PART NUMBER
312520301

VERSION NUMBER
4.0

EDITION NUMBER
1

ASM

Application Storage Manager ™

INSTALLATION AND CONFIGURATION GUIDE (UNIX)

For ASM, ASM-QFS, and ASM/QFS-Standalone

PRODUCT TYPE
SOFTWARE

StorageTek®

Application Storage Manager™ (ASM)

ASM, ASM-QFS, and ASM/QFS-Standalone Installation and Configuration Guide

for UNIX

Version 4.0

First Edition

Part Number 312520301

Information contained in this publication is subject to change without notice. Comments concerning the contents of this publication should be directed to:
Global Learning Solutions
Storage Technology Corporation
One StorageTek Drive
Louisville, CO 80028-3256
USA

URL: http://sts.stortek.com/sts/tis

Export Destination Control Statement
These commodities, technology or software were exported from the United States in accordance with the Export Administration Regulations. Diversion contrary to U.S. law is prohibited.

First Edition (May 2003)
This edition contains 104 pages.

Contents

Preface

This manual, the *ASM, ASM-QFS, ASM/QFS-Standalone Storage and Archive Management Guide*, explains the installation and upgrade procedures for the ASM/QFS-Standalone, ASM, and ASM-QFS software products, release 4.0. The 4.0 release can be installed on Sun Solaris™ 7, 8, and 9 operating environment (OE) platforms.

This manual is written for system administrators responsible for configuring and maintaining ASM/QFS-Standalone, ASM, and ASM-QFS software. You, the system administrator, are assumed to be knowledgeable about Sun Solaris procedures, including creating accounts, performing system backups, and other basic system administrator tasks.

How This Book Is Organized

This manual contains the following chapters:

- Chapter is an overview.
- Chapter 2 contains system requirements.
- Chapter 3 explains the ASM/QFS-Standalone initial installation procedure.
- Chapter 4 explains the ASM/QFS-Standalone upgrade procedure.
- Chapter 5 explains the ASM and ASM-QFS initial installation procedure.
- Chapter 6 explains the ASM and ASM-QFS upgrade procedure.

The glossary defines terms used in this and other QFS-Standalone, ASM, and ASM-QFS documentation.

Related Documentation

This manual is part of a set of documents that describes the operations of the QFS-Standalone, ASM, and ASM-QFS software products. Table 1. shows the complete release 4.0 documentation set for these products.

Table 1. ASM 4.0 Documentation

Title	Part Number
ASM, ASM-QFS, and ASM/QFS-Standalone Storage and Archive Management Guide	312520101
ASM-Remote Administrator's Guide	312520201
ASM, ASM-QFS, and ASM/QFS-Standalone Installation and Configuration Guide	312502301
ASM, ASM-QFS, and ASM/QFS-Standalone File System Administrator's Guide	312502401
ASM, ASM-QFS, and ASM/QFS-Standalone Disaster Recovery Guide	312502501

Note that the *ASM-Remote Administrator's Guide* has not been updated for the 4.0 release. An updated version of this manual will be provided at a later date.

How to Obtain Documentation

All the ASM publications are available from the following

sources:

- Contact StorageTek Publication Sales and Service at 800-436-5554 or send a fax to 303-661-7367.
- Online (for viewing and printing), at the StorageTek Customer Resource Center (CRC) website at: www.support.storagetek.com. Click on Software and go to the ASM Software list.

Access to the CRC site requires a password. To obtain a password, call StorageTek Customer Support at 1-800-678-4430.

Support

The publication "Requesting Software Support" is included in your media package. Please consult this book for the most information on your ASM support options, as well as regional phone numbers and procedures.

Using UNIX Commands

This document does not contain information on basic UNIX® commands and procedures such as shutting down the system, booting the system, and configuring devices.

See one or more of the following for this information:

- *Solaris Handbook for Sun Peripherals*
- AnswerBook2™ online documentation for the Sun Solaris OE
- Other software documentation that you received with your system

Typographic Conventions

Table 2. lists the typographic conventions used in this manual.

Table 2. Typographical Conventions

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

Table 2. Typographical Conventions

Typeface or Symbol	Meaning	Examples
`{ arg | arg}`	In syntax, braces and pipes indicate that one of the arguments must be specified.	`sndradm -b {phost | shost}`
`\`	At the end of a command line, the backslash (\) indicates that the command continues on the next line.	`atm90 /dev/md/rdsk/d5 \` `/dev/md/rdsk/d1`

Shell Prompts

Table 3. shows the shell prompts that this manual uses.

Table 3. Shell Prompts

Shell	Prompt
C shell	*machine-name%*
C shell superuser	*machine-name#*
Bourne shell and Korn shell	`$`
Bourne shell and Korn shell superuser	#

Overview

1

Before installing an ASM/QFS-Standalone, ASM, or ASM-QFS software package, you must set up and configure the hardware to be used. The necessary hardware differs, depending on the software product, as follows:

- The ASM/QFS-Standalone environment typically consists of the following:
  - A server based on SPARC™ technology running the 7, 8, or 9 Sun Solaris operating environment (OE).
  - A disk cache consisting of RAID devices, JBOD devices, or both.
- The ASM and ASM-QFS environment typically consists of the following:
  - A server based on SPARC technology running the 7, 8, or 9 Sun Solaris OE.
  - A disk cache consisting of RAID devices, JBOD devices, or both.
  - One or more libraries that include one or more tape or magneto-optical drives and/or one or more standalone drives. This requirement does not apply if you plan to archive to disk in another file system.

This chapter describes the characteristics of the ASM/QFS-Standalone, ASM, and ASM-QFS release packages. It contains the following topics:

This manual does not describe the installation or configuration procedures you need to use in order to create an ASM/QFS-Standalone shared file system. Use the procedures in this installation and configuration guide to create your ASM/QFS-Standalone or ASM-QFS environment before attempting to configure an ASM/QFS-Standalone shared file system. The configuration procedures for an ASM/QFS-Standalone shared file system are described in the *ASM, ASM-QFS, and ASM/QFS-Standalone File System Administrator’s Guide*.

Release Package Contents

You can obtain the ASM/QFS-Standalone, ASM, and ASM-QFS software packages from the StorageTek Customer Resource Center (CRC) or on a CD-ROM. The CRC is located at the following URL:

`http://www.support.storagetek.com.`

Contact your StorageTek sales representative if you have questions on obtaining the software in one of these ways.

After the release, upgrade patches are also available from the following URL:

https://www.support.storagetek.com/

All packages are in Sun Solaris `pkgadd`(1M) format. These packages reflect the Sun Solaris version for the platform upon which you will be installing the ASM/QFS-Standalone, ASM, or ASM-QFS software. For information on the directories and files provided with the release package, see the `README` file provided with the software.

Table 1. shows the release packages that include the software products.

Table 1. Release Packages

Installed Package	Description
`SUNWqfs`	an ASM/QFS-Standalone file system software package. This package is not needed if you are installing ASM or ASM-QFS software.
`SUNWsamfs`	ASM and ASM-QFS software package. Includes the standard file system and the storage and archive management (SAM) software.
SUNWcqfs, SUNWcsamf	Chinese localized packages.
SUNWjqfs, SUNWjsamf	Japanese localized packages.
SUNWfqfs, SUNWfsamf	French localized packages.

The releases are identified using characters arranged in the following format:

major.*minor*.*bugfix*

Table 2. explains the release numbering scheme.

Table 2. Release Numbering

Release Level Component	Meaning
major	The release level of a major release.
minor	The release level of a minor feature release.
bugfix	The bug fix number. A number between 1 and 99

Examples:

- 4.0 is a major release with no minor release revisions and no bug fixes.
- 4.1 is a minor release.
- 4.0.1 is a bug fix release and is typically delivered in a patch. This number appears in the patch's `README` file.

Directories and Files Created

This section describes the directories and files associated with the ASM/QFS-Standalone, ASM, and ASM-QFS products. Additional information about the files in this section can be obtained from the man pages after they have been installed.

Directories Created

Table 3. lists the directories created when the ASM/QFS-Standalone, ASM, and ASM-QFS software packages are installed.

Table 3. Directories Created

Directory	Content	Used By
`/dev/samst`	Device driver special files.	ASM, ASM-QFS
`/var/opt/SUNWsamfs`	Device catalogs, the catalog trace file, and log files. Also archiver data directory and queue files.	ASM, ASM-QFS
`/etc/fs/samfs`	Commands specific to QFS-Standalone, ASM-FS, and ASM-QFS software.	ASM/ QFS-Standalone, ASM, ASM-QFS
`/etc/opt/SUNWsamfs`	Configuration and license files.	ASM/ QFS-Standalone, ASM, ASM-QFS
`/opt/SUNWsamfs/bin`	User command binaries.	ASM/ QFS-Standalone, ASM, ASM-QFS
`/opt/SUNWsamfs/client`	Files for RPC API client.	ASM/ QFS-Standalone, ASM, ASM-QFS
`/opt/SUNWsamfs/examples`	Various example configuration files.	ASM/ QFS-Standalone, ASM, ASM-QFS

Table 3. Directories Created (Continued)

Directory	Content	Used By
/opt/SUNWsamfs/jre	Java runtime environment for the GUI tools. This is a symbolic link.	ASM, ASM-QFS
/opt/SUNWsamfs/include	API include files.	ASM/ QFS-Standalone, ASM, ASM-QFS
/opt/SUNWsamfs/lib	Relocatable libraries.	QFS-Standalone, ASM, ASM-QFS
/opt/SUNWsamfs/man	man(1) pages.	ASM/ QFS-Standalone, ASM, ASM-QFS
/opt/SUNWsamfs/sbin	System administrator commands and daemon binaries.	ASM/ QFS-Standalone, ASM, ASM-QFS
/opt/SUNWsamfs/doc	Documentation repository for any informational files included in the release. The README file is included in this directory. It summarizes the current release's features. You can access the README file for this release at any time from one of the documentation websites described in this manual's preface. It is available as the *ASM and ASM-QFS README File*. After your software is installed, the content of the README file is located in http:// www.support.storagetek.com .	ASM/ QFS-Standalone, ASM, ASM-QFS

Files Created

Table 4. lists the files created when the ASM/QFS-Standalone, ASM, or ASM-QFS software is installed.

Table 4. Files Created

File	Description	Used By
`/etc/opt/SUNWsamfs/inquiry.conf`	Vendor and product identification strings for recognized SCSI devices.	ASM, ASM-QFS
`/kernel/drv/samst`	Driver for SCSI media changers, optical drives, and non-motion I/O for tape drives.	ASM, ASM-QFS
`/kernel/drv/samst.conf`	Configuration file for the `samst` driver.	ASM, ASM-QFS
`/kernel/fs/samfs`	Sun Solaris 32-bit file system module.	ASM/QFS-Standalone, ASM, ASM-QFS
`/kernel/fs/sparcv9/samfs`	Sun Solaris 64-bit file system module.	ASM/QFS-Standalone, ASM, ASM-QFS
`/kernel/sys/samsys`	System call module.	ASM/QFS-Standalone, ASM, ASM-QFS
`/kernel/sys/sparcv9/samsys`	Sun Solaris 32- and 64-bit system call module.	ASM/QFS-Standalone, ASM, ASM-QFS

The ASM/QFS-Standalone, ASM, and ASM-QFS file systems have dynamically loadable components that are stored in the `Sun Solaris /kernel` directory (see Table 4.). You can determine the modules that are loaded by using the `modinfo`(1M) command. Typically, the file system module is loaded with the kernel at boot time. Alternatively, the file system module can be loaded when the file system is first mounted after the software is installed.

Site Files

Table 5. lists the files you create that are used by ASM and ASM-QFS. Some of the files in this list are optional, and others are required. The following are the required files:

- `/etc/opt/SUNWsamfs/LICENSE.4.0`
- `/etc/opt/SUNWsamfs/mcf`

Table 5. Site Files

File	Description	Used By
`/etc/opt/SUNWsamfs/LICENSE.4.0`	License file. For more information, see the licensing information pertinent to your installation in “Step 6: Verifying the Software License” on page 21. This is a required file.	ASM/ QFS-Standalone, ASM, ASM-QFS
`/etc/opt/SUNWsamfs/archiver.cmd`	Archiver command file. For more information, see the `archiver.cmd`(4) man page or see the *ASM, ASM-QFS, and ASM/QFS-Standalone Storage and Archive Management Guide*.	ASM, ASM-QFS
`/etc/opt/SUNWsamfs/samfs.cmd`	File system mount parameter command file. For more information, see the `samfs.cmd`(4) man page or see the *ASM, ASM-QFS, and ASM/QFS-Standalone File System Administrator’s Guide*.	ASM/ QFS-Standalone, ASM, ASM-QFS

Table 5. Site Files (Continued)

File	Description	Used By
/etc/opt/SUNWsamfs/recycler.cmd	Recycler command file. For more information, see the recycler.cmd(4) man page or see the *ASM, ASM-QFS, and ASM/QFS-Standalone Storage and Archive Management Guide*.	ASM, ASM-QFS
/etc/opt/SUNWsamfs/releaser.cmd	Releaser command file. For more information, see the releaser.cmd(4) man page or see the *ASM, ASM-QFS, and ASM/QFS-Standalone Storage and Archive Management Guide*.	ASM, ASM-QFS
/etc/opt/SUNWsamfs/preview.cmd	Previewer command file. For more information, see the preview.cmd(4) man page or see the *ASM, ASM-QFS, and ASM/QFS-Standalone Storage and Archive Management Guide*.	ASM, ASM-QFS
/etc/opt/SUNWsamfs/defaults.conf	Miscellaneous default values. For more information, see the defaults.conf(4) man page.	ASM, ASM-QFS
/etc/opt/SUNWsamfs/mcf	Master configuration file. For more information, see the mcf(4) man page. This is a required file.	ASM/QFS-Standalone, ASM, ASM-QFS

Modified System Files

During installation, ASM/QFS-Standalone, ASM, and ASM-QFS software adds information to certain Sun Solaris system files. These system files are ASCII

text files. Sun Solaris uses these files to identify loadable kernel modules by number rather than by name.

Table 6. lists the system files that are modified during the installation of software packages

.

Table 6. Modified System Files

File	Description	Used By
/etc/ name_to_sysnum	System call information file. The line added is as follows:[1] samsys 181	ASM/ QFS-Standalone, ASM, ASM-QFS
/etc/ name_to_major	Maps driver to major number, as follows:[2] samst 63 samrd 64	ASM, ASM-QFS

1. The samsys system call number can vary depending on the system call numbers in use by the Sun Solaris software.
2. The samst and samrd major numbers can vary depending on the major numbers in use by the Sun Solaris software.

As part of the installation process, the existing /etc/name_to_sysnum file is backed up to /etc/name_to_sysnum.SUNWsamfs.

To Install Sun Solaris Patches

Certain Sun Solaris patches overwrite the /etc/name_to_sysnum file. The /etc/name_to_sysnum file identifies loadable kernel modules by number rather than by name. If a patch overwrites this file, the ASM/QFS-Standalone, ASM, and ASM-QFS system call number does not exist. A system panic can occur when you attempt to start ASM/QFS-Standalone, ASM, or ASM-QFS processes and mount ASM/QFS-Standalone, ASM, or ASM-QFS file systems if the system call number is not present in this file.

One indication of the problem is the presence of the following message in the /var/adm/messages file:

```
WARNING: system call missing from bind file
```

To avoid this message, you need to check and perhaps to fix your /etc/ name_to_sysnum file if you install any Sun Solaris patches after your ASM/ QFS-Standalone, ASM, or ASM-QFS software is installed.

Note: The 4.0 releases use only one system call and only one system call number. By default, this is declared as samsys 181 in the /etc/

name_to_sysnum file. The following procedure assumes that you are using the default system call number of 181. If you are using a different number, modify the steps in this procedure to accommodate your system.

1. Use the cp(1) command to save your current /etc/name_to_sysnum file to an alternate location.

   For example:

```
# cp /etc/name_to_sysnum /etc/name_to_sysnum.old
```

2. Install one or more patches.
3. Use vi(1) or another editor to open your /etc/name_to_sysnum file and search for lines containing the string samsys.

   a. If no lines containing samsys exist in this file, add the following line:

   samsys 181

   b. If lines containing the samsys exist, examine the system call number declared on the line.

   - The following lines, if contained in the /etc/name_to_sysnum file, can be deleted:

```
samsys 180
samsys64 181
```

   samsys64 is no longer used in the release 4.0 software. samsys 181 is the release 4.0 entry.

   - If a samsys or samsys64 line is present and does not declare 180 or 181 as its system call number, you need to examine the file to determine whether it should declare default system call numbers. If your file uses a nondefault system call number, ensure that a samsys system call line is present with a unique number.

4. Reboot your system.
5. Delete the backup file when all is determined to be running correctly.

For example:

```
# rm /etc/name_to_sysnum.old
```

System Requirements

2

This chapter outlines the system requirements that must be met prior to the installation of ASM/QFS-Standalone, ASM, and ASM-QFS software packages. These requirements are as follows:

The following sections describe these requirements in more detail.

Step 1: Verifying the Environment

ASM/QFS-Standalone, ASM, and ASM-QFS software packages run on many workstations and servers. Before installation, you should verify the applicability of the hardware and the level of the Sun Solaris operating environment (OE). In addition, you must ensure that you have root-level access to your system.

To Verify the Environment

1. Verify that your system has a CD-ROM drive or that it can access the release package at StorageTek’s Customer Resource Center.
2. Verify your system’s Sun Solaris OE.

   The software relies on a properly configured Sun Solaris 7, 8, or 9 OE. Check to see that your server is running one of these levels by entering the following command:

```
# uname -sr
SunOS 5.9
```

SunOS 5.*x*.*y* levels correspond to Sun Solaris OE levels. The preceding system is running the Sun Solaris 9 OE.

3. Log in to your system as root.

   You must have superuser access to install the software.

Step 2: Verifying Disk Cache

The ASM/QFS-Standalone, ASM, and ASM-QFS software packages require a certain amount of disk cache in order for them to create and manage data files and directories. For the ASM/QFS-Standalone and ASM-QFS software, at least two disk devices or partitions are required, one for file data and one for metadata. Multiple disk devices or partitions increase I/O performance. For the ASM software, at least one disk device or partition is required.

The disk devices or partitions do not require any special formatting, nor do they need to have a UNIX file system created on them. Make sure that the disks and partitions that you are using are not currently in use and do not contain any existing data because any existing data is lost when you create the ASM/QFS-Standalone, ASM, or ASM-QFS file system.

For ASM and ASM-QFS software, the disk must be connected to the server using a Fibre Channel or SCSI controller. Individual disk partitions can be specified for a disk, or the entire disk can be used as a disk cache. Disk arrays, including those under the control of volume management software, such as Solstice DiskSuite™ software and other volume management software products, are supported.

To Verify the Available Disk Cache Space

1. Estimate the disk cache requirements for ASM/QFS-Standalone software (`ma` file systems).
   - Disk cache = largest file (in bytes) + amount of space needed for working files
   - Metadata cache = ((number of files + number of directories) * 512) + 16384 * number of directories
2. Estimate the disk cache requirements for ASM software (`ms` file systems).
   - Disk cache = largest file (in bytes) + ((number of files + number of directories) * 512) + 4096 * number of directories + amount of space needed for working files
3. Estimate the disk cache requirements for ASM-QFS software.
   - Disk cache = largest file (in bytes) + amount of space needed for working files

- Metadata cache = ((number of files + number of directories) * 512) + 16384 * number of directories

4. Use the `format`(1M) command to verify that you have sufficient disk cache space.

   Remember that you must use `<ctrl>-d` to exit the `format`(1M) command.

Example

Figure 1. shows four disks attached to a server. There are two internal disks connected by means of controller `0` on targets `0` (`c0t0d0`) and `1` (`c0t1d0`). There are two external disks connected by means of controller `3` on targets `0` (`c3t0d0`) and `2` (`c3t2d0`).

Figure 1. `format`**(1M) Command Example**

```
# format
Searching for disks...done


AVAILABLE DISK SELECTIONS:
       0. c0t0d0 <SUN9.0G cyl 4924 alt 2 hd 27 sec 133>
          /sbus@1f,0/SUNW,fas@e,8800000/sd@0,0
       1. c0t1d0 <SUN2.1G cyl 2733 alt 2 hd 19 sec 80>
          /sbus@1f,0/SUNW,fas@e,8800000/sd@1,0
       2. c3t0d0 <SUN36G cyl 24620 alt 2 hd 27 sec 107>
          /sbus@1f,0/QLGC,isp@0,10000/sd@0,0
       3. c3t2d0 <SUN36G cyl 24620 alt 2 hd 27 sec 107>
          /sbus@1f,0/QLGC,isp@0,10000/sd@2,0
Specify disk (enter its number): 1
selecting c0t1d0
[disk formatted]
Warning: Current Disk has mounted partitions.


FORMAT MENU:
        disk       - select a disk
        type       - select (define) a disk type
        partition  - select (define) a partition table
        current    - describe the current disk
        format     - format and analyze the disk
        repair     - repair a defective sector
        label      - write label to the disk
        analyze    - surface analysis
        defect     - defect list management
        backup     - search for backup labels
        verify     - read and display labels
        save       - save new disk/partition definitions
        inquiry    - show vendor, product and revision
        volname    - set 8-character volume name
```

Figure 1. `format`**(1M) Command Example (Continued)**

```
        !<cmd>     - execute <cmd>, then return
        quit
format> par

PARTITION MENU:
        0      - change ‘0’ partition
        1      - change ‘1’ partition
        2      - change ‘2’ partition
        3      - change ‘3’ partition
        4      - change ‘4’ partition
        5      - change ‘5’ partition
        6      - change ‘6’ partition
        7      - change ‘7’ partition
        select - select a predefined table
        modify - modify a predefined partition table
        name   - name the current table
        print  - display the current table
        label  - write partition map and label to the disk
        !<cmd> - execute <cmd>, then return
        quit
partition> pri
Current partition table (original):
Total disk cylinders available: 2733 + 2 (reserved cylinders)

Part      Tag    Flag     Cylinders        Size            Blocks
  0        var    wm       0 - 2732        1.98GB    (2733/0/0)
4154160
  1 unassigned    wm       0               0         (0/0/0)         0
  2     backup    wm       0 - 2732        1.98GB    (2733/0/0)
4154160
  3 unassigned    wm       0               0         (0/0/0)         0
  4 unassigned    wm       0               0         (0/0/0)         0
  5 unassigned    wm       0               0         (0/0/0)         0
  6 unassigned    wm       0               0         (0/0/0)         0
  7 unassigned    wm       0               0         (0/0/0)         0

partition> q
```

Step 3: Verifying Disk Space

The software requires a certain amount of disk space in the `/` (root), `/opt`, and `/var` directories. The actual amount needed varies depending on the packages you install. Table 7. shows the minimum amount of disk space required in these various directories.

Table 7. Minimum Disk Space Requirements

Directory	ASM, ASM-QFS Disk Space Needed	ASM/QFS-Standalone Disk Space Needed
`/ (root)` directory	1,775 – 1,779 kilobytes	1,700 kilobytes
`/opt` directory	18,323 – 19,093 kilobytes	8,000 kilobytes
`/var` directory	577 – 597 kilobytes	170 kilobytes

Note that the archiver data directory, archiver queue files, and the log files are written to the `/var` directory, so the sizes shown in Table 7. should be considered a minimum amount for the `/var` directory.

Table 8. shows the disk space requirements for the individual installation packages.

Table 8. Software Package Requirements (Required Packages)

Package	Space on `/` (root)	Space on `/opt`	Space on `/var`
`SUNWqfs`	1,700 kilobytes	8,000 kilobytes	170 kilobytes
`SUNWsamfs`	1,775 kilobytes	18,323 kilobytes	577 kilobytes

The following optional packages have been localized for Chinese, French, and Japanese language users. Table 9. shows the disk space requirements for these packages.

Table 9. Software Package Requirements (Optional Packages)

Package	Space on `/opt`	Space on `/var`
`SUNWcqfs`, `SUNWcsamf`	8 kilobytes	90 kilobytes
`SUNWfqfs`, `SUNWfsamf`	7 kilobytes	140 kilobytes
`SUNWjqfs`, `SUNWjsamf`	10 kilobytes	140 kilobytes

To Verify Disk Space

The following procedure shows how to verify whether or not there is enough disk space on your system to accommodate the `SUNWqfs` or `SUNWsamfs` packages. Note that these figures do not include the space needed for the optional localized packages described in Table 9..

1. Issue the `df`(1M) command as shown in Figure 2..

Figure 2. Using the `df`(1M) Command to Verify Disk Space

```
# df -k /
Filesystem        kbytes  used    avail capacity   Mounted on
/dev/dsk/c0t1dos0 76767   19826   49271    29%     /
# df -k /opt
Filesystem        kbytes  used    avail capacity   Mounted on
/dev/dsk/c0t1dos4 192423  59006  114177    35%     /opt
```

2. Verify that there are at least 2,000 kilobytes available in the `avail` column for the `/` directory.
3. Verify that there are at least 20,000 kilobytes in the `avail` column for the `/opt` directory.
4. Verify that there are at least 577 kilobytes available in the `/var` directory.

   A quantity of 30 megabytes or more is recommended to allow for the growth of log files and other system files.
5. If there is not enough room for the software under each directory, repartition the disk to make more space available to each file system.

   To repartition a disk, see your StorageTek ASM system administration documentation.

Step 4: Verifying Sun Solaris Patches

The latest patches for the Sun Solaris OE are required. An updated list of Sun Solaris patches required prior to installation is included in the *ASM/QFS-Standalone, ASM, and ASM-QFS README File*. After installation, the `README` is located in `/opt/SUNWsamfs/doc/README`.

To Verify Sun Solaris Patches

- Use the `showrev`(1M) command to determine which patches are installed on your system.

  For the sake of convenience, redirect the output to a file that can be examined easily. For example:

```
# showrev -p > outfile
```

If the required patches are not listed in the output from the `showrev`(1M) command, you must install them before installing any ASM/QFS-Standalone, ASM, or ASM-QFS release packages. Patches are provided to customers with a Sun Microsystems maintenance contract by means of CD-ROM and the Sun Microsystems web site (`http://www.sunsolve.com`).

To install a patch, load the CD-ROM or transfer the patch software to your system. Follow the instructions outlined in the *Patch Installation Instructions and Special Install Instructions* in the `README` file included in the patch or jumbo patch cluster.

Note: All patches should be installed before the ASM/QFS-Standalone, ASM, or ASM-QFS 4.0 software is installed. Certain Sun Solaris patches overwrite the `/etc/name_to_sysnum` file. If you install Solaris patches after you have installed the ASM/QFS-Standalone, ASM, or ASM-QFS 4.0 software, verify that your `/etc/name_to_sysnum` file is intact. The procedure for verifying your `/etc/name_to_sysnum` file is described in “Modified System Files” on page 111.

Step 5: Verifying Archive Media (Optional)

This section applies to ASM and ASM-QFS packages only.

The ASM and ASM-QFS environments support archiving to removable media devices or to disk space in another file system.

If you are archiving to disk space in another file system, also called *disk archiving*, verify that you have space available on a disk and that the disk is accessible from the server upon which the ASM or ASM-QFS software is installed. The archive copies can be written to any Sun Solaris file system, but they do not need to be written to a disk partition that is part of an ASM or ASM-QFS file system.

If you are archiving to removable media devices, your environment should include at least one such device for archiving files. This device can be a single tape or optical drive or it can be multiple devices such as the drives within an automated library.

Note: One or more logical library catalogs are associated with each physical library. The library catalog tracks the movement of removable media devices within the library. If you are using tape devices, be aware that it is not possible to mix different types of tape media within a single library catalog. All media recorded in one library catalog must have the same media type code. Most libraries attached to a host system through a direct (SCSI) connection have one library catalog, and all drives in the library must be of the same media type. A library that is logically partitioned into two or more libraries has one catalog for each logical library, and the catalogs for each of these logical libraries can contain information for only one media type. For more information on media type codes and the types of media included in each code, see the `mcf`(4) man page.

The ASM and ASM-QFS environments support a wide variety of removable media devices. A list of currently supported drives and libraries is available from StorageTek.

The device that you intend to use must be attached and recognized by the server. If the removable media device is already connected and communicating with the server, go to “Step 6: Verifying the Software License” on page 21.

Note: Instructions for attaching removable media devices to a server are presented in this section. These are general guidelines for attaching removable media hardware to a server. For explicit instructions on connecting these peripherals to a server, refer to the hardware installation guide supplied by the vendor with the automated library and drives.

To Attach Removable Media Devices

1. Power off the server before connecting devices.

   Typically, central components are powered off first and then the peripheral equipment. So, use the `init`(1M) command to power off the server, as follows:

```
# init 5
```

This command brings down the system to the PROM level. At this point it is safe to power off the server and peripherals. For specific instructions regarding your equipment, see the documentation from the hardware vendor for proper power-on and power-off sequences.

2. Ensure that the removable media devices and the disk(s) to be used for the ASM and ASM-QFS file system are connected and properly addressed.

3. Ensure that the SCSI target IDs are unique for each SCSI initiator (host adapter). (Optional)

   Perform this step if you have libraries attached to the host system through a SCSI interface.

   Avoid setting SCSI target IDs for peripherals to ID `7` because this ID is typically reserved for the initiator. For example, if you are using a SCSI host adapter with a previously attached disk drive set to use a target ID of `3`, any additional peripheral connected to this bus must not have an ID of `3`. Typically, the internal disk drive ID is `3` for SPARC systems and `0` for UltraSPARC™ systems.

4. Power on the peripherals and server according to the manufacturer’s recommended sequence.

   Typically, outermost peripherals are powered on first, working toward more central components in sequence.

5. Disable autobooting.

At the > ok prompt, disable autobooting by entering the following command:

```
>ok setenv auto-boot? false
```

6. Type reset at the next prompt.

   For example:

```
>ok reset
```

   Reenabling autobooting is described later in this procedure.

7. Conduct an inventory of target IDs and LUNs for each device connected to the host system through a SCSI interface. (Optional)

   Perform this step if you have libraries attached to the host system through a SCSI interface.

   Enter the following at the PROM >ok prompt:

```
{0} ok probe-scsi-all
/pci@1f,4000/scsi@3,1
Target 3
  Unit 0   Removable Device type 7     HP      C1716T
3404
Target 4
  Unit 0   Removable Device type 7     HP      C1716T
3404
Target 5
  Unit 0   Removable Device type 8     HP      C1710T
6.16

/pci@1f,4000/scsi@3
Target 0
  Unit 0   Disk     SEAGATE ST318404LSUN18G 4207
Target 6
  Unit 0   Removable Read Only device   TOSHIBA
XM6201TASUN32XCD1103
```

In some instances, SCSI-attached devices might use a target number greater than 6 or a Logical Unit Number (LUN) greater than 0. If this is the case with your system, you must edit the /kernel/drv/samst.conf file when installed.

Save this output. The information in this output is needed when configuring SCSI-attached libraries.

8. Conduct an inventory of target IDs and LUNs for each device connected to the host system through a Fibre Channel interface. (Optional)

Perform this step if you have libraries or tape drives attached to the host system through a Fibre Channel interface.Enter the following commands to locate the host adapter directory, select an item, and display the Fibre Channel host bus adapter (HBA) devices. The following is an example session:

:

```
{0} ok show-devs
/SUNW,ffb@1e,0
/SUNW,UltraSPARC-II@2,0
/SUNW,UltraSPARC-II@0,0
/counter-timer@1f,1c00
/pci@1f,2000
/pci@1f,4000
/virtual-memory
/memory@0,a0000000
/aliases
/options
/openprom
/chosen
/packages
/pci@1f,2000/SUNW,qlc@1
/pci@1f,2000/SUNW,qlc@1/fp@0,0
/pci@1f,2000/SUNW,qlc@1/fp@0,0/disk
/pci@1f,4000/SUNW,ifp@2
/pci@1f,4000/scsi@3,1
/pci@1f,4000/scsi@3
/pci@1f,4000/network@1,1
/pci@1f,4000/ebus@1
/pci@1f,4000/SUNW,ifp@2/ses
{0} ok select /pci@1f,2000/SUNW,qlc@1
{0} ok show-children
LiD HA LUN  --- Port WWN ---  ----- Disk description -----
 2  7e   0  500104f00041182b  STK     L700            0236
7c  7e   0  500104f00043abfc  STK     9840            1.28
7d  7e   0  500104f00045eeaf  STK     9840            1.28
6f  7e   0  5005076300416304  IBM     ULT3580-TD1     16E0
6e  7e   0  5005076300416303  IBM     ULT3580-TD1     16E0
```

If the server does not acknowledge all the devices (disk drives, tape or optical drives, the automated library, and so on), you should check the cabling. Cabling is often the problem when devices and controllers are not communicating. Do not proceed until all devices appear when probed.

Save this output. The information in this output is needed when configuring libraries attached through a Fibre Channel interface.

Note: “Step 4: Verifying Sun Solaris Patches” on page 24 explains the importance of installing all recommended patches before installing the ASM/QFS-Standalone, ASM, or ASM-QFS software. If you have

devices connected to the host system through a Fibre Channel interface, it is particularly important that you install the following Fibre Channel driver patches at the stated patch level or later:

- 111095-05 fctl/fp/fcp/usoc driver patch
- 111096-03 fcip driver patch
- 111097-04 qlc driver patch

In addition, you need to install the SAN Foundation Kit package (SUNWsan Package 1.0) and patch 111847-01 or later.

For more information, see the “ASM and ASM-QFS Initial Installation Procedure” on page 57.

9. Reenable autobooting.

At the `>ok` prompt, enable autobooting by entering the following:

```
>ok setenv auto-boot? true
```

10. Boot the system.

```
>ok boot
```

11. Review system files.

Review the following files:

- `/var/adm/messages` to ensure that all devices were recognized.
- `/dev/rmt` for expected tape devices.
- `/dev/dsk` and `/dev/rdsk` for expected disks.

Due to special driver requirements, no device information appears in `/var/adm/messages` for magneto-optical devices or libraries until after the ASM or ASM-QFS software package is installed.

Step 6: Verifying the Software License

If you do not have a StorageTek license key for the release level that you are installing, contact StorageTek. When you contact StorageTek for a license, you will be asked to provide information regarding your environment.

For ASM/QFS-Standalone software, you will need to provide information such as the following:

- Your StorageTek sales order number.

- The host IDs of the system(s) upon which you will be installing the ASM/QFS-Standalone software.
- The server class of each host system to be used in the ASM/QFS-Standalone environment.

For ASM or ASM-QFS, you will need to provide information such as the following:

- Your StorageTek sales order number.
- The host IDs of the system(s) upon which you will be installing the ASM/QFS-Standalone, ASM, or ASM-QFS software.
- The types of libraries to be used in the ASM or ASM-QFS environment.
- The media types of the drives to be used in the ASM or ASM-QFS libraries.
- The total number of slots to be available to the ASM or ASM-QFS software.

After your initial installation, if you upgrade your software or if you change your environment's configuration, you might need to change your software license. Changes to the environment that might necessitate upgrading your license include adding a library or changing a host system. If you have questions regarding your existing license, you can obtain license information by accessing the `samu`(1M) utility's `l` (for *license*) display. If you need to upgrade your license, contact your StorageTek sales representative.

The license keys for the ASM/QFS-Standalone, ASM, and ASM-QFS packages allow the system to run indefinitely unless one of the following conditions is present:

- You were issued a temporary license. When a temporary license expires, the system is no longer able to load and unload cartridges, or to archive, stage, or release files.
- You are using ASM or ASM-QFS software and you have exceeded the number of slots allowed for the license. In this case, you cannot import or label cartridges. Access continues unaffected for files already on a disk.

If your license expires, you can mount the file systems, but you cannot archive or stage files in a ASM or ASM-QFS environment.

Step 7: Verifying Compatibilities

The ASM/QFS-Standalone, ASM, and ASM-QFS software interoperates with many different hardware and software products from third-party vendors. Depending on your environment, you might need to upgrade other software or firmware before installing or upgrading the ASM/QFS-Standalone, ASM, or ASM-QFS package. Consult the README file distributed with this release for

information pertaining to library model numbers, firmware levels, and other compatability information.

Step 8: Verifying the Java Runtime Environment (Optional)

If you want to use the optional GUI tools that are included with the ASM and ASM-QFS software packages, make sure that you have the Java™ Runtime Environment (JRE) installed. If you do not have JRE software, you must download and install the JRE software yourself. Version 1.2.2 is recommended for use with the ASM and ASM-QFS 4.0 software. You can download the JRE from `java.sun.com`.

ASM/QFS-Standalone Initial Installation Procedure

3

This chapter describes the procedure for installing and configuring ASM/QFS-Standalone software for the first time. Use this procedure if this is the initial installation of the ASM/QFS-Standalone software package at your site. If you are upgrading ASM/QFS-Standalone software on an existing server, see the .

The procedure in this chapter describes obtaining the packages, installing the software packages on your server, and configuring the software to match the hardware at your site. The main sections are as follows:

In addition, this chapter describes how to initialize your ASM/QFS-Standalone file system and how to check the status of your system. For most of the procedures in this chapter, you must have superuser (`root`) access.

Step 1: Obtaining the Release Files

You can obtain the ASM/QFS-Standalone software package from the StorageTek Customer Resource Center (CRC) or on a CD-ROM. The CRC is located at the following URL:

`http://www.support.storagetek.com`. Contact your StorageTek sales representative if you have questions on obtaining the software in one of these ways.

After the release, upgrade patches are also available from the following URL:

`http://www.support.storagetek.com`

Caution: If you have not read the `README` file delivered with this release, please do so before continuing. You can access the `README` file for this release at any time from one of the documentation websites described in this manual's preface. It is available as the *ASM and ASM-QFS README File*. After your software is installed, the content of the `README` file is located in `/opt/SUNWsamfs/doc/README`.

To Install from a CD-ROM

1. Log in as root.

The ASM/QFS-Standalone file system uses the Sun Solaris operating environment (OE) packaging utilities for adding and removing software. You must be logged in as superuser (`root`) to make changes to software packages. The `pkgadd`(1M) utility prompts you to confirm various actions necessary to install the ASM/QFS-Standalone package.

2. Insert the CD into the CD drive.

   The system should automatically detect the CD's presence. If it does not, issue commands to stop and start the ASM Volume Manager and to change to the directory that contains the ASM/QFS-Standalone software package.

For example:

```
# /etc/init.d/volmgt stop
# /etc/init.d/volmgt start
# volcheck
# cd /cdrom/cdrom0
```

On the CD, the ASM/QFS-Standalone package resides in the `/cdrom/cdrom0` directory organized by Sun Solaris version.

Step 2: Adding the Packages

1. Use the `pkgadd`(1M) command to add the `SUNWqfs` package.

   For example:

```
# pkgadd -d SUNWqfs
```

2. Use the `pkgadd`(1M) command to add one or more localized packages. (Optional)

   Perform this step only if you want to install the packages localized for Chinese, French, or Japanese. To install these packages, enter one or more of the following commands:

```
# pkgadd -d SUNWcqfs
# pkgadd -d SUNWfqfs
# pkgadd -d SUNWjqfs
```

3. Reboot the server.

   For example:

```
# reboot
```

If you have added new equipment that should be recognized by the host system, use the following `reboot`(1M) command:

```
# reboot -- -r
```

Changes associated with adding the ASM/QFS-Standalone system call number to the `/etc/name_to_sysnum` file are enabled at this time.

Note: Failure to reboot the system at this time prevents the file system from being mounted.

Step 3: Setting Up PATH and MANPATH Variables

To Set Up PATH and MANPATH Variables

1. For users who need to access the ASM/QFS-Standalone user commands (for example, `sls`(1)), add `/opt/SUNWsamfs/bin` to the users' `PATH` variables.
2. For users who need to access the ASM/QFS-Standalone man pages, add `/opt/SUNWsamfs/man` to the `MANPATH` variable.

3. For users, such as superusers, who need to access the administrator commands, add `/opt/SUNWsamfs/sbin` to the `PATH` variable.
4. In the Bourne or Korn shells, edit the `.profile` file, change the `PATH` and `MANPATH` variables, and export the variables.

   For example:

```
PATH=$PATH:/opt/SUNWsamfs/bin:/opt/SUNWsamfs/sbin
MANPATH=$MANPATH:/opt/SUNWsamfs/man
export PATH MANPATH
```

5. In the C shell, edit the `.login` and `.cshrc` files.

   For example, the `path` statement in your `.cshrc` file might look like the following:

```
set path = ($path /opt/SUNWsamfs/bin /opt/SUNWsamfs/sbin)
```

   For example, the `MANPATH` statement in your `.login` file might look like the following:

```
setenv MANPATH /usr/local/man:opt/SUNWspro:/$OPENWINHOME/\
share/man:/usr/share/man:/opt/SUNWsamfs/man
```

Step 4: Licensing the ASM/QFS-Standalone Software

A license key is required to run the ASM/QFS-Standalone software. For information on license keys, see “Step 6: Verifying the Software License” on page 21.

The ASM/QFS-Standalone file system uses an encrypted license key. The license key consists of an encoded alphanumeric string.

To License the ASM/QFS-Standalone Software

1. Create the `/etc/opt/SUNWsamfs/LICENSE.4.0` file.
2. Starting in column one, place the license key you have obtained from StorageTek on the first line in the `/etc/opt/SUNWsamfs/LICENSE.4.0` file.

   The key must start in column one. No other keywords, host IDs, comments, or other information can appear in the `/etc/opt/SUNWsamfs/LICENSE.4.0` file. The license becomes effective when the ASM/QFS-Standalone file system is mounted.

Step 5: Configuring the ASM/QFS-Standalone File System

Each ASM/QFS-Standalone environment is unique. The system requirements and hardware used differ from site to site. It is up to you, the system administrator at your site, to set up the specific configuration for your ASM/QFS-Standalone environment.

The topology of the equipment managed by the ASM/QFS-Standalone file system is defined in the master configuration file, `/etc/opt/SUNWsamfs/mcf`. This file specifies the devices and file systems included in the environment. Each piece of equipment is assigned a unique equipment identifier in the `mcf` file.

To configure ASM/QFS-Standalone devices, create an `mcf` file in `/etc/opt/SUNWsamfs/mcf` that contains a line for each device and family set in your configuration. The `mcf` contains information that enables you to identify the disk slices to be used and to organize them into ASM/QFS-Standalone file systems.

There are examples of `mcf` files in `/opt/SUNWsamfs/examples`.

Note: For information on file system design considerations, see the *ASM, ASM-QFS, and ASM/QFS-Standalone File System Administrator's Guide*.

To Create the mcf File

When you create the `mcf` file, follow these guidelines:

- Delimit the fields in each line with spaces or tabs.
- Begin each comment line entered into this file with a pound sign (#).
- Use a dash (-) to indicate optional fields that are omitted.

The following shows the fields of each line entry in the `mcf` file.

```
#
# QFS-Standalone file system configuration
#
# Equipment       Equip  Equip Fam   Dev    Additional
# Identifier      Ord    Type  Set   State  Parameters
# ----------      -----  ----- ----  -----  ----------
```

Table 10 shows the information to be contained in each field and whether or not the field is a required or optional field.

Table 10. ASM/QFS-Standalone `mcf` File Fields

Field	Description
Equipment Identifier	*Required*. This field is either the name of the file system or a `/dev/dsk` entry for a disk partition or disk slice.
Equipment Ordinal	*Required*. Enter a unique integer from 1 to 65535.
Equipment Type	*Required*. Enter the code for the Equipment Type, as follows: • The `ma` Equipment Type defines a file system in an ASM/QFS-Standalone file system. • The `mm` Equipment Type defines a metadata device. • The `mr` and `md` Equipment Type defines a round-robin or striped data device. • The `g`*XXX* Equipment Type defines a striped group data device. Striped groups start with the letter `g`, followed by a 1-, 2-, or 3-digit integer. For example, `g2` or `g14` are both valid values for a striped group. For more information on Equipment Types, see the `mcf`(4) man page.
Family Set	*Required*. The Family Set organizes all devices with the same family set name together as an ASM/QFS-Standalone file system.
Device State	*Optional*. If specified, this field should contain either the keyword `on` or a dash character (`-`). Enter a state for the device for when the ASM/QFS-Standalone file system is initialized.
Additional Parameters	*Optional*. This field is used only if you are running the ASM/QFS-Standalone shared file system. The ASM/QFS-Standalone file system must be configured and tested on a host system prior to configuring the ASM/QFS-Standalone shared file system. For information on configuring the ASM/QFS-Standalone shared file system, see the *ASM, ASM-QFS, and ASM/QFS-Standalone File System Administrator's Guide*.

For more information on this file, see the `mcf`(4) man page. There is an example of an `mcf` file located in `/opt/SUNWsamfs/examples/mcf`.

Caution: Make sure you specify disk partitions that are not in use on your system. Do not use overlapping partitions. If an ASM/QFS-Standalone file system attempts to use a partition that is already in

use, the software issues a message to indicate that the device is busy.

If you give the wrong partition names, you risk damaging user or system data. This is true when creating any type of file system. The risk is greatest if the partition named contains a UFS file system that is not mounted currently.

The following example shows file system entries in an `mcf` file:

```
#
# QFS-Standalone file system configuration
#
# Equipment       Equip  Equip Fam   Dev    Additional
# Identifier      Ord    Type  Set   State  Parameters
# ----------      -----  ----- ----  -----  ----------
qfs1               1     ma    qfs1
/dev/dsk/c1t0d0s0 11     mm    qfs1  on
/dev/dsk/c1t1d0s4 12     mr    qfs1  on
/dev/dsk/c1t2d0s4 13     mr    qfs1  on
/dev/dsk/c1t3d0s4 14     mr    qfs1  on
```

To Reinitialize the mcf File

Reinitialization of the `mcf` file is not needed at this point in the configuration process. Be aware, however, that if you change the `mcf` file after the ASM/QFS-Standalone file system is in use, you must convey the new `mcf` specifications to the ASM/QFS-Standalone software. For information on reinitializing the `mcf` file, see the *ASM, ASM-QFS, and ASM/QFS-Standalone File System Administrator's Guide*.

Configuration Example 1

Figure 3. shows output from the Sun Solaris `format`(1M) command. It reports that the disks are partitioned as follows.

Figure 3. `format`**(1M) Command Output**

```
1. c1t0d0 <SEAGATE-ST15230W-0168 cyl 3974 alt 2 hd 19 sec
111>
/iommu@0,10000000/sbus@0,10001000/QLGC,isp@1,10000/sd@0,0
Current partition table (original) :
Total disk cylinders available: 3974 + 2 (reserved cylinders)
Part    Tag          Flag     Cylinders    Size       Blocks
0       unassigned   wm       0-3499       3.52GB     (3500/
0/0)
1       unassigned   wm       3500-3972    487.09MB   (473/0/
0)
```

Figure 3. `format`**(1M) Command Output (Continued)**

```
2       backup       wu       0-3973       4.00GB     (3974/
0/0)
3       unassigned   wm       0            0          (0/0/0)
4       unassigned   wm       0            0          (0/0/0)
5       unassigned   wm       0            0          (0/0/0)
6       unassigned   wm       0            0          (0/0/0)
7       unassigned   wm       0            0          (0/0/0)
2. c1t1d0 <SEAGATE-ST15230W-0168 cyl 3974 alt 2 hd 19 sec
111>
/iommu@0,10000000/sbus@0,10001000/QLGC,isp@1,10000/sd@1,0
Current partition table (original):
Total disk cylinders available: 3974 + 2 (reserved cylinders)
Part    Tag          Flag     Cylinders    Size       Blocks
0       unassigned   wm       1000-3973    2.99GB     (2974/
0/0)
1       unassigned   wu       0            0          (0/0/0)
2       backup       wu       0-3973       4.00GB     (3974/
0/0)
3       unassigned   wm       0            0          (0/0/0)
4       unassigned   wm       0            0          (0/0/0)
5       unassigned   wm       0-999        1.01GB     (1000/
0/0)
6       unassigned   wm       0            0          (0/0/0)
7       unassigned   wm       0            0          (0/0/0)
```

One file system (`qfs1`) is placed on slice `0` of disk `c1t0d0` and slice `0` of `c1t1d0`. Another file system (`qfs2`) is created on slice `1` of disk `c1t0d0` and slice `5` of disk `c1t1d0`.

To Write the mcf File

Begin writing the `mcf` file for this configuration example by defining the file system and its disk partitions, as follows:

1. Make an `ma` entry for the first file system.

   The name of this file system (`qfs1`) is used later when writing the `/etc/vfstab` entry for the file system and making the file system.

2. Make an `mm` entry listing the partition(s) that comprise the metadata for the `qfs1` file system.

3. Make a series of `mr` entries listing the partitions that comprise the file data for the `qfs1` file system.

4. Make similar entries for the second (`qfs2`) file system.

The resulting `mcf` file is as follows:

```
# Disk cache configuration for 2 file systems: qfs1, qfs2
# Equipment       Eq    Eq   Fam.  Dev.    Additional
# Identifier      Ord  Type  Set   State   Parameters
#-----------      ---   --  ------ ------  -----------------
#
qfs1                1   ma  qfs1
/dev/dsk/c2t0d0s0  11   mm  qfs1   on
/dev/dsk/c1t0d0s0  12   mr  qfs1   on
/dev/dsk/c1t1d0s0  13   mr  qfs1   on
#
#
qfs2               20   ma  qfs2
/dev/dsk/c2t0d0s1  21   mm  qfs2   on
/dev/dsk/c1t0d0s1  22   mr  qfs2   on
/dev/dsk/c1t1d0s5  23   mr  qfs2   on
```

Configuration Example 2

The server in this example has a StorageTek Clarion RAID device with four StorageTek OPENstorage 9153 disk drives. Each drive has 34 gigabytes of storage.

The Sun Solaris `format(1M)` command reports that the disks are partitioned as follows:

```
# format
Searching for disks…done

AVAILABLE DISK SELECTIONS:
    0. c0t0d0 <SUN4.2G cyl 3880 alt 2 hd 16 sec 135>
       /sbus@1f,0/SUNW,fas@e,8800000/sd@0,0
    1. c0t1d0 <SEAGATE-ST39140WC-1206 cyl 9004 alt 2 hd 8 sec 246>
       /sbus@1f,0/SUNW,fas@e,8800000/sd@1,0
    2. c2t4d0 <STK-OPENstorage9153-0205 cyl 17338 alt 2 hd 64 sec 64>
       /pseudo/rdnexus@2/rdriver@4,0
    3. c2t4d1 <STK-OPENstorage9153-0205 cyl 17338 alt 2 hd 64 sec 64>
       /pseudo/rdnexus@2/rdriver@4,1
    4. c2t2d2 <STK-OPENstorage9153-0205 cyl 34977 alt 2 hd 64 sec 64>
       /pseudo/rdnexus@2/rdriver@4,2
    5. c2t4d3 <STK-OPENstorage9153-0205 cyl 34977 alt 2 hd 64 sec 64>
       /pseudo/rdnexus@2/rdriver@4,3
    6. c3t2d0 <SEAGATE-ST15230W-0168 cyl 3974 alt 2 hd 19 sec 111>
       /sbus@1f,0/QLGC,isp@2,10000/sd@2,0
```

One file system named `qfs1` is created on disks `c2t4d0`, `c2t4d1`, `c2t4d2`, and `c2t4d3`. Each disk is partitioned identically, with slice `0` consuming the entire disk. The following is an example of a partition map for these disks.

```
Part    Tag          Flag   Cylinders    Size      Blocks
----    ---          ----   ---------    ----      ------
0       usr          wm     0-17377      33.86GB   (17337/0/0)
                                                   71012352
1       unassigned   wm     0            0         (0/0/0)
2       backup       wu     0-17377      33.86GB   (17337/0/0)
                                                   71012352
3       unassigned   wm     0            0         (0/0/0)
4       unassigned   wm     0            0         (0/0/0)
5       unassigned   wm     0            0         (0/0/0)
6       unassigned   wm     0            0         (0/0/0)
7       unassigned   wm     0            0         (0/0/0)
```

The file system entries in the `mcf` file are as follows:

```
# QFS-Standalone file system configuration example
#
# Equipment       Equip  Equip  Fam   Dev    Additional
# Identifier      Ord    Type   Set   State  Parameters
# ----------      -----  -----  ----  -----  ----------
qfs1              10     ma     qfs1
/dev/dsk/c1t1d0s0 11     mm     qfs1  on
/dev/dsk/c2t4d0s0 12     mr     qfs1  on
/dev/dsk/c2t4d1s0 13     mr     qfs1  on
/dev/dsk/c2t4d2s0 14     mr     qfs1  on
/dev/dsk/c2t4d3s0 15     mr     qfs1  on
```

In the preceding `mcf` file, the lines are as follows:

Line 1 defines an ASM/QFS-Standalone file system. The name of this file system, `qfs1`, is used later when writing the `/etc/vfstab` entry for the file system and when creating the file system.

Line 2 shows an `mm` device type entry for the metadata device. Note that this entry is not part of the RAID device described previously. A separate disk is used for caching inode information, leaving the RAID for high-speed data accesses.

Lines 3, 4, 5, and 6 are the data devices using the `mr` device type.

Configuration Example 3

This example illustrates an ASM/QFS-Standalone file system. If there is a low-latency device available, such as a solid-state disk, use that for the metadata. Round-robin allocation is used on four disk drives.

The following assumptions are used:

- The metadata device is a single partition (`s1`) used on controller `0`, LUN `0`.
- The data devices consist of four disks attached to controller 1. Each disk is on a separate target (`1-4`).

To Configure the System

This example introduced the round-robin data layout. For more information on data layout, see the *ASM, ASM-QFS, and ASM/QFS-Standalone File System Administrator's Guide*.

1. Write the `mcf` file.

   The following is a sample `mcf` file for a round-robin disk configuration:

```
# QFS-Standalone disk cache configuration - Round-robin mcf
sample
#
# Equipment       Eq   Eq    Fam.  Dev.    Additional
# Identifier      Ord  Type  Set   State   Parameters
#-----------      ---   --  ------ ------  ------------------
qfs1                1   ma   qfs1
/dev/dsk/c0t0d0s1  11   mm   qfs1    on
/dev/dsk/c1t1d0s1  12   mr   qfs1    on
/dev/dsk/c1t2d0s1  13   mr   qfs1    on
/dev/dsk/c1t3d0s1  14   mr   qfs1    on
/dev/dsk/c1t4d0s1  15   mr   qfs1    on
```

Note: For completeness, this example includes information on modifying the `/etc/vfstab` system file and on using the `sammkfs`(1M) command to initiate the file system. These steps are described in more detail later in this procedure.

2. Modify the `/etc/vfstab` file.

   The `/etc/vfstab` is edited.

To explicitly set round-robin on the file system, set the `stripe=0` mount parameter as follows:

```
qfs1    -    /qfs    samfs    -   yes   stripe=0
```

3. Run the `sammkfs`(1M) command.

Initialize the ASM/QFS-Standalone file system by using the `sammkfs`(1M). The default DAU is 64 kilobytes, but the following example sets the DAU size to 128 kilobytes:

```
# sammkfs -a 128 qfs1
```

Configuration Example 4

This example illustrates an ASM/QFS-Standalone file system that again separates the metadata onto a low-latency disk. File data is striped to four disk drives.

The following assumptions are used:

- The metadata device is a single partition (`s6`) used on controller `0`, LUN `0`.
- The data devices consist of four disks attached to controller `1`. Each disk is on a separate LUN (1-4).

To Configure the System

1. Write the `mcf` file.

   Write the `mcf` file using the disk configuration assumptions. The following is a sample `mcf` file for a striped disk configuration:

```
# QFS-Standalone disk cache configuration - Striped Disk mcf
sample
#
# Equipment       Eq   Eq    Fam.  Dev.    Additional
# Identifier      Ord  Type  Set   State   Parameters
#-----------      ---   --  ------ ------  ------------------
qfs1               10   ma  qfs1
/dev/dsk/c0t1d0s6  11   mm  qfs1     on
/dev/dsk/c1t1d0s6  12   mr  qfs1     on
/dev/dsk/c1t2d0s6  13   mr  qfs1     on
/dev/dsk/c1t3d0s6  14   mr  qfs1     on
/dev/dsk/c1t4d0s6  15   mr  qfs1     on
```

Note: For completeness, this example includes information on modifying the `/etc/vfstab` system file and on initializing the file system using the `sammkfs`(1M) command. These steps are described in more detail later in this procedure.

2. Modify the `/etc/vfstab` file.

Set the stripe width using the `stripe=` option. The following example sets the stripe width equal to one disk allocation unit (DAU):

```
qfs1    -    /qfs    samfs    -    yes    stripe=1
```

This setting stripes file data across all four of the `mr` data drives with a stripe width of one disk allocation unit (DAU). Note that the DAU is the allocation unit you set when you initialize the file system.

3. Run the `sammkfs`(1M) command.

   Initialize the ASM/QFS-Standalone file system by using the `sammkfs`(1M) command. The following examples set the DAU size to 128 kilobytes:

```
# sammkfs -a 128 qfs1
```

With this striped disk configuration, any file written to this file system is striped across all of the devices in increments of 128 kilobytes. Files less than the aggregate stripe width times the number of devices still use 128 kilobytes of disk space. Files larger than 128 kilobytes have space allocated for them as needed in total space increments of 128 kilobytes. Metadata is written to device 11 only.

Configuration Example 5

Striped groups allow you to group RAID devices together for very large files. Normally, a DAU is represented by one bit in the bitmaps. With striped groups, however, there is only one DAU per striped group. This method of writing huge DAUs across RAID devices saves bitmap space and system update time. Striped groups are useful for writing very large files to a group of RAID devices.

Note: A DAU is the minimum disk space allocated. The minimum disk space allocated in a striped group is as follows:

allocation_unit X *number of disks in the group*

Writing a single byte of data fills the entire striped group. The use of striped groups is for very specific applications. Make sure that you understand the effects of using striped groups with your file system.

The devices within a striped group must be the same size. It is not possible to increase the size of a striped group. You can add additional striped groups, however.

This example configuration illustrates an ASM/QFS-Standalone file system that separates the metadata onto a low-latency disk. Two striped groups are set up on four drives. The following assumptions are used:

- The metadata device is a single partition (s6) used on controller 0, LUN 0.
- The data devices consist of four disks (two groups of two identical disks) attached to controller 1. Each disk is on a separate LUN (1-4).

To Configure the System

1. Write the mcf file.

   Write the mcf file by using the disk configuration assumptions. The following is a sample mcf file for a striped group configuration:

```
# QFS-Standalone disk cache configuration - Striped Groups
mcf sample
#
# Equipment       Eq   Eq    Fam.  Dev.    Additional
# Identifier      Ord  Type  Set   State   Parameters
#-----------      ---   --  ------ ------  ------------------
qfs1               10   ma  qfs1
/dev/dsk/c0t1d0s6  11   mm  qfs1     on
/dev/dsk/c1t1d0s4  12   g0  qfs1     on
/dev/dsk/c1t2d0s4  13   g0  qfs1     on
/dev/dsk/c0t3d0s4  14   g1  qfs1     on
/dev/dsk/c0t4d0s4  15   g1  qfs1     on
```

Note: For completeness, this example includes information on modifying the /etc/vfstab system file and on initializing the file system by using the sammkfs(1M) command. These steps are described in more detail later in this procedure.

2. Modify the /etc/vfstab file.

   Set the stripe width by using the stripe= option. This sample sets the stripe width equal to zero, which essentially specifies a round-robin allocation from striped group g0 to striped group g1:

```
qfs1    -    /qfs    samfs    -   yes   stripe=0
```

3. Run the sammkfs(1M) command.

   Initialize the ASM/QFS-Standalone file system by using the sammkfs(1M) command. The -a option is not used with striped groups because the DAU is equal to the size of an allocation, or the size, of each group.

```
# sammkfs qfs1
```

In this example, there are two striped groups, g0 and g1. With stripe=0 in /etc/vfstab, devices 12 and 13 are striped; devices 14 and 15 are striped; and files are round robined around the two striped groups. You are actually

treating a striped group as a bound entity. That is, the configuration of a striped group, once it is created, cannot be changed.

You cannot change these groups without issuing another `sammkfs`(1M) command.

Step 6: Creating the samfs.cmd File (Optional)

You can create the `/etc/opt/SUNWsamfs/samfs.cmd` file as the place from which mount parameters are read. Creating this file can be beneficial if you are configuring multiple ASM/QFS-Standalone systems with multiple mount parameters.

The mount parameters can be provided in the `samfs.cmd` file, in the `/etc/vfstab` file, and on the `mount`(1M) command. Specifications in the `/etc/vfstab` file override the directives in the `samfs.cmd` file, and options to the `mount`(1M) command override specifications in the `/etc/vfstab` file.

Certain features can be managed more easily from a `samfs.cmd` file. These features include the following:

- Striping.
- Quotas.
- Readahead, which specifies the number of bytes that are read ahead when performing paged I/O.
- Writebehind, which specifies the number of bytes that are written behind when performing paged I/O.
- Qwrite, which enables simultaneous reads and writes to the same file from different threads.

For more information on the `samfs.cmd` file, see the *ASM, ASM-QFS, and ASM/QFS-Standalone File System Administrator's Guide*, or see the `samfs.cmd`(4) man page. For more information on the `/etc/vfstab` file, see "Step 7: Creating the Mount Point and Updating the /etc/vfstab FileEdit the /etc/vfstab file and create an entry for each ASM/QFS-Standalone file system." on page 42 . For more information on the `mount`(1M) command, see the `mount_samfs`(1M) man page.

Step 7: Creating the Mount Point and Updating the /etc/vfstab File

1. Edit the `/etc/vfstab` file and create an entry for each ASM/QFS-Standalone file system.

   An example entry follows:

```
qfs1   -   /qfs1   samfs   -   yes   stripe=1
```

   Table 11. shows the various fields and their content.

Table 11. Fields in the `/etc/vfstab` File

Field	Field Title and Contents
1	Device to mount. The name of the ASM/QFS-Standalone file system to mount. This must be the same name as the Family Set name specified in the `mcf` file.
2	Device to `fsck`(1M). Must be a dash (-) character. The dash indicates that there are no options. This prevents the system from performing an `fsck`(1M) on the ASM/QFS-Standalone file system. For more information on this process, see the `fsck`(1M) man page.
3	Mount point. For example, `/qfs1`.
4	File system type. Must be `samfs`.
5	`fsck`(1M) pass. Must be a dash (-) character. A dash indicates that there are no options.
6	Mount at boot. Specifying `yes` in this field causes the ASM/QFS-Standalone file system to be automatically mounted at boot time. Specifying `no` in this field indicates that you do not want to automatically mount the file system. For information on the format of these entries, see the `mount_samfs`(1M) man page.
7	Mount parameters. A list of comma-separated parameters (with no spaces) that are used in mounting the file system. For example, `stripe=1` specifies a stripe width of one DAU. For a list of available mount options, see the `mount_samfs`(1M) man page.

2. Create the mount point.

The example in this step assumes that `/qfs1` is the mount point of the `qfs1` file system. You can select a different name and substitute it for `/qfs1`, if you want. For example:

```
# mkdir /qfs1
```

3. Change the permissions, owner, or group owner of the `/qfs1` directory when it is not mounted. (Optional)

   This is done on a UFS file system. For example:

```
# chmod 555 /qfs1
# chown root /qfs1
# chgrp other /qfs1
```

Note: If you configured multiple mount points, repeat these steps for each mount point, using a different mount point (such as `/qfs2`) and Family Set name (such as `qfs2`) each time.

Step 8: Initializing the File System

To Initialize a File System

- Use the `sammkfs`(1M) command and the Family Set names that you have defined to create a file system for each Family Set.

  For example, the following command creates a file system for Family Set name `qfs1`:

```
# sammkfs -a 128 qfs1
```

  At this point, the system generates a message similar to the following:

```
# sammkfs qfs1
Building 'qfs1' will destroy the contents of devices:
                /dev/dsk/c1t0d0s0
                /dev/dsk/c3t1d0s6
                /dev/dsk/c3t1d1s6
                /dev/dsk/c3t2d0s6
Do you wish to continue? [y/N]
```

  Enter `y` in response to this message to continue the file system creation process.

Caution: Running `sammkfs`(1M) creates a new file system. It removes all references to the data currently contained in the partition associated with the file system in the `/etc/opt/SUNWsamfs/mcf` file.

Step 9: Mounting the File System

The mount(1M) command mounts a file system. For information on the mount(1M) command, see the mount_samfs(1M) man page.

The mount(1M) command mounts an ASM/QFS-Standalone file system and reads the /etc/vfstab configuration file. For more information on the mount(1M) command, see the mount_samfs(1M) man page. For more information on the /etc/vfstab file, see the vfstab(4) man page.

ASM/QFS-Standalone file systems can be mounted either automatically at boot time or manually. This step shows both methods.

To Mount the File System Automatically

- Edit the /etc/vfstab file and change the Mount at Boot field, which is the sixth field in the file, to yes.

The following example entry in the /etc/vfstab file specifies that the qfs1 file system should be mounted at system startup.

```
qfs1 - /qfs1 samfs - yes stripe=1
```

The previous line causes /etc/rc1.d/S01MOUNTFSYS to mount the qfs1 file system at system startup.

Proceed to “To Confirm that a File System is Mounted and Set Permissions” on page 43.

To Mount the File System Manually

1. Edit the /etc/vfstab file and change the Mount at Boot field, which is the sixth field in the file, to no. (Optional)

The advantage to making an /etc/vfstab entry is that it allows you to provide mount parameters in the /etc/vfstab file. When the file system is mounted, the mount(1M) command reads the mount parameters from the /etc/vfstab file and you do not have to specify them on the command line.

The following example entry in the /etc/vfstab file specifies that the qfs1 file system should not be mounted automatically at system startup.

```
qfs1 - /qfs1 samfs - no stripe=1
```

2. Issue the mount(1M) command to mount the file system after system startup.

If there is an entry for this file system in the `/etc/vfstab` file, issue the `mount`(1M) command and specify the file system mount point as the argument. For example:

```
# mount /qfs1
```

If there is no entry for this file system in the `/etc/vfstab` file, issue the `mount`(1M) command and specify the file system type, file system name, and the file system mount point as arguments. For example:

```
# mount -F samfs qfs1 /qfs1
```

To Confirm that a File System is Mounted and Set Permissions

1. Issue the `mount`(1M) command with no arguments.

Examine its output to see whether the file system is mounted. For example:

```
# mount
<<< information deleted >>>
/qfs on qfs6 read/write/setuid/dev=8001b1 on Mon Jan 14
12:21:03 2002
<<< information deleted >>>
```

2. Issue the `chmod`(1) and `chown`(1) commands to change the permissions and ownership of the file system's root directory. (Optional)

If this is the first time the file system has been mounted, it is typical to perform this step. For example:

```
# chmod 755 /qfs
# chown root:other /qfs
```

Step 10: Sharing the File System With NFS Client Systems (Optional)

Perform this step if you want the ASM/QFS-Standalone file system to be NFS shared.

The Sun Solaris `share`(1M) command must be run to make the file system available for mounting by remote systems. The `share`(1M) commands are typically placed in the `/etc/dfs/dfstab` file and are executed automatically by the Sun Solaris OE when you enter `init`(1M) state 3.

To NFS Share the File System

1. Use an editor to add a `share`(1M) command to the `/etc/dfs/dfstab` file.

   For example, add a line like the following:

```
share -F nfs -o rw=client1:client2 -d “QFS" /qfs1
```

2. Use the `ps`(1) command to determine whether or not `nfs.server` is running.

   Issue the following commands:

```
# ps -ef | grep nfsd
# ps -ef | grep mountd
```

3. If `nfs.server` is not running, enter the following command to start `nfs.server`:

```
# /etc/init.d/nfs.server start
```

4. If you want to NFS share the file system immediately, you must type the `share`(1M) command at a `root` shell prompt. (Optional)

If there are no NFS shared file systems when the Sun Solaris OE boots, the NFS server is not started. You must change to run level 3 after adding the first share entry to this file. For example:

```
# init 3
# who -r
.       run-level 3  Dec 12 14:39     3    2  2
# share
-          /qfs   -   “QFS”
```

Some NFS mount parameters can affect the performance of an NFS mounted ASM/QFS-Standalone file system. You can set these parameters in the `/etc/vfstab` file as follows:

- `timeo =` *n*. This value sets the NFS timeout to *n* tenths of a second. The default is 11 tenths of a second. For performance purposes, StorageTek recommends using the default value. You can increase or decrease the value appropriately to your system.
- `rsize =` *n*. This value sets the read buffer size to *n* bytes. In NFS 2, change the default value (8192) to 32768. In NFS 3, retain the default value of 32768.

- `wsize = n`. This value sets the write buffer size to *n* bytes. In NFS 2, change the default value (8192) to 32768. In NFS 3, retain the default value of 32768.

For more information on these parameters, see the `mount_nfs`(1M) man page.

Step 11: Mounting the File System on the NFS Client Systems (Optional)

Perform this step if you have shared the file system with NFS client systems.

1. On the client systems, mount the server's ASM/QFS-Standalone file system at a convenient mount point.

In the following example, `server:/qfs1` is mounted on `/qfs1`, and information is entered into the `/etc/vfstab` file:

```
server:/qfs1 - /qfs1  nfs - no intr,timeo=60
```

2. Enter the `mount`(1M) command.

   For example:

```
# mount /qfs1
```

The automounter can also do this, if it is preferred. Follow your site procedures for adding `server:/qfs1` to your automounter maps.

Step 12: Establishing Periodic Dumps Using qfsdump(1M)

File systems are made up of directories, files, and links. The ASM/QFS-Standalone file system keeps track of all the files in the `.inodes` file. The `.inodes` file resides on a separate metadata device. All file data is stored on the data devices.

It is important to periodically use the `qfsdump`(1M) command to create a dump file of metadata and file data. The `qfsdump`(1M) command saves the relative path information for each file contained in a complete file system or in a portion of a file system. This protects your data in the event of a disaster.

Dump files should be created at least once a day, but the frequency depends on your site's requirements. By dumping file system data on a regular basis, old files and file systems can be restored or moved from one file system to another, or even from one server to another.

The following are some guidelines to be used when creating dump files:

- The `qfsdump`(1M) command dumps file names, inode information, and data. This command creates full dumps of specified files and directories. An incremental dump is not available. Because of these factors, a `qfsdump`(1M) dump file can be very large. The `qfsdump`(1M) command does not have any tape management, size estimations, or incremental dump facilities as does `ufsdump`(1M). In addition, the `qfsdump`(1M) command does not support volume overflow, so you need to weigh space considerations and make sure that the size of the file system does not exceed the size of the dump media.
- The `qfsdump`(1M) command dumps all the data of a sparse file, and the `qfsrestore`(1M) command restores all the data. These commands do not, however, preserve file qualitites that enable sparse files to be characterized as sparse. This can lead to files occupying more space on dump files and on restored file systems than anticipated.
- The `qfsdump`(1M) command must be issued on a mounted file system. Inconsistencies can arise as new files are being created on disk. Dumping file systems during a quiet period (a time when files are not being created or modified) is a good idea and minimizes these inconsistencies.
- Ensure that you dump metadata and data for all ASM/QFS-Standalone file systems. Look in `/etc/vfstab` for all file systems of type `samfs`.

You can run the `qfsdump`(1M) command manually or automatically. Even if you implement this command to be run automatically, you might need to run it manually from time to time depending on your site's circumstances. In the event of a disaster, you can use the `qfsrestore`(1M) command to recreate your file system. You can also restore a single directory or file. For more information on these commands, see the `qfsdump`(1M) man page and see the *ASM, ASM-QFS, and ASM/QFS-Standalone Disaster Recovery Guide*.

For more information on creating dump files, see the `qfsdump`(1M) man page. The following sections describe procedures for issuing this command both manually and automatically.

To Run the qfsdump(1M) Command Automatically

1. Make an entry in `root`'s `crontab` file so that the `cron` daemon runs the `qfsdump`(1M) command periodically.

   For example:

```
10 0 * * * (cd /qfs1; /opt/SUNWsamfs/sbin/qfsdump -f /dev/
rmt/0cbn)
```

This entry executes the `qfsdump`(1M) command at 10 minutes after midnight. It uses the `cd`(1) command to change to the mount point of the `qfs1` file

system, and it executes the `/opt/SUNWsamfs/sbin/qfsdump` command to write the data to tape device `/dev/rmt/0cbn`.

2. Using the previous step as a guide, make similar `crontab` file entries for each file system. (Optional)

   Perform this step if you have more than one ASM/QFS-Standalone file system.

   Make sure you save each dump file in a separate file.

To Run the qfsdump(1M) Command Manually

1. Login as superuser.
2. Use the `cd`(1) command to go to the directory that contains the mount point for the file system.

   For example:

```
# cd /qfs1
```

3. Create a dump file by issuing the `qfsdump`(1M) command and writing the output to a file system outside of the one you are dumping.

   For example:

```
# qfsdump -f /save/qfs/dump_file
```

ASM/QFS-Standalone Upgrade Procedure

4

This chapter describes upgrading a server to a new release of the QFS software. Use this procedure if you are upgrading your ASM/QFS-Standalone file system.

You must perform all the steps in this chapter as superuser (`root`).

The main sections in this chapter are as follows:

Step 1: Obtaining the Release Files

You can obtain the ASM-QFS software from the StorageTek Customer Resource Center (CRC) or on a CD-ROM. The StorageTek CRC is located at the following URL:

`http://www.support.storagetek.com.` Contact your StorageTek sales representative if you have questions on obtaining the software in one of these ways.

After the release, upgrade patches are also available from the following URL:

`http://www.support.storagetek.com`

Caution: If you have not read the `README` file delivered with this release, please do so before continuing. You can access the `README` file for this release at any time from one of the documentation websites described in this manual's preface. It is available as the *ASM and ASM-QFS README File*. After your software is installed, the content of the `README` file is located in `/opt/SUNWsamfs/doc/README`.

To Install from a CD-ROM

1. Log in as root.

The ASM/QFS-Standalone file system uses the Sun Solaris operating environment (OE) packaging utilities for adding and removing software. You must be logged in as superuser (`root`) to make changes to software packages. The `pkgadd`(1M) utility prompts you to confirm various actions necessary to install the package.

2. Insert the CD into the CD drive.

   The system should automatically detect the CD's presence. If it does not, issue commands to stop and start the ASM Volume Manager and to change to the directory that contains the QFS software package.

For example:

```
# /etc/init.d/volmgt stop
# /etc/init.d/volmgt start
# volcheck
# cd /cdrom/cdrom0
```

On the CD, the ASM-QFS package resides in the `/cdrom/cdrom0` directory organized by the ASM version.

■ Step 2: Backing Up Each File System

To Back Up Each File System

1. Log in as the `root` user (Optional)

   If you have not already logged in as `root`, do so now.

2. Use the `boot`(1M) command to boot the system in single-user mode.

Use the following command:

```
# boot -s
```

3. Use the `mount`(1M) command to mount the QFS file system.

   For example:

```
# mount /qfs1
```

4. Use the `qfsdump`(1M) command to back up the file data and metadata of each ASM/QFS-Standalone file system.

   The `qfsdump`(1M) command dumps file names, inode information, and file data. The destination of the `qfsdump`(1M) output (generally a file) must be as large or larger than the ASM/QFS-Standalone file system you are backing up. The destination location must have enough space (disk or tape) to hold the amount of file data and metadata you are dumping. For more information on using the `qfsdump`(1M) command, see “Step 12: Establishing Periodic Dumps Using qfsdump(1M)” on page 46 or see the `qfsdump`(1M) man page.

The location to which you dump each file system should be outside of the ASM/QFS-Standalone file system. For more information, see the `qfsdump`(1M) man page.

Note: You must back up your file system now if you plan to use the following QFS 4.0 features:

- Access Control Lists (ACLs)
- QFS shared file system
- `md` devices in ASM/QFS-Standalone or ASM-QFS (`ma`) file systems
- Dual-sized disk allocation units (DAUs) on `mm` devices

In order to use any of these features, you need to reinitialize the file system. Reinitializing the file system is described in “Step 10: Reinitializing and Restoring the File Systems (Optional)” on page 54 of this installation process. After the file system is reinitialized, by using the `sammkfs`(1M) command, you can use the `qfsrestore`(1M) command to restore files to the new file system from the dump file created in this installation step.

For example, assume that you have a file system named `qfs1` (mounted at `/qfs1`) that you want to back up. Your choices are as follows:

a. You can write the `qfsdump`(1M) output as a file to a file system or a tape device.

For example, to write to a tape in device `/dev/rmt/1cbn`, use the following commands:

```
# cd /qfs1
# qfsdump -f /dev/rmt/1cbn
```

b. You can write the `qfsdump`(1M) output to a file in a UFS file system

For example, you can use the following commands:

```
# cd /qfs1
# qfsdump -f /save/qfs/qfs1.bak
```

c. You can intialize a new ASM/QFS-Standalone file system, using a ASM/QFS-Standalone 4.0 or later release, and perform the `qfsrestore`(1M) directly into that new ASM/QFS-Standalone file system.

This alternative is applicable only if you already have the ASM/QFS-Standalone software installed and operational as a file system somewhere in your environment. Pursuing this alternative assumes that you want to use the features included in the ASM/QFS-Standalone 4.0 release.

For example, assume that you want to write the dump file into a second ASM/QFS-Standalone file system called `qfs2` (mounted at `/qfs2`). Make sure that the `qfs2` file system has been initialized using ASM/QFS-Standalone release 4.0, or later. You can accomplish this by issuing the following commands:

```
# mount /qfs2
# cd /qfs1
# qfsdump -f - | (cd /qfs2; qfsrestore -f -)
```

5. Repeat these steps for each ASM/QFS-Standalone file system in your environment.

For more information on backing up your file systems, see the *Sun QFS, Sun SAM-FS, and Sun SAM-QFS Disaster Recovery Guide*.

Step 3: Unsharing the File Systems (Optional)

You must complete this step if your ASM/QFS-Standalone file systems are NFS shared file systems.

To Unshare the File Systems

- Use the `unshare`(1M) command on the ASM/QFS-Standalone file system.

  For example, the following command unshares the `qfs1` file system:

```
# unshare /qfs1
```

■ Step 4: Unmounting the File Systems

There are several ways to unmount a file system. Any of the following methods can accomplish this task. The easiest method is presented first. After the file system is unmounted, you can proceed to .

To Unmount Using the umount(1M) Command

- Using the `umount`(1M) command, unmount each ASM/QFS-Standalone file system.

  If you are unmounting a file system from a Sun Solaris 8 or later OE, you can use the `-f` option to the `umount`(1M) command. The `-f` option forces a file system to unmount.

To Unmount Using the fuser(1M), kill(1), and umount(1M) Commands

If using `umount`(1M) to unmount is not successful, it might be because you or another user are using files or because you or another user have changed to directories in the file system.

1. Use the `fuser`(1M) command to determine whether or not any processes are still busy.

   For example, the following command queries the `qfs1` file system:

```
# fuser -uc /qfs1
```

2. If any processes are still busy, use the `kill`(1M) command to terminate them.
3. Using the `umount`(1M) command, unmount each ASM/QFS-Standalone file system.

To Unmount by Editing the /etc/vfstab File and Rebooting

1. Edit the `/etc/vfstab` file.

   For all ASM/QFS-Standalone file systems, change the Mount at Boot field from `yes` or `delay` to `no`.

2. Reboot your system.

Step 5: Removing Existing ASM/QFS-Standalone Software

The following section describes how to remove software from a release prior to 4.0.

To Remove Software from a Release Prior to 4.0

1. Use the `pkginfo`(1) command to determine which ASM/QFS-Standalone software packages are installed on your system.

   For example:

```
# pkginfo | grep LSC
```

2. Use the `pkgrm`(1M) command to remove the existing ASM/QFS-Standalone software.

   You must remove all existing ASM/QFS-Standalone packages before installing the new packages. If you are using any optional ASM/QFS-Standalone packages, you should make sure that you remove these packages before removing the main `LSCqfs` package. The install script prompts you to confirm several of the removal steps.

   The following example command removes the `LSCdoc` package and the `LSCqfs` package:

```
# pkgrm LSCdoc LSCqfs
```

   The `LSCqfs` package must be the last package removed.

Note: If you are upgrading from a QFS release prior to the ASM/QFS-Standalone 4.0 release, you must remove both the `LSCdoc` and `LSCqfs` packages, in that order. When you install the new ASM/QFS-Standalone 4.0 package, you will install only a `SUNWqfs` package.

Step 6: Adding the Packages

The ASM/QFS-Standalone software package uses the Sun Solaris packaging utilities for adding and deleting software. The `pkgadd`(1M) command prompts you to confirm various actions necessary to upgrade the ASM/QFS-Standalone package.

On the CD-ROM, the ASM/QFS-Standalone packages and all optional products reside in the `/cdrom/cdrom0` directory organized by Sun Solaris version.

To Add the Packages

1. Issue the `pkgadd`(1M) command to upgrade the `SUNWqfs` package.

   Answer `yes` to each question.

For example:

```
# pkgadd -d SUNWqfs
```

2. Use the `pkgadd`(1M) command to add one or more localized packages. (Optional)

   Perform this step only if you want to install the packages localized for Chinese, French, or Japanese. To install these packages, enter one or more of the following commands:

```
# pkgadd -d SUNWcqfs
# pkgadd -d SUNWfqfs
# pkgadd -d SUNWjqfs
```

Step 7: Updating the License Keys

A license key is required to run the ASM/QFS-Standalone software. For information on license keys, see “Step 6: Verifying the Software License” on page 21.

The ASM/QFS-Standalone file system uses an encrypted license key. The license key consists of an encoded alphanumeric string.

To License the ASM/QFS-Standalone Software

1. Verify whether the license file exists.

   The license file is as follows:

`/etc/opt/SUNWsamfs/LICENSE.4.0`

2. If the /etc/opt/SUNWsamfs/LICENSE.4.0 file does not exist, create it.
3. Starting in column one, place the license key you have obtained from StorageTek on the first line in the /etc/opt/SUNWsamfs/LICENSE.4.0 file.

   The key must start in column one. No other keywords, host IDs, comments, or other information can appear in the /etc/opt/SUNWsamfs/LICENSE.4.0 file. The license becomes effective when the ASM/QFS-Standalone file system is mounted.

Step 8: Verifying the mcf File

The topology of the equipment managed by the ASM/QFS-Standalone file system is defined in the master configuration file, /etc/opt/SUNWsamfs/mcf. This file specifies the devices and file systems included in the environment. Each piece of equipment is assigned a unique Equipment Identifier in the mcf file.

To Verify the mcf File

1. Issue the cd(1) command to change to the /etc/opt/SUNWsamfs directory.

   This is the directory that contains your mcf file.
2. Verify that an mcf file exists.

   The correct location for this file is as follows:

   /etc/opt/SUNWsamfs/mcf
3. Issue the sam-fsd(1M) command to check the mcf file for errors.

For example:

```
# /usr/lib/fs/samfs/sam-fsd
```

Note: For information on file system design considerations, see the *Sun QFS, Sun SAM-FS, and Sun SAM-QFS File System Administrator's Guide*.

To Reinitialize the mcf File

Reinitialization of the mcf file is not needed at this point in the configuration process. Be aware, however, that if you change the mcf file after the ASM/QFS-Standalone file system is in use, you must convey the new mcf specifications to the ASM/QFS-Standalone software. For information on reinitializing the mcf file, see the *Sun QFS, Sun SAM-FS, and Sun SAM-QFS File System Administrator's Guide*.

Step 9: Modifying the /etc/vfstab File (Optional)

If you modified the `/etc/vfstab` file in “Step 4: Unmounting the File Systems” on page 51, you must complete this step.

To Modify the /etc/vfstab File

- Edit this file again, and change the Mount at Boot field for all ASM/QFS-Standalone file systems from `no` to `yes`.

Step 10: Reinitializing and Restoring the File Systems (Optional)

You must use the ASM/QFS-Standalone 4.0 `sammkfs`(1M) command to reinitialize your file systems if you want to use all of the ASM/QFS-Standalone 4.0 features. The features that require you to reinitialize your file systems are noted in “Step 2: Backing Up Each File System” on page 50.

In this step, you reinitialize your file systems and restore the saved data from “Step 2: Backing Up Each File System” on page 50 into the new file systems. To accomplish this, use the `sammkfs`(1M) and `qfsrestore`(1M) commands on each file system.

To Reinitialize and Restore the File Systems

1. Use the `sammkfs`(1M) command to initialize a new ASM/QFS-Standalone file system.

If you want to use all the ASM/QFS-Standalone 4.0 features, issue the `sammkfs`(1M) command without any options. The following example `sammkfs`(1M) command reinitializes a file system named `qfs1` with ASM/QFS-Standalone 4.0 feature capabilities.

```
# sammkfs qfs1
```

For more information on the options to the `sammkfs`(1M) command, see the `sammkfs`(1M) man page.

2. Use the `qfsrestore`(1M) command to restore the dumped data into the new file system.

For example, the following commands assume that you have a file system named `qfs1` (mounted at `/qfs1`) that you want to back up from files dumped to `qfs1.bak`, which exists outside of the ASM/QFS-Standalone file system:

```
# cd /qfs1
# qfsrestore -f /save/qfs/qfs1.bak
```

Note: If you do not want to use all of the ASM/QFS-Standalone 4.0 features, use the `-P` option to the `sammkfs`(1M) command. This creates a version 1 superblock. For more information on the options to the `sammkfs`(1M) command, see the `sammkfs`(1M) man page.

Step 11: Checking the File System (Optional)

If you did not perform “Step 10: Reinitializing and Restoring the File Systems (Optional)” on page 54, you are encouraged to complete this step.

- Use the `samfsck`(1M) command to check your existing file systems for inconsistancies.

Do this for each ASM/QFS-Standalone file system.

Step 12: Mounting the File System (Optional)

If you have *not* modified the `/etc/vfstab` file to have `yes` or `delay`, you must perform this step.

- Use the `mount`(1M) command to mount the file systems and continue operation with the upgraded ASM/QFS-Standalone software.

  In the following example, `qfs1` is the name of the file system to be mounted:

```
# mount qfs1
```

Step 13: Recompiling API-Dependent Applications (Optional)

If you are running applications that use the ASM/QFS-Standalone application programming interface (API), you need to complete this step.

Because file headers, the calling sequence, and other elements of the API can change from release to release, you should recompile all applications that depend on the API at this time.

Caution: Failure to recompile API-dependent applications at this point can cause your applications to generate unexpected results.

ASM and ASM-QFS Initial Installation Procedure

5

This chapter describes the procedure for installing and configuring the ASM and ASM-QFS software for the first time. Use this procedure if this is the initial installation of the ASM or ASM-QFS software package at your site.

If you are upgrading ASM or ASM-QFS software on an existing server, see .

The procedure in this chapter describes copying and installing the software packages onto your server and configuring the software to match the hardware at your site. You must have superuser (`root`) access to perform most of the steps described in this chapter.

The main sections of this chapter are as follows:

Step 1: Obtaining the Release Files

You can obtain the ASM and ASM-QFS software from the StorageTek Customer Resource Center (CRC) or on a CD-ROM. StorageTek CRC is located at the following URL:

`http://www.support.storagetek.com.`

Contact your StorageTek sales representative if you have questions on obtaining the software in one of these ways.

After the release, upgrade patches are available from the following URL:

`http://www.support.storagetek.com.`

Caution: If you have not read the `README` file delivered with this release, please do so before continuing. You can access the `README` file for this release at any time from one of the documentation websites described in this manual’s preface. It is available as the *ASM and ASM-QFS README File*. After your software is installed, the content of the `README` file is located in `/opt/SUNWsamfs/doc/README`.

To Install from a CD-ROM

1. Log in as root.

The ASM and ASM-QFS software uses the Sun Solaris operating environment (OE) packaging utilities for adding and removing software. You must be logged in as superuser (`root`) to make changes to software packages. The `pkgadd`(1M) utility prompts you to confirm various actions necessary to install the packages.

2. Insert the CD into the CD drive.

   The system should automatically detect the CD’s presence. If it does not, issue the following commands to stop and start the Sun Solaris Volume

Manager and to change to the directory that contains the ASM or ASM-QFS software packages.

For example:

```
# /etc/init.d/volmgt stop
# /etc/init.d/volmgt start
# volcheck
# cd /cdrom/cdrom0
```

On the CD, the ASM and ASM-QFS packages reside in the `/cdrom/cdrom0` directory organized by Sun Solaris version.

Step 2: Adding the Administrator Group (Optional)

By default, the ASM and ASM-QFS administrator commands can be executed by a `root` user only. However, during installation you can supply an administrator group name. The `pkgadd`(1M) process prompts you for this group name during ASM or ASM-QFS installation.

Members of the administrator group can execute all administrator commands except for `star`(1M), `samfsck`(1M), `samgrowfs`(1M), `sammkfs`(1M), and `samd`(1M). The administrator commands are located in `/opt/SUNWsamfs/sbin`.

You can also define an operator group that can access only the GUI tools, which are `libmgr`(1M), `samtool`(1M), `robottool`(1M), `previewtool`(1M), and `devicetool`(1M). This group can be defined in the `/etc/opt/SUNWsamfs/defaults.conf` file as described later in “Step 10: Setting Up Default Values (Optional)” on page 75 and in the `defaults.conf`(4) man page.

You can add or remove the administrator group after installing the package by running the `set_admin.sh`(1M) command. This action performs the same function that occurs when you select an administrator group during the package install. You must be logged in as superuser (`root`) to use the `set_admin.sh`(1M) command. You can also undo the effect of this selection and make the programs in `/opt/SUNWsamfs/sbin` executable only by a root user. For more information on this command, see the `set_admin.sh`(1M) man page.

To Add the Administrator Group

1. Choose a group name or select a group that already exists within your environment.
2. Use the `groupadd`(1M) command, or edit the `/etc/group` file.

The following is an entry from the group file designating an administrator group for the ASM or ASM-QFS software. In this example, the `samadm` group consists of both the `adm` and `operator` users.

```
samadm::1999:adm,operator
```

Step 3: Adding the Packages

The ASM and ASM-QFS software uses the Sun Solaris packaging utilities for adding and deleting software. You must be logged in as superuser (`root`) to make changes to software packages. The `pkgadd`(1M) utility prompts you to confirm various actions necessary to install the packages.

To Add the Packages

1. Issue the `pkgadd`(1M) command to add the `SUNWsamfs` package.

   For example:

```
# pkgadd -d SUNWsamfs
```

2. Enter `all` as the answer to the first question, and enter `yes` or `y` as the answer to each of the other questions.

When you install `SUNWsamfs`, you are asked if you want to define an administrator group. Select `y` to accept the default (no administrator group) or select `n` if you want to define an administrator group. You can also reset permissions on certain commands later by using the `set_admin.sh`(1M) command. For more information on this command, see the `set_admin.sh`(1M) man page.

3. Use the `pkgadd`(1M) command to add one or more localized packages. (Optional)

   Perform this step only if you want to install the packages localized for Chinese, French, or Japanese. To install these packages, enter one or more of the following commands:

```
# pkgadd -d SUNWcsamf
# pkgadd -d SUNWfsamf
# pkgadd -d SUNWjsamf
```

Step 4: Editing the st.conf and samst.conf Files

Some tape devices require you to make changes to the `st.conf` file. These changes are needed in order to enable the ASM and ASM-QFS software to

work with these devices. You must edit the `st.conf` file if you want certain drives, including the following, to work within the ASM or ASM-QFS environment:

- DLT 2000, 2200, 2500, 2700, 4000, 4500, 4700, 7000, 8000
- StorageTek 9940 and 9840; StorageTek RedWood SD-3; and StorageTek TimberLine 9490
- IBM 3590 Magstar, IBM 3570
- Sony DTF-2, Sony DTF-1, Sony Advanced Intelligent Tape (AIT), SDX-500C, SDX-300C
- Fujitsu M8100

Perform the steps in this section if you want any of the preceding drives, or any other devices not already present in the Sun Solaris kernel's `st` driver, to operate within your ASM or ASM-QFS environment, you need to complete the steps in this section.

The steps in this section show how to read lines from the `/opt/SUNWsamfs/examples/st.conf_changes` file into the `/kernel/drv/st.conf` file.

To Add a New Device Type

1. Examine the `/opt/SUNWsamfs/examples/st.conf_changes` file.

Decide if you want any of the devices in this file to be under the control of ASM or ASM-QFS. If so, go on to the next step in this procedure.

Devices that are not to be under ASM or ASM-QFS control do not need to be considered.

2. Copy your existing `st.conf` to a backup file.
3. Use an editor to open your `st.conf` file and add the pertinent lines from the `/opt/SUNWsamfs/examples/st.conf_changes` file.

Example 1

The following is an example of a `/kernel/drv/st.conf` file that has been modified to add support in the Sun Solaris kernel for both the StorageTek 9840 tape drive and the DLT 7000 tape drive:

```
tape-config-list =
"STK 9840", "STK 9840 Fast Access", "CLASS_9840",
"QUANTUM DLT7000", "DLT 7000 tape drive", "dlt7-tape";
CLASS_9840 = 1,0x36,0,0x1d679,1,0x00,0;
dlt7-tape = 1,0x36,0,0xd679,4,0x82,0x83,0x84,0x85,3;
```

The preceding file contains four names paired with values. These are called *name*=*value* pairs, and the following sections describe each one.

The First Pair

The first pair is a device triplet. It consists of one *name* specification and, in this example, two *value* triplets.

```
tape-config-list =
"STK 9840", "STK 9840 Fast Access", "CLASS_9840",
"QUANTUM DLT7000", "DLT 7000 tape drive", "dlt7-tape";
```

This pair has the following format:

name="*value*";

As this example shows, multiple *value* triplets can be specified for a single *name*. In this example, two sets of triplets follow the *name* specification. The elements of this format consist of a quoted, comma-separated string in a 3-element triplet format.

- Each element in the triplet must be enclosed in quotation marks (" ").
- Each element must be separated by a comma (,).
- More than one triplet can be specified for a single *name.*
- The final triplet in the specification must be followed by a semicolon (;).
- Each triplet contains the following three elements:
  - Element 1 contains the exact SCSI inquiry string returned from the device by means of the channel.
  - Element 2 contains a comment string for logging and text-based reporting.
  - Element 3 contains a *name* string that refers to another *name*=*value* pair elsewhere in this st.conf file.
- The final triplet must be followed by a semicolon (;).

The Second and Third Pairs

The second and third pairs refer to the third element in each of the device triplets previously defined in the st.conf file.

```
CLASS_9840 = 1,0x36,0,0x1d679,1,0x00,0;
dlt7-tape = 1,0x36,0,0xd679,4,0x82,0x83,0x84,0x85,3;
```

These *name*=*value* pairs have the following format:

- A name as defined previously in the `st.conf` file, which must match the third element of a previously defined *name="value"* device triplet.
- A comma-separated list of hexadecimal values that identifies proper `st` driver interaction with a particular device as specified in the `st`(7D) man page.
- A terminal semicolon (`;`). Do not use a comma (`,`).

Example 2

Errors can occur if the `st.conf` file is not configured properly during ASM or ASM-QFS software installation. The following examples show typical error messages and provides suggestions for problem resolution.

The following message is found in the `sam-log` file:

```
May 18 12:38:18 baggins genu-30[374]: Tape device 31 is
default type. Update ‘/kernel/drv/st.conf’.
```

Corresponding messages are found in the device log for an associated drive. These messages are as follows:

```
1999/05/18 12:34:27*0000 Initialized. tp
1999/05/18 12:34:28*1002 Device is QUANTUM , DLT7000
1999/05/18 12:34:28*1003 Serial CX901S4929, rev 2150
1999/05/18 12:34:28*1005 Known as Linear Tape(lt)
1999/05/18 12:34:32 0000 Attached to process 374
1999/05/18 12:38:18 1006 Slot 1
1999/05/18 12:38:18 3117 Error: Device is type default.
Update /kernel/drv/st.conf
```

The preceding messages indicate that the appropriate changes have not been made to `/kernel/drv/st.conf`.

To Add a Target Device or LUN

The procedure in this section shows how to add a target device or LUN for either a SCSI-attached library or a library attached through a Fibre Channel interface.

Note that in some cases, SCSI drives can use a target number greater than 6 or a LUN greater than 0. This occurs, for example, with DLT 2700 devices, which use a LUN of 1. In this case, you must edit both the `/kernel/drv/samst.conf` file and the `/kernel/drv/st.conf` file.

1. Use `vi`(1) or another editor to open the `/kernel/drv/st.conf` file. (Optional)

   Perform this step only if you are adding device support for a SCSI-attached device.

Add or uncomment the following line for each target/LUN combination, making the appropriate substitutions. For example, the following lines apply to target `4`, LUN `1`:

```
name=”st” class=”scsi” target=4 lun=1;
```

2. Exit the editor.
3. Use `vi`(1) or another editor to open the `/kernel/drv/samst.conf` file.

   The edits you need to make in the `/kernel/drv/samst.conf` file differ depending on whether the device support being added uses a SCSI attachment or a Fibre Channel interface.

   a. For a SCSI-attached device, un-comment or add the appropriate lines for each device.

   For example:

```
name=”samst” class=”scsi” target=4 lun=1;
```

   b. For a library device attached through a Fibre Channel interface, add a line for each device.

   For a library device attached directly through a Fibre Channel interface, the world-wide name used should be the world-wide name of the library device. This line has the following format:

```
name=”samst” parent=”fp” lun=1 fc-port-
wwn=”500104f00041182b”;
```

   For a library device attached through a Fibre-Channel-to-SCSI bridge device, the world-wide name used should be the world-wide name of the bridge device. This line has the following format:

```
name=”samst” parent=”fp” lun=1 fc-port-
wwn=”100000e00221b1a2”;
```

4. Save your changes.
5. Exit the editor.

6. If you have added new devices since running `pkgadd`(1M), run the `samdev`(1M) or `devfsadm`(1M) command to create the device entries in `/dev/samst`.

   In a Sun Solaris 7 OE, enter the following command:

```
# /opt/SUNWsamfs/sbin/samdev
```

   In a Sun Solaris 8 or 9 OE, enter the following command:

```
# /usr/sbin/devfsadm -i samst
```

Step 5: Rebooting the System

Note: Failure to reboot the system at this time can cause the `st` and `samst` modules to remain unconfigured.

To Reboot the System

- Reboot the server by issuing the `touch`(1) and `init`(1M) commands.

  For example:

```
# touch /reconfigure
# init 6
```

  Changes to the `st.conf`, `samst.conf`, and the `/etc/name_to_sysnum` files are enabled at this time.

Step 6: Setting Up PATH and MANPATH Variables

To Set Up PATH and MANPATH Variables

1. For users who need to access the ASM or ASM-QFS user commands (for example, `sls`(1)), add `/opt/SUNWsamfs/bin` to the users' `PATH` variables.
2. For users who need to access the ASM or ASM-QFS man pages, add `/opt/SUNWsamfs/man` to the `MANPATH` variable.
3. For users, such as superusers, who need to access the administrator commands, add `/opt/SUNWsamfs/sbin` to the `PATH` variable.
4. In the Bourne or Korn shells, edit the `.profile` file, change the `PATH` and `MANPATH` variables, and export the variables.

For example:

```
PATH=$PATH:/opt/SUNWsamfs/bin:/opt/SUNWsamfs/sbin
MANPATH=$MANPATH:/opt/SUNWsamfs/man
export PATH MANPATH
```

5. In the C shell, edit the `.login` and `.cshrc` files.

   For example, the `path` statement in your `.cshrc` file might look like the following:

```
set path = ($path /opt/SUNWsamfs/bin /opt/SUNWsamfs/sbin)
```

   For example, the `MANPATH` statement in your `.login` file might look like the following:

```
setenv MANPATH /usr/local/man:opt/SUNWspro:/$OPENWINHOME/\
share/man:/usr/share/man:/opt/SUNWsamfs/man
```

Step 7: Licensing the ASM and ASM-QFS Software

License keys are required to run the ASM and ASM-QFS software. For information on license keys, see “System Requirements” on page 11.

The ASM and ASM-QFS environments use encrypted license keys. The license keys consist of encoded alphanumeric strings. You receive one or more license keys depending on the system configuration and the products being licensed.

To License the ASM and ASM-QFS Software

1. Create the `/etc/opt/SUNWsamfs/LICENSE.4.0` file.
2. Starting in column one, place the license key you have obtained from StorageTEk on the first line in the `/etc/opt/SUNWsamfs/LICENSE.4.0` file.

The key must start in column one. No other keywords, host IDs, comments, or other information can appear in the `/etc/opt/SUNWsamfs/LICENSE.4.0` file.

Note: The license becomes effective the next time the `sam-initd` daemon is started.

3. View the current license settings by issuing the `samcmd`(1M) command with its `l` argument, as follows.

```
# samcmd l
```

This is the `l` (letter “ell” as in *license*) argument.

The license keys allow the system to run indefinitely unless one of the following conditions is present:

- You were issued a temporary license. When a temporary license expires, the system is no longer able to load or unload cartridges, or to archive, stage, or release files.
- You have exceeded the number of slots allowed for the license. If you exceed the number of slots for which the system is licensed, you cannot import or label media. Access continues unaffected for files already on disk.
- You have changed the hardware with which the software must interoperate. These types of changes include changes to drives, automated libraries, and servers. Licenses are assigned to a specific `hostid` and are not transferable.

Step 8: Configuring System Logging

The ASM and ASM-QFS systems log errors, cautions, warnings, and other messages using the standard Sun Solaris `syslog`(3) interface. By default, the ASM and ASM-QFS facility is `local7`.

To Enable Logging

1. Open the `/etc/syslog.conf` file in an editor.

   Read in the line from the following file:

   `/opt/SUNWsamfs/examples/syslog.conf_changes`

   The line is similar, if not identical, to the following line:

```
local7.debug  /var/adm/sam-log
```

Note: The preceding entry is all one line and has a TAB character (not a space) between the fields.

This step assumes that you want to use `local7`, which is the default. If you set logging to something other than `local7` in the `/etc/syslog.conf` file, you need to edit the `defaults.conf`(4) file and reset it there, too. For more information, see the `defaults.conf`(4) man page.

2. Use the following commands to append the logging line from `/opt/SUNWsamfs/examples/syslog.conf_changes` to your `/etc/syslog.conf` file:

```
# cp /etc/syslog.conf /etc/syslog.conf.orig
# cat /opt/SUNWsamfs/examples/syslog.conf_changes >> /
etc/syslog.conf
```

3. Create an empty log file and send the `syslogd` a HUP signal.

The following command sequence creates a log file in `/var/adm/sam-log` and sends the HUP to `syslogd`:

```
# touch /var/adm/sam-log
# pkill -HUP syslogd
```

For more information, see the `syslog.conf`(4) and `syslogd`(1M) man pages.

4. Use the `log_rotate.sh`(1M) command to enable log file rotation. (Optional)

   Log files can become very large, and the `log_rotate.sh`(1M) command can help in managing log files. For more information, see the `log_rotate.sh`(1M) man page.

Step 9: Configuring the Environment

Each ASM and ASM-QFS software environment is unique. The system requirements and hardware differ from site to site. The ASM and ASM-QFS environments support a wide variety of tape and optical devices, automated libraries, and disk drives. It is up to you, the system administrator at your site, to set up the specific configuration for your environment.

The topology of the equipment managed by the ASM or ASM-QFS file system is defined in the master configuration file, `/etc/opt/SUNWsamfs/mcf`. This file specifies the devices, automated libraries, and file systems included in the environment. Each piece of equipment is assigned a unique equipment identifier in the `mcf` file.

Note: For information on file system design considerations, see the *ASM, ASM-QFS, and ASM/QFS-Standalone File System Administrator's Guide*.

To configure ASM or ASM-QFS devices, create the `/etc/opt/SUNWsamfs/mcf` file. It must contain a line for each device and family set in your configuration. The `mcf` file contains information that enables you to perform the following tasks:

- Identify the disk devices to be used and organize them into one or more ASM or ASM-QFS file systems.
- Identify the drives to be used and, optionally, associate them with the automated libraries to which they are attached.

Note: The instructions for creating the `mcf` file differ depending on whether you are creating an ASM or ASM-QFS environment.

If you are installing the ASM software, all configuration instructions are contained in this section.

If you are installing the ASM-QFS software, the instructions for library and drive configuration are contained in this section. For information on configuring the file system portion of the ASM-QFS file system, see the .

When you create the `mcf` file, delimit the fields in each line with spaces or tabs. Comment lines entered into this file must start with a pound sign (#) in column 1. Some fields are optional, so use a dash (`-`) to indicate omitted fields.

The following format shows the fields of each line entry in the `mcf` file:

```
#
# ASM-QFS file system configuration
#
# Equipment       Equip  Equip Fam   Dev    Additional
# Identifier      Ord    Type  Set   State  Parameters
# ----------      -----  ----- ----  -----  ----------
```

The `mcf` file can contain both comment lines and lines that pertain to a device. The types of lines that can pertain to a device are as follows:

- Family Set parent identifiers and family set devices
- Family Set member devices
- Standalone devices

Table 12. shows the information to be contained in each field and whether or not the field is a required or optional field.

Table 12. `mcf`(4) File Fields

Field	Description
Equipment Identifier	*Required*. This field specifies the physical device with which to interact, as follows: • If the device is a file system, this is the file system name. • If the device is a member device, this field is the `/dev/dsk` entry. • If the device is an automated library or optical drive, this field is the `/dev/samst` entry. If you are using a network-attached automated library, see the information about managing automated libraries in the *ASM, ASM-QFS, and ASM/QFS-Standalone Storage and Archive Management Guide*. • If the device is a tape drive, it is the `/dev/rmt/`*`n`*`cbn` entry.
Equipment Ordinal	*Required*. This field allows each device to be identified by a unique number. All devices and parent identifiers must be assigned a unique Equipment Ordinal. Enter a unique integer from 1 to 65535. The number selected affects the order in which devices are displayed in the user interfaces. Lower numbers are displayed first.
Equipment Type	*Required*. This field provides information that the software uses to determine how to interact with a particular device. Enter the two- or three-character mnemonic for the device type. Most equipment can use the generic equipment types of `od` (optical disk), `tp` (tape), and `rb` (robot). For specific Equipment Types, see the `mcf`(4) man page.
Family Set	*Required for most devices*. A Family Set parent line identifies groups of devices configured to work together as a group. For a parent device, enter a descriptive name for the Family Set. The Family Set names bind members and parents together, as follows: • If the device is a member device that is associated with a Family Set parent device (that is, a file system or automated library), enter the Family Set name for the parent device. • If the device is a standalone device such as a manually loaded drive, use the dash (`-`) to indicate that this field is omitted.

Table 12. mcf(4) File Fields (Continued)

Field	Description
Device State	*Optional*. Enter a state for the device for when the file system is initialized. This state could be `on`, `off`, `unavail`, `down`, or a dash (`-`) for default behavior.
Additional Parameters	*Required for some devices*. The content of the Additional Parameters field varies by device type, as follows: • If the device is a disk, this field is optional. If this field contains information, the content must identify the `/dev/rdsk` entry. This is similar to the `/dev/dsk` entry in the Equipment Identifier field. • If the device is an automated library, this field is optional. If it contains information, the content can be an alternate path to the library's media catalog file or it can be the default path to the library catalog file (`/var/opt/SUNWsamfs/catalog/`*family_set_name*). For tape and magneto-optical devices, no entry is required in this field.

The following `mcf` entries define a StorageTek 9738 automated library with two 9840 drives:

```
# Equipment          Eq    Eq    Family    Dev    Additional
# Identifier         Ord   Ty    Set       St     Parameters
#
/dev/samst/c0t03u0   50    s9    9738      on
/dev/rmt/0cbn        51    sg    9738      on
/dev/rmt/1cbn        52    sg    9738      on
```

In the preceding example, the library catalog is written to the default location, `/var/opt/SUNWsamfs/catalog/9738`.

For more information about this file, see the `mcf`(4) man page. There is an example `mcf` file located in `/opt/SUNWsamfs/examples/mcf`.

To Reinitialize the mcf File

Reinitialization of the `mcf` file is not needed at this point in the configuration process. Be aware, however, that if you change the `mcf` file after the ASM or ASM-QFS software is in use, you must enter commands to convey the new `mcf` specifications to the system. For information on reinitializing the `mcf` file, see the *ASM, ASM-QFS, and ASM/QFS-Standalone File System Administrator's Guide*.

Example ASM Configuration

Assume that the following equipment is to be included in an ASM configuration:

- Two Seagate ST15230W 4-gigabyte disk drives used as cache for the file system
- One StorageTek 9730 30-slot automated library that contains two DLT tape drives
- One manually loaded DLT 2000 drive
- One HP Model C1710T magneto-optical automated library containing two HP Model C1716 magneto-optical drives
- One manually loaded HP Model C1716 magneto-optical drive

This equipment is connected to three SCSI buses with the following SCSI targets:

- The server's internal, single-ended, SCSI bus with targets shown in Table 13..

Table 13. Targets for Server's Internal Single-ended SCSI Bus

SCSI Target	Equipment
2	Manually loaded magneto-optical drive
3	The Sun Solaris internal hard disk
4	Manually loaded DLT drive

- A differential SCSI bus connected to the HP Model C1710T automated library and file system disks with targets shown in Table 14..

Table 14. Targets for SCSI Bus Connected to the HP Model C1710T Automated Library

SCSI Target	Equipment
0 and 1	Seagate 4-gigabyte disks
2	HP C1710T automated library
5	First optical drive
6	Second optical drive

- A differential SCSI bus connected to the StorageTek 9730 automated library and tape drives with targets shown in Table 15.

Table 15. Targets for SCSI Bus Connected to the StorageTek 9730 Automated Library

SCSI Target	Equipment
0	StorageTek 9730 automated library
1	First DLT 7000 drive
2	Second DLT 7000 drive

Example ASM Disk Cache Configuration

Figure 4. shows the output from the Sun Solaris `format`(1M) command. It reports how the disks are partitioned.

Figure 4. `format`**(1M) Command Example**

```
1. c1t0d0 <SEAGATE-ST15230W-0168 cyl 3974 alt 2 hd 19 sec 111>
/iommu@0,10000000/sbus@0,10001000/QLGC,isp@1,10000/sd@0,0
Current partition table (original) :
Total disk cylinders available: 3974 + 2 (reserved cylinders)
Part    Tag          Flag     Cylinders    Size        Blocks
0       root         wm       0-3499       3.52GB      (3500/0/0)
1       unassigned   wm       3500-3972    487.09MB    (473/0/0)
2       backup       wu       0-3973       4.00GB      (3974/0/0)
3       unassigned   wm       0            0           (0/0/0)
4       unassigned   wm       0            0           (0/0/0)
5       unassigned   wm       0            0           (0/0/0)
6       unassigned   wm       0            0           (0/0/0)
7       unassigned   wm       0            0           (0/0/0)
2. c1t1d0 <SEAGATE-ST15230W-0168 cyl 3974 alt 2 hd 19 sec 111>
 /iommu@0,10000000/sbus@0,10001000/QLGC,isp@1,10000/sd@1,0
Current partition table (original):
Total disk cylinders available: 3974 + 2 (reserved cylinders)
Part    Tag          Flag     Cylinders    Size        Blocks
0       root         wm       1000-3973    2.99GB      (2974/0/0)
1       unassigned   wu       0            0           (0/0/0)
2       backup       wu       0-3973       4.00GB      (3974/0/0)
3       unassigned   wm       0            0           (0/0/0)
4       unassigned   wm       0            0           (0/0/0)
5       unassigned   wm       0-999        1.01GB      (1000/0/0)
6       unassigned   wm       0            0           (0/0/0)
7       unassigned   wm       0            0           (0/0/0)
```

One ASM file system (`samfs1`) is to be placed on partition `0` of disk `c1t0d0` and partition `5` of `c1t1d0`. Another file system (`samfs2`) is to be created on partition `1` of disk `c1t0d0` and partition `0` of disk `c1t1d0`.

The following procedure explains how to begin writing the `mcf` file for this example configuration by defining the file systems and their disk partitions.

To Write the mcf File

1. Make an `ms` (mass storage) entry for the first file system.

   An `ms` entry is the Equipment Identifier for an ASM file system. The name of this file system (`samfs1`) is used later when writing the `/etc/vfstab` entry for the file system and creating the file system. Note that the name as specified in the Equipment Identifier field must be the same as the Family Set name for the file system.

2. Make a series of `md` (magnetic disk) entries listing the partitions that comprise the `samfs1` file system member devices.

3. Make similar entries for the second (`samfs2`) file system.

   The `mcf` file looks like this:

```
# Disk cache configuration for 2 file systems: samfs1, samfs2
#
# Equipment        Eq    Eq   Fam.  Dev.    Additional
# Identifier       Ord  Type  Set   State   Parameters
#-----------       ---   --  ------ ------  -----------------
samfs1              10   ms  samfs1
/dev/dsk/c1t0d0s0  11   md  samfs1   on
/dev/dsk/c1t1d0s5  12   md  samfs1   on
#
samfs2              20   ms  samfs2
/dev/dsk/c1t1d0s0  21   md  samfs2   on
/dev/dsk/c1t0d0s1  22   md  samfs2   on
```

Caution: If you give the wrong partition names, you risk damaging user or system data. This is true when initializing any type of file system. Make sure you specify only disk partitions that are not already in use on your system. Do not use overlapping partitions.

Identifying Peripherals Using the /var/adm/messages File

When your system boots, a series of messages are written to `/var/adm/messages`. These messages identify the Sun Solaris hardware path to each of the peripherals on your system. To display information from the latest system reboot, search backward from the end of the file.

Each peripheral has three lines, as follows (note that the third line wraps to the next line in this example and that the sixth field, `samst2`, indicates that these lines are associated with each other):

```
# tail -200 /var/adm/messages | more
Aug 23 11:52:54 baggins unix: samst2: Vendor/Product ID = HP
C1716T
Aug 23 11:52:54 baggins unix: samst2 at esp0: target 2 lun 0
Aug 23 11:52:54 baggins unix: samst2 is /iommu@0,10000000/
sbus@0,10001000/espdma@5,8400000/esp@5,8800000/samst@2,0
```

Note: For readability, the preceding output, and many other outputs that are wider than 80 characters, have been wrapped for inclusion in this manual.

The first line displays the vendor and product information that the SCSI peripheral reported to the Sun Solaris kernel.

The second line displays the SCSI bus, SCSI target, and LUN of the peripheral.

The third line displays the peripheral's hardware path. This path is reflected in the `/devices` directory. Symbolic links (symlinks) to the `/devices` directory are set up in the `/dev/samst` and `/dev/rmt` directories.

Matching the symbolic link to the correct peripheral is the key to configuring an ASM or ASM-QFS environment. Use the `ls`(1) command with the `-l` option in both the `/dev/samst` and `/dev/rmt` directories to display the path name of the peripheral.

Optionally, you can set up the device down notification script at this point. The `dev_down.sh`(4) man page contains information about setting up this script, which sends email to `root` when a device is marked `down` or `off`. For more information, see the `dev_down.sh`(4) man page.

Configuring a Manually Loaded Magneto-optical Drive

The HP Model C1716T is target 2 on the internal SCSI bus.

To Configure the Drive

1. Examine the `/var/adm/messages` file to find the messages for these devices.

The following information is located in the block of lines in `/var/adm/messages` associated with this device (note that the third line wraps to the next line in this example):

```
Aug 23 11:52:54 baggins unix: samst2: Vendor/Product ID = HP
C1716T

Aug 23 11:52:54 baggins unix: samst2 at esp0: target 2 lun 0
Aug 23 11:52:54 baggins unix: samst2 is /iommu@0,10000000/
sbus@0,10001000/espdma@5,8400000/esp@5,8800000/samst@2,0
```

2. Issue the `cd`(1) command to change to the `/dev/samst` directory.
3. Use the `ls`(1) and `grep`(1) commands to find the correct symbolic link.

   For example, use the following `ls`(1) command:

```
# ls -l | grep “samst@2”
```

   The preceding `ls`(1) command searches for a symbolic link that points to the following hardware path:

```
lrwxrwxrwx   1 root     other         88 Aug 23 12:27 c0t2u0
-> /devices/iommu@0,10000000/sbus@0,10001000/
espdma@5,8400000/esp@5,8800000/samst@2,0:a,raw
```

   The ASM `samst` driver uses the name `/dev/samst/c0t2u0` when referencing the device.

4. Use an editor to open the `/etc/opt/SUNWsamfs/mcf` file.
5. Add an entry for the drive to the `mcf` file.

   Add the following entry in `/etc/opt/SUNWsamfs/mcf`:

```
/dev/samst/c0t2u0  30  od  -  on
```

This entry contains the device name (`/dev/samst/c0t2u0`), a unique Equipment Ordinal (30), the Equipment Type of the drive (`od`), a dash (`-`) to indicate that a Family Set name is not associated with the drive, and the device state (`on`).

Configuring a Magneto Optical Library

The HP C1710T automated library has three SCSI devices: the robotic mechanism and the two magneto optical drives that the automated library loads and unloads.

To Configure the Library

1. Examine the `/var/adm/messages` file to find the messages for these devices.

```
Aug 23 11:52:56 baggins unix: samst16: Vendor/Product ID = HP
C1710T
Aug 23 11:52:56 baggins unix: samst16 at QLGC,isp0: target 2
lun 0
Aug 23 11:52:56 baggins unix: samst16 is /iommu@0,10000000/
sbus@0,10001000/QLGC,isp@1,10000/samst@2,0

Aug 23 11:52:56 baggins unix: samst19: Vendor/Product ID = HP
C1716T
Aug 23 11:52:56 baggins unix: samst19 at QLGC,isp0: target 5
lun 0
Aug 23 11:52:56 baggins unix: samst19 is /iommu@0,10000000/
sbus@0,10001000/QLGC,isp@1,10000/samst@5,0

Aug 23 11:52:56 baggins unix: samst20: Vendor/Product ID = HP
C1716T
Aug 23 11:52:56 baggins unix: samst20 at QLGC,isp0: target 6
lun 0
Aug 23 11:52:56 baggins unix: samst20 is /iommu@0,10000000/
sbus@0,10001000/QLGC,isp@1,10000/samst@6,0
```

2. Issue the `cd`(1) command to change to the `/dev/samst` directory.
3. Use `ls`(1) and `grep`(1) commands to find the correct symbolic links.

Use `ls`(1) commands, as follows, to search for the three symbolic links that point to the `/devices` files with the same Sun Solaris hardware paths shown in the `/var/adm/messages` file:

```
# ls -l | grep “samst@2”
lrwxrwxrwx   1 root     other         74 Aug 23 12:27 c1t2u0
-> /devices/iommu@0,10000000/sbus@0,10001000/
QLGC,isp@1,10000/samst@2,0:a,raw
# ls -l | grep “samst@5”
lrwxrwxrwx   1 root     other         74 Aug 23 12:27 c1t5u0
-> /devices/iommu@0,10000000/sbus@0,10001000/
QLGC,isp@1,10000/samst@5,0:a,raw
# ls -l | grep “samst@6”
lrwxrwxrwx   1 root     other         74 Aug 23 12:27 c1t6u0
-> /devices/iommu@0,10000000/sbus@0,10001000/
QLGC,isp@1,10000/samst@6,0:a,raw
```

4. Use an editor to open the `/etc/opt/SUNWsamfs/mcf` file.
5. Add entries for the library and drives to the `/etc/opt/SUNWsamfs/mcf` file.

   For example:

```
/dev/samst/c1t2u0  50  rb  hp30  on
/dev/samst/c1t5u0  51  od  hp30  on
/dev/samst/c1t6u0  52  od  hp30  on
```

The first line defines the automated library itself. It contains the `/dev/samst` name for the device `(/dev/samst/c1t2u0)` followed by a unique Equipment Ordinal (`50`), the Equipment Identifier (`rb`, for a generic SCSI-attached library), the Family Set identifier specified on all devices associated with this library (`hp30`), and the Device State (`on`).

The two remaining lines define the drives inside the library. They are similar to the manually loaded drives defined in the previous section except that instead of a dash, they include the Family Set name of the library where they reside (`hp30`).

Note: The order of drives as specified in the `mcf` file must match the logical order of the drives in a SCSI-attached automated library. The first drive defined in the `mcf` must be the first logical drive in the library, and so on. For more information, see your hardware documentation for logical orientation of drives. Incorrect configuration can cause cartridges to be mounted in the wrong drives, which will lead to total inoperability of this software.

Configuring a Manually Loaded DLT Drive

When configuring DLT drives, be sure to add the DLT definitions to the `/kernel/drv/st.conf` file (see). DLT drives are not part of the standard

Sun Solaris configuration, and they are not recognized if configured incorrectly.

To Configure the Drive

1. Examine the `/var/adm/messages` file to find the messages for these devices.

   The following lines from the `/var/adm/messages` file refer to the manual DLT drive:

```
Aug 23 11:52:54 baggins unix: samst4: Vendor/Product ID = DEC
DLT2000
Aug 23 11:52:54 baggins unix: samst4 at esp0: target 4 lun 0
Aug 23 11:52:54 baggins unix: samst4 is /iommu@0,10000000/
sbus@0,10001000/espdma@5,8400000/esp@5,8800000/samst@4,0
```

2. Issue the `cd`(1) command to change to the `/dev/samst` directory.

3. Use the `ls`(1) and grep(1) commands to search for the symbolic links that point to the `/devices` files with the same Sun Solaris hardware paths shown in the `/var/adm/messages` file.

   For example:

```
# ls -l | grep “samst@4”
lrwxrwxrwx   1 root     other         88 Aug 23 12:27 c0t4u0
-> /devices/iommu@0,10000000/sbus@0,10001000/
espdma@5,8400000/esp@5,8800000/samst@4,0:a,raw
```

For tape devices (automated tape libraries and tape drives), you can leave the Additional Parameters field empty. The system finds the proper `/dev/samst/*` symbolic link by using the Sun Solaris `st` driver.

Note: The Additional Parameters field is required if the Equipment Identifier field is not in the form `/dev/rmt/*` (the standard `st` device driver). In this case, the Additional Parameters field should be specified using the path to the `samst` special file (for example, `/dev/samst/c`*`n`*`t`*`n`*`u`*`n`*). This is true, for example, for Ampex libraries.

For a tape device, there is another symbolic link located in `/dev/rmt`. This symbolic link is the name that the Sun Solaris `st` driver (see `st`(7)) uses when referencing the device. There are many symbolic links in `/dev/rmt` that point to the hardware path. Each link has various combinations of the option letters `c`, `b`, and `n`. When making the `mcf` entry, always use the `b` and `n` options. If the drive supports compression, and if you want compression in the hardware, use `cbn` as the prefix. The symbolic link is as follows:

```
lrwxrwxrwx   1 root     other         85 Aug 15 11:37 /dev/
rmt/0cbn -> ../../devices/iommu@0,10000000/sbus@0,10001000/
espdma@5,8400000/esp@5,8800000 st@4,0:cbn
```

4. Use an editor to open the `/etc/opt/SUNWsamfs/mcf` file.
5. Add the `/etc/opt/SUNWsamfs/mcf` entry.

   For example:

```
/dev/rmt/0cbn  40  tp  -  on
```

   The first entry on the line is the `st` driver name for the device (`/dev/rmt/0cbn`), followed by a unique Equipment Ordinal (`40`), the Equipment Type (`tp` for a generic tape), a dash (`-`) to indicate that a Family Set name is not associated with the manually mounted device, and the Device State (`on`).

Configuring a DLT Library

The last piece of equipment to define is the STK 9730 automated library. This automated library has three SCSI devices: the robotic mechanism and the two DLT 7000 tape drives that the robot loads and unloads.

To Configure the Library

1. Examine the `/var/adm/messages` to find the messages for these devices.

   The `/var/adm/messages` file is as follows:

```
Aug 23 12:08:41 baggins unix: samst98: Vendor/Product ID =
STK     9730
Aug 23 12:08:41 baggins unix: samst98 at QLGC,isp2:
Aug 23 12:08:41 baggins unix:  target 0 lun 0
Aug 23 12:08:41 baggins unix: samst98 is
/iommu@f,e0000000/sbus@f,e0001000/QLGC,isp@1,10000/samst@0,0
Aug 23 12:08:41 baggins unix: samst99: Vendor/Product ID =
QUANTUM DLT7000
Aug 23 12:08:41 baggins unix: samst99 at QLGC,isp2:
Aug 23 12:08:41 baggins unix:  target 1 lun 0
Aug 23 12:08:41 baggins unix: samst99 is
/iommu@f,e0000000/sbus@f,e0001000/QLGC,isp@1,10000/samst@1,0
Aug 23 12:08:41 baggins unix: samst100:    Vendor/Product ID
= QUANTUM DLT7000
Aug 23 12:08:41 baggins unix: samst100 at QLGC,isp2:
Aug 23 12:08:41 baggins unix:  target 2 lun 0
Aug 23 12:08:41 baggins unix: samst100 is
/iommu@f,e0000000/sbus@f,e0001000/QLGC,isp@1,10000/samst@2,0
```

2. Issue the `cd`(1) command to change to the `/dev/samst` directory.
3. Use the `ls`(1) and `grep`(1) commands to search for the symbolic links that point to the `/devices` files with the same Sun Solaris hardware paths shown in the `/var/adm/messages` file.

   For example:

```
# ls -l | grep “samst@0”
lrwxrwxrwx   1 root           44 Aug 23 09:09 c2t0u0 ->
/devices/iommu@f,e0000000/sbus@f,e0001000/QLGC,isp@1,10000/
samst@0,0:a,raw
# ls -l | grep “samst@1”
lrwxrwxrwx   1 root           44 Aug 23 09:09 c2t1u0 ->
/devices/iommu@f,e0000000/sbus@f,e0001000/QLGC,isp@1,10000/
samst@1,0:a,raw
# ls -l | grep “samst@2”
lrwxrwxrwx   1 root           44 Aug 23 09:09 c2t2u0 ->
/devices/iommu@f,e0000000/sbus@f,e0001000/QLGC,isp@1,10000/
samst@2,0:a,raw
```

4. A tape device is involved, so find a symbolic link in `/dev/rmt` that points to the tape devices.

Note in the following that the automated library does not have this additional link.

```
lrwxrwxrwx   1 root           44 Aug 23 09:09 0cbn ->
../../devices/iommu@f,e0000000/sbus@f,e0001000/
QLGC,isp@1,10000/st@1,0:cbn
lrwxrwxrwx   1 root           44 Aug 23 09:09 1cbn ->
../../devices/iommu@f,e0000000/sbus@f,e0001000/
QLGC,isp@1,10000/st@2,0:cbn
```

Again, there are multiple symbolic links in the directory that point to the same hardware path. To enable hardware compression, choose the one with the `cbn` suffix. Had the drive not supported hardware compression, the symbolic link whose name ended with `bn` would have been the correct choice.

5. Use an editor to open the `/etc/opt/SUNWsamfs/mcf` file.
6. Add the `/etc/opt/SUNWsamfs/mcf` entries.

   For example:

```
/dev/samst/c2t0u0  60  rb  9730  on
/dev/rmt/0cbn      61  tp  9730  on
/dev/rmt/1cbn      62  tp  9730  on
```

The first line defines the automated library and includes the `/dev/samst` name (`/dev/samst/c2t0u0`). It also contains a unique Equipment Ordinal (`60`), the Equipment Type (`rb`, for the generic robot Equipment Type), a Family Set name for the robot and the drive (`9730`), and the Device State (`on`).

The second line defines the first DLT tape drive inside the library. These entries refer to the Equipment Identifier for this tape device (`/dev/rmt/0cbn`), the Equipment Ordinal for the device (`61`), the Equipment Type (`tp`), the Family Set Name (`9730`), and the Device State (`on`).

The third line defines the second DLT tape drive inside the automated library. These entries refer to the Equipment Identifier for this tape device (`/dev/rmt/1cbn`), the Equipment Ordinal for the device (`62`), the Equipment Type (`tp`), the Family Set name (`9730`), and the Device State (`on`).

7. Add the DLT definitions to the `/kernel/drv/st.conf` file. (Optional)

   Perform this step if you are configuring DLT drives.

   This is shown in . DLT drives are not part of the standard Sun Solaris configuration.

Table 16. shows the completed `mcf` file:

.

Table 16. Completed `mcf` File

```
# Equipment          Eq   Eq    Family  Dev   Additional
# Identifier         Ord  Type  Set     Sta   Parameters
#
samfs1               10   ms    samfs1
/dev/dsk/c1t0d0s0    11   md    samfs1  on
/dev/dsk/c1t1d0s5    12   md    samfs1  on
#
samfs2               20   ms    samfs2
/dev/dsk/c1t1d0s0    21   md    samfs2  on
/dev/dsk/c1t0d0s1    22   md    samfs2  on
#
/dev/samst/c0t2u0    30   od    -       on
#
/dev/rmt/0cbn        40   tp    -       on
#
/dev/samst/c1t2u0    50   rb    hp30    on
/dev/samst/c1t5u0    51   od    hp30    on
/dev/samst/c1t6u0    52   od    hp30    on
#
/dev/samst/c2t0u0    60   rb    9730    on
/dev/rmt/0cbn        61   tp    9730    on
/dev/rmt/1cbn        62   tp    9730    on
```

Step 10: Setting Up Default Values (Optional)

The `/opt/SUNWsamfs/examples/defaults.conf` file contains default settings for certain parameters in the ASM and ASM-QFS environment. These settings can be changed after the initial installation. To determine if you want to change any default settings now, examine the `defaults.conf`(4) man page.

To Set Up Default Values

1. Read the `defaults.conf`(4) man page and examine this file to determine which, if any, of the defaults should be changed.
2. Copy the example `/opt/SUNWsamfs/examples/defaults.conf` file to its functional location.

Its functional location is as follows:

```
/etc/opt/SUNWsamfs/defaults.conf
```

3. Edit the file.

Remove the comment characters from the entries to be enabled. Comment characters are preceded by a pound character (#).

Step 11: Creating the samfs.cmd File (Optional)

The `/etc/opt/SUNWsamfs/samfs.cmd` file can be created as the place from which mount parameters are read. Creating this file can be beneficial if you are configuring multiple ASM or ASM-QFS file systems with multiple mount parameters.

The mount parameters can be provided in the `samfs.cmd` file, in the `/etc/vfstab` file, and on the `mount`(1M) command. Specifications in the `/etc/vfstab` file override the directives in the `samfs.cmd` file. Options to the `mount`(1M) command override specifications in the `/etc/vfstab` file.

For more information about the `samfs.cmd` file, see the `samfs.cmd`(4) man page or see the *ASM, ASM-QFS, and ASM/QFS-Standalone File System Administrator's Guide*. For more information about the `/etc/vfstab` file, see "Step 12: Creating the Mount Point and Updating the /etc/vfstab File" on page 76. For more information on the `mount`(1M) command, see the `mount_samfs`(1M) man page.

Step 12: Creating the Mount Point and Updating the /etc/vfstab File

The example in this step assumes that `/sam` is the mount point of the `samfs1` file system. You can select a different name and substitute it for `/sam`.

To Create the Mount Point and Update the /etc/vfstab File

1. Edit the `/etc/vfstab` file and make an entry for each ASM and ASM-QFS file system.

   An example entry, showing the header fields, is as follows:

```
#DEVICE    DEVICE   MOUNT  FS    FSCK  MOUNT    MOUNT
#TO MOUNT  TO FSCK  POINT  TYPE  PASS  AT BOOT  PARAMETERS
#
samfs1     -        /sam   samfs  -    yes      high=80,low=60
```

Table 17. shows the various fields in the `/etc/vfstab` file and their content.

Table 17. `/etc/vfstab` File Fields

Field	Field Title	Content
1	Device to Mount	The name of the ASM or ASM-QFS file system to mount.
2	Device to `fsck`(1M)	A dash (`-`) indicates that there are no options. This prevents the system from performing an `fsck`(1M) on an ASM or ASM- QFS file system. For more information about this process, see the `fsck`(1M) or `samfsck`(1M) man page.
3	Mount Point	For example, `/sam`.
4	File System Type	Must be `samfs`.
5	`fsck`(1M) Pass	A dash (`-`) indicates that there are no options.
6	Mount at Boot	Specifying `yes` in this field requests that the ASM or ASM-QFS file system be automatically mounted at boot time. Specifying `no` in this field indicates that you do not want to automatically mount the file system. For information about the format of these entries, see the `mount_samfs`(1M) man page.
7	Mount Parameters	A list of comma-separated parameters (with no spaces) that are used in mounting the file system. Mount specifications can be provided in the `samfs.cmd` file, in the `/etc/vfstab` file, and on the `mount`(1M) command. Specifications in the `/etc/vfstab` file override the directives in the `samfs.cmd` file. Options to the `mount`(1M) command override specifications in the `/etc/vfstab` file. For a list of available mount options, see the `mount_samfs`(1M) man page.

2. Create the mount point.

   For example:

```
# mkdir /sam
```

3. Change the permissions, owner, or group owner of the `/sam` directory when it is not mounted. (Optional)

   This prevents writing into the mount point directory when the file system is offline for maintenance. For example:

```
# chmod 555 /sam
# chown root:other /sam
```

Note: If you configured multiple mount points, repeat these steps for each mount point, using a different mount point (such as `/sam2`) and family set name (such as `samfs2`) each time.

Step 13: Initializing the File System

This step shows how to use the `sammkfs`(1M) command and the Family Set names that you have defined to initialize a file system for each Family Set.

Note: One tuning parameter, the disk allocation unit (DAU), is set at this time. You cannot reset this parameter without reinitializing the file system. For information about how the DAU affects tuning, see the *ASM, ASM-QFS, and ASM/QFS-Standalone File System Administrator's Guide* or see the `sammkfs`(1M) man page.

To Initialize the File System

- Use the `sammkfs`(1M) command to initialize the file system.

For example, the following command initializes a file system with the Family Set name `samfs1`:

```
# sammkfs samfs1
total data kilobytes       = 31842048
total data kilobytes free  = 31841680
```

The actual numbers returned vary from file system to file system.

Caution: Running the `sammkfs`(1M) command initializes a new file system. It removes all data currently contained in the partitions associated with the file system in the `/etc/opt/SUNWsamfs/mcf` file.

Step 14: Mounting the File System

The `mount`(1M) command mounts a file system. For information on the `mount`(1M) command, see the `mount_samfs`(1M) man page.

The `mount`(1M) command mounts an ASM or ASM-QFS file system and reads the `/etc/vfstab` configuration file. For more information on the `mount`(1M) command, see the `mount_samfs`(1M) man page. For more information on the `/etc/vfstab` file, see the `vfstab`(4) man page.

ASM or ASM-QFS file systems can be mounted either automatically at boot time or manually. This step shows both methods.

To Mount the File System Automatically

- Edit the `/etc/vfstab` file and change the `Mount at Boot` field, which is the sixth field in the file, to `yes`.

The following example entry in the `/etc/vfstab` file specifies that the `samfs1` file system should be mounted at system startup.

```
samfs1 - /samfs1 samfs - yes -
```

The previous line causes `/etc/rc1.d/S01MOUNTFSYS` to mount the `samfs1` file system at system startup.

Proceed to “To Confirm that a File System is Mounted and Set Permissions” on page 90.

To Mount the File System Manually

1. Edit the `/etc/vfstab` file.
2. Change the Mount at Boot field, which is the sixth field in the file, to `no`.

The advantage to making an `/etc/vfstab` entry is that it allows you to provide mount parameters in the `/etc/vfstab` file. When the file system is mounted, the `mount`(1M) command reads the mount parameters from the `/etc/vfstab` file and you do not have to specify them on the command line.

The following example entry in the `/etc/vfstab` file specifies that the `samfs1` file system should not be mounted automatically at system startup.

```
samfs1 - /samfs1 samfs - no trace
```

3. Issue the `mount`(1M) command to mount the file system after system startup.

If there is an entry for this file system in the `/etc/vfstab` file, issue the `mount`(1M) command and specify the file system mount point as the argument. For example:

```
# mount samfs1
```

To Confirm that a File System is Mounted and Set Permissions

1. Issue the `mount`(1M) command with no arguments.

Examine its output to see whether the file system is mounted. For example:

```
# mount
<<< information deleted >>>
/samfs1 on /samfs1 read/write/setuid/dev=8001b1 on Mon Jan 14
12:21:03 2002
<<< information deleted >>>
```

2. Issue the `chmod`(1) and `chown`(1) commands to change the permissions and ownership of the file system's root directory. (Optional)

If this is the first time the file system has been mounted, it is typical to perform this step. For example:

```
# chmod 755 /samfs1
# chown root:other /samfs1
```

■ Step 15: Checking Drive Order

The drive order check procedure differs depending on whether your automated library has a front panel and whether it has tape or magneto-optical drives. Use one of the following procedures to check the drive order of each library:

- "To Check the Drive Order of Tape or Magneto-Optical Libraries With a Front Panel" on page 90
- "To Check the Drive Order of Tape Libraries Without a Front Panel" on page 91
- "To Check the Drive Order of Magneto-Optical Libraries Without a Front Panel" on page 92

To Check the Drive Order of Tape or Magneto-Optical Libraries With a Front Panel

Follow this procedure to check the order of tape and magneto-optical drive devices for libraries with a front panel.

1. Start the ASM or ASM-QFS software by mounting a file system or by using the `samd start` command.
2. Verify the order of the drives.

If the automated library contains more than one drive, the drives defined in the `mcf` file must be listed in the same order as the drives viewed by the automated library's controller. The drive order that is recognized by the controller can be different from the order of the devices as reported in the `/var/adm/messages` file.

3. Verify the order in which the drives are recognized by the automated library's controller.

   Check the SCSI target IDs or world wide numbers displayed by the control panel of the automated library. For optical drives, read the SCSI target IDs displayed on the control panel for your automated library. The order in which the drive targets are reported should be the order in which they are configured in the `mcf` file.

To determine whether the drives become active when loaded with a cartridge, you can visually inspect the drives or you can use the `samu`(1M) utility's `r` display.

Refer to your hardware maintenance manual for instructions on identifying and setting target addresses.

To Check the Drive Order of Tape Libraries Without a Front Panel

1. Start the ASM or ASM-QFS software.

   You can do this by mounting a file system or by using the `samd start` command.

2. Verify the order of the drives.

   If your automated library contains more than one drive, the drives defined in the `mcf` file must be listed in the same order as the drives viewed by the library controller. The drive order that is recognized by the media changer controller can be different from the order of the devices as reported in the `/var/adm/messages` file.

   Make sure you check *each* drive in a library.

Make the drive state unavailable to the ASM or ASM-QFS file system. You can do this by issuing the following `samcmd`(1M) command:

```
# samcmd unavail eq
```

For more information about the `samcmd`(1M) command's format, see the `samcmd`(1M) man page.

3. Load a cartridge into the drive by using the `samload`(1M) command.

Two possible formats for this command are as follows:

```
# samload mediatype.vsn eq
```

or:

```
# samload eq:slot[:partition] eq
```

For more information about the `samload`(1M) command's format, see the `samload`(1M) man page.

4. Determine if the correct drive responds while under ASM or ASM-QFS control.

Enter the following information, where *X* is the raw tape device entry in the `mcf` file:

```
# mt -f /dev/rmt/X status
```

The following example of a status message indicates that a tape is in the drive:

```
# mt -f /dev/rmt/0 status
DLT 7000 tape drive tape drive:
   sense key(0x2)= Not Ready   residual= 0   retries= 0
   file no= 0   block no= 0
```

If the tape did not load or the drive did not return a status, the drive might not be listed in the proper order in the `mcf`. Make sure the order is correct in the `mcf` file, and repeat this test. If you change any information in the `mcf` file, you must reinitialize the `mcf` file. For information on reinitializing the `mcf` file, see the *ASM, ASM-QFS, and ASM/QFS-Standalone File System Administrator's Guide*.

To Check the Drive Order of Magneto-Optical Libraries Without a Front Panel

1. Start the ASM or ASM-QFS software.

   You can do this by mounting a file system or by using the `samd start` command.

2. Verify the order of the drives.

   If your automated library contains more than one drive, the drives defined in the `mcf` file must be listed in the same order as the drives viewed by the library controller. The drive order that is recognized by the media changer

controller can be different from the order of the devices as reported in the `/var/adm/messages` file.

Make sure you check *each* drive in a library.

Make the drive state unavailable to the ASM or ASM-QFS file system. You can do this by issuing the following `samcmd`(1M) command:

```
# samcmd unavail eq
```

For more information about the `samcmd`(1M) command's format, see the `samcmd`(1M) man page.

3. Load a cartridge into the drive by using the `samload`(1M) command.

   Two possible formats for this command are as follows:

```
# samload mediatype.vsn eq
```

or:

```
# samload eq:slot[:partition] eq
```

For more information about the `samload`(1M) command's format, see the `samload`(1M) man page.

4. Determine if the correct drive reponds while under ASM or ASM-QFS control.

Enter the following information at the system prompt:

```
# /bin/dd if=device_path bs=2k iseek=3374 of=/tmp/foo count=10
```

In the `dd`(1M) command, *device_path* is the `samst` device entry in the `mcf` file.

Figure 5. shows a status message that indicates that an optical catridge is in the selected device:

Figure 5. `dd`**(1M) Showing a Cartridge in a Drive**

```
# dd if=/dev/samst/c0t3u0 bs=2k iseek=3374 of=/tmp/junk
count=10
10+0 records in
10+0 records out
```

Figure 6. shows a status message that indicates that an optical catridge is not in the selected device:

Figure 6. `dd`**(1M) Showing a Cartridge in a Drive**

```
# dd if=/dev/samst/c0t5u0 bs=2k iseek=3374 of=/tmp/
junk1 count=10
read: I/O error
0+0 records in
0+0 records out
```

If the optical catridge did not load or if the device returned messages like those in Figure 6., the drives might not be listed in the correct order in the `mcf` file. Make sure the order is correct in the `mcf` file, repeat this test. If you change any information in the `mcf` file, you must reinitialize the `mcf` file. For information on reinitializing the `mcf` file, see the *ASM, ASM-QFS, and ASM/QFS-Standalone File System Administrator's Guide*.

Step 16: Labeling Tapes or Optical Disks (Optional)

If you have standalone tape or optical devices, or if your automated library has no barcode reader, you must perform this step.

To prepare cartridges, use the `tplabel`(1M) command for tapes or use the `odlabel`(1M) command for optical disks. These commands create a label on the cartridge that can be used in the ASM and ASM-QFS environment.

The `tplabel`(1M) command has the following format:

```
# tplabel -new -vsn new_vsn eq:slot
```

where:

new_vsn	The new volume serial name.
eq	The Equipment Ordinal of the automated library or manually loaded drive being addressed as defined in the `mcf` file.
slot	The number of a storage slot in an automated library as recognized in the library catalog. This argument is not applicable for manually loaded drives.

The `odlabel`(1M) command has the following format:

```
# odlabel -new -vsn new_vsn eq:slot:partition
```

where:

new_vsn	The new volume serial name.
eq	The Equipment Ordinal of the automated library or manually loaded drive being addressed as defined in the `mcf` file.
slot	The number of a storage slot in an automated library as recognized in the library catalog. This argument does not apply to manually loaded drives.
partition	A side of a magneto-optical disk. The *partition* must be 1 or 2.

Cartridges are ready to be used after these commands are issued. Both the `tplabel`(1M) and the `odlabel`(1M) commands accept a `-old` option that can be used to relabel previously labeled cartridges. For more information about these commands, see the `tplabel`(1M) and `odlabel`(1M) man pages.

Example 1. The following command labels a tape.

```
# tplabel -vsn TAPE01 -new 50:0
```

Example 2. The following command labels one side of an optical disk.

```
# odlabel -vsn OPTIC01 -new 30:1:1
```

Step 17: Configuring the Archiver (Optional)

By default, the archiver automatically archives all files under all ASM and ASM-QFS mount points. The administrator is not required to take action. The archiver archives to all VSNs in all configured automated libraries.

If your site has additional requirements, you need to set up the archiver command file, `archiver.cmd`. Archiving commences at the end of this installation procedure. If you do not want the archiver to automatically being operating at the end of this installation procedure, you can insert a `wait` directive in the `archiver.cmd` file.

For additional information, see the `archiver.cmd`(4) man page and see the information on the archiver in the *ASM, ASM-QFS, and ASM/QFS-Standalone Storage and Archive Management Guide*.

Step 18: Enabling Disk Archiving (Optional)

Disk archiving is the process of writing archive copies of file data to online disk in another file system. The file system to which the archive copies are written can be any UNIX file system. The destination file system does not have to be an ASM or ASM-QFS file system, but the Sun Solaris host system to which archive files are written must have at least one ASM or ASM-QFS file system installed on it.

Disk archiving differs from traditional archiving in several ways. For example, no automated libraries or removable media cartridges are used. You can, however, specify that one set of archive copies be written to disk and another set be written to other archive media. For a more complete description of disk archiving, see the *ASM, ASM-QFS, and ASM/QFS-Standalone Storage and Archive Management Guide*.

If you plan to enable disk archiving, you must complete the following procedures for enabling disk archiving on the client and on the server.

To Enable Disk Archiving on the Client

1. Log on to the client system.

   This is the system upon which the source files reside.

2. Use the `cd`(1) command to change to the `/etc/opt/SUNWsamfs` directory in the ASM or ASM-QFS file system.

3. Edit the `archiver.cmd` file to add disk archive sets.

   If you did not configure an `archiver.cmd` file in “Step 17: Configuring the Archiver (Optional)” on page 83, you will create the `archiver.cmd` file in this step.

Figure 7. shows a fragment from an `archiver.cmd` file that defines disk archive sets.

Figure 7. `/etc/opt/SUNWsamfs/archive.cmd` **on the Client**

```
# This is the part of the archiver.cmd file that defines
# disk archive sets.
#
params
archset1.1 -disk_archive disk01
archset2.1 -disk_archive disk02
endparams
```

For more information on specifying archive sets, see the `archiver.cmd`(4) man page or see the *ASM, ASM-QFS, and ASM/QFS-Standalone Storage and Archive Management Guide*.

4. Use `vi`(1) or another editor to create a file named `diskvols.conf`.

Figure 8. shows an example `diskvols.conf` file:

Figure 8. Example `diskvols.conf` File on the Client

```
# This is file sourceserver:/etc/opt/SUNWsamfs/diskvols.conf
# on the client.
#
# VSN_name [host_name:] path
#
disk01 otherserver:/sam/archset1
disk02 otherserver:/sam/archset2
```

Figure 8. shows a `diskvols.conf` file that archives files from two archive sets. The disk volumes named `disk01` and `disk02` reside in a file system on the server system named `otherserver`.

As Figure 8. shows, the `diskvols.conf` file can contain comment lines that begin with a pound character (#), and it must contain data in two fields: the VSN name field and the path field. The host name field can be left blank if you are archiving to a file system that resides on the same host system as the source files, but if specified, it must be followed by a colon character (:). Table 18. shows the information that must appear in this file.

Table 18. Format of the `diskvols.conf` File

Field Name	Content
VSN Name	A unique alphanumeric name for the disk VSN to receive the archive copies. Can be up to 31 characters long.
Host Name	The name of the server to which archive copies are written. If you are archiving to disk on another server, you must specify the name of the destination server to receive the archive copies. If a host name is specified, it must be followed by a colon. If you are archiving to a file system that resides on the same server as the source file system, you do not need to specify the host name.
Path	The full path, relative to the mount point, to the directory to receive the archive files. This directory must be in place before archiving can commence, and the destination file system must be mounted.

For more information about the `diskvols.conf` file, see the `diskvols.conf`(4) man page.

To Enable Disk Archiving on the Server

1. Log on to the server system.

   This is the system to which the archive copies will be written.

2. Create directories in the file system to which the archive copies will be written.

3. Use the `cd`(1) command to change to the `/etc/opt/SUNWsamfs` directory in the ASM or ASM-QFS file system.

4. Use `vi`(1) or another editor to create a file named `diskvols.conf`.

   This file contains the `clients` and `endclients` directives and names the client system upon which the the files to be archived reside.

For more information on disk archiving, see the *ASM, ASM-QFS, and ASM/QFS-Standalone Storage and Archive Management Guide*.

Figure 9. Example `diskvols.conf` File on the Server

```
# This is
# file destination_server:/etc/opt/SUNWsamfs/diskvols.conf
# on the server
#
clients
sourceserver
endclients
```

Step 19: Sharing the File System With NFS Client Systems (Optional)

Perform this step if you want the ASM or ASM-QFS file system to be NFS shared.

The Sun Solaris `share`(1M) command must be run to make the file system available for mounting by remote systems. The `share`(1M) commands are typically placed in the `/etc/dfs/dfstab` file and are executed automatically by the Sun Solaris OE when you enter `init`(1M) state 3.

To NFS Share the File System

1. Use an editor to add a `share`(1M) command to the `/etc/dfs/dfstab` file.

   For example, add a line like the following:

```
share -F nfs -o rw=client1:client2 -d “ASM-FS" /samfs1
```

2. Use the `ps`(1) command to determine whether or not `nfs.server` is running.

   Issue the following commands:

```
# ps -ef | grep nfsd
# ps -ef | grep mountd
```

3. Start the NFS server. (Optional)

   Perform this step if `nfs.server` is not running.

Enter the following command to start `nfs.server`:

```
# /etc/init.d/nfs.server start
```

4. Type the `share`(1M) command at a `root` shell prompt. (Optional)

Perform this step if you want to NFS share the file system immediately.

If there are no NFS shared file systems when the Sun Solaris OE boots, the NFS server is not started. You must change to run level 3 after adding the first `share` entry to this file. For example:

```
# init 3
# who -r
.       run-level 3  Dec 12 14:39     3    2  2
# share
-          /samfs1 -   “ASM-FS”
```

Some NFS mount parameters can affect the performance of an NFS mounted ASM or ASM-QFS file system. You can set these parameters in the `/etc/vfstab` file as follows:

- `timeo =` *n*. This value sets the NFS timeout to *n* tenths of a second. The default is 11 tenths of a second. For performance purposes, StorageTek recommends using the default value. You can increase or decrease the value appropriately to your system.

- `rsize = n`. This value sets the read buffer size to *n* bytes. In NFS 2, change the default value (8192) to 32768. In NFS 3, retain the default value of 32768.
- `wsize = n`. This value sets the write buffer size to *n* bytes. In NFS 2, change the default value (8192) to 32768. In NFS 3, retain the default value of 32768.

For more information on these parameters, see the `mount_nfs`(1M) man page.

Step 20: Mounting the File System on NFS Client Systems (Optional)

Perform this step if you have shared the file system with NFS client systems.

On the client systems, mount the server's ASM or ASM-QFS file system at a convenient mount point.

To Mount the File System on Clients

1. Enter a line similar to the following into your `/etc/vfstab` file:

```
server:/sam   -   /sam   nfs   -   yes   hard,intr,timeo=60
```

   In this example, server `:/sam` is mounted on `/sam`, and information is entered into the `/etc/vfstab` file.

2. Issue the `mount`(1M) command and include the name of the file system to be mounted.

   For example:

```
client# mount /sam
```

   The automounter can also do this, if you prefer. Follow your site procedures for adding server `:/sam` to your automounter maps.

Note: It is strongly recommended that clients mount the file system with the `hard` option. At times, there might be a significant delay in the ASM or ASM-QFS file system's response to NFS client requests. This can occur in a situation, for example, when a requested file resides on a cartridge that must be loaded into a DLT tape drive, when all tape drives are full, or when drives are slow. If the `hard` option is not specified, the NFS client can receive an error instead of retrying the operation until it is completed.

Note:

If you use the `soft` option, make sure you set the value of `retrans` to a large number such as 120 (the default is 5). This sets the number of NFS retransmissions occurring at the `time=`*n* interval.

Step 21: Writing Periodic Dump Files By Using samfsdump(1M)

Periodically, the server should create a metadata dump file by using the `samfsdump`(1M) command. The `samfsdump`(1M) command supports dumping unarchived data. The `-u` option of the `samfsdump`(1M) command causes unarchived data to be interspersed with the metadata normally contained in a `samfsdump`(1M) dump.

When using the `samfsdump`(1M) command, please note the following:

- The `samfsdump`(1M) command dumps file names and inode information, not data. That is, the dump file does not include the archive data stored in your file system. The dump file does include the inode and directory structure information necessary to quickly locate the data on your archive media. This information is necessary for recovering from a file system failure. For more information on this, see the `samfsdump`(1M) man page.
- You can use the `-u` option to the `samfsdump`(1M) command to dump metadata and file data for files that have not yet been archived. A `samfsdump`(1M) dump taken using the `-u` option can be very large. The `samfsdump`(1M) command does not have any tape management or estimations as does `ufsdump`(1M). You need to weigh the tradeoffs of space and unarchived data when using the `-u` option. For more information about these commands, see the `samfsdump`(1M) and `ufsdump`(1M) man pages.
- You can use the `samfsrestore`(1M) command to restore the metadata dump file after initializing the file system if a failure occurs.

  For more information about using the `samfsdump`(1M) command, see the `samfsdump`(1M) man page. Also see the information on metadata, disaster preparation, and recovery in the *ASM, ASM-QFS, and ASM/QFS-Standalone Disaster Recovery Guide*.

The following sections describe procedures for issuing this command both automatically and manually.

To Run the samfsdump(1M) Command Automatically

1. Make an entry in the root user's `crontab` file so the `cron` daemon runs the `samfsdump`(1M) command periodically.

   Example 1:

```
0 0 * * * find /csd.directory/sam -type f -mtime +7 \
-print| xargs -l1 rm -f; cd /sam; \
/opt/SUNWsamfs/sbin/samfsdump -f \
/csd.directory/sam/‘date +\%y\%m\%d‘
```

   This example `crontab` entry uses an ASM file system mounted on `/sam`. Replace `/csd.directory` with an existing directory of your choice. This entry causes the commands to be executed each day at midnight. First, the old dumps are renamed and a new dump is created in `/csd.directory/sam/`*yymmdd*. After that, `cron`(1M) emails the `samfsdump`(1M) output to `root`.

2. Using the previous step as a guide, make similar `crontab` entries for each file system. (Optional)

   If you have multiple ASM or ASM-QFS file systems, make similar entries for each. Make sure you save each dump in a separate file.

To Run the samfsdump(1M) Command Manually

1. Log in as superuser.
2. Use the `cd`(1) command to go to the directory that contains the mount point for the file system.

For example:

```
# cd /samfs.mt.pt
```

In this example, `samfs.mt.pt` is an ASM mount point.

3. Create a dump file by issuing the `samfsdump`(1M) command and writing the output to a file system outside of the one you are dumping.

For example:

```
# samfsdump -T -u -f /dumpster/dump.file
```

In this example, `dump.file` is the newly created dump structure.

Step 22: Enabling Automatic Rotating of Log and Trace Files (Optional)

Perform this step if you want to enable log and trace file rotation.

Certain log and trace files can get very large. The `log_rotate.sh`(1M) script rotates the log and trace files generated by the ASM and ASM-QFS software systems. This script can be enabled at any time, but it is sometimes convenient to enable it at installation time. To enable log rotating, see the `log_rotate.sh`(1M) man page for instructions.

ASM and ASM-QFS Upgrade Procedure

6

This chapter explains how to upgrade a server to a new release of the ASM or ASM-QFS software. Use this procedure if you are upgrading your ASM or ASM-QFS environment.

You must perform all the steps in this chapter as superuser (root).

The main sections in this chapter are as follows:

Step 1: Obtaining the Release Files

You can obtain the ASM and ASM-QFS software from the StorageTek Customer Resource Center (CRC) or on a CD-ROM. The StorageTek CRC is located at the following URL:

```
http://www.support.storagetek.com
```

Contact your StorageTek sales representative if you have questions on obtaining the software in one of these ways.

After the release, upgrade patches are also available from the following URL:

`http://www.support.storagetek.com`

Caution: If you have not read the `README` file delivered with this release, please do so before continuing. You can access the `README` file for this release at any time from one of the documentation websites described in this manual's preface. It is available as the *ASM and ASM-QFS README File*. After your software is installed, the content of the `README` file is located in `/opt/SUNWsamfs/doc/README`.

To Install From a CD-ROM

1. Log in as root.

   The ASM and ASM-QFS software uses the Sun Solaris operating environment (OE) packaging utilities for adding and removing software. You must be logged in as superuser (`root`) to make changes to software packages. The `pkgadd`(1M) utility prompts you to confirm various actions necessary to install the packages.

2. Insert the CD into the CD drive.

   The system should automatically detect the CD's presence. If it does not, issue commands to stop and start the Sun Solaris Volume Manager and to change to the directory that contains the ASM/QFS-Standalone software package.

```
# /etc/init.d/volmgt stop
# /etc/init.d/volmgt start
# volcheck
# cd /cdrom/cdrom0
```

   On the CD, the ASM/QFS-Standalone package resides in the `/cdrom/cdrom0` directory organized by Sun Solaris version.

Step 2: Backing Up Each ASM and ASM-QFS File System

If you do not have current backup files for each of your ASM and ASM-QFS file systems, create them now using the `samfsdump`(1M) command.

Perform this step as a precautionary measure if you suspect that your current `samfsdump`(1M) file is incorrect or outdated.

Note: You must back up your file system now if you plan to use the following ASM or ASM-QFS 4.0 features:

- Access Control Lists (ACLs)
- ASM/QFS-Standalone shared file system
- `md` devices in ASM/QFS-Standalone or ASM-QFS (`ma`) file systems
- Dual-sized disk allocation units (DAUs) on `mm` devices

In order to use any of these features, you need to reinitialize the file system. Reinitializing the file system is described in “Step 11: Reinitializing and Restoring the File Systems (Optional)” on page 114 of this installation process. After the file system is reinitialized, by using the `sammkfs`(1M) command, you can use the `samfsrestore`(1M) command to restore files to the new file system from the dump file created in this installation step.

To Back Up File Systems

1. Make sure that all files are archived.

   For example, assume that `sam1` is the mount point of the file system. You can complete this step by entering a command similar to the following:

```
# sfind /sam1 ! -archived
```

Examine this command’s output. The files in this output are those that have not been archived. If you want any of these files to appear in the dump file, archive them now before you proceed to the next step. Alternatively, you can use the `-u` option to the `samfsdump`(1M) command to dump unarchived data if you suspect some files have not yet been archived. The `-u` option can create very large dump files, however, so you need to weigh space considerations when using this option.

2. Back up each ASM or ASM-QFS file system’s metadata by using `samfsdump`(1M).

   The location to which you dump each file system’s metadata should be outside the ASM or ASM-QFS file system.

   The following example assumes that you have a file system mounted at `/sam1` that you want to back up to `samfs1.dump`, which exists outside of the ASM file systems:

```
# cd /sam1
# samfsdump -f /csd_dump_dir/samfs1.dump
```

The `samfsdump`(1M) command dumps file names and inode information, not data. For more information, see the `samfsdump`(1M) man page.

You must back up the metadata information for each file system, so repeat the preceding steps for each file system in your ASM and ASM environment.

For more information about backing up your file systems, see the *ASM, ASM-QFS, and ASM/QFS-Standalone Disaster Recovery Guide*.

Step 3: Stopping the ASM or ASM-QFS File System

To Stop the File System

1. Use the `samcmd`(1M) command to idle the drives in your system.

To idle the drives, enter the following for each removable media drive *eq* configured in your `mcf` file:

```
# samcmd idle eq
```

Enter the a `samcmd idle` command for each removable media drive *eq* configured in your `mcf` file.

You can also idle the drives by using the `samu`(1M) operator utility or by using either the `robottool`(1M) or `libmgr`(1M) graphical user interface (GUI) tools. For more information about the `samcmd`(1M) command, see the `samcmd`(1M) man page.

2. Stop the archiver.

Use an editor to open the `/etc/opt/SUNsamfs/archiver.cmd` file and add a `wait` directive near the top of this file. The purpose of the `wait` directive is to stop the archiver.

3. Issue the `samd`(1M) stop command to stop all operations.

Make sure to idle the drives in your ASM or ASM-QFS environment before you issue the `samd stop` command. This enables the archiver, stager, and other processes to complete current operations.

For example:

```
# samd stop
```

Step 4: Unsharing the File Systems (Optional)

You must complete this step if your ASM or ASM-QFS file systems are NFS shared file systems.

To Unshare the File Systems

- Use the `unshare`(1M) command on the ASM or ASM-QFS file system.

For example, the following command unshares the `samqfs1` file system:

```
# unshare samqfs1
```

■ Step 5: Unmounting the File Systems

There are several ways to unmount a file system. Any of the following methods can accomplish this task. The easiest method is presented first. After the file system is unmounted, you can proceed to“Step 6: Removing Existing ASM or ASM-QFS Software” on page 92.

To Unmount Using the umount(1M) Command

- Using the `umount`(1M) command, unmount each ASM or ASM-QFS file system.

  If you are unmounting a file system from a Sun Solaris 8 or later OE, you can use the `-f` option to the `umount`(1M) command. The `-f` option forces a file system to unmount.

To Unmount Using the fuser(1M), kill(1), and umount(1M) Commands

If using `umount`(1M) to unmount is not successful, it might be because you or another user are using files or because you or another user have changed to directories in the file system.

1. Use the `fuser`(1M) command to determine whether or not any processes are still busy.

For example, the following command queries the `samqfs1` file system:

```
# fuser -uc /samqfs1
```

2. If any processes are still busy, terminate them by using the `kill`(1M) command.
3. Using the `umount`(1M) command, unmount each ASM or ASM-QFS file system.

To Unmount by Editing the /etc/vfstab File and Rebooting

1. Edit the `/etc/vfstab` file.

   For all ASM or ASM-QFS file systems, change the Mount at Boot field from `yes` or `delay` to `no`.

2. Reboot your system.

Step 6: Removing Existing ASM or ASM-QFS Software

The following section describes how to remove software from a release prior to 4.0.

To Remove Software from a Release Prior to 4.0

1. Use the `pkginfo`(1) command, as follows, to determine which ASM and ASM-QFS software packages are installed on your system.

```
# pkginfo | grep LSC
```

2. Use the `pkgrm`(1M) command to remove the existing ASM or ASM-QFS software.

   You must remove all existing ASM and ASM-QFS packages before installing the new packages. If you are using any optional ASM or ASM-QFS packages, you should make sure that you remove these packages before you remove the main `LSCsamfs` package. The install script prompts you to confirm several of the removal steps.

The following example removes all of the ASM-FS and ASM-QFS 3.5.0 packages:

```
# pkgrm LSCibm LSCstk LSCdst LSCsony LSCgui LSCjre
LSCdoc LSCmigkit \LSCtools LSCremote LSCsamfs
```

The `LSCsamfs` package must be the last package removed.

As part of the installation process, the existing master configuration file and the catalog files are copied to `/etc/opt/SUNWsamfs/samfs.old.`*date*.

Note: If you are upgrading from an ASM-FS or ASM-QFS release prior to the ASM or ASM-QFS 4.0 release, you must remove any installed `LSCibm`, `LSCstk`, `LSCdst`, `LSCsony`, `LSCgui`, `LSCjre`, `LSCdoc`, `LSCmigkit`, `LSCtools` or `LSCremote` packages *prior to* removing the `LSCsamfs` package. The new ASM or ASM-QFS 4.0 package

consists of the SUNWsamfs package, which includes most of the packages distributed individually prior to the 4.0 release. The only package that is still available separately is the tools package, which is now available as the SUNWsamtp package. Contact your StorageTek sales representative for information on obtaining the 4.0 version of the SUNWsamtp package.

Step 7: Adding the Packages

The ASM and ASM-QFS software packages use the Sun Solaris packaging utilities for adding and deleting software. You must be logged in as superuser (root) to make changes to software packages. The pkgadd(1M) command prompts you to confirm various actions necessary to upgrade the packages.

During the installation, the system detects the presence of conflicting files and prompts you to indicate whether or not you want to continue with the installation. You can go to another window and copy the files you want to save to an alternate location.

On the CD-ROM, all products reside in the /cdrom/cdrom0 directory organized by Sun Solaris version.

To Add the Packages

1. Run the pkgadd(1M) command to upgrade the SUNWsamfs package.

   The format of this command is as follows:

```
# pkgadd -d SUNWsamfs
```

2. Answer yes to each question.
3. Use the pkgadd(1M) command to add one or more localized packages. (Optional)

   Perform this step only if you want to install the packages localized for Chinese, French, or Japanese. To install these packages, enter one or more of the following commands:

```
# pkgadd -d SUNWcsamf
# pkgadd -d SUNWfsamf
# pkgadd -d SUNWjsamf
```

Step 8: Updating the License Keys

License keys are required to run the ASM and ASM-QFS software. For information on license keys, see “Step 6: Verifying the Software License” on page 21.

The ASM and ASM-QFS software uses encrypted license keys. The license keys consist of encoded alphanumeric strings. You receive one or more license keys depending on the system configuration and the products being licensed.

To License the Software

1. Verify whether the license file exists.

   The license file is as follows:

`/etc/opt/SUNWsamfs/LICENSE.4.0`

2. If the `/etc/opt/SUNWsamfs/LICENSE.4.0` file does not exist, create it.
3. Starting in column one, place the license key you have obtained from StorageTek on the first line in the `/etc/opt/SUNWsamfs/LICENSE.4.0` file.

The key must start in column one. No other keywords, host IDs, comments, or other information can appear in the `/etc/opt/SUNWsamfs/LICENSE.4.0` file. The license becomes effective when a ASM or ASM-QFS file system is mounted.

4. View the current license settings by issuing the `samcmd`(1M) command with its `l` argument, as follows.

```
# samcmd l
```

This is the `l` (letter “ell” as in *license*) argument.

The license keys allow the system to run indefinitely unless you were issued a temporary license.

■ Step 9: Verifying System Files

The topology of the equipment managed by the ASM or ASM-QFS file system is defined in the master configuration file, `/etc/opt/SUNWsamfs/mcf`. This file specifies the devices, automated libraries, and file systems included in the environment. Each piece of equipment is assigned a unique Equipment Identifier in the `mcf` file.

Note: For information about file system design considerations, see the *ASM, ASM-QFS, and ASM/QFS-Standalone File System Administrator’s Guide*.

To Verify System Files

If tracing is implemented in your `archiver.cmd` file or `stager.cmd` file, edit these files now to remove the tracing directives. All system tracing in ASM and

ASM-QFS 4.0 is implemented in the `defaults.conf` file. The presence of tracing directives in the `archiver.cmd` file or `stager.cmd` file on a 4.0 system will inhibit archiving and staging. For more information about these files, see the `archiver.cmd`(4), `stager.cmd`(4), and `defaults.conf`(4) man pages.

1. Edit your `archiver.cmd` and `stager.cmd` files to identify and then remove the `trace =` directives.
2. Edit your `defaults.conf` file to reenable tracing.

   For information on how to enable tracing, see the `defaults.conf`(4) man page.

To Verify the mcf File

1. Issue the `cd`(1) command to change to the `/etc/opt/SUNWsamfs` directory.

   This is the directory that contains your `mcf` file.
2. Verify that an `mcf` file exists.

   The correct location for this file is as follows:

   `/etc/opt/SUNWsamfs/mcf`
3. Issue the `sam-fsd`(1M) command to check the `mcf` file for errors.

For example:

```
# sam-fsd
```

To Reinitialize the mcf File

Reinitialization of the `mcf` file is not needed at this point in the configuration process. Be aware, however, that if you change the `mcf` file after the ASM or ASM-QFS software is in use, you must enter commands to convey the new `mcf` specifications to the system. For information on reinitializing the `mcf` file, see the *ASM, ASM-QFS, and ASM/QFS-Standalone File System Administrator's Guide*.

■ Step 10: Modifying the /etc/vfstab File (Optional)

If you modified the `/etc/vfstab` file in "Step 5: Unmounting the File Systems" on page 92, you must complete this step.

To Modify the /etc/vfstab File

- Edit this file again, and change the Mount at Boot field for all ASM and ASM-QFS file systems from `no` to `yes`.

Step 11: Reinitializing and Restoring the File Systems (Optional)

You must use the ASM or ASM-QFS 4.0 `sammkfs`(1M) command to reinitialize your file systems if you want to use all of the ASM/QFS-Standalone 4.0 features. The features that require you to reinitialize your file systems are noted in “Step 2: Backing Up Each ASM and ASM-QFS File System” on page 90.

In this step, you reinitialize your file systems and restore the saved data from “Step 2: Backing Up Each ASM and ASM-QFS File System” on page 90 into the new file systems. To accomplish this, use the `sammkfs`(1M) and `samfsrestore`(1M) commands on each file system.

To Reinitialize and Restore the File Systems

1. Use the `sammkfs`(1M) command to initialize a new ASM or ASM-QFS file system.

If you want to use all the ASM/QFS-Standalone 4.0 features, issue the `sammkfs`(1M) command without any options. The following example `sammkfs`(1M) command reinitializes a file system named `samqfs1` with ASM-QFS 4.0 feature capabilities.

```
# sammkfs samqfs1
```

For more information on the options to the `sammkfs`(1M) command, see the `sammkfs`(1M) man page.

2. Use the `samfsrestore`(1M) command to restore the dumped data into the new file system.

For example, the following commands assume that you have a file system named `samqfs1` (mounted at `/samqfs1`) that you want to back up from files dumped to `samqfs1.bak`, which exists outside of the ASM/QFS-Standalone file system:

```
# cd /samqfs1
# samfsrestore -f /save/qfs/samqfs1.bak
```

Note: If you do not want to use all of the ASM/QFS-Standalone 4.0 features, use the `-P` option to the `sammkfs`(1M) command. This creates a

version 1 superblock. For more information on the options to the `sammkfs`(1M) command, see the `sammkfs`(1M) man page.

Step 12: Checking the File System (Optional)

If you did not perform “Step 11: Reinitializing and Restoring the File Systems (Optional)” on page 114, you are encouraged to complete this step.

- Use the `samfsck`(1M) command to check your existing file systems for inconsistancies.

Do this for each ASM-FS and ASM-QFS file system.

Step 13: Mounting the File Systems (Optional)

If you have *not* modified the Mount at Boot field in the `/etc/vfstab` file to say `yes` or `delay` instead of `no`, you must perform this step.

- Use the `mount`(1M) command to mount the file systems and continue operation with the upgraded software.

  In the following example, `samqfs1` is the name of the file system to be mounted:

```
# mount samqfs1
```

Step 14: Verifying the archiver.cmd File

- Use the `archiver`(1M) command with its `-l` and `-v` options to verify the validity of the `archiver.cmd` file.

For example:

```
# archiver -lv
```

If this command reports errors in the `archiver.cmd` file, correct them now. For more information on the `archiver`(1M) command, see the `sam-archiverd`(1M) man page.

Step 15: Recompiling API-Dependent Applications (Optional)

If you are running applications that use the ASM or ASM-QFS application programming interface (API), you need to complete this step.

Because file headers, the calling sequence, and other elements of the API can change from release to release, you should recompile all applications that depend on the API at this time.

Caution: Failure to recompile API-dependent applications at this point can cause your applications to generate unexpected results.

Glossary

A

addressable storage

The storage space encompassing online, nearline, offsite, and offline storage that is user-referenced through an ASM/QFS-Standalone, ASM, or ASM-QFS file system.

archive media

The media to which an archive file is written. Archive media can be removable tape or magneto-optical cartridges in a library. In addition, archive media can be a mount point on another system.

ASM

The Application Storage Manager™. The ASM software controls the access to all files stored and all devices configured in the master configuration file (`mcf`).

ASM-QFS

The ASM-QFS software combines the Application Storage Manager™ with the ASM/QFS-Standalone file system. ASM-QFS offers a high-speed, standard UNIX file system interface to users and administrators in conjunction with the storage and archive management utilities. It uses many of the commands available in the ASM command set as well as standard UNIX file system commands.

ASM-Remote client

An ASM-Remote client is an ASM or ASM-QFS system that establishes an ASM-Remote client daemon that contains a number of pseudodevices. It might or might not have its own library devices. The client depends on an ASM-Remote server for archive media for one or more archive copies.

ASM-Remote server

The ASM-Remote server is both a full-capacity ASM or ASM-QFS storage management server and an ASM-Remote server daemon that defines libraries to be shared among ASM-Remote clients.

archive storage

Copies of file data that have been created on archive media.

archiver

The archive program that automatically controls the copying of files to removable cartridges.

audit (full)

The process of loading cartridges to verify their VSNs. For magneto-optical cartridges, the capacity and space information is determined and entered into the automated library's catalog.

automated library

A robotically controlled device designed to automatically load and unload removable media cartridges without operator intervention. An automated library contains one or more drives and a transport mechanism that moves cartridges to and from the storage slots and the drives.

B

backup storage

A snapshot of a collection of files for the purpose of preventing inadvertent loss. A

backup includes both the file's attributes and associated data.

block allocation map

A bitmap representing each available block of storage on a disk and indicating whether the block is in use or free.

block size

See DAU.

C

cartridge

A physical entity that contains media for recording data. A tape or optical disk. Sometimes referred to as *a piece of media*, *a volume*, or *the medium*.

catalog

A record of the VSNs in an automated library. There is one catalog for each automated library, and at a site, there is one historian for all automated libraries.

client-server

The model of interaction in a distributed system in which a program at one site sends a request to a program at another site and awaits a response. The requesting program is called the client. The program satisfying the response is called the server.

connection

The path between two protocol modules that provides reliable stream delivery service. A TCP connection extends from a TCP module on one machine to a TCP module on the other.

D

data device

For an ASM/QFS-Standalone, ASM, or ASM-QFS file system, a device or group of devices upon which file data is stored.

DAU (disk allocation unit)

The basic unit of online storage. Also called block size.

The ASM and ASM-QFS file systems support both a small and a large DAU. The small DAU is 4 kilobytes (2^{14} or 4096 bytes). The large DAU is 16, 32, or 64 kilobytes. The available DAU size pairs are 4/16, 4/32, and 4/64.

In addition, the ASM/QFS-Standalone and ASM-QFS file systems support a fully adjustable DAU, sized from 16 kilobytes through 65,528 kilobytes. The DAU you specify must be a multiple of 8 kilobytes.

device logging A configurable feature that provides device-specific error information used to analyze device problems.

device scanner

Software within the ASM or ASM-QFS file system that periodically monitors the presence of all manually mounted removable devices and that detects the presence of mounted cartridges that can be requested by a user or other process.

direct access

A file attribute (stage never) designating that a nearline file can be accessed directly from the archive media and need not be retrieved to disk cache.

direct-attached library

An automated library connected directly to a server using a SCSI interface. A SCSI attached library is controlled directly by the ASM or ASM-QFS software by using the SCSI standard for automated libraries.

direct I/O

An attribute used for large block-aligned sequential I/O. The `setfa`(1) command's `-D` option is the direct I/O option. It sets the direct I/O attribute for a file or directory. If applied to a directory, the direct I/O attribute is inherited.

directory

A file data structure that points to other files and directories within the file system.

disk allocation unit

See DAU.

disk buffer

When using ASM-Remote software, the disk buffer is a buffer on the server system that is used when archiving data from the client to the server.

disk cache

The disk-resident portion of the ASM and ASM-QFS file system software. It is used to create and manage data files between online disk cache and archive media. Individual disk partitions or an entire disk can be used as disk cache.

disk space thresholds

An administrator-defined amount of disk space that is available to a user. This defines the range of desirable disk cache utilization. The high threshold indicates the maximum level of disk cache utilization. The low threshold indicates the minimum level of disk cache utilization. The releaser controls disk cache utilization based on these predefined disk space thresholds.

disk striping

The process of recording a file across several disks, thereby improving access performance and increasing overall storage capacity. Also see entries for striping.

drive

A mechanism for transferring data to and from a removable media volume.

E

Ethernet

A local-area, packet-switched network technology. Originally designed for coaxial cable, it is now found running over shielded, twisted-pair cable. Ethernet is a 10- or 100-megabytes-per-second LAN.

extent array

The array within a file's inode that defines where each data block assigned to the file is located on the disk.

F

family device set

See family set.

family set

A storage device that is represented by a group of independent physical devices, such as a collection of disks or the drives within an automated library. Also see disk cache family set.

FDDI

Fiber distributed data interface. A 100-megabytes-per-second fiber-optic LAN.

fibre channel

The ANSI standard that specifies high-speed serial communication between devices. Fibre channel is used as one of the bus architectures in SCSI-3.

fibre-distributed data interface

See FDDI.

file system

A hierarchical collection of files and directories.

file system specific directives

Archiver and releaser directives that follow global directives, are specific to a particular file system, and begin with `fs =`. File system specific directives apply until the next `fs =` directive line or until the end of file is encountered. If multiple directives affect a file system, the file system-specific directives override the global directives.

FTP

File Transfer Protocol. An internet protocol for transferring files between two hosts over a TCP/IP network.

G

global directives

Archiver and releaser directives that apply to all file systems and that appear before the first `fs =` line.

grace period

For disk quotas, this is the amount of time that can elapse during which a user is allowed to create files and/or allocate storage after a user reaches their soft limit.

H

hard limit

For disk quotas, a maximum limit on file system resources (blocks and inodes) that users cannot exceed.

I

indirect block

A disk block that contains a list of storage blocks. The ASM/QFS-Standalone, ASM, and ASM-QFS file systems have up to three levels of indirect blocks. A first-level indirect block contains a list of blocks used for data storage. A second-level indirect block contains a list of first-level indirect blocks. A third-level indirect block contains a list of second-level indirect blocks.

inode

Index node. A data structure used by the file system to describe a file. An inode describes all the attributes associated with a file other than the name. The attributes include ownership, access, permission, size, and the file location on the disk system.

inode file

A special file (`.inodes`) on the file system that contains the inode structures for all files resident in the file system. All ASM/QFS-Standalone, ASM, and ASM-QFS inodes are 512 bytes long. The inode file is a metadata file, which is separated from file data in the ASM/QFS-Standalone and ASM-QFS file systems.

K

kernel

The central controlling program that provides basic system facilities. The UNIX kernel creates and manages processes, provides functions to access the file system, provides general security, and supplies communication facilities.

L

LAN

Local area network.

lease

In a ASM/QFS-Standalone shared file system, a lease grants a client host permission to perform an operation on a file for as long as the lease is valid. The metadata server issues leases to each client host. The leases are renewed as necessary to permit continued file operations.

library

See automated library.

library catalog

See catalog.

LUN

Logical unit number.

M

mcf

Master configuration file. The file that is read at initialization time that defines the relationships between the devices (the topology) within a ASM/QFS-Standalone, ASM, and ASM-QFS environment.

media

Tape or optical disk cartridges.

media recycling

The process of recycling or reusing archive media with low use (that is, archive media with few active files).

metadata

Data about data. Metadata is the index information needed to locate the exact data position of a file on a disk. It consists of information about files, directories, access control lists, symbolic links, removable media, segmented files, and the indexes of segmented files. Metadata must be protected because if data is lost, the metadata that locates the data must be restored before the lost data can be retrieved.

metadata device

A separate device (for example, a solid-state disk or mirrored device) upon which ASM/QFS-Standalone and ASM-QFS file system metadata is stored. Separating file data from metadata can increase performance. In the `mcf` file, a metadata device is declared as an `mm` device within an `ma` file system.

mirror writing

The process of maintaining two copies of a file on disjointed sets of disks to prevent loss from a single disk failure.

mount point

The directory on which a file system is mounted.

multireader file system

The ASM/QFS-Standalone multireader file system is a single-writer, multireader capability that enables you to specify a file system that can be mounted on multiple hosts. Multiple hosts can read the file system, but only one host can write to the file system. Multiple readers are specified with the `-o reader` option on the `mount`(1M) command. The single-writer host is specified with the `-o writer` option on the `mount`(1M) command. For more information on the `mount`(1M) command, see the `mount_samfs`(1M) man page.

N

name space

The metadata portion of a collection of files that identifies the file, its attributes, and its storage locations.

nearline storage

Removable media storage that requires robotic mounting before it can be accessed. Nearline storage is usually less expensive than online storage, but it incurs a somewhat longer access time.

network-attached automated library

A library, such as those from StorageTek, ADIC/Grau, IBM, or Sony, that is controlled using a software package supplied by the vendor. The ASM and ASM-QFS file

systems interface with the vendor software using an ASM or ASM-QFS media changer daemon designed specifically for the automated library.

NFS

Network file system. An ASM distributed file system that provides transparent access to remote file systems on heterogeneous networks.

NIS

The SunOS 4.0 (minimum) Network Information Service. A distributed network database containing key information about the systems and the users on the network. The NIS database is stored on the master server and all the slave servers.

O

offline storage

Storage that requires operator intervention for loading.

offsite storage

Storage that is remote from the server and is used for disaster recovery.

online storage

Storage that is immediately available (for example, disk cache storage).

P

partition

A portion of a device or a side of a magneto-optical cartridge.

preallocation

The process of reserving a contiguous amount of space on the disk cache for writing a file. This ensures that the space is contiguous. Preallocation can be performed only on zero-sized files. That is, the `setfa -l` command can be specified only for a file that is size zero. For more information, see the `setfa`(1) man page.

prioritizing preview requests

Assigning priority to archive and stage requests that cannot be immediately satisfied.

pseudo device

A software subsystem or driver with no associated hardware.

Q

quota

The amount of system resources that a user is allowed to consume. Quotas are not supported for removable media or disk archive resources.

R

RAID

Redundant array of inexpensive/ independent disks. A disk technology that uses several independent disks to reliably store files. It can protect against data loss from a single disk failure, can provide a fault-tolerant disk environment, and can provide higher throughput than individual disks.

recycler

An ASM and ASM-QFS utility that reclaims space on cartridges that is occupied by expired archive copies.

release priority

A method of calculating the release priority of a file within a file system by multiplying various weights by the corresponding file properties and then summing the results.

releaser

An ASM and ASM-QFS component that identifies archived files and releases their disk cache copies, thus making more disk cache space available. The releaser automatically regulates the amount of online disk storage to high and low thresholds.

remote procedure calls

See RPC.

removable media file

A special type of user file that can be accessed directly from where it resides on a removable media cartridge, such as magnetic tape or optical disk cartridge. also used for writing archive and stage file data.

robot

The portion of an automated library that moves cartridges between storage slots and drives. Also called a transport.

round robin

A data access method in which entire files are written to logical disks in a sequential fashion. When a single file is written to disk, the entire file is written to the first logical disk. The second file is written to the next logical disk, and so on. The size of each file determines the size of the I/O.

By default, ASM/QFS-Standalone, ASM, and ASM-QFS file systems implement striped data access unless striped groups are present. Files are round-robined if round robin access is specified. If the file system contains mismatched striped groups, striping is not supported and round robin is forced.

Also see glossary entries for disk striping and striping.

RPC

Remote procedure calls. The underlying data exchange mechanism used by NFS to implement custom network data servers.

S

samfsdump

A program that creates a control structure dump and copies all the control structure information for a given group of files. It is analogous to the UNIX `tar`(1) utility, but it does not generally copy file data.

samfsrestore

A program that restores inode and directory information from a control structure dump.

SCSI

Small Computer System Interface. An electrical communication specification commonly used for peripheral devices such as disk and tape drives and automated libraries.

small computer system interface

See SCSI.

soft limit

For disk quotas, a threshold limit on file system resources (blocks and inodes) that you can temporarily exceed. Exceeding the soft limit starts a timer. When you exceed the soft limit for the specified time (default is one week), no further system resources can be allocated until you reduce file system use to a level below the soft limit.

staging

The process of copying a nearline or offline file from archive storage back to online storage.

storage family set

A set of disks that are collectively represented by a single disk family device.

storage slots

Locations inside an automated library in which cartridges are stored when not being used in a drive. If the library is direct-

attached, the contents of the storage slots are kept in the automated library's catalog.

stripe size

The number of disk allocation units (DAUs) to allocate before moving to the next device of a stripe. If `stripe=0`, the file system uses round-robin access, not striped access.

striped group

A collection of devices within an ASM/QFS-Standalone or ASM-QFS file system and defined in the `mcf` file as one (usually two) or more `g`*XXX* devices. Striped groups are treated as one logical device and are always striped with a size equal to the disk allocation unit (DAU). You can specify up to 128 striped groups within a file system, but you can specify no more than 252 total devices.

striping

A data access method in which files are simultaneously written to logical disks in an interlaced fashion. All ASM/QFS-Standalone, ASM, and ASM-QFS file systems enable you to declare either striped or round robin access for each individual file system. The ASM/QFS-Standalone and ASM-QFS file systems enable you to declare striped groups within each file system. Also see the glossary entry for round robin.

superblock

A data structure in the file system that defines the basic parameters of the file system. It is written to all partitions in the storage family set and identifies the partition's membership in the set.

T

tar

Tape archive. A standard file/data recording format used by the ASM and ASM-QFS software for archive images.

TCP/IP

Transmission Control Protocol/Internet Protocol. The internet protocols responsible for host-to-host addressing and routing, packet delivery (IP), and reliable delivery of data between application points (TCP).

thresholds

A mechanism for defining the desirable available storage window for online storage. Thresholds set the storage goals for the releaser. Also see disk space thresholds.

timer

Quota software that keeps track of the time elapsed between a user reaching a soft limit and a hard limit being imposed on the user.

V

volume

A named area on a cartridge for sharing data. A cartridge has one or more volumes. Double-sided cartridges have two volumes, one on each side.

volume overflow

A capability that enables the system to span a single file over multiple volumes. Volume overflow is useful for sites using very large files that exceed the capacity of their individual cartridges.

VSN

Volume serial name. If you are archiving to removable media cartridges, the VSN is a logical identifier for magnetic tape and optical disk that is written in the volume label. If you are archiving to disk cache, this is the unique name for the disk archive set.

W

WORM

Write once read many. A storage classification for media that can be written only once but read many times.

Index

Symbols

A

B

C

D

E

F

G

H

I

J

K

L

M

N

O

P

Q

R

S

T

U

V

Reader’s Comment Form

Contact Us

Submit your questions, comments, and suggestions to StorageTek’s Global Learning Solutions. We appreciate your correspondence and are committed to responding to you.

Publication Information

Publication Name:

Publication Part Number:

Questions and Comments:

Note: Staples can cause problems with automated mail sorting equipment. Please use pressure sensitive or other gummed tape to seal this form. If you would like a reply, please supply your name and address on the reverse side of this form.

Thank you for your cooperation. No postage stamp is required if mailed in the U.S.A.

TO COMPLY WITH POSTAL REGULMATIONS, FOLD EXACTLY ON DOTTED LINES AND TAPE (DO NOT

NO POSTAGE
NECESSARY
IF MAILED
IN THE
UNITED STATES

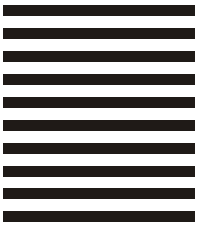

BUSINESS REPLY CARD

FIRST CLASS PERMIT NO. 2 LOUISVILLE, CO U.S.A.

POSTAGE WILL BE PAID BY ADDRESSEE

GLOBAL LEARNING SOLUTIONS MS 3256
STORAGE TECHNOLOGY CORPORATION
ONE STORAGETEK DRIVE
LOUISVILLE CO 80028-9989
USA

FOLD HERE AND TAPE DO NOT STAPLE FOLD HERE AND TAPE

If you would like a reply, please print:

Your Name:_____

Company Name:_____ Department:_____

Street Address:_____

City:_____

State:_____ Zip Code:_____

Storage Technology Corporation
One StorageTek Drive
Louisville, CO 80028-3256
USA

NEED MORE INFORMATION?

www.storagetek.com

ABOUT STORAGETEK®

StorageTek® (NYSE:STK), a $2 billion world-wide company with headquarters in Louisville, Colo., delivers a broad range of storage solutions for digitized data. StorageTek solutions are easy to manage and allow universal access to data across servers, media types and storage networks. StorageTek is the innovator and global leader in virtual storage solutions for tape automation, disk storage systems and storage networking. Because of StorageTek, customers can manage and leverage their digital assets as their businesses grow, and can maximize IT productivity to ensure enterprise-class business continuity.

WORLD HEADQUARTERS
Storage Technology Corporation
One StorageTek Drive
Louisville, Colorado 80028 USA
Phone: **1.800.525.0369**

StorageTek®