Sun StorEdge™ 5210 and 5310 NAS Appliance Release Notes

Sun Microsystems, Inc.
www.sun.com

Part No. 819-2857-10
May 2005, Revision A

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

Please
Recycle

Adobe PostScript

Contents

Sun StorEdge 5210 and 5310 NAS Appliance Release Notes

This document contains important information about the Sun StorEdge 5210 and 5310 NAS Appliance 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 might impact the installation and operation of the Sun StorEdge 5210 and 5310 NAS Appliance.

These release notes contain information for the Sun StorEdge 5310 software version 4.5.0.14, the GA release, which is backwards compatible with the Sun StorEdge 5210 software. You must upgrade to this version to have full Compliance (5310 only) and File Replicator functionality. Refer to "Installing Patches" on page 2 for details.

To determine your software and build version, from the Web Administrator navigation panel select System Operations -> Update Software. Alternatively, from the Command Line Interface (CLI), type `version`, and build 14 will display as `4.05 M0 (build 14)`.

When installing the Sun StorEdge 5310 Cluster, use these release notes with the *Sun StorEdge 5310 Cluster Setup Instructions - Read This First* (part no. 819-2355-11).

These release notes contain the following sections:

Note – Compliance-specific issues are relevant to the Sun StorEdge 5310 NAS Appliance only.

Important – The system supports only a single administrative login at one time.

System Requirements

The Sun StorEdge 5210 and 5310 NAS Appliance system ships with the Web Administrator software already installed. You do not need to install any software to manage the Sun StorEdge 5210 and 5310 NAS Appliance.

To access the Web Administrator management interface, you must have a network-attached computer running one of the following browsers. You must use a Java™ technology-enabled browser with Java Plug-In 1.3.1 (or later).

- Internet Explorer
- Mozilla™
- Netscape Navigator™

Note – To download the latest Java Plug-in software, go to `http://java.com`.

Installing Patches

Upgrade your system by downloading the latest version of software from `http://sunsolve.sun.com`. Select the `Patchfinder` link and then enter the patch number appropriate for your system.

- 118216 software for the Sun StorEdge 5210
- 119351 software for the Sun StorEdge 5310

Resolved Issues

- The SNMP attribute `system.sysDescr.0` is now set correctly.
- All checkpoints now list the Backup option.
- EMU boards in the 5210 JBOD are now properly instrumented and monitored.
- The NAS head no longer sends false battery errors under any circumstances.

- The FTP module in the NAS operating system now loads automatically at startup but remains disabled by default.
- The GUI can now show as many as three external JBODs instead of two.
- The console now displays all file volumes, even if more than 50 were created.
- Create and Delete bonding (port aggregate) occasionally caused an unresponsive system; this has been fixed.
- Large NDMP backup will not fill up `/dvol` with NDMP job files.
- SNMP attribute `se5210RaidBBUStatus` is set to “normal”.
- On a CSM RAID controller hardware failure, pulling out the controller with active I/Os no longer causes volumes to become read-only.
- Occasional disk and/or Fibre Channel errors while running I/Os no longer cause some volumes to be marked Read-Only.
- A RAID volume will now automatically start rebuilding on a 5310EU F during use of existing hot-spares.

Known Issues

- The date/time shown on the System Events Panel is the client's machine time and not the Sun StorEdge 5X10 NAS Appliance system time.
- NDMP recover has a mismatch on timestamp on the intermediate directories.

  Workaround: NDMP recover may restore the mid-level directories with an incorrect “creation” timestamp. However, the target directories and files will be recovered with the correct timestamp.
- NDMP DAR recovery fails with Backup Type set to “tar”.

  Workaround: NDMP recovery may fail if the backup type is set to “tar”. It is recommended that you set the backup type to “dump” for all NDMP backups.
- If the NAS is shut down in a method other than through the Web Administrator, or if the Web Administrator loses contact with the NAS, the browser may hang.

  Workaround: Close all instances of the Web Administrator and web browsers. After the system reboots, reopen a web browser and relaunch the Java browser interface.
- The Web Administrator does not indicate read-only volumes if there is a LUN failure.

  Workaround: You can check the volume status using the Telnet Menu/CLI.

- Setting the time zone from the Telnet Menu/CLI displays the following message on the Web Administrator:

  `Current time zone locale not found`

  Workaround: Ignore this message when using the Telnet Menu/CLI.

- The Notification Email URL field shows the hostname URL, and you may not be able to connect to the Web Administrator by clicking this URL.

  Workaround: If the DNS does not resolve the hostname, use the IP address to connect to the Sun StorEdge 5310 NAS Appliance.

- Help contents in the Web Administrator include RAID management topics.

  Workaround: In-band RAID management from the Sun StorEdge 5310 NAS Appliance Web Administrator is not supported in this release. Ignore these topics.

- The HA and Port aggregation bond IP address may not restore properly after you delete a bond.

  Workaround: Select a different IP address for the bond.

- Poor RX/TX optical signal strength may result in degraded performance.

  Workaround: If there are no other critical hardware errors and you see significant performance degradation, this degradation could be related to Fibre Channel link errors. Contact Sun Services for assistance (see "Service Contact Information" on page 12).

- When you run Configure NFS -> Setup Hosts -> Add user , the window contents don't refresh, and the system appears to be hung because of many entries in the NIS/NIS+ mappings.

  Please wait for the system to finish processing and repaint the screen. Do not reboot your system.

- Mounting a checkpoint volume causes a `pwd` command malfunction if you mount it using the following command:

  `mount -F nfs se5k:/vol01.chkpnt /z/vlcp` (incorrect)

  When a checkpoint volume is mounted as above, running the `pwd` command anywhere under `/z/v1cp` incorrectly displays the following result:

  `/`

  Workaround: Do not mount any `/vol*.chkpnt` volume itself; instead, mount its subdirectories. For example, the following mount command example avoids the problem:

  `mount -F nfs se5k:/vol01.chkpnt/cp2 /z/v1cp/cp2` (correct)

- Old exports appear when there are no associated volumes.

  Workaround: Manually delete the stale exported file system using the Web Administrator. In the navigation panel, select UNIX Configuration->Configure NFS->Configure Exports panel, and then delete the stale exported file system.

Cluster-Specific Issues

- Manual movement of LUNs between heads results in a zero capacity reading. This occurs during initial cluster setup or during the addition of new trays.

  Workaround: Run a manual disk scan from either the Web Administrator or the Telnet Menu/CLI, and the head will refresh the LUN capacity.

- When a cluster is in failover mode, if a volume is created from the Alone head on a LUN that was originally owned by the Quiet head, applications accessing that volume might get an EACCESS error during the cluster recovery process.

  Workaround: Restart the application once head recovery has been completed.

- When a cluster is in a degraded state (for example, one head is in the Alone state and the other head is in the Quiet state), the NFS/CIFS clients can mount file systems owned by Head1 using Head2's IP address or vice versa. This can cause the client I/O to stop after the cluster is back in the Normal state.

  Workaround: When mounting the file systems from the clients, always use Head1's IP address to mount Head1's file systems and Head2's IP address to mount Head2's file systems.

- The cluster may lose time and get out of sync when under extreme load.

  Workaround: NTP should be enabled when using clusters so that time is kept up to date and in sync between cluster nodes. If you find that your times are not up to date, enabling NTP and resetting the time will allow NTP to keep the clocks in sync.

- In a cluster configuration, if the QUIET head experienced system problems during recovery, some of its volumes may fail to mount on the ALONE head.

  Workaround: Using the Telnet Menu/CLI, perform the following command:

  ```
  mount  -f  /volname
  ```

- In a cluster configuration, before doing a recovery, check the partner head using the LCD to see if the head is in QUIET mode. Then do the recovery from the Web Administrator or Telnet Menu of the ALONE head.

- If you add a new tray, you cannot assign ownership of an unowned LUN with the Web Administrator.

  Workaround: Perform LUN ownership tasks using the Telnet Menu/CLI.

Compliance-Specific Issues (5310 Only)

- When a compliance license expires or is removed, the system needs to be rebooted for the changes to take effect. Once this occurs, files that were previously wormed will maintain compliance rules, but no new files may be wormed.
- The compliance feature of worming a file through Windows is turned off by default.

  Workaround: To turn the Windows trigger on, from the CLI use the command `fsctl compliance wte on`.

File Replicator Specific Issues

- If there is a system failure (such as a power failure) within 10 seconds of the start of a change role process, both systems may be set as the TARGET and there will be no MASTER, causing loss of the mirror.

  Workaround: Contact Technical Support for help in re-establishing your mirror.
- When any of the mirrors are in the Pause state, if the NAS server encounters a Link Failover or a Path Failover (or a Head Failover in a cluster configuration), the PAUSED mirrors resume the mirror operations.

  Workaround: Pause the mirrors again using the Telnet Menu interface.
- If you do a Change Role operation while there is heavy I/O activity on the master volume, the master might time out, and you might lose CIFS access to the volume.

  Workaround: Do a manual remount of the file volume from the CLI. For example, if the volume name is `volx`, do this from the CLI:

  ```
  nas-5310> umount /volx
  nas-5310> mount /volx
  ```
- During creation of a new mirror, if the Target/Mirror system does not have enough space and partitions, then the Source/Master system continously retries until enough space and partitions are available.

  Workaround: You can break the mirror, and recreate the mirror after enough space and partitions are available on the target system.
- RESYNC and PAUSE options are not available on the Web Administrator GUI.

  Workaround: These options are available from the Telnet Menu.
- When a mirrored volume is promoted using the Web Administrator GUI, there is no status message displayed on the GUI.

  Workaround: Check the system log to see the status of the operation.

- During mirroring with heavy I/O activity, or during mirroring with cluster systems, you might see the following messages in the logs of the target/mirror server:

  ```
  nmir: deseq_recv: The mirror log appears to be full.
  ```

  Workaround: These messages are for informational purposes only and the mirrors will continue with the normal operation. These messages can be safely ignored.

Addenda to the Documentation

This section includes information that is additional to or overrides information in the documentation.

Exporting a File Volume

You can export a file volume only to a set of hosts with root permission (like Solaris or UNIX) by adding the hosts to the "trusted group" using the Set Up Hosts screen.

Another way of doing this is to add the set of hosts to a Host Group and then export the required file volume against this group using the "with Map Root User set to Root User" option.

MIB files

You can retrieve the latest MIB files from `http://sunsolve.sun.com`.

NAS System Log Messages

If your system log contains error messages stating "`Unowned SFS2`" volumes, call Technical Support for assistance.

Identification of Specific Disks for Replacement

If you have a disk drive failure, use the log entry to help you identify the specific disk (you can interpret disk locations in both the system log and diagnostic reports the same way). Look at the following log entry example:

```
Controller 0 enclosure 0 row 0 column 6
```

To interpret such log entries, be aware of the following:

- Ignore any channel and target numbers.
- Controller numbering starts at 0. For example, the controllers in the first array are 0 (slot A) and 1 (slot B), and the controllers in the second array are 2 and 3.
- Enclosure numbering starts at 0 and is relative to the array to which it belongs. For example, if the first array has 2 enclosures they are identified as enclosure 0 and 1.
- Row numbering is always 0 for the Sun StorEdge 5310 Cluster system.
- Column numbering starts at 0 and specifies the slot number in the enclosure.

Thus, you can interpret the example as indicating slot 7 of the 1st enclosure in the 1st array.

Note – There is no standard way to identify which array is the first one and which is the second one. Typically, the first HBA port is connected to the first array, the second HBA port is connected to the second array, and so on.

Scheduling Multiple Checkpoints Per Volume

Scheduled checkpoints per volume are limited to 5, but multiple checkpoints may be specified per schedule.

The following is an example:

```
                          Days      Hours AM      Hours PM        Keep
    Enabled  Description   SMTWTFS   M1234567890E  M1234567890E  Days + Hours
1.     Y     MTWTF5am5pm   -*****-   -----*------  -----*------  1      0
2.     Y     SunWed1pm     *--*---   ------------  -*----------  0      12
3.     Y     MWFmidnight   -*-*-*-   *-----------  ------------  0      3
4.     Y     Weekend       *-----*   *-----*-----  *-----*-----  0      6
5.     Y     FriEvery2hrs  -----*-   *-*-*-*-*-*-  *-*-*-*-*-*-  0      2
```

File Replicator Best Practices

This section includes some “best practices” instructions for the file replicator software, including the following:

- Network Link Connection
- Mirroring
- Mirror Buffer
- Mirror Operations

Network Link Connection

You should have a reliable, private network connection between the two servers. This will reduce the network latency during synchronization of all the file system changes to the mirror system.

The interface type connecting these two servers can be 100Mb Ethernet or 1000Mb Ethernet. The servers may be directly connected using a cross-over cable, or connected through a switch or router.

If you are connecting the servers to a router, be sure to configure the static route setting to ensure that the mirroring data is directed through the private route.

If you are connecting the servers to a switch, create a virtual LAN (VLAN) for each server to isolate network traffic.

Mirroring

Sun StorEdge File Replicator software allows you to duplicate any or all of the file volumes of one Sun StorEdge NAS server onto another Sun StorEdge NAS server. The source server is called the active server and the target server is called the mirror server.

Before you begin, make sure you have the following:

- Two Sun StorEdge NAS servers are required for mirroring. The servers can be different models.
- The mirror server must contain an equal or larger amount of storage space than the file volumes to be mirrored.
- There must be a reliable, continuously available network connection with sufficient capacity between the active and mirror servers.

- Both servers must have the same version of the operating system installed.
- The active file volumes to be mirrored must be at least 1GB.

Mirror Buffer

Mirroring is performed on a per-volume basis. The mirror buffer stores file system write transactions while they are being transferred to the mirror server. The file volume free space on the active server is reduced by the allocation size of the mirror buffer.

The size of the mirror buffer depends on a variety of factors, but must be at least 100 MB, and the mirror buffer can never be more than half of the remaining free space on any given file volume.

In a normal scenario, it is recommended that you create a mirror buffer that is approximately 10% of the size of the file volume you are mirroring.

The size you choose should depend on how much information is being written to the file volume rather than the size of the file volume.

As a rule of thumb, the size of mirror buffer is directly proportional to the frequency of writes to the file volume and inversely proportional to the speed of the network connection between the two servers.

If there is high write activity to the file volume and a slow network connection between the two mirror servers, it is recommended to create a mirror buffer that is approximately 25-30% of the size of the file volume you are mirroring.

The size of the mirror buffer cannot be dynamically increased. To increase the size of the mirror buffer, you would have to break the existing mirror and create the mirror again with the new mirror buffer size.

Mirror Operations

You should have no I/O activity to the source file volume on the active server while the mirror is being created.

In the event a mirror cracks, (this happens if the connection between the two servers is down for sometime or if the mirror buffer is too small and there are lots of writes to the master volume), do the following:

1. **Establish a faster network connection between the two servers.**
2. **Quiesce all the I/O activity to the master file system, until the mirror reaches the In Sync state.**

3. **Mount the target file system on the mirror server as read-only from either the CIFS or NFS client. This file system can be used for backup or any read-only activity.**

   You can also combine checkpoints with the Mirroring functionality. When a checkpoint is created on the active server, the checkpoint also gets mirrored to the mirrored server. This can be used for scheduled backups or to give read-only checkpoint access to other users and applications.

Requirements and Limitations of File Replicator with a Cluster Configuration

Both the heads in the cluster configuration should have "Sun StorEdge File Replicator" license enabled.

Mirrors should be established only from and to a Head1 (serial number ending in "-H1"). The active server or mirror server of any mirror (when using the cluster configuration) should always be Head1, *not* Head2.

To perform any mirror management operations (these operations include New Mirror creation, Change Role, Promote, and Break), both heads in the cluster should be in the NORMAL state.

When the cluster is in failover mode, that is, one Head is in the ALONE state and the other Head is in the QUIET state, *do not* perform any mirror management operations. You should bring the cluster to the NORMAL state before doing any of the mirror management operations.

Existing mirrors will continue mirroring, even when the cluster configuration fails over. Also, the existing mirrors will continue mirroring when the cluster is restored after a failover.

There should not be any mirror created from Head1 to Head2 of the same cluster.

Release Documentation

The following documentation is included with this release:

Sun StorEdge 5310 Cluster Setup Instructions - Read This First

The remaining documentation is posted on the documentation Web site at:

`http://www.sun.com/hwdocs/Network_Storage_Solutions/nas`

The remaining documentation includes:

Sun StorEdge 5310 NAS Appliance Quick Reference Manual

Sun StorEdge 5310 NAS Appliance Hardware Installation, Configuration, and User Guide

Sun StorEdge 5310 NAS Appliance Software Installation, Configuration, and User Guide

Sun StorEdge 5310 NAS Appliance Safety and Compliance Guide

Sun StorEdge 5300 RAID Expansion Unit and Sun StorEdge 5300 Expansion Unit Safety and Compliance Guide

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`