Sun™ Cluster 3.0 and Sun StorEdge™ Software Release Note Supplement

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

Part No. 816-5128-12
May 2003, Revision A

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

Sun StorEdge Availability Suite 3.1 Software Release Note Supplement

Note – The Sun StorEdge™ Availability Suite 3.1 point-in-time copy and remote mirror software is not supported in a Sun Cluster 2.2 operating environment.

This release note supplement contains important last-minute information about the Availability Suite 3.1 point-in-time copy and remote mirror software operating in a Sun™ Cluster 3.0 Update 1, 2, or 3 environment.

This Sun Cluster Release...	Is Also Known As...
Sun Cluster 3.0 07/01	Sun Cluster 3.0 Update 1
Sun Cluster 3.0 12/01	Sun Cluster 3.0 Update 2
Sun Cluster 3.0 05/02	Sun Cluster 3.0 Update 3

This supplement includes the following topics:

Typographic Conventions

Typeface	Meaning	Examples
`AaBbCc123`	The names of commands, files, and directories; on-screen computer output	Edit your `.login` file. Use `ls -a` to list all files. `% You have mail.`
`AaBbCc123`	What you type, when contrasted with on-screen computer output	`% su` `Password:`
AaBbCc123	Book titles, new words or terms, words to be emphasized	Read Chapter 6 in the *User's Guide*. These are called *class* options. You *must* be superuser to do this.
	Command-line variable; replace with a real name or value	To delete a file, type `rm` *filename*.
[ ]	In syntax, brackets indicate that an argument is optional.	`scmadm [-d` *sec*`] [-r` *n*`[:`*n*`][,`*n*`]...] [-z]`
{ *arg* | *arg*}	In syntax, braces and pipes indicate that one of the arguments must be specified.	`sndradm -R b {p|s}`

\	At the end of a command line, the \ (backslash) indicates that the command continues on the next line.	`scrgadm -a -L \`
		`-g` *groupname*`-stor-rg \`
		`-l` *lhost1*`,`*lhost2* `\`
		`-n nafo0@`*node*`,nafo0@`*node*

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 version of the suite does not contain or support any `SUNWnvm` version.

TABLE 1 Supported Software and Hardware

Operating Environment Software	Solaris™ 8 and Solaris 9; all releases that are supported by the Sun Cluster 3.0 Update 1, 2, and 3 software
Sun Cluster Software	Sun Cluster 3.0 07/01 software (also known as Update 1) Sun Cluster 3.0 12/01 software (also known as Update 2) Sun Cluster 3.0 05/02 software (also known as Update 3)
Volume Manager Software	Solstice DiskSuite 4.2.1, Solaris Volume Manager VERITAS Volume Manager (VxVM) 3.1 The Sun StorEdge software does not support metatrans (metapartition) devices created by using the Sun Solstice DiskSuite and Sun Volume Manager.
Sun StorEdge Software	Sun StorEdge Availability Suite 3.1 remote mirror and point-in-time copy software.
Supported Cluster Configuration	The Sun Cluster 3.0 07/01, 12/01, 05/02 releases and Sun StorEdge Availability 3.1 software are supported in a two-node cluster environment only.
Hardware	A CD-ROM drive connected to the host server where the Sun software is to be installed. Disk space requirements: • The remote mirror software requires approximately 1.4 Mbytes • The point-in-time copy software requires approximately 1 Mbyte • The Sun StorEdge configuration location requires 5.5 Mbytes • Supporting Sun StorEdge core packages require approximately 3 Mbytes

Product Notes

This section describes the following topics:

Using the Quorum Disk to Store the Configuration Location

The *Sun Cluster 3.0 and Sun StorEdge Software Integration Guide* states that the Sun StorEdge Availability Suite configuration information cannot be located on the quorum disk. This configuration location is used by the Sun StorEdge Availability Suite remote mirror and point-in-time copy software to store information about volumes and other features. You specify this location during software installation.

If you are using the Sun Cluster and Sun StoEdge Availability Suite software in a supported two-node cluster environment with the Oracle Parallel Server software, you may use the quorum disk to store the configuration location. See the *Sun Cluster 3.0 and Sun StorEdge Software Integration Guide* for other configuration location requirements.

Using the Point-In-Time Copy Software Import, Export, and Join Features in a Sun Cluster Environment

The point-in-time copy software import, export, and join features are not supported in a Sun Cluster environment at this time.

Shutting Down Nodes After You Install the Sun StorEdge Software

Because the Sun StorEdge installation process requires you to shut down and restart each node in the cluster, *make sure that you install the Sun StorEdge 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 Sun StorEdge Volume Sets

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

In a clustered environment, two or more administrators should not be writing to the Sun StorEdge Availability Suite software configuration at the same time. The operations that access the configuration include but are not limited to:

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

Switching Over Global Devices Only

The `scswitch`(1M) command enables you to manually switch all resource groups and device groups from the primary mastering node to the next preferred node. The *Sun Cluster 3.0 System Administration Guide* describes how to perform these tasks.

Local devices do not fail over and switch back; 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.

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.

Workarounds to Known Bugs

This section provides workarounds to the following known bugs:

None as of this release.

Upgrading The Availability Suite Software in a Sun Cluster Environment

Note – Perform these upgrade procedures during your scheduled maintenance time. See the Sun Cluster 3.0 documentation for high availability software upgrade procedures and also "Shutting Down Nodes After You Install the Sun StorEdge Software" on page 5.

You can upgrade your Sun StorEdge Availability Suite 3.0 software currently installed on a cluster node to the version 3.1 software. You can upgrade the following suite software components:

- Sun StorEdge Instant Image software: version 3.0.1 and version 3.0 (version 3.0 requires patches 111945-*xx* through 111948-*xx*, where *xx* is the revision level)
- Sun StorEdge Network Data Replicator software: version 3.0.1 and version 3.0 (version 3.0 requires patches 111945-*xx* through 111948-*xx*)

TABLE 2 lists the manuals containing complete installation and upgrade instructions for the Sun StorEdge Availability Suite software. To upgrade the software on non-cluster nodes, follow the procedures in the documentation in TABLE 2.

TABLE 2 Sun StorEdge Availability Suite Installation Documentation

Title	Part Number
Sun Cluster 3.0 and Sun StorEdge Software Integration Guide	816-5127
Sun StorEdge Availability Suite 3.1 Point-in-time Copy Software Installation Guide	816-4312
Sun StorEdge Availability Suite 3.1 Remote Mirror Software Installation Guide	816-4413

Upgrade Steps Summary

The general steps to upgrade the Sun StorEdge Availability Suite software in a Sun Cluster 3.0 05/02 software environment are as follows:

1. If possible, perform the upgrade procedure during scheduled maintenance time. If not, you can still perform this procedure while the cluster is live and online.
2. Place any currently-configured Sun StorEdge Availability Suite resource groups in an offline state. See the `scswitch`(1M) man page.
3. Reboot the cluster node.
4. Remove the Sun StorEdge Availability Suite 3.0.1 or 3.0 software.
5. Shut down and restart the cluster node in single-user mode.
6. Install the Sun StorEdge Availability Suite 3.1 software.
7. Place any currently-configured Sun StorEdge Availability Suite resource groups in an online state.
8. Shut down and restart the node into cluster mode.
9. Repeat Step 1 through Step 8 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 any currently-configured Sun StorEdge Availability Suite resource groups from this 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 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 into non-cluster mode.**

```
# reboot -- "-x"
```

5. **Remove the Sun StorEdge Availability Suite software according the procedures in the manuals listed in** TABLE 2**.**

Note – The software removal steps described as part of the upgrade procedure require you to shut down and restart your system in single-user mode to avoid volume data corruption.

---

6. **Install the Sun StorEdge Availability Suite 3.1 software according the procedures in the manuals listed in** TABLE 2**.**

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

```
The Sun StorEdge Data Services database configuration location
has already been set.
Current location: /dsfile

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

7. **Type Y.**

   The following message is displayed:

```
Run dscfg -u to upgrade database after all nodes of this cluster have been
upgraded
```

   In this case, type the configuration location upgrade command **`dscfg -u`** *after* you upgrade the Sun StorEdge Availability Suite software on *both nodes*. You execute this command in Step 11.

8. **Shut down and restart the node**:

```
# /etc/shutdown -y -g0 -i 6
```

9. **Check that the software is upgraded**:

```
# /usr/opt/SUNWesm/sbin/sndradm -v
SNDR version 3.1
# /usr/opt/SUNWesm/sbin/iiadm -v
Instantimage version 3.1
```

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

11. **After you upgrade the software on both nodes, type the following *on one node only*:**

```
# /usr/opt/SUNWscm/sbin/dscfg -u
```

The HAStoragePlus Resource Type

Note – The *Sun Cluster 3.0 and Sun StorEdge Software Integration Guide* contains configuration rules for the software and HAStorage resource type.

As described in the *Sun Cluster 3.0 5/02 Supplement* (part number 816-3380), the Sun Cluster 3.0 05/02 release introduced the HAStoragePlus resource type. This resource type enables you to make local file system partitions residing on global disk groups highly available. It performs the same functions as the HAStorage resource type and coordinates the startup order between resource groups and disk device groups.

With HAStoragePlus, the local file system partitions must reside on global disk groups with affinity switchovers enabled. The Sun Cluster environment must also be configured for failover.

The HAStorage resource enables you to define one device group (and resource group for the remote mirror software) that allows all volume sets in the group to failover. The HAStoragePlus resource type also enables you to failover individual volume sets that you specify using the `scrgadm`(1M) command. The HAStoragePlus resource type is useful for I/O-intensive applications such as the Sun StorEdge Availability Suite software. The *Sun Cluster 3.0 5/02 Supplement* describes the resource type in more detail.

See the following sections:

- "To Configure a HAStoragePlus Resource Type" on page 12
- "Configuring the HAStoragePlus Resource Types with Volume Sets" on page 14

▼ To Configure a HAStoragePlus Resource Type

Note – Before using this procedure, ensure that you have upgraded to the Sun Cluster 3.0 05/02 release software

---

1. **Log on as superuser to the first 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.HAStoragePlus as a resource type:**

```
# scrgadm -a -t HAStoragePlus
```

4. **Create a failover resource group for the disk device group** *devicegroup***:**

```
# scrgadm -a -g devicegroup-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 logical-hostname resource to the resource group:**

```
# scrgadm -a -L -g devicegroup-rg -l lhostname
```

`-l` *lhostname*	specifies a logical hostname by which clients communicate with the Sun StorEdge Availability Suite software in the resource group. This option can also be a comma-separated list of UNIX hostnames.

6. **Create the resource of type** `SUNW.HAStoragePlus`:

```
# scrgadm -a -j resource-name -g devicegroup-rg -t HAStoragePlus \
-x FilesystemMountPoints=/global/devicegroup \
-x AffinityOn=True
```

resource-name	is the resource name to add.
devicegroup-rg	is the resource group name you created in Step 4.
`-t HAStoragePlus`	specifies the HAStoragePlus resource type.
`-x FileSystemMountPoints=`	specifies the extension property that the Sun StorEdge Availability Suite software relies on. In this case, use the `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 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.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.

7. **Enable the resource group and place it online.**

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

Configuring the HAStoragePlus Resource Types with Volume Sets

"To Configure a HAStoragePlus Resource Type" on page 12 provides an example configuration. 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/rdsk/iidg/ii01,/dev/vx/rdsk/ii02, \
/dev/vx/rdsk/iidg/ii11,/dev/vx/rdsk/iidg/ii12,/dev/vx/rdsk/iidg/iibitmap1, \
/dev/vx/rdsk/iidg/iibitmap2
```

where:

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