# Sun StorEdge™ Availability Suite 3.2 Software Release Notes

Sun Microsystems, Inc.
www.sun.com

Part No. 817-2782-10
December 2003, Revision A

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

Please
Recycle

# Contents

# Sun StorEdge Availability Suite 3.2 Software Release Notes

This document contains important information about the Sun StorEdge™ Availability Suite 3.2 software that was not available at the time the product documentation was published. Read this document so that you are aware of issues or requirements that can impact the installation and operation of the Sun StorEdge Availability Suite 3.2 software.

# Remote Mirror Software New Features

The following features are new to the remote mirror software with the Sun StorEdge Availability Suite 3.2 software release.

## Disk-based Asynchronous Queues

Data can be queued on disk as well as in memory. Memory-based queues are the default. Disk-based queues allow:

- Larger local queues in asynchronous replication mode
- Larger bursts of I/O activity without impacting application response time

If a disk-based queue fills up, the remote mirror software goes to non-blocking, or scoreboarding mode.

### *Blocking Mode*

Blocking mode ensures write ordering of the packets to the secondary site. It is the default mode when operating in asynchronous mode.

If the asynchronous queue fills up when the software is running in blocking mode, the response time to the application can be affected adversely because the software must acknowledge each write operation before removing it from the queue. New write opertions are blocked from the queue until space is available.

### *Non-Blocking Mode*

Non-blocking mode is optional in asynchronous operation. In this mode, if the asynchronous queue fills up, software discards the queue is discarded and goes into logging mode.

In logging mode, the application's write operations are not blocked, but write ordering is lost because scoreboarding keeps track of changed sectors, not the order of changes. However, the application sees no significant degradation in response time.

To synchronize data on the primary and secondary sites after the filling of the queue and subsequent entry into logging mode, perform an update synchronization.

## Multiple Asynchronous Flusher Threads

The software now has the ability to use multiple flusher threads to increase the drain rate from the asynchronous queues. This allows multiple I/Os per consistency group or set on the network at one time. The default number of queue-flushing threads is two. If you prefer to use a remote mirror operation similar to the one in version 3.1, set the flusher threads to one. For enhanced performance in a low-latency network environment, increase the number of threads to greater than two.

When using multiple threads, write operations often arrive at the secondary site out of sequence. To prevent any problems at the secondary site, sequence numbers are added to all data write operations at the primary site. The secondary site manages the incoming data based on the sequence numbers, essentially restoring the write order. Write operations that arrive out of order are stored in memory until previous write operations arrive.

The use of multiple asynchronous flusher threads on the primary site requires more memory at the secondary site. Each set can result in a maximum of 64 pending requests in memory at the secondary site. The memory requirements for the secondary site depend on the number of groups or sets tracked, the number of requests, and the size of the write operations.

When the number of requests for a group or set reaches 64, the secondary site prevents the primary site from issuing any more requests for that group or set. If memory is not available when a packet arrives, the packet is rejected and all groups and sets go into logging mode at the secondary site.

## Write Coalescing

If possible, the remote mirror software combines, or coalesces, multiple sequential write operations to the primary volume into a single network write operation. The size of the write operations and the network packet size affect the remote mirror software's ability to coalesce the write operations. Write coalescing provides two important advantages:

- Improves the asynchronous queue's drain rate
- Improves network bandwidth utilization

## Protocol Changes

The protocol now takes advantage of the software's improved asynchronous flushing rate and its improved usage of network bandwidth. The protocol has been enhanced to work efficiently with the new disk-based asynchronous queues and the associated multiple flusher threads.

# Point-in-Time Copy Software New Features

The following features are new to the point-in-time copy software with the Sun StorEdge Availability Suite 3.2 software release.

## Caching of Bitmaps Only

In Sun StorEdge Availability Suite 3.2 point-in-time copy software, data is no longer cached. Bitmaps are cached. To simplify the data path, the read cache is available only for bitmap volumes. The ability to use the read cache on data volumes is no longer provided. Because many disk arrays, applications, and file systems provide read and write caching, this change does not adversely affect product performance.

## EBS PowerSnap Integration

The point-in-time copy software is now integrated with the Sun StorEdge EBS (Legato NetWorker) 7.1 PowerSnap module.

# Other New Features

The following features are also in the Sun StorEdge Availability Suite 3.2 software:

- `dsstat` utility: The following commands provided in previous versions have been removed and are replaced with the `/usr/opt/SUNWesm/sbin/dsstat` utility:
  - `/usr/opt/SUNWesm/sbin/sndrstat`
  - `/usr/opt/SUNWesm/sbin/scmadm -S`
- HA StoragePlus: This resource type in Sun™ Cluster 3.*n* environments is now supported.
- Increased Storage Volume Limit: The default limit for storage volumes in use is increased from 1024 in the earlier versions of the software to 4096 storage volumes in the current version. To increase the number of storage volumes, see the *Sun StorEdge Availability Suite 3.2 Software Installation Guide*.

# System Requirements

## Supported Software and Hardware In a Nonclustered Environment

TABLE 1 shows the supported software in a nonclustered environment.

TABLE 2 shows the supported hardware in a nonclustered environment.

If you have a SunSolve$^{SM}$ service subscription, patches are available at `http://sunsolve.sun.com/`

**TABLE 1** Supported Software, Nonclustered Environments

| Operating Environment and Software | Patches Required |
|---|---|
| Solaris™ 8 Operating System (Solaris OS) | None |
| Solaris 9 OS (update 3 or higher) | None |
| Sun StorEdge Availability Suite 3.2 remote mirror software | None |
| TCP/IP network transport such as SunATM™ or Gigabit Ethernet transports | None |
| Sun StorEdge Availability Suite 3.2 point-in-time copy software | None |
| Volume manager software | • Sun Solstice DiskSuite™ software<br>• Solaris Volume Manager<br>• VERITAS Volume Manager<br>The Sun StorEdge software does not support metatrans devices created by using the Sun Solstice DiskSuite and Solaris Volume Manager. |

**TABLE 2** Supported Hardware, Nonclustered Environments

| | |
|---|---|
| **Hardware** | A CD-ROM drive connected to the host server where the Availability Suite software is to be installed.<br>If you plan to export shadow volumes, you must store the shadow volume on a dual-ported drive.<br><br>The Sun StorEdge Availability Suite 3.2 software is supported on any Sun server or workstation that has an UltraSparc II or later processor and that is running a supported version of the Solaris OS. Hosts include but are not limited to:<br>• Sun Enterprise™ 220R, 250, 420R, and 450 servers<br>• Sun Enterprise 3500, 4500, 5500, 6500, and 10000 servers<br>• Sun Fire™ 3800, 4800, 4810, and 6800 servers<br>• Sun Fire 15K server<br>• Sun Ultra ™ 60 and 80 workstations<br>• Sun Blade™ 100 and 1000 workstations<br>• Sun Netra™ t 1400/1405 and 1120/1125 servers |
| **Disk Space** | Allocate approximately 15 Mbytes for the installation.<br>• The remote mirror software requires approximately 1.7 Mbytes.<br>• The point-in-time copy software requires approximately 1.9 Mbytes.<br>• The Sun StorEdge configuration location requires 5.5 Mbytes.<br>• Supporting packages require approximately 5.4 Mbytes. |
| **Supported Attached Storage** | The remote mirror software is storage-hardware independent. |

# Supported Sun Cluster Software

The Sun StorEdge Availability Suite 3.2 software is compatible with the following versions of Sun Cluster software:

- Sun Cluster 3.0 Update 3
- Sun Cluster 3.1

**Note –** If you are using any version of the Solaris 8 operating system with Sun Cluster 2.2, you cannot install Sun StorEdge Availability Suite 3.2 software. The products are incompatible.

See the *Sun Cluster 3.0/3.1 and Sun StorEdge Availability Suite 3.2 Software Integration Guide* for more information.

# Known Issues and Bugs

This section discusses issues that affect how you can use the software effectively.

## Known Issues

This section describes issues when using the Sun StorEdge Availability Suite 3.2 software.

### Shadow Volume Size

The software returns the following error message in several circumstances when the shadow volume of a point-in-time copy volume set is not the same size as the master volume.

```
Another package would not allow target to be changed at this moment
```

Whenever a point-in-time copy snapshot is taken, the volume that is used to create the snapshot, the shadow volume, is made to look exactly like the master volume, including matching the number of blocks. If the master volume is larger or smaller than the shadow volume's physical size, the shadow volume appears to be resized at the moment the snapshot is taken. Physically, the shadow volume has not changed size, but the point-in-time copy kernel module always reports its size to be the same size as the master volume. This can present several problems with the remote mirror software, which does not expect the size of the volume to change:

- If you plan to use a remote mirror primary or secondary volume as a point-in-time copy shadow volume, the master and the shadow of that volume set must be exactly the same size. If they are different, you get an error when attempting to enable the point-in-time copy volume set with the `iiadm -e` command.
- You might also see this error during a disable `iiadm -d` command. If the shadow volume of the existing point-in-time copy volume set is not the same size as the master volume and a remote mirror volume set was enabled using that shadow volume as the primary or secondary, the point-in-time copy software stops you from performing the disable operation.

## Accessing Statistics

The `dsstat` utility is now the only source of statistics for the operation of the software. The utility, `sndrstat`, is no longer supported. See the *Sun StorEdge Availability Suite 3.2 Administration and Operations Guide* for details of the `dsstat` utility.

## Bitmap Volume Placement

Due to the read-modify-write behavior of point-in-time bitmap volumes, placing many of them on a single RAID-5 volume introduces high I/O contention involving the disks associated with the RAID-5 volume. This I/O contention is noticeable at volume unmounting time, as seen when an application suite or system is being shutdown. Place bitmap volumes on RAID-1 sets, multiple RAID-5 sets, or on a cached-array disk controller.

## Rebooting Your Server Using the `shutdown` Command

When you install, remove, or reinstall the software, shutdown the system and restart in single-user mode. This provides the following protection while you are working:

- Prevents other users from getting access to data volumes
- Prevents the volumes from unmounting automatically

When you have completed these procedures, shut down and restart in multi-user mode.

**Caution –** Do not use the `reboot` command. Always use the `shutdown` command. The `shutdown` command ensures that any shutdown scripts in the `/etc/init.d` directory are executed.

## Upgrading the Solaris OS with the Availability Suite 3.2 Software Installed

If you installed the Sun StorEdge Availability Suite 3.2 software in one Solaris OS and now wish to upgrade to a newer release of the Solaris OS, follow this procedure:

1. **Remove the Sun StorEdge Availability Suite 3.2 software.**
2. **Upgrade the operating environment.**
3. **Reinstall the Sun StorEdge Availability Suite 3.2 software.**

## Creating and Configuring Sun StorEdge Volume Sets

**Caution –** Only one system administrator or root user at a time can create and configure Sun StorEdge shadow volume sets. This helps avoid inconsistency in the Sun StorEdge configuration.

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

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

## Fast Write Cache Not Supported

Fast Write Cache is no longer supported so when you use the `scmadm -o` command, the write hints, `wrthru` and `nowrthru`, are no longer displayed. In addition, the write hints are no longer displayed when the `scmadm` command is run with no arguments.

## Configurable Maximum Number of Devices That Can be Handled by StorEdge Data Services

In this release, the default limit of sv-enabled devices has increased to 4096 devices from the default limit of 1024 devices in the Availability Suite 3.1 release. If you are using VERITAS Volume Manager volumes and you have more than approximately 2250 volumes in any one disk group, the VERITAS Volume Manager might be unable to start the volumes after a reboot and then report that the configuration request is too large. If the VERITAS volumes fail to start, the Availability Suite software is also unable to resume the devices.

Resolve this problem by adding a second volume manager disk group and by dividing the volumes between the two disk groups. For additional information, see the *Sun Cluster 3.0/3.1 and Sun StorEdge Availability Suite 3.2 Software Integration Guide*, and the *Sun Cluster 3.0/3.1 and Sun StorEdge Availability Suite 3.2 Software Release Note Supplement*.

## Bugs

- **4930424: The software allows you to specify a disk queue of greater than 1 TB although this creates an error condition.**

  **Work around:** None. Do not specify a disk queue larger than the maximum size, 1 TB.

- **4942385: Long volume names cause warning messages to be cut off due to a buffer limitation. If you are using long volume names, some errors may be cut off early.**

- **4942997: sndr: sndradm unknown host: vol printed in ds.log**

  `diskq` operations are printing "unknown" when logging to `ds.log`.

- **4943430: Availability Suite 3.2 TNF racing panics a Solaris 8 system**.

  Sun StorEdge Availability Suite 3.2 TNF tracing on any system running these operating environments cause the system to panic:

  - Solaris 8, FCS initial release
  - Solaris 8, 10/00 release (also known as Update 1)

  **Work around:** Do not initiate TNF tracing with the `prex -k` command on any system running an operating environment prior to Solaris 8, Update.

- **4943317: Availability Suite 3.2 on a Solaris 8 system outputs unnecessary error message during shut down.**

  Executing `/etc/shutdown` on a system running Solaris 8 and Sun StorEdge Availability Suite 3.2 causes the following message to be displayed on the console:

```
scmadm: cache has been deconfigured
Cannot stop nskernd: 2 kernel threads still running
```

  As the Solaris 8 system is shutting down, this irrelevant message will be displayed. the `nskernd` daemon will be stopped successfully and will not prevent a clean system shut down.

# Release Documentation

The product documentation is located on the product CD in Adobe® Acrobat (PDF) format. To get access to this documentation:

1. **Change to the root user.**
2. **Insert the product CD into the CD-ROM drive that is connected to your system.**

3. **If the Volume Manager daemon `vold`(1M) is not started, use the following command to start it. This allows the CD to automount the `/cdrom` directory.**

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

4. **Change to the `Docs` directory.**

```
# cd /cdrom/cdrom0/Docs
```

From this location, you can view the documentation using the free Adobe Acrobat Reader software. If you do not have the Adobe Acrobat Reader software, this CD provides it in the `/cdrom/cdrom0/Acro_Read` directory. Install this to your local machine. The software is also available from Adobe Systems at `www.adobe.com`.

# Service Contact Information

If you need help installing or using this product, call 1-800-USA-4SUN, or go to:

`http://www.sun.com/service/contacting/index.html`