Shutting Down and Restarting the Nodes	19
▼	To Shut Down and Restart a Node	19
	Supported Configurations for the Sun SNDR Software	20
	The Sun SNDR Primary Host is Configured On a Cluster Node	20
	The Sun SNDR Secondary Host is Configured On a Cluster Node	21
	The Sun SNDR Primary and Secondary Hosts are Configured On a Cluster Node	21
	Supported Configurations for the Instant Image Software	22
	Configuration Rules for the Instant Image and Sun SNDR Software	22
	Configuring the Sun Cluster 3.0 U1 Environment for Use with the Sun StorEdge Services Software	24
▼	To Configure the Sun Cluster Environment	25
<b>3.</b>	<b>Using The Sun StorEdge Services <code>iiadm</code> and <code>sndradm</code> Commands</b>	<b>29</b>
	Configuring Sun SNDR and Instant Image Volume Sets	30
	Mounting and Replicating Global Volume File Systems	31
	The <code>iiadm</code> and <code>sndradm</code> Command Syntax	32
	Global Device Command Syntax	32
	Local Device Command Syntax	34
	The <code>C tag</code> and <code>-C tag</code> Options	36
	Which Host Do I Issue Commands From?	37
	Putting All Cluster Volume Sets in an I/O Group	39
	Preserving Instant Image Volume Data During a System Failure or Failover	40
	A Directory Named <code>._</code> is Created After Mounting a Secondary Volume	41

# Tables

---

TABLE 1-1	Supported Software, Patches, and Hardware	3
TABLE 1-2	Cluster Terminology and Status	5
TABLE 2-1	Installation and Configuration Steps	11
TABLE 3-1	Which Host to Issue Sun SDR Commands From	37



# Preface

---

The *Sun Cluster 3.0 U1 and Sun StorEdge Software 3.0 Integration Guide* describes the Sun StorEdge™ Network Data Replicator (Sun Sندر) and Instant Image software products and integrating them in a Sun™ Cluster 3.0 Update 1 environment.

---

**Note** – Sun Cluster 3.0 Update 1 is also known as the Sun Cluster 3.0 07/01 release.

---

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 software as described in the installation and release documentation in [“Related Documentation” on page x](#), see [Chapter 2](#).

---

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

- *Sun Cluster 3.0 U1 Concepts*
- *Sun Cluster 3.0 U1 Installation Guide*
- *Sun Cluster 3.0 U1 System Administration Guide*
- *Sun Cluster 3.0 U1 Release Notes*
- *Sun Cluster 3.0 U1 Release Notes Supplement*
- *Sun StorEdge Network Data Replicator 3.0 Installation Guide*
- *Sun StorEdge Network Data Replicator 3.0 Release Notes*
- *Sun StorEdge Network Data Replicator 3.0 System Administrator’s Guide*
- *Sun StorEdge Instant Image 3.0 Installation Guide*

- *Sun StorEdge Instant Image 3.0 Release Notes*
  - *Sun StorEdge Instant Image 3.0 System Administrator's Guide*
- 

## How This Book Is Organized

**Chapter 1** is an overview of the Sun Cluster and Sun StorEdge software integration.

**Chapter 2** describes installing and configuring the Sun StorEdge software for use in a Sun Cluster environment.

**Chapter 3** describes using the Sun StorEdge software commands in a Sun Cluster environment.

---

## Using UNIX Commands

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

See the software documentation that you received with your system.

---

# Typographic Conventions

Typeface	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.
<b>AaBbCc123</b>	What you type, when contrasted with on-screen computer output	% <b>su</b> Password:
<i>AaBbCc123</i>	Book titles, new words or terms, words to be emphasized	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.
	Command-line variable; replace with a real name or value	To delete a file, type <code>rm filename</code> .
[ ]	In syntax, brackets indicate that an argument is optional.	<code>scmadm [-d sec] [-r n[:n][,n]...] [-z]</code>
{ arg / arg }	In syntax, braces and pipes indicate that one of the arguments must be specified.	<code>sndradm -R b {p s}</code>
\	At the end of a command line, the \ (backslash) indicates that the command continues on the next line.	<code>scrgadm -a -L \ -g <i>groupname</i>-stor-rg \ -l <i>host1</i>, <i>host2</i> \ -n <i>nafo0@node</i>, <i>nafo0@node</i></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
Hardware	<i>Sun Cluster 3.0 U1 Hardware Guide</i>	806-7070
Software Installation	<i>Sun Cluster 3.0 U1 Installation Guide</i>	806-7069
	<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
Data Services	<i>Sun Cluster 3.0 U1 Data Services Installation and Configuration Guide</i>	806-7071
Concepts	<i>Sun Cluster 3.0 U1 Concepts</i>	806-7074
Error Messages	<i>Sun Cluster 3.0 U1 Error Messages Manual</i>	806-7076
Release Notes	<i>Sun Cluster 3.0 U1 Release Notes</i>	806-7078
	<i>Sun Cluster 3.0 U1 Release Notes Supplement</i>	806-7079
	<i>Sun StorEdge Network Data Replicator 3.0 Release Notes</i>	806-7513
	<i>Sun StorEdge Instant Image 3.0 Release Notes</i>	806-7678
System Administration	<i>Sun Cluster 3.0 U1 System Administration Guide</i>	806-7073
	<i>Sun StorEdge Network Data Replicator 3.0 System Administrator's Guide</i>	806-7512
	<i>Sun StorEdge Instant Image 3.0 System Administrator's Guide</i>	806-7677

---

## Accessing Sun Documentation Online

A broad selection of Sun system documentation is located at:

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

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

<http://docs.sun.com>

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

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

---

## Ordering Sun Documentation

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

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

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

---

## Sun Welcomes Your Comments

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

[docfeedback@sun.com](mailto:docfeedback@sun.com)

Please include the part number (816-1544-10) of your document in the subject line of your email.



# Overview

---

---

**Note** – This guide assumes that you have already installed the volume manager software and the Sun Cluster 3.0 Update 1 software on each node in your cluster.

---

The Sun Cluster 3.0 Update 1 and Sun StorEdge 3.0 services software combine to provide a highly available environment for cluster storage. The Sun StorEdge Network Data Replicator (Sun SNDR) software is a data replication application that provides access to data as part of business continuance and disaster recovery plans. The Sun StorEdge Instant Image software is a point-in-time copy application that enables you to create copies of application or test data.

---

**Note** – Sun Cluster 3.0 Update 1 is also known as the Sun Cluster 3.0 07/01 release.

---

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 Services 3.0 Software in a Sun Cluster Environment” on page 5](#)
- [“Sun StorEdge Fast Write Cache” on page 8](#)
- [“Sun StorEdge Services Software Version Compatibility” on page 8](#)

---

# 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.

In the Sun StorEdge software documentation, the term *data services* is used to describe the Sun SNDR and Instant Image software, running in a noncluster environment. This guide uses the term *Sun StorEdge services* to describe the Sun SNDR and Instant Image software configured to run on a cluster.

## *Primary and secondary hosts and nodes*

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

- The primary and secondary hosts are physically-separate servers running the Sun SNDR 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

---

**Note** – You cannot use the Sun StorEdge Fast Write Cache (FWC) product (all versions, including the SUNWnvm Version 3.0 software) in any Sun Cluster environment because cached data is inaccessible from other machines in a cluster. To compensate, you can use a Sun caching array. See “[Sun StorEdge Fast Write Cache](#)” on page 8.

---

TABLE 1-1 Supported Software, Patches, and Hardware

---

<b>Operating Environment Software</b>	<p>Solaris™ 8 (initial release) Solaris 8 10/00 (also known as Update 2) Solaris 8 01/01 (Update 3) Solaris 8 04/01 (Update 4)</p> <ul style="list-style-type: none"><li>• Solaris 8 6/00 (Update 1) is not supported.</li></ul>
<b>Sun Cluster Software</b>	Sun Cluster 3.0 Update 1 software (also known as the 07/01 release)
<b>Volume Manager Software</b>	Solaris Disk Suite 4.2.1, Solaris Volume Manager Veritas Volume Manager (VxVM) 3.1
<b>Sun StorEdge Services Software and Patches</b>	<p>Sun StorEdge Version 3.0 software includes the Sun StorEdge core services and one or both of the Sun SNDR and Sun StorEdge Instant Image Version 3.0 software.</p> <p>The initial release of the Version 3.0 software requires that the following patches are installed. Patches are available from SunSolve at <a href="http://sunsolve.sun.com/">http://sunsolve.sun.com/</a></p> <p><b>Patch number</b> 111945-xx - Storage Cache Manager 111946-xx - Storage Volume Driver 111947-xx - Instant Image 111948-xx - Sun SNDR</p>
<b>Supported Cluster Configuration</b>	The Sun Cluster 3.0 Update 1 and Sun StorEdge 3.0 services software with patches are supported in a two-node cluster environment only.
<b>Hardware</b>	A CD-ROM drive connected to the host server where the Sun software is to be installed.

---

**TABLE 1-1** Supported Software, Patches, and Hardware

---

Disk space requirements:

- The Sun SNDR software requires approximately 1.4 Mbytes
  - The Instant Image software requires approximately 1 Mbyte
  - The Sun StorEdge configuration location requires 4.5 Mbytes
  - Supporting Sun StorEdge services packages require approximately 3 Mbytes
-

---

# Using the Sun StorEdge Services 3.0 Software in a Sun Cluster Environment



---

**Caution – Do not install or use the Sun StorEdge Version 3.0 Core Services, Sun SNDR, and Instant Image software with or without patches on servers operating the initial release of the Sun Cluster 3.0 software. The Version 3.0 software and patches are not cluster aware or co-existent in this Sun Cluster 3.0 environment.**

---

To use cluster failover features with the Sun StorEdge services software, your software environment requires the Sun Cluster Update 1 (07/01 release) software.

You can install and use the Sun StorEdge Version 3.0 Core Services, Sun SNDR, and Instant Image software with [patches](#) on servers running the Sun Cluster 3.0 Update 1 software environment. In this environment, the Sun StorEdge services software is *cluster aware*. See TABLE 1-2.

The Sun SNDR `sndradm` and Instant Image `iiadm` commands are used to control the Sun StorEdge services. You can use the command options `C tag` and `-C tag` in a Sun Cluster 3.0 Update 1 environment only, as described in the Sun SNDR and Instant Image system administrator's guides. 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 Services 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 Sun Cluster provides.	The Sun StorEdge Version 3.0 services software with <a href="#">patches</a> is cluster aware in a two-node, Sun Cluster 3.0 Update 1 software 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 Version 3.0 services software is cluster tolerant in a Sun Cluster 2.2 environment with the Sun Solaris 8 operating environment patch 109210-05 installed. It will not interfere with failover.

# Global and Local Use of the Sun StorEdge Services Software

---

**Note** – See [“Configuration Rules for the Instant Image and Sun SDR Software”](#) on page 22.

---

The Sun StorEdge services software can use volumes that are local or global devices. Global devices are those Sun StorEdge services or other volumes accessible from any cluster node and which will fail over under the control of the SunCluster 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` and `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 system administrator’s guides listed in [“Related Documentation”](#) on page x.

# Volumes Eligible for Use with the Sun SNDR and Instant Image Software

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 (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 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).

Instant Image supports Sun StorEdge and 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 Sun StorEdge Instant Image 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

---

# Sun StorEdge Fast Write Cache

---

**Note** – The Sun StorEdge Fast Write Cache product, all versions, is not supported in any cluster environment.

---

**All versions of the Sun StorEdge Fast Write Cache product are not supported** when using the Sun StorEdge Version 3.0 services software in a Sun Cluster 3.0 U1 environment because cached data is inaccessible from other machines in a cluster. To compensate, you can use a Sun caching array.

For example, the Sun StorEdge Core Services Version 3.0 CD contains the Sun StorEdge `SUNWnvm` Version 3.0 software package. This package is intended 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 with Sun SNDR and Instant Image Version 3.0 services software *in a nonclustered environment*.

---

# Sun StorEdge Services Software Version Compatibility

The Version 1.x, 2.0, and Version 3.0 Sun StorEdge services are binary incompatible. If your system includes Versions 1.x and 2.0 of the Sun StorEdge Instant Image software (including Instant Image 2.0.1 with Sun StorEdge Target Emulation 1.2), Sun SNDR software, or Sun StorEdge Fast Write Cache, you must remove them before installing the Version 3.0 services.

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 service, you must uninstall all Version 1.X, 2.0, and 2.0.1 services.

See the Sun SNDR and Instant Image installation guides listed in [“Related Documentation” on page x](#).

# Installing and Configuring The Sun StorEdge Services Software

---

---

**Note** – This chapter assumes that you have already installed the volume manager software and the Sun Cluster 3.0 Update 1 software on each node in your cluster.

---

The *Sun StorEdge Network Data Replicator 3.0 Installation Guide* and *Sun StorEdge Instant Image Version 3.0 Installation Guide* describe how to install the Sun StorEdge services software in a noncluster environment.

The installation steps to install this software in a Sun Cluster 3.0 Update 1 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 3.0 U1 environment. The topics in this chapter include:

- [“Shutting Down Nodes After Installing the Sun StorEdge Services Software” on page 10](#)
- [“Task Summary: Steps to Configure the Sun Cluster 3.0 U1 and Install the Sun StorEdge Services Software” on page 10](#)
- [“Disk Device Groups and the Sun StorEdge Services Software” on page 12](#)
- [“Specifying the Configuration Location During the Sun StorEdge Installation Process” on page 13](#)
  - [“Requirements for the Configuration Location” on page 13](#)
- [“Editing the Sun StorEdge Bitmap Parameter Files” on page 16](#)
- [“Shutting Down and Restarting the Nodes” on page 19](#)
- [“Supported Configurations for the Sun SNDR Software” on page 20](#)
- [“Supported Configurations for the Instant Image Software” on page 22](#)
- [“Configuration Rules for the Instant Image and Sun SNDR Software” on page 22](#)
- [“Configuring the Sun Cluster 3.0 U1 Environment for Use with the Sun StorEdge Services Software” on page 24](#)

---

# Shutting Down Nodes After Installing the Sun StorEdge Services Software

Because the Sun StorEdge services installation process requires you to shut down and restart each node in the cluster, *make sure that you install the Sun StorEdge services 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 U1 Concepts* manual describes this mechanism and the Cluster Membership Monitor (CMM).

See [“Shutting Down and Restarting the Nodes” on page 19.](#)

---

## Task Summary: Steps to Configure the Sun Cluster 3.0 U1 and Install the Sun StorEdge Services Software

The recommended installation order to use the Sun StorEdge services software in the Sun Cluster 3.0 U1 environment is as follows for each node:

1. Install the volume manager software
2. Install the Sun Cluster 3.0 software (including the Update 1 07/01 release)
3. Install the Sun StorEdge service software.

See TABLE 2-1.

**TABLE 2-1** Installation and Configuration Steps

Installation Steps	For More Information, See...
1. Install the Sun StorEdge Core Services software from CD on a cluster node.	Sun SNDR and Instant Image installation guides listed in <a href="#">“Related Documentation”</a> on page x.  <a href="#">“Supported Software and Hardware”</a> on page 3.  <a href="#">“Requirements for the Configuration Location”</a> on page 13.
2. Specify the configuration location during installation.	<a href="#">“Specifying the Configuration Location During the Sun StorEdge Installation Process”</a> on page 13
3. Install the Sun StorEdge services from CD on a cluster node.	Sun SNDR and Instant Image installation guides.
4. Edit the <code>/usr/kernel/drv/rdc.conf</code> or <code>/usr/kernel/drv/ii.conf</code> files, if necessary.	<a href="#">“Editing the Sun StorEdge Bitmap Parameter Files”</a> on page 16
5. Shut down and restart the node.	<a href="#">“Shutting Down and Restarting the Nodes”</a> on page 19
6. Repeat Step 1 through Step 5 for each additional cluster node.	
7. Configure the Sun Cluster software for use with the Sun StorEdge services.	<a href="#">“Supported Configurations for the Sun SNDR Software”</a> on page 20  <a href="#">“Configuration Rules for the Instant Image and Sun SNDR Software”</a> on page 22  <a href="#">“Configuring the Sun Cluster 3.0 U1 Environment for Use with the Sun StorEdge Services Software”</a> on page 24

---

# Disk Device Groups and the Sun StorEdge Services 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 3.0 U1 Environment for Use with the Sun StorEdge Services Software”](#) on page 24.

The SDS and VxVM device paths contain the disk device group. When operating in a Sun Cluster 3.0 U1 environment, the Sun StorEdge services software commands `sndradm` and `iiadm` automatically detect and use the disk device group as configured in [“Configuring the Sun Cluster 3.0 U1 Environment for Use with the Sun StorEdge Services Software”](#) on page 24.

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 [“The `iiadm` and `sndradm` Command Syntax”](#) on page 32.

---

# Specifying the Configuration Location During the Sun StorEdge Installation Process

When you install the Sun StorEdge Core Services software on the first cluster node, the installation process asks you to specify a block device for the single configuration location used by all Sun StorEdge services software you plan to install.

The following section describes how to specify the configuration location; the complete installation process is described in the *Sun StorEdge Network Data Replicator 3.0 Installation Guide* and *Sun StorEdge Instant Image Version 3.0 Installation Guide*.

## Requirements for the Configuration Location



---

**Caution** – When selecting a volume to be used as 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 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).

---

Procedures for installing the Sun StorEdge services in a cluster or noncluster environment are similar, except that in a Sun Cluster 3.0 U1 environment, the Sun StorEdge configuration location has different requirements.

In a Sun Cluster 3.0 U1 environment, the Sun StorEdge services configuration location must be available to all nodes running the Sun StorEdge services software. The requirements for this configuration location are:

- The configuration location must be on a shared disk; it cannot be on the quorum disk or any disk that might be locked during the boot process.
- The shared disk volume cannot be under volume manager control, like a disk device group that is mastered to a particular node.
- The shared disk volume must be under device ID (DID) control only.
- The configuration location requires 4.5 Mbytes of storage.

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

## ▼ To Specify the Configuration Location

---

**Note** – When you install the Sun StorEdge services on the second cluster node, specify the same configuration location that you specified on the first node. A warning message about the configuration already being initialized might appear. Ignore this message.

---

1. **Log on as the root user at the first cluster node where you are installing the Sun StorEdge services software.**  
You can install the software in single user or multiuser state.
2. **Insert the Sun StorEdge Core Services software CD into the CD-ROM drive that is connected to your system.**
3. **Start the Volume Manager daemon `vold(1M)` (if needed) and install the Sun StorEdge Core Services software.**

---

**Note** – If you are installing more than one Sun StorEdge service, you only need to start the Volume Manager daemon and install the Core Services software once. Do not start the daemon and install the Core Services software more than once.

---

```
# /etc/init.d/volmgt start
# cd /cdrom/cdrom0
# ./install_core
```

The Core Services software package installation starts. If you did not use a response file, the installation process prompts you as follows.

```
Do you want to specify the Sun StorEdge data services configuration
location? [y,n,?]
```

4. **Type `y` and press Return.**

The Core Services software prompts you as follows:

```
Where should the Sun StorEdge data service configuration be
located?
```

5. Enter a block device name for the single configuration location used by all Sun StorEdge service software you plan to install.

For example, you might specify `/dev/did/dsk/d39s0`.

6. Remove the Sun StorEdge Core Services software CD from the CD-ROM drive:

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

7. Continue installing the Sun StorEdge services software as described in the Sun SNDR and Instant Image installation guides.

---

**Note** – As described in the *Sun StorEdge Network Data Replicator 3.0 Installation Guide*, place the names and IP addresses of all machines you plan to use with the Sun SNDR 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 Sun SNDR software in the `/etc/hosts` file. Edit this file on each machine where you are installing and running the Sun SNDR software.

---

8. Shutdown and restart this node.

See [“Shutting Down Nodes After Installing the Sun StorEdge Services Software”](#) on page 10 and [“Shutting Down and Restarting the Nodes”](#) on page 19.

9. Log on as the root user at the next cluster node where you are installing the software and repeat [Step 2](#) through [Step 7](#).
10. See [“Mounting and Replicating Global Volume File Systems”](#) on page 31 for information about global file systems.

---

# Editing the Sun StorEdge Bitmap Parameter Files

---

**Note** – After editing the `/usr/kernel/drv/rdc.conf` file, shutdown and restart your node as described in [“Shutting Down and Restarting the Nodes” on page 19](#).

---

Bitmap volumes are used by the Sun SNDR and Instant Image software to track differences between volumes and provide information for volume updates. The Sun StorEdge services software documentation listed in [“Related Documentation” on page x](#) in describes the bitmap size and other requirements.

In a Sun Cluster 3.0 U1 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 Sun SNDR and Instant Image software include two configuration files that determine how bitmap volumes are written to and saved:

- Sun SNDR - `/usr/kernel/drv/rdc.conf`
- Instant Image - `/usr/kernel/drv/ii.conf`



---

**Caution** – **The Sun SNDR and Instant Image Version 3.0 software does not support bitmap files.** The software uses regular raw devices to store bitmaps. These raw devices should be located on a disk separate from the disk that contains your data.

---

## The /usr/kernel/drv/rdc.conf File

The /usr/kernel/drv/rdc.conf file contains one field that sets the Sun SNDR bitmap save mode:

- `rdc_bitmap_mode` - modify to change how the bitmap volume is saved during a shut down or system crash. In a Sun Cluster 3.0 U1 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 rdc.conf File

1. Open the /usr/kernel/drv/rdc.conf file using a text editor such as `vi`(1).
2. In a Sun Cluster 3.0 U1 environment, set the bitmap mode to 1.

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=1;
```

3. Save and exit the file.
4. Shut down and restart your server as described in [“Shutting Down and Restarting the Nodes”](#) on page 19.

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

The /usr/kernel/drv/ii.conf file contains one setting that sets the Instant Image 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 3.0 U1 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 3.0 U1 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 the Nodes”](#) on page 19.**

---

# Shutting Down and Restarting the Nodes



---

**Caution** – Because the Sun StorEdge services installation process requires you to shut down and restart each node in the cluster, *make sure that you install the Sun StorEdge services 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 U1 Concepts* manual describes this mechanism and the Cluster Membership Monitor (CMM).

---

After performing the steps listed in “[Task Summary: Steps to Configure the Sun Cluster 3.0 U1 and Install the Sun StorEdge Services Software](#)” 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 3.0 U1 System Administration Guide*.

---

## ▼ 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 services 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 Sun SNDR Software

---

**Note** – As described in the *Sun StorEdge Network Data Replicator 3.0 Installation Guide*, place the names and IP addresses of all machines you plan to use with the Sun SNDR 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 Sun SNDR software in the `/etc/hosts` file. Edit this file on each machine where you are installing and running the Sun SNDR software.

---

Three configurations for the Sun SNDR software are supported:

- [“The Sun SNDR Primary Host is Configured On a Cluster Node” on page 20](#)
- [“The Sun SNDR Secondary Host is Configured On a Cluster Node” on page 21](#)
- [“The Sun SNDR Primary and Secondary Hosts are Configured On a Cluster Node” on page 21](#)

## The Sun SNDR Primary Host is Configured On a Cluster Node

In this configuration, the Sun SNDR primary host is the logical host you created in the Sun SNDR resource group for the Sun SNDR disk group using the `scrgadm` command; for example, see [Step 6](#) in [“To Configure the Sun Cluster Environment” on page 25](#).

If you have configured the Sun SNDR autosynchronization feature on the primary host, the Sun SNDR software starts an update resynchronization from the primary host for all affected Sun SNDR volume sets following a switchover or failover event. This operation is performed after the resource group and network switchover operation is complete. See the `sndradm` man page and the *Sun StorEdge Network Data Replicator System Administrator’s Guide* for a description of the `sndradm -a` command to set the autosynchronization feature.

## The Sun SNDR Secondary Host is Configured On a Cluster Node

In this configuration, the Sun SNDR secondary host is the logical host you created in the Sun SNDR resource group for the Sun SNDR disk group using the `scrgridm` command; for example, see [Step 6](#) in “[To Configure the Sun Cluster Environment](#)” on [page 25](#).

Operations such as update resynchronizations occur and are issued from the primary host machine. Following a switchover (or failover) event, the Sun SNDR software attempts to start an update resynchronization for all affected Sun SNDR volume sets. However, the Sun SNDR secondary host in a Sun SNDR 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 Sun SNDR secondary host switchover appears to be a short network outage to the Sun SNDR primary host.

If you have configured the Sun SNDR 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 Network Data Replicator System Administrator's 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.

## The Sun SNDR Primary and Secondary Hosts are Configured On a Cluster Node

**Sun SNDR 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 SNDR volume set reside in the same disk device group. See “[Configuration Rules for the Instant Image and Sun SNDR Software](#)” on [page 22](#).

However, if the Sun SNDR primary and secondary hosts are configured in different clusters, see “[The Sun SNDR Primary Host is Configured On a Cluster Node](#)” on [page 20](#) and “[The Sun SNDR Secondary Host is Configured On a Cluster Node](#)” on [page 21](#) for operating considerations.

---

# Supported Configurations for the Instant Image Software

See [“Configuration Rules for the Instant Image and Sun SNDR Software”](#) on page 22.

---

## Configuration Rules for the Instant Image and Sun SNDR Software



---

**Caution** – Make sure that only one system administrator or root user at a time is creating and configuring Sun StorEdge volume sets to help avoid corrupting the Sun StorEdge services configuration. See [“Configuring Sun SNDR and Instant Image Volume Sets”](#) on page 30.

---

The information in this guide, including this section, only applies to the Sun SNDR and Instant Image Version 3.0 software installed in a Sun Cluster 3.0 U1 environment.

### Rules for the Instant Image Software

- **All Instant Image volume set components must reside in the same disk device group.** (An Instant Image volume set includes the master, shadow, bitmap, and optional overflow volumes.)
- **With the Instant Image 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 system failure or Sun Cluster failover occurs during an Instant Image 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 Instant Image Volume Data During a System Failure or Failover”](#) on page 40 describes how to avoid this situation.

## Rules for the Sun SNDR Software

- **The primary volume and its bitmap volume or the secondary volume and its bitmap volume in a Sun SNDR volume set must reside in the same disk device group per node.** (A Sun SNDR 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 Sun SNDR volume set.

- **With the Sun SNDR 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 Sun SNDR 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 Sun SNDR software requires that the `SUNW.HASStorage` resource is configured in the same resource group as the logical host, as described in the procedures in [“To Configure the Sun Cluster Environment” on page 25](#).

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
```

- **Sun SNDR 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 SNDR volume set reside in the same disk device group.

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

---

# Configuring the Sun Cluster 3.0 U1 Environment for Use with the Sun StorEdge Services Software

---

**Note** – The Sun StorEdge services software is supported in a two-node Sun Cluster 3.0 U1 environment only.

---

The procedures in this section describe how to configure the Sun Cluster software for use with the Sun SNDR and Instant Image software. The *Sun Cluster 3.0 U1 Data Services 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 steps are:

- 
1. Log on to any node in the cluster.
  2. Configure a disk device group using your volume manager.
  3. Register the `SUNW.HASStorage` resource type.
  4. Create a resource group.
  5. Add `SUNW.HASStorage` to the disk device group.
  6. (Sun SNDR step only) Add a logical failover host to the resource group.
  7. Enable the resource group and bring it online.
- 

When you complete the procedure described in [“To Configure the Sun Cluster Environment” on page 25](#), the resource group is configured and ready to use.

## ▼ To Configure the Sun Cluster Environment

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)`, `vxvg`, or `vxprint` commands, depending on your volume manager software.

3. Register `SUNW.HAStorage` as a resource type.

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

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.

## 5. Add a SUNW.HAStorage resource to the resource group.



---

**Caution** – Do not add additional resources to this lightweight resource group. Failure to follow this rule might cause the Sun StorEdge services software to not properly failover or switchover.

---

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

*devicegroup* is the disk device group name.

-x ServicePaths= specifies the extension property that the Sun StorEdge services software relies on. In this case, use the disk device *groupname*.

-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.

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

---

**Note** – Perform this step for the Sun SNDR volumes only. This step is not needed for Instant Image volumes.

---

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

`-j lhost-stor` is the 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 services 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.



## Using The Sun StorEdge Services `iiadm` and `sndradm` Commands

---

This chapter describes using the Sun StorEdge services commands `iiadm` and `sndradm` in a Sun Cluster 3.0 U1 environment. The topics in this chapter include:

- “Configuring Sun SNDR and Instant Image Volume Sets” on page 30
- “The `iiadm` and `sndradm` Command Syntax” on page 32
- “Putting All Cluster Volume Sets in an I/O Group” on page 39
- “Preserving Instant Image Volume Data During a System Failure or Failover” on page 40
- “A Directory Named `._` is Created After Mounting a Secondary Volume” on page 41

---

# Configuring Sun SNDR and Instant Image Volume Sets

The Sun SNDR and Instant Image administrator's guides listed in [“Related Documentation” on page x](#) describe how to configure volume sets. Make sure that only one system administrator or root user at a time is creating and configuring these volume sets from a single server or node to help avoid corrupting the Sun StorEdge services configuration.

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

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

# 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 services 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 3.0 U1 System Administration Guide* 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.**

---

# The `iiadm` and `sndradm` Command Syntax

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

- *Global devices* are Sun StorEdge services or other volumes accessible from any cluster node and which fail over and switch back under the control of the SunCluster framework.
- *Local devices* are Sun StorEdge services 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 Sun StorEdge services administrator's guides listed in [“Related Documentation” on page x](#) describe the full command syntax and options for `iiadm` and `sndradm`.

## Global Device Command Syntax

---

**Note** – During the initial enable of the Sun SNDR or Instant Image 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 36](#).

---

The Sun StorEdge services software automatically derives the disk device group name from the volume path when you first enable volume sets. During this initial enable operation, the Sun SNDR and Instant Image software creates a configuration entry for each volume set. Part of the entry is the disk device group name for use in a cluster.

Sun SNDR shows this name as `C tag`, where `tag` is the disk device group name. Instant Image shows this name as `Cluster tag: tag`.

## Sun SNDR Example

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

```
# sndradm -e host1 /dev/vx/rdisk/sndrdg/datavol /dev/vx/rdisk/sndrdg/datavolbm1 \  
host2 /dev/rdisk/clt3d0s0 /dev/rdisk/clt2d0s4 ip sync
```

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

```
# sndradm -i  
  
host1 /dev/vx/rdisk/sndrdg/datavol /dev/vx/rdisk/sndrdg/datavolbm1 \  
host2 /dev/rdisk/clt3d0s0 /dev/rdisk/clt2d0s4 ip sync \  
C sndrdg
```

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

## Instant Image Example

When you enable an Instant Image volume set on a cluster node (logical failover host):

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

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

```
# iiadm -i  
  
/dev/vx/rdisk/iidg/clt3s0d0: (master volume)  
/dev/vx/rdisk/iidg/clt3d0s4: (shadow volume)  
/dev/vx/rdisk/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 an Instant Image 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 Sun SNDR volume set, use the `C local` option as part of the *SNDR-set* volume set definition:

```
sndradm -e SNDR-set
```

where *SNDR-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 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 Sun SNDR volume set cannot reside in a shared disk device group or metaset.

---

## Instant Image Example

When you enable this Instant Image volume set where `local` indicates a disk device group:

```
# iiadm -C local -e ind /dev/rdisk/clt90d0s5 /dev/rdisk/clt90d0s6 \  
/dev/rdisk/clt90d0s7
```

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

```
# iiadm -i  
  
/dev/rdisk/iidg/clt90d0s5: (master volume)  
/dev/rdisk/iidg/clt90d0s6: (shadow volume)  
/dev/rdisk/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/rdisk/clt3d0s0  
  
ii: /dev/rdisk/clt90d0s5 /dev/rdisk/clt90d0s6 /dev/rdisk/clt90d0s6 I - l.localhost  
- -
```

## 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 32](#).

Typically, the Sun StorEdge services 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.

You can use more than one disk device group for cluster switchover and failover; ensure that your configuration follows the [“Configuration Rules for the Instant Image and Sun Sندر Software” on page 22](#).

### Example

This command makes an Instant Image volume set in the `iigrp2` disk device group wait for all copy or update operations to finish before you can issue other Instant Image commands.

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

# Which Host Do I Issue Commands From?

The Sun StorEdge services 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 “[To Configure the Sun Cluster Environment](#)” on page 25.

When you enable the Sun SNDR 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 Sun SNDR Commands From

Task	Where Command Is Issued	Comments
Assign a new bitmap to a volume set.	Primary and secondary host	Perform this command first on the host where the new bitmap resides and is being assigned, and then perform it on the other host.
Disable the Sun SNDR software.	Primary or secondary host	You can disable on one host, leave the other host enabled, and then re-enable the disabled host.  Perform this operation on both hosts if you are deleting a volume set.
Enable the Sun SNDR software.	Primary and secondary host	When enabling the Sun SNDR software for the first time, issue the command from both hosts.
Full forward or reverse synchronization (copy).	Primary host	Ensure that both hosts are enabled.
Forward or reverse synchronization (update).	Primary host	Ensure that both hosts are enabled.

**TABLE 3-1** Which Host to Issue Sun SNDR Commands From *(Continued)*

<b>Task</b>	<b>Where Command Is Issued</b>	<b>Comments</b>
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	
Update an I/O group.	Primary and secondary	

---

# 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.

---

The Sun SDR and Instant Image services 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. If you have created more than one disk device group, you can also assign volume sets from different disk device groups into one I/O group.

## Example

1. **Enable three Instant Image 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 Instant Image Volume Data During a System Failure or Failover

If a Solaris operating environment system failure or Sun Cluster failover occurs during an Instant Image 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 Instant Image control or return it to your storage pool.

---

## A Directory Named `._` is Created After Mounting a Secondary Volume

After you synchronize the primary and secondary Sun SNDR volumes, you might notice a directory named `._` if you perform a directory listing. This directory is created by the cluster file system. For example:

```
secondary_hostname# ls -a
.
..
._
.profile
bin
classes
[and so on]
```

You can ignore this directory or delete it. When you unmount the cluster file system, the directory disappears.



# Index

---

## B

- bitmap volume
  - parameter files, 16
- bitmap volumes
  - ii.conf file, 18
  - rdc.conf file, 17

## C

- C local, 34
- commands
  - iiadm and sndradm, 29
- configuration
  - location, 13
  - procedures, 24
  - rules, 22
  - supported for Instant Image software, 22
  - supported for Sun SNDR software, 20
- configuring the Sun StorEdge services, 9, 24

## D

- data service
  - defined, 2
- disk device groups, 12

## F

- Fast Write Cache, 8

## G

- global devices, 6
  - command syntax, 32
- grouping volume sets, 39

## H

- High Availability applications
  - see data services, 2

## I

- I/O group
  - grouping volume sets in a cluster, 39
- ii.conf file, 18
- Instant Image
  - configuration rules, 22
  - data in a system failover, 40
  - I/O groups, 39
  - ii.conf file, 18
  - iiadm command, 5, 32
  - supported configurations, 22

## L

- local devices, 6
  - command syntax, 34

## R

- rdc.conf file, 17
- reboot
  - shutdown and restart node, 19
- requirements
  - bitmap volumes, 16

## S

- services
  - data, 2
  - Sun StorEdge, 2
- shutdown and restart node, 19
- Sun SNDR
  - sndradm command, 32
- Sun Cluster
  - configuration rules with Sun StorEdge services, 22
  - using Sun StorEdge services in, 5
- Sun SNDR
  - configuration rules, 23
  - I/O groups, 39
  - rdc.conf file, 17
  - sndradm command, 5
  - supported configurations, 20
- Sun StorEdge services
  - bitmap volumes, 16
  - configuration
    - location, 13
    - rules, 22
  - configuring, 9
  - defined, 2
  - Fast Write Cache
    - SUNWnvm, 8
  - iiadm command, 6
  - in a Sun Cluster environment, 5
  - installing, 9
  - Instant Image, 2
  - reboot node, 19
  - sndradm command, 6
  - Sun SNDR, 2
- SUNWnvm, 5
- supported configurations, 20, 22
- syntax
  - global device command, 32
  - iiadm and sndradm commands, 32

- local device command, 34

## V

- volume sets
  - I/O group, 39