



Sun StorageTek™ QFS Linux Client Guide

Version 4, Update 6

Sun Microsystems, Inc.
www.sun.com

Part No. 819-7937-10
April 2007, Revision A

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

Copyright 2007 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, California 95054, U.S.A. All rights reserved.

Sun Microsystems, Inc. has intellectual property rights relating to technology that is described in this document. In particular, and without limitation, these intellectual property rights may include one or more of the U.S. patents listed at <http://www.sun.com/patents> and one or more additional patents or pending patent applications in the U.S. and in other countries.

This document and the product to which it pertains are distributed under licenses restricting their use, copying, distribution, and decompilation. No part of the product or of this document may be reproduced in any form by any means without prior written authorization of Sun and its licensors, if any.

Third-party software, including font technology, is copyrighted and licensed from Sun suppliers.

Parts of the product may be derived from Berkeley BSD systems, licensed from the University of California. UNIX is a registered trademark in the U.S. and in other countries, exclusively licensed through X/Open Company, Ltd.

Sun, Sun Microsystems, the Sun logo, AnswerBook2, docs.sun.com, Solaris, SunOS, SunSolve, Java, JavaScript, Solstice DiskSuite, and Sun StorageTek are trademarks or registered trademarks of Sun Microsystems, Inc. in the U.S. and in other countries.

All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the U.S. and in other countries. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc.

Mozilla is a trademark or registered trademark of Netscape Communications Corporation in the United States and other countries.

The OPEN LOOK and Sun™ Graphical User Interface was developed by Sun Microsystems, Inc. for its users and licensees. Sun acknowledges the pioneering efforts of Xerox in researching and developing the concept of visual or graphical user interfaces for the computer industry. Sun holds a non-exclusive license from Xerox to the Xerox Graphical User Interface, which license also covers Sun's licensees who implement OPEN LOOK GUIs and otherwise comply with Sun's written license agreements.

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

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 2007 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, Californie 95054, Etats-Unis. Tous droits réservés.

Sun Microsystems, Inc. a les droits de propriété intellectuelle relatants à la technologie qui est décrit dans ce document. En particulier, et sans la limitation, ces droits de propriété intellectuelle peuvent inclure un ou plus des brevets américains énumérés à <http://www.sun.com/patents> et un ou les brevets plus supplémentaires ou les applications de brevet en attente dans les Etats-Unis et dans les autres pays.

Ce produit ou document est protégé par un copyright et distribué avec des licences qui en restreignent l'utilisation, la copie, la distribution, et la décompilation. Aucune partie de ce produit ou document ne peut être reproduite sous aucune forme, par quelque moyen que ce soit, sans l'autorisation préalable et écrite de Sun et de ses bailleurs de licence, s'il y en a.

Le logiciel détenu par des tiers, et qui comprend la technologie relative aux polices de caractères, est protégé par un copyright et licencié par des fournisseurs de Sun.

Des parties de ce produit pourront être dérivées des systèmes Berkeley BSD licenciés par l'Université de Californie. UNIX est une marque déposée aux Etats-Unis et dans d'autres pays et licenciée exclusivement par X/Open Company, Ltd.

Sun, Sun Microsystems, le logo Sun, AnswerBook2, docs.sun.com, Solaris, SunOS, SunSolve, Java, JavaScript, Solstice DiskSuite, et Sun StorageTek 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.

Mozilla est une marque de Netscape Communications Corporation aux Etats-Unis et à d'autres pays.

L'interface d'utilisation graphique OPEN LOOK et Sun™ a été développée par Sun Microsystems, Inc. pour ses utilisateurs et licenciés. Sun reconnaît les efforts de pionniers de Xerox pour la recherche et le développement du concept des interfaces d'utilisation visuelle ou graphique pour l'industrie de l'informatique. Sun détient une licence non exclusive de Xerox sur l'interface d'utilisation graphique Xerox, cette licence couvrant également les licenciées de Sun qui mettent en place l'interface d'utilisation graphique OPEN LOOK et qui en outre se conforment aux licences écrites de Sun.

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.



Contents

Sun StorageTek QFS	
Linux Client Guide	v
Introduction to the Linux Client	v
Installing and Uninstalling the Linux Client	vi
▼ To Install the Linux Client	vii
▼ To Configure Access to man Pages	viii
▼ To Uninstall the Linux Client	viii
Linux Differences	viii
Differences in Function	ix
mcf File Differences	ix
Mounting Differences	x
Execution Differences	x
Performance Differences	xi
Linux Kernel Patch	xi
SELinux	xi
Troubleshooting	xii
Troubleshooting Tools	xii
Frequently Asked Questions	xiii

Sun StorageTek QFS Linux Client Guide

This document contains important information about Linux clients in the Sun StorageTek™ QFS file system in Release 4, Update 6 (4U6). The document is written for system administrators and programmers who are already familiar with the Sun StorageTek QFS software running on the Solaris™ Operating System (OS). It explains the differences between installation and configuration on the Solaris OS and installation and configuration on the Linux OS.

Before installing the Sun StorageTek QFS software on a Linux client, you should be familiar with the detailed installation procedure for Sun StorageTek QFS software as outlined in the *Sun StorageTek QFS Installation and Upgrade Guide*. You should also have an experienced system administrator's knowledge of the supported Red Hat Enterprise or SuSE Linux operating environment in which you are installing the Sun StorageTek QFS software.

This document contains the following sections:

- [“Introduction to the Linux Client” on page v](#)
- [“Installing and Uninstalling the Linux Client” on page vi](#)
- [“Linux Differences” on page viii](#)
- [“Troubleshooting” on page xii](#)

Introduction to the Linux Client

The following Linux versions are supported in the 4U6 software release:

- Red Hat Enterprise Linux 3.0 (UD-5 and UD8) for x86/x64 platforms
- Red Hat Enterprise Linux 4.0 (UD-2 and UD-4) for x64 platforms

- SuSE Linux Enterprise Server 8 (service pack 4) for x64 platforms
- SuSE Linux Enterprise Server 9 (service pack 2) for x64 platforms and SGI Altix Itanium systems
- SuSE Linux Enterprise Server 10 for x64 platforms

Note – This is the last version of Sun StorageTek QFS that will support SuSE Enterprise Linux 8. It will not be supported in the next version of the software.

Unlike shared Sun StorageTek QFS Solaris clients, the Linux client is restricted to client-only behavior. It cannot be configured as a potential metadata server. The Linux client supports interaction with Sun StorageTek Storage Archive Manager (SAM) software, but does not support Sun StorageTek SAM commands such as `stage`, `archive`, `release`, and `samu`. It has Sun StorageTek QFS file system functionality only.

The Sun StorageTek QFS software functionality is largely the same for the Solaris and Linux clients. Sun StorageTek QFS software components are stored in the `/opt/SUNWsamfs`, `/etc/opt/SUNWsamfs`, and `/var/opt/SUNWsamfs` directories on the Linux client, just as they are on a Solaris system. The `sam-fsd` and `sam-sharefsd` commands function the same way on both platforms.

Installing and Uninstalling the Linux Client

This section provides instructions for installing and uninstalling the Sun StorageTek QFS Linux Client software.

The following are the differences in installation between the Sun StorageTek QFS Solaris client and the Sun StorageTek QFS Linux client:

- The Linux software comes in two folders on the Sun StorageTek QFS 4.6 DVD. The installation files can be found in the `linux1` directory on that DVD (or in the downloaded software package).
- The Linux installation packages are in the form of RPMs (Red Hat Package Manager) and SRPMs (Source RPM). RPMs typically contain binary commands and loadable modules, whereas SRPMs contain source code RPMs.
- Because so many variations of the Linux kernel exist, the installation program also includes custom kernel capability, which enables the software to work with as many kernel variations as possible.

▼ To Install the Linux Client

1. **Edit the Solaris metadata server's `/etc/opt/SUNWsamfs/hosts.fsname` file to add the Linux client's address, and run `samsharefs` to update the information on disk.**

In the following example, `hiball-mn` is the Solaris metadata server:

```
hiball-mn      10.1.170.213  1 0 server
linux-mn       10.1.170.210   0 0
rollerball-mn 10.1.170.132    0 0
sandiego-mn   10.1.170.8           0 0
```

2. **Insert the installation DVD into the Linux system or locate the downloaded software package.**
3. **Do one of the following to install the software:**
 - a. **If you installing from the DVD, type the following commands as `root` on the Linux system:**

```
# mount -o ro -t iso9660 /dev/cdrom /mnt
# /mnt/linux1/Install
```

- b. **If you are installing from a downloaded software package, type the following commands as `root` on the Linux system:**

```
# mount -o ro,loop -t iso9660 StorageTek_QFS_4.6.iso /mnt
# /mnt/linux1/Install
```

The installation program installs the software.

If the installation program does not recognize the Linux kernel version, it will display the following message:

```
A direct match for your kernel wasn't found. Attempt creating a
custom rpm for your kernel (yes/no)?
```

Select `yes` to enable the installation program to adapt the Sun StorageTek QFS loadable binary kernel module to the Linux system. If you select `no`, the installation cannot be completed.

4. **Review the `/etc/opt/SUNWsamfs/mcf` file to verify that it contains the correct paths for your installation.**

Note – An `mcf` file is automatically generated for the Linux client. For more information, see [“mcf File Differences” on page ix](#).

For most installations, the `mcf` file that is created will be correct, but it may have to be edited to reflect your specific environment.

5. **Verify that the Solaris metadata server is running, and then run the `mount` command to mount the file system.**

Note – The `mount` command requires that the `shared` mount option be set, either on the command line or in `/etc/fstab`.

When you mount the file system, the following message is displayed:

```
Warning: loading SUNWqfs will taint the kernel: SMI license
See http://www.tux.org/lkml/#export-tainted for information
about tainted modules. Module SUNWqfs loaded with warnings
This is a normal result of the installation and can be disregarded.
```

▼ To Configure Access to man Pages

Sun StorageTek QFS man pages are located in section 1m. In order to ensure access to the man pages from the Linux client, carry out the following procedures:

- On Red Hat Linux clients, add 1m to `MANSECT` in the `/etc/man.config` file.
- On SuSE clients, add 1m to `SECTION` in the `/etc/manpath.config` file.

▼ To Uninstall the Linux Client

Use the `uninstall` script that resides in the `/var/opt/SUNWsamfs` directory to uninstall the Linux client software.

Caution – Do not use other processes, such as `rpm -e`, to uninstall the software. They can cause unexpected results and problems with uninstalling or reinstalling the software.

Linux Differences

The following sections describe the functionality, mounting, and execution differences between the Linux and Solaris clients.

Differences in Function

There are a few areas in which the Linux client does not match the functionality of the Solaris client. The Sun StorageTek QFS Linux client does not support the following functions:

- Remote file sharing
- `samaio`
- Quotas
- Access control lists (ACLs)
- Buffer-cache readahead
- Browser interface (BI) support
- Forced unmount
- IPv6

The Linux Shared QFS client supports striping, but striped performance is generally limited to the performance of a single device.

`mcf` File Differences

The `mcf` file, `/etc/opt/SUNWsamfs/mcf`, defines the topology of the equipment that the Sun StorageTek QFS file system manages.

Unlike a Solaris client, the Linux client software automatically generates an `mcf` file. If an `mcf` file does not exist, the Linux client will create one when the system is booted or when `samd config` is run. This file should be examined to verify that it correctly reflects the Sun StorageTek QFS environment. It may need to be manually edited to match your specific file system configuration. Note, however, that if you create your own `mcf` file or edit the auto-generated `mcf` file, the system will no longer auto-generate the file.

When an existing disk is repartitioned or a new file system is created, run `samd config` or reboot the system to create a new `mcf` file that reflects the changes. When new hardware is added, the SCSI bus needs to be rescanned with the “add single device” mechanism, or the equivalent, and then `samd config` must be run. Rebooting the system will also cause the SCSI bus to be rescanned.

If you need to modify the Linux client `mcf` file, run the `samfsconfig` command first. This prints the existing Sun StorageTek QFS file system names and the Linux device path names that are associated with them. When editing the `mcf` file, you must take into account the differences between Solaris and Linux logical unit numbers (LUNs) in the device paths.

Instead of specifying an actual device path for the metadata device (mm), specify nodev in the Linux Sun StorageTek QFS mcf file. This is the default in the auto-generated mcf file.

```
belmont      40  ma  belmont  on  shared
nodev        43  mm  belmont  on
/dev/sdb5    44  mr  belmont  on
```

For more information about creating mcf files, see the mcf(4) man page or the *Sun StorageTek QFS File System Configuration and Administration Guide*.

Mounting Differences

Mounting functions are almost identical on the Linux and Solaris clients. However, the mount options available for the Linux client are limited to the following:

- rw, ro
- retry
- shared
- meta_timeo
- rdlease, wrlease, aplease
- minallopsz, maxallopsz
- min_pool (ignored outside of samfs.cmd)
- noauto, auto (ignored outside of /etc/fstab)

No other mount options are available for the Linux client.

Execution Differences

The Linux application programming interface (API) differs from the Solaris API in many respects, so users will see some differences. For example, upon lookup of a directory that has been removed, the Solaris system returns an ENOENT message, whereas the Linux system returns an ESTALE message.

There are several host bus adapter (HBA) failover methods that are supported in the various Linux distributions, and failover is included in QLogic drivers. None of these methods have been exhaustively tested with the Sun StorageTek QFS Linux client.

Caution – Do not use the `mdadm` (multiple devices admin) package for path failover on a Sun StorageTek QFS Linux client. The `mdadm` package writes a superblock to devices that it uses. The result is that `mdadm` has the potential to corrupt data that Solaris has written to those devices. Furthermore, Solaris can also corrupt the superblock that `mdadm` has written to the devices.

Performance Differences

Performance of the Sun StorageTek QFS Linux client is affected by the fact that direct I/O is more restrictive on the Linux client. The Linux device layer supports a maximum of 4 kilobytes per I/O operation. This means that large requests in a Linux environment must be broken up and serviced individually.

Linux Kernel Patch

Solaris 10 added a new partition ID for storage management initiative (SMI) labels on x86 and 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 StorageTek QFS Linux client interaction with a Sun StorageTek QFS metadata server running Solaris 10 on an x64 platform.

Two patches are included with the Sun StorageTek QFS Linux client software: one for 2.4.x series kernels, and one for 2.6.x series kernels (only required for kernels before 2.6.10). They are located in the `linux1/patches` directory. To install the patch, 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>/2.4_patch
```

SELinux

Unexpected results may occur when you run the Sun StorageTek QFS Linux client software on a system with SELinux enabled, depending on your site's specific SELinux configuration. For example, the `df` command may return `Permission denied` when run by a non-root user.

If you encounter such a similar error, check the security context settings on the `/dev/samsys` file. The correct settings should be similar to the following:

```
crw-r--r--  1 user_u:object_r:device_t root root 253, 0 Jan 31
17:19 /dev/samsys
```

You might also want to check the security label of the user and verify that the user has a sufficient security label to access the object, given the object's security label.

Troubleshooting

Linux clients and Solaris clients use different procedures to store system information that might be used to diagnose Sun StorageTek QFS issues.

Files that contain system information from the Linux kernel are in the `/proc` directory. For example, the `/proc/cpuinfo` file contains hardware information. The table below describes some files that contain useful troubleshooting information.

TABLE 1

File Name	Information Provided
<code>version</code>	Running kernel version
<code>cpuinfo</code>	Hardware information
<code>uptime</code>	Time in seconds since boot time, and total time used by processes
<code>modules</code>	Information about the modules that are loaded
<code>cmdline</code>	Command-line parameters that are passed to the kernel at boot time
<code>filesystems</code>	Existing file system implementations
<code>scsi/scsi</code>	Attached SCSI devices
<code>fs/samfs/<QFS file system>/fsid</code>	File system ID, which must be included in the share options for network file system (NFS)

Note – Linux kernel log messages go to the `/var/log/messages` file.

Troubleshooting Tools

Because the Linux kernel has many variations, troubleshooting problems can be very challenging. A few tools are available that might help in debugging:

- Projects such as `lkcd` and `kgdb` can provide kernel crash dump information.
- Projects such as `kdb`, `kgdb`, and `icebox` are kernel debuggers.

Note – These projects are not present by default in Red Hat or SuSE Linux. You must obtain the appropriate RPMs or SRPMs and might have to reconfigure the kernel to use them.

- The `strace` command traces system calls and signals. It is similar to the Solaris `truss` command.
- The Sun StorageTek QFS `samtrace` command dumps the trace buffer.
- The Sun StorageTek QFS `samexplorer` command generates a Sun StorageTek QFS diagnostic report script.

Note – Trace files are placed in the `/var/opt/SUNWsamfs/trace` directory on the Linux client, just as they are on the Solaris client.

Frequently Asked Questions

The following are some questions about the Linux client that are frequently asked by users who are familiar with Sun StorageTek QFS on the Solaris platform.

Q: The Linux install script reports that I got a negative score and cannot install the software. Is there any way I can still install the software?

A: You can try the `-force-custom` and `-force-build` install options. However, this may cause a system panic when installing the modules. This is especially risky if your kernel is built with some of the kernel hacking options, such as `spinlock debugging`, enabled.

Q: Can I use commands such as `vmstat`, `iostat`, `top`, `truss`, and `sar` on Linux?

A: The `vmstat`, `top`, `sar`, and `iostat` commands are found in many Linux Red Hat installations. If they are not installed, they can be added using the `sysstat` and `procps` RPMs. The Linux equivalents of `truss` are `ltrace` and `strace`.

Q: How can multipath failover be done with QFS on Linux?

A: If you are using the Sun StorageTek 6130, you can use the multipath tool found at <http://www.sun.com/download/products.xml?id=432f43a5>. Otherwise, several path failover methods are possible, depending on your Linux distribution and type of HBAs. These include but are not limited to Sun StorageTek Traffic Manager (SSTM), multipulse, device mapper, and `qlogic` path failover. None of these have been tested exhaustively with Sun StorageTek QFS.

Do not use the `mdadm` (multiple devices admin) package for path failover on a Sun StorageTek QFS Linux client. The `mdadm` package writes a superblock to devices that it uses. The result is that `mdadm` has the potential to corrupt data that Solaris has written to those devices. Furthermore, Solaris can also corrupt the superblock that `mdadm` has written to the devices.

Q: Can Sun StorageTek™ Traffic Manager be used with the Sun StorageTek QFS Linux client?

A: Yes. First build a custom kernel with multipathing support as described in the Sun StorageTek Traffic Manager documentation. Then install the Linux client software.

Q: Can Extensible Firmware Interface (EFI) labels be used on the Sun StorageTek QFS Linux client?

A: Most Linux Red Hat kernels are not built with support for EFI labels with GPT (GUID Partition Table) partitions. Therefore, to use EFI labels, you must rebuild the kernel with the `CONFIG_EFI_PARTITION` option set. For more information about building a custom kernel, see the distribution documentation.

Q: Can I use other Linux volume managers such as logical volume management (LVM), Enterprise Volume Management System (EVMS), or Device Mapper with the Sun StorageTek QFS Linux client software?

A: No. To use a file system with EVMS, you need to have a File System Interface Module (FSIM) for that file system. No FSIM exists for the Sun StorageTek QFS product. For you to use LVM, the partition type that `fdisk` shows must be LVM(8e). Partitions that Sun StorageTek QFS uses must be SunOS.

Q: Can I use file systems that are larger than 2 TB?

A: Yes, but some utilities that provide file system information, such as `df`, might return incorrect information when run on Linux. In addition, there may be problems when sharing the file system with NFS or Samba.

Q: Are there any differences between the mount options supported on the Linux client and those supported on the Solaris client?

A: There are many `samfs` mount options that are not supported on the Linux client. Two to be aware of are `nosuid` and `forcedirectio`. See [“Mounting Differences” on page x](#) for a complete list of supported mount options on the Linux client.

Q: How can I mount filesystems that have SMI labels generated by X64 Solaris 10 metadata server?

A: Apply one of the patches available from Sun for the 2.4 or 2.6 kernel.

Q: What Linux versions and distributions were used in testing Sun StorageTek QFS 4.6?

A: Sun StorageTek QFS was tested with the following Linux distributions:

- 2.4.21-32 on x86 and X64 - RH 3 Update 5
- 2.4.21-47 on x86 and X64 - RH 3 Update 8
- 2.6.16-8 on X64 - RH 4 Update 2
- SuSE 8 SP4
- 2.6.5-7.191 on X64 and Altix - SuSE 9 SP2
- 2.6.16.21-0.8 - SuSE 10 FCS

Q: The File System Manager user interface does not show Sun StorageTek QFS installed on my Linux clients. Why?

A: The File System Manager software does not support Linux clients.

Q: Is there a Linux version of `cfgadm` or `devfsadm` to rescan/reprobe the SCSI devices?

A: On Linux you must reboot the system in order to guarantee that the SCSI devices are rescanned. You can also try removing and reinstalling the HBA driver. For example:

```
rmmod qlaXXXX; modprobe qlaXXXX
```

If you are using a Qlogic driver that supports LUN hot add, you can try the following:

```
echo "scsi-qlascan" > /proc/scsi// (qlogic driver will re-scan)
```

Q: There's no forced unmount on the Linux client like there is on Solaris. What can I do when I encounter a busy filesystem when attempting to unmount?

A: First try `fuser -k`. If the busy filesystem message persists, use `lsof` to find any open files and their associated PIDS, then kill the processes.

