



Sun™ Cluster 3.0/3.1 and Sun StorEdge™ Availability Suite 3.2 Software Release Note Supplement

Sun Microsystems, Inc.
www.sun.com

Part No. 817-4225-10
December 2003, [Revision A](#)

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

Typographic Conventions	2
Supported Software and Hardware	3
Product Notes	4
Choosing the Configuration Location	4
Shutting Down Nodes	4
Creating and Configuring Volume Sets	4
Multiple Writers Not Supported	5
A Directory Named <code>._</code> is Created After Mounting a Secondary Volume	5
Import, Export, and Join Features Not Supported in a Sun Cluster Environment	6
Group Commands Fail For Volume Sets In a Sun Cluster Environment	7
Workarounds to Known Bugs	8
Upgrading the Availability Suite Software in a Sun Cluster Environment	8
Installation Documentation	9
Upgrade Steps Summary	9
▼ To Upgrade the Availability Suite Software	9

Release Note Supplement

This release note supplement contains important information about the Availability Suite 3.2 point-in-time copy and remote mirror software operating in a Sun™ Cluster 3.0 Update 3 or a Sun Cluster 3.1 environment.

This Sun Cluster Release...	Is Also Known As...
Sun Cluster 3.0 05/02	Sun Cluster 3.0 Update 3
Sun Cluster 3.1	The initial release

Note – The Sun StorEdge™ Availability Suite 3.2 software is not supported in a Sun Cluster 2.2 operating environment.

This supplement includes the following topics:

- [“Supported Software and Hardware” on page -3](#)
- [“Product Notes” on page -4](#)
- [“Workarounds to Known Bugs” on page -8](#)
- [“Upgrading the Availability Suite Software in a Sun Cluster Environment” on page -8](#)

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. % You have mail.
AaBbCc123	What you type, when contrasted with on-screen computer output	% su Password:
<i>AaBbCc123</i>	Book titles, new words or terms, words to be emphasized	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>lhost1</i>, <i>lhost2</i> \ -n <i>nafo0@node</i>, <i>nafo0@node</i></code>

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

Supported Software and Hardware

Note – You cannot use the Sun StorEdge Fast Write Cache (FWC) product (all versions) 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.

Note – The Sun StorEdge Network Data Replicator and Instant Image software (versions 3.0 and 3.0.1) and the Availability Suite 3.0 software contained the SUNWnvm package for Sun StorEdge Fast Write Cache 2.0 users. The 3.1 and 3.2 versions of the software do not contain or support any SUNWnvm version.

TABLE 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 and the Sun Cluster 3.1 software
Sun Cluster Software	Sun Cluster 3.0 Update 3 and Sun Cluster 3.1 initial release
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.
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: 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

Product Notes

This section describes the following topics:

- [“Choosing the Configuration Location” on page -4](#)
- [“Shutting Down Nodes” on page -4](#)
- [“Creating and Configuring Volume Sets” on page -4](#)
- [“A Directory Named ._ is Created After Mounting a Secondary Volume” on page -5](#)
- [“Import, Export, and Join Features Not Supported in a Sun Cluster Environment” on page -6](#)
- [“Group Commands Fail For Volume Sets In a Sun Cluster Environment” on page -7](#)

Choosing the Configuration Location

In a Sun Cluster environment, place the configuration database on a slice of the cluster quorum device.

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).

Creating and Configuring Volume Sets



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

The operations that access the configuration include but are not limited to:

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

Multiple Writers Not Supported

The Sun StorEdge Availability Suite software is a volume-based snapshot and remote replication suite that supports any application and database that has a single Solaris host writing data.

Applications or databases that are configured to allow multiple Solaris hosts to write data to a shared volume are not supported, for example, Oracle 9iRAC and Oracle Parallel Server software.

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

After you synchronize the primary and secondary remote mirror software 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.

Import, Export, and Join Features Not Supported in a Sun Cluster Environment

The export/import/join facility is *not supported* in a Sun Cluster environment.

The Sun StorEdge Instant Image/Availability Suite Point-In-Time Copy software is a point-in-time snapshot utility. When installed and configured, the Sun StorEdge Instant Image software provides the system administrator with the ability to make and to maintain replicated data storage systems.

A Sun StorEdge Instant Image snapshot set is comprised of three or more volumes; a master volume, one or more shadow and bitmap volumes, and for compact dependent shadows, an optional overflow volume. The original copy of the data is maintained on the master volume, where the primary business application stores data. Shadow volumes contain a snapshot of the data in the master volume that was taken at a given point in time. As the data changes in the master volume and the shadow volume, a bitmap volume keeps track of the blocks that change to be able to update the shadow or master volume. Only the blocks marked as changed by bitmap entries need to be copied.

Sun StorEdge Instant Image 3.x software has a feature to allow for a dual ported shadow volume to be deported back and forth between a primary and secondary host while under control of Sun StorEdge Instant Image software. This capability, called import, export, join (E/I/J) allows the shadow volume to be accessed by another host for processing that does not affect the original application host. Applications that often leverage this feature are data mining and backup.

Sun StorEdge Instant Image software's import, export, join functionality is not supported in a Sun Cluster environment. when deployed in a Sun Cluster environment, the Sun StorEdge Availability Suite DataServices are loosely integrated into the SunCluster failover framework. The Sun StorEdge Availability Suite DataServices are suspended or resumed by the SunCluster Membership Monitor after volume manager volumes are deported or imported and before file systems are mounted and applications started.

Failover within SunCluster software is on a resource group granularity. It is a requirement that all constituents of an Sun StorEdge Instant Image software set are in the same volume manager disk group. this allows SunCluster software to guarantee that these volumes will be available when the DataServices are started or resumed on a cluster mode. If constituents of an Sun StorEdge Availability Suite DataServices set are not available when they are resumed, then that set is off lined from the perspective of the Sun StorEdge Availability Suite DataServices software. If or when an application writes to this volume, which is no longer in the Sun StorEdge Availability Suite DataServices stack, then the consistency of that Sun StorEdge Availability Suite DataService for that volume is lost.

Volume managers, such as VxVM and LVM do not allow deportation of a single volume within a disk group. In a non cluster configuration where E/I/J is deployed, the shadow volumes that are to be exported are either not in a volume manager group, or in a group different than the constituents of the `iiadm` command volume set that will not be deported (master, bitmap, overflow).

As a result of this conflict in configuration rules, E/I/J is not supported in a Sun Cluster environment.

For additional information, see the Sun StorEdge Availability Suite product documentation.

Group Commands Fail For Volume Sets In a Sun Cluster Environment

When a volume set is created using the `iiadm` command, a group name can be specified. This feature makes it possible to use the group name to perform a command on more than one volume set at a time. For example, to update all shadows in sets that have the name `mygroup`, specify:

```
# iiadm -g mygroup -u s
```

However, in a cluster environment, the ability to perform a command on more than one volume set can run into problems. The cluster resource group to which a volume set belongs is stored separately from the `iiadm` command created group to which it belongs. This can create a situation where a group contains volume sets that span many cluster resource groups.

As a result, the `iiadm` command is not able to properly process the persistence data base (the `dscfg` file) for commands like `disable`, `copy`, `update`, or `wait`. If an attempt is made to disable all volume sets in the `mygroup` group with the command,

```
# iiadm -g mygroup -d
```

an error results. The solution is use commands like `disable`, `copy`, `update`, or `wait` on the individual sets instead of using the `-g` command syntax.

Workarounds to Known Bugs

This section provides workarounds to the following known bugs:

- **4898593:** In a two-node cluster with a one-to-one remote mirror system configuration; remote mirror with autosync ON; remote mirror reverse sync is not able to resume upon completion of Sun Cluster failover and network recovery.

Work Around: Run the `sndradm -n -r -m` command to resume remote mirror reverse sync manually.

- **4943413: Cluster failover during reverse sync makes mounted volume unusable.**

If a cluster failover occurs during a reverse sync to the primary volume, that volume is unusable on the second cluster node.

Work Around: Unmount the primary volume and remount it.

FFIN:I0830-1: Sun StorEdge Instant Image 3.0/3.1/3.2 Import/Export/Join functionality is not supported in a SunCluster environment.

URL:

<http://sunsolve.east.sun.com/cgi/retrieve.pl?type=0&doc=fins/I0830-1>

Upgrading the Availability Suite Software in a Sun Cluster Environment

To upgrade the software on non-cluster nodes, follow the procedures in the *Sun StorEdge Availability Suite 3.2 Software Installation Guide*, which is listed in the documentation in TABLE 2.

Installation Documentation

TABLE 2 Sun StorEdge Availability Suite Installation Documentation

Title	Part Number
<i>Sun Cluster 3.0/3.1 and Sun StorEdge Availability Suite 3.2 Software Integration Guide</i>	817-4224
<i>Sun StorEdge Availability Suite 3.2 Software Installation Guide</i>	817-2783

Upgrade Steps Summary

The general steps to upgrade the Sun StorEdge Availability Suite software in a Sun Cluster environment are as follows:

1. If possible, perform the upgrade procedure during scheduled maintenance time, although you can still perform this procedure while the cluster is live and online. See the Sun Cluster documentation for high availability software upgrade procedures and also [“Shutting Down Nodes” on page -4](#).
2. Place any Sun StorEdge Availability Suite resource groups in an offline state. See the `scswitch(1M)` man page.
3. Shutdown and restart the cluster node in single-user mode.
4. Remove earlier versions of the Sun StorEdge Availability Suite software.
5. Install the Sun StorEdge Availability Suite 3.2 software.
6. Shut down and restart the node into cluster mode.
7. Repeat [Step 1](#) through [Step 6](#) for the second node in your two-node cluster environment.

▼ To Upgrade the Availability Suite Software

1. Log on to the cluster node as the superuser.

2. Evacuate Sun StorEdge Availability Suite resource groups from the primary node:

```
# scswitch -S -h node
```

where:

-S Evacuate all resource and device groups under cluster control from the specified node.

-h *node* Specifies the name of the primary node.

3. Ensure that the resource groups have evacuated from the node successfully:

```
# scstat -D -g
```

where:

-D -g Shows the status for all resource and device groups under cluster control.

4. Reboot the cluster node into single-user mode.

```
# /etc/shutdown -i0 -g0 -y  
# ok boot -s
```

5. If you have the following patches, use `patchrm(1M)` to remove them in the order listed where *nn* specifies the patch revision.

Patch	Description
113057- <i>nn</i>	Availability Suite remote mirror patch
113056- <i>nn</i>	Availability Suite point-in-time copy patch
113055- <i>nn</i>	Storage Volume driver patch
113054- <i>nn</i>	Storage Cache Manager and Volume Driver patch

To find the exact patch revision, run the command:

```
# showrev -p | grep 11305
```

6. Execute the `install.sh -a` script to determine which packages must be removed.

7. Use `pkgrm` to remove all packages listed by `install.sh` script in the order listed. For example:

```
# pkgrm SUNWiiu SUNWiir SUNWrdcu SUNWrder SUNWnvm SUNWspsvu
SUNWspsvr SUNWscmu SUNWscmr
```

8. Install the Sun StorEdge Availability Suite 3.2 software according to the procedures in the manuals listed in TABLE 2.

- If the installation script finds an existing configuration location as in the case of an upgrade, the location is displayed and the script prompts you as follows.

```
The Sun StorEdge Data Services database configuration location
has already been set.
Current location: /dev/did/rdisk/d4s1

Would you like to keep its current location [y,n,?]
```

9. Type `Y`.

10. Shut down and restart the node:

```
# /etc/shutdown -i6 -g0 -y
```

11. Check that the software is upgraded:

```
# /usr/opt/SUNWesm/sbin/sndradm -v
SNDR version 3.2
# /usr/opt/SUNWesm/sbin/iiadm -v
Instantimage version 3.2
```

12. Repeat all steps in this section for the next node where you are upgrading the Sun StorEdge Availability Suite software.

