



File System Manager User's Guide v 3.0

Sun Microsystems, Inc.
www.sun.com

Part No. 819-7939-10
April 2007, Revision A

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

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

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

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

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

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

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

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

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

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

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

DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID.

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

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

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

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

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

Sun, Sun Microsystems, le logo Sun, AnswerBook2, docs.sun.com, Solaris, Java, Solstice DiskSuite, SunPlex, et Sun StorageTek sont des marques de fabrique ou des marques déposées de Sun Microsystems, Inc. aux Etats-Unis et dans d'autres pays.

Toutes les marques SPARC sont utilisées sous licence et sont des marques de fabrique ou des marques déposées de SPARC International, Inc. aux Etats-Unis et dans d'autres pays. Les produits portant les marques SPARC sont basés sur une architecture développée par Sun Microsystems, Inc.

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

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

LA DOCUMENTATION EST FOURNIE "EN L'ÉTAT" ET TOUTES AUTRES CONDITIONS, DECLARATIONS ET GARANTIES EXPRESSES OU TACITES SONT FORMELLEMENT EXCLUES, DANS LA MESURE AUTORISEE PAR LA LOI APPLICABLE, Y COMPRIS NOTAMMENT TOUTE GARANTIE IMPLICITE RELATIVE A LA QUALITE MARCHANDE, A L'APTITUDE A UNE UTILISATION PARTICULIERE OU A L'ABSENCE DE CONTREFAÇON.



Contents

Preface xv

- 1. About the File System Manager Software 1**
- 2. How the Software Works 2**
- 3. Assigning Administrative Privileges to Users 3**
- 4. Logging In to and Out of the Software 6**
- 5. Setting the Session Timeout Value 7**
- 6. Registering the Software 7**
- 7. About Interface Elements 9**
 - About the Banner and Content Pane 10
 - Reference: Banner Elements 11
 - Reference: Content Pane Navigational Elements 12
 - Reference: Content Pane Icons and Buttons 13
- 8. About Wizards 15**
- 9. Enabling Pop-Up Windows 15**

10. Changing the Current Server	16
11. About Report Tables	16
Reference: Report Table Controls	17
Sorting Tables	18
Filtering Tables	19
12. Using Help	19
13. Task Overview: Configuring a Server	23
14. Reference: Server Configuration Files	25
15. Adding a Server	26
16. Removing a Server	28
17. Monitoring Servers	31
Displaying the Configuration of Servers in Your Site	32
Displaying the Configuration of a Server	32
Displaying Software Packages on a Server	33
Displaying Log or Trace Files on a Server	33
Displaying Configuration Files on a Server	34
Displaying Warning Messages in Server Files	34
Displaying the Minimum Server Release Supported by File System Manager	35
18. Monitoring File Systems	35
Displaying the Status of a File System	36
Displaying File System Devices	36
Displaying NFS Shared Directories in a File System	37
Displaying File Match Criteria for an Archiving File System	37
Displaying the Status of Recovery Point Activities	38

Displaying the Status of a Restore Process for an Archiving File System	38
Displaying the Metadata Server for a Shared File System	39
Displaying All Hosts in a Shared File System	39
About System Metrics	40
Media Utilization and Status Page	40
File Data Distribution Page	41
File System Utilization Page	42
Creating Media Utilization and Status Reports	42
Viewing File Data Distribution	43
Creating File System Utilization Reports	43
19. Monitoring Policies	44
Displaying File Match Criteria in a Policy	44
Displaying File Systems Using Specific File Match Criteria	45
Displaying the Copy Details of a Policy	45
Displaying the VSNs Assigned to a Policy Copy	46
20. Monitoring Libraries	46
Displaying VSNs for a Library	47
Displaying the Drives in a Library	47
Displaying the Status and State of a Library	48
Displaying Historian Catalog Information	48
Displaying Mount Requests in a Library	49
21. Monitoring Archiving, Releasing, and Staging Activity	49
Displaying Archiving Activity	50
Displaying Releasing Activity	50
Displaying Staging Activity	51
22. Monitoring Jobs	51

	About Jobs	52
	Displaying Jobs	53
	Canceling a Job	54
	Using the Monitoring Console	55
23.	Monitoring Faults	56
	About Faults	56
	Displaying Faults	57
	Acknowledging Faults	58
	Deleting Faults	58
24.	About File System Management	59
25.	About Discovering Available Devices	60
26.	About File Systems	61
	About Supported File Systems	62
	About Stand-Alone, Non-archiving File Systems	63
	About Archiving File Systems	65
	About Shared File Systems	66
	About HA-QFS File Systems	67
27.	Administering Stand-Alone, Non-archiving File Systems and Archiving File Systems	68
	Planning a Stand-Alone, Non-archiving or Archiving File System	68
	Creating a Stand-Alone, Non-archiving or Archiving File System	70
	Growing a Stand-Alone, Non-archiving or Archiving File System	71
	Checking a Stand-Alone, Non-archiving or Archiving File System	72
	Mounting a Stand-Alone, Non-archiving or Archiving File System	73
	Editing Mount Options for a Stand-Alone, Non-archiving or Archiving File System	73
	Unmounting a Stand-Alone, Non-archiving or Archiving File System	74

	Deleting a Stand-Alone, Non-archiving or Archiving File System	75
	About Creating an Archiving File System and VSN Association	75
28.	Managing Policies for Archiving File Systems	76
	Adding a Policy to an Archiving File System	77
	Adding Existing File Match Criteria to an Archiving File System	77
	Reordering File Match Criteria for an Archiving File System	78
	Removing File Match Criteria From an Archiving File System	79
29.	Protecting Archiving File System Data	80
	About Protecting Archiving File System Data	80
	Scheduling and Creating Recovery Points	81
	About Recovery Points	81
	Planning Recovery Points	82
	Scheduling the Creation of Recovery Points	83
	Creating a Recovery Point on Demand	84
	Editing a Recovery Point Schedule	85
	Removing a Recovery Point Schedule	85
	Restoring Archiving File System Data	86
	Restoring an Archiving File System	86
	Filtering Recovery Point Files	90
	Filter File Name Pattern Values	91
	Deleting the Index for a Recovery Point	92
	Deleting a Recovery Point	93
30.	Administering Shared File Systems	95
	Planning a Shared File System	95
	Creating a Shared File System	96
	Creating the Shared Hosts File	97
	Adding a Client or Potential Metadata Server	98

Deleting a Client or Potential Metadata Server	99
Deleting the Metadata Server	100
Mounting and Unmounting Shared File Systems	100
Mounting a Shared File System	101
Mounting the Metadata Server	102
Mounting a Client or Potential Metadata Server	102
Editing Mount Options for a Shared File System	103
Unmounting a Shared File System	103
Unmounting a Client or Potential Metadata Server	104
Unmounting the Metadata Server	105
Deleting a Shared File System	105
31. Managing NFS Sharing for File Systems	106
About NFS Sharing	106
Adding an NFS Directory	107
Editing NFS Directory Options	108
Changing the Status of an NFS Directory	109
Removing an NFS Directory	110
32. About Archive Management	111
33. About Archiving, Releasing, and Staging	113
About the Sun StorageTek SAM Software	113
About Archiving	114
Planning for Archiving	115
About Continuous Archiving	117
About Associative Archiving	118
About the Releaser	119
About the Stager	119

34. Controlling Archiving for Archiving File Systems	120
Configuring Archiver Scanning for an Archiving File System	120
Stopping or Idling Archiving for an Archiving File System	121
Running Archiving for an Archiving File System	121
35. Managing Archive Policies	122
Administering Archive Policies	122
About Archive Policies	123
Reference: Creating a Policy <code>archiver.cmd</code> Example	126
Creating a Policy	127
Creating a No-Archive Policy	128
Editing a Policy	128
Deleting a Policy	129
Enabling Archiving Options in a Policy	129
Enabling Disk Archiving	130
For a New Policy	130
For an Existing Copy in an Existing Policy	130
For a New Copy in an Existing Policy	131
Enabling Archiving to Tape	131
For a New Policy	131
For an Existing Copy in an Existing Policy	131
For a New Copy in an Existing Policy	132
Enabling Associative Archiving	132
Enabling Continuous Archiving	133
Enabling Archiving to Pools of Disk or Tape VSNs	134
For a New Policy	134
For an Existing Copy in an Existing Policy	134
For a New Copy in an Existing Policy	134
Administering File Match Criteria	135

	About File Match Criteria	135
	Editing File Match Criteria	136
	Adding a File System to a Set of File Match Criteria	136
	Removing a File System From a Set of File Match Criteria	137
	Planning File Match Criteria	138
	Adding File Match Criteria to a Policy	138
	Removing File Match Criteria From a Policy	139
	Administering Policy Copies	139
	About Policy Copies	140
	Editing a Policy Copy	140
	Modifying VSN Assignments in a Policy Copy	140
	Adding a Copy to a Policy	141
	Removing a Copy From a Policy	141
36.	Managing VSN Pools for Archiving	142
	About VSN Pools	142
	Creating a VSN Pool	143
	Editing a VSN Pool	143
	Deleting a VSN Pool	144
37.	Managing Disk VSNs for Archiving	145
	About Archiving to Disk-Based VSNs	145
	Creating a Disk VSN	146
	Creating an Sun StorageTek 5800 VSN	147
	Editing the Media Attributes of a Disk VSN	147
38.	Setting Up General Archiving	148
	About Archiver Scanning Methods	148
	Configuring Global Archiving Values	149
	Configuring Global Staging Values	150

	Configuring Global Releasing Values	150
	Configuring Global Media Type Parameters	151
39.	Configuring the Recycler	152
	About the Recycler	152
	Selecting Eligible Volumes	152
	Recycling Volumes	153
	Running Recycling	154
	Configuring Library Recycling Values	154
	Configuring Policy Recycling Values	155
40.	Managing Archiving and Staging Activity	156
	Controlling Global Archiving Activity	156
	Controlling Global Staging Activity	157
	Staging a File or Directory on Demand	157
41.	About Removable Media Management	159
42.	About Removable Media	160
43.	About Supported Libraries and Drives	161
44.	About the Historian	162
45.	Managing Libraries	163
	About Libraries	163
	Adding a Library	164
	Changing the State of a Library	165
	Unloading a Library	166
	Removing Archiving, Staging, and Recycling Configurations	166
	Removing Archiving and Staging Configurations	166
	Removing Recycling Configurations	167

Deleting a Library	167
Managing Library Drives	168
Changing the State of a Library Drive	168
Idling a Drive	169
Unloading a Drive	169
Cleaning a Drive	170
Managing VSNs in a Library	170
Searching for a VSN	171
Importing a VSN	171
Exporting a VSN	172
Loading a VSN Into a Drive in a Library	173
Unloading a VSN From a Drive in a Library	173
Labeling a VSN	174
Reserving a VSN in a Library	175
Unreserving a VSN in a Library	176
Auditing the Slot for a VSN in a Library	176
Editing the Media Attributes for a VSN in a Library	177
46. About Email Alerts	179
47. Creating an Email Alert	181
48. Changing Email Alert Events for a Recipient	181
49. Deleting an Email Alert	182
50. Using SAMreports	185
About SAMreports	186
Generating SAMreports	186
Displaying SAMreports	187

- 51. Incompatible Software Release 187
- 52. One Device Displayed Multiple Times 188
- 53. Inaccessible Server 189
- 54. No Remaining Space for a VSN 190
- 55. Damaged or Stale Files in a Recovery Point 191
- 56. Existing Fault Conditions Not Displayed 191

Glossary 193

Index 207

Preface

The *File System Manager User's Guide* provides detailed procedures for using the File System Manager browser-based graphical user interface. This document is written for system administrators and users who use File System Manager to configure, control, protect, and monitor one or more file systems in the network from a central location.

Before You Read This Document

This document assumes that the management station software and the Sun StorageTek QFS or Sun StorageTek Storage Archive Manager (SAM) software packages are already installed.

For installation procedures, see the *Sun StorageTek QFS Software Installation and Upgrade Guide* and the *Sun StorageTek Storage Archive Manager (SAM) Software Installation and Upgrade Guide*.

Using UNIX Commands

This document might not contain information about basic UNIX® commands and procedures such as shutting down the system, booting the system, and configuring devices. Refer to the following for this information:

- Software documentation that you received with your system
- Solaris™ Operating System documentation, which is at:

<http://docs.sun.com>

Shell Prompts

Shell	Prompt
C shell	<i>machine-name%</i>
C shell superuser	<i>machine-name#</i>
Bourne shell and Korn shell	\$
Bourne shell and Korn shell superuser	#

Typographic Conventions

Typeface*	Meaning	Examples
AaBbCc123	The names of commands, files, and directories; on-screen computer output	Edit your <code>.login</code> file. Use <code>ls -a</code> to list all files. % You have mail.
AaBbCc123	What you type, when contrasted with on-screen computer output	% su Password:
<i>AaBbCc123</i>	Book titles, new words or terms, words to be emphasized. Replace command-line variables with real names or values.	Read Chapter 6 in the <i>User's Guide</i> . These are called <i>class</i> options. You <i>must</i> be superuser to do this. To delete a file, type <code>rm filename</code> .

* The settings on your browser might differ from these settings.

Related Documentation

This manual is part of a set of documents that describes the operations of the Sun StorageTek QFS and Sun StorageTek SAM software products. [TABLE P-1](#) shows the complete release 4U6 documentation set for these products.

TABLE P-1 Related File System Manager Documentation

Title	Part Number
<i>Sun StorageTek QFS File System Configuration and Administration Guide</i>	819-7935-10
<i>Sun StorageTek Storage Archive Manager (SAM) File System Configuration and Administration Guide</i>	819-7934-10
<i>Sun StorageTek Storage Archive Manager (SAM) Archive Configuration and Administration Guide</i>	819-7931-10
<i>Sun StorageTek Storage Archive Manager (SAM) Installation and Upgrade Guide</i>	819-7932-10
<i>Sun StorageTek QFS Installation and Upgrade Guide</i>	819-7936-10
<i>Sun StorageTek QFS Linux Client Guide</i>	819-7937-10
<i>Sun StorageTek Storage Archive Manager (SAM) Troubleshooting Guide</i>	819-7933-10
<i>Sun StorageTek QFS and Sun StorageTek SAM 4.6 Release Notes</i>	819-7938-10

If you are configuring a File System Manager file system in a Sun Cluster environment, the following additional documents might interest you:

- *Sun Cluster Concepts Guide for Solaris OS*
- *Sun Cluster Software Installation Guide for Solaris OS*
- *Sun Cluster Data Services Planning and Administration Guide for Solaris OS*
- *Sun Cluster Data Service for Oracle Real Application Clusters Guide for Solaris OS*

Accessing Sun Documentation From docs.sun.com

This web site contains documentation for Solaris and many other Sun software products.

1. **Go to the following URL:**

`http://docs.sun.com`

The docs.sun.com page appears.

2. **Find the documentation for your product by searching for File System Manager in the search box.**

Third-Party Web Sites

Sun is not responsible for the availability of third-party web sites mentioned in this document. Sun does not endorse and is not responsible or liable for any content, advertising, products, or other materials that are available on or through such sites or resources. Sun will not be responsible or liable for any actual or alleged damage or loss caused by or in connection with the use of or reliance on any such content, goods, or services that are available on or through such sites or resources.

Sun Welcomes Your Comments

Sun is interested in improving its documentation and welcomes your comments and suggestions. You can submit your comments by going to:

<http://docs.sun.com/app/docs/form/comments>

Please include the title and part number of your document with your feedback:

File System Manager User's Guide, part number 819-7939-10

Product Overview

This chapter provides an overview of the File System Manager. It contains the following sections:

- [“About the File System Manager Software” on page 1](#)
- [“How the Software Works” on page 2](#)
- [“Assigning Administrative Privileges to Users” on page 3](#)
- [“Logging In to and Out of the Software” on page 6](#)
- [“Setting the Session Timeout Value” on page 7](#)
- [“Registering the Software” on page 7](#)

About the File System Manager Software

The File System Manager software is a browser-based graphical user interface that enables you to configure, control, protect, and monitor one or more file systems in your network from a central location. To access this central location, you can use the web browser on any host in your network.

The goal of the software is to provide a less complex way than the command line of performing the most common tasks associated with these file systems. To configure options that are unavailable through the browser interface, use the command-line interface, configuration files, and so on, that are associated with the file systems.

Note: If you manually edit configuration files outside of the browser interface, you are responsible for any errors that might occur due to incorrect configuration.

How the Software Works

The File System Manager software is installed on a server called the management station. The name of the management station is always displayed in the banner of the browser interface.

When you first run the browser interface, you add servers upon which the file systems that you want to administer will reside. You can add the management station to the browser interface to administer file systems on the management station.

Note: To add a server to the browser interface, you must install the appropriate packages locally. For example, if you want to manage archiving file systems through the browser interface, you must have the `SUNWsamfsr` and `SUNWsamfsu` packages installed locally. For information about Sun StorageTek™ Storage Archive Manager (SAM) and Sun StorageTek QFS software releases that are compatible with the current File System Manager browser interface, click a server release on the Servers page or see the README file.

The servers that you add are displayed in the Servers menu. By clicking a server name on this page, you can access web pages where you can configure and administer the file systems on the server. Depending on the packages installed locally on the server that you select, different web pages, fields, and options are displayed in the browser interface.

Each time you perform a configuration action through the browser interface, the File System Manager software edits the appropriate configuration files that are located on the server that you are administering. This enables you to configure the file systems on the server remotely and without using a command-line interface.

When you perform any action in the browser interface, including clicking the Refresh button in the software, the latest data is updated. For example, mounting a file system updates the browser interface with data about the file system, its drives, and files.

Use the Current Jobs page to view the progress of user and system processes. For example, to view the archiving status of files, select Current Jobs and choose a category from the Filter menu. All jobs from the selected category are displayed.

For information about how to perform initial configuration steps, see [“Task Overview: Configuring a Server” on page 23](#).

For information about installing the management station and about installing the Sun StorageTek QFS and Sun StorageTek SAM software packages, see the *Sun StorageTek QFS Software Installation and Upgrade Guide* and the *Sun StorageTek Storage Archive Manager (SAM) Software Installation and Upgrade Guide*.

Assigning Administrative Privileges to Users

By default, the root user has privileges to perform all operations available from the File System Manager software.

You can assign other users full access to all File System Manager operations, or access to only a subset of operations. The following table lists the five levels of privileges that you can assign to File System Manager users.

TABLE 1-1 Administrative Privileges

Administrative Privilege Level	Privileges
<code>com.sun.netstorage.fsmgr.config</code>	Unlimited.
<code>com.sun.netstorage.fsmgr.operator.media</code>	Import, export, and assign VSNs (volume serial names).
<code>com.sun.netstorage.fsmgr.operator.sam.control</code>	Perform operations relating to faults and jobs; generate SAMreports; and start, stop, and idle the archiving function.
<code>com.sun.netstorage.fsmgr.operator.file</code>	Perform staging and restoring operations.
<code>com.sun.netstorage.fsmgr.operator.filesystem</code>	Mount and unmount file systems, and check and repair file systems.

These privilege levels and the functions assigned to each level are defined in the `/etc/security/auth_attr` file.

Caution: For proper system operation, do not make any changes to the `/etc/security/auth_attr` file.

To add users and assign privileges, complete the following procedure.

1. Outside of the browser interface, log in to the management station server as `root`.
2. Use the `useradd` and `passwd` commands to add each user.

For example, to add a user with account name `bobsmith`, type the following:

```
# /usr/sbin/useradd/useradd bobsmith
# /usr/bin/passwd bobsmith
```

Each user account that you add in this way has read-only viewing privileges for File System Manager functions.

3. (Optional) To specify full or partial configuration privileges for a user, add the following line to the `/etc/user_attr` file:

```
account-name:::auths=privilege-level
```

Where *account-name* is the name of the user's account and *privilege-level* is the level of authorization that you want to assign to the user.

For example, to assign full privileges (privilege level `com.sun.netstorage.fsmgr.config`) for user account `bobsmith`, add the following line to the `/etc/user_attr` file:

```
bobsmith:::auths=com.sun.netstorage.fsmgr.config
```

To assign bobsmith privileges only for staging and restoring file systems (privilege level `com.sun.netstorage.fsmgr.operator.file`) and exporting, importing, and assigning VSNs (privilege level `com.sun.netstorage.operator.media`), add the following line to the `/etc/user_attr` file:

```
bobsmith:::auths=com.sun.netstorage.fsmgr.operator.file,  
com.sun.netstorage.fsmgr.operator.media
```

You can also create an account for multiple users that includes a role with privileges that only some of those users can access. To do this:

1. Use the `useradd` and `passwd` commands to add the account.
2. Use the `roleadd` and `passwd` commands to add the role.
3. Specify the privilege levels in the `/etc/user_attr` file.

For example, to add a user account called `guest` for multiple users and a role called `admin` with special privileges within that account, type the following:

```
# /usr/sbin/useradd/useradd guest
```

```
# /usr/bin/passwd guest
```

```
# /usr/sbin/roleadd admin
```

```
# /usr/bin/passwd admin
```

To assign privileges to the `admin` role to restore, stage, archive and release files on archiving file systems, add the following lines to the `/etc/user_attr` file:

```
admin:::auths=com.sun.netstorage.fsmgr.operator.file
```

```
guest:::type=normal;roles=admin
```

In this example, when a user logs in as `guest`, File System Manager prompts the user to select either `No Role` or `Admin`. If users know the `Admin` role password, they can select `Admin`, enter the `Admin` password, and have privileges to restore and stage file systems. All other users must select `No Role` and have read-only privileges.

Because multiple users with the same privilege level can be logged in to the software concurrently, there is a risk of one user's changes overwriting another user's previous changes. To prevent this, develop policies about who can make changes and how to notify others.

Logging In to and Out of the Software

To log in to the File System Manager software:

1. Open a web browser window and type the following address in the Address field:

https://management-station-name:6789

The Sun Java™ Web Console is displayed.

2. In the User Name field, type either `root` or your user name.

Note: If you have upgraded the File System Manager software from an earlier version, the `samadmin` user account is still available. You can type `samadmin` in the User Name field and then type the `samadmin` password to gain full access to all File System Manager operations.

3. In the Password field, type the root password or your user password.

If you do not know the password, contact your system administrator.

4. Click Log In.

The Sun Java Web Console displays a list of the software applications that are available to you.

5. In the Storage section, click File System Manager.

After logging in to the File System Manager software, perform operations (such as monitoring) in a single session. Opening multiple sessions could cause the overwriting of changes made in either session.

To log out of the File System Manager software, click Log Out in the upper right corner of the page. If you are using a personal computer and do not log out using this method, the session might remain open, and you might not be prompted for a user name the next time you try to log in.

Setting the Session Timeout Value

By default, a Sun Java Web Console session times out after 15 minutes of idleness. The File System Manager installation process changes the session timeout value to 60 minutes.

You can change the session timeout value, but setting a value greater than 60 minutes may pose a security risk.

To change the timeout value:

1. Outside of the browser interface, log in to the management station server as `root`.
2. Enter the command `/opt/SUNWfsmgr/bin/fsmgr session timeout`, where *timeout* is the timeout value in minutes. For example, to set the timeout to 45 minutes, enter the following command:

```
# /opt/SUNWfsmgr/bin/fsmgr session 45
```

Registering the Software

The Registration page in the File System Manager allows you to register the software with Sun Connection services. Registration is optional.

To register the software:

1. In the File System Manager navigation tree, select Registration under System Administration.

The Sun Connection Services page is displayed.

2. In the Sun Online Account Information section, type your Sun Online Account Name and Password.
3. In the Contact Information section, type the Contact Name and Email Address for the person who will receive subscription information from Sun. This must be the same person whose online account information you have provided.
4. In the Internet Connection Settings section, choose whether you use a direct connection or a proxy server to connect to the Internet. If you use a proxy server, type in the Proxy Host Name, the Proxy Port, and, if required, the User Name and Password for proxy server authentication.
5. Click Submit to save this information and submit it to Sun.

Note: If this registration procedure fails, it will have no affect on your ability to use the software. It is an optional procedure.

Using the Browser Interface

This chapter provides information about using the browser interface. It contains the following sections:

- [“About Interface Elements” on page 9](#)
- [“About Wizards” on page 15](#)
- [“Enabling Pop-Up Windows” on page 15](#)
- [“Changing the Current Server” on page 16](#)
- [“About Report Tables” on page 16](#)
- [“Using Help” on page 19](#)

About Interface Elements

This section provides information about the elements of the browser interface. It contains the following subsections:

- [“About the Banner and Content Pane” on page 10](#)
- [“Reference: Banner Elements” on page 11](#)
- [“Reference: Content Pane Navigational Elements” on page 12](#)

- “Reference: Content Pane Icons and Buttons” on page 13

About the Banner and Content Pane

The File System Manager browser interface pages are divided into three sections:

- Banner
- Navigation tree
- Content pane

The banner displays the name of the application and other status information, such as the last time data was updated in the browser interface. The navigation tree contains the server menu and a list of all of the file system content pages. The content pane contains the currently selected page.




FIGURE 2-1 Banner, Navigational Elements, and Content Pane


You can navigate through the File System Manager browser interface as you would through a typical web page. Use the navigational tree on the left side of the content pane to move to different pages throughout the application. Click a link in the content pane to view more detailed information about that selected item. Do not use the browser’s Back button to return to a previous browser interface page. Instead, click the link that you want to return to in the navigational tree.

In addition, you can place your mouse over a portion of the browser interface to display a tool tip, which is a line of text that describes that portion of the page.

Reference: Banner Elements

The banner spans the top of the application. The following table describes the elements in the banner.

TABLE 2-1 Banner Elements

Button	Description
	Returns you to the Sun Java Web Console page.
	Displays the File System Manager software release and copyright information.
	Refreshes the current page.
	Logs you out of the Sun Java Web Console and the current application.
	Opens the online help in a separate window.
Status Information	
	Displays the name and role of the user who is currently logged in to the software.
	Displays the name of the management station system that is hosting the File System Manager software.
	Displays the latest time that data was retrieved from the file systems that you are administering. The latest data is collected and displayed each time you refresh the browser window or perform an action in the browser.
	Displays the current number of each type of unacknowledged fault on the server. There are four fault types:  Down,  Critical,  Major, and  Minor.
For more information about faults, see “About Faults” on page 56 .	

Reference: Content Pane Navigational Elements

The content pane is located below the banner of the application. It displays report information and enables you to enter data on forms. In addition, you can navigate through the pages of the application by using the navigational tree.

The following table describes the navigational elements that are displayed in the navigation tree.

TABLE 2-2 Navigational Elements

Navigational Element	Description
Current Server menu	Displays the server for which you are currently viewing information. To change the current server, select a new server from the drop-down menu.
Manage Servers link	Displays the Servers page, which provides a summary of the servers that you have added to the browser interface. Clicking a server on the Servers page takes you to the File Systems Summary page, from which you can manage the file systems on that server.
File Browsing & Recovery	Contains links to the File Browser and Recovery Points pages. The File Browser page displays a list of files and directories on the currently selected server, along with their size and current availability. The Recovery Points page shows any previously saved recovery point files for archiving file systems. Use these links to initiate archiving, staging, releasing, or restoring operations on a per-file or per-directory basis for archiving file systems.
File Systems & NFS	Contains links to the File Systems Summary and NFS pages. The File Systems Summary page provides a summary of the file systems that are located on the selected server. You can create and administer file systems from this page. The NFS page allows you to share files through Network File System (NFS). It contains a list of file systems that are currently shared.

TABLE 2-2 Navigational Elements (*Continued*)

Navigational Element	Description
Archive Administration	Contains links to the Policies, Storage Recycling, General Setup, and Archive Activity pages. The Archive Policies Summary page provides a list of the current archive policies for the selected server. From this page you can configure and manage your archive policies. The other three pages provide additional tools for monitoring and configuring archiving on the selected server.
Storage Administration	Contains links to the Tape Libraries, Tape VSNs, Disk VSNs, and VSN Pools pages. These pages help you configure and manage tape libraries, volume serial name (VSN) pools, and tape and disk VSNs used for archiving the file systems.
System Administration	Contains links to a number of general management pages including System Details, Faults, Jobs, Scheduled Tasks, Monitoring Console, Email Alerts, and System Metrics. From these pages you can configure, schedule, and monitor the various activities on all of your managed servers.




Reference: Content Pane Icons and Buttons

The following table describes icons and buttons that are commonly displayed on reports and forms in the content pane.

TABLE 2-3 Icons and Buttons for Reports and Forms

Icon/Button	Description
	Displays the top of a report or form
	Displays the part of the report or form that is indicated by the text next to this icon
	Indicates that the message is an error message
	Indicates that the message is a warning message

TABLE 2-3 Icons and Buttons for Reports and Forms *(Continued)*

Icon/Button	Description
	Indicates that the message is informational in nature
	Enables you to add the selected object to the browser interface
	Exits a procedure without applying changes
	Closes the selected window or, from the Faults tab, closes the selected faults
	Deletes the selected object from the File System Manager browser interface
	Enables you to modify the selected object
	Creates a new instance of an object
	Removes the selected object from the browser interface
	Sets all page elements to the original selections that were displayed when the page was first accessed
	Saves the selections and entries that you have made
	Applies the sort selections from the Advanced Sort window to the report table

About Wizards

A wizard is a tool that enables you to perform actions that are composed of multiple steps, such as creating a policy or file system.

Before you perform an action that opens a wizard, make sure your browser settings are correct. For more information, see [“Enabling Pop-Up Windows” on page 15](#).

Enabling Pop-Up Windows

To use all the functions in the File System Manager application, you must have pop-up windows enabled. By default, the Mozilla web browser blocks pop-up windows.

To enable pop-up windows in the Mozilla web browser:

1. Launch a web browser session.
2. Choose Edit > Preferences.

The Preferences window is displayed.

3. In the Category pane, expand the Privacy & Security node.
4. In the Privacy & Security expanded list, click Pop-up Windows.

The Pop-up Windows settings are displayed in the right pane of the window.

5. Deselect the Block Unrequested Pop-up Windows check box, and click OK.

6. Close the web browser session and open a new session.
Pop-up windows are enabled in the new browser session.

Changing the Current Server

You can administer only one server at a time through the browser interface.

To change the current server, select the server you want to administer from the Current Server menu at the top of the navigation tree. The information on the page changes to reflect the newly selected server.

About Report Tables

Certain reports and forms display report tables. This section provides information about how you control the display of tables and sort and filter the information that is displayed in report tables. It contains the following subsections:

- [“Reference: Report Table Controls” on page 17](#)
- [“Sorting Tables” on page 18](#)
- [“Filtering Tables” on page 19](#)

Reference: Report Table Controls

Report tables have the following controls that enable you to manipulate the display of the tables and to work with items listed in the tables.

TABLE 2-4 Report Table Controls

Table Control	Description
	Applies the selected operation to the objects that are selected on the current page of a summary report table or to the object displayed on a detailed report. Note: Available operations in this menu vary, depending on the objects that are selected. If you are viewing this from a summary report, and one of the selected objects in the table cannot have a particular operation applied to it, that operation is disabled in this menu.
	Indicates how many pages are in the report, and displays the page you are currently viewing. To view a different page, type the page number in the Page field and click Go.
	Enables you to select the entries that you want to display. Click the button on the left to display the next 25 report entries. Click the button on the right to display the last 25 report entries.
	Enables you to select the entries that you want to display. Click the button on the left to display the first 25 report entries. Click the button on the right to display the previous 25 report entries.
	Indicates that the column in the report table is sorted in ascending order. The ascending sort order is by number, by uppercase letter, and then by lowercase letter. If the icon has a gray background and a purple arrow, it indicates that this is the table's current ascending sort order. Click the icon on the left to redisplay the table and re-sort the column in descending order. Click the icon on the right to redisplay the table and change the table's current ascending sort order.

TABLE 2-4 Report Table Controls (*Continued*)

Table Control	Description
	Indicates that the column in the report is sorted in descending order. The descending sort order is by lowercase letter, by uppercase letter, and then by number. A column that has the gray and purple version of this icon indicates that this is the table's current descending sort order. Click the icon on the left to redisplay the table and re-sort the column in ascending order. Click the icon on the right to redisplay the table and change the table's current descending sort order.
	Enables you to select or deselect all of the check boxes on the current page of the report table. Click the icon on the left to select all of the check boxes. Click the icon on the right to deselect all of the check boxes.
	Click to deselect the selected radio button in the report table.
	Opens the Advanced Sort window. Making the table sort selections and clicking Sort re-sorts the table with your selections. The ascending sort order is by number, by uppercase letter, and then by lowercase letter. The descending sort order is by lowercase letter, by uppercase letter, and then by number.
	Enables you to toggle between displaying 1000 table rows one page at a time, and displaying 25 table rows one page at a time. When the top icon is displayed on a report table, click the icon to page through 1000 rows of data in the table. When the bottom icon is displayed on a report table, click the icon to page through 25 rows of data in the table.
	Enables you to display only the information that interests you in a report table.

Sorting Tables

To sort a report table:

1. Click the Advanced Sort icon () in the table's banner.

The Advanced Sort window is displayed.

2. Select the columns that you want to sort.
3. Select the sort order you want, and click Sort.

Alternatively, you can sort the information in individual columns in ascending or descending order by clicking either the purple or white triangle in the column heading.

Filtering Tables

You can filter the contents of certain report tables so that only the information that interests you is displayed.

When filtering tables, keep the following in mind:

- A filter must have at least one defined criterion.
- A filter applies to the current server only. You cannot apply a filter to report tables across multiple servers.
- You can use filters with the sorting function to further narrow your results.

To filter a report table, choose the filter criterion from the Filter menu in the table.

Using Help

To view additional information about the File System Manager software, click Help in the banner of the File System Manager browser interface. The help window consists of a Navigation pane on the left and a Topic pane on the right.

The following table describes the tabs in the help window.

TABLE 2-5 Help Window Tabs

Tab	Description
Contents	Click a triangle icon to display topics and subtopics. Click a topic page icon to display the help page for that topic in the Topic pane.
Index	Click the Index tab to display an alphabetized list of subjects. Click an index entry to display the help page for that topic.
Search	Click the Search tab and type the words for which you want to search in the Search field and press Enter or Return. The Navigation pane displays a list of topics that match your search criteria in order of relevancy. Click a topic link to display the help page for that topic.

The following table describes the buttons in the help window.

TABLE 2-6 Help Window Buttons

Button	Description
	Click to return to the previous help topic that you were viewing.
	Click to go forward to the next help topic that you viewed in the current session. This button is enabled only if you have clicked the Previous button.
	Click to print the help topic that is displayed in the right pane.
	Click to close the help window.

Getting Started: Task Overview

The following table provides an overview of the tasks that you can perform to get started with the software.

For a brief overview of the software, see [“How the Software Works”](#) on page 2.

TABLE 3-1 Getting Started Tasks

Task	See This Information
Perform the initial configuration.	• “Task Overview: Configuring a Server” on page 23
Monitor the environment.	• “Monitoring Servers” on page 31 • “Monitoring File Systems” on page 35 • “Monitoring Policies” on page 44 • “Monitoring Libraries” on page 46 • “Monitoring Archiving, Releasing, and Staging Activity” on page 49 • “Monitoring Jobs” on page 51 • “Monitoring Faults” on page 56
Manage file systems.	• “Administering Stand-Alone, Non-archiving File Systems and Archiving File Systems” on page 68 • “Managing Policies for Archiving File Systems” on page 76 • “Protecting Archiving File System Data” on page 80 • “Administering Shared File Systems” on page 95 • “Managing NFS Sharing for File Systems” on page 106

TABLE 3-1 Getting Started Tasks (*Continued*)

Task	See This Information
Manage libraries (archiving file systems only).	• “Adding a Library” on page 164 • “Changing the State of a Library Drive” on page 168 • “Unloading a Library” on page 166 • “Deleting a Library” on page 167 • “Managing VSNs in a Library” on page 170 • “Managing Library Drives” on page 168
Set up and administer archiving (archiving file systems only).	• “Controlling Archiving for Archiving File Systems” on page 120 • “Managing VSN Pools for Archiving” on page 142 • “Managing Disk VSNs for Archiving” on page 145 • “Setting Up General Archiving” on page 148 • “Configuring the Recycler” on page 152 • “Managing Archiving and Staging Activity” on page 156
Schedule recovery points to protect file system data.	• “Protecting Archiving File System Data” on page 80 • “Scheduling and Creating Recovery Points” on page 81
Perform general configuration of the software.	• “Creating an Email Alert” on page 181

Performing Initial Configuration Tasks

This chapter provides information about initial configuration tasks. It contains the following sections:

- [“Task Overview: Configuring a Server” on page 23](#)
- [“Reference: Server Configuration Files” on page 25](#)
- [“Adding a Server” on page 26](#)
- [“Removing a Server” on page 28](#)

Task Overview: Configuring a Server

If the Sun StorageTek SAM software (`SUNWsamfsr` and `SUNWsamfsu` packages) has already been installed locally, the local server automatically appears in the servers list when you first launch the browser interface.

If you want to add more servers, you can add and configure them using the browser interface. To host and administer archiving file systems, the servers on which the file systems reside must have the Sun StorageTek SAM software (`SUNWsamfsr` and

SUNWsamfsu packages) installed locally. Servers with only the Sun StorageTek QFS software (SUNWqfsr and SUNWqfsu packages) installed locally can host and administer only non-archiving file systems.

For information about the file systems that can be managed using the File System Manager browser interface, see [“About Supported File Systems” on page 62](#).

The following table lists the tasks you must perform to initially configure a server through the browser interface. Complete the configuration tasks in the order listed in the table.

You must have the appropriate privileges to perform these tasks.

For more information, see [“Assigning Administrative Privileges to Users” on page 3](#).

TABLE 4-1 Initial Configuration Tasks

Task	Applicable File System	See This Information
1. Add each server to the browser interface.	Archiving, non-archiving	• “Adding a Server” on page 26
2. If you are using disk archiving, create disk-based VSNs.	Archiving	• “About Archiving to Disk-Based VSNs” on page 145 • “Creating a Disk VSN” on page 146
3. If you are using tape libraries, bring the libraries under Sun StorageTek SAM software control.	Archiving	• “About Removable Media Management” on page 159 • “About Supported Libraries and Drives” on page 161 • “About Archiving” on page 114 • “Adding a Library” on page 164
Note: This task is not applicable for sites that perform disk archiving only.		
4. (Optional) Create VSN pools for each type of archive media.	Archiving	• “About VSN Pools” on page 142 • “Creating a VSN Pool” on page 143
5. If necessary, create a file system on each newly added server.	Archiving, non-archiving	• “About File System Management” on page 59 • “About File Systems” on page 61 • “Creating a Stand-Alone, Non-archiving or Archiving File System” on page 70 • “Creating a Shared File System” on page 96
6. Schedule the creation of recovery points for any newly created archiving file systems.	Archiving	• “About Recovery Points” on page 81 • “Scheduling the Creation of Recovery Points” on page 83

TABLE 4-1 Initial Configuration Tasks *(Continued)*

Task	Applicable File System	See This Information
7. (Optional) Create a custom policy and add the policy's file match criteria to the newly created archiving file systems.	Archiving	• “About Archiving, Releasing, and Staging” on page 113 • “About Archive Policies” on page 123 • “Creating a Policy” on page 127 • “Adding Existing File Match Criteria to an Archiving File System” on page 77

For information about how these tasks affect the configuration files on a server, see [“Reference: Server Configuration Files” on page 25](#).

Reference: Server Configuration Files

When you add a server to the browser interface and configure Sun StorageTek QFS file systems, the File System Manager software creates or edits the appropriate Sun StorageTek QFS configuration files on that server. If any file systems are configured to be archiving, the File System Manager software also creates or edits the appropriate Sun StorageTek SAM configuration files. You can use the command-line interface for the file systems to further edit these files later.

The following table describes the browser interface actions that affect the Sun StorageTek QFS and Sun StorageTek SAM configuration files on the server.

TABLE 4-2 Browser Interface Actions and their Effect on Files and Commands

Browser Interface Action	Affected Configuration File or Files	Associated Man Commands
Create, edit, grow, or remove a file system.	• <code>mcf</code> • <code>samfs.cmd</code> • <code>/etc/vfstab</code>	• <code>mcf(4)</code> • <code>samfs.cmd(4)</code> • <code>vfstab(4)</code>
If you are using tape libraries, add the library and bring the library under Sun StorageTek SAM software control or remove the library from Sun StorageTek SAM software control.	• <code>mcf</code> • <code>samst.conf</code> Note: If the library is a small computer system interface (SCSI) direct-attached library and its target number is larger than 6 or its logical unit number (LUN) identifier is larger than 1, the software does not update the <code>samst.conf</code> file. In this situation, you must manually configure the <code>samst.conf</code> file.	• <code>mcf(4)</code> • <code>samst.conf(4)</code>
Create, edit, or delete archive policies, and apply a policy to an archiving file system.	• <code>archiver.cmd</code> • <code>diskvols.conf</code>	• <code>archiver.cmd(4)</code> • <code>diskvols.conf(4)</code>
Make configuration changes on the General Setup page under the Archive Management tab.	• <code>stager.cmd</code> • <code>recycler.cmd</code> • <code>releaser.cmd</code>	• <code>stager.cmd(4)</code> • <code>recycler.cmd(4)</code> • <code>releaser.cmd(4)</code>
Schedule or create a recovery point and restore files from the recovery point.	<code>/var/spool/cron/crontabs/root</code>	<code>crontab(1)</code>

Adding a Server

Adding a server enables you to administer its file systems, archiving processes, media, and so on, through the browser interface.

To add a server:

1. Outside of the browser interface, use `telnet` to connect to the server.
2. Log in to the server as `root`.
3. Use the `fsmadm(1M) add` command to add the management station (the system on which the File System Manager software is installed) to the list of hosts that can remotely administer the server.

fsmadm add *management-station-name:domain-name*

Only hosts that are added to the list through this command can remotely administer the server.

4. To ensure that the management station has been successfully added, use the `fsmadm(1M) list` command.
5. Launch a browser window and log in to the File System Manager browser interface as an administrator user.
6. Perform either of the following:
 - If the Servers page is displayed, click Add.
 - If any other page is displayed, click the Manage Servers link at the top of the navigation tree to go to the Servers page and click Add.The Add Server window is displayed.
7. Type the name of the server or type the IP address of the server in the Server Name field or IP Address field.
8. Click Add Server.
9. If you typed the name or address of a server that is part of a SunTM Cluster environment, a list of the other nodes in the cluster is displayed. Select the other nodes to also add them as managed servers for an HA (highly available)-QFS configuration.

Note: To create an HA-QFS configuration, you must also use the Sun Cluster management tool, SunPlexTM Manager, to configure the devices as part of a Sun Cluster resource group. File System Manager includes *ha* in the description of file systems that are configured on cluster devices even if the devices are not part of a resource group and the file systems are not therefore truly highly available.

When you add a server, the software detects any file systems that are already configured on the server. These file systems are automatically displayed on the File Systems Summary page.

If no file systems have already been created on the server, you can create and configure a new file system.

For more information, see [“Creating a Stand-Alone, Non-archiving or Archiving File System” on page 70.](#)

Removing a Server

If you no longer want to administer the file systems on a server through the browser interface, you can manually remove the server.

Note: Removing a server does not uninstall any software from the server. It only causes the server to no longer be displayed in the browser interface.

To remove a server from the browser interface:

1. Perform either of the following:

- If the Servers page is displayed, select the radio button next to the server that you want to remove and click Remove.
- If any other page is displayed, click the Manage Servers link at the top of the navigation tree to go to the Servers page, select the radio button next to the server that you want to remove, and click Remove.

A message box prompts you to confirm the removal.

2. Click OK.

3. If you want to permanently disable remote administration of the server from the management station, do the following:

- a. Outside of the browser interface, use `telnet` to connect to the server.
- b. Log in to the server as `root`.
- c. Use the `fsmadm(1M) remove` command to remove the management station from the list of hosts that can remotely manage the server.

fsmadm remove *management-station-name.domain-name*

- d. To verify that the management station has been removed, use the `fsmadm(1M) list` command.

Monitoring the Environment

This chapter provides information about monitoring your environment. It contains the following sections:

- [“Monitoring Servers” on page 31](#)
- [“Monitoring File Systems” on page 35](#)
- [“Monitoring Policies” on page 44](#)
- [“Monitoring Libraries” on page 46](#)
- [“Monitoring Archiving, Releasing, and Staging Activity” on page 49](#)
- [“Monitoring Jobs” on page 51](#)
- [“Monitoring Faults” on page 56](#)

Monitoring Servers

This section provides information about monitoring servers. It contains the following subsections:

- [“Displaying the Configuration of Servers in Your Site” on page 32](#)
- [“Displaying the Configuration of a Server” on page 32](#)

- [“Displaying Software Packages on a Server” on page 33](#)
 - [“Displaying Log or Trace Files on a Server” on page 33](#)
 - [“Displaying Configuration Files on a Server” on page 34](#)
 - [“Displaying Warning Messages in Server Files” on page 34](#)
 - [“Displaying the Minimum Server Release Supported by File System Manager” on page 35](#)
-

Displaying the Configuration of Servers in Your Site

You can view configuration information about all servers in your site.

If at least one server is hosting archiving file systems, you can view the media devices that can be used for archiving.

In addition to viewing server configuration information, you can also use the Monitoring Console page to see an overview of server activities.

To display configuration information about the servers in your site:

- From the Servers page, click the Site Information tab. The Site Information page lists the servers and their media devices.
-

Displaying the Configuration of a Server

You can view configuration information about an individual server, such as its log, trace files, or packages.

To display configuration information about a server:

1. From the Servers menu or the Servers page, select the server that contains the configuration information you want to view.
2. From the navigation tree, click System Administration.

The Server Configuration page is displayed.

Displaying Software Packages on a Server

To display the software packages that are installed locally on a server:

1. From the Servers menu or the Servers page, select the server that contains the software packages you want to view.

2. From the navigation tree, click System Administration.

The Server Configuration page is displayed.

3. Scroll down to the Packages section.

The table lists any File System Manager, Sun StorageTek QFS, or Sun StorageTek SAM software packages that are installed on the server.

Displaying Log or Trace Files on a Server

You can view the contents of the Sun StorageTek QFS and Sun StorageTek SAM log and trace files on the current server. This is useful when a fault condition occurs and you want to investigate its cause.

To view the contents of a log or trace file:

1. From the Servers menu or the Servers page, select the server that contains the log or trace files you want to view.

2. From the navigation tree, click System Administration.

The Server Configuration page is displayed.

3. Scroll down to the Log and Trace section and click the path name of the log or trace file that you want to view.

A window is displayed that enables you to view the contents of the selected file. This window periodically refreshes so that the latest information in the file is displayed.

Displaying Configuration Files on a Server

You can view the contents of the Sun StorageTek QFS and Sun StorageTek SAM configuration files on the current server.

To view the contents of a configuration file:

1. From the Servers menu or the Servers page, select the server that contains the configuration files you want to view.
2. From the navigation tree, select System Administration.

The Server Configuration page is displayed.

3. Scroll down to the Configuration File Status section and click the path name of the configuration file that you want to view.

A window is displayed that enables you to view the contents of the selected file. This window periodically refreshes so that the latest information in the file is displayed.

Displaying Warning Messages in Server Files

You can display warning messages for a server.

If there are warning messages related to a Sun StorageTek QFS or Sun StorageTek SAM configuration file on a server, the configuration file is displayed as a link in the browser interface.

To display the warning messages related to a server configuration file:

1. From the Servers menu or the Servers page, select the server that contains the configuration file warning messages you want to view.
2. From the navigation tree, click System Administration.

The Server Configuration page is displayed.

3. Scroll down to the Configuration File Status section and click the configuration status of the file that you want.

A window is displayed that lists the warning messages related to the configuration file.

Displaying the Minimum Server Release Supported by File System Manager

The features and functionality supported through the browser interface depend on the release level of the Sun StorageTek QFS and Sun StorageTek SAM software installed locally on a server.

To display the minimum server release supported by the browser interface:

- From the Servers page, click the release and update level in the Release column of the table. The File System Manager Release Highlights page is displayed.

This page lists the features and functionality available through the browser interface.

Monitoring File Systems

This section provides information about monitoring file systems. It contains the following subsections:

- [“Displaying the Status of a File System” on page 36](#)
- [“Displaying File System Devices” on page 36](#)
- [“Displaying NFS Shared Directories in a File System” on page 37](#)
- [“Displaying File Match Criteria for an Archiving File System” on page 37](#)
- [“Displaying the Status of Recovery Point Activities” on page 38](#)
- [“Displaying the Status of a Restore Process for an Archiving File System” on page 38](#)
- [“Displaying the Metadata Server for a Shared File System” on page 39](#)

- [“Displaying All Hosts in a Shared File System” on page 39](#)
- [“About System Metrics” on page 40](#)

Displaying the Status of a File System

To display the status of a file system:

1. From the navigation tree, click File Systems & NFS.
The File Systems Summary page is displayed.
2. To display a specific file system, page through the table until the specified file system is displayed, or filter the table so that the file system that meets the filter criteria is displayed.
3. View the information in the columns to determine whether the file system is mounted, whether it is shared, how much space is consumed, and so on.

In addition to the information shown here, you can also view file system status information using the Monitoring Console.

Displaying File System Devices

To display the devices on a file system:

1. On the File Systems Summary page, click the name of the file system that contains the devices you want to view.
The File System Details page is displayed.
2. If you are displaying the device of a UNIX File System (UFS) file system, review the Device Name field. Otherwise, click View Devices.
The Devices page for that file system is displayed.

Displaying NFS Shared Directories in a File System

You can display Network File System (NFS) directories that are shared out to the network from a local file system.

To display the NFS shared directories in a file system:

- From the navigation tree, click NFS under File Systems & NFS.

The NFS Shared Resources page is displayed.

Displaying File Match Criteria for an Archiving File System

File match criteria are defined in a policy and can be applied to an archiving file system. When the files in the archiving file system meet the file match criteria defined in the policy, the files are eligible to be archived.

To display the sets of file match criteria that are applied to an archiving file system:

1. From the Servers menu or the Servers page, select the name of the server on which the archiving file system is located.
2. From the navigation tree, choose Policies under Archive Administration.
The Archive Policies Summary page displays the policy and file match criteria combinations that apply to this archiving file system.
3. To view details about a set of file match criteria:
 - a. Click the name of the policy.
The Policy Details page for the selected policy is displayed.
 - b. Click the file match criteria that you want to view.
The Policy Criteria Details page is displayed.

Displaying the Status of Recovery Point Activities

To display the status of activities related to recovery points on a server:

1. From the Servers menu or the Servers page, select the name of the server for which you want to display the status of a recovery point.
2. From the navigation tree, choose Jobs under System Administration.
The Current Jobs Summary page is displayed.
3. From the Filter menu, choose one of the following:
 - Create Recovery Point – When a recovery point is created on demand, this job enables you to view the status of the recovery point being created.
 - Create Recovery Point Index – When you want to restore a file from archive media, you might need to create a search index for the recovery point so you can browse through it. This job enables you to view the status of this process.
4. (Optional) To display more detailed information about a recovery point job, click its job ID.

Displaying the Status of a Restore Process for an Archiving File System

In the event of a disaster, you can restore lost files from archive media to online disk cache of an archiving file system.

To display the status of a restore process for an archiving file system:

1. From the Servers menu or the Servers page, choose the name of the server for which you want to display the status of a file restore process.
2. From the navigation tree, choose Jobs under System Administration.
The Current Jobs Summary page is displayed.
3. Choose whether you want to display current, pending, or all jobs by choosing the appropriate item under Jobs in the navigation tree.
4. From the Filter menu, choose Restore.

Displaying the Metadata Server for a Shared File System

To display the metadata server for a shared file system:

1. From the Servers menu or the Servers page, choose a server on which the shared file system is configured.
2. On the File Systems Summary page, click the shared file system that contains the metadata server you want to display.

The File System Details page is displayed.

3. Do one of the following:
 - If the current server is acting as a client for the shared file system, scroll down to the Server field to view the metadata server.
 - If the current server is acting as a potential metadata server or as the metadata server, the Participating Hosts table is displayed. Look at the values in the Type column to determine which server is acting as the metadata server.

Displaying All Hosts in a Shared File System

To display all participating hosts in a shared file system:

1. From the Servers menu or the Servers page, choose the name of the server that is acting as the metadata server or as the potential metadata server for the shared file system.
2. From the File Systems Summary page, click the name of the shared file system that contains the hosts you want to display.

The File System Details page is displayed.

3. Look at the Participating Hosts table to view all participating hosts in the shared file system.

About System Metrics

You can use the System Metrics pages to run and view reports on file system utilization, data distribution, and media status. The reports can be saved, allowing you to view this data over time. System Metrics are only available if **version 4, update 6** of the software is installed on the server.

The following sections describe the System Metrics pages:

- [“Media Utilization and Status Page” on page 40](#)
- [“File Data Distribution Page” on page 41](#)
- [“File System Utilization Page” on page 42](#)
- [“Creating Media Utilization and Status Reports” on page 42](#)
- [“Viewing File Data Distribution” on page 43](#)
- [“Creating File System Utilization Reports” on page 43](#)

Media Utilization and Status Page

Use this page to create media utilization and status reports, or to view reports that have previously been created. This report is not available if Sun StorageTek QFS and Sun StorageTek SAM version 4.5 software is installed on the server.

Click New to create a new report. To view a report, click on the date and time in the Time Generated column. To delete a report that is no longer needed, select the radio button next to the report and click Remove.

The following information can be provided in the media utilization and status report:

- **VSNs Requiring Operator Attention:** This table lists any VSNs that had issues requiring attention at the time the report was created. When you create the report, you can select the VSN errors that you would like to see listed.
- **Archive Media Utilization:** This table lists all of the archive media for an archiving file system and gives the total space, the space consumed, and the percentage consumed.
- **VSN is blank:** This table lists all VSNs that were blank at the time the report was created.
- **VSN is in use in the last 24 hours:** This table lists all VSNs that were used in the 24 hours before the report was created.
- **Archiver Pool Levels:** This table lists any VSN pools and gives the VSN count and the free space for each one.

- **VSN Summary:** This table lists all of the media types in the files system and gives the VSN count, the total space, and the free space for each type.

If one or more Sun StorageTek Automated Cartridge System Library Software (STK ACSLS) libraries are network-attached or direct attached to the managed server, you can also generate an ACSLS report. The following information is provided in the ACSLS report:

- **Drives:** The table lists the drives that are attached to the storage module, providing information regarding their physical location in the library and its status.
- **Locks:** This table lists the lock status of the drives.
- **Scratch Pools:** This table lists the pool of scratch VSNs. Scratch is a VSN attribute that indicates it is blank or contains no useful data.
- **Accessed or Imported Volumes:** This table lists the VSNs that were accessed or imported in the last 24 hours.

Note: The terminology used in these reports is the same as that used in the Sun StorageTek Automated Cartridge System Library Software.

File Data Distribution Page

Use this page to graphically view data distribution within the file system. This data can help you gauge whether you are making the best use of your primary file system disk space, and whether your archiving policies are working efficiently. Before you can view data you must either have indexed recovery points or set up scheduled data collection from the selected file system. The more data points that are collected over time, the more useful these graphs become for detecting file data distribution trends.

You can move your mouse over points on the graph to view additional information displayed in pop-up tooltips, such as Number of files, Files Online, or Files Offline.

To schedule data collection, click [Schedule Collection from File System](#).

When you have data available to view, select the file system name, the start date, the end date, and then click [Go](#) to see the data in the graph.

The following data views are available:

- **Storage Tiers** – This option displays a time series graph showing the total amount of storage used in each tier of storage by the file system.
- **Files by Access** – This metric provides time series line plots for six different ranges of the time since the file was accessed. Each line shows the cumulative size of all files from the range. The graph shows how much of the file system's data is

being accessed on a regular basis. This information can be used to feed back into decisions about the relative sizing of the storage tiers. Roll your mouse over different points in the graphs to see more data.

- Files by Age – This metric provides time series line graphs of the cumulative size of files in five age ranges. Roll your mouse over different points in the graphs to see more data.
- Files by Owner – This option displays a bar graph showing the percentage of files owned by the top ten users.
- Files by Group– This option displays a bar graph showing the percentage of files owned by the top ten groups.

File System Utilization Page

Use this page to create and view file system utilization reports. This report is not available if Sun StorageTek QFS and Sun StorageTek SAM version 4.5 software is installed on the server.

Click New to create a new report. To view a report, click on the date and time in the Time Generated column. To delete a report that is no longer needed, select the radio button next to the report and click Remove.

The following information is displayed in the report:

- File system name
- Mounted – yes or no
- Total space
- Free space
- High water mark
- Low water mark
- How many times the Releaser has run in the last 24 hours

Creating Media Utilization and Status Reports

The Media Utilization report provides an overview of archiving storage capacity, a summary of unusable media, and information about VSN usage.

To create and view a report, follow these instructions:

1. From the Servers menu or the Servers page, select the server for which you want to create the report.
2. From the navigation tree, choose System Metrics under System Administration.

The Media Utilization & Status page is displayed.

3. Click New to create a Media Utilization report.

The New Metric page is displayed.

4. Select the volume serial name (VSN) failures that you want to have reported.
5. Click Submit to run the report.
6. To view the report, click the report date in the Media Utilization Historic Summary table.

Viewing File Data Distribution

1. From the Servers menu or the Servers page, select the server that contains the data you want to view.
2. From the navigation tree, choose System Metrics under System Administration, and then choose File Data Distribution.

The File Data Distribution page is displayed.

3. To schedule the collection of data for the File Data Distribution report, click Schedule Collection from File System.

The Schedule Data Collection page is displayed.

4. Select the time and interval for the data collection process, and then click Submit.
5. After data has been collected, select the file system name and a date range, and then click Go to view the data. Click on the links to the left of the graph to see different data views.

Creating File System Utilization Reports

1. From the Servers menu or the Servers page, select the server that contains the data you want to view.
2. From the navigation tree, choose System Metrics under System Administration, and then choose File System Utilization.

The File System Utilization page is displayed.

3. Click New to create a File System Utilization report.

The new report is displayed in the File System Utilization Historic Summary table.

4. To view the report, click the report date in the File System Utilization Historic Summary table.

Monitoring Policies

This section provides information about monitoring policies. It contains the following subsections:

- [“Displaying File Match Criteria in a Policy” on page 44](#)
- [“Displaying File Systems Using Specific File Match Criteria” on page 45](#)
- [“Displaying the Copy Details of a Policy” on page 45](#)
- [“Displaying the VSNs Assigned to a Policy Copy” on page 46](#)

Displaying File Match Criteria in a Policy

You can display the details of a specific set of file match criteria that is defined in a policy. When files meet the file match criteria specified in the policy, the files are eligible for archiving.

To display file match criteria in a policy:

1. From the Servers menu or the Servers page, choose the name of the server that is configured with the policy that you want to view.
2. From the navigation tree, choose Policies under Archive Administration.
The Archive Policies Summary page is displayed.
3. Click the name of the policy that contains the file match criteria you want to view.
The Policy Details page lists the file match criteria that is defined in the policy.

4. If you want to view the details of a specific set of file match criteria, click the name of the set.

The Policy Criteria Details page is displayed.

Displaying File Systems Using Specific File Match Criteria

File match criteria can be applied to numerous archiving file systems.

To display the archiving file systems to which a selected set of file match criteria applies:

1. From the Servers menu or the Servers page, choose the name of the server on which the file systems are located.
2. From the navigation tree, choose Policies under Archive Administration.
The Archive Policies Summary page is displayed.
3. Click the name of the policy that you want to view.
The Policy Details page for the selected policy is displayed.
4. Click the file match criteria that you want to view.
The Policy Criteria Details page is displayed.
5. Scroll down the page to the File Systems Using Criteria section.
A list of file systems that use the selected file match criteria is displayed.

Displaying the Copy Details of a Policy

To display the copy details of a policy:

1. From the Servers menu or the Servers page, choose the name of the server that is configured with the policy that you want to view.
2. From the navigation tree, choose Policies under Archive Administration.
The Archive Policies Summary page is displayed.
3. Click the name of the policy that you want to view.

The Policy Details page for the selected policy is displayed.

4. Scroll down to the Copy Information section and click a copy number.

The Policy Copy Details page is displayed.

Displaying the VSNs Assigned to a Policy Copy

A policy copy defines the volumes or pool of volumes to which an archive copy of a file can be sent. A volume serial name (VSN) identifies a volume.

To display the VSNs and VSN pools of a policy copy:

1. From the Servers menu or the Servers page, choose the name of the server that is configured with the policy that you want to view.

2. From the navigation tree, choose Policies under Archive Administration.

The Archive Policies Summary page is displayed.

3. In the table, click the name of the policy that contains the VSN information you want to display.

The Policy Details page for the selected policy is displayed.

4. Scroll down the page to the Copy Information section and click the VSNs listed in the VSNs Available field for a specified copy.

The Policy Copy VSNs page lists the VSNs in the Available VSNs field.

Monitoring Libraries

This section provides information about monitoring libraries. It contains the following subsections:

- [“Displaying VSNs for a Library” on page 47](#)
 - [“Displaying the Drives in a Library” on page 47](#)
 - [“Displaying the Status and State of a Library” on page 48](#)
 - [“Displaying Historian Catalog Information” on page 48](#)
 - [“Displaying Mount Requests in a Library” on page 49](#)
-

Displaying VSNs for a Library

Volume serial names (VSNs) are the named volumes associated with a selected library.

To display VSN information for a library:

1. From the Servers menu or the Servers page, choose the name of the server on which the file system that you want is located.
 2. From the navigation tree, choose Tape VSNs under Storage Administration.
The VSN Summary page is displayed.
 3. If you have more than one library, select the library that you want to view from the drop-down list in the upper right part of the window.
The VSNs for that library are displayed.
 4. For information about a VSN, click the name of the VSN.
The VSN Details page for the selected VSN is displayed.
-

Displaying the Drives in a Library

To display the drives in a library:

1. From the Servers menu or the Servers page, choose the name of the server that you want to monitor.
2. From the navigation tree, choose Tape Libraries under Storage Administration.
The Library Summary page is displayed.
3. Select the radio button next to the library that is configured with the drives that you want to view, and click View Drives.

The Library of Drives Detail page for the selected library is displayed.

Displaying the Status and State of a Library

You can determine the status and state of libraries through the browser interface. The status indicates whether any fault alarms are present. The state indicates whether the library is on, off, or unavailable.

To display the status and state of a library:

1. From the Servers menu or the Servers page, choose the name of the server that you want to monitor.
2. From the navigation tree, choose Tape Libraries under Storage Administration.
The Library Summary page is displayed.
3. View the State column to see the state of the library.
4. View the Faults column to see the status of the library.

If the library is in a fault condition, click the corresponding fault icon to display the Fault Summary page for that library.

Displaying Historian Catalog Information

The Historian is a catalog that keeps track of volumes that have been exported from an automated library or that have been unloaded from manually loaded devices.

To display the Historian catalog:

1. From the Servers menu or the Servers page, choose the name of the server that you want to monitor.
2. From the navigation tree, choose Tape Libraries under Storage Administration.
The Library Summary page is displayed.
3. Click Historian.

The Historian page is displayed.

Displaying Mount Requests in a Library

To display mount requests:

1. From the Servers menu or the Servers page, choose the name of the server that you want to monitor.
2. From the navigation tree, choose Jobs under System Administration.
The Current Jobs Summary page is displayed.
3. Choose whether you want to display current, pending, or all jobs requests by clicking the appropriate item under Jobs in the navigation tree.
4. To view all media mount request jobs, choose Media Mount Request from the Filter menu.
5. (Optional) To display more detailed information about a media mount request job, click its job ID.

Monitoring Archiving, Releasing, and Staging Activity

If you are managing a server that is hosting the Sun StorageTek SAM software, you can monitor the current archiving, staging, and releasing activity for any archiving file systems on the server. This section provides information on completing these tasks. It contains the following subsections:

- [“Displaying Archiving Activity” on page 50](#)
- [“Displaying Releasing Activity” on page 50](#)
- [“Displaying Staging Activity” on page 51](#)

Displaying Archiving Activity

You can view the archiving activity of archiving file systems on the current server.

To display archiving activity:

1. From the Servers menu or the Servers page, choose the name of the server that contains the archiving activity you want to display.
2. From the navigation tree, choose Jobs under System Administration.
The Current Jobs Summary page is displayed.
3. Choose whether you want to display current, pending, or all jobs by choosing the appropriate item under Jobs in the navigation tree.
4. From the Filter menu, choose Archive Copy or Archive Scan to view all jobs of either type.
5. (Optional) To display more detailed information about an archive job, click its job ID.

Note: Archive Scan jobs do not show data for file systems that use continuous archiving.

Displaying Releasing Activity

You can view the releasing activity of archiving file systems on the current server.

To display releasing activity:

1. From the Servers menu or the Servers page, choose the name of the server that has the releasing activity you want to view.
2. From the navigation tree, choose Jobs under System Administration.
The Current Jobs Summary page is displayed.
3. Choose whether you want to display current, pending, or all jobs by choosing the appropriate item under Jobs in the navigation tree.
4. To view all releasing jobs, choose Release from the Filter menu.
5. (Optional) To display more detailed information on a release job, click its job ID.

Displaying Staging Activity

You can view the staging activity of archiving file systems on the current server.

To display staging activity:

1. From the Servers menu or the Servers page, choose the name of the server that has the staging activity you want to view.
2. From the navigation tree, choose Jobs under System Administration.
The Current Jobs Summary page is displayed.
3. Choose whether you want to display current, pending, or all jobs by choosing the appropriate item under Jobs in the navigation tree.
4. To view all staging jobs, from the Filter menu, choose Stage.
5. (Optional) To display more detailed information about a stage job, click a job ID.

Monitoring Jobs

This section provides information about monitoring jobs. It contains the following subsections:

- [“About Jobs” on page 52](#)
- [“Displaying Jobs” on page 53](#)
- [“Canceling a Job” on page 54](#)
- [“Using the Monitoring Console” on page 55](#)

About Jobs

Jobs are generated when you perform an action through the browser interface or when the server system performs a process. Some jobs are generated and completed so quickly that you might not ever see the job in the Jobs table.

All jobs in the browser interface, with the exception of the File System Check job, are generated only if the current server has the `SUNWsamfsr` and `SUNWsamfsu` packages installed locally.

The following table describes the different types of jobs.

TABLE 5-1 Job Types

Job	Description
Tape Label	A volume is being labeled or relabeled.
Archive Scan	The archiver is scanning directories for files to be archived.
Release	The releaser is freeing up disk cache after an archive copy is made.
Archive Copy	The archiver is archiving the file.
File System Check	The file system is being checked and repaired.
Stage	The stager is copying file data to the disk cache.
Create Recovery Point	The software is creating a recovery point for the metadata of an archiving file system.
Create Recovery Point Index	The software is performing at least one of the following actions, as necessary: • Moving the recovery point file from removable media to online disk • Decompressing the recovery point file • Creating a search index for the recovery point file

TABLE 5-1 Job Types (*Continued*)

Job	Description
Recovery Point Search	The software is searching the recovery point file for files to restore.
Restore	The software is reviewing the pointers to the archive locations of files or directories on archive media. In addition, the files or directories might optionally be restored to online disk, depending on how you initiated the restore job.
Media Mount Request	You have manually requested that media be mounted.

The categories by which you can display jobs are as follows:

- Current – Jobs that are running
- Pending – Jobs that are scheduled to run
- All – Both current and pending jobs

Displaying Jobs

To display jobs:

1. From the Servers menu or the Servers page, choose the name of the server that has the jobs you want to view.
2. From the navigation tree, choose Jobs under System Administration.
The Current Jobs Summary page lists all currently running jobs.
3. Choose whether you want to display current, pending, or all jobs by choosing the appropriate item under Jobs in the navigation tree.
4. (Optional) To display detailed information about a specific job, click its job ID.
The Details page for that job is displayed.

Canceling a Job

You can cancel certain jobs, depending on whether they are current or pending. Some jobs can only be in the current or pending state, but not both. Unavailable states are indicated by the NA (not applicable) values in the following table.

TABLE 5-2 Job Cancellation Options

Job Type	Can Cancel In Current State?	Can Cancel In Pending State?
Archive Copy	Yes	No
Archive Scan	Yes	No
Stage	Yes	Yes
Release	No	NA
Media Mount Request	NA	Yes
Create Recovery Point	Yes	NA
Create Recovery Point Index	Yes	NA
Recovery Point Search	Yes	NA
Restore	Yes	NA
File System Check	Yes	NA
Note: The File System Check job cannot be canceled if the Repair option is specified when the job is initiated.		
Tape Label	No	NA

To cancel a job:

1. From the Servers menu or the Servers page, choose the name of the server for which you want to cancel a job.
2. From the navigation tree, choose Jobs under System Administration.
The Current Jobs Summary page is displayed.
3. Use the Filter menu or navigate through the table to find the job that you want to cancel.
4. Select the radio button next to the job and click Cancel.
A window is displayed that confirms your cancellation.
5. Click OK.

Using the Monitoring Console

As an alternative to the Current Jobs Summary page, you can use the Monitoring Console to monitor system activities, if the server that has either of the following installed locally:

- Release 4, update 6, SUNWsamfsr and SUNWsamfsu packages
- Release 4, update 6, SUNWqfsr and SUNWqfsu packages

The following items are available in the Monitoring Console:

- Application Daemons
- File Systems
- Archive Media Utilization
- Tape Libraries
- Library Drives
- VSN Pending Load Request
- Unusable VSNs
- Archiving Copy Queue
- Staging Queue

To monitor system activity using the Monitoring Console:

1. From the Servers menu or the Servers page, choose the name of the server you want to monitor.
2. From the navigation tree, choose Monitoring Console under System Administration.

The Monitoring Console window is displayed.

3. Click one of the selections in the left navigation pane to see the current data.

Click the Refresh button in the content pane to manually refresh the information. A time stamp in each table title shows the last time the data was retrieved. You can also use the auto-refresh menu in the navigation pane to choose how often the Monitoring Console data is refreshed.

4. When you are finished viewing the data, click Close to close the Monitoring Console window.

Monitoring Faults

This section provides information about monitoring faults. It contains the following subsections:

- [“About Faults” on page 56](#)
- [“Displaying Faults” on page 57](#)
- [“Acknowledging Faults” on page 58](#)
- [“Deleting Faults” on page 58](#)

About Faults

Faults are displayed in the browser interface when situations occur that meet the fault criteria defined by the Sun StorageTek QFS and Sun StorageTek SAM software installed locally on the current server.

If there are more than 25 faults on the server, you can acknowledge and delete all faults on the system at one time by clicking the Show Data in a Single Page icon on the Fault Summary page. After you click this icon, you can select all faults and then click Acknowledge or Delete.

If a fault condition is discovered more than once by the software, the number of times it has been discovered and the initial time it was discovered are displayed in the browser interface:

```
No space available. Repeated 3 times since Mon Apr 18, 2005
2:09 p.m.
```

The following table describes the fault severities.

TABLE 5-3 Fault Severities

Severity Level	Description
Down 	Indicates that the File System Manager software cannot reach the server or object. This might be due to a faulty network connection, a power outage, and so on. Down faults are only displayed on the Servers page. For more information about possible causes for this fault condition, see “Inaccessible Server” on page 189 .
Critical 	Indicates a critical fault condition, such as a superblock count mismatch or a library being unavailable. This is a very serious condition and requires some administrative action.
Major 	Indicates a major fault condition, such as a library being full. This is a serious condition and requires some administrative action.
Minor 	Indicates a minor fault condition, such as a tape drive needing to be cleaned, or a library needing to be maintained. This is a serious condition, but you can choose whether to perform some administrative action.

Displaying Faults

To display faults:

1. From the Servers menu or the Servers page, choose the name of the server that has the faults you want to view.
2. From the navigation tree, choose Faults under System Administration.
The Faults Summary page lists the faults on that server.
3. (Optional) To display more detailed information about a device that is in a fault condition, click the device name.
The Fault Summary page for that device lists the faults for that device.

Acknowledging Faults

You can optionally acknowledge a fault to inform other administrators that you are addressing the fault condition.

To acknowledge a fault:

1. From the Servers menu or the Servers page, choose the name of the server that has the fault you want to acknowledge.
2. Click on one of the fault icons in the upper-right corner of the browser window. The Faults Summary page is displayed.
3. Select the check box next to the fault that you want to acknowledge, and click Acknowledge.

Acknowledged is displayed in the State field for the fault.

Deleting Faults

Faults are automatically deleted upon reaching a system-defined threshold. However, you can also manually delete faults.

To delete a fault:

1. From the Servers menu or the Servers page, choose the name of the server that has the fault you want to delete.
2. From the navigation tree, choose Faults under System Administration. The Fault Summary page is displayed.
3. Select the check box next to the fault that you want to delete, and click Delete. A message box prompts you to confirm your choice.
4. Click OK.

Managing File Systems

This chapter provides information about managing file systems. It contains the following sections:

- [“About File System Management” on page 59](#)
- [“About Discovering Available Devices” on page 60](#)
- [“About File Systems” on page 61](#)
- [“Administering Stand-Alone, Non-archiving File Systems and Archiving File Systems” on page 68](#)
- [“Managing Policies for Archiving File Systems” on page 76](#)
- [“Protecting Archiving File System Data” on page 80](#)
- [“Administering Shared File Systems” on page 95](#)
- [“Managing NFS Sharing for File Systems” on page 106](#)

About File System Management

You can perform the following file system-level operations through the browser interface:

- Create Sun StorageTek QFS file systems and configure them to be stand-alone, archiving, or shared

- Create non-archiving HA (highly available) shared or stand-alone Sun StorageTek QFS file systems on nodes within a Sun Cluster configuration

Note: To create an HA-QFS configuration, you must also use the Sun Cluster management tool, SunPlex Manager, to configure the devices as part of a Sun Cluster resource group. File System Manager includes *ha* in the description of file systems that are configured on cluster devices even if the devices are not part of a resource group and the file systems are not therefore truly highly available.

- Create Sun StorageTek SAM file systems, which are automatically configured to be archiving
- Create stand-alone UFS file systems
- Add clients and potential metadata servers to and remove them from shared file systems
- Grow Sun StorageTek SAM file systems and grow stand-alone and archiving Sun StorageTek QFS file systems
- Check and repair Sun StorageTek SAM file systems and check and repair stand-alone and archiving Sun StorageTek QFS file systems
- Mount and unmount stand-alone, shared, and archiving file systems
- Mount and unmount VERITAS File Systems (VxFS)
- Control archiving for archiving file systems
- Manage archive policies for archiving file systems
- Protect the data of mounted archiving file systems

If file systems already exist on the current server, they are automatically displayed on the File Systems Summary page.

About Discovering Available Devices

When you create or grow a file system, the File System Manager software attempts to discover available devices that can be used in the process.

The software discovers available devices on formatted disks that are properly labeled. SMI and EFI labels are supported. A device is considered unavailable if it meets any of the following restrictions:

- Its size is equal to zero.
- It is listed in the `/etc/vfstab` file.
- It is listed in the selected server's `mcf` file.
- It starts at sector 0. Because of this, `s2` is excluded for disks that use the default Solaris™ Operating System formatting scheme.
- It cannot be opened in exclusive-read mode, which eliminates slices or volumes that are used by mounted file systems.
- It is used by a Solaris Volume Manager volume or by a `metadb` slice.
- It is used by a VERITAS Volume Manager volume.

Caution: A device is also unavailable if another host is using the device. However, the browser interface lists such devices as available. You are responsible for knowing which shared devices are being used by other hosts.

There might be additional circumstances under which the devices listed as available in the browser interface are actually unavailable. For example, the browser interface may list devices as available because the software expects disks to be partitioned without overlapping slices, or because the software was unable to invoke volume manager commands when performing discovery.

If you need to format a disk, use the `format(1M)` command outside of the browser interface. If you use the `format(1M)` command, you are responsible for determining whether devices on the disk are in use.

About File Systems

This sections provides information about file systems. It contains the following subsections:

- [“About Supported File Systems” on page 62](#)
- [“About Stand-Alone, Non-archiving File Systems” on page 63](#)
- [“About Archiving File Systems” on page 65](#)

- [“About Shared File Systems” on page 66](#)
- [“About HA-QFS File Systems” on page 67](#)

About Supported File Systems

File systems can be stand-alone, non-archiving file systems; archiving file systems; or shared file systems. Through the browser interface you can configure Sun StorageTek SAM and Sun StorageTek QFS file systems that are optionally archiving or shared.

Sun StorageTek SAM file systems store metadata and data on the same file system devices. You create a Sun StorageTek SAM file system through the browser interface by selecting Metadata and Data on Same Device in the New File System wizard.

Sun StorageTek QFS file systems store metadata and data on separate file system devices. They are created in the File System wizard by selecting the advanced option “Metadata and data on separate devices.”

Both Sun StorageTek SAM and Sun StorageTek QFS file systems can be made archiving or shared.

You can also configure Sun StorageTek QFS on Sun Cluster nodes for a stand-alone or shared, non-archiving, HA (highly available) file system.

Note: To create an HA-QFS configuration, you must also use the Sun Cluster management tool, SunPlex Manager, to configure the devices as part of a Sun Cluster resource group. File System Manager includes *ha* in the description of file systems that are configured on cluster devices even if the devices are not part of a resource group and the file systems are not therefore truly highly available.

You can create and administer the following file systems through the browser interface, based on the packages installed on the server.

TABLE 6-1 Supported File Systems

Packages	Product File Systems	Archive?	Share?	Notes
• SUNWsamfsr • SUNWsamfsu	• Sun StorageTek QFS • Sun StorageTek SAM	Yes	Yes	Remote or local management of these file systems is supported.
• SUNWqfsr • SUNWqfsu	Sun StorageTek QFS	No	Yes	Remote or local management of this file system is supported.

TABLE 6-1 Supported File Systems (Continued)

Packages	Product File Systems	Archive?	Share?	Notes
Either: • UFS software package, which is included in Solaris • SUNWsamfsr • SUNWsamfsu Or: • UFS software package, which is included in Solaris • SUNWqfsr • SUNWqfsu	UFS	No	No	
Either: • VERITAS file system software package • SUNWsamfsr • SUNWsamfsu Or: • VERITAS software package • SUNWqfsr • SUNWqfsu	VERITAS	No	No	VERITAS file systems are supported in view only mode. They cannot be created.

About Stand-Alone, Non-archiving File Systems

Stand-alone, non-archiving file systems are file systems that are not distributed and that do not use the storage and archiving software (Sun StorageTek SAM software) to archive files.

The stand-alone, non-archiving file systems that are supported by the browser interface are the Sun StorageTek QFS, Sun StorageTek SAM, UFS, and VERITAS file systems.

For information about the packages required on a server in order for it to host stand-alone, non-archiving file systems, see [“About Supported File Systems” on page 62](#).

This topic describes the Sun StorageTek QFS file system. For information about the UFS file system or the VERITAS file system, see the product documentation for those file systems.

The stand-alone, non-archiving Sun StorageTek QFS file system is a high-performance UNIX file system that enables you to store metadata on separate devices from file data.

Advantages of storing metadata on one or more devices separate from file data are that it reduces device head movement and rotational latency, improves RAID cache utilization, and enables the mirroring of metadata without the mirroring of file data.

Sun StorageTek QFS file systems are virtually unlimited with regard to file size, the number of files that can reside in a file system, and the number of file systems that you can create on a server.

These file systems support files up to 2^{63} bytes in length. Such large files can be striped across many disks or RAID devices, even within a single file system. This is because Sun StorageTek QFS file systems use true 64-bit addressing. A typical UFS file system is not a true 64-bit file system.

You can configure Sun StorageTek QFS file systems to support multiple RAID devices by defining striped groups in the file systems. Disk block allocation can be optimized for a striped group, reducing the overhead for updating the on-disk allocation map. You can assign a file to a striped group.

The number of Sun StorageTek QFS file systems you can configure on a server is unlimited. The volume manager enables each file system to include up to 252 device partitions (typically disk). Each partition can include up to 1 terabyte of data.

There is no predefined limit for the number of files on a Sun StorageTek QFS file system. This is because the inode space, which holds information about files, is dynamically allocated. Therefore, the maximum number of files is limited only by the amount of disk storage consumption in the file system.

If you want to configure the file system for archiving, there are guidelines to follow so that you don't overextend your environment.

For more information, see [“About Archiving, Releasing, and Staging” on page 113](#).

For more detailed information about the Sun StorageTek QFS file system, see the *Sun StorageTek QFS Installation and Upgrade Guide*.

About Archiving File Systems

Archiving file systems are file systems that use the Sun StorageTek SAM software to archive files. Using this software, you can restore files and the file systems themselves from remote media, in the event of a disaster. You can also manage the libraries or stand-alone tape drives that contain archive copies of the files.

The browser interface supports the creation and management of the Sun StorageTek SAM and the Sun StorageTek QFS file systems as archiving file systems.

For information about the packages required on a server in order for it to host archiving file systems, see [“About Supported File Systems” on page 62](#).

You create a Sun StorageTek SAM file system through the browser interface by selecting Metadata and Data on Same Device in the New File System wizard. By default, this creates an archiving Sun StorageTek SAM file system. This file system stores metadata and file data on the same devices. This file system is labeled as “qfs - archiving” in the browser interface.

You create an archiving Sun StorageTek QFS file system through the browser interface by selecting Metadata and Data on Separate Devices and Archive This File System in the New File System wizard. This automatically configures the file system to use the Sun StorageTek SAM software for archiving files.

With archiving file systems, the file systems can read and write files directly from magnetic disk, or can access archive copies of files as if the files were all on primary disk storage.

When managing archiving file systems through the browser interface, you can use the Archive Administration section of the navigation tree to configure how, where, and when to archive files.

For recommendations about configuring archiving based on your environment, see [“About Archiving, Releasing, and Staging” on page 113](#).

For information about the Sun StorageTek SAM and Sun StorageTek QFS file systems, see the *Sun StorageTek Storage Archive Manager (SAM) File System Configuration and Administration Guide* and the *Sun StorageTek QFS File System Configuration and Administration Guide*.

About Shared File Systems

Shared file systems are file systems that are distributed on multiple hosts. The advantage of a shared file system is that file data is passed directly from the Fibre Channel disks to the hosts. Data travels by way of local path I/O (direct access I/O). This is in contrast to the NFS service, which transfers data over the network.

The browser interface supports the creation and management of the Sun StorageTek QFS file system as a shared file system.

For information about the packages required on the server in order for it to host shared file systems, see [“About Supported File Systems” on page 62](#).

You configure a Sun StorageTek QFS file system to be shared by selecting Shared in the New File System wizard. When you configure the Sun StorageTek QFS file system to be shared, you configure it to be a distributed file system that can be mounted on multiple host systems. One host acts as the metadata server for the file system, and additional hosts act as clients.

More than one host can be configured as a potential metadata server, but only one host can be configured as the metadata server at any one time. There is no limit to the number of shared file system mount points.

Archiving shared file systems can be configured through the browser interface by selecting both the archiving and shared options in the File System wizard. When both of these options are enabled, the software manages archiving for the file system on the metadata server. If you later change the metadata server to another host, you must manually set up the archiver configuration on the new host.

Note: When you create a shared file system through the browser interface, the software updates the hosts configuration for the file system. However, the software does not create or update the `hosts.family-set-name` file for the file system. If you want to create the `hosts.family-set-name` file, see [“Creating the Shared Hosts File” on page 97](#).

The shared file system does not support the following file types:

- Blocked special files
- Character special files
- FIFO (named pipe) special files

Through the browser interface, you can perform the following tasks for a shared file system:

- Create a metadata server for a new file system
- Add clients to an existing file system

- Add a potential metadata server to an existing file system
- Mount or unmount the metadata server and the clients
- Change mount options
- Delete a client
- Delete the metadata server

All other operations related to shared file systems must be performed outside of the browser interface. For more detailed information about the shared Sun StorageTek QFS file system, see the *Sun StorageTek QFS Installation and Upgrade Guide*.

About HA-QFS File Systems

An HA-QFS (highly available QFS) file system is a non-archiving HA shared or stand-alone Sun StorageTek QFS file system running on cluster devices and configured to be part of a Sun Cluster resource group. Using File System Manager, you can configure Sun StorageTek QFS on cluster devices. You must use the SunPlex Manager tool to configure the devices as part of a Sun Cluster resource group.

Note: File System Manager includes *ha* in the description of file systems that are configured on cluster devices even if the devices are not part of a resource group and the file systems are therefore not truly highly available file systems.

Note: The browser interface does not support shared archiving HA file systems.

For stand-alone HA-QFS, the Sun StorageTek QFS file system is configured on several Sun Cluster nodes. Only one node at a time performs file system input and output processing.

For shared HA-QFS, all hosts within the shared Sun StorageTek QFS configuration must be nodes within the Sun Cluster environment. The Sun Cluster environment controls metadata server failover.

Administering Stand-Alone, Non-archiving File Systems and Archiving File Systems

This section provides information about administering stand-alone, non-archiving file systems and archiving file systems. It contains the following subsections:

- [“Planning a Stand-Alone, Non-archiving or Archiving File System” on page 68](#)
- [“Creating a Stand-Alone, Non-archiving or Archiving File System” on page 70](#)
- [“Growing a Stand-Alone, Non-archiving or Archiving File System” on page 71](#)
- [“Checking a Stand-Alone, Non-archiving or Archiving File System” on page 72](#)
- [“Mounting a Stand-Alone, Non-archiving or Archiving File System” on page 73](#)
- [“Editing Mount Options for a Stand-Alone, Non-archiving or Archiving File System” on page 73](#)
- [“Unmounting a Stand-Alone, Non-archiving or Archiving File System” on page 74](#)
- [“Deleting a Stand-Alone, Non-archiving or Archiving File System” on page 75](#)
- [“About Creating an Archiving File System and VSN Association” on page 75](#)

Planning a Stand-Alone, Non-archiving or Archiving File System

Before creating a stand-alone, non-archiving file system or an archiving file system, do the following:

- Decide whether you want the file system to be a stand-alone, non-archiving file system or an archiving file system. For information, see [“About Stand-Alone, Non-archiving File Systems” on page 63](#) and [“About Archiving File Systems” on page 65](#).
- Make sure the appropriate packages are installed on the server on which the stand-alone, non-archiving file system or archiving file system will reside. For more information, see [“About Supported File Systems” on page 62](#).
- If you are creating a Sun StorageTek QFS file system, determine the separate devices you want to use for metadata and file data.
- If you are creating an highly available QFS file system, determine which nodes in the Sun Cluster environment you want to use for the file system.
- If you are creating a Sun StorageTek SAM file system, determine the devices you want to use for both metadata and file data.
- Decide where you want the file system to be mounted on the current server. If the mount point is a shared NFS directory, you must mount the file system at boot time in the New File System wizard. Otherwise, NFS automatically shares the directory, possibly preventing you from mounting the file system at a later time.
- Decide the method by which you want files allocated on the file system: by the single allocation method or by the dual allocation method. For more information, see the allocation definitions in the [Glossary](#).

The following considerations apply only to archiving file systems:

- Before creating an archiving file system, determine whether your environment can support it. For information about archiving guidelines based on your environment, see [“About Archiving, Releasing, and Staging” on page 113](#).
- Ensure that there are available archive media and (optionally) VSN pools to which you can send archive copies of files from the archiving file system. You will be prompted to configure the default archive policy for the file system during creation. For more information see [“Creating a Disk VSN” on page 146](#), [“Adding a Library” on page 164](#), and [“Creating a VSN Pool” on page 143](#). Be aware of the default archive copy considerations as described in [“Creating a Stand-Alone, Non-archiving or Archiving File System” on page 70](#).
- Decide whether you want to create recovery points for the archiving file system, and if so, where you want the recovery point files stored. For more information, see [“Planning Recovery Points” on page 82](#).

Creating a Stand-Alone, Non-archiving or Archiving File System

Note: Before you create a stand-alone, non-archiving file system or an archiving file system, you must enable pop-up windows in your browser. For more information, see [“Enabling Pop-Up Windows” on page 15](#).

To create a stand-alone, non-archiving file system or an archiving file system:

1. From the Servers menu or the Servers page, choose the name of the server on which you want the file system to reside.

The File Systems Summary page is displayed.

2. Click New File System.

The New File System wizard is displayed.

3. Specify the type of file system that you want to create.

- To create a file system that stores metadata and file data on the same device, select the QFS option. This creates a default file system type with a default DAU size. If you want to modify the DAU size or specify metadata and data on separate devices, click Advanced Options.

Note: For general-purpose workloads, storing metadata and file data on the same device provides the best system performance.

- (Optional) If the server has the `SUNWsamfsu` package installed, you can configure the file system as archiving by selecting the Archiving option.
- (Optional) If the server is a node in a Sun Cluster configuration, select HA to configure highly available QFS (HA-QFS).

Note: To create an HA-QFS configuration, you must also use the Sun Cluster management tool, SunPlex Manager, to configure the devices as part of a Sun Cluster resource group. File System Manager includes *ha* in the description of file systems that are configured on cluster devices even if the devices are not part of a resource group and the file systems are not therefore truly highly available.

4. Complete the steps in the New File System wizard.

Click the Help tab in the wizard for more information.

Growing a Stand-Alone, Non-archiving or Archiving File System

When you grow a file system, you increase storage capacity by adding disk devices.

Through the browser interface, you can grow only the following file systems:

- An unmounted Sun StorageTek SAM file system
- An unmounted, stand-alone, non-archiving Sun StorageTek QFS file system
- An unmounted, archiving Sun StorageTek QFS file system

Note: To grow a UFS or shared file system, use the command-line interface for that file system.

To grow a file system:

1. From the Servers menu or the Servers page, choose the name of the server on which the file system is located.
The File Systems Summary page is displayed.
2. Select the radio button next to the file system you want to grow, and click Grow.
The Grow File System wizard is displayed.
3. If the file system has metadata and file data on the same device, skip to Step 6. Otherwise, select the metadata devices you want to add to the file system, and click Next.
Note: The software discovers available devices to the best of its ability. There might be circumstances in which the devices displayed are in use. For more information, see [“Administering Stand-Alone, Non-archiving File Systems and Archiving File Systems” on page 68.](#)
4. If the file system has striped groups, do the following:
 - a. Type the number of striped groups that you want to add to the file system and click Next. The software automatically assigns a name to each striped group that you add. A file system cannot contain more than 126 striped groups.
 - b. For each striped group that you add, select the available devices that you want as members of the striped group, and click Next. The devices must be the same size as each other.
5. From the list of available data devices, select the devices you want to add to the file system, and click Next.

6. Review the specified information for the new file system. If you are satisfied, click Finish.

Checking a Stand-Alone, Non-archiving or Archiving File System

The file checking process looks for, and optionally repairs, file corruption on the disk partitions belonging to the selected file system.

Note: You cannot check a UFS or shared file system through the browser interface. To perform this operation, use the command-line interface for that file system.

To check a stand-alone, non-archiving file system or an archiving file system:

1. From the Servers menu or the Servers page, choose the name of the server on which the file system is located.

The File Systems Summary page is displayed.

2. Select the radio button next to the file system whose files you want to check.
3. From the Operations menu, choose Check File System.

The Check File System window is displayed.

4. Select an option to either check the file system, or to check the file system and automatically make repairs.

Note: The check and repair option is only available if the file system is not mounted.

5. Specify a location to which the file system check report will be sent.
6. Click Submit.

Mounting a Stand-Alone, Non-archiving or Archiving File System

By mounting a stand-alone, non-archiving file system or an archiving file system, you attach the file system to the file system hierarchy at the specified mount point, which is the path name of a directory. Only an unmounted stand-alone, non-archiving file system or unmounted archiving file system can be mounted.

To mount a stand-alone, non-archiving file system or archiving file system:

1. From the Servers menu or the Servers page, choose the name of the server on which the file system is located.

The File Systems Summary page is displayed.

2. Select the radio button next to the file system that you want to mount.

3. From the Operations menu, choose Mount.

The file system is redisplayed on the File Systems Summary page with a disk usage value in its Disk Usage (Total) column.

Editing Mount Options for a Stand-Alone, Non-archiving or Archiving File System

Through the browser interface, you can edit the mount options of a Sun StorageTek SAM file system; a stand-alone, non-archiving Sun StorageTek QFS file system; or an archiving Sun StorageTek QFS file system.

When you edit mount options, some changes cannot take effect until you unmount and then mount the stand-alone or archiving file system. If the file system is currently mounted, some mount options are unavailable for editing.

Follow these additional guidelines when editing mount options:

- If you want to change a mount option value to the default value, delete the existing value in the field that you want to change and click Save.

- If you want to enable or disable mounting of the file system at boot time, you must manually edit the `mount at boot` field in the `/etc/vfstab` file located on the current server. For more information about this file, see the *Sun StorageTek QFS Installation and Upgrade Guide* or the *Sun StorageTek Storage Archive Manager (SAM) Installation and Upgrade Guide*.

To edit the mount options of a stand-alone, non-archiving file system or archiving file system:

1. From the Servers menu or the Servers page, choose the name of the server on which the file system is located.
The File Systems Summary page is displayed.
2. Select the radio button next to the file system whose mount options you want to edit.
3. From the Operations menu, choose Edit Mount Options.
The file system Edit Mount Options page is displayed.
4. Edit the fields that you want to change and click Save.

Unmounting a Stand-Alone, Non-archiving or Archiving File System

You must unmount a stand-alone, non-archiving file system or an archiving file system before performing certain operations such as growing the file system, deleting the file system, or upgrading hardware. Before you unmount an archiving file system, you must idle or stop any archiving operation that is in process.

To unmount a stand-alone, non-archiving file system or archiving file system:

1. From the Servers menu or the Servers page, choose the name of the server on which the file system is located.
The File Systems Summary page is displayed.
2. Select the radio button next to the file system that you want to unmount.
3. From the Operations menu, choose Unmount.
A message box prompts you to confirm your choice.
4. Click OK.
The file system is redisplayed on the File Systems Summary page with the unmounted value displayed in its Disk Usage (Total) column.

Deleting a Stand-Alone, Non-archiving or Archiving File System

You can delete selected stand-alone, non-archiving file systems or archiving file systems from the server when they are no longer required. Before you can delete an archiving file system, you must stop archiving, and unmount the archiving file system.

Caution: The process of deleting a file system removes all data from the file system. If you want to preserve any of the data on this file system, make sure to copy it to another file system before proceeding.

Note: You cannot delete an archiving file system if the file system is the last file system of a policy or of a policy's file match criteria. If this is the case, you must delete that policy or file match criteria before deleting the file system.

To delete a stand-alone, non-archiving file system or an archiving file system:

1. From the Servers menu or the Servers page, choose the name of the server on which the file system is located.

The File Systems Summary page is displayed.

2. Select the radio button next to the file system that you want to delete.

3. From the Operations menu, choose Delete.

A message box prompts you to confirm your choice.

4. Click OK.

About Creating an Archiving File System and VSN Association

Note: This information applies to **version 4 update 5** servers only.

Each archiving file system on a server has a default archive copy that archives all files that are not explicitly members of an archive policy. If an `archiver.cmd` file exists on the server, there must be a volume association for each file system's default archive copy. A volume serial name (VSN) identifies a volume.

If an `archiver.cmd` file already exists on the server when you create an archiving file system through the browser interface, the File System Manager software automatically creates a VSN association with an available or valid media type for the default archive copy.

If an `archiver.cmd` file does not already exist on the server when you create an archiving file system, the VSN association is not explicitly created, and the default archiving behavior is retained. In this situation, you can create an archive policy from the Policies page and apply the policy to the file system. This action creates the `archiver.cmd` file and the necessary VSN association for the default archive copy of the file system.

To change these default copy definitions, you can manually edit the `archiver.cmd` file at a later time.

Managing Policies for Archiving File Systems

This section provides information about managing policies for an archiving file system. It contains the following subsections:

- [“Adding a Policy to an Archiving File System” on page 77](#)
- [“Adding Existing File Match Criteria to an Archiving File System” on page 77](#)
- [“Reordering File Match Criteria for an Archiving File System” on page 78](#)
- [“Removing File Match Criteria From an Archiving File System” on page 79](#)

Adding a Policy to an Archiving File System

Adding a policy to an archiving file system creates a new policy that is applied directly to the selected file system. This enables you to change archiving behavior by defining new file match criteria and policy copy directives for the file system.

To add a policy to a file system:

1. From the Servers menu or the Servers page, choose the server on which the file system that you want resides.

The File Systems Summary page is displayed.

2. Select the radio button next to the file system to which you want to add a new policy, and click View Policies.

The File System Archive Policies page is displayed.

3. Click Add New Policy.

The New Policy wizard is displayed.

4. Complete the steps in the New Policy wizard as described in the Help tab of the wizard.

Adding Existing File Match Criteria to an Archiving File System

To change archiving behavior for a selected archiving file system, you can add one or more sets of file match criteria from a list of existing file match criteria.

When you add file match criteria to an archiving file system, the files in that file system will be archived according to the archiving directives defined by the new file match criteria.

To add existing file match criteria to an archiving file system:

1. From the Servers menu or the Servers page, choose the name of the server on which the archiving file system is located.

The File Systems Summary page is displayed.

2. Select the radio button next to the archiving file system to which you want to add file match criteria, and click View Policies.

The File System Archive Policies page is displayed.

3. In the Archive Policy Criteria table, click Add Existing Criteria.

The Available Criteria window is displayed.

4. From the list of available file match criteria, select the file match criteria you want to add to the file system, and click Add.

Reordering File Match Criteria for an Archiving File System

You can reorder file match criteria associated with an archiving file system so that the file match criteria archive directives are applied in a new sequence.

To reorder file match criteria for an archiving file system:

1. From the Servers menu or the Servers page, choose the name of the server on which the archiving file system is located.

The File Systems Summary page is displayed.

2. Select the radio button next to the archiving file system for which you want to reorder file match criteria, and click View Policies.

The File System Archive Policies page is displayed.

3. In the Archive Policy Criteria table, click Reorder.

The Reorder Match Criteria window is displayed.

4. Select the file match criteria that you want, and click either Move Up or Move Down to change the order of the file match criteria in the list.

Repeat this step until all of the file match criteria are in the order you want.

5. Click Submit.

Removing File Match Criteria From an Archiving File System

You can remove the relationship between file match criteria in a policy and an archiving file system. Removing the relationship between file match criteria and an archiving file system is not the same as removing the file match criteria from a policy. When you remove file match criteria from an archiving file system, the file system's files are no longer archived according to the directives defined by that set of criteria. An archiving file system can be associated with more than one set of file match criteria from several different policies.

File match criteria cannot be removed from an archiving file system if the file system is the sole file system associated with the file match criteria. In this case, the file match criteria must be deleted for the association to be removed. If the set of file match criteria is the sole set of file match criteria in a policy, you must delete the policy.

To remove file match criteria from an archiving file system:

1. From the Servers menu or the Servers page, choose the name of the server on which the archiving file system is located.

The File Systems Summary page is displayed.

2. Select the radio button next to the archiving file system from which you want to remove the file match criteria, and click View Policies.

The File System Archive Policies page is displayed.

3. In the Archive Policy Criteria table, select the radio button next to the policy and file match criteria combination that you want to remove from the file system.

4. Click Remove.

A message box prompts you to confirm the removal.

5. Click OK.

Protecting Archiving File System Data

You can protect the files of an archiving file system by regularly creating recovery points and by using media. This section provides information about protecting files on an archiving file system. It contains the following subsections:

- [“About Protecting Archiving File System Data” on page 80](#)
- [“Scheduling and Creating Recovery Points” on page 81](#)
- [“Restoring Archiving File System Data” on page 86](#)

About Protecting Archiving File System Data

To protect the data on archiving file systems, you can schedule recovery points for the file systems and you can archive copies of files from the file systems to archive media. A recovery point contains all the metadata for a complete archiving file system at a single point in time.

File data consists of the actual files and directories of a file system. Metadata consists of information about that file data. In particular, metadata is data about files, file paths, the inodes for files, directories, access control limits, symbolic links, removable media, segmented files, and the indexes of segmented files. Before lost file data can be retrieved from archive media, the metadata must be restored.

When a recovery point is created through the browser interface, it contains all the metadata for the file system, including pointers to the archive locations for each file. By using this metadata, the File System Manager software can recover file data. Store recovery point files separately from the file system to ensure that they are available if the file system is lost.

With archiving file systems, the archiver, which is a component of the Sun StorageTek SAM software, can copy file data to archive media. This protects the data because it enables you to retrieve the data from the archive media at a later time, as long as the media onto which archive copies are written is not erased and as long as File System Manager can access the recovery point file.

Scheduling and Creating Recovery Points

This section provides information about scheduling and creating recovery points for an archiving file system. It contains the following subsections:

- [“About Recovery Points” on page 81](#)
- [“Planning Recovery Points” on page 82](#)
- [“Scheduling the Creation of Recovery Points” on page 83](#)
- [“Creating a Recovery Point on Demand” on page 84](#)
- [“Editing a Recovery Point Schedule” on page 85](#)
- [“Removing a Recovery Point Schedule” on page 85](#)

About Recovery Points

You can schedule recovery points for the archiving file systems on the server.

Scheduling a recovery point through the browser interface is the equivalent of creating an entry in the `root crontab(1)` file to automate the Sun StorageTek SAM software `samfsdump(1M)` process. This process enables you to back up metadata separate from the file data of an archiving file system. A recovery point contains all the metadata for a complete archiving file system at a single point in time.

Note: Files written to an archiving file system after the creation of a recovery point starts might not be archived, and archive copies on cartridges might not be reflected in the recovery point file. Because of this, the files might not be known to the system if the recovery point is used to restore the archiving file system. Files written to the archiving file system or archived after the recovery point is created will be included in the next recovery point.

Recovery points can be run only on mounted, archiving file systems. A recovery point does the following:

- Saves the relative path for each file.

- Contains all information required for restoring an archiving file system, such as the inode for each file, directory information, and symbolic links.
- Enables you to restore an entire archiving file system, a directory hierarchy, or a single file. See [“Restoring an Archiving File System” on page 86](#) for more information about the restore process.

Planning Recovery Points

The following table provides planning recommendations that can help you protect the data of an archiving file system.

TABLE 6-2 Recommendations for Planning Recovery Points

Recommendation	Reason
Keep data from each archiving file system on its own unique set of tapes or other archive media, and do not mix data from multiple file systems.	The path names of the files on archive media do not include the originating file system’s name, and this might cause inconsistencies when you restore files. For more information, see the <i>Sun StorageTek Storage Archive Manager (SAM) Troubleshooting Guide</i> .
Make sure the archiving file system for which you want to schedule recovery points is mounted.	Recovery points cannot be scheduled for unmounted file systems.
Store recovery point files for each file system in a separate directory.	Keeping the recovery points in separate files helps you locate the recovery point you need to use to restore a file system.
Schedule recovery points for a time when files are not being created or modified.	Creating a recovery point when files are not being created or modified will help ensure that the file data being captured is not stale and will minimize the creation of damaged files.
Archive your files regularly, and schedule recovery points at least once a day.	To ensure that files can be restored, configure archiver scan intervals so that files are archived before a recovery point is created.
Make sure that the location on which you are creating and storing recovery point files has sufficient storage space for these files.	Recovery point files can be quite large, depending on the number of files being stored on the archiving file system.
Periodically check the recovery point log file to ensure that recovery points are being created without any errors.	Checking the log file for issues can prevent the creation of recovery points that contain damaged files. You specify the location of the recovery point log file on the Recovery Point Schedule page.

TABLE 6-2 Recommendations for Planning Recovery Points (*Continued*)

Recommendation	Reason
In anticipation of creating a recovery point on demand, determine in advance the approximate length of time it will take to create the recovery point and the approximate size of the recovery point file.	Determining the length of time and the size of the file can help you decide whether to wait for the scheduled recovery point creation. You can use the following values to make this determination: • Length of time = 5 to 10 Gbytes of data per hour • Recovery point file size = 666 * number of files on the file system These values will vary, depending on the disk transfer and seek speeds of the systems being used.
Decide on a retention period for your recovery point files.	Retaining all recovery point files indefinitely uses up space on your system. You can specify a retention period on the Recovery Point Schedule page. For example, you can specify that each time a recovery point is created, all recovery point files older than two days (48 hours) should be deleted.
Determine who should receive notifications of recovery point errors.	Recovery points are critical in maintaining a reliable restore capability of the file system. If a recovery point fails, system administrators should be notified. You can specify addresses to which you want error notifications sent on the Email Alert Summary page, which is accessible from the Recovery Point Schedule page.

Scheduling the Creation of Recovery Points

By default, archiving file systems do not have recovery points scheduled. For information about planning for recovery points, see [“Planning Recovery Points” on page 82](#).

To schedule the creation of recovery points:

1. From the Servers menu or the Servers page, choose the server on which the archiving file system is located.
The File Systems Summary page is displayed.
2. Select the radio button next to the mounted archiving file system for which you want to schedule a recovery point.
3. From the Operations menu, choose Schedule Recovery Points.
The Schedule Recover Points page is displayed.
4. Specify the appropriate values, and click Save.

After recovery point files have been created, periodically check the recovery point log file to ensure that recovery points are being created without errors or damaged files.

Creating a Recovery Point on Demand

You might want to manually create a recovery point instead of creating a recovery point at the regularly scheduled time, in order to capture the metadata from an archiving file system at a specific point in time.

To create a recovery point on demand if you have the **version 4, update 6** software installed:

1. From the Servers menu or the Servers page, choose the server on which the archiving file system is located.
2. In the navigation tree, select Recovery Points under File Browsing & Recovery.
3. Select the file system from the drop-down menu and click Create Recovery Point Now.

To create a recovery point on demand if you have the **version 4, update 5** software installed:

1. From the Servers menu or the Servers page, choose the server on which the archiving file system is located.
The File Systems Summary page is displayed.
2. Select the radio button next to the mounted archiving file system for which you want to create a recovery point.
3. From the Operations menu, choose Take Recovery Point.
The Take Recovery Point window is displayed.
4. In the Fully Qualified Recovery Point File Name field, type the path and the name of the recovery point file that you want to create.
5. Click Submit.

For information about how long it might take to create a recovery point file when taken on demand, see [“Planning Recovery Points” on page 82](#).

Editing a Recovery Point Schedule

To edit the recovery point schedule for an archiving file system:

1. From the Servers menu or the Servers page, choose the server that has the recovery point schedule you want to edit.
The File Systems Summary page is displayed.
2. Select the radio button next to the mounted archiving file system for which you want to edit the recovery point schedule.
3. From the Operations menu, choose Schedule Recovery Points.
The Schedule Recovery Points page is displayed.
4. Edit the fields that you want to change, and click Save.

Note: If the server is running version 4, update 6 of the software, you can also edit a recovery point schedule by selecting Scheduled Tasks under System Administration in the navigation tree.

Removing a Recovery Point Schedule

Removing a recovery point schedule from an archiving file system does not remove the existing recovery point files from the current server.

You can remove recovery point files using the Delete Recovery Point button on the Restore File System page. For more information, see [“Deleting a Recovery Point” on page 93](#).

To remove a recovery point schedule from an archiving file system if the **version 4, update 6** software is installed:

1. From the Servers menu or the Servers page, choose the server that has the recovery point schedule you want to delete.
2. In the navigation tree, select Scheduled Tasks under System Administration.
The Scheduled Tasks page is displayed.
3. Select the radio button next to the recovery point schedule that you want to delete.
4. Click Remove.

To remove a recovery point schedule from an archiving file system if the **version 4, update 5** software is installed:

1. From the Servers menu or the Servers page, choose the server that has the recovery point schedule you want to delete.

The File Systems Summary page is displayed.

2. Select the radio button next to the mounted archiving file system for which you want to remove the recovery point schedule.
3. From the Operations menu, choose Remove Recovery Point Schedule.
A message box prompts you to confirm the removal.
4. Click OK.

Restoring Archiving File System Data

This section provides information about restoring data for an archiving file system. It contains the following subsections:

- [“Restoring an Archiving File System” on page 86](#)
- [“Filtering Recovery Point Files” on page 90](#)
- [“Filter File Name Pattern Values” on page 91](#)
- [“Deleting the Index for a Recovery Point” on page 92](#)
- [“Deleting a Recovery Point” on page 93](#)

Restoring an Archiving File System

With archiving file systems, you can configure the archiver to write copies of files to archive media. After files are archived, information about the files can be captured in a recovery point. In the event of a disaster, you can then use the recovery point file to restore the files or the file system itself.

You can restore file systems using any recovery point, including recovery points generated by other file systems, as long as the archived file data from the file system is accessible.

You can also use multiple recovery points to restore files into a single file system. Conversely, you can use a single recovery point to restore file data into multiple file systems. For example, you could restore directory `dir1` from recovery point file `snapshot1` into file system `mysam`, and also restore directory `dir2` from `snapshot2` into the same file system (`mysam`). You could also restore directory

`dira` from recovery point file `snapshot3` into file system `sam1` and directory `dirb` from the same recovery point file (`snapshot3`) to a different file system, for example `sam2`.

For information about planning recovery points, see [“Planning Recovery Points” on page 82](#).

To restore a file system if the software **version 4 update 6** is installed:

1. From the Servers menu or the Servers page, choose the name of the server on which the file system is located.
2. In the navigation tree, select File Browser under File Browsing & Recovery.
3. Select the file system mount point in the Current Directory drop-down list.
The file system files and directories are displayed in the table.
4. Select the Recovery Point radio button at the top of the table, and then select a recovery point date from the drop-down list.
5. If you cannot find the recovery point date in the drop-down list, follow these steps to bring the recovery point file online for browsing:
 - a. In the navigation tree, select Recovery Points under File Browsing & Recovery.
 - a. Select the radio button next to the recovery point.
 - b. Click Create Index.
An index file is created. If the recovery point file was compressed using the `compress` utility, a new decompressed version of the recovery point file is also created. The original version of the recovery point file remains on the server. Files compressed with `gzip` do not requiring decompressing.
Note: If the recovery point is a headerless file created using the Sun StorageTek SAM `samfsdump(1M)` command with the `-H` option, the system cannot index the recovery point, and you cannot restore the file system using that recovery point with the File System Manager software. For information on restoring data using these files, see the *Sun StorageTek Storage Archive Manager Troubleshooting Guide*.
 - c. Return to the File Browser page and follow steps 3 and 4 above.
6. (Optional) Click Browse to select a different directory from which to display the recovery point.
7. Do one or more of the following to find the files that you want to restore:
 - Navigate through the recovery point by browsing through the files and directories in the table.

- Click Filter to limit the items in the table so that only the information that interests you is displayed. Then click Apply to apply the filter criteria. For more information, see [“Filtering Recovery Point Files” on page 90](#).
 - Type a value in the Maximum Entries field to limit the total number of items that are read from the recovery point file.
8. To restore the whole file system, select Restore Entire Recovery Point from the Operations menu. Alternatively, to restore a file or directory, follow these steps:
 - a. Select the radio button next to the file or directory in the Recovery Point Entries table and select Restore from the Operations menu.
 - b. Specify the location to which you want to restore the file or directory. By default, the location is the path of the original file or directory, relative to the mount point of the file system. You can specify a different path relative to the mount point, or you can specify an absolute path on any archiving file system.
 - c. Specify the action you want the system to take if the file already exists in the location to which you are restoring files.
 9. From the Online Status After Restoring menu, choose the actions that you want the file system to take after completing the restore process.
 10. Click Submit.

To restore a file system if the software **version 4 update 5** is installed:

1. From the Servers menu or the Servers page, choose the name of the server on which the file system is located.
The File Systems Summary page is displayed.
2. Select the radio button next to the mounted archiving file system that contains the files you want to restore, and choose Restore from the Operations menu.
The Restore File System page is displayed with a Recovery Points table listing the recovery points available in the current directory.
3. (Optional) Click Browse to select a different directory from which to display the Recovery Point Summary table.
4. If the recovery point file is not displayed as a link in the Recovery Point Summary table, follow these steps to bring the recovery point file online for browsing:
 - a. Select the radio button next to the recovery point.
 - b. Click Create Index.
An index file is created. If the recovery point file was compressed using the `compress` utility, a new decompressed version of the recovery point file is also created. The original version of the recovery point file remains on the server. Files compressed with `gzip` do not requiring decompressing.

Note: If the recovery point is a headerless file created using the Sun StorageTek SAM `samfsdump(1M)` command with the `-H` option, the system cannot index the recovery point, and you cannot restore the file system using that recovery point with the File System Manager software. For information on restoring data using these files, see the *Sun StorageTek Storage Archive Manager Troubleshooting Guide*.

5. When the recovery point is displayed as a link in the Recovery Points table, do one of the following:
 - Click the recovery point file to browse its contents.
 - Select the radio button next to the recovery point file and click Show Contents.

The Restore File System page is refreshed, and the top-level items in the selected recovery point file are displayed in the Recovery Points table.

Note: Top-level items are displayed unless you previously selected a recovery point file and the directory that was displayed from that recovery point is found in the newly selected recovery point file.

6. Do one or more of the following to find the files that you want to restore:
 - Navigate through the recovery point by clicking directories and by using the Up One Level or Go to Directory buttons.
 - Click Filter to limit the items in the table so that only the information that interests you is displayed. For more information, see [“Filtering Recovery Point Files” on page 90](#).
 - Click Search Recovery Point to search the whole recovery point file and display results based on your search criteria.
 - Type a value in the Maximum Recovery Point Entries Retrieved field to limit the total number of items that are read from the recovery point file.
7. To restore the whole file system, select Entire Contents of Recovery Point in the Restore section of the page. Alternatively, to restore a file or directory, follow these steps:
 - a. Select the radio button next to the file or directory in the Recovery Point Entries table.

The file or directory that you select is displayed in the File or Directory Selected in Table field.
 - b. Specify the location to which you want to restore the file or directory. By default, the location is the path of the original file or directory, relative to the mount point of the file system. You can specify a different path relative to the mount point, or you can specify an absolute path on any archiving file system.
 - c. Specify the action you want the system to take if the file already exists in the location to which you are restoring files.

8. From the Online Status After Restoring menu, choose the actions that you want the file system to take after completing the restore process.
9. Click Restore.

Filtering Recovery Point Files

You can control which files and directories in a recovery point file are displayed in the browser interface. Filtering a recovery point file displays the files and directories that meet the filter criteria in a hierarchical fashion and enables you to navigate through the hierarchy.

When filtering the contents of a recovery point file, be aware of the following:

- You must choose at least one filter criterion.
- A filter applies to the contents of the selected recovery point file only, not to multiple recovery point files.
- A filter remains in effect as you navigate into directories and subdirectories of a recovery point file. For example, you might find a directory named `user` by specifying `*us` as a filter criterion. When you navigate through the `user` directory, the filter is still applied, and so not all files are displayed. To view all files in that directory, reset or remove the filter criterion.
- To remove filter criteria from the table, click the Clear Filter button.
- Specifying multiple values in the Filter Recovery Point window is the equivalent of using AND in a search.

The following table describes the filter values that you can specify.

TABLE 6-3 Filter Values

Filter Value	Description
File Name Pattern field	Type a value to filter for files and directories with the specified name. For more information about wildcard characters that you can use in this field, see “Filter File Name Pattern Values” on page 91 .
File Size greater than or equal to	Type a value to filter for files that are larger than the specified size.
File Size less than or equal to	Type a value to filter for files that are smaller than the specified size.
File Date	Filter for files that are created or that are modified during the specified time frame.
Owner	Type a value to filter for files that are owned by the specified user, as defined on the server from which the files originated.

TABLE 6-3 Filter Values (*Continued*)

Filter Value	Description
Group	Type a value to filter for files that are owned by users who are members of the specified group, as defined on the server from which the files originated.
Is Damaged	Filter for files that are damaged.
Is Online	Filter for files that were online at the time the recovery point was created.

Filter File Name Pattern Values

When you filter a recovery point file, you can filter for files and directories based on the file or directory name. You do this by typing a value in the File Name Pattern field. For more information about the patterns, see the `fnmatch(3c)` man page.

The following table lists valid wildcard characters that you can type in the File Name Pattern field.

TABLE 6-4 File Name Wildcards

Wildcard	Description
?	Filters for files that contain any character in place of the question mark. For example, <code>a?d</code> returns files named <code>aad</code> , <code>abd</code> , <code>acd</code> , and so on.
*	Filters for files that contain any number of characters in place of the asterisk. For example, <code>*a*d</code> returns files named <code>ad</code> , <code>abcd</code> , <code>efabcd</code> , <code>aaaad</code> , and <code>adddd</code> .
[]	Filters for files that contain one of the alternatives specified within the brackets. For example, <code>a[bc]</code> returns files named <code>ab</code> or <code>ac</code> .
""	Filters for files that contain only the specified characters. For example, <code>"abc"</code> returns only files named <code>abc</code> .
\	Removes the wildcard distinction from any wildcard character that you use with the backslash. For example, <code>"*\(\?"</code> returns files with <code>(?</code> in their names. To filter for files that have a backslash in their name, type two backslashes <code>\\</code> .
/	Filters for files in a path name. The slash character in a path name must be explicitly matched by one or more slashes in the File Name Patter field (for example, <code>/var</code>). It cannot be matched by an asterisk, question mark, or bracket. Slashes in the File Name Pattern field are identified before brackets. This means that, for example, <code>a[b/c]d</code> does not return files named <code>abd</code> or <code>a/d</code> . It will return only a file named <code>a[b/c]d</code> .

The following table shows examples of different File Name Pattern values and the files or directories they return.

TABLE 6-5 File Name Patterns

File Name Pattern	Some Possible Files or Directories Returned
abc	abc, abcd, gabc
"abc"	abc
a?c	abc, accd, gabc, gaggc
a*c	abbc, afilenameec, gabc, abc, gaggc
"a?c"	a?c
a*c	a*c
*/var	abc/var, path/var

Deleting the Index for a Recovery Point

If you used a recovery point file to restore a file system, an index file was created for the recovery point. If the recovery point file was compressed using the `compress` utility, a new decompressed version of the recovery point file was also created. (A decompressed version of the file is not necessary if it was compressed using `gzip`.) Because these files can be large, you might want to delete them after the restore process is finished.

Deleting the index and, if applicable, the decompressed version of a recovery point file does not delete the original recovery point file.

You can remove recovery point files using the Delete Recovery Point button on the Restore File System page. For more information, see [“Deleting a Recovery Point” on page 93](#).

To delete the index and the decompressed version of a recovery point file from a server that has **version 4, update 6** of the Sun StorageTek SAM software:

1. From the Servers menu or the Servers page, choose the name of the server that contains the files you restored.
2. In the navigation tree, select Recovery Points under File Browsing & Recovery.
The Recovery Points page is displayed.
3. Select the file system from the drop-down list at the top of the table.
The recovery points for the selected file system are displayed.

4. Select the radio button next to the recovery point file for which you manually created an index.
5. Click Delete Index.

A message box prompts you to confirm your choice.

6. Click OK.

The index file is deleted. If the recovery point file had a decompressed version of itself created by the index creation process, the decompressed version of the recovery point is also deleted. The original recovery point file remains on the server.

To delete the index and the decompressed version of a recovery point file from a server that has **version 4, update 5** of the Sun StorageTek SAM software:

1. From the Servers menu or the Servers page, choose the name of the server that contains the files you restored.

The File Systems Summary page is displayed.

2. Select the radio button next to the mounted archiving file system for which you have restored files, and choose Restore from the Operations menu.

The Restore File System page is displayed.

3. Select the radio button next to the recovery point file for which you manually created an index.

4. Click Delete Index.

A message box prompts you to confirm your choice.

5. Click OK.

The index file is deleted. If the recovery point file had a decompressed version of itself created by the index creation process, the decompressed version of the recovery point is also deleted. The original recovery point file remains on the server.

Deleting a Recovery Point

To delete a recovery point file if you have **version 4, update 6** of the Sun StorageTek SAM software installed:

1. From the Servers menu or the Servers page, choose the name of the server that contains the recovery point.

2. In the navigation tree, select Recovery Points under File Browsing & Recovery.

The Recovery Points page is displayed.

3. Select the file system from the drop-down list at the top of the table.

The recovery points for the selected file system are displayed.

4. Select the radio button next to the recovery point file that you want to delete.
5. Click Delete Recovery Point.

A message box prompts you to confirm your choice.

6. Click OK.

The recovery point file is deleted. Any associated index or decompressed files are also deleted.

To delete a recovery point file if you have **version 4, update 5** of the Sun StorageTek SAM software installed:

1. From the Servers menu or the Servers page, choose the name of the server that contains the recovery point.

The File Systems Summary page is displayed.

2. Select the radio button next to the mounted archiving file system that contains the files you have restored, and choose Restore from the Operations menu.

The Restore File System page is displayed.

3. Select the radio button next to the recovery point file you want to delete.
4. Click Delete Recovery Point.

A message box prompts you to confirm your choice.

5. Click OK.

The recovery point file is deleted. Any associated index or decompressed files are also deleted.

Administering Shared File Systems

This section provides information about administering shared file systems. It contains the following subsections:

- [“Planning a Shared File System” on page 95](#)
- [“Creating a Shared File System” on page 96](#)
- [“Creating the Shared Hosts File” on page 97](#)
- [“Adding a Client or Potential Metadata Server” on page 98](#)
- [“Deleting a Client or Potential Metadata Server” on page 99](#)
- [“Deleting the Metadata Server” on page 100](#)
- [“Mounting and Unmounting Shared File Systems” on page 100](#)
- [“Deleting a Shared File System” on page 105](#)

Planning a Shared File System

Before you can create, add a client to, or add a potential metadata server to a shared file system, the following prerequisites must be met:

- All participating hosts must have either of the following packages installed locally:
 - SUNWqfsr and SUNWqfsu
 - SUNWsamfsr and SUNWsamfsu
- All Sun StorageTek QFS and Sun StorageTek SAM software packages installed on the participating hosts must be at the same release level.
- All participating hosts must be added to the File System Manager browser interface, as described in [“Adding a Server” on page 26](#).

- The File System Manager service, `fsmgmtd`, must be running and available on each host that will participate in the shared file system. For more information about the `fsmgmtd` service, see the *Sun StorageTek QFS Installation and Upgrade Guide*.

Consider the following before creating a shared file system:

- Determine the devices you want to use for metadata and file data in the file system. Metadata devices must be accessible from every host that is a potential metadata server in the file system. Data devices must be accessible from every host in the file system (clients, metadata server, and potential metadata servers).
- Decide where you want the file system to be mounted on the current server. If the mount point is a shared Network File System (NFS) directory, you must choose to mount the file system at boot time in the New File System wizard. Otherwise, NFS automatically shares the directory, possibly preventing you from mounting the file system at a later time.
- Decide the method by which you want files allocated on the file system: by the single allocation method or by the dual allocation method. For more information, see the allocation definitions in the [Glossary](#).

Creating a Shared File System

Before you can create a shared file system, you must have prepared your configuration as described in [“Planning a Shared File System” on page 95](#).

You might also want to change how your configuration entries are saved. For more information, see [“Enabling Pop-Up Windows” on page 15](#).

To create a shared file system:

1. From the Servers menu or the Servers page, choose the name of the server on which you want the shared file system to reside.

The File Systems Summary page is displayed.

2. Click New File System.

The New File System wizard is displayed.

3. Specify the type of file system that you want to create.

- To create a file system that stores metadata and file data on the same device, select the QFS option. This creates a default file system type with a default DAU size. If you want to modify the DAU size or specify metadata and data on separate devices, click Advanced Options.

Note: For general-purpose workloads, storing metadata and file data on the same device provides the best system performance.

- Select the Shared option.
- (Optional) If the server has the `SUNWsamfsu` package installed, you can configure the file system as archiving by selecting the Archiving option.

Note: You cannot configure an archiving file system to be HA through the browser interface.

- (Optional) If the server is a node in a Sun Cluster configuration, select HA to configure highly available QFS (HA-QFS).

Note: To create an HA-QFS configuration, you must also use the Sun Cluster management tool, SunPlex Manager, to configure the devices as part of a Sun Cluster resource group. File System Manager includes *ha* in the description of file systems that are configured on cluster devices even if the devices are not part of a resource group and the file systems are not therefore truly highly available.

4. Complete the steps in the New File System wizard.

Click the Help tab in the wizard for more information.

Creating the Shared Hosts File

If you want to modify the host configuration of a shared file system outside of the browser interface, you can create the shared hosts file (`hosts.family-set-name`) for the file system and then edit the file manually with the new hosts. You might want to do this in order to add Linux clients to the shared file system.

The procedure in this topic describes how to create the shared hosts file. For information about how to edit the file, see “To Change the Shared Hosts File Information on an Unmounted File System” in the *Sun StorageTek QFS File System Configuration and Administration Guide*.

To create the shared hosts file for a shared file system:

1. Use `telnet` to connect to the metadata server of the shared file system, and log in as `root`.
2. Navigate to the `/etc/opt/SUNWsamfs` directory.
3. Use the `samsharefs(1M) -R` command to create the shared hosts file:

```
# samsharefs -R shared-file-system-name > hosts.family-set-name
```

Adding a Client or Potential Metadata Server

Before you can add a client or potential metadata server to an existing shared file system, you must have prepared your configuration as described in [“Planning a Shared File System” on page 95](#).

After you add a potential metadata server to the shared file system, the File System Manager software automatically updates the `hosts` file on the potential metadata server system.

To add a client or potential metadata server to a shared file system:

1. From the Servers menu or the Servers page, choose a server that is acting as the metadata server or as the potential metadata server for the shared file system.

The File Systems Summary page is displayed.

2. Scroll through the table and click the shared file system name.

The File System Details page is displayed.

3. Click Add.

The Add a Potential Metadata Server or a Shared Client window is displayed.

4. Choose the type of host you want to add:

- If you want to add a client to the shared file system, choose Shared Client.
- If you want to add a potential metadata server to the shared file system, choose Potential Metadata Server.

5. From the Host menu, choose the host that you want to add to the shared file system.

The menu lists hosts that have been added to the browser interface, that have the same Sun StorageTek QFS and Sun StorageTek SAM release levels as the metadata server, and that have the Sun StorageTek QFS package installed locally.

6. If you are adding a potential metadata server, select or deselect Read Only as follows:

- If you want the client to have full access to data in the file system, deselect it.
- If you want the client to have read-only access to data in the file system, select it.

7. In the Mount Point field, type the full directory path to which you want the file system attached (mounted) within your system's directory tree. Use the following syntax:

/directory-name/file-system-name

Note: If the directory that you specify already exists, you can access existing files in that directory only when the file system is unmounted.

8. If you want to create the mount point that you specified or you want the file system to be mounted at boot time, select the appropriate check boxes.
9. Do one of the following:
 - To create and mount the file system, select the After Addition, Mount option. This option is displayed only if the metadata server for the shared file system is mounted.
 - To create the file system without mounting it, select the After Addition, Do Not Mount option.
10. Click OK.

Deleting a Client or Potential Metadata Server

Before you can delete a client or potential metadata server from a shared file system, the shared file system must be unmounted on all participating hosts. For information about unmounting a shared file system, see [“Unmounting a Shared File System” on page 103](#).

To delete a client or potential metadata server:

1. From the Servers menu or the Servers page, choose the name of the server that is acting as the metadata server or as the potential metadata server for the shared file system.

The File Systems Summary page is displayed.
2. Scroll through the table and click the shared file system name.

The File System Details page is displayed.
3. Select the radio button next to the client or potential metadata server that you want to delete, and click Delete.

A message box prompts you to confirm the deletion.
4. Click OK.

Deleting the Metadata Server

Before you can delete the metadata server of a shared file system, you must perform the following tasks:

1. Unmount the clients and potential metadata servers, as described in [“Unmounting a Client or Potential Metadata Server” on page 104](#).
2. Unmount the metadata server, as described in [“Unmounting the Metadata Server” on page 105](#).
3. Delete the clients and potential metadata servers, as described in [“Deleting a Client or Potential Metadata Server” on page 99](#).

To delete the metadata server:

1. From the Servers menu or the Servers page, choose the name of the server that is acting as the metadata server for the shared file system.

The File Systems Summary page is displayed.

2. Scroll through the table and click the shared file system name.

The File System Details page is displayed.

3. Select the radio button next to the metadata server that you want to delete, and click Delete.

A message box prompts you to confirm the deletion.

4. Click OK.

Mounting and Unmounting Shared File Systems

This section provides information about mounting shared file systems. It contains the following subsections:

- [“Mounting a Shared File System” on page 101](#)
- [“Mounting the Metadata Server” on page 102](#)
- [“Mounting a Client or Potential Metadata Server” on page 102](#)
- [“Editing Mount Options for a Shared File System” on page 103](#)
- [“Unmounting a Shared File System” on page 103](#)

- [“Unmounting a Client or Potential Metadata Server” on page 104](#)
- [“Unmounting the Metadata Server” on page 105](#)

Mounting a Shared File System

By mounting a shared file system, you attach the shared file system to the file system hierarchy at the specified mount point, which is the full path name of a directory (*/directory-name/file-system-name*).

To mount a shared file system:

1. Ensure that the `fsmgmtd` daemon is running and available on the metadata server:
 - a. On the metadata server, use the `fsmadm(1M)` `status` command to check whether the daemon is running.
 - b. If it is not, manually restart the daemon by typing the following:

```
/opt/SUNWsamfs/sbin/fsmadm restart
```
2. Return to the File System Manager interface.
3. From the Servers menu or the Servers page, choose the name of the server that is acting as the metadata server or as the potential metadata server for the shared file system.

The File Systems Summary page is displayed.

4. Scroll through the table and click the shared file system for which you want to mount the metadata server.

The File System Details page is displayed.

5. Select the radio button next to the metadata server that you want to mount, and choose Mount from the Operations menu.

The metadata server is redisplayed with the `mounted` value in its Mount State column.

6. Select the radio button next to the client or potential metadata server that you want to mount, and choose Mount from the Operations menu.

The client or potential metadata server is redisplayed with the `mounted` value in its Mount State column.

Mounting the Metadata Server

Before you can mount the metadata server of a shared file system, the `fsmgmtd` daemon must be running and available on the metadata server. For more information about this daemon, see the *Sun StorageTek QFS Installation and Upgrade Guide*.

To mount the metadata server:

1. From the Servers menu or the Servers page, choose the name of the server that is acting as the metadata server for the shared file system.

The File Systems Summary page is displayed.

2. Scroll through the table and click the shared file system for which you want to mount the metadata server.

The File System Details page is displayed.

3. Select the radio button next to the metadata server that you want to mount.

4. From the Operations menu, choose Mount.

The metadata server is redisplayed with the `mounted` value in its Mount State column.

Mounting a Client or Potential Metadata Server

Before you can mount a client or potential metadata server of a shared file system, the `fsmgmtd` service must running and available on the metadata server and on the client or potential metadata server that you are mounting. Additionally, the metadata server must be mounted.

To mount a client or potential metadata server:

1. From the Servers menu or the Servers page, choose the name of the server that is acting as the metadata server or as the potential metadata server for the shared file system.

The File Systems Summary page is displayed.

2. Scroll through the table and click the shared file system for which you want to mount a client or potential metadata server.

The File System Details page is displayed.

3. Select the radio button next to the client or potential metadata server that you want to mount.

4. From the Operations menu, choose Mount.

The client or potential metadata server is redisplayed with the `mounted` value in its Mount State column.

Editing Mount Options for a Shared File System

To edit the mount options for a shared file system, you must edit mount options on its metadata server. When you make changes to the mount options on the metadata server, the changes are propagated to all the participating hosts in the file system. This synchronizes the hosts's mount options, unless a client host is flagged as read-only.

To edit a shared file system's mount options:

1. From the Servers menu or the Servers page, choose the name of the server that is acting as the metadata server or as the potential metadata server for the shared file system.

The File Systems Summary page is displayed.

2. Scroll through the table and click the shared file system for which you want to edit mount options.

The File System Details page is displayed.

3. Select the radio button next to the metadata server on which you want to edit mount options.

4. From the Operations menu, choose Edit Mount Options.

The File System Edit Mount Options page is displayed.

5. Edit the fields that you want to change, and click Save.

Unmounting a Shared File System

You must unmount a shared file system before upgrading hardware.

Note: The main reasons to unmount a shared file system are to validate its configuration or to upgrade hardware. You cannot grow a shared file system through the browser interface.

To unmount a shared file system:

1. From the Servers menu or the Servers page, choose the name of the server that is acting as the metadata server for the shared file system.

The File Systems Summary page is displayed.

2. Scroll through the table and click the shared file system that you want to unmount.

The File System Details page is displayed.

3. For each client or potential metadata server participating in the shared file system, do the following:
 - a. Select the radio button next to the client or potential metadata server.
 - b. From the Operations menu, choose Unmount.

The client or potential metadata server is redisplayed with the `unmounted` value in the Mount State column.

4. Select the radio button next to the metadata server, and choose Unmount from the Operations menu.

The metadata server is redisplayed with the `unmounted` value in the Mount State column.

Unmounting a Client or Potential Metadata Server

You must unmount a shared file system before upgrading hardware. Before you unmount a file system, you must idle or stop any archiving operation that is in process.

To unmount a client or potential metadata server in a shared file system:

1. From the Servers menu or the Servers page, choose the name of the server that is acting as the metadata server or as the potential metadata server for the shared file system.

The File Systems Summary page is displayed.

2. Scroll through the table and click the shared file system that contains the client or potential metadata server that you want to unmount.

The File System Details page is displayed.

3. Select the radio button next to the client or potential metadata server.
4. From the Operations menu, choose Unmount.

The client or potential metadata server is redisplayed with the `unmounted` value in its Mount State column.

Unmounting the Metadata Server

You must unmount a shared file system before upgrading hardware. Before you unmount a file system, you must idle or stop any archiving operation that is in process.

Note: Before you can unmount the metadata server for a shared file system, you must unmount each client and potential metadata server that is participating in the shared file system.

To unmount a metadata server for a shared file system:

1. From the Servers menu or the Servers page, choose the name of the server that is acting as the metadata server or as the potential metadata server for the shared file system.

The File Systems Summary page is displayed.

2. Scroll through the table and click the shared file system for which you want to unmount its metadata server.

The File System Details page is displayed.

3. Select the radio button next to the metadata server that you want to unmount.

4. From the Operations menu, choose Unmount.

The metadata server is redisplayed with the unmounted value in its Mount State column.

Deleting a Shared File System

Before you delete a shared file system, you must perform the following tasks:

1. Unmount all clients and potential metadata servers, as described in [“Unmounting a Client or Potential Metadata Server” on page 104](#).
2. Unmount the metadata server, as described in [“Unmounting the Metadata Server” on page 105](#).
3. Delete all clients and potential metadata servers, as described in [“Deleting a Client or Potential Metadata Server” on page 99](#).
4. Delete the metadata server, as described in [“Deleting the Metadata Server” on page 100](#).

To delete the shared file system:

1. From the Servers menu or the Servers page, choose the name of the server on which the shared file system resides.

The File Systems Summary page is displayed.

2. Scroll through the table and select the radio button next to the shared file system that you want to delete.
3. From the Operations menu, choose Delete.
A message box prompts you to confirm the deletion.
4. Click OK.

Managing NFS Sharing for File Systems

This section provides information about sharing file systems and editing Network File System (NFS) properties. It contains the following subsections:

- [“About NFS Sharing” on page 106](#)
- [“Adding an NFS Directory” on page 107](#)
- [“Editing NFS Directory Options” on page 108](#)
- [“Changing the Status of an NFS Directory” on page 109](#)
- [“Removing an NFS Directory” on page 110](#)

About NFS Sharing

You can manage Network File System (NFS) directories in file systems.

If at least part of a stand-alone, archiving, or shared file system is shared with the NFS service, the Description column on the File Systems Summary page displays the `nfs-shared` value. For more information about the NFS service, see the NFS documentation.

For a whole file system to be NFS shared, the mount point of the file system must be a shared NFS directory. For part of a file system to be NFS shared, a subdirectory of its mount point must be a shared NFS directory.

Note: If you create a file system and you specify a shared NFS directory as the mount point, you must mount the file system at boot time in the New File System wizard. Otherwise, NFS automatically shares the directory, possibly preventing you from mounting the file system at a later time.

When a file system is NFS shared, remote servers in the network can mount the NFS directories on the file system, assuming the servers have the NFS service enabled. This is the equivalent of using the `mount_nfs` command. For more information, see the `mount_nfs(1M)` man page.

The browser interface supports the addition and removal of NFS directories to and from file systems. In addition, the browser interface supports the ability to change the options of the NFS directories and change the status of the directories to shared or unshared.

The browser interface does not support the NFS automount feature (`autofs`). For more information, see the `NFS automount(1M)` man page.

Adding an NFS Directory

You can add a Network File System (NFS) directory to a mounted local file system. If you add an NFS directory and choose to share it, other remote servers in the network can mount the NFS directory.

Before adding an NFS directory to a file system, the NFS service must be enabled on the server where the file system is located. For more information, see the NFS documentation.

To add an NFS directory:

1. From the Servers menu or the Servers page, choose the server on which the file system that you want to administer resides.
2. In the navigation tree, select NFS under File Systems & NFS.
The NFS Shared Resources page is displayed.
3. In the NFS Shared Resources table, click Add.

The Add Directory/File window is displayed.

4. Click Browse to find the directory or file that you want.

The Browse window is displayed.

5. Use the fields and buttons on the Browse window to navigate through the directories on the server and do one of the following:

- Upon finding the directory that you want, click Choose Folder.
- If you cannot find a directory that you want, type a name in the Folder Name field and click Create Folder.

The Browse window is closed and the directory name that you specified is displayed on the Add Directory window.

6. To immediately share the directory, click Save. If you do not want to share the directory, deselect Share now and click Save.

Editing NFS Directory Options

You can edit Network File System (NFS) directory options for a local file system that has the NFS service enabled. Some NFS options that you can edit include the read and write permissions of an NFS directory. Editing these permissions controls whether remote servers have read-only, read and write, or root access to the shared NFS directory.

To edit NFS directory options for a file system:

1. From the Servers menu or the Servers page, choose the server on which the file system that you want to administer resides.

2. From the navigation tree, choose NFS.

The NFS Shared Resources page is displayed.

3. Select the radio button next to the directory that you want to edit and click Edit.

The Editing NFS Options window is displayed.

4. Specify one or more of the following:

- Read only – Select to enable read-only access to this directory. In the Access List field, optionally specify which servers have read-only access by typing the server or network group names, separated by a colon. If you leave the Access List field blank, all servers have read-only access.

- Read/Write – Select to enable read and write access to this directory. In the Access List field, optionally specify which servers have read and write access by typing the server or network group names, separated by a colon. If you leave the Access List field blank, all servers have read and write access.
- Root – Select to enable executive privileges on this directory. In the Access List field, specify which servers have executive privileges by typing the server or network group names, separated by a colon.

5. Click Save.

Changing the Status of an NFS Directory

You can change the status of an NFS directory to shared or unshared. When you change the status of an NFS directory to unshared, any remote file system that has mounted the NFS directory can no longer access the files or subdirectories in the file system.

Note: Changing the status of an NFS directory to unshared does not permanently disable the sharing of the NFS directory. Instead, it changes the status of the NFS directory to inactive. To permanently disable sharing, you must remove the NFS directory. For more information, see [“Removing an NFS Directory” on page 110](#).

To change the status of an NFS directory:

1. From the Servers menu or the Servers page, choose the server on which the file system that you want to administer resides.
2. From the navigation tree, choose NFS.
The NFS Shared Resources page is displayed.
3. In the NFS Shared Resources table, select the radio button next to the directory that has the status you want to change.

4. From the Operations menu, choose Share or Unshare.

The directory is redisplayed with the newly changed status value in its Current Status column.

Removing an NFS Directory

You can remove an NFS directory from a local file system. Removing an NFS directory does not delete the directory from the server on which it resides. Instead, the NFS directory is permanently unshared and is removed from the NFS `/etc/dfs/dtstab` configuration file.

To remove an NFS directory from a file system:

1. From the Servers menu or the Servers page, choose the server on which the file system that you want to administer resides.
2. From the navigation tree, choose NFS.

The NFS Shared Resources page is displayed.

3. Select the radio button next to the directory that you want to remove, and click Remove.

A message box prompts you to confirm the removal.

4. Click OK.

Managing Archiving

This chapter provides information about managing archiving. It contains the following sections:

- [“About Archive Management” on page 111](#)
- [“About Archiving, Releasing, and Staging” on page 113](#)
- [“Controlling Archiving for Archiving File Systems” on page 120](#)
- [“Managing Archive Policies” on page 122](#)
- [“Managing VSN Pools for Archiving” on page 142](#)
- [“Managing Disk VSNs for Archiving” on page 145](#)
- [“Setting Up General Archiving” on page 148](#)
- [“Configuring the Recycler” on page 152](#)
- [“Managing Archiving and Staging Activity” on page 156](#)

About Archive Management

If you are managing a server that has the Sun StorageTek SAM software (SUNWsamfsr and SUNWsamfsu packages) installed locally, you can create archiving Sun StorageTek SAM and Sun StorageTek QFS file systems on the server. Otherwise, the Archive Management options are not displayed in the browser interface and you cannot perform archiving options.

If the current server has the `SUNWsamfsr` and `SUNWsamfsu` packages installed locally, you can perform the following procedures:

- Create a stand-alone, non-archiving Sun StorageTek QFS file system and configure it to be archiving.
- Create a Sun StorageTek SAM file system that is automatically configured to be archiving.
- Create an archive policy, which is a collection of file system directives, copy directives, copy parameters, and volume associations that determines how groups of files can be archived.
- Perform disk archiving by creating a policy and selecting Disk as your media type in one of the copies in the policy. You can specify several disk volumes, or you can choose a pool of disk volumes to be used for archiving.
- Add file match criteria and policy copies to a policy and remove them from a policy.
- Apply file match criteria in a policy to a file system. This defines how and when groups of files in the file system are archived.
- Remove file match criteria in a policy from a file system, which stops the archiving of files in the file system, but does not delete the file match criteria from the configuration (`archiver.cmd` file).
- Delete a policy from the configuration (`archiver.cmd` file).
- Configure continuous archiving on a global basis by choosing No Scan on the General Archiving Setup page.
- Configure continuous archiving on a file system basis by choosing No Scan on the File System Archive Policies page.
- Configure metadata archiving by editing the policy settings that are defined in the default policy for a file system.
- Create a no-archive policy, which enables groups of files to never be archived.
- Create, edit, and delete pools of disk or media volumes.
- Configure global archiving directives.

About Archiving, Releasing, and Staging

This section provides information about archiving, staging, and releasing. It contains the following subsections:

- [“About the Sun StorageTek SAM Software” on page 113](#)
- [“About Archiving” on page 114](#)
- [“Planning for Archiving” on page 115](#)
- [“About Continuous Archiving” on page 117](#)
- [“About Associative Archiving” on page 118](#)
- [“About the Releaser” on page 119](#)
- [“About the Stager” on page 119](#)

About the Sun StorageTek SAM Software

The Sun StorageTek SAM software enables you to archive copies of files to and retrieve copies from remote media or partitions of another file system.

The following table describes the components of the Sun StorageTek SAM software.

TABLE 7-1 Components of Sun StorageTek SAM Software

Component	Description
Archiver	Software that automatically copies online disk cache files to archive media, such as to a disk or a tape.
Releaser	Software that automatically maintains the online disk cache of an archiving file system by freeing disk blocks occupied by eligible archived files. The software determines when to release disk blocks by using the High Water Mark and Low Water Mark threshold values that you define when you create an archiving file system. You can edit these values on the Edit Mount Options page for the file system.
Stager	Software that restores files to online disk cache. When a user or process requests file data that has been released from disk cache, the stager automatically copies the file data back to the online disk cache.
Recycler	Software that clears archive volumes of expired copies and makes volumes available for reuse.

When you create a file system through the browser interface and you select Archive This File System, the File System Manager software configures the file system to use the Sun StorageTek SAM software that is installed locally on the current server.

About Archiving

Archiving is the process of copying a file from a file system to a volume that is located on a removable media cartridge or on a disk partition of another file system. The archiver, a component of the Sun StorageTek SAM software, is software that automatically copies online disk cache files to archive media, unless you use a no-archive policy to configure it to do otherwise.

Before the archive process is initiated, the archiver searches for archiving directives that are defined in archive policies. An archive policy is a collection of file system directives, copy directives, copy parameters, and volume associations that determines how groups of files can be archived. Several archive policies can be associated with an archiving file system. For more information about policies, see [“About Archive Policies” on page 123](#).

Based on the directives defined in the default policy for a file system, the archiver automatically creates one archive copy of all files and sends the copy to archive media. You can configure the archiver to create up to four archive copies on a variety of archive media by modifying the default policy for the file system or by creating a custom policy and applying it to the file system.

Before a file is considered a candidate for archiving, the data in the file must be modified. A file is not archived if it is only accessed. A file is selected for archiving based on its archive age, which is the amount of time since the file’s last modification. The archive age can be defined for each archive copy in an archive policy.

If a large file is segmented into smaller pieces, each segment is treated as a file, and each segment is archived separately.



Planning for Archiving

The following table provides general planning recommendations that can help improve archiving performance.

TABLE 7-2 Archiving Recommendations

Recommendation	Reason
Save archive log files.	The archive log files provide information that is essential to recovering data. Store the log files in a safe place in the case of a disaster or in case the Sun StorageTek SAM software is unavailable. You can configure global archive log files on the General Setup page. You can also configure individual log files for each archiving file system on the File System Archive Policies page for the file system.
Specify volume ranges in archive copies, rather than specific volume names.	Let the software place files on many different volumes. Volume ranges enable the system to run continuously. Typing specific volume names in archive copies can rapidly fill a volume, causing undue workflow problems as you remove a piece of media and replace it with another. For information about creating archive copies, see “Adding a Copy to a Policy” on page 141 .
Base archive intervals on how often files are created and modified, and whether you want to save all modification copies.	The archive interval is the time between file system scans. A very short archive interval keeps the archiver scanning almost continuously.

TABLE 7-2 Archiving Recommendations (*Continued*)

Recommendation	Reason
Consider the number of file systems you are using.	Multiple small archiving file systems result in better archiver performance than a single large archiving file system. The archiver uses a separate process for each file system. Multiple file systems can be scanned in considerably less time than a single large file system. For information about the number of file systems to create based on your environment, see “Archiving Configuration Guidelines” on page 117 .
Use directory structures to organize files in an archiving file system as you would a UNIX file system.	For performance considerations, do not place more than 10,000 files in a directory. For information about other archiving configuration guidelines, see “Archiving Configuration Guidelines” on page 117 .
Always make a minimum of two file copies on two separate volumes.	Placing data on a single media type puts your data at risk if physical problems with the media occur. Do not rely on a single archive copy. For information about creating archive copies, see “Adding a Copy to a Policy” on page 141 .
Schedule recovery points to be created on a regular basis.	You can use the information stored in a recovery point to recover an archiving file system in the event of a disaster. Make sure that files are archived before the recovery point is created. For more information, see “About Protecting Archiving File System Data” on page 80 .

If you are managing a server that has the Sun StorageTek SAM software (SUNWsamfsr and SUNWsamfsu packages) installed locally and you are configuring file systems on the server to be archiving, you should have at least one tape library associated with the current server.

The following table describes archiving configuration guidelines, on a per-tape-library basis, that can prevent you from over-extending your environment.

TABLE 7-3 Archiving Configuration Guidelines

Number of Tape Drives	Number of Custom Policies to Create	Maximum Number of File Systems to Create	Maximum Number of Files in Each File System	Library Recycler Values
2–3	1	4	6 million	• Minimum Gain - 90% • VSN Limit (#) - 2 • High Water Mark - 50% • Size Limit - 30 Gbytes
4–5	1	6	6 million	• Minimum Gain - 90% • VSN Limit (#) - 3 • High Water Mark - 50% • Size Limit - 40 Gbytes
6–7	2	10	8 million	• Minimum Gain - 90% • VSN Limit (#) - 5 • High Water Mark - 50% • Size Limit - 50 Gbytes
8–10	4	10	10 million	• Minimum Gain - 90% • VSN Limit (#) - 8 • High Water Mark - 50% • Size Limit - 70 Gbytes

About Continuous Archiving

Continuous archiving eliminates the need for the archiver to periodically scan an archiving file system. With continuous archiving, the archiving file system notifies the archiver when files have changed and the archiver determines how and when to schedule archiving based on the Start Age, Start Count, and Start Size values defined in the copies of the policy associated with those files.

To enable continuous archiving for all archiving file systems on the current server, choose the No Scan scan method on the General Setup page. To enable continuous archiving for an individual archiving file system, choose the No Scan method on the File System Archive Policies page for that file system. Any setting that you make on the File System Archive Policies page overrides the equivalent setting on the General Setup page.

To specify Start Age, Start Count, and Start Size values for all archiving file systems on the current server, edit the values in the copies of the configurable defaults (allsets) archive policy. To override the settings in the configurable defaults policy, specify Start Age, Start Count, and Start Size values in the individual copies of any other archive policy. You do this on the Advanced Copy Options page of an archive policy.

The following examples describe how the policy copy values affect continuous archiving:

- **Start Age** – If it takes one hour to create files that must be archived together, you can create a policy copy and can define the Start Age value in the copy to be 1 hour. This ensures that all necessary files are created before archiving begins.
- **Start Size** – If you want the archiver to wait until a specific amount of data is ready to be archived, you can create a policy copy and specify the Start Size value to be 150 Gbytes, for example. This directs the archiver to wait until 150 Gbytes of data are ready to be archived.
- **Start Count** – If you know that 3000 files will be generated for archival, you can create a policy copy and specify the Start Count value to be 3000. This ensures that all 3000 files will be archived together.

About Associative Archiving

Associative archiving is useful if you want to archive an entire directory to one volume and if all the contents in that directory will fit on a single archive volume. You can enable associative archiving by selecting the Force files in a directory to be archived together check box in an archive policy copy.

When files are archived, they are grouped together in one or more archive files to efficiently pack the volume. Subsequently, when accessing files from the same directory, you might experience delays as the stage process repositions through a volume to read the next file. To alleviate delays, you can choose to archive files from the same directory paths contiguously within an archive file. The process of associative archiving overrides the space efficiency algorithm to keep files from the same directory together. You can specify that these files are to be archived contiguously within a copy of an archive policy.

Associative archiving is useful when the file content does not change and you always want to access the group of files together at the same time. For example, you might use associative archiving at a hospital for accessing medical images. Images associated with a single patient can be kept in a single directory, enabling a doctor to access those images together at one time.

However, because associative archiving specifies that all files from the same directory be archived on a single volume, it is possible that a group of files might not fit on any available volume. In this case, the files are not archived until more volumes are assigned to the archive policy. Another issue is that the group of files to be archived might be so large that it can never fit on a single volume. In such a case, the files are never archived.

An alternative to associative archiving is the sorting of files by path. Sorting keeps files together, but does not force the files to be archived together. You can specify sorting in the copy of an archive policy.

About the Releaser

The releaser is a component of the Sun StorageTek SAM software. After a file is archived, it becomes eligible to be released. The releaser frees primary (disk) storage that is used by the archived file's data. Two threshold values, High Water Mark and Low Water Mark, manage online disk cache free space. You can define these thresholds values when you create an archiving file system or when you edit the mount options of an archiving file system.

When online disk consumption exceeds the High Water Mark threshold, the system automatically begins releasing the disk space occupied by eligible archived files. Disk space occupied by archived file data is released until the Low Water Mark threshold is reached. Files are selected for release depending on file size and age.

Optionally, the first portion of a file can be retained on disk for speedy access and for masking staging delays. If a file has been archived in segments, portions of the file can be released individually.

About the Stager

The stager is a component of the Sun StorageTek SAM software. When a file whose data blocks have been released is accessed, the stager automatically stages the file or file segment data block back to online disk cache. The read operation tracks along directly behind the staging operation, allowing the file to be immediately available to the application before the entire file is completely staged.

The Sun StorageTek SAM software processes stage request errors automatically. If a stage error is returned, the system attempts to find the next available archive copy of the file. Stage errors that can be automatically processed include media errors, unavailability of media, and unavailability of an automated library, among others.

Controlling Archiving for Archiving File Systems

This section provides information about configuring archiver scanning, and about stopping, idling, and running archiving for archiving file systems. It contains the following subsections:

- [“Configuring Archiver Scanning for an Archiving File System” on page 120](#)
- [“Stopping or Idling Archiving for an Archiving File System” on page 121](#)
- [“Running Archiving for an Archiving File System” on page 121](#)

Configuring Archiver Scanning for an Archiving File System

If you are administering an archiving file system, you can configure the archiver scan method, scan interval, and log file location for the file system. This controls when the software scans for files on the file system to be archived.

To configure archiver scanning for a file system:

1. From the Servers menu or the Servers page, choose the name of the server that you want to configure.

The File Systems Summary page is displayed.

2. Click the radio button next to the file system for which you want to configure archiving values.

3. Click View Policies.

The File System Archive Policies page is displayed.

4. Specify the following:

- Scan Method – Choose how the archiver scans for files in the file system to be archived. For more information, see [“About Archiver Scanning Methods” on page 148](#).
- Interval – Choose how often the archiver log file is updated.
- Log File – Type the location of the archiver log file.

5. Click Save.

Stopping or Idling Archiving for an Archiving File System

When you stop archiving, all archiving immediately stops on the selected archiving file system. If you want to stop archiving at a logical point in the process, choose to idle archiving instead.

To stop archiving for a file system:

1. From the Servers menu or the Servers page, choose the name of the server on which the file system is located.
 2. In the navigation tree, select Archive Activity under Archive Administration.
 3. Select one of the following actions:
 - To immediately stop archiving, select the Stop radio button.
 - To stop archiving at a logical point in the process, select the Idle radio button.
 4. Click Submit.
-

Running Archiving for an Archiving File System

To run archiving for an archiving file system:

1. From the Servers menu or the Servers page, choose the name of the server on which the file system is located.
2. In the navigation tree, select Archive Activity under Archive Administration.
3. Select the Run Now radio button.
4. Choose the file system from the file system drop-down list.
5. Click Submit.

Managing Archive Policies

This section provides information about managing archive policies. It contains the following subsections:

- [“Administering Archive Policies” on page 122](#)
- [“Enabling Archiving Options in a Policy” on page 129](#)
- [“Administering File Match Criteria” on page 135](#)
- [“Administering Policy Copies” on page 139](#)

Administering Archive Policies

This section provides information about administering archive policies. It contains the following subsections:

- [“About Archive Policies” on page 123](#)
- [“Reference: Creating a Policy archiver.cmd Example” on page 126](#)
- [“Creating a Policy” on page 127](#)
- [“Creating a No-Archive Policy” on page 128](#)
- [“Editing a Policy” on page 128](#)
- [“Deleting a Policy” on page 129](#)

About Archive Policies

An archive policy is a collection of file system directives, copy directives, copy parameters, and volume associations that determine how groups of files can be archived. When the documentation refers to a policy, it is referring to a custom policy unless otherwise specified.

The following table describes the types of archive policies.

TABLE 7-4 Archive Policy Types

Policy Type	Description
File system default	A policy that is automatically created when you create an archiving file system. It archives all files that do not meet any file match criteria defined in any custom or no-archive policies that are applied to the archiving file system. You can edit these policies, but you cannot delete them. They are automatically deleted when you delete its archiving file system. You also cannot edit the file match criteria in these policies. There are two types of default archive policy. One type has the same name as the file system. This type of archive policy only applies to the file system with the same name, and archives metadata for that file system. The other type of default archive policy has the name “all.” It can be shared between multiple file systems by selecting All in the File System wizard when you create an archiving file system. This type of default policy does not archive the file system’s metadata.

TABLE 7-4 Archive Policy Types (*Continued*)

Policy Type	Description
Custom	A policy that you can create for the purpose of archiving a group of files according to a different set of rules and, optionally, to different archive media than those specified in the file system default policy. You define the name of the policy and one or more sets of file match criteria that determine which files can be archived. You can edit and delete a custom policy, and the file match criteria of a custom policy can be applied to more than one archiving file system. Note: When the documentation refers to a policy, it is referring to a custom policy unless otherwise specified.
No-archive	A policy that you can create for the purpose of preventing files from being archived. There can be only one no-archive policy on a server at a time, and the policy always has the name of <code>no_archive</code>. This policy does not have any copies, but you can add file match criteria to the policy to prevent more groups of files from being archived. The file match criteria in this policy can be applied to numerous archiving file systems. You can edit and delete the no-archive policy.
Configurable defaults	A policy that always exists and has the name of <code>allsets</code>. This policy enables you to set default values for advanced copy options and volume serial name (VSN) assignments. These default values are applied to all custom and file system default policies on the current server. For example, if you set a value in copy 1 of the configurable defaults policy, this value is set in copy 1 of all the custom and file system default policies on the server. In addition to having copies 1 through 4, the configurable defaults policy has an <code>allsets</code> copy. The settings in the <code>allsets</code> copy apply to all other copies on the server. However, any setting in a specific policy overrides the values in the configurable defaults policy. You cannot delete this policy.

A policy consists of the following attributes:

- A policy name
- One or more sets of file match criteria that can contain:
 - A starting directory, which is a path relative to the mount point of the file system
 - Optional search criteria based on the following options:
 - File name that uses pattern matching
 - User name or group name or both

- Minimum or maximum file size or both
- Access age
- Release or stage attributes
- Archive age for each copy
- Unarchive age and release options for each copy
- At least one file system (no maximum limit) to which the set of match criteria is applied
- From one to four archiving copies of the files from the member file systems that match the file match criteria specified in the policy

When you create an archive policy, you must specify whether each copy to be made will be made to disk or to tape.

Reference: Creating a Policy `archiver.cmd` Example

After you create an archive policy through the browser interface and you associate the policy with file systems, the system adds additional lines to the `archiver.cmd` file on the selected server. In the following example, the policy `policy1` is associated with the file systems `samfs1` and `samfs2`.

```

...
#
#       File System Directives
#
...
fs = samfs1
policy1 dir1 -minsize 10k -maxsize 10M
    1 4m 7d
    2 4m 7d
...
fs = samfs2
policy1 dir1 -minsize 10k -maxsize 10M
    1 4m 7d
    2 4m 7d
#
#       Copy Parameters Directives
#
params
policy1.1 -join path
policy1.1 -startage 301m
policy1.1 -startsize 7M
policy1.2 -bufsize 9
policy1.2 -drives 2
policy1.2 -ovflmin 2M
...
endparams
...
#
#       VSN Directives
#
vsns
...
policy1.1 1t BL794[2-4]
policy1.2 1t BL7942 -pool pool2
...
endvsns

```

Creating a Policy

To create a policy:

1. From the Servers menu or the Servers page, choose the name of the server for which you want to create a policy.

2. From the navigation tree, choose Archive Administration.

The Archive Policies Summary page is displayed.

3. Click New.

The New Archive Policy wizard is displayed.

4. Complete the steps in the New Archive Policy wizard.

Click the Help tab in the wizard for more information.

Creating a No-Archive Policy

With a no-archive policy, you create a policy that does not contain any policy copies. Files that meet the file match criteria defined in the no-archive policy are never archived. This set of file match criteria in the policy maps directly to the `no_archive` set, which you can define by using the Sun StorageTek SAM command-line interface.

Only one no-archive policy can exist on a server at a time. If one already exists, you cannot create a new one without first deleting the existing one. Alternatively, you can add additional file match criteria to the existing policy if you want to prevent the archiving of additional files.

To create a no-archive policy, follow the directions in [“Creating a Policy” on page 127](#), and choose to prevent files from being archived in the first step of the New Archive Policy wizard. Upon creation, the no-archive policy is displayed on the Archive Policies Summary page.

Editing a Policy

To edit a policy:

1. From the Servers menu or the Servers page, choose the name of the server where the policy that you want to edit is configured.

2. From the navigation tree, choose Archive Administration.

The Archive Policies Summary page is displayed.

3. Click the policy that you want to edit.

The Policy Details page is displayed.

4. Do one of the following (not all of these options are available for every policy):

- To add a set of file match criteria to the policy, click Add in the Archive Criteria table, and then follow the steps in the New Criteria wizard.

- To add a policy copy to the policy, click Add in the Copy Information table, and follow the steps in the New Copy wizard.
- To edit an existing set of file match criteria, select the radio button next to the file match criteria that you want, click Edit, and make your changes.
- To edit an existing policy copy, click Edit Advance Options in the Copy Information table and edit the fields that you want to change.

Deleting a Policy

Deleting a policy deletes the policy from the server's configuration (`archiver.cmd` file). This is different from removing the relationship between a file match criteria set and a file system.

To delete a policy:

1. From the Servers menu or the Servers page, choose the name of the server that has the policy you want to delete.
2. From the navigation tree, choose Archive Administration.
The Archive Policies Summary page is displayed.
3. Select the radio button next to the policy that you want to delete, and click Delete.
A message box prompts you to confirm the deletion.
4. Click OK.

Enabling Archiving Options in a Policy

This section provides information about enabling archiving options in a policy. It contains the following subsections:

- [“Enabling Disk Archiving” on page 130](#)
- [“Enabling Archiving to Tape” on page 131](#)
- [“Enabling Associative Archiving” on page 132](#)
- [“Enabling Continuous Archiving” on page 133](#)
- [“Enabling Archiving to Pools of Disk or Tape VSNs” on page 134](#)

Enabling Disk Archiving

When you create an archive policy, you can specify archiving to disk-based VSNs for one or more copies of that archive policy.

The File System Manager software can only add information to the `diskvols.conf` file. It does not edit previously written data, even if you attempt to edit the policy through the browser interface.

You can enable disk archiving under the circumstances described in the following sections:

- [“For a New Policy” on page 130](#)
- [“For an Existing Copy in an Existing Policy” on page 130](#)
- [“For a New Copy in an Existing Policy” on page 131](#)

For a New Policy

To enable disk archiving for a new policy:

1. If you want to archive to a new disk-based VSN, follow the instructions in [“Creating a Disk VSN” on page 146](#) or [“Creating an Sun StorageTek 5800 VSN” on page 147](#).
2. Follow the instructions in [“Creating a Policy” on page 127](#).
3. When you configure media parameters for a copy in the New Archive Policy wizard, choose Disk or STK 5800 from the Media Type menu.
4. Continue through the wizard and click Finish.

For an Existing Copy in an Existing Policy

To enable disk archiving for an existing copy in an existing policy:

1. If you want to archive to a new disk-based volume, follow the instructions in [“Creating a Disk VSN” on page 146](#) or [“Creating an Sun StorageTek 5800 VSN” on page 147](#).
2. On the Policy Details page, select the radio button next to the copy for which you want to enable disk archiving. Click Modify VSN Assignments.
The Policy Copy Number VSN/VSN Pools page is displayed.
3. Choose Disk or STK 5800 from the Media Type menu.
4. Specify the disk volume serial names (VSNs) that you want to use for archiving.

5. Click Save.

For a New Copy in an Existing Policy

To enable disk archiving for a new copy in an existing policy, perform the following:

1. If you want to archive to a new disk-based VSN, follow the instructions in [“Creating a Disk VSN” on page 146](#) or [“Creating an Sun StorageTek 5800 VSN” on page 147](#).
2. Follow the instructions in [“Adding a Copy to a Policy” on page 141](#).
3. When you configure media parameters for the copy in the New Copy wizard, choose Disk or STK 5800 from the Media Type menu.
4. Continue through the wizard and click Finish.

Enabling Archiving to Tape

When you create an archive policy, you can specify archiving to tape for one or more copies of that archive policy.

You can enable archiving to tape under the circumstances described in the following sections:

- [“For a New Policy” on page 131](#)
- [“For an Existing Copy in an Existing Policy” on page 131](#)
- [“For a New Copy in an Existing Policy” on page 132](#)

For a New Policy

To enable archiving to tape for a new policy:

1. Follow the instructions in [“Creating a Policy” on page 127](#).
2. When you configure media parameters for a copy in the New Archive Policy wizard, choose one of the tapes listed in the Media Type menu.
3. Continue through the wizard and click Finish.

For an Existing Copy in an Existing Policy

To enable archiving to tape for an existing copy in an existing policy:

1. On the Policy Details page, select the radio button next to the copy for which you want to enable disk archiving. Click VSN Assignments.

The Policy Copy Number VSN/VSN Pools page is displayed.

2. Choose one of the tapes listed in the Media Type menu.
3. Specify the tape VSNs that you want to use for archiving.
4. Click Save.

For a New Copy in an Existing Policy

To enable archiving to tape for a new copy in an existing policy:

1. Follow the instructions in [“Adding a Copy to a Policy” on page 141](#).
2. When you configure media parameters for the copy, choose one of the tapes listed in the Media Type menu.
3. Continue through the wizard and click Finish.

Enabling Associative Archiving

After you have created a policy, you can specify associative archiving for one or more copies of that policy.

To enable associative archiving for a policy:

1. From the Servers menu or the Servers page, choose the name of the server where the policy that you want to edit is configured.
2. From the navigation tree, choose Archive Administration.
The Archive Policies Summary page is displayed.
3. Click the policy that you want to edit.
The Policy Details Page is displayed.
4. In the Copy Information table, click the copy that you want to edit.
The Copy Details page is displayed.
5. In the Archive Organization section, select Force files in a directory to be archived together.
6. Click Save.

Enabling Continuous Archiving

You can enable continuous archiving for all archiving file systems on the current server or for a specific file system.

To enable continuous archiving:

1. Do either of the following:
 - To enable continuous archiving for all archiving file systems on the current server, navigate to the General Setup page under Archive Administration in the navigation tree, and choose No Scan as the scan method.
 - To enable continuous archiving for a specific file system, navigate to the File System Archive Policies page for that file system and choose No Scan as the scan method.
2. Set the Start Age, Start Count, or Start Size values in the appropriate policy:
 - a. From the Servers menu or the Servers page, choose the name of the server that you want to configure.
 - b. From the navigation tree, choose Archive Administration.

The Archive Policies Summary page is displayed.
 - c. Do either of the following:
 - To set values for all archiving file systems on the current server, click the configurable defaults policy.
 - To set values for a specific file system, click the file system default policy for that file system.

The Policy Details page is displayed.
 - d. In the Copy Information table, click the copy for which you want to enable continuous archiving.

The Policy Advanced Copy Options page is displayed.
 - e. Edit the values in one or more of the following fields:
 - Start Age
 - Start Count
 - Start Size
 - f. Click Save.

Archiving is scheduled according to the specified values.

Enabling Archiving to Pools of Disk or Tape VSNs

You can enable archiving to pools of disk or tape volume serial names (VSNs).

You can enable archiving to pools of disk or tape VSNs under the circumstances described in the following sections:

- [“For a New Policy” on page 134](#)
- [“For an Existing Copy in an Existing Policy” on page 134](#)
- [“For a New Copy in an Existing Policy” on page 134](#)

For a New Policy

To enable archiving to pools of disk or tape volumes in a new policy, follow the directions in [“Creating a Policy” on page 127](#), and do one of the following in the New Archive Policy wizard:

- On the Copy x - Media Parameters step, choose a VSN pool from the VSN Pool menu.
- On the Copy x - Media Parameters step, specify a range of VSNs or type a comma-separated list of VSNs in the Specify VSNs section.

For an Existing Copy in an Existing Policy

To enable archiving to pools of disk or tape volumes in an existing copy of an existing policy, follow steps 1 through 5 in [“Modifying VSN Assignments in a Policy Copy” on page 140](#) and then do one of the following:

- Choose a VSN pool from the VSN Pool menu, and click Save.
- Specify a range of VSNs, or type a comma-separated list of VSNs in the Specify VSNs section, and click Save.

For a New Copy in an Existing Policy

To enable archiving to pools of disk or tape volumes in a new copy of an existing policy, follow the directions in [“Adding a Copy to a Policy” on page 141](#), and do one of the following in the New Copy wizard:

- On the Copy Media Parameters step, choose a VSN pool from the From the VSN Pool menu.
- On the Copy Media Parameters step, specify a range of VSNs or type a comma-separated list of VSNs in the Specify VSNs section.

Administering File Match Criteria

Policies consist of policy copies and sets of file match criteria. The sets of file match criteria define archiving directives for the files that meet the criteria.

This section provides information about file match criteria. It contains the following subsections:

- [“About File Match Criteria” on page 135](#)
- [“Editing File Match Criteria” on page 136](#)
- [“Adding a File System to a Set of File Match Criteria” on page 136](#)
- [“Removing a File System From a Set of File Match Criteria” on page 137](#)
- [“Planning File Match Criteria” on page 138](#)
- [“Adding File Match Criteria to a Policy” on page 138](#)
- [“Removing File Match Criteria From a Policy” on page 139](#)

About File Match Criteria

When you create a policy, you define a single set of file match criteria that identifies groups of files to be archived. After you create the policy, you can add additional sets of file match criteria to the policy to further identify groups of files to be archived.

File match criteria can be based on file size, ownership, group, and directory location. Once a set of criteria is defined, you can associate it with archiving file systems. When files on the associated archiving file systems meet the criteria defined in a policy, the files are copied to the volumes that are defined in the copies of the policy.

You can add a set of file match criteria to or remove it from an existing policy. However, a policy must have at least one set of file match criteria at any time. If a policy has only one set of file match criteria, the criteria cannot be removed from the policy. You must delete the policy itself.

You can add and remove the relationship between an archiving file system and a set of file match criteria. However, a set of file match criteria cannot exist without at least one member file system. You cannot remove a relationship between an archiving file system and a set of file match criteria if the file system is the sole relationship with the set of criteria. In this case, you must delete the set of file match criteria itself.

Additionally, an archiving file system cannot be deleted if it is the only archiving file system associated with an existing set of file match criteria. In this case, you must remove the set of file match criteria from the policy and then you can delete the archiving file system.

Editing File Match Criteria

To edit file match criteria:

1. From the Servers menu or the Servers page, choose the name of the server with the policy that you want to edit.
2. From the navigation tree, choose Archive Administration.
The Archive Policies Summary page is displayed.
3. Click the policy that contains the file match criteria set that you want to edit.
The Policy Details Page is displayed.
4. Select the radio button next to the file match criteria set that you want and click Edit.
The Policy Criteria Details page is displayed.
5. Edit the fields that you want to change, and click Save.

Adding a File System to a Set of File Match Criteria

Adding an archiving file system to a set of file match criteria enables the files in that file system to be archived according to the archiving directives defined in the set. All files that match the defined file match criteria will be archived.

By default, when you add an archiving file system to a set of file match criteria, the set of file match criteria is placed at the bottom of the list of criteria that are already applied to the file system. For information about reordering file match criteria for an archiving file system, see [“Reordering File Match Criteria for an Archiving File System” on page 78](#).

To add an archiving file system to a set of file match criteria:

1. From the Servers menu or the Servers page, choose the name of the server on which the set of file match criteria is defined.
2. From the navigation tree, choose Archive Administration.

The Archive Policies Summary page is displayed.

3. Click the policy that contains the file match criteria for which you want to add a file system.

The Policy Details page is displayed.

4. Click the radio button next to the file match criteria that you want and click Edit.

The Policy Criteria Details page is displayed.

5. Scroll down to the File Systems Using Criteria table and click Add.

The Apply Criteria to File Systems window is displayed.

6. Select the file system that you want to add to the file match criteria, and click Save.

Removing a File System From a Set of File Match Criteria

Removing an archiving file system from a set of file match criteria does not delete the archiving file system or the file match criteria set. Instead, it removes the relationship between the file system and the file match criteria set. Once this relationship is removed, files in the file system are no longer archived according to the archiving directives defined in the set of file match criteria.

To remove an archiving file system from a set of file match criteria:

1. From the Servers menu or the Servers page, choose the name of the server on which the file system is located.

2. From the navigation tree, choose Archive Administration.

The Archive Policies Summary page is displayed.

3. Click the policy that contains the file match criteria that you want to edit.

The Policy Details page is displayed.

4. Click the radio button next to the set of file match criteria that you want, and click Edit.

The Policy Criteria Details page is displayed.

5. In the File Systems Using Criteria table, select the radio button next to the file system that you want to remove, and click Remove.

A message box prompts you to confirm the removal.

6. Click OK.

Planning File Match Criteria

The New Criteria wizard guides you through the process of adding file match criteria to an existing policy. See [“Adding File Match Criteria to a Policy” on page 138](#) for instructions on accessing the New Criteria wizard.

The wizard requires that you enter the following information, which you should have available before you begin the process of adding file match criteria:

- The name of the directory, relative to the mount point of the file system, to which this set of file match criteria should be applied. All files and subdirectories in this directory will be archived according to the directives defined in this set of file match criteria.
- A regular expression that identifies a file as a member of this archive policy. For example, if you want all .gif files to be archived according to the directives in this archive policy, .gif would be the value of the regular expression.
- The user name of the owner of the files that you want archived.
- The group to which the owner of the files belongs.
- The smallest size that an archive file can be in order to be archived.
- The largest size that an archive file can be in order to be archived.
- The length of time during which a file must be accessed for it to meet the file match criteria.
- Whether you want to stage when files are accessed, and if so, whether you want all files in the policy to be staged, to stage only when files are accessed directly from archive media, or to reset all files to the default staging configuration.
- When you want to release a file from disk cache. You can choose not to release the file, to free up online disk cache after the first archive copy is made, or to release a file after all archive copies are made.
- When a file should be unarchived, based on the last time the file was modified or the last time the file was accessed.

Adding File Match Criteria to a Policy

To add a set of file match criteria to an existing policy:

1. From the Servers menu or the Servers page, choose the name of the server with the policy that you want to edit.
2. From the navigation tree, choose Archive Administration.
The Archive Policies Summary page is displayed.
3. Click the policy that you want to edit.

The Policy Details Page is displayed.

4. In the Archive Criteria table, click Add.

The New Criteria wizard is displayed.

5. Complete the steps in the New Criteria wizard.

Click the Help tab in the wizard for more information.

Removing File Match Criteria From a Policy

Removing a set of file match criteria removes the criteria from the policy. This is different from removing the file match criteria association with a file system.

A policy must have at least one set of file match criteria at any time. You cannot delete a set of file match criteria if it is the only remaining set in a policy. Instead, you must delete the policy itself.

To remove a set of file match criteria from a policy:

1. From the Servers menu or the Servers page, choose the name of the server with the policy that you want to edit.

2. From the navigation tree, choose Archive Administration.

The Archive Policies Summary page is displayed.

3. Click the policy that you want to edit.

The Policy Details Page is displayed.

4. In the Archive Criteria table, select the radio button next to the set of file match criteria that you want to remove, and click Remove.

A message box prompts you to confirm the removal.

5. Click OK.

Administering Policy Copies

This section provides information about administering policy copies. It contains the following subsections:

- [“About Policy Copies” on page 140](#)
- [“Editing a Policy Copy” on page 140](#)
- [“Modifying VSN Assignments in a Policy Copy” on page 140](#)

- [“Adding a Copy to a Policy” on page 141](#)
- [“Removing a Copy From a Policy” on page 141](#)

About Policy Copies

A policy must have between one and four archiving copies. When a file meets the file match criteria defined in a policy, the file is archived to the media defined in the policy copy.

The values in the Start Age, Start Count, and Start Size fields in a policy copy are very important in determining the way in which the software copies files to archive media. For more information, see [“About Continuous Archiving” on page 117](#).

Editing a Policy Copy

To edit an existing copy in a policy:

1. From the Servers menu or the Servers page, choose the name of the server with the policy that you want to edit.
2. From the navigation tree, choose Archive Administration.
The Archive Policies Summary page is displayed.
3. Click the policy that you want to edit.
The Policy Details Page is displayed.
4. In the Copy Information table, click the copy that you want to edit.
The Policy Copy Details page for that copy is displayed.
5. Edit the fields that you want to change, and click Save.

Modifying VSN Assignments in a Policy Copy

You can change the volume serial names (VSNs) that are assigned to a policy copy for archiving.

To modify the VSN assignments in a policy copy:

1. From the Servers menu or the Servers page, choose the name of the server with the policy that you want to edit.
2. From the navigation tree, choose Archive Administration.

The Archive Policies Summary page is displayed.

3. Click the policy that you want to edit.

The Policy Details Page is displayed.

4. In the Copy Information table, select the radio button next to the policy copy that you want to edit, and click Modify VSN Assignments.

The Policy Copy VSNs page is displayed.

5. Edit the fields that you want to change, and click Save.

Adding a Copy to a Policy

To add a new copy to an existing policy:

1. From the Servers menu or the Servers page, choose the name of the server with the policy that you want to edit.

2. From the navigation tree, choose Archive Administration.

The Archive Policies Summary page is displayed.

3. Click the policy that you want to edit.

The Policy Details Page is displayed.

4. In the Copy Information table, click Add.

The New Copy wizard is displayed.

5. Complete the steps in the New Copy wizard.

Click the Help tab in the wizard for more information.

Removing a Copy From a Policy

Deleting a copy removes the copy from the policy.

A policy must have at least one copy. You cannot delete a copy if it is the only remaining copy in a policy. Instead, you must delete the policy itself.

To remove a copy from a policy:

1. From the Servers menu or the Servers page, choose the name of the server with the policy that you want to edit.

2. From the navigation tree, choose Archive Administration.

The Archive Policies Summary page is displayed.

3. Click the policy that you want to edit.

The Policy Details Page is displayed.

4. In the Copy Information table, select the radio button next to the copy that you want to remove, and click Remove.

A message box prompts you to confirm the removal.

5. Click OK.

Managing VSN Pools for Archiving

This section provides information about managing volume serial name (VSN) pools. It contains the following subsections:

- [“About VSN Pools” on page 142](#)
- [“Creating a VSN Pool” on page 143](#)
- [“Editing a VSN Pool” on page 143](#)
- [“Deleting a VSN Pool” on page 144](#)

About VSN Pools

A volume serial name (VSN) pool specifies a collection of individual VSNs, or volumes. This collection is used as archive media for an archive policy. As such, VSN pools provide a useful method for assigning VSNs and reserving VSNs to archive policies.

You can use VSN pools to define separate groups of VSNs for use by departments within an organization, users within a group, data types, and other groupings. The pool is assigned a name, media type, and a set of VSNs. A scratch pool is a set of volumes used when specific volumes in a VSN association are exhausted or when another VSN pool is exhausted.

If a VSN is reserved, it is no longer available to the pool in which it originated. Therefore, the number of VSNs within a named pool changes as VSNs are used.

A VSN pool definition requires at least three fields: the pool name, the media type, and at least one VSN.

Creating a VSN Pool

To create a volume serial name (VSN) pool:

1. From the Servers menu or the Servers page, choose the name of the server on which you want to create a VSN pool.
2. From the navigation tree, choose VSN Pools under Storage Administration.
The VSN Pool Summary Page is displayed.
3. Click New.
The New VSN Pool window is displayed.
4. Enter a name for the new pool.
5. Select a media type from the Media Type menu.
6. Select one of the following VSN options, and enter the information for that option.
 - Start – To specify a single range of VSNs, type the name of a VSN in the Start field and type the name of a VSN in the End field.
 - Range of VSNs – To specify more than one range of VSNs, type a range of VSNs separated by a comma or type individual VSN names separated by a comma. Use the following syntax:
`EW4276-EW4288,KQ7086,FV8434`
7. Click Save.

Editing a VSN Pool

To edit a volume serial name (VSN) pool:

1. From the Servers menu or the Servers page, choose the name of the server with the VSN pool that you want to edit.

2. From the navigation tree, choose VSN Pools under Storage Administration.
The VSN Pool Summary Page is displayed.
 3. Select the radio button next to the VSN pool that you want to edit.
 4. Click Edit.
The the Edit VSN Pool window is displayed.
 5. (Optional) From the Media Type menu, choose a different media type.
 6. (Optional) In the Specify VSNs section, choose one of the following:
 - Start – To specify a new range of VSNs, type a new name of a VSN in the Start field and type a new name of a VSN in the End field.
 - Range of VSNs – To specify more than one new range of VSNs, type a new range of VSNs separated by a comma or type new individual VSN names separated by a comma. Use the following syntax:
`EW4276-EW4288,KQ7086,FV8434`
 7. Click Save.
-

Deleting a VSN Pool

A volume serial name (VSN) pool can be deleted only if its member VSNs are not being used by a policy. If a VSN pool has a member that is being used by a policy, the Delete button is unavailable.

To delete a VSN pool:

1. From the Servers menu or the Servers page, choose the name of the server with the VSN pool that you want to delete.
2. From the navigation tree, choose VSN Pools under Storage Administration.
The VSN Pool Summary Page is displayed.
3. Select the radio button next to the VSN pool that you want to delete, and click Delete.
A message box prompts you to confirm the deletion.
4. Click OK.

Managing Disk VSNs for Archiving

This section provides information about managing disk volume serial names (VSNs) for archiving. It contains the following subsections:

- [“About Archiving to Disk-Based VSNs” on page 145](#)
- [“Creating a Disk VSN” on page 146](#)
- [“Creating an Sun StorageTek 5800 VSN” on page 147](#)
- [“Editing the Media Attributes of a Disk VSN” on page 147](#)

About Archiving to Disk-Based VSNs

In an archive policy, you can choose to send archive copies of your files to disk-based volume serial names (VSNs). In Sun StorageTek SAM version 4, update 6, there are two types of disk-based VSNs. The first type, called simply a disk VSN, allows archiving to a directory in a mounted file system that has been assigned a volume name. These are referred to by the media type “Disk” in the archive pages of the File System Manager. The second type, called an STK5800 VSN, allows archiving to the Sun StorageTek 5800. These are referred to by the media type “STK5800.”

Disk VSNs can be located on a file system on the current server or on a different server. The file system on which the disk VSNs are located can be any UNIX file system. However, if the file system is located on a server other than the current server, there must be at least one archiving file system configured on that other server.

There are no restrictions on the location of disk VSNs on a file system, but disk VSNs should reside on a disk other than where the original files that you are archiving reside. Also, commonly an entire file system is used as a single disk VSN. It is then clear that all of the remaining space on the volume can be consumed for the purpose of storing archive copies which makes planning simpler.

STK5800 VSNs allow archiving to a Sun StorageTek 5800. The Sun StorageTek 5800 is an online appliance with clustered disk-based storage nodes.

When configuring archiving, you should make more than one copy in a policy and write to more than one type of archive media. Also if you are configuring archiving to disk VSNs and tape VSNs, you should configure copy 1 to be disk VSNs and the higher numbered copies to use tape VSNs. This will help minimize tape loading because the Sun StorageTek SAM software will preferentially stage from copy 1.

The Disk-Based VSNs tab is displayed under the Archive Management tab in the application. From the Disk-Based VSNs tab, you can create disk VSNs and edit their media attributes and you can create STK5800 VSNs.

Creating a Disk VSN

To create a disk volume serial name (VSN):

1. From the Servers menu or the Servers page, choose the server for which you want to create a disk VSN.
2. From the navigation tree, choose Disk-Based VSNs under Storage Administration.
The Disk-Based VSNs page is displayed.

3. Click New.

The New Disk-Based VSN window is displayed.

4. Select the media type "Disk."
5. Specify the following:
 - Name – Type the name of the new disk VSN. This can be from one to 31 characters from the following categories:
Uppercase letters from A to Z
Numerals from 0 (zero) to 9
The following special characters: ! " % & ' () * + , . / : ; < > = ? _
 - Host – Choose the server on which the file system that will contain the disk VSN.
 - Path – Type the full path to the directory in which you want to store archive copies. If you do not know the path, click Browse to find it.
6. Click Submit.

Creating an Sun StorageTek 5800 VSN

To create a STK5800 volume (VSN):

1. From the Servers menu or the Servers page, choose the server for which you want to create an STK5800 VSN.
2. From the navigation tree, choose Disk-Based VSNs under Storage Administration.

The Disk-Based VSNs page is displayed.

3. Click New.

The New Disk-Based VSN window is displayed.

4. Select the media type STK5800 from the media type drop-down
5. Specify the following:
 - Name – Type the name of the new STK5800 VSN. This can be from one to 31 characters from the following categories:
 - Uppercase letters from A to Z
 - Numerals from 0 (zero) to 9
 - Note:** If you are running multiple connections to a STK 5800 disk archive, the name you choose here must be unique across all Sun StorageTek SAM environments. One way to achieve this would be to use the server you are currently managing as part of the VSN name.
 - IP address – Enter the IP address and port number for the Sun StorageTek STK5800.

Editing the Media Attributes of a Disk VSN

To edit the media attributes for a disk volume serial name (VSN):

1. From the Servers menu or the Servers page, choose the server for which you want to edit a disk VSN.
2. From the navigation tree, choose Disk VSNs under Storage Administration.

The Disk VSNs page is displayed.

3. Select the disk VSN that you want to edit and click Edit Media Flags.
The Edit Disk Media Flags window is displayed.
4. Select the check boxes next to the attributes that you want for the media on which a disk volume is located, and click Submit.

Setting Up General Archiving

This section provides information about setting up general archiving. It contains the following subsections:

- [“About Archiver Scanning Methods” on page 148](#)
- [“Configuring Global Archiving Values” on page 149](#)
- [“Configuring Global Staging Values” on page 150](#)
- [“Configuring Global Releasing Values” on page 150](#)
- [“Configuring Global Media Type Parameters” on page 151](#)

About Archiver Scanning Methods

If you are managing a server that has the Sun StorageTek SAM package installed locally, you can define how the software scans the archiving file systems on the current server for archiving. Each scanning method determines which files on the archiving file systems need to be archived and determines the archiving schedule. Scanning occurs at the time intervals that you specify on the General Setup page or on the File System Archive Policies page.

Scanning methods defined on the General Setup page are overridden by the scanning method defined on the File System Archive Policies page for a file system.

You can choose the scanning methods described in the following table.

TABLE 7-5 Scanning Methods

Scan Method	Description
Traditional Scan	The software performs several scans of each file system on the current server. The first file system scan descends recursively through the directory tree. Each file is examined, and files that do not need to be archived are flagged. During successive scans, the <code>.inodes</code> file is examined. The software examines only the inodes that are not flagged. Upon scan completion, the software sends the archive requests to the archiver daemon to be scheduled for file copying to media.
Scan Directories	The software performs a directory scan of each file system on the current server. The directory scan descends recursively through the directory tree of each file system. Files that do not need to be archived are flagged. The software assumes that all remaining files need archiving. Upon scan completion, the software generates archive requests based on its finding and sends the requests to the archiver daemon to be scheduled for file copying to media. This method might be useful if the majority of your files always need to be archived.
Scan inodes	The software scans the <code>.inodes</code> file. All inodes that are not flagged are examined. A file is flagged if it does not need to be archived. Upon scan completion, the software generates archive requests based on its findings. This method might be useful if you know which files are changed on a regular basis. This scan method will examine only those files, without performing a larger scan.
No Scan	The software does not scan the file systems on the current server to determine the files to be archived. Instead, the software works with the file system to determine what files need to be archived. The archiver determines when to begin archiving by using the Start Age, Start Count, and Start Size values that you define in your policy copies. This is the default archiving method, otherwise known as continuous archiving.

Configuring Global Archiving Values

If you are managing a server that has the Sun StorageTek SAM package installed locally, you can configure global archiving values for all archiving file systems on the current server. When you configure global archiving values, the `archiver.cmd` file is updated.

To configure global archiving values:

1. From the Servers menu or the Servers page, choose the name of the server that you want to configure.

2. From the navigation tree, choose General Setup under Archive Administration.
The General Setup page is displayed.
 3. Edit the archiver-specific fields that you want to change, and click Save.
-

Configuring Global Staging Values

If you are managing a server that has the Sun StorageTek SAM package installed locally, you can configure global staging values for all archiving file systems on the current server. The stager restores file data to the disk cache. When a user or process requests file data that has been released from the disk cache, the stager automatically copies the file data back to the online disk cache.

When you configure global staging values, the `stager.cmd` file is updated.

To configure the stager:

1. From the Servers menu or the Servers page, choose the name of the server that has the stager you want to configure.
 2. From the navigation tree, choose General Setup under Archive Administration.
The General Setup page is displayed.
 3. Edit the stager-specific fields that you want to change, and click Save.
-

Configuring Global Releasing Values

If you are managing a server that has the Sun StorageTek SAM package installed locally, you can configure global releasing values for all archiving file systems on the current server. Releasing is the process of freeing primary (disk) storage that is used by an archived file's data. Files are selected for release according to criteria such as age.

When you configure global releasing values, the `releaser.cmd` file is updated.

To configure the releaser:

1. From the Servers menu or the Servers page, choose the name of the server that has the releaser you want to deliver.
2. From the navigation tree, choose General Setup under Archive Administration.

The General Setup page is displayed.

3. Edit the releaser-specific fields that you want to change, and click Save.

Configuring Global Media Type Parameters

If you are managing a server that has the Sun StorageTek SAM package installed locally, you can configure global archiving and staging parameters for each available media type that is associated with the current server. When you configure global archiving and staging parameters, the `archiver.cmd` and the `stager.cmd` files are updated.

To configure media type parameters:

1. From the Servers menu or the Servers page, choose the name of the server that you want to configure.
2. From the navigation tree, choose General Setup under Archive Administration.
The General Setup page is displayed.
3. Edit the fields that you want to change in the Media Type Parameters and Drive Limits section of the page, and click Save.

Configuring the Recycler

This section provides information about the recycler. It contains the following subsections:

- [“About the Recycler” on page 152](#)
- [“Running Recycling” on page 154](#)
- [“Configuring Library Recycling Values” on page 154](#)
- [“Configuring Policy Recycling Values” on page 155](#)

About the Recycler

The recycler is a component of the Sun StorageTek SAM software that prevents you from overextending the archive disk volumes or removable media in your environment. It does this by freeing up space on volumes that contain expired archive copies.

As users modify files on archiving file systems, archive copies associated with the old versions of these files are considered to be expired on their archive disk volumes or removable media. These copies can be purged from the system to make room for new copies.

The recycler attempts to purge the expired copies first by selecting volumes that are eligible to be recycled and then by recycling the eligible volumes.

Selecting Eligible Volumes

The recycler selects volumes for recycling based on the amount of space used by expired archive copies as a percentage of total space on the volume. The volumes with the highest percentages of expired copies are selected to bring the storage usage below the configured High Water Mark threshold.

After the volumes are selected, the recycler determines whether the volumes would gain enough space after recycling to warrant the operation. The recycler determines this by examining the value in the Minimum Gain field either on the Recycler page (for volumes in a library) or in the archive copy (for volumes associated with a policy). If the selected volumes do not meet this criterion, they are not recycled.

Recycling Volumes

If the selected volumes are considered eligible for recycling, the recycler starts the recycling process. Volumes are recycled differently, depending on whether the volume is a disk volume or whether it is a removable cartridge in a library:

- When a disk volume is selected for recycling, additional archive copies can still be written to it. However, expired archive copies on the disk volume are identified and removed. Valid archive copies are left alone.
- When a tape or magneto-optical volume is selected for recycling, the system prevents additional archive copies from being written to it. If you are recycling to cartridges in a library, all files with active archive copies in volumes on the cartridge are marked to be rearchived. The archiver moves these copies to other volumes. In subsequent runs, the recycler checks these volumes to ensure that all valid archive copies have been relocated.

The recycler then determines whether there are volumes that contain only expired archive copies. If it finds a volume in this condition, the recycler performs one of the following actions, according to the user-defined recycler settings:

- Relabels the volume for immediate reuse.
- Exports the volume from the library (ejects the cartridge from the library, which makes the that slot available for a new volume). The volume can then be moved to offsite storage. A historical record of file changes in the volume is preserved in the Historian.

You define library recycling directives on the Recycler page, which controls recycling for all volumes in the library. Alternatively, you can define archive policy-based recycling for disk or tape volumes in the Archive Policy Based Recycling section of an archive policy copy. Archive policy-based recycling is the only way to recycle disk volumes.

For information about recycler settings for libraries, see [“About Archiving, Releasing, and Staging” on page 113](#).

Running Recycling

If you are managing a server that has the `SUNWsamfsr` and `SUNWsamfsu` packages installed locally, you can automatically run recycling at off-peak hours by logging in to the current server as `root` and specifying a schedule in the `crontab(1)` file. For more information, see the `crontab(1)` man page. You can also run recycling on demand at any time.

When you run recycling, the recycler looks in the `recycler.cmd(1M)` and the `archiver.cmd(1M)` files for recycling directives. The `recycler.cmd(1M)` file specifies recycling directives for the libraries associated with the server you are administering.

To define recycler settings for libraries in the browser interface, go to the Recycler page. The values you specify are reflected in the `recycler.cmd(1M)` file.

The `archiver.cmd(1M)` file specifies recycling directives for the volumes that are associated with archive policies. To define recycler settings for volumes associated with archive policies in the browser interface, go to the Archive Policy Based Recycling section of an archive policy copy. The values you specify are reflected in the `archiver.cmd(1M)` file.

To run recycling on demand:

1. Log in to the server for which you want to run recycling.
2. At the command line, type the following:

```
sam-recycler
```
3. Log out of the server.

Configuring Library Recycling Values

The recycler clears volume serial names (VSNs) in a library of expired archive copies and makes the volumes available for reuse. For information about recycler configuration guidelines for a library, see [“Planning for Archiving” on page 115](#).

If you are managing a server that has the `SUNWsamfsr` and `SUNWsamfsu` packages installed locally, you can configure the recycler for volumes in a library.

To configure the recycler for volumes in a library:

1. From the Servers menu or the Servers page, choose the name of the server for which you want to configure the recycler.

The File Systems Summary page is displayed.

2. From the navigation tree, choose Storage Recycling under Archive Administration.

The Recycler page is displayed.

3. Configure recycling values for each library listed on this page, and click Save.

Note: If you configure the recycler to automatically export a volume after recycling it, manually move the exported volume to offsite storage. This ensures that old data is maintained in the event of a disaster.

Configuring Policy Recycling Values

The recycler clears expired archive copies from disk or tape volumes that are associated with archive policies. Archive policy-based recycling is the only way to recycle disk volumes.

If you are managing a server that has the `SUNWsamfsr` and `SUNWsamfsu` packages installed locally, you can configure values for volumes associated with an archive policy.

To configure recycling values for volumes associated with an archive policy:

1. From the Servers menu or the Servers page, choose the name of the server with the policy that you want to edit.
2. From the navigation tree, choose Archive Administration.
The Archive Policies Summary page is displayed.
3. Click the policy that you want to edit.
The Policy Details Page is displayed.
4. In the Copy Information table, click the name of the policy copy that you want to edit.
The Policy Copy Options page is displayed.
5. Edit the recycling-related fields that you want to change in the Archive Policy Based Recycling section, and click Save.

Managing Archiving and Staging Activity

This section provides information about controlling archiving and staging activity for all archiving file systems on the current server. It contains the following subsections:

- [“Controlling Global Archiving Activity” on page 156](#)
- [“Controlling Global Staging Activity” on page 157](#)
- [“Staging a File or Directory on Demand” on page 157](#)

Controlling Global Archiving Activity

You can control archiving activity for all archiving file systems on the current server.

To control archiving activity:

1. From the Servers menu or the Servers page, choose the name of the archiving server you want to control.
2. From the navigation tree, choose Archive Activity under Archive Administration.
The Activity Management page is displayed.
3. In the Archiving section of the page, select the radio button next to the option that you want to control and click Submit.

To view the progress of the archiving activity, choose Jobs under System Administration in the navigation tree.

Controlling Global Staging Activity

You can control staging activity for all archiving file systems on the current server.

To control staging activity:

1. From the Servers menu or the Servers page, choose the name of the server that has the staging activity you want to control.
2. From the navigation tree, choose Archive Activity under Archive Administration. The Activity Management page is displayed.
3. In the Staging section of the page, select the radio button next to the option that you want to control and click Submit Job.

To view the progress of the staging activity, click the Jobs tab and look for the job that you submitted.

Staging a File or Directory on Demand

The system automatically stages files and directories when you access them, but large files or sets of files can require a great deal of time to come online. You can avoid waiting for the system to stage files or directories by explicitly staging them.

If **version 4, update 6** of the Sun StorageTek SAM software is installed on the server, do the following to stage a file or directory:

1. From the Servers menu or the Servers page, choose the name of the server that contains the file or directory you want to stage.
2. From the navigation tree, choose File Browser.
The File Browser page is displayed.
3. Select the radio button next to the file or directory that you want to stage, and choose Stage from the Operations menu.
The Stage File window is displayed.
4. From the drop-down menu, choose the copy of the file or directory that you want to stage.
5. (Optional) To recursively stage all files within a directory, select Recursively Stage Files in this Directory.

6. Click Submit.

The system displays a message indicating that the file or directory has been successfully staged. Depending on the available system resources, the staging process can take some time.

If **version 4, update 5** of the Sun StorageTek SAM software is installed on the server, do the following to stage a file or directory:

1. From the Servers menu or the Servers page, choose the name of the server that contains the file or directory you want to stage.

The File Systems Summary page is displayed.

2. Select the radio button next to the file system for which you want to stage the file or directory, and choose Stage from the Operations menu.

The Stage File Page page is displayed.

3. (Optional) In Maximum Files/Directories, specify the maximum number of files or directories you want to retrieve, and click Apply.

A list of files and directories is displayed.

4. (Optional) To see a list of files within a directory, click the directory name.

5. Select the radio button next to the file or directory that you want to stage.

The Stage File window is displayed.

6. From the drop-down menu, choose the copy of the file or directory that you want to stage.

7. (Optional) To recursively stage all files within a directory, select Recursively Stage Files in this Directory.

8. Click Submit.

The system displays a message indicating that the file or directory has been successfully staged. Depending on the available system resources, the staging process can take some time.

Managing Removable Media

This chapter provides information about managing removable media. It contains the following sections:

- [“About Removable Media Management” on page 159](#)
- [“About Removable Media” on page 160](#)
- [“About Supported Libraries and Drives” on page 161](#)
- [“About the Historian” on page 162](#)
- [“Managing Libraries” on page 163](#)

About Removable Media Management

If you are managing a server that has the `SUNWsamfsr` and `SUNWsamfsu` packages installed locally, you can associate removable media with the archiving file systems on the server. Otherwise, the Media Management tab is not displayed in the browser interface, and you cannot configure removable media.

If the current server has the `SUNWsamfsr` and `SUNWsamfsu` packages installed locally, you can bring automated libraries under the control of Sun StorageTek SAM software by adding the tape drives and libraries to the browser interface. Adding libraries to the browser interface automates the drive ordering of certain libraries. For more information, see [“About Supported Libraries and Drives” on page 161](#).

When you add libraries to the browser interface, the libraries are automatically added to the `mcf` file on the current server.

After removable media devices are added to the `mcf` file, you can use the browser interface to view the status of the removable media, view the details of libraries (drives and catalogs), and search for volume serial names (VSNs) across multiple libraries. You can also load and unload a VSN, import and export a VSN, and modify VSN media attributes.

When you perform any removable media operation, the software displays a message telling you whether the operation was started successfully. Some operations take a long time to complete. To view the status of these operations, refresh the Library Details or the Drive Details page.

About Removable Media

If you want the archiving file systems on the current server to be archived to removable media devices, you must have at least one available removable media device, either an automated library with multiple drives or a stand-alone tape drive. For information about guidelines for archiving to a library, see [“About Archiving, Releasing, and Staging” on page 113](#).

An automated library can be attached directly to a server through a small computer system interface (SCSI) cable or Fibre Channel, or it can be attached through the network. Devices from vendors such as Sony, Fujitsu, and IBM that attach to the network require third-party software tools.

One or more library catalogs are associated with each automated library. The library catalog tracks the movement of removable media devices within the library. For libraries that are directly attached, a record of the volume serial names (VSNs) in the

library is maintained. There is one catalog for each automated library, and there is one Historian for all automated libraries at a site. All media recorded in a library catalog must be of the same media type, such as IBM 3570 or SONY AIT.

Archive media devices that you intend to use must be attached to and recognized by the current server. For more information, see the *Sun StorageTek Storage Archive Manager (SAM) Installation and Upgrade Guide*.

About Supported Libraries and Drives

When you add a library or drive to the browser interface, you bring the library or drive under the control of the Sun StorageTek SAM software. A library or drive cannot be listed in the Sun StorageTek SAM `mc.f` file on the current server before being added to the browser interface.

Using the File System Manager interface, you can support bringing the following types of libraries and drives under the control of Sun StorageTek SAM software:

- **Small computer system interface (SCSI) direct attached libraries** – The File System Manager software can discover and display the tape libraries with the correct drive ordering after the library is attached to, and recognized by, the current server. The libraries must meet the following criteria:
 - The target must be less than or equal to 6, and the library must have no logical unit number identifiers (LUN IDs) or have only one LUN ID.
 - If the target is larger than 6, or the library has more than one LUN ID, you must configure the library properly so that it is attached to, and recognized by, the current server. This involves manually updating the `samst.conf` file on the current server and then adding the library to the browser interface. For more information, see the *Sun StorageTek Storage Archive Manager (SAM) Installation and Upgrade Guide*.

- **Fibre Channel direct attached libraries** – The File System Manager software can discover and display the tape libraries with the correct drive ordering after the library is attached to, and recognized by, the current server. In addition, the software updates the `samst.conf` file for the server with the World Wide Name (WWN), if the WWN does not already exist in the `samst.conf` file.
- **Network attached libraries** – You can enter the name, type, and parameter file location of this library type. You are responsible for correctly ordering the drives in the `mcf` file on the server, which you do by editing the parameter file of the library. The following libraries of this type are supported:
 - Sun StorageTek ACSLS
 - IBM 3494
 - Sony PetaSite
 - Fujitsu
 - ADIC and Grau

About the Historian

The Historian is a catalog that keeps track of cartridges exported from an automated library or unloaded from manually loaded devices. The Historian is similar to the catalog of an automated library but has no defined hardware devices associated with it.

Like an automated library, the Historian is configured in the `mcf` file, records entries for all cartridges associated with it, can import and export cartridges, and is displayed in the browser interface as an automated library.

By default, the Historian is created with 32 entries and can grow to be larger. Each time the catalog becomes full, 32 new entries of approximately 200 bytes each are added. Make sure that the file system on which the Historian resides has enough storage space to support this potential growth.

Managing Libraries

This section provides information about managing libraries:

- [“About Libraries” on page 163](#)
 - [“Adding a Library” on page 164](#)
 - [“Changing the State of a Library” on page 165](#)
 - [“Unloading a Library” on page 166](#)
 - [“Removing Archiving, Staging, and Recycling Configurations” on page 166](#)
 - [“Deleting a Library” on page 167](#)
 - [“Managing Library Drives” on page 168](#)
 - [“Managing VSNs in a Library” on page 170](#)
-

About Libraries

Libraries are also known as media changers, jukeboxes, robots, automated libraries, or media libraries. The File System Manager software enables you to manage certain libraries, assuming the current server has the `SUNWsamfsr` and `SUNWsamfsu` packages installed locally. For information about the types of libraries that are supported, see [“About Removable Media Management” on page 159](#).

Libraries are robotically controlled devices designed to load (import) and unload (export) removable cartridges without operator intervention. A library contains a robotic arm that is used to transport cartridges, slots in which to store cartridges, and one or more drives on which to mount cartridges.

Each cartridge in a library has one or more volumes. A volume is a named area on a cartridge that is used for storing data. Double-sided cartridges have two volumes, one on each side. A volume serial name (VSN) identifies a volume.

You can use the File System Manager software to add libraries to or remove libraries from the archiving file system configuration. Archive policies for the site determine which libraries and drives to use for archiving and staging.

Adding a Library

Before adding a library to the browser interface, you must connect the library directly either to the current server, or through the network. For more information, see the *Sun StorageTek Storage Archive Manager (SAM) Installation and Upgrade Guide*.

For information about the types of libraries that you can add to the configuration, see [“About Supported Libraries and Drives” on page 161](#).

Before you can add a network attached library through this wizard, ensure that the parameter file for the library exists and that the drives are ordered correctly in the parameter file.

Note: If you are adding one or more STK ACSLS (Sun StorageTek Automated Cartridge System Library Software) libraries and are also adding one or more libraries that are attached directly to the current server, complete this procedure first for the STK ACSLS libraries and then go back and complete the procedure for the directly attached libraries. This sequence is necessary for proper operation of the system.

To add a library:

1. From the Servers menu or the Servers page, choose the name of the server to which you want to add a library.
2. From the navigation tree, choose Storage Administration.
The Library Summary page is displayed.
3. Click Add.
The Add Library wizard is displayed.
4. Specify the type of library you want to add and click Next.
 - Select Direct Attached if the library is attached directly to the current server.
 - Select STK ACSLS if the library is an ACSLS library.
 - Select Network Attached if the library is attached to the server through the network.
5. Enter information about the library as requested by the wizard. (The information requested varies depending on which type of library you select.)

6. Review the specified information for the new library. If you are satisfied, click Finish. Otherwise, click Previous to make changes.

Changing the State of a Library

A library can be in a state of On, Off, or Unavailable.

- When a library is On, it is under the control of the Sun StorageTek SAM software. This enables the Sun StorageTek SAM software to perform the following actions:
 - Discover the internal state of the library, such as the location of tapes or whether barcodes are used
 - Update the library catalog and other internal structures
- When a library is Off, it is no longer under the control of the Sun StorageTek SAM software, which stops I/O operations and the automatic movement of cartridges. The drives in the library continue to be in the On state.

You might want to turn off a library for the following reasons:

- To stop Sun StorageTek SAM operations for this library only
- To power off the automated library
- When a library is Unavailable, you can reconfigure the library, such as configuring its cartridges.

To change the state of a library:

1. From the Servers menu or the Servers page, choose the name of the server that you want to administer.
2. From the navigation tree, choose Storage Administration.
The Library Summary page is displayed.
3. Select the radio button next to the library whose state you want to change.
4. From the Operations menu, choose Change State.
The Change State window is displayed.
5. Select the state that you want, and click Submit.

Unloading a Library

When you unload a library, all the volume serial names (VSNs) in the library are removed and placed into the Historian catalog. This changes the state of the library to Off.

Note: Make sure you are aware of the consequences before performing this procedure. Unloading the library stops all archiving activity for that library.

You can import the VSNs from the Historian catalog back into the library at a later time.

To unload a library:

1. From the Servers menu or the Servers page, choose the name of the server that you want to administer.
2. From the navigation tree, choose Storage Administration.
The Library Summary page is displayed.
3. Select the radio button next to the library that you want to unload.
4. From the Operations menu, choose Unload.
A message box prompts you to confirm the unloading.
5. Click OK.

Removing Archiving, Staging, and Recycling Configurations

A library cannot be deleted if you have changed the default values to configure archiving, staging, and recycling for the library. Before deleting the library, you must remove these configurations.

Removing Archiving and Staging Configurations

To remove the archiving and staging configuration settings for a library:

1. From the Servers menu or the Servers page, choose the name of the server that you want to administer.

2. From the navigation tree, choose General Setup under Archive Administration.
The General Setup page is displayed.
3. In the Drive Limits table, delete the values in the Maximum Drives for Archiving and the Maximum Drives for Staging columns for the library that you want to delete. These columns must be left blank.
4. Click Save.
The default drive count values for the library are displayed in the columns that you just edited. This can be ignored.

Removing Recycling Configurations

To remove recycling configuration settings for a library:

1. From the Servers menu or the Servers page, choose the name of the server that you want to administer.
2. From the navigation tree, choose Storage Recycling under Archive Administration.
The Recycler page is displayed.
3. In the Recycling Characteristics table, delete all of the values for the library. For example, clear the check box, delete all values in the text fields, and choose - - from the Size Limit unit of measure menu.
4. Click Save.
The default recycling parameters for the library are displayed in the table. These values can be ignored.

Deleting a Library

Deleting a library logically removes it from the server to which it is attached. Data contained in the library is no longer accessible to the archiving file systems associated with the library, but the data itself is not destroyed.

If you have not configured archiving, staging, and recycling for the library, you can immediately delete the library. Otherwise, you must first remove these configurations. For instructions, see [“Removing Archiving, Staging, and Recycling Configurations” on page 166](#).

To delete a library:

1. From the Servers menu or the Servers page, choose the name of the server that you want to administer.
2. From the navigation tree, choose Storage Administration.
The Library Summary page is displayed.
3. Select the radio button next to the name of the library you want to delete.
4. From the Operations menu, choose Delete from the list of commands.
A message box prompts you to confirm the deletion.
5. Click OK.

Managing Library Drives

This section provides information about managing drives in a library. It contains the following subsections:

- [“Changing the State of a Library Drive” on page 168](#)
- [“Idling a Drive” on page 169](#)
- [“Unloading a Drive” on page 169](#)
- [“Cleaning a Drive” on page 170](#)

Changing the State of a Library Drive

You might want to change the state of a drive in a library. For example, if a cartridge becomes lodged in the drive, turn the state of the drive to Off before addressing the situation.

To change the state of a library drive:

1. From the Servers menu or the Servers page, choose the name of the server that you want to administer.
2. From the navigation tree, choose Storage Administration.
The Library Summary page is displayed.
3. Select the radio button next to the library on which the drive that you want to change is located, and click View Drives.
The Library and Drive Details page is displayed.

4. Select the radio button next to the row in which the drive that you want to change is located, and click Change State.

The Change State window is displayed.

5. Select the drive state that you want, and click Submit.

Idling a Drive

Idling a drive stops any archiving or staging activity that is in process. You might idle drives in a library if you wanted to manually manipulate cartridges in a library without disturbing archiving or staging operations.

To idle a drive:

1. From the Servers menu or the Servers page, choose the name of the server that you want to administer.
2. From the navigation tree, choose Storage Administration.

The Library Summary page is displayed.

3. Click the name of the library on which the drive that you want to idle is located.

The Library and Drive Details page is displayed.

4. Select the radio button next to the row in which the drive that you want to idle is located, and click Idle.

To restart archiving again after idling a drive, click Restart in the Archiving section of the Activity Management page. For more information, see [“Controlling Global Archiving Activity” on page 156](#).

Unloading a Drive

If a volume is no longer needed, its cartridge is automatically unloaded. However, you can also manually unload a drive. When you unload a drive, the mounted cartridge is removed from the drive and placed in its storage slot.

Before unloading a drive, you must idle the drive. For more information, see [“Idling a Drive” on page 169](#).

To unload a drive:

1. From the Servers menu or the Servers page, choose the name of the server that you want to administer.
2. From the navigation tree, choose Storage Administration.

The Library Summary page is displayed.

3. Click the name of the library on which the drive that you want to unload is located.

The Library and Drive Details page is displayed.

4. Select the radio button next to the row in which the drive that you want to unload is located, and click Unload.

To load a different cartridge containing a new volume into a drive, see [“Loading a VSN Into a Drive in a Library” on page 173](#).

Cleaning a Drive

You can clean the read/write heads of a tape drive by mounting a cleaning tape on the drive and running it.

To clean a drive:

1. From the Servers menu or the Servers page, choose the name of the server that you want to administer.

2. From the navigation tree, choose Storage Administration.

The Library Summary page is displayed.

3. Click the name of the library on which the drive that you want to clean is located.

The Library and Drive Details page is displayed.

4. Select the radio button next to the row in which the drive that you want to clean is located, and click Clean.

Managing VSNs in a Library

This section provides information about managing volume serial names (VSNs) in a library. It contains the following subsections:

- [“Searching for a VSN” on page 171](#)
- [“Importing a VSN” on page 171](#)
- [“Exporting a VSN” on page 172](#)
- [“Loading a VSN Into a Drive in a Library” on page 173](#)
- [“Unloading a VSN From a Drive in a Library” on page 173](#)

- [“Labeling a VSN” on page 174](#)
- [“Reserving a VSN in a Library” on page 175](#)
- [“Unreserving a VSN in a Library” on page 176](#)
- [“Auditing the Slot for a VSN in a Library” on page 176](#)
- [“Editing the Media Attributes for a VSN in a Library” on page 177](#)

Searching for a VSN

You can search for a specific volume serial name (VSN) from two different pages in the browser interface:

- Historian page – Initiates a search within the Historian
- VSN Summary page – Initiates a search within the selected library

This topic provides an example of how to search for a VSN from the Library Summary page.

To search for a VSN:

1. From the Servers menu or the Servers page, choose the name of the server where you want to search for a VSN.
2. From the navigation tree, choose Tape VSNs under Storage Administration.
The VSN Summary page is displayed.
3. Click Search for VSN.
The Search VSNs page is displayed.
4. Type the VSN number you want to find, and click Search VSN.
If an exact match is found, the VSN Details page for that VSN is displayed. If several matches are found, the Search Results page is displayed.

Importing a VSN

When you import a volume serial name (VSN), the cartridge is brought into the library and placed in the first available slot.

For a direct attached library with mail slots, the File System Manager software automatically imports one VSN at a time or imports all the VSNs at one time, depending on the library. If you import one VSN, you must repeat this procedure for each of the remaining VSNs.

Note: If you import a VSN for a network attached library, and the VSN is unavailable, the software imports the VSN to the Historian instead. For information about when the software considers a VSN unavailable, see [“Administering Stand-Alone, Non-archiving File Systems and Archiving File Systems”](#) on page 68.

To import a VSN:

1. From the Servers menu or the Servers page, choose the name of the server to which you want to import a VSN.
2. From the navigation tree, choose Storage Administration.
The Library Summary page is displayed.

3. Select the radio button next to the name of the library into which you want to import the VSN, and click Import.

Depending on the library driver configured with the system, one of the following occurs:

- SAMST – The appropriate VSN is automatically imported.
- ACSLS – The Specify VSN Name window is displayed.

Specify either a starting VSN and an ending VSN, or enter one VSN name, and click Submit.

The Import VSNs page provides advanced filtering options for importing VSNs.

- Other – This is a library driver other than SAMST or ACSLS. The Specify VSN Name window is displayed.

Either specify a starting VSN and an ending VSN, or enter one VSN name, and click Submit.

Exporting a VSN

Exporting a volume serial name (VSN) removes the cartridge containing the selected volume from the library and places it in the Historian. For direct attached libraries, at least one mail slot is needed for a VSN to be exported.

To export a VSN:

1. From the Servers menu or the Servers page, choose the name of the server that has the VSN you want to export.
2. From the navigation tree, choose Storage Administration.

The Library Summary page is displayed.

3. Select the radio button next to the library from which you want to export the VSN, and click View VSNs.

The VSN Summary page is displayed.

4. Select the radio button next to the slot that contains the VSN you want to export.
5. From the Operations menu, choose Export.

The cartridge is ejected from the library device.

Loading a VSN Into a Drive in a Library

When a volume serial name (VSN) is requested for archiving or staging, a cartridge containing the volume is automatically loaded into a drive in the library. You can also manually perform this operation. The software determines the drive into which the VSN is loaded.

To load a VSN into a drive in a library:

1. From the Servers menu or the Servers page, choose the name of the server that you want to administer.
2. From the navigation tree, choose Storage Administration.

The Library Summary page is displayed.

3. Select the radio button next to the library into which you want to load, and click View VSNs.

The VSN Summary page is displayed.

4. Select the radio button next to the slot for the VSN you want to load.
5. From the Operations menu, choose Load.

The cartridge is drawn into the drive unit.

Unloading a VSN From a Drive in a Library

When a volume serial name (VSN) is no longer needed, its cartridge is automatically removed (unloaded) from the drive in the library and placed in its storage slot. You can also manually perform this operation.

To unload a VSN from a drive in a library:

1. From the Servers menu or the Servers page, choose the name of the server that you want to administer.

2. From the navigation tree, choose Storage Administration.
The Library Summary page is displayed.
3. Select the radio button next to the library that you want, and click View Drives.
The Library and Drive Details page is displayed.
4. Select the radio button next to the drive containing the VSN that you want to unload, and click Unload.
A message tells you whether the unload operation is issued successfully.

Labeling a VSN

Labeling a volume serial name (VSN) enables you to name it for the first time. You can also relabel a volume that already has data on it.

If you have volumes in a library that has no bar code reader, you must label the volumes to name them for the first time.

Caution: If you are relabeling a volume, all data on the volume is rendered inaccessible. To preserve the data in this volume, make a copy of it before proceeding.

To label a VSN:

1. Outside of the browser interface, make a copy of the volume to preserve its data.
2. From the Servers menu or the Servers page, choose the name of the server that you want.
3. From the navigation tree, choose Storage Administration.
The Library Summary page is displayed.
4. Select the radio button next to the library where the VSN you want to label is located, and click View VSNs.
The VSN Summary page is displayed.
5. Select the radio button next to the row in which the VSN you want to label is located, and click Label.
The Label Tape window is displayed.
6. Provide the following information, and click Submit Job.
 - Label Type – Select Label (name for the first time) or Relabel.
 - Label Name – Type a VSN name from one to six characters from the following categories:

Uppercase letters from A to Z

Numerals from 0 to 9

The following special characters: ! " % & ' () * + , . / : ; < > = ? _

- Block Size – Select a block size from the list. The block size determines the basic unit of online storage for the volume.

A message box prompts you to confirm the label operation.

7. Click OK.

A message tells you whether the label operation was issued successfully. If it was, a tape label job is created, which you can view from the Jobs tab. After the job is completed, the VSN is updated with the new label.

Reserving a VSN in a Library

You can reserve a specific volume serial name (VSN) to be used for archiving of designated files.

To reserve a VSN:

1. From the Servers menu or the Servers page, choose the name of the server that you want to administer.

2. From the navigation tree, choose Storage Administration.

The Library Summary page is displayed.

3. Select the radio button next to the library where the VSN you want to reserve is located, and click View VSNs.

The Library VSN Summary page is displayed.

4. Select the radio button next to the row where the VSN you want to reserve is located, and click Reserve VSN.

The Reserve VSN wizard is displayed.

5. Select one or more reservation methods, as prompted, and click Next.

- File System – You want the VSN reserved for archiving files from a specific file system.
- Archive Policy – You want the VSN reserved for archiving files that are associated with a specific archive policy.
- Owner – You want the VSN reserved for archiving files that are owned by a specific user, that reside in a specific directory path, or that are owned by a user in a specific group.

6. Depending on the options you selected in the previous step, provide the following information as prompted, and click Next.
 - File System Name – Select a file system from the list.
 - Archive Policy – Select an archive policy from the list.
 - Owner – Select Owner, Group, or Directory, and type the user name, owner name, or the directory path in the field next to the option you selected.
7. Review the specified information for the reservation method. If you are satisfied, click Finish.

Unreserving a VSN in a Library

Unreserving a volume serial name (VSN) removes its reserved status.

To unreserve a VSN:

1. From the Servers menu or the Servers page, choose the name of the server that you want to administer.
2. From the navigation tree, choose Storage Administration.

The Library Summary page is displayed.
3. Select the radio button next to the library where the VSN you want to unreserve is located and click View VSNs.

The Library VSN Summary page is displayed.
4. Select the radio button next to the row where the VSN you want to unreserve is located, and click Unreserve VSN.

Auditing the Slot for a VSN in a Library

Occasionally, the reported space remaining on a tape or magneto-optical cartridge might need to be updated in the library catalog. When you audit a slot in a library, the cartridge containing the volume serial name (VSN) is loaded, the label of the volume is read, and the library catalog is updated with the slot capacity information.

To audit the slot for a volume in a library:

1. From the Servers menu or the Servers page, choose the name of the server that you want to administer.
2. From the navigation tree, choose Storage Administration.

The Library Summary page is displayed.

3. Select the radio button next to the library whose slot you want to audit and click View VSNs.

The Library VSN Summary page is displayed.

4. Select the radio button next to the slot that you want to audit.
5. From the Operations menu, choose Audit.

The library catalog is updated with the slot capacity information.

Editing the Media Attributes for a VSN in a Library

You can edit various attributes for the media on which a volume serial name (VSN) is located.

To edit the media attributes for a VSN:

1. From the Servers menu or the Servers page, choose the name of the server that you want to administer.

2. From the navigation tree, choose Storage Administration.

The Library Summary page is displayed.

3. Select the radio button next to the library where the VSN you want to edit is located, and click View VSNs.

The Library VSN Summary page is displayed.

4. Select the radio button next to the row in which the VSN you want to edit is located, and click Edit Media Attributes.

The Edit Media Attributes window is displayed.

5. Select or deselect the check boxes next to the listed attributes, as desired, and click OK:

- Damaged Media – Indicates that the VSN might be damaged.
- Duplicate VSN – Indicates that another VSN exists with the same name.
- Read-only – Indicates that the VSN is read-only.
- Write-protected – Indicates that the VSN is write-protected.
- Foreign media – Indicates that the VSN is foreign; that is, it is not created in a Sun StorageTek SAM or Sun StorageTek QFS environment.
- Recycle – Specifies if the VSN will be recycled.
- Volume is Full – Indicates that the VSN has no space available for storage.
- Unavailable – Indicates that the VSN is not available for storage use.

- Need Audit – Indicates that the library catalog needs to be updated with the slot capacity information of the VSN.

Configuring Email Alerts

This chapter provides information about performing general administration tasks. It contains the following sections:

- [“About Email Alerts” on page 179](#)
- [“Creating an Email Alert” on page 181](#)
- [“Changing Email Alert Events for a Recipient” on page 181](#)
- [“Deleting an Email Alert” on page 182](#)

About Email Alerts

Email alerts are available only to servers that have the `SUNWsamfsr` and `SUNWsamfsu` packages installed locally.

The Email Alerts system informs a user of certain events or conditions by generating a message and sending it to the user’s email address.

You can configure the File System Manager software to automatically send email notifications to recipients when certain events occur. When you create a notification, you must provide an email address for the intended recipient of the notification. Every event that triggers a notification, except for the file system overflow event, applies to archiving file systems.

The 'root' user is the default subscriber for all email alerts.

If the error condition persists, a repeat email is sent every two hours. This prevents flooding the recipient's mailbox with emails sent repeatedly for the same error condition.

You can set up the system so that one or more of the events described in the following table triggers a notification.

TABLE 9-1 Events

Event	Description
Library or Drive is Down	One or more devices are unusable by the file system.
Archiving is Interrupted	The software is unable to complete archiving operations. An email message is sent if the archiver is interrupted for any of the following reasons: • No volumes that are associated with the policy are available. • A file being archived is too large for volume overflow on the remaining volumes that are associated with the policy. • A file being archived is too large for any remaining volume that is associated with the policy. • Joined files being archived are too large for any remaining volume that is associated with the policy.
Requested Volume is Unavailable	Media devices need to be imported or manually loaded to support archiving or staging.
Recycling is Complete	The recycler has identified one or more tape volumes that contain large proportions of expired archive copies. The recycler has finished removing the expired copies from the volumes. You can now either export or relabel the tape volume.
Recovery Points Error Recovery Points Warning	The creation of a recovery point for a file system was interrupted. An email message is sent if the recovery point fails or if the software attempts to run a disaster recovery script before creating the recovery point and the software encounters any of the following conditions: • It cannot find the specified script. • The script does not contain any content. • The script does not have root executable permissions. • The file system for which a recovery point is scheduled is not mounted. • The recovery point file itself does not contain any content. • The recovery point file size is truncated. This occurs when the percentage of content change in the current recovery point is less than 10 percent, when compared against the three previous recovery points. • The compression or indexing of the recovery point file fails.
No Space Available on File System	The file system is full. There is no space available for storage use.
File System Exceeds its High Water Mark	The software has detected that the space consumed by the file system exceeds the high water mark setting.

Note: If you have the `SUNWsamfsr` and `SUNWsamfsu` packages installed locally on a server, and you wish to enable notifications through the browser interface, be aware that the software automatically defines the following value in the `notify=filename` directive of the `archiver.cmd` file:

```
/etc/opt/SUNWsamfs/scripts/archiver.sh
```

If you manually change this value, the browser interface cannot send archiver-related notifications.

Creating an Email Alert

To create an email alert:

1. From the Servers menu or the Servers page, choose the name of the server for which you want to create a notification.
2. From the navigation tree, choose Email Alerts under System Administration. The Email Alert Summary page is displayed.
3. In the Subscribers field, type the email address of the recipient that you want notified.
4. Select the check boxes next to the events that you want to trigger notifications. For information about these events, see [“About Email Alerts” on page 179](#).
5. Click Submit.

Changing Email Alert Events for a

Recipient

You can change the types of events that will send notifications to a specific recipient.

Note: If you want to change the recipient who receives the notifications, you must create a new notification and specify the new recipient's email address.

To change notification events for a recipient:

1. From the Servers menu or the Servers page, choose the name of the server on which you want to change notification events.
2. From the navigation tree, choose Email Alerts under System Administration
The Email Alert Summary page is displayed.
3. Choose the recipient for whom you want to change events from the Subscribers menu.
4. Select or deselect the check boxes next to the events that you want to automatically trigger notifications to this recipient.
5. Click Submit.

Deleting an Email Alert

To delete an email alert:

1. From the Servers menu or the Servers page, choose the name of the server from which you want to delete a notification.
2. From the navigation tree, choose Email Alerts under System Administration.
The Email Alert Summary page is displayed.
3. Choose the recipient notification that you want to delete from the Subscribers menu.

4. Click Deselect All to deselect all of the notification events.
5. Click Submit.
A message box prompts you to confirm the deletion.
6. Click OK.

Troubleshooting

This chapter provides information about troubleshooting the File System Manager software. It contains the following sections:

- [“Using SAMreports” on page 185](#)
- [“Incompatible Software Release” on page 187](#)
- [“One Device Displayed Multiple Times” on page 188](#)
- [“Inaccessible Server” on page 189](#)
- [“No Remaining Space for a VSN” on page 190](#)
- [“Damaged or Stale Files in a Recovery Point” on page 191](#)
- [“Existing Fault Conditions Not Displayed” on page 191](#)

Using SAMreports

This section provides information about using SAMreports. It contains the following subsections:

- [“About SAMreports” on page 186](#)
- [“Generating SAMreports” on page 186](#)
- [“Displaying SAMreports” on page 187](#)

About SAMreports

The Sun StorageTek QFS and Sun StorageTek SAM Explorer feature collects information from a SAM-QFS environment and writes it to file `/var/tmp/SAMreport.<hostname>.<SAM/QFS version>.<date>`. The information contained in the SAMreport is an important aid to diagnosing complex problems, and it is needed by an engineer in the event of an escalation.

The SAMreport includes the following information:

- Packaging, revision levels, and licensing information
- Configuration files [`mcf(4)`, `archiver.cmd(4)`, `recycler.cmd(4)`, `inquiry.conf(4)`, `defaults.conf(4)`]
- Log files (`sam-log`, `messages`, `archiver.log`, `recycler.log`, `releaser.log`, trace files)
- Memory dump information

If log files are not routinely collected, an important source of diagnostic information is missing from the SAMreport. Ensure that sites implement a comprehensive logging policy as part of their standard system administration procedures.

Generate a SAMreport in the following circumstances:

- Whenever there is a system panic, core dump, crash, hang, or stall
- As close to any system event as possible

Generating SAMreports

To generate a SAMreport:

1. From the Servers menu or the Servers page, choose the server for which you want to generate a SAMreport.
2. From the navigation tree, choose System Administration.
The Server Configuration page is displayed.
3. Scroll down to the SAM Explorer section and click Generate Report.
The Run SAM Explorer page is displayed.
4. Select the directory where you want to store the SAMreport and the number of lines from each log file that should be included in the SAMreport.

5. Click Run.

The system generates the SAMreport in the directory you specified. The current date and time are appended to the file name.

Displaying SAMreports

To view the contents of a SAMreport:

1. From the Servers menu or the Servers page, choose the server for which you want to view SAMreports.

2. From the navigation tree, choose System Administration.

The Server Configuration page is displayed.

3. Scroll down to the SAM Explorer section and click the path name of the SAMreport that you want to view.

A window is displayed that enables you to view the contents of the selected file. Click the Next and Previous buttons to page through the report, or enter a line number in the Jump To box and click Go to navigate directory to a specific line.

Note: If the SAMreport path name is not an active link, the SAMreport file might be compressed. Browse to the file on the server and uncompress it so that you can view the contents.

Incompatible Software Release

The File System Manager management station can collect or discover data from any server that has the corresponding or prior update of the following packages installed locally:

- SUNWsamfsr and SUNWsamfsu
- SUNWqfsr and SUNWqfsu

For example, the File System Manager release 3.0 software was released at the same time as the Sun StorageTek QFS and Sun StorageTek SAM release 4, update 6, software. Therefore, the File System Manager release 3.0 browser interface can discover data from a server that has either of the following installed locally:

- Release 4, update 5 or update 6, `SUNWsamfsr` and `SUNWsamfsu` packages
- Release 4, update 5 or update 6, `SUNWqfsr` and `SUNWqfsu` packages

If the server has an older release level of the packages installed, “Not Supported” is displayed in the Release column on the Servers page.

If the server has a newer release level of the packages installed, a Down icon is displayed on the Servers page.

In either of these situations, do the following to obtain compatible release levels:

1. On the management station and on the server, use the `pkginfo(1M) -l` (the letter ell) command to verify the package release on the system.
2. Depending on the results of the `pkginfo(1M) -l` command, do the following:
 - If the packages on the management station and server are not compatible, upgrade the system that has the older release.
 - If the packages on the management station and server are compatible, use the `showrev(1M) -p` command on both the management station and server to view the patch release levels.

```
showrev -p /usr/xpg4/bin/grep -E 'SUNWsam|SUNWqfs'
```

For more information about compatible software and patch release levels, see the README document.

One Device Displayed Multiple Times

If a virtual device is exported from a storage box on two different ports, the current server might view this device as two separate devices. If it does, the File System Manager browser interface will display the device two times.

To address this situation, you must install multipathing software on the server and then click Refresh in the browser interface.

Inaccessible Server

If the File System Manager software cannot communicate with a server, a Down icon is displayed next to the server in the browser interface. There are several reasons this might occur:

- The `SUNWqfsr` and `SUNWqfsu` packages or the `SUNWsamfsr` and `SUNWsamfsu` packages might not be installed on the server or they might be at incompatible release levels. Search for these packages on the server to verify that they are installed. For information about compatible release levels, see [“Incompatible Software Release” on page 187](#).
- The File System Manager daemon (`fsmgmtd`) might not be running on the server. Do the following:
 - a. On the server, use the `fsmadm(1M)` `status` command to check whether the daemon is running.
 - b. If it is not, manually restart the daemon by typing the following:

```
/opt/SUNWsamfs/sbin/fsmadm restart
```

You can configure the daemon to start automatically by using the `fsmadm(1M)` `config -a` command on the server.
- The `rpcbind` process might not be running on the server. Do the following:
 - a. Log in to the server as `root`.
 - b. Restart the `rpcbind` process by typing the following:

```
/usr/sbin/rpcbind &
```

- The management station might not be able to remotely administer the server. All hosts that are listed in `/opt/SUNWsamfs/sbin/fsmadm` on the server can remotely administer the server. To add the management station to the list, log in to the server as `root` and use the `fsmadm(1M)` `add management-station-name.domain-name` command.
- The server is unavailable, it is not connected to the network, or the network itself is unavailable. Check the server's cabling and investigate any possible networking issues.

No Remaining Space for a VSN

If you receive a notification that there is no space available for archiving on a volume (VSN), do one of the following:

- Add more VSNs to the copies of the appropriate archive policy. The appropriate archive policy is the policy that is applied to the files being archived to that VSN. For information about how to perform this procedure, see [“Modifying VSN Assignments in a Policy Copy” on page 140](#).
- Recycle the expired archive copies on the VSN by running the recycler. For more information, see [“Running Recycling” on page 154](#).

Note: You can use the Monitoring Console to track VSN usage. One of the available reports lists archive copies with the highest usage. For more information, see [“Using the Monitoring Console” on page 55](#).

Damaged or Stale Files in a Recovery Point

If you encounter either of the following situations, you must archive the specified files and then create another recovery point of the archiving file system:

- You review the recovery point log file and you find a message that says a file is damaged and cannot be restored with the specified recovery point.
- You receive an email notification that a stale files exists and that the file cannot be restored using the specified recovery point.

Existing Fault Conditions Not Displayed

If you are aware that a server is in a fault condition, but there is no indicator of this fault condition in the browser interface, you might need to restart the Solaris sysevent daemon (`syseventd`) on the server.

The following explanation applies only to servers with version 4, update 5, or earlier versions of the software.

The reason this might occur is that `/opt` might not be part of the root file system on the server. If this is the case, the Solaris sysevent daemon (`syseventd`) starts before `/opt` is available when the server is rebooted. Because `/opt` is unavailable to the

daemon upon rebooting, the daemon cannot find some of the `SUNWsamfs` binaries that are required for the File System Manager software to report recent fault conditions or to send Recovery Point Interrupted or File System Overflow notifications.

To address this issue, perform the following:

1. After rebooting the server, log in to the server as `root`.
2. Type the following at the command prompt:

```
# pkill -HUP syseventd
```

The Solaris `sysevent` daemon is restarted. The daemon can now find the required `SUNWsamfs` binaries.

Glossary

Definitions obtained from the Storage Networking Industry Association (SNIA) Dictionary are indicated with “(SNIA)” at the end. For the complete SNIA Dictionary, go to www.snia.org/education/dictionary.

A

- allocation scheme** The method by which files are allocated to disk. There are two file allocation schemes: dual allocation and single allocation.
- archive age** The amount of time since a file’s last modification.
- archive files** Files that are copied by the archiver to archive media. Archive files are compatible with the standard UNIX `tar(1)` format. This compatibility ensures data compatibility with the Sun Solaris Operating System and other UNIX systems. The file format of archive files includes the file access data (inode) and the path to the file. If a complete loss of the Sun StorageTek SAM environment occurs, the `tar(1)` format enables file recovery using standard UNIX tools and commands.
- archive interval** The amount of time between complete archiving processes.
- archive media device** The media device to which an archive file is written. Archive media can be removable tape or magneto-optical cartridges in a library. In addition, archive media can be a mount point on another file system.
- archive policy** See *policy*.
- archiver** A program that automatically copies files from online disk to offline storage, such as a disk or tape.
- archive set** A group of files to be archived. The groups can be defined across any group of file systems. Files in an archive set share common criteria that pertain to size, ownership, group, and directory location. The archive set controls the

destination of the archive copy, the length of time for which the archive copy is kept, and the length of time the system waits before archiving the data. All files in an archive set are copied to the volumes associated with the archive set. A file in a file system can be a member of only one archive set. In the browser interface, archive sets are called archive policies.

archive storage	Copies of file data that have been created on archive media.
archiving file system	A file system that uses the Sun StorageTek SAM software for storage and archive management of its files. The Sun StorageTek SAM software controls access to all files stored and all devices configured in the master configuration file (<code>mcf</code>).
associative staging	A file attribute that you can define in a policy. When this attribute is set, all files to which the policy is applied are staged.
audit	The process of loading of cartridges to verify their volume serial numbers (VSNs). For magneto-optical cartridges, the capacity and space information is determined and entered into the automated library's catalog.
automated library	A robotically controlled device designed to load and unload removable media cartridges without operator intervention. Cartridges are imported to and exported from the library and they are loaded and unloaded automatically. The archiving and staging processes use a site-defined scheme for allocating the number of drives to be used.

B

block size	See <i>disk allocation unit</i> .
-------------------	-----------------------------------

C

capacity	The total amount of space available for storage on a file system, device, or object.
cartridge	A physical entity that contains a media device for recording data, such as a tape or optical disk.
catalog	A record of the VSN pools 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 file 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 another machine.
continuous archiving	A scheduled archiving process based on the Start Age, Start Size, and Start Count values in a policy copy. When a file meets the specified criteria, it is scheduled to be archived.
critical fault	A severe error condition. For more information about the different conditions that can cause a critical fault to occur in the browser interface, see the <code>/opt/SUNWsamfs/mibs/SUN-SAM-MIB.mib</code> file on the server on which the fault condition occurred.
current job	A File System Manager, Sun StorageTek QFS, or Sun StorageTek SAM process that is currently running or is attempting to run.
current management station	The management station that is hosting the current session of the File System Manager software. The name of the current management station is displayed in the banner of the File System Manager browser interface.
current server	The system for which you are monitoring and administering file systems through the browser interface. This server's name is displayed next to the Change Server button above the global navigational tabs in the browser interface.

D

data device	A device or group of devices upon which file data is stored.
DAU	See <i>disk allocation unit</i> .
direct access	A file staging option (stage never) designating that a file can be accessed directly from the archive media and need not be retrieved to disk cache.
direct attached library	An automated library connected directly to a server through a simple computer system interface (SCSI) interface. A SCSI attached library is controlled directly by the Sun StorageTek SAM software using the SCSI standard for automated libraries.

direct copy method	An offline copy method that copies files directly from the offline volume to the archive volume without using the cache. For this copy method, the source volume and the destination volume must be different, and two drives must be available.
disk	A nonvolatile, randomly addressable, rewriteable data storage device. This definition includes both rotating magnetic disks, optical disks, and solid-state disks, or nonvolatile electronic storage elements. (SNIA)
disk allocation unit (DAU)	The basic unit of online storage. The Sun StorageTek QFS file system uses an adjustable DAU. You can use this configurable DAU to tune the file system to the physical disk storage device. For Sun StorageTek QFS file systems that use the dual allocation scheme, the DAU is 16, 32, or 64 Kbytes. For Sun StorageTek QFS file systems that use the single allocation scheme, the DAU is a multiple of 8 Kbytes within the range of 16 Kbytes and 64 Mbytes.
disk archiving	The process of archiving of data to disk volumes on another file system.
disk cache	The disk-resident portion of an archiving file system. Disk cache is used to create and manage data files between an online disk and archive media. Individual disk partitions or an entire disk can be used as disk cache.
disk space threshold	An administrator-defined range of disk space that is available to a user. 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 according to these predefined disk space thresholds.
disk striping	The process of recording a file across several disks, thereby improving access performance and increasing overall storage capacity. See also <i>striping</i> .
down fault	An indicator that the File System Manager software is unable to communicate with a server or object, for any reason. For information about the types of conditions that can generate faults in the browser interface, see the <code>/opt/SUNWsamfs/mibs/SUN-SAM-MIB.mib</code> file on the management station.
drive	A device that transfers data to and from a removable media volume.
dual allocation	A method by which you define a small and large allocation sizes for files being allocated to storage devices. When a file is created on a storage device, the system allocates the first eight addresses of the file in the small allocation. If more space is needed, the file system uses one or more large disk allocation units (DAUs) to expand the file. For Sun StorageTek QFS file systems, the small allocation size is 4 Kbytes and the large allocation size is the DAU, which can be 16, 32, or 64 Kbytes.

E

- eq** See *equipment ordinal number*.
- equipment ordinal number** The status code for a removable media device.
- export** The unloading of removable cartridges from a library.

F

- failover** The process of moving metadata server control to a host other than the primary metadata server that is associated with a shared file system.
- family set** A storage device that is represented by a group of independent physical devices, such as a collection of disks or the drives within an automated library.
- fault** An error condition on the current server.
- file system** A software component that imposes structure on the address space of one or more logical or physical disks so that applications can deal more conveniently with abstract named data objects of variable size (files). (SNIA)
- file system resources** Sun StorageTek SAM software components, volumes, or disks that are associated with a file system. The Sun StorageTek SAM resources include the archiver, stager, releaser, and recycler.

G

- global directives** Archiver and releaser directives that apply to all file systems on the current server.

H

- hard limit** For disk quotas, a maximum limit on file system resources that users cannot exceed.

high water mark When defined as a mount option for an archiving file system, a threshold that triggers the release of disk space occupied by eligible archived files on the file system. When defined in a policy copy or on the Recycler page, a threshold that triggers recycling, expressed as a percentage of the total capacity of the volumes associated with a policy or an automated library.

Historian A catalog that keeps track of cartridges exported from an automated library or a manually mounted device. The Historian records entries for all cartridges associated with it, and can import and export cartridges. There is one Historian for all automated libraries at a site.

I

import The loading of removable cartridges into a library.

inode (index node) A 512-block of information that describes the characteristics of a file or directory. This information is allocated dynamically within the Sun StorageTek QFS file system.

involuntary failover An operation in which, when the current metadata server of a shared file system becomes inaccessible, metadata server control is switched from the inaccessible metadata server to the potential metadata server. Involuntary failover is performed from the potential metadata server.

J

job A process from a user or from the File System Manager, Sun StorageTek QFS, or Sun StorageTek SAM software.

L

last update The most recent time that the browser interface was updated with data.

library See *automated library*.

library catalog See *catalog*.

M

major fault	An error condition that requires some administrator action. For more information about the different conditions that can cause a major fault to occur in the browser interface, see the <code>/opt/SUNWsamfs/mibs/SUN-SAM-MIB.mib</code> file on the server on which the fault condition occurred.
management station	The fully qualified host name of the system on which the File System Manager software is installed locally. You can connect to this system's File System Manager browser interface from other hosts in the network. From the File System Manager browser interface, you can configure and administer file systems that are located on other servers in the network.
master configuration file (mcf)	The file, read at initialization time, that defines the relationships between devices within a Sun StorageTek SAM or Sun StorageTek QFS environment.
mcf	See <i>master configuration file</i> .
media	In the Sun StorageTek SAM environment, tape or optical disk cartridges.
media type	The type of tape or optical disk cartridge, such as, SONY AIT or IBM 3570.
metadata	Data about data. In the Sun StorageTek SAM or Sun StorageTek QFS environments, metadata is the index information needed to locate the exact data position of a file on a disk. Metadata consists of information about files, directories, access control lists, symbolic links, removable media, segmented files, and the indexes of segmented files. Metadata must be protected, because if the data is lost, the metadata that locates the data must be restored before the lost data can be retrieved.
metadata archiving	An archiving process in which metadata is archived but file data is not. In an archiving file system, the file system default policy archives the metadata for the archiving file system.
metadata device	A device, such as a solid-state disk or mirrored device, upon which a Sun StorageTek QFS file system's metadata is stored, separate from the system's file data. Separating file data from metadata can increase performance.
metadata server	A host system in a shared file system that is used by the file system for metadata, the <code>.inodes</code> file, and file data. The metadata server also controls storage and archive management operations for the shared file system, if the shared file system is configured to be archiving.
metadata snapshot	A file that captures all the metadata for a complete archiving file system at a single point in time. This file can be used by the File System Manager software to recover lost file data in the event of a disaster.

metadata snapshot schedule	The automatic creation of metadata snapshot files for archiving file systems. It is recommended that metadata snapshots are taken at least once a day.
minor faults	An error condition that is not as severe as a critical error condition, but that is more severe than a minor error condition. For information about the different conditions that cause minor faults to occur, see the <code>/opt/SUNWsamfs/mibs/SUN-SAM-MIB.mib</code> file on the server on which the fault condition occurred.
mount point	(1) A method used to attach a file system to the hierarchy at a path name location directory. (2) The directory on which a file system is mounted.

N

network-attached library	A library, such as IBM or Sony, controlled through a software package that is supplied by the vendor, IBM or Sony. The Sun StorageTek SAM file system interfaces with the vendor software using a Sun StorageTek SAM media changer daemon designed specifically for the library.
none copy method	An offline copy method that stages files as needed for each archive file before copying to the archive volume.
NFS (network file system) shared directories	Network file system directories in a file system that enable you to share files across servers that are connected to the network. Through the browser interface, you can manipulate the read and write permissions of these shared directories, and can control whether the directories are shared or unshared.
notification	A mechanism by which an email message is sent to the specified recipient when a certain condition is detected by the File System Manager software.

O

offline storage	Storage that requires operator intervention for loading.
online storage	Storage that is immediately available upon request, such as, disk cache storage.

P

- partition** A portion of a device or a side of a magneto-optical cartridge.
- policy** A mechanism that indicates how and when files in an archiving file system are archived to remote media. When the documentation refers to a policy, it is referring to a custom policy unless otherwise specified.
- pool** A collection of disks or virtual disks that can be carved into volumes.
- port** A connection point on a host system through which connected devices pass data into and out of the host.
- potential metadata server** A host system that has the ability to become the metadata server for a shared file system.
- primary metadata server** The original, preferred host system for a shared file system's metadata, `.inodes` file, and file data. The primary metadata server also controls the storage and archive management operations for the shared file system, if the shared file system is configured for archiving. When you create a shared file system, the server that you are managing at the time automatically becomes the primary metadata server for the shared file system.

Q

- quota** The amount of system resources that a user is allowed to consume. Quotas are not supported for removable media or disk archive resources.

R

- recycler** A utility that inspects all archive files that are on offline storage and determines which of these are no longer needed. Upon completing its assessment, the recycler utility enables you to reuse the space on the media.
- releaser** A software 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 between specified high and low thresholds.

- report** The result of a File System Manager query. Reports are built from the results of the most recent update of discovered data.
- restore** The process of recovering lost files or an archiving file system itself from archive media.
- robot** The portion of an automated library that moves cartridges between storage slots and drives. Also called a transport.
- round-robin** A data access method in which entire files are written to logical disks in a sequential fashion. The entire first 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.
- In Sun StorageTek QFS file systems, files are round-robin if round-robin access is specified or if the file system contains mismatched striped groups.

S

- SCSI** See *small computer system interface*.
- segment** A section of a large files that has been divided into smaller pieces.
- server** (1) A system that is hosting the file systems that you are managing through the browser interface. (2). As displayed in the browser interface's banner, the management station.
- severity level** An indicator of the seriousness of a fault condition.
- single allocation** A method by which you define one allocation size for files that are being allocated to storage devices. The allocation size is the disk allocation unit (DAU) value, which can be a multiple of 8 Kbytes within the range of 16 Kbytes and 64 Mbytes. Only stand-alone Sun StorageTek QFS file systems can use the single-allocation scheme.
- small computer system interface (SCSI)** A collection of ANSI standards and proposed standards which define I/O buses primarily intended for connecting storage subsystems or devices to hosts through host bus adapters. Originally intended primarily for use with small (desktop and desk-side workstation) computers, SCSI has been extended to serve most computing needs, and is arguably the most widely implemented I/O bus in use today. (SNIA)
- SNMP community name** The name of the simple network management protocol (SNMP) community within which the management station system and the SNMP host system are members.

SNMP port	The port number of the SNMP host.
soft limit	For disk quotas, a threshold limit on file system resources that you can temporarily exceed. Exceeding the soft limit starts a timer. When you exceed the soft limit for the specified time (the default is one week), no further system resources can be allocated until you reduce file system to a level below the soft limit.
stage ahead method	An offline copy method that stages the next archive file as the current archive file is written to the destination. For stage ahead copying, two drives must be available and space must be available in cache for all files in one archive file.
stage all method	An offline copy method that stages all files before archiving. For stage all copying, one drive must be available, and space must be available in cache for all files.
stager	A software component that copies files from offline storage back to online disk as they are needed.
staging	The process of copying an offline file from archive storage back to online storage.
stand-alone tape drive	A device that is similar to an automated library but that requires manual intervention for loading media into, and removal of media from, its drive.
storage area network (SAN)	A network whose primary purpose is the transfer of data between computer systems and storage elements, and among storage elements. A SAN consists of a communication infrastructure, which provides physical connections, and a management layer, which organizes the connections, storage elements, and computer systems so that data transfer is secure and robust. (SNIA)
storage slots	Locations inside an automated library in which cartridges are stored when they are not being used in a drive. If the library is direct attached, the contents of the storage slots are kept in the automated library's catalog.
striped group	A collection of devices within a Sun StorageTek QFS file system that are defined in the <code>mcf</code> file as one or more (usually two) <code>gxxx</code> devices. Striped groups are treated as one logical device and are always striped with a size equal to the disk allocation unit. You can specify up to 128 striped groups within a file system, but you can specify no more than 252 total devices. Files can be allocated across the striped groups.
stripe size	The specified number of disk allocation units (DAUs) to be allocated before allocation to the next device of a stripe.
striping	A data access method in which files are simultaneously written to logical disks in an interlaced fashion. In a Sun StorageTek QFS environment you can declare either striped or round-robin access for each file system, and you can declare striped groups within each file system.

**Sun StorageTek QFS
file system**

A high-speed UNIX file system that can store file data and metadata on the same device or on separate devices. The Sun StorageTek QFS software controls access to all files stored and all devices configured in the master configuration file (`mcf`).

**Sun StorageTek SAM
software**

Software that can be configured with a file system to provide storage, archive management, and archive retrieval capabilities. This software archives files by copying the files from online disk to archive media.

T

threshold

A mechanism that defines the desirable available storage window for online storage. Thresholds set the storage goals for the releaser. Also see *disk space threshold*.

U

unarchiving

The process of deleting archive entries for one or more files or directories.

V

volume

A named area on a cartridge or a disk for sharing data. A cartridge has one or more volumes. Double-sided cartridges have two volumes, one on each side.

volume overflow

The spanning of a single file over multiple volumes. Volume overflow is useful for sites that use very large files that exceed the capacity of their individual cartridges.

**volume serial name
(VSN)**

In the context of archiving to removable media cartridges, a logical identifier for magnetic tape and optical disk that is written in the volume label. In the context of archiving to disk cache, the unique name for the disk archive set. An individual volume is identified by media type and VSN.

voluntary failover

An operation in which, when the primary metadata server is running and available, the metadata server control is switched from the primary metadata server to a potential metadata server.

VSN See *volume serial name*.

W

World Wide Name

(WWN) A 16-digit, hexadecimal number that uniquely identifies a peripheral device, such as a switch or a fabric.

WWN See *worldwide name*.

Index

A

activity

- archiving, displaying, 50
- global archiving, controlling, 156
- global staging, controlling, 157
- recovery point, displaying status of, 38
- releasing, displaying, 50
- staging, displaying, 51

administrative privileges, assigning, 3

archive policies

- about, 123
- adding file match criteria to, 138
- adding to file systems, 77
- administering, 122
- archiver.cmd file for, 126
- archiving options for, enabling, 129
- copies of
 - about, 140
 - adding, 141
 - administering, 139
 - editing, 140
 - removing, 141
- VSN assignments to, 46, 140
- creating, 127, 128
- deleting, 129
- editing, 128
- file match criteria in, displaying, 44
- managing, 76, 122
- monitoring, 44
- no-archive type
 - about, 114, 125
 - creating, 112, 128
- removing file match criteria from, 139

archiver.cmd file, 126

archiving

- about, 111, 113, 114
- activity, 49, 50, 156
- administering policies for, 122
- associative
 - about, 118
 - enabling, 132
- configuring global values for, 149
- continuous
 - about, 117
 - enabling, 133
- controlling, 120
- to disk VSNs, 145
- to disk, enabling, 130
- file systems. *See* archiving file systems
- general, setting up, 148
- global activity for, controlling, 156
- idling for a file system, 121
- planning for, 115
- to pools of VSNs, enabling, 134
- running for a file system, 121
- scan methods, about, 148
- stopping for a file system, 121
- to tape, 131

archiving file systems

- about, 65
- adding policies to, 77
- administering, 68
- archiver scanning, configuring for, 120
- archiving for, controlling, 120
- checking, 72
- configuring scanning for, 120

- creating, 70
- deleting, 75
- file match criteria for, 77, 78, 79
- growing, 71
- idling archiving for, 121
- mount options for, editing, 73
- mounting, 73
- planning, 68
- protecting data in, 80
- restoring, 86
- running archiving for, 121
- unmounting, 74
- VSN association for, creating, 75

associative archiving

- about, 118
- enabling, 132

auditing, slots for VSNs in a library, 176

B

banner elements, 11

banner pane, 10

browser

- banner and content panes in, 10
- navigating through, 10
- overview of, 9

C

catalog, displaying for Historian, 48

cleaning drives, 170

clients

- adding, 98
- deleting, 99
- mounting, 102
- unmounting, 104

configuration files

- overview of, 25
- server, displaying, 34

content pane, 12

continuous archiving

- about, 117
- enabling, 133

controls, report table, 17

conventions, typographic, xvi

copies of archive policies. *See* archive policies, copies of

current server, changing, 16

D

damaged files in recovery point file, troubleshooting, 191

devices, discovering available, 60

directories, displaying NFS, 37

discovering available devices, 60

disk archiving, enabling, 130

disk VSNs

- archiving to, 145
- archiving to a Sun StorageTek 5800 VSN, 147
- archiving to pools of, 134
- creating, 146
- editing the media attributes of, 147

documentation

- accessing online, xviii
- related, xvii

E

email alerts

- about, 179
- changing events for a recipient of, 181
- creating, 181
- deleting, 182

explorer files or reports, SAM. *See* SAMreports

F

fault conditions not displayed, troubleshooting, 191

faults

- about, 56
- acknowledging, 58
- deleting, 58
- displaying, 57
- monitoring, 56

file match criteria

- about, 135
- adding existing to archiving file systems, 77
- adding file systems to, 136
- adding to a policy, 138
- administering, 135
- defined in a policy, displaying, 44
- displaying for a file system, 37
- editing, 136
- removing file systems from, 137
- removing from archiving file system, 79
- removing from policy, 139
- reordering for archiving file systems, 78
- using to display file systems, 45

file name patterns, 91

File System Manager software, installing, 2

file systems

- about, 61

- adding file match criteria to, 136

- adding NFS directories to, 107

- archiving. *See* archiving file systems

- devices on, displaying, 36

- file match criteria for, displaying, 37

- file match criteria, displaying using, 45

- HA-QFS. *See* highly available QFS (HA-QFS) file systems

- managing, 59

- NFS directories for, displaying, 37

- protecting data on, 80

- removing file match criteria from, 137

- restore files process for, displaying, 38

- shared. *See* shared file systems

- stand-alone, non-archiving. *See* stand-alone, non-archiving file systems

- status, displaying, 36

filter and search file name patterns, 91

filtering

- recovery point content, 90

- tables, 19

G

general archiving, setting up, 148

getting started (overview of tasks), 21

H

help system, using, 19

highly available QFS (HA-QFS) file systems, 67

- about, 67

- adding servers for, 27

- creating, 60, 69

Historian

- about, 162

- displaying catalog information for, 48

host file, shared, creating, 97

hosts, displaying for a shared file system, 39

I

incompatible software version message,
troubleshooting, 187

initial configuration tasks, 23

installation, of File System Manager software, 2, 3

J

jobs

- about, 52

- canceling, 54

- displaying, 53

- monitoring, 51

- using to view progress of processes, 3

L

labeling a VSN, 174

libraries

- about, 163

- adding, 164

- auditing slots for VSNs in, 176

- changing drive states, 168

- changing the state of, 165

- deleting, 167

- managing, 163

- managing VSNs in, 170

- media attributes of VSNs in, editing, 177

- monitoring, 46

- mount requests in, displaying, 49

- status and state, displaying, 48

- supported, 161

- unloading, 166

- unloading VSNs from drives in, 173

- unreserving VSNs in, 176

- VSNs associated with, displaying, 47

log and trace files on servers, displaying, 33

logging in and out, 6

M

management station, software installed on, 2

media attributes, editing for disk VSNs, 147

media type parameters, configuring global, 151

media, removable, managing, 159

metadata servers

- adding, 98

- deleting, 99, 100

- displaying for a shared file system, 39

- mounting, 102

- unmounting, 104, 105

metadata snapshots. *See* recovery points

minimum server release, displaying, 35

monitoring

- archiving activity, 49

- environment, 21

- faults, 56
- file systems, 35
- jobs, 51
- libraries, 46
- policies, 44
- releasing activity, 49
- servers, 31
- staging activity, 49
- mount options
 - editing, 73, 103
- mount requests, displaying for a library, 49
- mounting
 - archiving file systems, 73
 - clients, 102
 - media, requests, 49
 - metadata servers, 102
 - requests in a library, displaying, 49
 - shared file systems, 100, 101
 - stand-alone, non-archiving file systems, 73
 - user privileges to perform, 4

N

- navigation tree, 12
- navigational elements, 12
- Network File System (NFS) directories
 - about, 106
 - adding, 107
 - changing the status of, 109
 - displaying for a file system, 37
 - editing options for, 108
 - removing, 110
- NFS directories *See* Network File System (NFS) directories
- notifications. *See* email alerts

O

- online help system, using, 19

P

- passwords, entering, 6
- planning
 - archiving, 115
 - archiving file systems, 68
 - file match criteria, 138
 - recovery points, 82
 - shared file systems, 95
 - stand-alone, non-archiving file systems, 68

- policies, archive. *See* archive policies
- pools, archiving to, 134
- pop-up windows, overview of, 15
- potential metadata servers
 - adding, 98
 - deleting, 99
 - mounting, 102
 - unmounting, 104
- privileges, assigning user, 3
- product overview, 1

R

- recovery points
 - about, 81
 - damaged or stale files in, troubleshooting, 191
 - displaying status of activities for, 38
 - editing schedules for, 85
 - filtering content of, 90
 - planning for, 82
 - removing schedules for, 85
 - scheduling, 81, 83
 - searching content of, 90
 - taking on demand, 84
- recycler
 - about, 152
 - configuring, 152
 - configuring library recycling values for, 154
 - configuring policy recycling values for, 155
 - running, 154
- related documentation, xvii
- releaser
 - about, 119
 - configuring global values for, 150
- releases, supported levels of software, 33
- releasing
 - about, 113
 - displaying activity of, 50
 - monitoring activity of, 49
- removable media
 - about, 160
 - managing, 159
- report icons, 13
- report tables, 16, 17
- reserving a VSN, 175
- restoring
 - displaying status of process, 38
 - file system data, 86

file systems, 86

S

SAM explorer files or reports. *See* SAMreports

SAMreports

- about, 186

- displaying, 187

- generating, 186

- troubleshooting with, 185

scanning methods, archiver, 148

searching, for a VSN, 171

security

- limiting user privileges for, 3

- setting timeout value for, 7

- through protecting system data, 80

servers

- adding, 23, 26

- changing current, 16

- configuration files on

 - displaying, 34

 - overview of, 25

- configuration of, displaying, 32

- configuring, 23

- file warning messages for, displaying, 34

- inaccessible, troubleshooting, 189

- log and trace files installed on, displaying, 33

- minimum release supported for, displaying, 35

- monitoring, 31

- removing, 28

- software packages on, displaying, 33

session timeout value, setting, 7

shared file systems

- about, 66

- administering, 95

- creating, 96

- deleting, 105

- editing mount options for, 103

- hosts in, displaying, 39

- metadata server for, displaying, 39

- mounting and unmounting, 100, 101, 103

- planning, 95

shared hosts file, creating, 97

software packages installed on server,

- displaying, 33

software version, incompatible message,

- troubleshooting, 187

sorting tables, 18

stager

- about, 119

- configuring global values for, 150

staging

- about, 113

- controlling global activity for, 157

- displaying activity of, 51

- file or directory on demand, 157

- managing activity for, 156

- monitoring activity of, 49

stand-alone, non-archiving file systems

- about, 63

- administering, 68

- checking, 72

- creating, 70

- deleting, 75

- editing mount options for, 73

- growing, 71

- mounting, 73

- planning, 68

- unmounting, 74

status and state

- changing for a library drive, 168

- of file systems, 36

- of libraries, displaying, 48

- of library, changing, 165

- of recovery point activities, 38

- of NFS directories, changing, 109

- of restore processes, 38

STK ACSLS (StorageTek Automated Cartridge System Library Software), adding library for, 162, 164

StorageTek 5800. *See* Sun StorageTek 5800

StorageTek SAM software. *See* Sun StorageTek SAM software

Sun Cluster environment

- adding servers that are part of, 27

- as relates to HA-QFS file systems, 60, 67, 69

Sun StorageTek 5800

- archiving to, 147

Sun StorageTek SAM software, about, 113

supported release levels, 35

T

tables

- controls, 17

- filtering, 19

- sorting, 18
- tape archiving, enabling, 131
- tape drives
 - cleaning, 170
 - idling, 169
 - in library, displaying, 47
 - in library, managing, 168
 - supported, 161
 - unloading, 169
- task overview
 - for configuring servers, 23
 - getting started, 21
- timeout value, setting, 7
- trace files on a server, displaying, 33
- troubleshooting
 - damaged or stale files in recovery point, 191
 - of device displayed multiple times, 188
 - of fault conditions not displayed, 191
 - of inaccessible server, 189
 - of incompatible software version message, 187
 - using SAMreports, 185
 - VSNs with no available space, 190
- typographic conventions, xvi

U

- UNIX commands, using, xv
- unloading
 - drives, 169
 - libraries, 163, 166
 - VSNs from drives of a library, 160, 173
- unmounting
 - archiving file systems, 74
 - clients, 104
 - metadata servers, 104, 105
 - shared file systems, 100, 103
 - stand-alone, non-archiving file systems, 74
 - user privileges to perform, 4
- unreserving VSNs, 176
- user privileges, assigning, 3

V

- volume serial names (VSNs)
 - about creating associations with archiving file systems, 75
 - assigned to policy copies, displaying, 46
 - auditing slots in a library for, 176
 - creating disk, 146

- disk or tape, enabling archiving to, 134
 - editing attributes for, 177
 - exporting, 172
 - importing, 171
 - in a library, managing, 170
 - labeling, 174
 - library, displaying, 47
 - loading into a library, 173
 - managing disk for archiving, 145
 - modifying in policy copies, 140
 - no available space on, troubleshooting, 190
- pools
 - about, 142
 - creating, 143
 - deleting, 144
 - editing, 143
 - managing for archiving, 142
- reserving, 175
- searching, 171
- unloading from drives of a library, 173
- unreserving, 176

VSNs. *See* volume serial names (VSNs)

W

- warning messages in server files, displaying, 34
- wizards, 15