



Sun StorEdge™ QFS and Sun StorEdge SAM-FS Release Notes

Version 4, Update 4

Sun Microsystems, Inc.
www.sun.com

Part No. 819-2759-10
September 2005

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

Copyright 2005 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, California 95054, U.S.A. All rights reserved.

U.S. Government Rights - Commercial software. Government users are subject to the Sun Microsystems, Inc. standard license agreement and applicable provisions of the FAR and its supplements.

Use is subject to license terms.

This distribution may include materials developed by third parties.

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, Jiro, Solaris, Sun StorEdge and QFS Linux Client Software are trademarks or registered trademarks 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 architecture developed by Sun Microsystems, Inc.

Red Hat® is a registered trademark of Red Hat, Inc.

Linux is a registered trademark of Linus Torvalds.

This product is covered and 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 lists 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 2005 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, California 95054, Etats-Unis. Tous droits réservés.

L'utilisation est soumise aux termes de la Licence.

Cette distribution peut comprendre des composants développés par des tierces parties.

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, Jiro, Solaris, Sun StorEdge et QFS Linux Client Software 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.

Red Hat® est une marque déposée de Red Hat, Inc.

Linux est une marque déposée de Linus Torvalds.

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 en matière de 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 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.

Sun StorEdge™ QFS and Sun StorEdge SAM-FS Release Notes, Version 4, Update 4

This document contains important information about the Sun StorEdge QFS and Sun StorEdge SAM-FS Version 4, Update 4 (4U4) release. This release incorporates design changes, feature changes, and function enhancements. The release also includes fixes to the software. System administrators and programmers who are familiar with these software products will see changes that can affect daily operations and automated scripts that were written to work with previous releases of this software. Study these release notes before upgrading to the Sun StorEdge QFS and Sun StorEdge SAM-FS 4U4 release.

If you are installing this product's base release and its software patches, study these release notes and the patch README files that are included with the software patches. The patch README files contain information that supplements the information in this document.

You can obtain a copy of the Sun StorEdge QFS and Sun StorEdge SAM-FS 4U4 software through Sun Microsystems or through your authorized service provider.

This document contains the following sections:

- "Features in This Release" on page 2
- "System Requirements" on page 14
- "Installing Packages and Patches" on page 16
- "Known Problems" on page 18
- "Release Documentation" on page 24
- "Service Contact Information" on page 25

Features in This Release

The following sections describe the new features in this release:

- “Support for Solaris on SPARC” on page 2
- “Support for Solaris OS on x64 Platforms” on page 2
- “Additional Linux Client Support” on page 5
- “Archive Support for Multiple Disk Volumes” on page 6
- “Additional Archiving Options” on page 7
- “Support for Multi-Owner Disksets in Solaris Volume Manager for Sun Cluster” on page 8
- “New Mount Options” on page 9
- “New `samu` Commands and Display” on page 9
- “File System Manager Changes” on page 10
- “Improved List I/O” on page 10
- “Modified WORM Functionality” on page 11
- “Improved Tape Drive Cleaning Management” on page 12
- “Additional `sfind` Options” on page 13
- “Newly Qualified Robot and Media” on page 14

Support for Solaris on SPARC

Release 4U4 of the Sun StorEdge QFS and Sun StorEdge SAM-FS software is supported for both Solaris 9 and Solaris 10 Operating Systems (OSs) on SPARC platforms. It is no longer supported for the Solaris 8 OS.

Support for Solaris OS on x64 Platforms

The 4U4 software version has added support for the Solaris 10 OS running on x64 AMD Opteron platforms.

The following are some important considerations for users who are combining or changing between SPARC and x64 hardware platforms:

- Sun StorEdge QFS and Sun StorEdge SAM-FS software is supported only for the Solaris 10 OS on x64 platforms (AMD64 architecture), not for the EM64T architecture. With the exception of the Sun StorEdge QFS shared Linux client, it is also not supported for any 32-bit x86 architectures.
- All functionality that is supported by Sun StorEdge QFS and Sun StorEdge SAM-FS on the SPARC platform is also supported on the x64 platform, except for the following:
 - Sun StorEdge QFS software for the Solaris 10 OS on x64 platforms does not support Sun Cluster software.
 - The ADIC/Grau, Fujitsu LMF, IBM3494, and Sony network-attached libraries are not supported on x64 platforms. StorageTek (STK) ACSLS-attached automated libraries are supported on x64 platforms.
 - Optical (MO and UDO) storage libraries and drives are not supported on x64 platforms.
 - SANergy software is not supported on x64 platforms.
 - The multireader file system mount options are not supported in mixed architecture environments (SPARC and x64).
- EFI labels are required on all shared disks if your Sun StorEdge QFS shared file system configuration contains both the Solaris 10 OS on x64 platforms and the Solaris 9 or Solaris 10 OS on SPARC platforms. See “Configuration of EFI Labels for Shared x64 and SPARC Volumes” on page 4 for information on relabeling disks.
- You can add Sun StorEdge QFS shared file system Linux clients to Solaris SPARC configurations that are using SMI VTOC8 disk labels and to Solaris AMD64 configurations that are using SMI VTOC16 disk labels. You can also add Sun StorEdge QFS shared file system Linux clients to these configurations when they are using EFI disk labels, but you may need to rebuild the Linux kernel for this capability. This is dependent on the particular Linux distribution. See “Linux Kernel Patch” on page 5 for more details.
- You must exercise caution when accessing the same SAN attached storage from a Solaris environment on both SPARC and x64 platforms. The Solaris OS on x64 platforms cannot interpret the SMI VTOC8 disk label created by the Solaris OS on SPARC platforms, and the Solaris OS on SPARC platforms cannot interpret the SMI VTOC16 disk label created by the Solaris OS on x64. This can make it appear as though a disk is unlabeled, when in fact it is labeled and in use by a platform of a different architecture type. For example, a disk that is labeled with SMI VTOC8 may have mounted partitions in use by Solaris on a SPARC platform, but will appear as unlabeled when viewed with the `format(1M)` partition command by Solaris on an x64 platform. If you make the mistake of running `fdisk(1M)` as prompted by the `format(1M)` command, you will destroy the contents of that disk.

- You cannot change the architecture type of the server responsible for control of the file system metadata operations (i.e. the server that was used to create the file system with the `sammkfs(1M)` command). For a Sun StorEdge QFS standalone file system, this means that you cannot mount the file system on a server that has a different architecture type from the one that created it. For a Sun StorEdge QFS shared file system, this means that you cannot change the architecture type of the metadata server or any potential metadata servers. This is because the different architectures use different byte-ordering schemes (endianness). However, you can migrate data from one architecture type to the other by copying the file system to temporary storage using either `qfsdump(1M)` or `samfsdump(1M)`, re-creating the file system using `sammkfs(1M)`, and then repopulating the file system with `qfsrestore(1M)` or `samfsrestore(1M)`.
- The Sun StorEdge Traffic Manager I/O multipathing feature (MPxIO) is disabled by default for the Solaris 9 OS and 10 OS on the SPARC platform and enabled by default for the Solaris 10 OS on x64. This feature should be configured the same way for all systems in your Sun StorEdge QFS shared file system configuration. It is configured in `/kernel/drv/scsi_vhci.conf` for the Solaris 9 OS and in `/kernel/drv/fp.conf` for the Solaris 10 OS.
- In a Sun StorEdge QFS shared file system environment, a configuration error will be generated if you have potential metadata servers of different architecture types (SPARC and x64) defined in the `/etc/opt/SUNWsamfs/hosts.fs` file.

Configuration of EFI Labels for Shared x64 and SPARC Volumes



Caution – Relabeling a disk will destroy the contents of that disk.

Use the Solaris `prtvtoc(1M)` command to determine whether a disk contains SMI or EFI labels. Under the Dimensions section of the output, SMI labels list the number of accessible cylinders, whereas EFI labels list the number of accessible sectors.

To convert disk labels from the default SMI VTOC8 to EFI, copy the file system to temporary storage using `qfsdump(1M)` or `samfsdump(1M)`, relabel the disks with EFI labels using the Solaris `format -e` command, re-create the file system using `sammkfs(1M)`, and repopulate the file system with `qfsrestore(1M)` or `samfsrestore(1M)`.

When using the Solaris `format -e` command to create EFI labels, you can select the partition command from the menu to create and modify partitions (slices). When doing this, you must specify a tag id name of `usr`, rather than `stand` or `unassigned`, for EFI labels.

Note that EFI labels reserve the first 34 sectors, which misaligns Sun RAID-5 storage from a performance perspective. Unless you realign the storage, you will incur a RAID-5 read/modify/write performance penalty when writing. You can avoid this performance penalty by selecting the proper starting sector for all disk partitions for your particular storage configuration. For example, an 8+P Sun StorEdge T3 array with a 64K block size should have starting sectors that are multiples of 1024 for all disk slices ($(8 * 64 * 1024) / 512 = 1024$). Similarly, a 5+P Sun StorEdge FC 3510 array with a 128K block size should have starting sectors that are multiples of 1280 for all disk slices ($(5 * 128 * 1024) / 512 = 1280$).

Additional Linux Client Support

The 4U4 release has added support for SuSE 8 Linux clients with shared Sun StorEdge QFS software. The software now supports both SuSE and Red Hat Linux clients for Sun StorEdge QFS shared environments.

Note – The Linux Shared QFS client was qualified with both SMI and EFI disk label types. A Linux kernel patch is needed for SMI disk label support. See “Linux Kernel Patch” on page 5.

The following versions of Linux were qualified for x64 64 bit architecture systems:

- Red Hat 3 Update 4 (2.4.21-27 x86_64 kernel) x64 64 bit AS and ES SMP; QLogic HBA failover and mdadm path failover were also tested.
- SUSE 8 Service Pack 4 (2.4.21-278 x86_64 kernel) x64 64 bit SMP; QLogic HBA failover was also tested.

The following versions of Linux were qualified for x86 32 bit architecture systems:

- Red Hat 3 FCS (2.4.21-4 kernel) x86 32 bit AS and ES SMP; Traffic Manager was also tested.
- Red Hat 3 Update 1 (2.4.21-9 kernel) x86 32 bit AS and ES SMP
- Red Hat 3 Update 2 (2.4.21-15 kernel) x86 32 bit AS and ES SMP
- Red Hat 3 Update 3 (2.4.21-20 kernel) x86 32 bit AS and ES SMP
- Red Hat 3 Update 4 (2.4.21-27 kernel) x86 32 bit AS and ES SMP; QLogic HBA failover and mdadm path failover were also tested.

For more information about the Linux clients, see the README file that is included on the Linux client software CD.

Linux Kernel Patch

The Solaris 10 OS includes a new partition ID for SMI labels on x86 or x64 platforms. Linux kernels with versions before 2.6.10 do not recognize this partition ID, so those kernels will need a patch in order to support Sun StorEdge QFS Linux client interaction with a Sun StorEdge QFS metadata server running the Solaris 10 OS on an x64 platform.

A patch is included with the Sun StorEdge QFS Linux client software (only required for kernels before 2.6.10). To install the patch, located on CD 2 of the software, navigate to the top level of the Linux kernel source directory and run `patch`. For example:

```
# cd /usr/src/linux-2.4
# patch -p1 < /<path to patch>/linux_2.4_smi.patch
```

Archive Support for Multiple Disk Volumes

During archiving to an online disk, the `-disk_archive` parameter specified in the `archiver.cmd` file allowed only one disk volume per archive set. In version 4U4 of the software, this restriction has been removed. You can now specify multiple disk volumes for an archive set. In addition, the `-disk_archive` parameter is no longer a valid parameter for the `archiver.cmd` file.

The volume serial number (VSN) associations section of the `archiver.cmd` file now assigns volumes to archive sets. Disk volumes are defined with a `dk` media type and the volumes are noted by one or more `vsn_expression` keywords, which are regular expressions. Each VSN in the list is used, without leading or trailing spaces, as input to `regcmp(3C)`. The compiled form is saved with the archive set copy definition. When a volume is needed for an archive set copy, each VSN, or each library or manual drive that has sufficient space and can be used for the archives, is used as the `subject` argument to `regex(3C)`. The archive set copy VSN expressions are used as the `re` argument to `regex(3C)`. If `regex(3C)` returns with a successful match, the volume is used for the archive set copy. You can run the `archiver` command to evaluate whether the new command file and disk volumes are assigned to the expected archive set copy definition.

In addition to a regular expression, you can also specify VSN pools from which disk volumes are to be selected. A VSN pool is a named collection of volumes. The pool is assigned a name, media type, and a set of volumes. For example:

```
vsnpools
data_pool dk disk0[0-5]
endvsnpools
```

```
vsns
arset0.1 dk disk10 disk1[2-5]
arset1.1 dk -pool data_pool
endvsns
```

Since disk volumes are supported in the VSN associations and pools section of the `archiver.cmd` file, the `-disk_archive` parameter is no longer needed to define a disk archive set.

Note – Use of the `-disk_archive` parameter will result in an invalid `archiver.cmd` file and the archiver will not run. Any `archiver.cmd` files created with earlier software versions must be edited to remove the `-disk_archive` parameter in order for archiving to work correctly in the 4U4 software. See the `archiver.cmd(4)` man page for details.

Additional Archiving Options

The following sections describe the new options that were added for the archiving process in the 4U4 release.

New Archive Set Validation Option

A new `-a` option has been added to the `archiver` command that validates a file's archive set. The `-a` option produces a line of output for each file found in an inode scan of the file system. The line lists present and future archive activity for the file. This option can be used to verify the `archiver.cmd` configuration before it is put into production and thus prevent data from being archived to the wrong media or being otherwise handled incorrectly.

The output is described in detail in the `archiver(1M)` man page. The basic format for each line of output follows:

```
type name ino.gen archive_set archive_info (last four fields)
```

The `archive_info` information is generally one of these four values:

- - (no archive copy required)
- `media.VSN`
- time when archiving is to begin
- time when unarchiving is to begin

New `-s` Release Attribute

The archiver now supports the release `-s` attribute in the `archiver.cmd` file. This attribute allows you to specify a specific partial release size rather than using the release `-p` default value. The syntax for this attribute is:

```
-release sXX
```

where `XX` is the specified partial release size.

New `-nftv` `archiver.cmd` Parameter

A new `archiver.cmd -nftv` (no file time validation) parameter has been added to prevent the validation of file access and modification times for files that have been "migrated" into a directory.

During determination of age, the access and modification times for files are validated to ensure that these times are greater than or equal to the file creation time, and less than or equal to the time at which the file is examined. This validation is intended to provide proper archiving and unarchiving. However, for files that have been migrated into a directory, this validation might not result in the desired behavior. The `-nftv` parameter can be used for such files to prevent the validation of file access and modification times.

Change to the `archiver.cmd` File `-release` and `-norelease` Parameters

The `archiver.cmd -release` and `-norelease` parameters are no longer mutually exclusive. If you want to make sure that the disk space is released immediately after all copies of an archive set have been archived, you can use the `-release` and `-norelease` options together. The combination of `-release` and `-norelease` will cause the archiver to release the archive set when all the copies having this combination are made. With this usage, the archive set is released immediately, rather than waiting for the releaser to be invoked, as is the case with the `-norelease` option alone.

New `archiver.cmd -recycle_minobs` Parameter

A new `archiver.cmd -recycle_minobs` parameter has been added to provide greater control for disk archive recycling. This option is used to set a threshold for the recycler's rearchiving process. When the percentage of obsolete files within an archived tar file on the disk reaches this threshold, the recycler begins moving the

valid files from the archive into a new tar file. Once all of the valid files have been moved, the original tar file is marked as a candidate to be removed from the disk archive. This option is ignored for removable media recycling.

Support for Multi-Owner Disksets in Solaris Volume Manager for Sun Cluster

Sun Cluster software now includes support for mirroring of disk volumes in software (RAID-1 only) through its multi-owner diskset feature with Sun StorEdge QFS and Solaris Volume Manager. This support is included with Sun Cluster software version 3.1 8/05 and later and with Sun StorEdge QFS software version 4U4 and later, and it requires Solaris 10, Update 1, which is currently pending release.

New Mount Options

The following new Sun StorEdge QFS mount options have been added:

- `noabr` and `abr`
For use in an Oracle RAC environment with Sun StorEdge QFS AIO only. These mount options disable or enable Application Based Recovery (ABR) of software mirrors. They apply only to Sun StorEdge QFS file systems built on Solaris Volume Manager mirrored volumes that support ABR.
- `nodmr` and `dmr`
For use in an Oracle RAC environment with Sun StorEdge QFS AIO only. These mount options disable or enable Directed Mirror Reads (DMR) of software mirrors. They apply only to Sun StorEdge QFS file systems built on Solaris Volume Manager mirrored volumes that support DMR.
- `nodio_szero` and `dio_szero`
The `dio_szero` option causes uninitialized areas of sparse files written with direct I/O to be zeroed when the area is accessed. This makes the sparse file behavior the same as that for paged I/O. By default, sparse files written by direct I/O do not have the uninitialized areas zeroed for performance reasons. The default is `nodio_szero`.

New samu Commands and Display

This section describes changes made to the `samu(1M)` utility.

Mount Point Option Commands

The following commands were added in order to set or clear mount point options:

- `abr eq` – Enables Application Based Recovery
- `noabr eq` – Disables Application Based Recovery
- `dmr eq` – Enables Directed Mirror Reads
- `nodmr eq` – Disables Directed Mirror Reads
- `dio_szero eq` – Turns on direct I/O sparse zeroing
- `nodio_szero eq` – Turns off direct I/O sparse zeroing

For more information, see the preceding section, “New Mount Options” on page 9.

Disk Archiving Support

A new `samu D` display has been added to support multiple disk volumes for an archive set. The `D` display keeps track of the disk media for disk archiving that has been defined in the `diskvols.conf` file. The dictionary contains information about each VSN, including the capacity, space remaining, and flags indicating the status of the VSN. The flag definitions are as follows:

- `l----` Volume is labeled; `seqnum` file has been created. This is set by the administrator to prevent the Sun StorEdge SAM-FS software from creating a new `seqnum` file.
- `-r---` Volume is defined on a remote host.
- `--U--` Volume is unavailable.
- `---R-` Volume is read only.
- `----E` Media error. Set when the Sun StorEdge SAM-FS software detects a write error on the disk archive directory.

A new `diskvols samu(1M)` command has also been added. This command can be used to set or clear the disk volume dictionary flags. It has the following format:

```
:diskvols volume [+flag | -flag]
```

File System Manager Changes

- The name of the SAM-QFS Manager web-based interface has been changed to File System Manager. Version 2.0 is the current release.
- The File System Manager daemon has been renamed from `sam-mgmt rpcd` to `fsmgmtd`.
- The File System Manager now has the ability to schedule and run metadata dumps. Using the metadata snapshot feature of File System Manager, you can schedule `samfsdump` to run at a certain time, or you can run it manually at any time. You can also use File System Manager to restore files or entire file systems from the dump files.

Improved List I/O

A new Sun StorEdge QFS interface has been added that permits a more efficient implementation of list I/O than the Solaris POSIX equivalent. The new interface specifies a list of I/O requests for a single file rather than using the POSIX standard, which is a list of requests for multiple files. Processing of I/O requests is handled completely inside the kernel. See the `qfs_listio(3)` man page for details.

Modified WORM Functionality

In the 4U4 software release the WORM-FS interface has been modified to be compatible with the new Compliance Archiving software for the Sun StorEdge 5310 NAS appliance. The existing WORM-FS interface using `ssum` is no longer supported, but previous enhancements to the `sls` and `sfind` utilities supporting the WORM-FS interface remain unchanged. With this new interface, utilities such as `touch` or programs incorporating the function `utimes()` can now update the access time, which modifies the WORM retention period. To maintain compatibility with the existing WORM-FS functionality, the request to set the retention period (update the access time) is intercepted by the Sun StorEdge QFS or Sun StorEdge SAM-FS software. The retention period is displayed, as in previous releases, in the `attribute_time` field of the `sls` output with the duration displayed below it.

With these enhancements, the `setuid` bit, when set by itself, can set the WORM bit on a file or directory. For example, you can use `chmod 4000` to set the WORM bit:

```
# chmod 4000 test
# sfs -D
test:
mode: -r-Sr--r--  links: 1  owner: root      group: other
length:          0  admin id: 0  inode: 1034.3
access:          Jul 13 2005  modification: Jul 13 15:45
changed:         Jul 13 15:45  retention-end: Jul 13 2005 15:46
creation:        Jul 13 15:45  residence: Jul 13 15:45
retention:       active      retention-period: 0y, 0d, 0h, 1m
```

The new access mode for a WORM file is the old mode read bits logically or'd with the SUID bit. In the example above, the old mode `-rw-r--r--` becomes `-r-Sr--r--` once the WORM trigger is applied.

The retention period semantics of WORM files remain the same as in previous releases. The retention period can be extended only after the WORM trigger is used on a file. Updates to the access time can either initially set or extend the retention period. If the new `atime` value is greater than the file's existing retention period, an attempt to extend the period will succeed.

If the new `atime` value is less than or equal to the existing retention period, the request will fail and return an error (`EPEERM`). When used on a WORM file, the `mtime` argument as passed by `touch` is ignored.

If a file with an `atime` in the past is made into a WORM file, the default retention period is applied. If no default period is provided at mount time, the system-wide default of 30 days applies. This functionality remains unchanged from previous releases. There is a new method for setting permanent retention, however. Applying a value of `INT_MAX` to the access time changes the retention period to permanent.

With this modification, there are new restrictions on WORM files. Executable files cannot be converted to WORM files. Applying the WORM trigger to a file with one or more execute permissions will fail and return an error (`EACCES`). Likewise, any attempt to enable write permission on a WORM file, to set the `setgid` bit (`S_ISGID`) or sticky bit (`S_ISVTX`), or to clear the `setuid` bit on a WORM file will fail and return an error (`EPEERM`). In addition, access control lists (ACLs) of WORM files cannot be modified. However, read access bits can be set or cleared, and the User ID (UID) or Group ID (GID) can be changed on WORM files.

For more information on the WORM-FS package and creating WORM files, see the *Sun StorEdge QFS Configuration and Administration Guide*.

Improved Tape Drive Cleaning Management

The tape drive TapeAlert (2E) clean now (20), clean periodic (21), expired cleaning media (22), and invalid cleaning media (23) flags have been integrated into the Sun StorEdge SAM-FS direct-attached auto-cleaning feature, along with the Sequential Device Access (3E) cleaning bit in parameter 256.

Beginning with the 4U4 release, the default setting for software-initiated tape drive cleaning is `off`, which forces the user to choose between a hardware- or software-initiated tape drive cleaning strategy.

A hardware-initiated cleaning strategy uses the media changer's built-in auto-cleaning feature. If the library tape drives are shared between Sun StorEdge SAM-FS and another backup application, the library's built-in auto-cleaning feature should be used.

A software-initiated cleaning strategy uses the Sun StorEdge SAM-FS auto-cleaning feature. The `logsense` option is an enhancement to this existing feature that prevents a drive from using expired cleaning media. To enable the Sun StorEdge SAM-FS auto-cleaning feature, disable the hardware-initiated cleaning and add the following line to the `defaults.conf` file:

```
tapeclean = all autoclean on logsense on
```

To invoke the legacy Sun StorEdge SAM-FS auto-cleaning feature which relies on sense data only for drive cleaning status, the following line should appear in the `defaults.conf` file:

```
tapeclean = all autoclean on logsense off
```

Note – When using the auto-cleaning feature with a library with more than two drives, it is recommended that you have at least two cleaning cartridges per Sun StorEdge SAM-FS catalog. If there are drives that need to be cleaned, and not enough cleaning cartridges, those drives will be put into a `DOWN` state.

Sun made the following additional auto-cleaning enhancements in 4U4:

- Fixed the auto-cleaning bug in which the `cleaning required` bit was cleared.
- Added Sun StorEdge SAM-FS auto-cleaning for STK ACSLS network attached libraries.
- Added STK ACSLS `cleandrive` command support.
- Integrated auto-cleaning with TapeAlert and the Sequential Device Access log sense cleaning flags.
- Added expired cleaning media detection via TapeAlert.
- Added SDLT320 expired cleaning media detection.

- Added the ability for Sun StorEdge SAM-FS to take down a drive if it remains dirty after cleaning.
- Changed the `drive needs cleaning` status so that this status no longer persists after a `samd start/stop` or a `drive down/on` command.
- Added the ability for expired cleaning media to be moved to an export slot, if supported by the device.

Additional `sfind` Options

The options `-archpos n`, `-archpos1 n`, `-archpos2 n`, `-archpos3 n`, and `-archpos4 n` have been added to the `sfind` command. These options allow searches for files that have any archive copy, or a specific archive copy, at a particular position on removable media or in a particular disk archive file. The position `n` may be specified as a decimal or hexadecimal number with `+` or `-` options, or as a disk archive relative path such as `d3/d2/d1/f201`. See the `sfind(1)` man page for more details and an example.

Newly Qualified Robot and Media

- The Quantum SDLT600 tape drive has been qualified for use with Sun StorEdge SAM-FS software.

Note – WORM support for the SDLT600 drive is not yet available.

- The SDLT-2 media type is now supported by `sam-stkd/acsls`.

System Requirements

This section describes some of the system requirements that must be met before you can use the Sun StorEdge QFS and Sun StorEdge SAM-FS 4U4 release.

Note – For more information about system requirements, see the *Sun StorEdge QFS Installation and Upgrade Guide* or the *Sun StorEdge SAM-FS Installation and Upgrade Guide*.

Operating System Requirements

The Sun StorEdge QFS and Sun StorEdge SAM-FS 4U4 releases require the following minimum operating system levels:

- Solaris 9 OS 4/03 for SPARC platforms
- Solaris 10 OS for SPARC and x64 platforms
- Red Hat Enterprise 3.0, AS, ES, and WS for x86/x64 platforms (Sun StorEdge QFS shared client only)
- SuSE Enterprise Server 8 for x64 platforms (Sun StorEdge QFS shared client only)

Some features might require specific operating system support levels. For information, see the feature descriptions.

Sun Cluster Requirements

The Sun Cluster 3.1 4/04 release supports configuration of Sun StorEdge QFS shared and highly available file systems. This capability requires a minimum operating system level of Solaris 9 4/03 or Solaris 10 and is supported only on SPARC platforms.

The Sun StorEdge QFS software supports ORACLE® Real Application Cluster (RAC) software at the ORACLE 9i Release 2 (9.2.0.1), Patchset 5 (9.2.0.5), patch 3566420 release level and above, and ORACLE 10g Release 1 (10.1.0.3 and above).

Note – ORACLE RAC 10g (10.1.0.3) on Sun StorEdge QFS 4U4 does not support NUMA optimization on entry-level servers (v210 or v240) with US-IIIi based processors.

Note – When Oracle 9i RAC databases are deployed on shared Sun StorEdge QFS using the support for Solaris Volume Manager for Sun Cluster, performance can be affected. Oracle 9i does not support the Oracle redo log files in ABR mode. If redo log files are deployed on Solaris Volume Manager with database system files, optimum performance can be affected by very high write activities to these files.

To avoid this situation, use separate underlying Solaris Volume Manager mirrors for each set of the redo log files used by a given RAC instance.

File System Manager Requirements

File System Manager (previously called SAM-QFS Manager) has been tested only with the Sun Java Web Console version 2.2.4 applications. Installing this product with any older Sun Java Web Console version breaks both applications. This product might work on newer Sun Java Web Console versions, but this has not been tested.

Tivoli SANergy Requirements

If you plan to configure the Sun StorEdge QFS file system with SANergy, verify that you have Tivoli SANergy File Sharing software at release level 2.2.4 with fixpack 3.2.5.0 or later. This software is available through Sun when you order the Sun StorEdge Multiplatform Client Software 2.2.4 for QFS. The current fixpack for this software can also be obtained from SunSolve. For more information about the SAN-QFS file system, see the *Sun StorEdge QFS Configuration and Administration Guide*.

Note – The current Tivoli SANergy product is classified as SANergy version 2, release 2.4, rather than version 3, release 2.0, as indicated in the SANergy documentation.

Installing Packages and Patches

The *Sun StorEdge QFS Installation and Upgrade Guide* and the *Sun StorEdge SAM-FS Installation and Upgrade Guide* provide information about installing version 4U4 of the software. If you are installing a patch release, see the README file that is distributed with the patch.

This section contains information about installation packages.

Upgrading File System Manager

File System Manager (previously called SAM-QFS Manager) version 2.0 supports the Sun StorEdge QFS and Sun StorEdge SAM-FS 4U3 and 4U4 releases. To upgrade from an earlier version of File System Manager, run the `fsmgr_setup(1M)` command. The *Sun StorEdge QFS Installation and Upgrade Guide* and *Sun StorEdge SAM-FS Installation and Upgrade Guide* include instructions for using this command.

The `fsmgr_setup(1M)` command removes all the packages that pertain to the previous release and installs the new versions of the packages. If you are upgrading from 1.1 to 2.0 and if you have turned tracing on in 1.1, the `fsmgr_setup` script will automatically preserve the trace settings and return tracing to the same level in 2.0. If you are upgrading from 1.0 to 2.0 and you have tracing turned on in 1.0, you need to re-enable tracing in 2.0 by following the instructions in the *Sun StorEdge QFS Installation and Upgrade Guide* or *Sun StorEdge SAM-FS Installation and Upgrade Guide*.

Required Solaris OS Patches

Sun Microsystems provides patches for the Sun products that are described in this section. Go to the following web site for a list of recommended patches:

<http://sunsolve.sun.com>

The following patches are required for Sun StorEdge QFS and Sun StorEdge SAM-FS support:

This Product	Running With	Patch Requirement
Sun StorEdge QFS	Sun Cluster and Solaris Volume Manager (SVM) and Multi-Owner Disksets	SVM patches containing fixes for Multi-Owner Diskset bugs 6220234, 6240367, 6274840, 6281300, and 6290815.
Sun StorEdge SAM-FS	The X6757A SBus Fibre Channel HBA, used to access either tape drives or libraries	Patch 112244-04 for Solaris 9 on SPARC platform: Hardware/FCode: SBus Dual Fibre Channel Host Adapter.
Sun StorEdge SAM-FS	Solaris 10 on x64 platforms using Fibre-attached or SCSI-attached tape drives	Solaris patch containing the fix for bug 6250131 (patch number not yet set).
Sun StorEdge SAM-FS	Solaris 10 on x64 platforms using SCSI-attached tape drives attached to Adaptec HBAs	Solaris patches containing the fixes for bugs 6254081 and 6290383 (in addition to the fix for bug 6250131) (Patch number not yet set)
Sun StorEdge QFS and Sun StorEdge SAM-FS	Fibre channel devices with Solaris 9 on a SPARC platform	Patch 111847-08 for Sun StorEdge SAN Foundation Kit 4.4.x.

This Product	Running With	Patch Requirement
Sun StorEdge QFS and Sun StorEdge SAM-FS	Solaris 10	<p>Solaris 10 patch 118822-18 (SPARC) which contains the fix for CR 6264443. No x64 patch is available at this time. This bug affects both SPARC and x64 systems and results in data corruption.</p> <p>A workaround is available, which can be used until the appropriate Solaris kernel update can be applied. The workaround should be removed when the kernel update is applied. The workaround is to disable <code>segkpm</code> by adding the following to <code>/etc/system</code>:</p> <pre>set kpm_enable=0</pre> <p>A reboot is required to activate this change.</p>

Known Problems

This section contains information about known software problems.

File System Manager Limitations

The Sun StorEdge QFS and Sun StorEdge SAM-FS 4U4 release includes File System Manager 2.0. File System Manager 2.0 does not interoperate with all Sun StorEdge QFS and Sun StorEdge SAM-FS 4U4 features. Specifically, File System Manager does not support the following features:

- Sun Cluster software
- Sun SAM-Remote software
- Linux clients
- Optical media
- WORM file system

In addition, File System Manager has the following limitations:

- In localized environments, File System Manager limits input text to ASCII text only.

- File System Manager supports Sun StorEdge QFS shared file systems with the following restrictions:
 - Before File System Manager can configure a new shared file system or collect information about an existing one, the participating hosts must be added on the Servers page.
 - You cannot create a shared file system that uses devices that are not SCSI compliant. File System Manager cannot identify these devices as available on multiple hosts.
 - You cannot configure archiving for a SAM-QFS shared file system through File System Manager. When you create a shared file system through File System Manager, the `nosam` mount option is enabled in the file system.
 - The number of clients (and potential metadata servers) that can be specified when the shared file system is created is limited to four. If more hosts are needed, they can be added at any time after the file system is created.

Note – You can use the Sun StorEdge QFS and Sun StorEdge SAM-FS command-line interface (CLI) to configure and monitor the features that File System Manager does not support.

File System Manager Problems

The following known problems exist in File System Manager:

- Do not close the pop-up window by clicking the X button in the upper left corner of the wizard screen in Internet Explorer or by clicking the hyphen (-) button in the Netscape™ or Mozilla™ browser. Always click the Cancel button to exit.
- The wizard button is sometimes incorrectly disabled. When this happens, click another tab in the application and then click the tab for the action that you want to perform. The wizard button should then be enabled.
- Each archiving file system on a Sun StorEdge SAM-FS server has a default archive copy that archives all files that are not explicitly members of an archive policy. If an `archiver.cmd(4)` file exists on the server, a volume serial name (VSN) must be associated with each file system's default archive copy.

When you create a file system through the web browser and an `archiver.cmd(4)` file already exists, File System Manager automatically creates a VSN association to an available or valid media type for the default archive copy.

When you create a file system and an `archiver.cmd(4)` file does not exist on the server, the system does not explicitly create the VSN association, and the default archiving behavior is retained. The Default Policy Copy details page may show incorrect information. In this situation, you can create an archive policy from the

Archive Management tab and apply the policy to the file system. This action creates an `archiver.cmd` file and creates the necessary VSN association for the file system's default archive copy.

To change these default copy definitions, go to the Archive Policies Summary page and click the link for the default archive policy for the file system (the default policy has the same name as the file system). Select copy 1 from the copy information table and configure disk archiving or VSNs for the copy.

- If `/opt` is not a part of the root file system and the SAM-QFS server is rebooted, the Solaris `sysevent` daemon (`syseventd`) starts up before `/opt` is available and does not find some of the `SUNwsamfs` binaries required for File System Manager to report any recent fault conditions or to send Snapshot Interrupted or File System Overflow notifications for that server. To enable the necessary binaries to be found, restart the Solaris `sysevent` daemon (`syseventd`) on the SAM-QFS server after reboot by issuing a `pskill -HUP syseventd` command.
- The following messages sometimes appear if you try to open the File System Manager in a browser:

```
Connect to hostname.domain:6789 failed (connection refused)
The connection was refused when attempting to contact
hostname.domain:6789
```

The system generates these messages under the following conditions:

- You typed the wrong URL (such as the wrong host name, domain, or port).
- The web server is not running on the specified host.
- The Java Virtual Machine running the web server has become unavailable. When this happens, the `/var/log/webconsole/console_debug_log` file on the host that is running the web server (as specified by *hostname*) contains lines that are similar to the following.

```
#
# The exception above was detected in native code outside the VM
#
# Java VM: Java HotSpot(TM) Server VM (1.4.1_03-b02 mixed mode)
#
# An error report file has been saved as /tmp/hs_err_pid24360.log.
# Please refer to the file for further information.
#
```

To remedy this problem, become superuser on the host that was supposed to run the web server (as specified in *hostname*) and issue the following command:

```
# /usr/sbin/smcwebserver restart
```

- Running software that blocks Internet pop-up windows while you are using File System Manager causes unpredictable and unexpected results. Disable such software.
- If you upgrade your Sun StorEdge QFS and Sun StorEdge SAM-FS software from a 4.0 release and you use File System Manager to make configuration changes, the software creates new versions of the `mcf`, `archiver.cmd`, and other configuration files. If you had comments in your existing files, they are removed in the new files. To retrieve these comments, see the configuration file backup copy in `/etc/opt/SUNWsamfs/cfg_backups`. The system retains the most recent 25 versions of each configuration file.
- To prevent script errors from popping up on your screen if you use the Internet Explorer browser, go to the Tools menu, select Internet Options, and then select Advanced. Turn off "Display a notification about every script error" and then click OK.
- If File System Manager runs extremely slowly, check whether two instances of the web server are running. To do this, run `ps -ef |grep noaccess`:

```
# ps -ef |grep noaccess
noaccess 21449      1  0 11:38:17 pts/1      2:31
/usr/j2se/bin/java -server -XX:+BackgroundCompilation
-Djava.security.manager -
```

If the system output shows more than one process, run `kill -9 <PID>` to kill all processes that are running by user `noaccess`. Then issue the command `smcwebserver start`.

- When installing File System Manager, you must enter a password for all users and roles as prompted by the install script. If you leave a password blank, the Java Web Console will display an "Authentication Error" message when you try to log in.
- Opening a new tab or window by right clicking on an HREF link is not supported within File System Manager. If you want to view two pages simultaneously, open a new window and navigate from the Server Selection page.

Sun StorEdge QFS and Sun StorEdge SAM-FS Problems

- Segmented files are not supported on Sun StorEdge QFS shared file systems.
- Sun StorEdge QFS shared file system clients do not support the `stage(1) -n` command. The entire file is staged back to the disk cache.

- If you use Sun Storage Automated Diagnostic Environment to collect diagnostic information in Sun StorEdge SAM-FS environments, be aware that either you or the Sun StorEdge SAM-FS software can take a device down. Sun Storage Automated Diagnostic Environment initiates an alarm for devices that are down regardless of whether a user or the software took the device down.
- The NFS export semantics on Linux SuSE 8 do not allow the export of Sun StorEdge QFS file systems.
- Sparse files are not supported on Linux clients.

Note – For more information about the Linux client, see the README file that is included on the Linux client software CD.

- When using large tape block sizes (generally above 128 kilobytes) on x64 platforms with less than 4 gigabytes of memory, you may get tape write failures with any of the following messages:

```
Aug 16 22:20:55 amur-mn scsi: [ID 107833 kern.warning]
WARNING:
/pci@0,0/pci1022
,7450@b/pci1077,10a@1,1/fp@0,0/tape@w500507630061fa00,0 (st5):
Aug 16 22:20:55 amur-mn alloc contig_mem failure: not enough
mem
Aug 19 14:33:09 amur-mn scsi: [ID 107833 kern.warning]
WARNING:
/pci@0,0/pci1022
,7450@b/pci1077,10a@1,1/fp@0,0/tape@w500507630061fa00,0 (st5):
Aug 19 14:33:09 amur-mn Cannot alloc contig buf for I/O for
2097152 blk size
Aug 19 14:33:09 amur-mn rootnex: [ID 561485 kern.warning]
WARNING: fp: coding error detected, the driver is using
ddi_dma_attr(9S) incorrectly. There is a small risk of data
corruption in particular with large I/Os. The driver should
be replaced with a corrected version for proper system
operation. To disable this warning, add 'set
rootnex:rootnex_bind_warn=0' to /etc/system(4).
Aug 19 14:33:09 amur-mn scsi: [ID 107833 kern.warning]
WARNING:
/pci@0,0/pci1022
,7450@b/pci1077,10a@1,1/fp@0,0/tape@w500507630061fa00,0 (st5):
Aug 19 14:33:09 amur-mn transport rejected
```

This will cause the tape to be marked with the `f` flag in the volume catalog. There is nothing wrong with the tape, and you can clear the `f` flag with the `chmed(1M)` command. A system reboot may be necessary if you continually encounter this problem.

The current workaround is to increase the system memory to at least 4 gigabytes.

- `snmp` traps do not work on Solaris 10 servers on the x64 platform. To receive traps from an x64 server, set the value of the `TRAP_PROGRAM` variable in `/etc/opt/SUNWsamfs/scripts/sendtrap` to `TRAP_PROGRAM="/usr/sfw/bin/snmptrap"`. By default it is set to `TRAP_PROGRAM="/opt/SUNWsamfs/sbin/snmptrap"`.
- Linux clients running SuSE 8 SP4 do not support writing to `mdadm` devices. Customers interested in multi-pathing support for SuSE 8 SP4 clients should use HBA failover with the appropriate driver from QLogic.

Tivoli SANergy File-Sharing Problems

The following problems can arise with SANergy:

- System panics can occur when you do either of the following:
 - Uninstall SANergy while using a fused mountpoint
 - Stop SANergy services while SANergy is in use
- Read operations on files that are marked with the `stage(1) -n` staging attribute, which suppresses staging, sometimes hang.
- During writes from a SANergy client, the software does not adhere to the `segment(1)` file attribute.
- Read operations on files that are marked with `release(1) -p`, which specifies partial staging, sometimes fail.
- An `ftruncate(3C)` operation does not always properly extend a file when issued from a SANergy client.
- An `mmap(2)` system call, if issued over SANergy during the use of the LAN option, sometimes aborts.

Localization Problems

In localized environments, be aware of the following product limitations:

- Sun StorEdge QFS and Sun StorEdge SAM-FS software does not support message displays in more than one language.
- Configuration files do not accept non-ASCII names or descriptions.

Note – No localized packages are available for the version 4U4 release.

Bugs

TABLE 1 shows the bugs that are known to exist in the Sun StorEdge QFS and Sun StorEdge SAM-FS software. TABLE 2 shows the bugs that are known to exist in the File System Manager software.

TABLE 1 Known Bugs in Sun StorEdge QFS and SAM-FS Software

Bug Number	Description
6276347	Synopsis: QFS shared file system hangs during Oracle RAC workload - truncate down lease starvation. User Action: Use the <code>samaio</code> driver on the database files.
6284606	Synopsis: QFS client hung while running <code>iozone</code> direct I/O test. User Action: Ensure that all shared QFS systems have 1 GByte or more of memory for version 4U4.
6313480	Synopsis: Null pointer dereference in <code>sam_get_acl</code> when running <code>samcheck</code> . Status: Results in a panic. Issue resolved in patch 120971-01 through 120976-01.
6317078	Synopsis: ASCII case differences prevent a shared file system from mounting. Status: Issue resolved in patch 120971-01 through 120976-01.
6320210	Synopsis: Panic in <code>__free_pages</code> when running <code>fsstress</code> . User Action: Use <code>fsync()</code> rather than <code>fdatasync()</code> .
6320250	Synopsis: QFS from NFS <code>modtime</code> updates incorrectly at sync. Status: NFS and shared QFS client acquires write and append leases, performs the write, and waits for the flush thread to walk the inodes and sync them with the metadata server. When the sync happens, the metadata server sets the <code>modtime</code> at the time of the sync. This results in unacceptable warnings from common programs (such as <code>emacs</code> or <code>vim</code>) about file <code>modtime</code> changes. This issue will be addressed in a future software release. User Action: Mitigate the <code>modtime</code> updates by setting the <code>meta_timeo</code> and <code>lease_timeo</code> QFS mount parameters to "1" (one) on both the metadata server and the shared QFS clients, and mount the NFS clients with the <code>-o noac</code> mount parameter.

TABLE 2 Known Bugs in File System Manager Software

Bug Number	Description
6254210	In the New File System wizard, pressing the Enter key takes you back to the previous page. Rather than using the Enter key, use the buttons that are provided in the wizard interface.

Release Documentation

Information on the version 4U4 release is available from the following sources, in addition to these Release Notes:

TABLE 3 Sun StorEdge QFS and Sun StorEdge SAM-FS Documentation

Title	Part Number
<i>Sun StorEdge SAM-FS Installation and Upgrade Guide</i>	819-2754-10
<i>Sun StorEdge SAM-FS File System Configuration and Administration Guide</i>	819-3635-10
<i>Sun StorEdge SAM-FS Storage and Archive Management Guide</i>	819-2755-10
<i>Sun StorEdge SAM-FS Troubleshooting Guide</i>	819-2756-10
<i>Sun StorEdge QFS Configuration and Administration Guide</i>	819-2758-10
<i>Sun StorEdge QFS Installation and Upgrade Guide</i>	819-2757-10
Sun StorEdge QFS Linux client README file (on the Linux client software CD)	

Service Contact Information

To contact Sun technical support, go to the following web site:

<http://www.sun.com/service/contacting>

For installation and configuration services, please contact Sun Global Customer Service at 1-800-USA4SUN, or contact your local Sun sales representative.

