Application Storage Manager™ (ASM) for Unix

Installation and Configuration Guide

Version 3.5.0

Part Number 313498501

Proprietary Information Statement

The information in this document is confidential and proprietary to Storage Technology Corporation and may be used only under the terms of the product license or nondisclosure agreement. The information in this document, including any associated software program, may not be disclosed, disseminated, or distributed in any manner without the written consent of Storage Technology Corporation.

Limitations on Warranties and Liability

This document neither extends nor creates warranties of any nature, expressed or implied. Storage Technology Corporation cannot accept any responsibility for your use of the information in this document or for your use of any associated software program. You are responsible for backing up your data. You should be careful to ensure that your use of the information complies with all applicable laws, rules, and regulations of the jurisdictions in which it is used**. Warning**: No part or portion of this document may be reproduced in any manner or in any form without the written permission of Storage Technology Corporation.

Restricted Rights

Use, duplication, or disclosure by the U.S. Government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.227-7013 or subparagraphs (c)(1) and (2) of the Commercial Computer Software – Restricted Rights at 48 CFR 52.227-19, as applicable.

Application Storage Manager™ (ASM) for Unix Installation and Configuration Guide, Version 3.5.0 Edition 2, June 3, 2002, Part Number 313498501

This edition applies to the Application Storage Manager™ (ASM) for Unix product and to all modifications of that product until otherwise indicated in new editions or revisions pages. If there are changes in the product or improvements in the information about the product, this document will be revised and reissued.

Comments concerning the contents of the manual should be directed to:

ASM Product Management
Storage Technology Corporation
One StorageTek Drive, MS 2139
Louisville, CO 80028-2138

New Features

The *ASM Installation and Configuration Guide*, revision 3.5.0-20, supports the ASM-QFS, ASM, and ASM-QFS 3.5.0-20 releases running on the Solaris 2.6, 2.7, and 2.8 platforms. This manual was derived from the *ASM System Administrator's Guide*, revision 3.5.0, and from the *ASM-QFS Standalone Administrator's Guide*, revision 3.4.

The ASM for Unix v3.5.0-20 releases support the following new features in the ASM-QFS Standalone and ASM-QFS environments:

- Solaris 2.8 operating system. StorageTek products can now be installed on servers running the Solaris 2.8 operating system.
- ASM-QFS SHARE file system. The ASM-QFS Standalone file system can be used in conjunction with fiber-attached devices in a Storage Area Network (SAN). When the ASM-QFS SHARE license key from StorageTek is enabled, the ASM-QFS Standalone file system enables high-speed access to data using software such as Tivoli SANergy File Sharing.
- ACSLS 5.4. The ASM and ASM-QFS environments now support the StorageTek ACSLS 5.4 release for network-attached StorageTek automated libraries.
- New devices. The ASM and ASM-QFS environments now support the following devices:

Device	Comments
Exabyte Mammoth-2 drive	Equipment type xm in your mcf file.
Fujitsu M8100 drive	Equipment type fd in your mcf file.
Sony DTF-2 drive	Equipment type so in your mcf file.
Sony network-attached automated library and drives	Sites with Sony network-attached automated libraries use the Sony DZC-8000S interface. Such automated libraries are equipment type pe in your mcf file. The drives are equipment type so in your mcf file. Sites with automated libraries attached through this interface need to add the LSCsony (samsony) package at installation time.

StorageTek 9940 drive	Equipment type sf in your mcf file.

Documentation restructuring. The StorageTek manual set has been restructured in order to make the documentation more modular. Certain parts of the *ASM System Administrator's Guide* have been moved into new manuals or man pages.

The following table indicates the topic that was affected by this restructuring, where it used to reside in the StorageTek documentation set prior to the 3.5.0-20 release, and where it resides in the documentation set as of the 3.5.0 release:

Topic	Pre-3.5.0-20 Location	Post-3.5.0-20 Location
ASM-QFS Standalone installation and configuration.	ASM-QFS Standalone Administrator's Guide.	ASM Installation and Configuration Guide.
ASM installation and configuration.	ASM System Administrator's Guide, chapter 2.	ASM Installation and Configuration Guide.
ASM-QFS Standalone file system overview, reference, and operations information.	ASM-QFS Standalone Administrator's Guide.	StorageTek File System Administrator's Guide.
ASM file system overview, reference, and operations information	*ASM System Administrator's Guide*, publication SG-0001, chapter 5 and part of chapter 14.	StorageTek File System Administrator's Guide.
Application Programmer Interface (API) overview	ASM System Administrator's Guide, chapter 12.	intro_libsam(3) man page.

Topic	Pre-3.5.0-20 Location	Post-3.5.0-20 Location
Storage and archive management operational, reference, and disaster recovery information	*ASM System Administrator's Guide*, chapters 3, 4, 6, 7, 8, 9, 10, 11, 13, and 14; appendixes A and B.	ASM Administrator's Guide

Record of Revision

Revision	Description
3.3	June 1998. Original printing.
3.5.0	July 2001. This rewrite supports version 3.5.0-20 and later releases.
3.5.0	July 2002. Document update.

Table of Contents

Preface

This manual describes the installation and upgrade procedures for the ASM-QFS Standalone, ASM, and ASM-QFS software products, release 3.5.0-20, running on the Solaris 2.6, 2.7, and 2.8 operating system platforms. It is written for system administrators responsible for setting up and maintaining StorageTek software. You, the system administrator, are assumed to be knowledgeable about Solaris operating system procedures, including creating accounts, performing system backups, and other basic Solaris system administrator tasks.

Other StorageTek software products, such as ASM Remote, can be licensed for use within the ASM-QFS Standalone, ASM, and ASM-QFS environments. For more information on these products, see the Licensing subsection in this preface.

This manual describes how to install, upgrade, and configure the ASM-QFS Standalone, ASM, and ASM-QFS release packages. It is organized as follows:

Section	Title
Chapter 1	Overview
Chapter 2	System requirements
Chapter 3	ASM-QFS Standalone initial installation procedure
Chapter 4	ASM-QFS Standalone upgrade procedure
Chapter 5	ASM and ASM-QFS initial installation procedure
Chapter 6	ASM and ASM-QFS upgrade procedure
Appendix	StorageTek Product Support

In addition to the preceding sections, the glossary section defines terms used in StorageTek documentation.

Conventions

The following conventions and terms are used throughout this manual:

Convention	Meaning
Courier	The fixed-space courier font denotes literal items such as commands, files, path names, system prompts, system output, and messages. For example: /etc/opt/LSCsamfs/mcf
Bold courier	The **bold courier** font denotes text you enter at the shell prompt. For example: server**# sls –D**
[ ]	Brackets enclose optional portions of commands or optional arguments to commands.
Italic	Italics indicate either a variable or a term being defined. For a variable, you must replace the variable with a real name or value. For example: server# **mount** *mnt_pt*
|	The pipe symbol indicates that one of two or more optional arguments must be specified.

Certain terms are used throughout this manual. Many terms can be found in the glossary, but some of the most commonly used ones are as follows:

Term	Meaning
Archiving	Automatically copying online, magnetic disk cache files to archive media.
Automated library	An automated device for storing tape and optical cartridges.
Cartridge	A tape or magneto optical cartridge.
Partition	A side of a magneto optical disk or a partition on an Ampex tape.
Staging	Automatically copying files located on archive media back to online disk.
Volume	A named area on a cartridge for storing data. A cartridge has one or more volumes. Double-sided cartridges have two volumes, one on each side.

Other StorageTek Publications

In addition to this manual, the following StorageTek publications might be useful to you:

- ASM Migration Toolkit Guide

- ASM Remote Administrator's Guide
- ASM Man Page Reference Manual
- ASM File System Administrator's Guide
- ASM Administrator's Guide

To order additional manuals, please send us a request using one of the methods described in the "Reader Comments" subsection.

Other File System Publications

In addition to publications from StorageTek, the following publications on UNIX file systems might interest you:

- Filesystem Hierarchy Standard (FHS) web pages at the following URL: http://www.pathname.com/fhs/2.0/fhs-toc.html
- StorageTek online documentation web pages at the following URL: http://www.storagetek.com. Go to the Customer Resource Center (CRC) in the Services and Support section.

Licensing

Licenses for StorageTek products can be obtained from StorageTek. In some cases, the capabilities that these additional licenses can provide are described in this document because these products can be used in conjunction with this product. For information on obtaining licenses for StorageTek software, contact your sales representative, your Authorized Service Provider (ASP), or StorageTek.

The following StorageTek products are licensed separately:

- Migration Toolkit
- ASM-QFS Standalone
- ASM
- ASM-QFS
- ASM Remote client
- ASM Remote server
- ASM Segment
- ASM-QFS Share

Reader Comments

If you have comments about the technical accuracy, content, or organization of this document, please let us know. We value your comments and will respond to them promptly. You can contact us by one of the following methods:

- Send email to info@storagetek.com
- Send a facsimile (FAX) with your comments to ASM for Unix Product Management, fax number 303-661-7949.
- Send written comments to the following address:

ASM Product Management
StorageTek
One StorageTek Drive, MS 2138
Louisville, CO 80028-2138
USA

Chapter 1 - Overview

Before installing an ASM software package, you must set up and configure the hardware to be used. The environment typically consists of the following:

- A server based on Sparc technology running the 2.6, 2.7, or 2.8 Solaris operating system.
- A disk cache consisting of RAID devices, SCSI disks, or both.
- One or more automated libraries with one or more tape or optical drives. This requirement applies to the ASM and ASM-QFS release packages. This requirement does not apply to the QFS release package.

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

- Release package contents
- Directories and files created

Release Package Contents

The ASM software packages are made available via anonymous FTP and CD-ROM in Solaris pkgadd(1M) format. These packages reflect the Solaris version for the platform upon which you will be installing ASM software. For information on the directories and files provided with the release package, see the README and CHANGES files provided with the software.

The ASM software releases include the following release packages:

Package To Install	Installed Package	Description	Used By (Licensed Product)
samqfs	LSCqfs	ASM QFS Standalone software package. This package is not needed if you are installing the ASM-QFS software package.	ASM-QFS
sampkg	LSCsamfs	Application Storage Manager (ASM) software package. This package must be installed prior to all other packages if you are installing the ASM or ASM-QFS software.	ASM, ASM-QFS

Package To Install	Installed Package	Description	Used By (Licensed Product)
samgui	LSCgui	Graphical User Interface (GUI) tools based on Java technology. (optional)	ASM, ASM-QFS
samjre	LSCjre	Java runtime environment. This optional package needs to be installed only if you want to use the GUI tools. (optional)	ASM, ASM-QFS
samibm	LSCibm	IBM 3494 automated library package. (optional)	ASM, ASM-QFS
samstk	LSCstk	StorageTek (ACSLS) automated library package. (optional)	ASM, ASM-QFS
samdst	LSCdst	Ampex tape drive package. (optional)	ASM, ASM-QFS
samsony	LSCsony	Sony network-attached automated library support package. (optional)	ASM, ASM-QFS
samdoc	LSCdoc	ASM documentation in PDF format. (optional)	ASM, ASM-QFS

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

platform.major.minor-bugfix

platform	The leftmost and platform release number.
major	The release level of a major feature release.
minor	The release level of a minor feature release.
bugfix	The bugfix number. A number of 1 through 99 indicates a bugfix release. A letter from A through Z indicates a beta release. The base release of a first feature release of a major release might not contain a bugfix level.

Examples:

3.3.1 is a minor release.
3.3.1-4 is bugfix 4 of the 3.3.1 minor release.
3.5.0 is a major release with no minor release revisions and no bugfixes.
3.5.0-3 is bugfix 3 of the 3.5.0 major release.

Directories and Files Created

This subsection describes the directories and files associated with ASM products. Additional information about the files in this subsection can be obtained from the man pages after they have been installed.

Directories Created

The following table lists the directories created when ASM products are installed:

Directory	Content	Used By
/dev/samst	Device driver special files.	ASM, ASM-QFS
/var/opt/LSCsamfs	Device catalogs, the catalog trace file, and log files.	ASM, ASM-QFS
/etc/opt/LSCsamfs	Configuration files.	ASM, ASM-QFS
/opt/LSCsamfs/bin	User command binaries.	ASM, ASM-QFS
/opt/LSCsamfs/client	Files for RPC API client.	ASM, ASM-QFS
/opt/LSCsamfs/examples	Various example configuration files.	ASM, ASM-QFS

/opt/LSCsamfs/jre	Java runtime environment for the GUI tools.	ASM, ASM-QFS

Directory	Content	Used By
/opt/LSCsamfs/include	API include files.	ASM, ASM-QFS
/opt/LSCsamfs/lib	Relocatable libraries.	ASM, ASM-QFS
/opt/LSCsamfs/man	man(1) pages.	ASM, ASM-QFS
/opt/LSCsamfs/sbin	System administrator commands and daemon binaries.	ASM, ASM-QFS
/opt/LSCsamfs/doc	Contains the CHANGES file, the README file, the README_PUBS file, and all ASM manuals. The ASM manuals are in PDF format. CHANGES summarizes software changes since the last release. README summarizes the current release's features. README_PUBS explains the documentation included in the LSCdoc package.	ASM, ASM-QFS

Files Created

The following table lists the files created when ASM products are installed:

File	Description	Used By
/etc/opt/LSCsamfs/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 samst driver.	ASM, ASM-

File	Description	Used By
		QFS
/kernel/fs/samfs	Solaris 32-bit file system module.	ASM, ASM-QFS
/kernel/fs/sparcv9/samfs	Solaris 64-bit file system module.	ASM, ASM-QFS
/kernel/sys/samsys	System call module.	ASM, ASM-QFS
/kernel/sys/sparcv9/samsys	Solaris 32-bit system call module.	ASM, ASM-QFS
/kernel/sys/sparcv9/samsys64	Solaris 64-bit system call module.	ASM, ASM-QFS

The ASM file systems have dynamically loadable components that are stored in the Solaris /kernel directory (see preceding list). You can determine the modules that are loaded by using the modload(1M) and modinfo(1M) commands. Typically, the file system module is loaded with the kernel at boot time using directives in /etc/system. Alternatively, the file system module can be loaded when the file system is first mounted after the ASM software is installed. The file system module can be unloaded when no ASM file systems are mounted by using the modunload(1M) command.

Site Files

ASM products use certain files that you create. The following table lists the files you create that are used by ASM products:

File	Description	Used By
/etc/opt/LSCsamfs/LICENSE.3.5	License file. For more information, see the licensing information pertinent to your installation in this chapter.	ASM, ASM-QFS
/etc/opt/LSCsamfs/archiver.cmd	Archiver command file. For more information, see the archiver.cmd(4) man page or see the *ASM and ASM-QFS Storage and Archive Management Guide*.	ASM, ASM-QFS

File	Description	Used By
/etc/opt/LSCsamfs/samfs.cmd	File system mount parameter command file. For more information, see the samfs.cmd(4) man page or see the *ASM File System Administrator's Guide*.	ASM, ASM-QFS
/etc/opt/LSCsamfs/recycler.cmd	Recycler command file. For more information, see the recycler.cmd(4) man page or see the *ASM and ASM-QFS System Administration Guide*.	ASM, ASM-QFS
/etc/opt/LSCsamfs/releaser.cmd	Releaser command file. For more information, see the releaser.cmd(4) man page or see the *ASM and ASM-QFS System Administration Guide*.	ASM, ASM-QFS
/etc/opt/LSCsamfs/preview.cmd	Previewer command file. For more information, see the preview.cmd(4) man page or see the *ASM and ASM-QFS System Administration Guide*	ASM, ASM-QFS
/etc/opt/LSCsamfs/defaults.conf	Miscellaneous default values. For more information, see the defaults.conf(4) man page.	ASM, ASM-QFS
/etc/opt/LSCsamfs/mcf	Master configuration file. For more information, see the mcf(4) man page.	ASM, ASM-QFS

Modified System Files

During installation, ASM software adds information to certain Solaris system files. These system files are ASCII text files. Solaris uses these files to identify loadable kernel modules by number rather than by name.

The following table lists the system files that are modified during the installation of ASM software packages:

File	Description	Used By
/etc/name_to_sysnum	System call information file. The lines added are as follows: samsys 180 samsys64 181 (for Solaris 2.7 and above)	ASM, ASM-QFS
/etc/name_to_major	Maps driver to major number.	ASM, ASM-QFS

Chapter 2 - System Requirements

This chapter outlines the system requirements that must be met prior to the installation of ASM software packages. These requirements are as follows:

- Verify the environment
- Obtain superuser access
- Verify disk cache
- Verify disk space
- Verify Solaris patches
- Verify removable media devices (ASM and ASM-QFS packages only)
- Verify the ASM software license

The following subsections describe these requirements in more detail.

WARNING

If you have not read the README file delivered with this release, please do so before continuing. The ASM 3.5.0 releases included significant restructuring changes as compared to previous revisions. Failure to recognize these changes could cause dramatic changes in script behavior. The README file is included in the FTP instructions and is on the CD-ROM distribution. If your ASM software is already installed, it is located in /opt/LSCsamfs/doc/README.

Requirement 1: Verify the Environment

ASM products run on the following platforms, all of which contain a SPARC processor:

- SPARCstation 4 workstations and above
- Sun Ultra 1 workstations and above
- Sun Enterprise 1000 servers and above

Verify that you are installing your software package on one of the preceding systems and that the system is up and running prior to installing any ASM software. The system must be capable of reading the release CD-ROM, be network-attached to another system capable of reading CD-ROMs, or be capable of accessing the release package via FTP.

Although not officially supported by ASM, all ASM software has been installed on SPARC clones.

ASM software relies on a properly configured Solaris 2.6, 2.7, or 2.8 operating system. Check to see that your server is running one of these levels of Solaris by entering the following command:

server# **uname -sr**

SunOS 5.7

SunOS 5.*x*.*y* levels correspond to Solaris 2.*x*.*y* levels. The above system is running Solaris 2.7.

Requirement 2: Obtain Superuser Access

You must have super user (root) access to the system upon which the ASM software is to be installed.

Requirement 3: Verify Disk Cache

The ASM file system requires a certain amount of disk cache for the creation and management of data files and directories. For ASM and ASM-QFS, 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.

The disk devices or partitions do not require any special formatting, nor do they need to have a UNIX file system made 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 make the ASM file system.

For ASM and ASM-QFS, the disk must be attached to the server using a fiber 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 third-party volume management software, are supported.

Use the format(1M) command to see the disks attached to your system. The following example shows three disks attached to a server, one internal disk connected via controller 0 on the first target (c0t1d0) and two external disks connect via controller 1 on targets 1 and 2 (c1t1d0 and c1t2d0):

server# **format**

c0t1d0 <SUN1.05 cyl 2036 alt 2 hd 14 sec 72>
/iommu@f,e0000000/sbus@f,e0001000/espdma@f,400000/esp@f,800000/sd@1,0

c1t1d0 <SEAGATE-ST424-0116 cyl 2604 alt 2 hd 19 sec 84>
/iommu@f,e0000000/sbus@f,e0001000/dma@2,81000/esp@2,80000/sd@1,0

c1t2d0 <SEAGATE-ST424-0119 cyl 2604 alt 2 hd 19 sec 84>
/iommu@f,e0000000/sbus@f,e0001000/dma@2,81000/esp@2,80000/sd@2,0

Note that the format(1M) command requires you to enter CTRL-D to exit.

The amount of disk cache needed depends upon the size of the files being managed, the frequency of use for the files, the applications using the files, and other factors. Rough estimates differ depending on the environment, too. These differences are as follows:

For ASM-QFS, an estimate can be determined using the following algorithm:

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

For ASM, an estimate can be determined using the following algorithm:

Disk Cache = Largest File (in bytes) +

((Number of Files + Number of Directories) * 512) +

4096 * Number of Directories +

Amount of space needed for working files

For ASM-QFS, an estimate can be determined using the following algorithm:

Disk Cache = Largest File (in bytes) +

((Number of Files + Number of Directories) * 512) +

4096 * Number of Directories +

Amount of space needed for working files

Metadata Cache = ((Number of Files + Number of Directories) * 512) +

16384 * Number of Directories

Requirement 4: Verify Disk Space

Directory	ASM, ASM-QFS Disk Space Needed	ASM-QFS Disk Space Needed
/ (root) directory	1,328 – 1,335 kilobytes	1,268 kilobytes
/opt directory	9,833 – 22,087 kilobytes	4,400 kilobytes
/var directory	163 – 231 kilobytes	122 kilobytes

The ASM 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. The following list shows the minimum amount of disk space required in these various directories:

Note that log files are written to the /var directory, so the sizes shown in the preceding two lists should be considered a minimum amount for the /var directory. 30 megabytes or more is recommended.

The individual installation packages require the following amounts of space:

Package	Space on / (root)	Space on /opt	Space on /var
LSCqfs (samqfs)	1,268 kilobytes	4,400 kilobytes	122 kilobytes
LSCsamfs (sampkg)	1,328 kilobytes	9,833 kilobytes	163 kilobytes
LSCjre (samjre)	0 kilobytes	6,563 kilobytes	7 kilobytes
LSCgui (samgui)	4 kilobytes	753 kilobytes	28 kilobytes
LSCdst (samdst)	0 kilobytes	176 kilobytes	6 kilobytes
LSCibm (samibm)	0 kilobytes	231 kilobytes	6 kilobytes
LSCstk (samstk)	2 kilobytes	1,714 kilobytes	6 kilobytes
LSCsony (samsony)	1 kilobyte	420 kilobytes	7 kilobytes
LSCdoc (samdoc)	0 kilobytes	2,397 kilobytes	8 kilobytes

Determine the amount of space you have by issuing the df(1M) command, as follows:

```
server# df -k /
Filesystem        kbytes  used    avail capacity   Mounted on
/dev/dsk/c0t1dos0 76767   19826   49271    29%     /
```

server# **df -k /opt**

```
Filesystem        kbytes  used    avail capacity   Mounted on
/dev/dsk/c0t1dos4 192423  59006  114177    35%     /opt
```

For the / (root) directory, the number in the avail column needs to be 4096 or higher. For the /opt directory, the number in the avail column needs to be 7168 or higher.

If there is not enough room for the software under each directory, either remove the existing files under each directory or re-partition the disk to make more space available to each file system. To repartition a disk, see the your Solaris system administration documentation.

Requirement 5: Verify Solaris Patches

The latest patches for the Solaris operating system are required. An updated list of Solaris patches required prior to installation is included with the ASM software in a file called README. After installation, the README is located in /opt/SUNWsamfs/doc/README.

To determine which patches are installed on your system, enter the following:

server# **showrev -p | more**

If the required patches are not listed in the output from the showrev(1M) command, you need to install them before installing any ASM release packages. Patches are provided to customers with a Sun Microsystems maintenance contract via CD-ROM, anonymous FTP, and the Sun Microsystems website (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.

NOTE

Installation of certain Solaris patches results in the overwriting of the /etc/name_to_sysnum file. The /etc/name_to_sysnum file identifies loadable kernel modules by number rather than by name. If this file is overwritten by a patch installation, the ASM system call number does not exist. A system panic can occur if you attempt to start ASM processes and mount ASM file systems if the system call number is not present in this file.

To resolve this problem, make a copy of the /etc/name_to_sysnum file prior to installing Solaris patches. After patch installation, check the file to ensure that the samsys entry (for example, samsys 180) is present in the file. If the entry is not present, add it to the file or reinstall the ASM software.

Requirement 6: Verify Removable Media Devices (ASM and ASM-QFS Packages Only)

The ASM and ASM-QFS environments should include at least one removable media 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.

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 the StorageTek website (http://www.StorageTek.com) in the Customer Resource Center (CRC).

The device that you are using must be attached and recognized by the server. If the removable media device is already connected and communicating with the server, skip to the next requirement.

Instructions for attaching removable media devices to a server are presented in this subsection. 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 with the automated library and drives.

The general connection guidelines are as follows:

Power down the server for connecting devices. Use the init(1M) command as follows:

server**# init 0**

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

Ensure that the removable media devices and the disk to be used for disk cache are connected. Ensure that the SCSI target IDs are unique for each controller. For example, if you are using the internal SCSI host adapter, the internal disk drive ID is usually 3; therefore, any peripheral connected to the internal bus must not have an ID of 3. Typically the internal disk drive ID is 3 for Sparc systems and 0 for Ultra systems.

Power on the peripherals and server according to the manufacturer's recommended sequence. Before the server boots, press the Stop key and the A key simultaneously to interrupt the boot process. Then enter the following at the PROM ok prompt:

>ok probe-scsi-all

This command returns a series of entries for each device connected to the system through a SCSI interface.

You also need to verify devices attached through fiber channel interfaces. Enter the following commands to locate the host adaptor directory, select an item, and display the fiber channel host bus adaptor (HBA) devices:

>ok **show-devs**
/*output_line1*
/*output_line2*
>ok **select select** /*output_line1*
>ok **show children**
output

If the server does not acknowledge all of 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. Ensure again that each device has a unique target ID. Do not proceed until all devices appear when probed.

In some instances, SCSI devices might use a target number greater than 6 or a Logical Unit Number (LUN) greater than 0. This occurs when using DLT 2700 drives, which use a LUN of 1. If this is the case with your system, you must edit the /kernel/drv/samst.conf file when installed.

For more information, see chapter 5, “ASM and ASM-QFS Initial Installation Procedure”.

Boot the server using the reconfiguration option as follows:

>ok boot -r

Requirement 7: Verify the ASM Software License

If you do not have an ASM license key for the release level that you are installing, contact your Authorized Service Provider (ASP) or StorageTek. For information on contacting your ASP or StorageTek, see appendix A, “StorageTek Product Support”.

You will need the following identification information:

- Company purchase order (PO) number
- Company name, address, phone, and contact information
- Host ID upon which the ASM software is to be licensed. To display the host ID on your system, use the hostid(1) command.
- The ASM product you are installing.
- The server upon which the software is to be installed. (ASM-QFS package only)

To install the ASM or ASM-QFS packages, you need the following information for each automated library to be used in the ASM or ASM-QFS environment:

- The vendor name and the model of the automated library and the type of cartridge used in the automated library.
- The number of slots for the automated library and the media type.
- ASM optional products to be used with this license. For more information on additional ASM products, see the “Licensing” subsection in this manual’s preface.

The license keys for the 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 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 cartridges. Access continues unaffected for files already on disk.

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

Chapter 3 - ASM-QFS Standalone Initial Installation Procedure

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.

The step-by-step procedure in this chapter describes obtaining the files, installing the software packages on your server, and configuring the software to match the hardware at your site. This subsection also describes the steps needed to initialize the ASM-QFS file system and procedures for checking the status of your system. For most of the procedures in this subsection, you must have superuser (root) access.

If you are upgrading ASM-QFS software on an existing server, see chapter 4, “ASM-QFS Upgrade Procedure”.

Step 1: Obtain the Release Files

The ASM-QFS Standalone software can be obtained on a CD-ROM or by anonymous FTP. Contact your ASP or StorageTek for information on obtaining the software in one of these ways.

If you have a CD-ROM, run the Solaris Volume Manager, insert the CD-ROM, and change the directory to the ASM-QFS Standalone software files by using the following command:

server# **cd /cdrom/cdrom0**

WARNING

If you have not read the README file delivered with this release, please do so before continuing. The ASM-QFS Standalone 3.5.0 release includes significant restructuring changes as compared to previous revisions. Failure to recognize these changes could cause dramatic changes in script behavior. The README file is included in the FTP instructions and is on the CD-ROM distribution. After your ASM-QFS Standalone software is installed, it is located in /opt/LSCsamfs/doc/README.

Step 2: Add the Packages

The ASM-QFS Standalone file system uses the Solaris packaging utilities for adding and deleting software. As such, 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 packages.

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

To satisfy product dependencies, the samqfs package must be installed first. Run the pkgadd(1M) command to install all packages. Answer all to the first question, and answer yes to each of the others:

server# **pkgadd -d samqfs** (must be first)

If you want to install ASM documentation, install it now. Manuals are available in PDF format. Add this package as follows:

server# **pkgadd -d samdoc** (optional)

Step 3: Reboot the System

Reboot the server using the reconfiguration option, as follows:

server# reboot -- -r

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

NOTE

Failure to reboot the system at this time can cause a system panic.

Step 4: Set Up PATH and MANPATH Variables

For users who will be executing the ASM-QFS Standalone user commands (for example, sls(1)), add /opt/LSCsamfs/bin to the users' PATH variables.

For users who need to access the ASM-QFS Standalone man pages, add /opt/LSCsamfs/man to the MANPATH variable.

For users, such as superusers, who need to access the administrator commands, add /opt/LSCsamfs/sbin to the PATH variable.

In the Bourne or Korn shells, edit the .profile file, change the PATH and MANPATH variables, and export the variables. For example:

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

In the C shell, change your .login and/or .cshrc file. For example, the path statement in your .cshrc file might look like this:

set path = ($path /opt/LSCsamfs/bin /opt/LSCsamfs/sbin)

For example, in the C shell, the MANPATH statement in your .login file might look like this:

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

Step 5: License the ASM-QFS Standalone Software

License keys are required to run the ASM-QFS Standalone software and associated products from StorageTek. For information on license keys, see chapter 2, “System Requirements”.

The ASM-QFS Standalone file system 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.

Verify whether or not the following file exists:

/etc/opt/LSCsamfs/LICENSE.3.5

If the /etc/opt/LSCsamfs/LICENSE.3.5 file does not exist, create it.

Starting in column one, place the license keys you have obtained from your ASP or from ASM on the first and succeeding lines in the /etc/opt/LSCsamfs/LICENSE.3.5 file.

Each license key must be on a separate line, and all keys must start in column one. No other keywords, host IDs, comments, or other information can appear in the LICENSE.3.5 file. The license becomes effective when the ASM-QFS Standalone file system is mounted.

The license keys allow the system to run indefinitely unless you were issued a temporary license.

Step 6: Configure 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/LSCsamfs/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 a master configuration file in /etc/opt/LSCsamfs/mcf that contains a line for each device and/or family set in your configuration. The mcf contains information that enables you to identify the disk slices to be used and organize them into ASM-QFS Standalone file systems.

NOTE

For information on file system design considerations, see the *ASM File System Administrator's Guide*.

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 (#). Some fields are optional, so use a dash (–) to indicate omitted fields. The following line shows the format for the fields of each line entry in the mcf file:

Equipment identifier	Equipment ordinal	Equipment type	Family set	Device state	Additional parameters

The following list shows the information to be contained in each field and whether or not the field is a required or optional field:

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 32757.

Field	Description
Equipment type	Required. Enter a 2-character or 3-character code for the device type, as follows: The ma device type defines an ASM-QFS Standalone file system. The mm device type defines a metadata device. The mr device type defines a round robin or striped data device. The g*XXX* device 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.

Field	Description
	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 unspecified, this field should contain a dash (-) character. If specified, this field should contain either on or off. Enter a state for the device for when the ASM-QFS Standalone file system is initialized.
Additional parameters	Required. For a disk slice, this field points to the /dev/rdsk entry.

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

CAUTION

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 a UFS file system is not mounted.

Make sure you specify disk partitions that are not in use on your system. Do not use overlapping partitions. If an ASM file system attempts to use a partition that is already in use, the ASM software issues a message to indicate that the device is busy.

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

```
#
# QFS 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/rdsk/c1t0d0s0
/dev/dsk/c1t1d0s2 12     mr    qfs1  on     /dev/rdsk/c1t1d0s2
/dev/dsk/c1t2d0s2 13     mr    qfs1  on     /dev/rdsk/c1t2d0s2
/dev/dsk/c1t3d0s2 14     mr    qfs1  on     /dev/rdsk/c1t3d0s2
```

Note that all ASM-QFS Standalone configurations could have automated libraries and other removable media devices defined as well, essentially extending the size of the disk cache. Removable media device configurations are not shown. For information on configuring removable media devices, chapter 5, “ASM and ASM-QFS Standalone Initial Installation Procedure”.

Example Configuration 1

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

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	root	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	root	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 5 of c1t1d0. Another file system (qfs2) is created on slice 1 of disk c1t0d0 and slice 0 of disk c1t1d0.

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

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.

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

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

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/c1t0d0s0  11   mm  qfs1   on    /dev/rdsk/c1t0d0s0
/dev/dsk/c2t0d0s0  12   mr  qfs1   on    /dev/rdsk/c2t0d0s0
/dev/dsk/c3t0d0s0  13   mr  qfs1   on    /dev/rdsk/c3t0d0s0
#
#
qfs2              20   ma  qfs2
/dev/dsk/c1t1d0s1  21   mm  qfs2   on    /dev/rdsk/c1t1d0s1
/dev/dsk/c2t1d0s1  22   mr  qfs2   on    /dev/rdsk/c2t1d0s1
/dev/dsk/c3t1d0s1  23   mr  qfs2   on    /dev/rdsk/c3t1d0s1
```

NOTE

Be sure that the /dev/dsk and /dev/rdsk names on each line reference the same c*n*t*n*d*n*s*n* partition.

Example Configuration 2

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

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

```
server# 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 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 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/rdsk/c1t1d0s0
/dev/dsk/c2t4d0s0 12    mr     qfs1  on  /dev/rdsk/c2t4d0s0
/dev/dsk/c2t4d1s0 13    mr     qfs1  on  /dev/rdsk/c2t4d1s0
/dev/dsk/c2t4d2s0 14    mr     qfs1  on  /dev/rdsk/c2t4d2s0
/dev/dsk/c2t4d3s0 15    mr     qfs1  on  /dev/rdsk/c2t4d3s0
```

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

- Line 1 defines the 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 making 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, and 5 are the data devices using the mr device type.

Note that you must make certain that the /dev/dsk and the /dev/rdsk names on each line reference the same c*n*t*n*d*n*s*n* partition.

Example Configuration 3

This example configuration illustrates an ASM-QFS Standalone file system that separates the metadata on to a low-latency disk. Round-robin allocation is used on four disk drives.

The following assumptions are used:

The metadata device is a single partition (s0) 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).

NOTE

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

Step 1: Write the mcf File

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

```
# QFS 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/c0t0d0s0  11   mm   qfs1    on    /dev/rdsk/c0t0d0s0
/dev/dsk/c1t1d0s0  12   mr   qfs1    on    /dev/rdsk/c1t1d0s0
/dev/dsk/c1t2d0s5  13   mr   qfs1    on    /dev/rdsk/c1t2d0s5
/dev/dsk/c1t3d0s0  14   mr   qfs1    on    /dev/rdsk/c1t3d0s0
/dev/dsk/c1t4d0s1  15   mr   qfs1    on    /dev/rdsk/c1t4d0s1
```

Step 2: Modify the /etc/vfstab File

The /etc/vfstab is edited. Because the ASM-QFS Standalone file system uses round-robined allocation as a default, no stripe width is necessary.

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

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

Step 3: Run the sammkfs(1M) Command

Initialize the ASM-QFS Standalone file system by using the sammkfs(1M). The default DAU is 16 kilobytes, but the following example sets the DAU size to 64 kilobytes:

```
server# sammkfs -a 64 qfs1
```

Example Configuration 4

This sample configuration 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). All partitions are used on the entire disk drive (s2).

NOTE

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

Step 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 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/rdsk/c0t1d0s6
/dev/dsk/c1t1d0s2  12   mr  qfs1    on    /dev/rdsk/c1t1d0s2
/dev/dsk/c1t2d0s2  13   mr  qfs1    on    /dev/rdsk/c1t2d0s2
/dev/dsk/c1t3d0s2  14   mr  qfs1    on    /dev/rdsk/c1t3d0s2
/dev/dsk/c1t4d0s2  15   mr  qfs1    on    /dev/rdsk/c1t4d0s2
```

Step 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 DAU. Note the DAU is the allocation unit you set when you initialize the file system (see “Step 3”, following).

Step 3: Run the sammkfs(1M) Command

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

```
server# 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.

Example Configuration 5

Striped groups allow you to group RAID devices together for very large files. Normally, a DAU is represented by one bit in the bit maps. With striped groups, however, there is only one DAU per striped group. This method of writing huge DAUs across RAID devices saves bit map 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 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 sample 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). All partitions are used on the entire disk drive (s2).

NOTE

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

Step 1. Write the mcf File

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

```
# QFS 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/rdsk/c0t1d0s6
/dev/dsk/c1t1d0s2  12   g0  qfs1    on    /dev/rdsk/c1t1d0s2
/dev/dsk/c1t2d0s2  13   g0  qfs1    on    /dev/rdsk/c1t2d0s2
/dev/dsk/c0t3d0s2  14   g1  qfs1    on    /dev/rdsk/c0t3d0s2
/dev/dsk/c0t4d0s2  15   g1  qfs1    on    /dev/rdsk/c0t4d0s2
```

Step 2: Modify the /etc/vfstab File

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

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

Step 3: Run the sammkfs(1M) Command

Initialize the ASM-QFS Standalone file system 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.

```
server# 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 2 striped groups. You are really 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 7: Create the samfs.cmd File (Optional)

The /etc/opt/LSCsamfs/samfs.cmd file can be created as the place from which mount parameters are read. Creating this file may be beneficial if you are configuring multiple ASM-QFS Standalone systems with multiple mount parameters. Information can be provided in the samfs.cmd file, in the /etc/vfstab file, and on the mount(1M) command. The directive lines in the samfs.cmd file serve as defaults, but they can be overridden by options on the mount(1M) command.

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

- Striping.

- The shared reader, which allows a file system to be shared by multiple servers.
- 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 File System Administrator's Guide* or see the samfs.cmd(4) man page. For more information on the /etc/vfstab file, see "Step 8: Create the Mount Point and Update the /etc/vfstab File" later in this procedure. For more information on the mount(1M) command, see the mount_samfs(1M) man page.

Step 8: Create the Mount Point and Update the /etc/vfstab File

Edit the /etc/vfstab file and make an entry for each ASM-QFS Standalone file system. An example entry follows:

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

The various fields and their content are as follows:

Field	Field Title and Content
1	Device to mount. The name of the ASM-QFS Standalone file system to mount. This must be the same name as specified in the mcf file.
2	Device to fsck. Must be a dash (-) character. The dash indicates that there are no options. This prevents the system from performing an fsck 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

Field	Field Title and Content
	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.

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. First, create the mount point, as follows:

server# **mkdir /qfs1**

Next, if you want, change the permissions, owner, or group owner of the /qfs1 directory when it is not mounted:

server# **chmod 555 /qfs1**

server# **chown root /qfs1**

server# **chgrp other /qfs1**

NOTE

If you configured multiple mount points, repeat these steps for each mount point, using a different mount point (such as /qfs1) and family set name (such as qfs1) each time.

Step 9: Make the File System

Using the sammkfs(1M) command and the family set names that you have defined, create a file system for each family set.

For example, the following command creates a file system for family set name qfs1:

server# **sammkfs –a 128 qfs1**

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

server# **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 data currently contained in the partition associated with the file system in the /etc/opt/LSCsamfs/mcf file.

Step 10: Start Up and Shut Down the ASM-QFS Standalone File System

The mount(1M) command mounts an ASM-QFS Standalone file system. For information on the mount(1M) command, see the mount_samfs(1M) man page.

Change the /etc/vfstab mount at boot parameter to yes. This specifies that ASM-QFS Standalone file systems be mounted by /etc/rc2.d/S01MOUNTFSYS.

Manual Start Up and Shut Down

Several commands are used to manually start up and shut down the ASM-QFS Standalone system. The following examples assume that /qfs1 is the mount point and qfs1 is the file system.

Enter the following commands to perform manual startup and to set the permissions on /qfs1 to allow appropriate read and write access:

server# **mount /qfs1**

server# **chmod 755 /qfs1**

server# **df –k** # *Allows you to determine whether* qfs1 *is mounted*

Enter the following commands to perform manual shutdown:

server# **umount /qfs1**

server# **df –k** # *Allows you to determine whether* qfs1 *is unmounted*

Automatic Start Up and Shut Down

To perform an automated start up, perform the following tasks:

Ensure that the mount at boot parameter is yes in the /etc/vfstab file.

Boot the system. Go to run level 3 (multiuser mode) using the init(1M) command with the 3 option.

Step 11: Share the File System with Client Machines

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

For example, on the server, enter a line like the following:

server# **share -F nfs -o rw=client1:client2 -d "QFS" /qfs1**

NOTE

If you write a share(1M) command like the preceding example command into the /etc/dfs/dfstab file, Solaris shares the file system after the next system reboot. If you want to share the file system immediately, you must type the share(1M) command at a root shell prompt. If there are no shared file systems when Solaris boots, the NFS server is not started. You must reboot after adding the first share entry to this file.

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, ASM 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 12: Mount the File System on the Client Machines

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 hard,intr,timeo=60

Next, on the command line, enter the mount command, as follows:

client# **mount /qfs1**

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

NOTE

It is strongly recommended that clients mount the file system with the hard option. At times, there may be a significant delay in the ASM-QFS Standalone file system's response to client requests. This can occur when a requested file resides on a cartridge that must be loaded into a DLT tape drive. If the hard option is not specified, the client can return an error instead of retrying the operation until it completes.

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.

Step 13: Establish Periodic Dumps Using qfsdump(1M)

The server should periodically create a control structure dump using qfsdump(1M). When using the qfsdump(1M) command, please note the following:

- A qfsdump(1M) dump taken under release 3.5.0 cannot be restored at earlier releases of the ASM-QFS Standalone software because of new data structures needed to support data inclusion.
- A qfsdump(1M) dump can be very large. The qfsdump(1M) command does not have any tape management or estimations as does ufsdump(1M). You need to weigh space considerations when using qfsdump(1M). For more information on these commands, see the qfsdump(1M) and ufsdump(1M) man pages.

This dump does not include the data stored in your file system, but it does include information necessary to quickly locate the data on your removable media devices. This information is necessary to recover from a cache disk failure. Use qfsrestore(1M) to restore the control structure dump after initializing the file system if such a failure occurs.

For example, you can make an entry in root's crontab file so the cron daemon runs qfsdump periodically:

10 0 * * * (cd /samfs1; qfsdump –B 512 –f /dev/rmt/0cbn)

If you have multiple ASM-QFS Standalone file systems, make similar entries for each. Make sure you save each dump in a separate file.

For more information on using qfsdump(1M), see the qfsdump(1M) man page and the *ASM File System Administrator's Guide*.

Chapter 4 - ASM-QFS Standalone Standalone Upgrade Procedure

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

All steps in this subsection must be performed as superuser (root).

Step 1: Obtain the Release Files

The ASM-QFS Standalone software can be obtained on a CD-ROM or by anonymous FTP. Contact your ASP or StorageTek for information on obtaining the software in one of these ways.

WARNING

If you have not read the README file delivered with this release, please do so before continuing. The ASM-QFS Standalone 3.5.0 release includes significant restructuring changes as compared to previous revisions. Failure to recognize these changes could cause dramatic changes in script behavior. The README file is included in the FTP instructions and is on the CD-ROM distribution. After your ASM-QFS Standalone software is installed, it is located in /opt/LSCsamfs/doc/README.

If you have a CD-ROM, run the Solaris Volume Manager, insert the CD-ROM, and change the directory to the ASM-QFS Standalone software files using the following command:

server# **cd /cdrom/cdrom0**

Step 2: Back Up Each ASM-QFS Standalone File System

If you do not have current backup files for each of your ASM-QFS Standalone file systems, create them now using qfsdump(1M) or by copying the inodes file using dd(1M), as follows:

Back up each ASM-QFS Standalone file system using qfsdump(1M). The location to which you dump each system should be outside the ASM-QFS Standalone file system.

The following example assumes that you have a file system named qfs1 (mounted at /qfs1) that you want to back up to qfs1.bak, which exists outside of the ASM-QFS Standalone file system:

server# cd /qfs1

server# qfsdump –B 512 –f /dev/rmt/0cbn

The qfsdump(1M) command dumps file names and inode information, not data. For more information on this, see the qfsdump(1M) man page.

Back up any site-defined scripts. Because of the directory restructuring in the ASM-QFS Standalone 3.5.0 system, any scripts stored in /etc/fs/samfs could be destroyed during an upgrade. Scripts in directories created by the ASM-QFS Standalone system in an earlier version could also be destroyed. After the ASM-QFS Standalone 3.5.0 software is installed, you can move the scripts to a location in /var/opt.

You need to back up files for each file system, so repeat the preceding steps for each file system in your ASM-QFS Standalone environment. For more information on this, see chapter 3, “ASM-QFS Standalone Initial Installation Procedure.”

Step 3: Unmount the File Systems

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

If you encounter difficulty unmounting a file system, 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. Use the fuser(1M) command to determine whether or not any processes are still busy. If any are still busy, you must terminate them by using the kill(1M) command. For example, you can enter the following command to determine whether or not processes are still running on the qfs1 file system:

server# **fuser –uc /qfs1**

If you are still unable to unmount the file system, issue the Solaris unshare(1M) command on the ASM-QFS Standalone file system as follows:

server # **unshare** *pathname*

In the preceding format, *pathname* is the path to the ASM-QFS Standalone file system that you are trying to unmount. After issuing the unshare(1M) command on the path to the ASM-QFS Standalone file system, try unmounting the file system again. For more information on this, see the unshare(1M) man page.

If all previous attempts to unmount the ASM-QFS Standalone file system still fail, edit the /etc/vfstab file. When editing this file, change all ASM-QFS Standalone file systems from yes or delay to no. Then reboot your system.

Step 4: Remove Existing ASM-QFS Standalone Software

Use the pkginfo(1) command, as follows, to determine which ASM-QFS Standalone software packages are installed on your system:

server# **pkginfo | grep LSC**

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 of the optional ASM packages as described at the beginning of this chapter, you should make sure that you remove these packages prior to the main LSCqfs package. The install script prompts you to confirm several of the removal steps.

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

server# **pkgrm LSCdoc LSCqfs**

The LSCqfs package must be the last package removed.

Step 5: Add the Packages

The ASM-QFS Standalone software package uses the Solaris packaging utilities for adding and deleting software. As such, 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 ASM packages.

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

To satisfy product dependencies, you must upgrade samqfs first. Run the pkgadd(1M) command to upgrade all packages, answering yes to each question:

server# **pkgadd -d samqfs** (must be first)

If you want to install ASM documentation, install it now. Manuals are available in PDF format. Add this package as follows:

server# **pkgadd -d samdoc** (optional)

Step 6: Update the License Keys

You must update the license keys for the ASM-QFS Standalone 3.5.0 release. If you are upgrading from an ASM-QFS Standalone release prior to 3.5.0, you need to place a new license key in the following file:

/etc/opt/LSCsamfs/LICENSE.3.5

To obtain new license keys, contact your ASP or StorageTek.

For more information on ASM-QFS Standalone license keys, see the information on licensing new ASM-QFS Standalone software in chapter 3, “ASM-QFS Standalone Initial Installation Procedure”.

Step 7: Verify the New Master Configuration File

The topology of the equipment managed by the ASM-QFS Standalone file system is defined in the master configuration file, /etc/opt/LSCsamfs/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.

Verify that a new mcf file exists in /etc/opt/LSCsamfs/mcf.

NOTE

For information on file system design considerations, see the *ASM File System Administrator’s Guide*.

Step 8: Modify the /etc/vfstab File (Optional)

If you modified the /etc/vfstab file in “Step 3: Unmount the File Systems”, edit this file again and change all ASM-QFS Standalone file systems from no to yes or delay.

Step 9: Reboot the System (Optional)

If you are upgrading from the ASM-QFS Standalone 3.3.1 release on Solaris 2.7, which is a 32-bit release, and are now upgrading to a 64-bit release, you need to use the reboot(1M) command to reboot at this point in order to build the 64-bit kernel.

Step 10: Mount the File System(s) (Optional)

You must perform this step if you have not modified the /etc/vfstab file to have yes or delay.

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 file system name to be mounted:

server# mount qfs1

Step 12: Relink API-dependent Applications (Optional)

If you are running applications that use the ASM-QFS Standalone application programmer interface (API) and you are using static linking, you should relink these applications at this time.

Chapter 5 - ASM and ASM-QFS Initial Installation Procedure

This chapter describes the step-by-step 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.

This step-by-step procedure describes copying and installing the software packages to your server and configuring the software to match the hardware at your site. This subsection also describes the steps needed to initialize the ASM and ASM-QFS file systems and procedures for checking the status of your system.

For most of the procedures in this subsection, you must have superuser (root) access.

If you are upgrading ASM or ASM-QFS software on an existing server, see chapter 6, “ASM and ASM-QFS Upgrade Procedure.”

Step 1: Obtain the Release Files

The ASM and ASM-QFS software can be obtained on a CD-ROM or by anonymous FTP. Contact your ASP or StorageTek for information on obtaining the software in one of these ways.

If you have a CD-ROM, run the Solaris Volume Manager, insert the CD-ROM, and change to the directory containing the ASM or ASM-QFS software files by using the following command:

server# **cd /cdrom/cdrom0**

WARNING

If you have not read the README file delivered with this release, please do so before continuing. The ASM and ASM-QFS 3.5.0 releases include significant restructuring changes as compared to previous revisions. Failure to recognize these changes could cause dramatic changes in script behavior. The README file is included in the FTP instructions and is on the CD-ROM distribution. After your software is installed, it is located in /opt/LSCsamfs/doc/README.

Step 2: Add the Administrator Group (Optional)

By default, ASM and ASM-QFS administrator commands can be executed by root only. However, during installation you can supply an administrator group name. This allows members of the administrator group to execute all administrator commands except for star(1M), samfsck(1M), samgrowfs(1M), sammkfs(1M), and samd(1M). The administrator commands are located in /opt/LSCsamfs/sbin.

If you want to enable an administrator group, choose a group name. The pkgadd(1M) process prompts you for this group name. This task is performed in "Step 3: Add the Packages".

Add the administrator group name. Do this by either using your site's procedure and the groupadd(1M) command or by editing the /etc/group file. The following is an entry from the group file designating an administrator group for ASM software. In this example, the samadm group consists of both the adm and operator users:

samadm::1999:adm,operator

You can also define an operator group that is allowed access only to 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/LSCsamfs/defaults.conf file as described later in "Step 10: Set Up Default Values" 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 can also undo the effect of this selection and make the programs in /opt/LSCsamfs/sbin executable only by root. For more information on this command, see the set_admin.sh(1M) man page.

Step 3: Add the Packages

ASM software uses the Solaris packaging utilities for adding and deleting software. As such, 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 packages.

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

To satisfy product dependencies, sampkg must be installed first. Run the pkgadd(1M) command to install all packages. Answer all to the first question, and answer yes to each of the others:

server# **pkgadd –d sampkg** (must be installed first)

When you install sampkg, 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.

If your ASM environment includes certain network-attached automated libraries, such as certain models from StorageTek, ADIC, IBM, or Sony, you may need to install one or more vendor-specific media changer packages. ASM supplies these packages. For information on whether or not you need to install a vendor-specific daemon package, see the *ASM and ASM-QFS System Administration Guide*.

To install one or more of these packages, install them using pkgadd(1M), as follows:

server# **pkgadd –d samstk** (optional StorageTek package)

server# **pkgadd –d samdst** (optional Ampex package)

server# **pkgadd –d samibm** (optional IBM package)

server# **pkgadd –d samsony** (optional Sony package)

If you want to use the GUI tools included with the ASM or ASM-QFS software package, install them now. The GUI tools require the presence of a Java runtime environment. Add the Java runtime environment and the GUI tool package as follows:

server# **pkgadd –d samjre** (optional)

server# **pkgadd –d samgui** (optional)

If you want to install ASM documentation, install it now. Manuals are available in PDF format. Add this package as follows:

server# **pkgadd –d samdoc** (optional)

Step 4: Add Tape Support to the st.conf File

Some tape devices are not, by default, supported in the Solaris kernel. If your environment uses any of the devices listed in /opt/LSCsamfs/examples/st.conf_changes, you must modify the /kernel/drv/st.conf file. The st(7D) tape driver configuration file for all supported tape drives is st.conf. By modifying this file and using the ASM or ASM-QFS software, you enable the normally unsupported drives to work with the software.

The following drives are not officially supported in all Solaris revisions:

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

If you want any of the preceding devices to operate within your environment, read the /opt/LSCsamfs/examples/st.conf_changes file into /kernel/drv/st.conf.

Notes on Target and LUN Numbers

In some instances, SCSI devices may 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 and the /kernel/drv/st.conf files.

For LUNs 1 through 7, make the following changes:

Edit /kernel/drv/st.conf. Add the following lines for each target/LUN combination, making the appropriate substitutions. For example, the following uses target 4, LUN 1:

name=”st” class=”scsi”

target=4 lun=1

Edit /kernel/drv/samst.conf. Un-comment or add the appropriate lines for each device, as follows:

name=”samst” class=”scsi”

target=4 lun=1

If you want to use targets 8 through 15, edit /kernel/drv/st.conf and find the following line:

In case there are wide tape drives, one can use these targets.

Un-comment the pair of lines for each target following this comment.

Then, edit /kernel/drv/samst.conf. Add lines for each target/LUN combination, substituting appropriately. For example, the following uses target 9, LUN 2:

name=”samst” class=”scsi”

target=9 lun=2

If you have added new devices since running pkgadd(1M), run the following command to create the device entries in /dev/samst:

server# **/opt/LSCsamfs/sbin/samdev**

Examples

The following examples show various st.conf files and error conditions.

Example 1: The following is an example of a /kernel/drv/st.conf file that has been modified to add support in the 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;
```

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

Example 2: 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.

Step 5: Reboot System

Reboot the server using the reconfiguration option, as follows:

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

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

NOTE

Failure to reboot the system at this time can cause a system panic.

Step 6: Set Up PATH and MANPATH Variables

For users running the ASM or ASM-QFS user commands (for example, sls(1)), add /opt/LSCsamfs/bin to the users' PATH variables.

To run the administrator commands, add /opt/LSCsamfs/sbin to the PATH variable.

To use ASM man pages, add /opt/LSCsamfs/man to the MANPATH variable.

In the Bourne or Korn shells, edit the .profile file, change the PATH and MANPATH variables, and export the variables. For example:

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

In the C shell, change your .login and/or .cshrc file. For example, the path statement in your .cshrc file might look like this:

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

For example, in the C shell, the MANPATH statement in your .login file might look like this:

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

Step 7: License the ASM Software

License keys are required to run all ASM software products. Licenses are assigned to specific hostid identifiers and are not transferable. For information on license keys, see chapter 2.

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.

Verify whether or not the following file exists:

/etc/opt/LSCsamfs/LICENSE.3.5

If the /etc/opt/LSCsamfs/LICENSE.3.5 file does not exist, create it.

Starting in column one, place the license keys you have obtained from your ASP or from StorageTek on the first and succeeding lines in the /etc/opt/LSCsamfs/LICENSE.3.5 file.

Each license key must be on a separate line and all keys must start in column one. No other keywords, host IDs, comments, or other information can appear in the LICENSE.3.5 file. The license becomes effective the next time sam-initd is started.

The license keys allow the system to run indefinitely unless one of the followi

- 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, and 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 ASM 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.

After the system is running, you can view the current license settings from the samu(1M) utility's l (the letter l, for license) display.

Step 8: Configure System Logging

The ASM and ASM-QFS systems log errors, cautions, warnings, and other messages using the standard Solaris syslog(3) interface. By default, the ASM facility is local7. Add a line similar to the following example to the /etc/syslog.conf file:

local7.debug /var/adm/sam-log

You can read this line from /opt/LSCsamfs/examples/syslog.conf_changes.

NOTE

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

After adding this line, create an empty log file, identify the process identifier (PID) for syslogd(1M), and send the syslogd PID a HUP signal. The following command sequence creates a log file in /var/adm/sam-log, identifies the PID, and sends the HUP:

server# **touch /var/adm/sam-log**
server# **ps -ef | grep syslogd**
server# **kill -HUP** *syslogd-pid*

For more information, see the syslog.conf(4) and syslogd(1M) man pages. A different logging facility can be set in the /etc/opt/LSCsamfs/defaults.conf file.

Step 9: Configure the Environment

Each ASM software environment is unique. The system requirements and hardware used 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/LSCsamfs/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 File System Administrator's Guide*.

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

- Identify the disk slices to be used and organize them into file systems.
- Identify the media drives to be used and organize them into libraries.

NOTE

The instructions for creating the mcf file differ depending on whether you are creating a ASM or ASM-QFS environment.

If you are installing the ASM environment, all configuration instructions are contained in this subsection.

If you are installing the ASM-QFS environment, the instructions for storage and archive management device configuration are contained in this subsection. The instructions for configuring the QFS file system, however, are contained in chapter 3, "ASM-QFS Initial Installation Instructions".

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 (#). Some fields are optional, so use a dash (–) to indicate omitted fields. The following line shows the format for the fields of each line entry in the mcf file:

Equipment identifier	Equipment ordinal	Equipment type	Family set	Device state	Additional parameters

The following list shows the information to be contained in each field and whether or not the field is a required or optional field:

Field	Description
Equipment identifier	Required. If the device is an automated library or optical drive, this field is the /dev/samst entry. If the device is a

Field	Description
	disk slice, this field is the /dev/dsk entry. If you are using a network-attached automated library, see the information on managing vendor-specific automated libraries in the *ASM Administrator's Guide*.
Equipment ordinal	Required. The automated libraries and all associated drives must be assigned a unique equipment ordinal. Enter a unique integer from 1 to 32757.
Equipment type	Required. Enter a 2- or 3-character mnemonic for the device type. Most equipment can use the generic equipment types of od (optical disk), tp (tape), and rb (robot). See the mcf(4) man page for specific equipment types.
Family set	Optional. If the device is associated with a family set (that is, a file system or automated library), enter the family set name for this device. If the device is a manually loaded drive, use the dash (-) to indicate that this field is omitted.
Device state	Optional. Enter a state for the device for when the file system is initialized.
Additional parameters	Optional. If the device is a disk slice, this field points to the /dev/rdsk entry. If the device is an automated library, this field is the path name to the library catalog file. The default library catalog file is /var/opt/LSCsamfs/catalog/*family_set_name*.

Example: 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/c0t3u0	50	s9	9738	on	9738
/dev/rmt/0cbn	51	sg	9738	on	
/dev/rmt/1cbn	52	sg	9738	on	

In the preceding example, the library catalog is written to /var/opt/LSCsamfs/catalog/9738.

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

Example ASM Configuration

Assume that the following equipment is available on an example server:

- Two Seagate ST15230W 4 gigabyte disk drives used as cache
- One StorageTek 9730 30-slot automated library that contains two DLT tape drives
- One manually loaded DLT 2000 DLT 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 the following targets:

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

- A differential SCSI bus connects to the HP Model C1710T automated library and cache disk with the following targets:

Equipment	SCSI target
Cache disks	0 and 1
HP C1710T automated library	2
first optical drive	5
second optical drive	6

- A differential SCSI bus connects to the StorageTek 9730 automated library and tape drives with the following targets:

Equipment	SCSI target
StorageTek 9730 automated library	0
first DLT 7000 drive	1
second DLT 7000 drive	2

Example ASM Disk Cache Configuration

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

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	root	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	root	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 (samfs1) is placed on slice 0 of disk c1t0d0 and slice 5 of c1t1d0. Another file system (samfs2) is created on slice 1 of disk c1t0d0 and slice 0 of disk c1t1d0.

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

Make an ms (for mass storage) entry for the first file system. The name of this file system (samfs1) is used later when writing the /etc/vfstab entry for the file system and making the file system.

Make a series of md (for magnetic disk) entries listing the partitions that comprise the samfs1 file system.

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

Figure 2-1 shows the file system entries in the mcf file.

NOTE

Be sure that the /dev/dsk and /dev/rdsk names on each line reference the same c*n*t*n*d*n*s*n* partition.

```
# 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/rdsk/c1t0d0s0
/dev/dsk/c1t1d0s5  12   md  samfs1   on    /dev/rdsk/c1t1d0s5
#
#
samfs2           20   ms  samfs2
/dev/dsk/c1t1d0s0  21   md  samfs2   on    /dev/rdsk/c1t1d0s0
/dev/dsk/c1t0d0s1  22   md  samfs2   on    /dev/rdsk/c1t0d0s1
```

Figure 2-1. File System Entries in an mcf

CAUTION

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

How to Identify Peripherals Using /var/adm/messages

When your system boots, a series of messages are written to /var/adm/messages. These messages identify the 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):

```
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

The first line displays the vendor and product information that the SCSI peripheral reported to the 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 peripheral is the key to configuring a ASM or ASM-QFS environment. Use the ls(1) command with the -l option in both the /dev/samst and /dev/rmt directories to point to 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 on 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. 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

Change directories to /dev/samst and use the following ls(1) command:

server# 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. Make the following entry in /etc/opt/LSCsamfs/mcf:

/dev/samst/c0t2u0 30 od - on

This entry contains the device name (/dev/samst/c0t2u0), a unique ordinal (20), 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. Look in /var/adm/messages 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

Change directories to /dev/samst and use ls(1) commands, as follows, to search for the three symbolic links that point to the /devices files with the same Solaris hardware paths shown in the /var/adm/messages file.

```
server# 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
server# 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
server# 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
```

Make the following entries in /etc/opt/LSCsamfs/mcf.

```
/dev/samst/c1t2u0  50  rb  hp30  on  /var/opt/LSCsamfs/catalog/hp30

/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 ordinal (30), the equipment identifier (rb, for a generic library), the family set identifier specified on all devices associated with this library (hp30), the device state (on), and the path name to the library catalog.

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

Configuring a Manually Loaded DLT Drive

When configuring DLT drives, make sure to add the DLT definitions to the /kernel/drv/st.conf file (see “Step 4: Add Tape Support to the st.conf File”). DLT drives are not part of the standard Solaris configuration.

The following lines from /var/adm/messages 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
```

Find the matching /dev/samst symbolic link:

```
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 you can leave the additional parameters field empty. The system finds the proper /dev/samst/* symbolic link using the 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 is the path to the special file (for example, /dev/samst/cntnun).

For a tape device, there is another symbolic link located in /dev/rmt. This symbolic link is the name that the 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 prefixed with c if the drive supports compression. 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
```

Using this information, construct the /etc/opt/LSCsamfs/mcf entry:

```
/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 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. Look in /var/adm/messages to find the messages for these devices. /var/adm/messages 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
```

Find the /dev/samst symbolic links which point to these hardware paths:

```
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
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
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
```

A tape device is involved, so find a symbolic link in /dev/rmt that points to the tape devices (note that the robot 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. The drive supports compression, so choose the one with the cbn suffix. Had it not, you would have chosen the symbolic link whose name ended with bn.

Make the following entries in /etc/opt/LSCsamfs/mcf:

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

The first line defines the automated library, displaying the /dev/samst name (/dev/samst/c2t0u0), a unique ordinal (50), the equipment type (rb**,** for the generic robot equipment type), a family set name for the robot and the drive (9730), the device state (on), and the path name to the automated library's catalog (/var/opt/LSCsamfs/catalog/9730).

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 ordinal for the device (51), 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 (52), the equipment type (tp), the family set name (9730), and the device state (on).

Again, when configuring DLT drives, make sure that you add the DLT definitions to the /kernel/drv/st.conf file, as shown in "Step 4: Add Tape Support to the st.conf File". DLT drives are not part of the standard Solaris configuration.

Completed mcf *File*

Figure 2-2 shows the complete mcf file for the configuration example:

# Equipment	Eq	Eq	Family	Dev	Additional
# Identifier	Ord	Type	Set	Sta	Parameters
#					
samfs1	10	ms	samfs1		
/dev/dsk/c1t0d0s0	11	md	samfs1	on	/dev/rdsk/c1t0d0s0
/dev/dsk/c1t1d0s5	12	md	samfs1	on	/dev/rdsk/c1t1d0s5

samfs2	20	ms	samfs2		
/dev/dsk/c1t1d0s0	21	md	samfs2	on	/dev/rdsk/c1t1d0s0
/dev/dsk/c1t0d0s1	22	md	samfs2	on	/dev/rdsk/c1t0d0s1
/dev/samst/c0t2u0	30	od	-	on	
/dev/samst/c1t2u0	50	rb	hp30	on	/var/opt/LSCsamfs/catalog/hp30
/dev/samst/c1t5u0	51	od	hp30	on	
/dev/samst/c1t6u0	52	od	hp30	on	
/dev/rmt/0cbn	40	tp	-	on	
/dev/samst/c2t0u0	60	rb	9730	on	/var/opt/LSCsamfs/catalog/9730
/dev/rmt/0cbn	61	tp	9730	on	
/dev/rmt/1cbn	62	tp	9730	on	

Figure 2-2. Completed Sample mcf File

Step 10: Set up Default Values

The /opt/LSCsamfs/examples/defaults.conf file contains default settings for certain parameters in the ASM environment. Read the defaults.conf(4) man page and examine this file to determine which, if any, of the defaults should be changed. Copy the example defaults.conf to /etc/opt/LSCsamfs/defaults.conf, then edit the file, removing comments (preceded by a # sign) from the entries to be enabled.

Step 11: Create a Volume Serial Name (VSN) Catalog (Optional)

If you have a network-attached automated library, you must perform this step.

For SCSI-attached libraries, a VSN catalog is created automatically when the ASM or ASM-QFS environment is initialized. The Additional parameters field for the automated library entry in the mcf file gives the path name to the catalog.

However, for a network-attached automated library, you must create a catalog at this time. For instructions on how to build a VSN catalog, see the information on managing vendor-specific automated libraries in the *ASM and ASM-QFS System Administrator's Guide.*

Step 12: Create the samfs.cmd File (Optional)

The /etc/opt/LSCsamfs/samfs.cmd file can be created as the place from which mount parameters are read. Creating this file may be beneficial if you are configuring multiple ASM file systems with multiple mount parameters. Information can be provided in the samfs.cmd file, in the /etc/vfstab file, and on the mount(1M) command. The directive lines in the samfs.cmd file serve as defaults, but they can be overridden by options on the mount(1M) command.

For more information on the samfs.cmd file, see the samfs.cmd(4) man page or see the *ASM File System Administrator's Guide*. For more information on the /etc/vfstab file, see "Step 13: Create the Mount Point and Update the /etc/vfstab File". For more information on the mount(1M) command, see the mount_samfs(1M) man page.

Step 13: Create the Mount Point and Update 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, if you want.

Edit the /etc/vfstab file and make an entry for each ASM file system. An example entry follows:

```
samfs1  -  /sam  samfs  -  yes  high=80,low=60
```

The various fields and their content are as follows:

Field	Field Title and Content
1	Device to mount. The name of the ASM file system to mount.

Field	Field Title and Content
2	Device to fsck. A dash (-) indicates that there are no options. This prevents the system from performing an fsck on an ASM file system. For more information on this process, see the fsck(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 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 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 a list of available mount options, see the mount_samfs(1M) man page.

Create the mount point:

server**# mkdir /sam**

Next, if you want, change the permissions, owner, or group owner of the /sam directory when it is not mounted:

server**# chmod 555 /sam**

server**# chown root /sam**

server**# chgrp other /sam**

NOTE

If you configured multiple mount points, repeat these steps for each mount point, using a different mount point (such as /sam) and family set name (such as samfs1) each time.

Step 14: Make the File System

Using the sammkfs(1M) command and the family set names that you have defined, create a file system for each family set.

For example, the following command creates a file system for family set name samfs1:

server# **sammkfs samfs1**

The output from the preceding command is as follows:

total data kilobytes = 31842048

total data kilobytes free = 31841680

CAUTION

Running sammkfs(1M) creates a new file system. It removes all data currently contained in the partition associated with the file system in the /etc/opt/LSCsamfs/mcf file.

Step 15: Start Up and Shut Down the File System

The mount(1M) command mounts an ASM file system and starts sam-initd if the system is at run level 3 or above. For information on run levels, see the init(1M) man page. For information on the mount(1M) command, see the mount_samfs(1M) man page.

/etc/rc3.d/S95samd starts up sam-initd if an mcf file exists and sam-initd is not already running.

Change the /etc/vfstab mount at boot parameter to yes. This specifies that ASM file systems be mounted by /etc/rc2.d/S01MOUNTFSYS. It is possible to stop sam-initd and leave the file systems mounted. When sam-initd is restarted, pending stages are reissued and archiving is resumed.

Manual Start Up and Shut Down

Several commands are used to manually start up and shut down the file system. The following examples assume that /sam is the mount point and samfs1 is the file system.

Enter the following commands to perform manual startup and to set the permissions on /sam to allow appropriate read and write access:

server# **mount /sam**

server# **chmod 755 /sam**

Enter the following commands to perform manual shutdown:

server# **samcmd idle** *eq* # see **NOTE**

server# **samd stop**

server# **umount /sam**

NOTE

The drives in your ASM or ASM-QFS environment must be idled prior to issuing the samd stop command, so enter a samcmd idle *eq* command for each *eq* configured in your mcf file. Alternatively, 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 on the samcmd(1M) command, see the samcmd(1M) man page.

Automatic Start Up and Shut Down

To perform an automated start up, perform the following tasks:

Edit /etc/vfstab to ensure that the mount at boot parameter is yes.

Boot the system. Go to run level 3 (multiuser mode) using the init(1M) command with the 3 option.

If sam-initd is not running, stages are not queued and files are not archived until sam-initd is started.

Step 16: Drive Order Check Procedure

The drive order check procedure differs depending on whether or not your automated library has a front panel. The procedure for systems with a front panel appears first.

Drive Order Check Procedure – Systems With a Front Panel

Start the ASM or ASM-QFS software by mounting the file system or by using the samd start command.

Verify the order of the drives. If your automated library contains more than one drive, the drives defined in the mcf file must be 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 than the order of the SCSI targets or LUNs.

Verify the order in which the drives are recognized by the media changer controller. Check the SCSI target IDs displayed by the control panel of the media changer.

Drive Order Check Procedure – Systems Without a Front Panel

Start the ASM or ASM-QFS software by mounting the file system or by using the samd start command.

Verify the order of the drives. If your automated library contains more than one drive, the drives defined in the mcf file must be 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 than the order of the SCSI targets or LUNs.

Make sure you check *each* drive in a library. Make the drive state unavailable to the ASM file system. You can do this either by using libmgr(1M) to select the drives and change the drive state to unavailable or by using the samu(1M) utility's :unavail command.

Load a cartridge into the drive using the load(1M) command. Two possible formats for this command are as follows:

server# **load** mt**.**vsn device_num

OR

server# **load** eq**:**slot device_num

For more information on the load command format, see the load(1M) man page.

Determine if the correct drive responds while under ASM or ASM-QFS control. The procedure for this step differs depending on whether you have a tape drive or an optical drive. The procedure for tape drives is presented first.

For tape drives, enter the following information, where *X* is the raw tape device entry in the mcf file:

server# mt -f /dev/rmt/*X* status

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

server# 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 tapes did not load or the drives did not return a status, the drives may not be listed in the proper order in the mcf. Make sure the order is correct in the mcf and repeat this test.

After modifying the mcf, you must reinitialize the sam-initd daemon. Use the samd stop command to stop sam-initd and the samd start command to start the sam-initd. daemon.

For optical drives, because they are not shared and do not return a status, you should 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 they are configured in the mcf file.

To determine whether the drives become active, you can visually inspect the drives or you can use the samu(1M) utility's o display.

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

Step 17: Label 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*[:*partition*]

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.

partition — A partition on a tape. This is an optional argument. For a partitioned tape, this argument must be specified, and it must be in the range 0 ≤ *partition* ≤ 256. For an nonpartitioned tape, this argument is 0 by default.

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

odlabel –new –vsn new_vsn eq:slot:partition

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.

The following examples show the odlabel(1M) and tplabel(1M) commands being used to label new cartridges:

```
server# tplabel –vsn TAPE01 -new 50:0
server# odlabel –vsn OPTIC01 –new 30:1:1
```

Cartridges are ready to be used after these commands are issued. For more information on these commands, see the tplabel(1M) and odlabel(1M) man pages.

Step 18: Configure the Archiver

By default, the archiver archives all files under all ASM and ASM-QFS mount points. Note that 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. For additional information, see the archiver.cmd(4) man page and see the information on the archiver in the *ASM Administrator's Guide.*

Step 19: Share the File System with Client Machines

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

For example, on the server, enter a line like the following:

server# **share -F nfs -o rw=client1:client2 -d "ASM" /sam**

NOTE

If you write a share(1M) command like the preceding example into the /etc/dfs/dfstab file, Solaris shares the file system after the next system reboot. If you want to share the file system immediately, you must type the share(1M) command at a root shell prompt. If there are no shared file systems when Solaris boots, the NFS server is not started. You must reboot after adding the first share entry to this file.

Some NFS mount parameters can affect the performance of an NFS mounted ASM 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, especially when staging files, ASM 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: Mount the File System on the Client Machines

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

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

```
server:/sam - /sam nfs – yes hard,intr,timeo=60
```

Next, on the command line, issue the mount(1M) command:

```
client# mount /sam
```

The automounter can also do this, if it is preferred. 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 may be a significant delay in the ASM or ASM-QFS file system's response to client requests. This can occur when a requested file resides on a cartridge that must be loaded into a DLT tape drive. If the hard option is not specified, the client can return an error instead of retrying the operation until it completes.

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.

Step 21: Establish Periodic Dumps Using samfsdump(1M)

The server should periodically create a control structure dump using samfsdump(1M). The samfsdump(1M) command supports dumping unarchived data. The –u option on the samfsdump(1M) command causes unarchived data to be interspersed with the control structure data normally contained in a samfsdump(1M) dump. When using the samfsdump(1M) command, please note the following:

- A samfsdump(1M) dump taken under release 3.5.0 using the -u option cannot be restored at earlier releases of the ASM or ASM-QFS software because of new data structures needed to support data inclusion.

- 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 on these commands, see the samfsdump(1M) and ufsdump(1M) man pages.

This dump does not include the data stored in your file system, but it does include information necessary to quickly locate the data on your removable media devices. This information is necessary to recover from a cache disk failure. Use samfsrestore(1M) to restore the control structure dump after initializing the file system if such a failure occurs.

For example, you can make an entry in root's crontab file so the cron daemon runs samfsdump periodically:

**10 0 * * * (find /csd.directory/sam -type f -mtime +3 **

**-print| xargs -l1 rm -f); cd /sam; **

**/opt/LSCsamfs/sbin/samfsdump -f **

/csd.directory/sam/‘date +\%y\%m\%d‘

This example crontab entry uses a 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.

If you have multiple ASM file systems, make similar entries for each. Make sure you save each dump in a separate file.

For more information on using samfsdump(1M), see the samfsdump(1M) man page and see the information on control structures, disaster preparation, and recovery in the *ASM Administrator's Guide.*

Step 22: Establish Periodic Backups of the .inodes File (Optional)

As an alternative to samfsdump(1M), you can periodically save the inode information, which is stored in file .inodes under the mount point. The .inodes file is used to recover a file system after a cache disk failure.

For example, if /sam is the mount point, the following dd(1M) command can be added to the crontab file:

0 0, 8, 16 * * * (find /csd.directory/inodes/sam \
-type f –mtime +3 –print | xargs –l1 rm –f); \
/bin/dd if=/sam/.inodes \
of=/csd.directory/inodes/sam/’date +\%y\%m\%d’ bs=128k

After a disk failure, the backup .inodes file can be used as input to sammkfs(1M). sammkfs(1M) constructs the file system and restores the inodes. All the files, directories, symbolic links, and removable media files are offline. You can back up the .inodes file more frequently since it takes less time than running samfsdump(1M). Specifying the block size in integer multiples of 16 kilobytes increases the performance of the dd(1M) copy (for example bs=128k or bs=512k).

For more information on using the dd(1M) command to back up the inodes, see the dd(1M) man page and see the information on control structures, disaster preparation, and recovery in the *ASM and ASM-QFS System Administration Guide*.

Chapter 6 - ASM and ASM-QFS Upgrade Procedure

This chapter describes upgrading 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.

All steps in this chapter must be performed as superuser (root).

Step 1: Obtain the Release Files

The ASM and ASM-QFS software can be obtained on a CD-ROM or by anonymous FTP. Contact your ASP or StorageTek for information on obtaining the software in one of these ways.

If you have a CD-ROM, run the Solaris Volume Manager, insert the CD-ROM, and change to the directory containing the ASM or ASM-QFS software files using the following command:

server# **cd /cdrom/cdrom0**

WARNING

If you have not read the README file delivered with this release, please do so before continuing. The ASM and ASM-QFS 3.5.0 releases include significant restructuring changes as compared to previous revisions. Failure to recognize these changes could cause dramatic changes in script behavior. The README file is included in the FTP instructions and is on the CD-ROM distribution. After your software is installed, it is located in /opt/LSCsamfs/doc/README.

Step 2: Back 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 samfsdump(1M) or by copying the inodes file using dd(1M), as follows:

Back up each ASM file system using samfsdump(1M). The location to which you dump each system should be outside the ASM file systems.

The following example assumes that you have a file system named samfs1 (mounted at /sam) that you want to back up to samfs1.bak, which exists outside of the ASM file systems:

server# **cd /sam**

server# samfsdump –f /csd_dump_dir/sam.csd

The samfsdump command dumps file names and inode information, not data. For more information on this, see the samfsdump(1M) man page.

If your site is archiving file system data, use the dd(1M) command to write this information to a file for safe keeping. The following example uses the dd(1M) command to write the file system at the /sam mount point:

server# dd if=/sam/.inodes of=/inode_dump_dir/sam.inodes bs=128k

Back up any site-defined scripts. Because of the directory restructuring in the ASM and ASM-QFS 3.5.0 software, any scripts stored in /etc/fs/samfs could be destroyed during an upgrade. Scripts in directories created by the ASM software in an earlier version could also be destroyed. After the ASM or ASM-QFS 3.5.0 software is installed, you can move the scripts to a location in /var/opt.

You need to back up files for each file system, so repeat the preceding steps for each file system in your ASM environment.

For more information on backing up your file systems, see chapter 5, “ASM and ASM-QFS Initial Installation Procedure”.

Step 3: Stop the ASM or ASM-QFS File System

To stop the ASM or ASM-QFS file system, enter the following command:

server# **samcmd idle** *eq* # see **NOTE**

server# samd stop

NOTE

The drives in your ASM or ASM-QFS environment must be idled prior to issuing the samd stop command, so enter a samcmd idle *eq* command for each *eq* configured in your mcf file. Alternatively, 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 on the samcmd(1M) command, see the samcmd(1M) man page.

The samd(1M) command is installed in /opt/LSCsamfs/sbin.

Step 4: Unmount the File Systems

Using the Solaris umount(1M) command, unmount each ASM file system.

If you encounter difficulty unmounting a file system, 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. Use the fuser(1M) command to determine whether or not any processes are still busy. If any are still busy, you must terminate them by using the kill(1M) command. For example, you can enter the following command to determine whether or not processes are still running on the samfs1 file system:

server**# fuser –uc** */mountpoint*

If you are still unable to unmount the file system, issue the Solaris unshare(1M) command on the file system as follows:

server **# unshare** *pathname*

In the preceding format, *pathname* is the path to the ASM file system that you are trying to unmount. After issuing the unshare(1M) command on the path to the ASM file system, try unmounting the file system again. For more information on this, see the unshare(1M) man page.

If all previous attempts to unmount the file system still fail, edit the /etc/vfstab file. When editing this file, change all ASM file systems from yes or delay to no. Then reboot your system.

Step 5: Remove Existing ASM or ASM-QFS Software

Use the pkginfo(1) command, as follows to determine which ASM software packages are installed on your system:

server**# pkginfo | grep LSC**

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 of the optional ASM packages as described at the beginning of this chapter, you should make sure that you remove these packages prior to the main LSCsamfs package. The install script prompts you to confirm several of the removal steps.

The following example removes all of the ASM packages:

server**# pkgrm LSCibm LSCstk LSCdst LSCsony LSCgui LSCjre LSCdoc 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/LSCsamfs/samfs.old.*date*.

Step 6: Add the Packages

ASM software packages use the Solaris packaging utilities for adding and deleting software. As such, 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 ASM packages.

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

To satisfy product dependencies, you must upgrade the sampkg first. Run the pkgadd(1M) command to upgrade all packages, answering yes to each question:

server# **pkgadd –d sampkg** (must be first)

If your ASM or ASM-QFS environment includes certain network-attached automated libraries, such as certain models from StorageTek, Ampex, IBM, or Sony, you may need to install one or more vendor-specific media changer packages. These packages are supplied by ASM. For information on whether or not you need to install a vendor-specific package, see the information on managing vendor-specific automated libraries in the *ASM and ASM-QFS System Administration Guide*.

To install one or more of these packages, install them using pkgadd(1M), as follows:

server# **pkgadd –d samstk** (optional StorageTek package)

server# **pkgadd –d samdst** (optional Ampex package)

server# **pkgadd –d samibm** (optional IBM package)

server# **pkgadd –d samsony** (optional Sony package)

If you want to use the GUI tools included with the ASM and ASM-QFS software package, install them now. The GUI tools require the presence of a Java runtime environment. Add the Java runtime environment and the GUI tool package as follows:

server# **pkgadd –d samjre** (optional)

server# **pkgadd –d samgui** (optional)

If you want to install ASM documentation, install it now. Manuals are available in PDF format. Add this package as follows:

server# **pkgadd –d samdoc** (optional)

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 wish to save to an alternate location.

Step 7: Restore File Changes (inquiry.conf and samst.conf)

ASM does not guarantee correct operation with peripherals other than those included in the /opt/LSCsamfs/examples/inquiry.conf file supplied with the release. The installation script compares this file with an existing one, if any, in /etc/opt/LSCsamfs. If these files differ, the following message is issued:

inquiry.conf has been updated. /opt/LSCsamfs/examples/inquiry.conf is the latest version; please add your changes and copy this file to /etc/opt/LSCsamfs/inquiry.conf.

If you have modified /kernel/drv/samst.conf, you need to merge any changes to it that might be needed for your configuration. The installation script compares this file with an existing one, if any, in /kernel/dev/samst.conf. If these files differ, the following message is issued:

samst.conf has been updated. /opt/LSCsamfs/examples/samst.conf is the latest version; please add your changes and copy it to /kernel/drv/samst.conf. When you have done this, you may need to run /usr/sbin/add_drv samst and /opt/LSCsamfs/sbin/samdev again.

Step 8: Update the License Keys

You must update the license keys for the ASM and ASM-QFS 3.5.0 release. If you are upgrading from an ASM release prior to 3.5.0, you need to place a new license key in the following file:

/etc/opt/LSCsamfs/LICENSE.3.5

To obtain new license keys, contact your ASP or StorageTek.

For more information see the licensing information in chapter 5, “ASM and ASM-QFS Initial Installation Procedure”.

Step 9: Verify the New Master Configuration File

The topology of the equipment managed by the ASM or ASM-QFS file system is defined in the master configuration file, /etc/opt/LSCsamfs/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.

Verify that a new mcf file exists in /etc/opt/LSCsamfs/mcf.

NOTE

For information on file system design considerations, see the *ASM File System Administrator’s Guide*.

Step 10: Modify the /etc/vfstab File (Optional)

If you modified the /etc/vfstab file in “Step 4: Unmount the File Systems”, edit this file again and change all ASM file systems from no to yes or delay.

Step 11: Reboot the System

You must reboot the system at this time using the reboot(1M) command.

Rebooting is particularly important if you are upgrading from a ASM or ASM-QFS 3.3.1 release on Solaris 2.7, which is a 32-bit release, and are now upgrading to a 64-bit release. The reboot is necessary to build the 64-bit kernel.

Step 12: Mount the File System(s) (Optional)

You must perform this step if you have not modified the /etc/vfstab file to have yes or delay.

Use the mount(1M) command to mount the file systems and continue operation with the upgraded ASM software.

In the following example, samfs1 is the file system name to be mounted:

```
server# mount samfs1
```

Step 13: Relink API-dependent Applications (Optional)

If you are running applications that use the ASM application programmer interface (API) and you are using static linking, you should relink these applications at this time.

Appendix - Requesting Help from Software Support

NOTE

Information for this Appendix was taken from StorageTek's Software Support Manual, Part Number 112124005.

In this Appendix, "software problem" refers to problems with and/or questions concerning both the software and the supporting documentation.

About This Appendix

This Appendix outlines the customer Software Support services provided by StorageTek's technical support organizations, and details the process and guidelines for getting help from StorageTek for problems with StorageTek software.

1. Contacting Software Support and Other Services

Around the world, StorageTek Software Support experts are committed to helping you use your software as effectively as possible. StorageTek can be contacted for software support problems via the following methods:

U.S. and Canada

Telephone	800-678-4430 or 303-678-4430
Fax	303–673–8015 or 303–661–4772
World Wide Web	http://www.storagetek.com or www.support.storagetek.com
E–mail	softsupt@louisville.stortek.com

Learning Products Organization

Telephone	303–673–6262
World Wide Web	http://wfd.stortek.com/

International

To obtain a list of International Subsidiary and Distributor Locations and phone numbers, use the following URL:

World Wide Web http://www.storagetek.com/global

International customers should contact their local Distributor or Subsidiary office for software support. Additionally, international customers may contact the U.S. Software Support center for assistance in the event that local Distributor or Subsidiary support cannot be reached.

Hours of Operation

Software Support operates from StorageTek headquarters in Louisville, Colorado. StorageTek provides software support seven days a week, 24 hours a day, 365 days a year.

Eligibility for Software Support

You are entitled to receive support for all software products for which there is a current maintenance contract. StorageTek's Software Support provides assistance for software problems or for answering questions based on your contract.

The hours that you can call Software Support are based on this contract. For example, you may choose to contract for Software Support services only between the hours of 8:00 a.m. – 5:00 pm, Monday through Friday, or contract for the full coverage of seven days a week, 24 hours a day, 365 days a year.

Determining Technical Severity for Problem Reports

Severity codes are assigned to each software problem reported to Customer Support. The following guidelines are used by Software Support in assigning a severity level to your problem.

Severity Level 1–The software product is non-operational, resulting in a critical system condition requiring immediate resolution. Support personnel will require continuous availability of a qualified customer technical contact until a circumvention or resolution is provided.

Severity Level 2–The software product is operational, but with severely restricted functionality or system degradation.

Severity Level 3–The software product is operational, with functional limitations or restrictions that are not critical to the overall system operations.

Severity Level 4–Questions associated with product usage, implementation, performance, or any other inquiries for the Software Solutions Support Organization.

NOTE
Software Support may downgrade or upgrade a severity code after analysis of your software problem.

Customer Resource Center (CRC) Web Page

Through the CRC (http://www.support.storagetek.com), StorageTek has developed the following Web tools to assist you in being more productive, as well as assist you in resolving problems or answering your product and application questions.

Software support through the Online Request Form

Online software product information

Online product user guides

Frequently asked questions and technical tips

Online software fixes

Online ordering of software documentation

Online customer training information and registration

2. Providing Help and Technical Assistance

This section defines how Software Support operates to achieve customer satisfaction.

Scope of Support

The goal of Software Support is to assist you with the use of StorageTek software by answering questions and resolving problems specifically related to the software. The underlying assumption is that you are generally knowledgeable about the product but have additional questions or problems not answered in the documentation or training.

However, there are several situations which are beyond the scope of Software Support. Some examples are:

Your need for product training

Code or scripts that you have written or modified

Fixes to prior releases of software versions that are no longer supported

Your computing environments, such as network configuration

Operations and performance tuning

Your user code exits

For problems where modified code or user code exits are involved, Software Support will confirm that the linkage between the StorageTek program/subroutine and your code is set up properly. Problems encountered thereafter that are not reproducible in the StorageTek source programs, and are attributed to your application or other software interfaces, are outside the scope of Software Support. You will be referred to our Consulting Services, who can provide assistance on either a fee basis or charge on a time and material basis.

Debugging

In most cases, problems and code defects can be fixed through standard software support. It is possible, however, for Software Support to determine that a product problem is caused by conditions at your site rather than by a defect in the StorageTek product. In this case, continued support can only be provided through Consulting Services on a fee basis or a time and material basis.

Support of Product Releases

StorageTek continues to enhance and modify its products to meet the changing demands and requirements of the markets it serves. StorageTek provides product phone support for all releases of software until we notify you that we are discontinuing support. However, we will only make fixes to reported software problems for the current supported software releases. StorageTek mails notification letters well in advance of the non–support date.

Platform Environments

StorageTek software is warranted to perform as documented on hardware platforms running versions/releases of operating systems that are currently supported by their manufactures.

StorageTek will not assign software support resources to correct technical problems for systems that the manufacturer does not currently support. In these cases, you will be referred to Consulting Services.

Operating System Support Policy

StorageTek supports its software products on operating systems that are:

Currently supported and maintained by hardware vendors, and

StorageTek certified products as defined in the StorageTek product documentation.

When an operating system is no longer supported by the hardware vendor, StorageTek will honor current maintenance contracts, and provide ample notice for any operating system for which we will discontinue support

Support of User Modified Code

StorageTek programs that you have modified, such as user exits and script files, will not be supported by Software Support. You will be referred to Consulting Services for assistance with problems related to your modified code.

3. Product Training and StorageTek Support

StorageTek Software Support has found that it cannot provide product training over the telephone. It is neither effective nor efficient to train this way. We make every effort to devote support time to the exclusive purpose of addressing problem issues and answering specific application questions. Software Support is not a substitute for product training.

StorageTek's Learning Products Organization offers product training for all of its products. If you are unfamiliar with basic product concepts and need product training, you will be referred to the StorageTek Learning Products Organization.

You can contact StorageTek's Learning Products Organization through the following:

Telephone	303–673–6262
World Wide Web	Web http://wfd.stortek.com/

4. Reporting Your Problem

StorageTek's Software Support provides assistance for software problems or for answering questions that are not covered in the product manuals. This section describes the steps to take in reporting your software problem to StorageTek. You should perform the following steps for each software problem you encounter:

a. Use the Customer Resource Center (CRC).

First, try to resolve the problem by using the CRC located at http://www.support.storagetek.com. The CRC provides you with the following types of resources:

Product User Guides

Release Notes

Frequently Asked Questions

Maintenance/Fixes

b. Be prepared with information about your problem.

If the online WEB resources do not provide a solution, be prepared to provide Software Support the following information:

c. Customer Information

Issue number if you are working on an open issue or a recurring problem

Site location number

Company name

Your name and phone number.

Alternate contact number

d. System Information

Software product release version numbers and maintenance level

Hardware platform and operating system level

System configuration information, such as peripheral devices and network configuration

e. Problem Information

What were you trying to do? For example, what command were you using?

Were there error messages, application logs and trace files, or system error logs? If yes, what were they?

Is the problem new or recurring? Can the problem be reproduced? If yes, under what conditions?

Were there any changes made to the system recently? If yes, did the problem occur prior to the change?

f. Open a software issue with Software Support.

Please have your product manuals available, and be prepared to provide any other information that may help Software Support assist you. Software Support recommends that you designate a single technical contact, by product group, to communicate with StorageTek Software Support.

There are two methods for opening an issue with Software Support:

Use the StorageTek Customer Resource Center (CRC) online request form for non–Severity 1 issues. The online request form can be accessed through the web at http://www.support.storagetek.com and clicking Product and CRC Support. Select the Online Request Form.

Call Software Support at 1–800–678–4430 or 303–673–4430.

g. Open a separate software issue for each software problem.

Provide Software Support with a detailed description and information about each software problem.

h. Record the assigned issue number for future reference.

All of the problem details reported are entered into Software Support's problem tracking database. A unique issue number is then assigned.

i. Software Support may ask for product specific diagnostic material or additional documentation.

Collect the additional information and use one of the methods described in "Sending Documentation, Error Logs/Trace Files, or Problem Records" on page 15 to send the information to StorageTek's Software Support Center. Always reference the issue number and the support representative's name that requested the information on all supporting items of documentation. For envelopes or packages, the issue number should be noted in the lower–left corner.

j. If, at any time, you are unhappy with the support provided by StorageTek Software Support, please contact Software Support and ask to speak to a Software Support Manager.

5. Responding to and Resolving Customer Calls

Call response time and problem resolution time are non–related.

Call response time is the amount of time it takes a StorageTek support representative to return your initial call if your call cannot be taken immediately.

As a guideline, the support representative will return your call within 15 minutes.

Problem resolution time refers to the total amount of time between when StorageTek receives your initial call and when StorageTek provides a resolution to your problem.

If we have to call you back, or your problem has been escalated to StorageTek's Engineering Solutions Support Team for advanced diagnosis, the following occurs:

We will make three attempts to contact you using primary and alternate telephone numbers and contacts.

Your issue will be closed if after three tries we cannot reach you or your alternate contact and we have not received a callback. Software Support will re–open your issue if you call back after your issue is closed.

6. Problem Resolution

Some software problems can be resolved immediately when you call Software Support and some problems require more time. While it is our intent to resolve calls as quickly as possible, we cannot project exact resolution times for each call. Resolution times are dependent on the following:

Complexity of the problem

Speed with which we receive the requested problem documentation and logs

Number of calls received by Software Support

Immediate Resolution

A Software Support representative will take your initial software support call. The support representative checks the problem tracking database to determine whether the problem was previously reported by other customers. If the software problem was previously reported and resolved, the support representative may have an immediate answer to your problem. Depending on your product, a work-around or PTF/Patch may be available.

If the problem was previously reported by other customers, but not yet resolved, the support representative links your problem issue to the already existing issue. A separate issue number for your specific software problem is also maintained. When a resolution becomes available, you are then provided the resolution.

When Investigation Is Needed

If your software problem cannot be resolved immediately, the support representative escalates your problem to an Engineering Solutions Support Team (SST). An SST engineer will contact you to further diagnose and resolve your problem. In some cases, problem diagnosis may require that SST have access to your system through a modem or the internet.

The SST engineer takes responsibility for researching the problem and determining whether a solution is available. Depending on your product, a work-around or PTF/Patch may be available. If there is no work–around, PTF/Patch, or resolution available, new code may need to be generated to resolve the issue.

If you are operating on a down level release of software, you may be required to move to the most current release or wait for the next product release to resolve your problem.

Closing Your Issue

Once a proposed issue resolution is provided, StorageTek closes the issue. If you require further explanation of the resolution, contact Software Support.

NOTE

International customers should contact their local software support representative.

Recurring Software Problems

If the issue has been closed with Software Support, yet you experience the same software problem, the issue can be reopened. Give Software Support the issue number of the closed issue and state that the problem has recurred.

If the fix you received creates a new problem, you must open a new issue. Give Software Support the issue number of the closed issue and state that the fix created a new problem. You will be given a new issue number to use for future reference**.**

7. Sending Documentation, Error Logs/Trace Files, or Problem Records to Software Support

Use the following methods to send problem information to StorageTek Software Support.

NOTE

International customers should follow instructions given by the local support representative when dealing directly with the local Distributor or Subsidiary.

MVS Software

Use the following procedure for reporting MVS software problems:

Place information on an 18–track or 36–track Standard Label tape.

Use IEBGENER and/or IEBCOPY to place the datasets on the tape, unless instructed by your support representative to use a different utility.

Include the StorageTek problem issue number on any documentation sent.

If you can only send a non–labeled tape, you must also send a copy of the DCB information so that Software Support can offload the tape.

NOTE

Failure to send the DCB information slows down the process of problem determination. StorageTek no longer returns documentation tapes.

VM Software

Use the following procedure for reporting VM software problems:

Place datasets on an 18–track tape using VMFPLC2 format.

Include the StorageTek problem issue number on any documentation sent.

NOTE

StorageTek no longer returns documentation tapes.

Other Software

Email

Small log/trace files can be emailed to:

ssrhql@blackcat.storagetek.com

NOTE

Be sure to include the StorageTek problem issue number as part of the file name.

FTP

Log/trace files can be FTP'd to:

ftp.storagetek.com

Login: Support

Password: <contact Software Support for password>

Directory = /incoming

Mail or Overnight

To send mail or overnight shipment, such as FED–EX, Airborne, or UPS:

Storage Technology Corp.

Attn: Software Support

issue # _____

One StorageTek Drive Louisville, Colorado 80028–4280

Fax

Small logs can be faxed to:

303–673–8015

Storage Technology Corp.

Attn: Software Support

issue # _____

Glossary

addressable storage

The storage space encompassing online, nearline, and offline storage that is user referenced through an ASM file system.

archiver

The archive program that automatically controls the copying of files to removable cartridges.

archive storage

Copies of file data that have been created on removable cartridges for long-term offline storage.

ASM

The ASM File System. 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 ASM with the QFS 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 containing a number of pseudo devices (/samdev/rd). It may or may not have its own library devices. The client depends on an ASM Remote server for cartridges.

ASM Remote daemon

A process initiated by the ASM Remote client that establishes and controls the network connection between ASM Remote and the ASM or ASM-QFS server. This daemon, named sam-clientd, also establishes pseudo device connections to be used for data transfer.

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.

audit (full)

The process of reading the VSNs from each cartridge in an automated library. For non-tape cartridges, the capacity and space information is determined and entered into the automated library's catalog.

automated library

See library.

backup storage

A snapshot of a collection of files for the express purpose of preventing inadvertent loss. A backup includes both the file's attributes and associated data.

block allocation map

A bit map representing each available block of storage on a disk and indicating whether the block is in use or free.

cartridge

The 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 request 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.

data device

For a file system, a device or group of devices upon which file data is stored.

data space

The portion of a collection of files that is the actual data information.

DAU (Disk Allocation Unit)

The basic unit of online storage.

The ASM file system uses several sizes. The small DAU is 4 kilobytes (217 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.

The ASM-QFS file systems support a fully adjustable DAU, sized from 16 kilobytes through 65528 kilobytes. The DAU you specify must be multiple of 8 kilobytes.

device logging

A feature that provides device-specific error information used to analyze device problems.

device scanner

Software within the ASM file system that periodically monitors the presence of all manually mounted removable devices and detects the presence of mounted cartridges that may be requested by a user or other process.

devicetool

An ASM and ASM-QFS administrative tool with a graphical user interface for viewing information about and managing individual devices.

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.

disk allocation unit

See DAU.

disk buffer (also called 'cache')

The disk buffer is used to buffer files when writing data to the ASM Remote server. This is frequently referred to as *cache*, but it is not to be confused with *disk cache*.

disk cache

The disk cache is used by ASM and ASM-QFS to create and manage data files between online disk and removable cartridges. Individual disk partitions or an entire disk can be used as disk cache.

disk cache family set

The definition for the devices that make up a family set. The name of the disk cache family set is found in the equipment identifier field of the Master Configuration File (mcf file). This is sometimes referred to as a *metadevice* in industry literature. Also see family set.

disk striping

The process of recording a file across several disks, thereby improving access performance and increasing overall storage capacity.

direct access

A file attribute (stage never) designating that a nearline file can be accessed directly from the archive cartridges and need not be staged for online access.

directory

A file data structure that points to other files and directories within the file system.

disk space thresholds

User-defined disk space thresholds that define 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 the pre-defined disk space thresholds.

drive

A mechanism for transferring data to and from a cartridge.

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 megabyte-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.

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 mounted within an automated library.

Also see disk cache family set.

FDDI

Fiber Distributed Data Interface. FDDI is a 100 megabytes-per-second fiber optic LAN.

file system-specific directives

Directives that follow global directives 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.

file system

A hierarchical collection of files and directories.

FTP

File Transfer Protocol. An internet protocol for transferring files between two hosts over a TCP/IP network.

global commands

Commands that apply to all file systems and appear before the first “fs = ” line.

indirect block

A disk block that contains a list of storage blocks. The ASM 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.

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 inodes are 512 bytes long. The inode file is a metadata file, which is separated from file data in the ASM-QFS file systems.

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.

LAN

Local Area Network.

library

A robotically controlled device designed to automatically load and unload removable media cartridges without operator intervention. An library contains one or more drives and a robot that moves cartridges to and from the storage slots and the drives.

library catalog

See catalog.

LUN

Logical Unit Number.

mcf

Master Configuration File. The file that is read at initialization time that defines the relationships between the devices (the topology) within an ASM and ASM-QFS environment.

media

Tape or optical disk cartridges.

media recycling

The process of recycling or reusing archive cartridges with low use (that is, archive cartridges with few active files).

metadata

Data about data. The index information needed to locate the exact data position of a file on a disk. Metadata contains information pertaining to the directory, symbolic link, removable media, segmented file index, and .inodes.

metadata device

A separate device (for example a solid-state disk or mirrored device) upon which 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 disjoint sets of disks to prevent loss from a single disk failure. It is often referred to as shadowing.

mount point

The path to a directory where a file system is mounted.

name space

The portion of a collection of files that identifies the file, its attributes, and its storage locations.

nearline storage

Removable 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 network-attached automated library, such as those from ASM, ADIC/Grau, IBM, or Sony, 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 media changer daemon specifically designed for the automated library.

NFS

Network File System. A standard protocol that allows a UNIX file system to be remotely mounted via a network.

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).

partition

A portion of a device.

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 only be performed on zero-sized files. That is, the setfa –l command can only be specified for a file that is size zero. For more information, see the setfa(1) man page.

prioritizing preview requests

A method of assigning priority to archive and stage requests that cannot be immediately satisfied.

pseudo device

A network connection to an actual device on the ASM or ASM-QFS server.

RAID

Redundant Array of Inexpensive/Independent Disks. A disk technology that uses several inexpensive disks to reliably store files. It may protect against data loss from a single disk failure, may provide a fault-tolerant disk environment, and may provide higher throughput than individual disks.

recycler

An ASM and ASM-QFS component that reclaims space on cartridges that is occupied by unused 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.

robot

The portion of an library that moves cartridges between storage slots and drives.

robottool

An ASM and ASM-QFS administrative tool with a graphical user interface (GUI) for viewing and managing automated libraries.

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 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 striping.

RPC

Remote Procedure Calls. The underlying data exchange mechanism used by NFS to implement custom network data servers.

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 copy data.

samfsrestore

A program that restores a control structure dump.

samtool

An ASM and ASM-QFS administrative tool with a GUI for invoking robottool, devicetool, and previewtool.

SCSI

Small Computer System Interface. An electrical communication specification commonly used for peripheral devices such as disk and tape drives and automated libraries.

SCSI-attached Library

An automated library connected directly to a server using the SCSI interface. These libraries are controlled directly by the ASM or ASM-QFS software by using the SCSI standard for automated libraries.

shared writer/shared reader

The ASM-QFS shared reader/shared writer capability allows you to specify a file system that can be shared by multiple servers. Multiple hosts can read the file system while only one host can write to the file system. Shared readers are specified with the –o shared_reader option on the mount(1M) command. The one-writer host is specified with the –o shared_writer option on the mount(1M) command. For more information on the mount(1M) command, see the mount_samfs(1M) man page.

small computer system interface

See SCSI.

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. 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 file system and defined in the mcf file as 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.

striping

A data access method in which files are simultaneously written to logical disks in an interlaced fashion.

All ASM file systems allow you to declare either striped or round robin access for each individual file system. The ASM-QFS file systems allow you to declare striped groups within each file system.

Also see the glossary entry for round robin.

super block

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.

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.

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

Allows 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. A logical identifier for magnetic tape and optical disk that is written in the volume label.

WORM

Write Once Read Many. A storage classification for media that can be written only once but read many times.

Index