



Sun Customer Ready Infinite Archive System

User Operations Guide for Version 3.6

Sun Microsystems, Inc.
www.sun.com

Part No. 316196801
Oct 2009, Revision AD

Submit comments about this document at: dsc-customer-feedback@sun.com

Copyright ©2009 Sun Microsystems, Inc. All rights reserved.

Use is subject to license terms.

This distribution may include materials developed by third parties.

Portions may be derived from Berkeley BSD systems, licensed from U. of CA. Sun, Sun Microsystems, the Sun logo, Jiro and Sun StorEdge are trademarks or registered trademarks of Sun Microsystems, Inc. in the U.S. and other countries. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the U.S. and other countries.

Copyright © 2009 Sun Microsystems, Inc. Tous droits réservés.

L'utilisation est soumise aux termes du contrat de licence.

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

Des parties de ce produit pourront être dérivées des systèmes Berkeley BSD licenciés par l'Université de Californie. Sun, Sun Microsystems, le logo Sun, Jiro et Sun StorEdge 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.

About this book

This book is intended for those who are responsible for the routine monitoring and administration of Sun Customer Ready Infinite Archive System (IAS) appliances. It focuses on using the Java Web Console graphical user interface to check system health, gather performance metrics, and administer users and IAS file system sharing.

The remainder of this preface explains how you can use this book most effectively and how you can obtain updated or supplementary documentation.

Understanding the intended scope of this book

The IAS console provides substantially more functionality than monitoring and routine administration, of course, and this document provides basic instructions for using the additional capabilities. But IAS is, first and foremost, an *appliance* that is meant to shield the user from the underlying complexities of administering archiving file systems and file servers. An appliance should not require reconfiguration or, indeed, any significant operator intervention at all. Accordingly, this document assumes that tasks like handling outages and implementing upgrades are primarily the responsibility of Sun service personnel and are thus out of scope in a user operations guide. While basic instructions for many of these tasks are included herein, they are by no means comprehensive or suitable in all circumstances.

As a user guide, this document presumes that those who choose to undertake *expert* appliance-maintenance tasks have the expertise to carry them through with minimal guidance. So, before undertaking any of the expert tasks that are outlined in this book, you should be thoroughly trained in the nuances of SAM-QFS file systems, entirely familiar with the extensive SAM-QFS documentation available from Sun, and completely comfortable with UNIX system administration. Clustered, multinode IAS appliances require Sun Cluster expertise as well. Do not rely exclusively on the

task outlines below. Instead, if you are at all unsure of your capabilities with regard to SAM-QFS, clustering, the Solaris operating system, or any other aspect of IAS configuration and maintenance, contact Sun support services for assistance.

How this book is organized

The document is *task-oriented*, organized around the work you have to do rather than around the features or components of the product. Chapters and sections generally begin with a list of tasks covered. Tasks are presented in order, and the steps in each process are numbered, in the sequence in which they are to be performed.

Conditional steps (steps that you perform only in specified circumstances) begin with the condition (“If A ...”) and end with the corresponding action (“... do B”); if the condition does not apply, you simply skip the step. Each task ends with a reference to the next task in the sequence:

Next task: “Installing ...” on page 3.

When the setup process branches, the task ends with conditional alternatives:

Next task:

- If the customer does not plan to run the management console from a host on the local area network (LAN), press `skip`, and go to the next task.
- Otherwise, carry out the procedure “Configuring the Ethernet LAN” on page 57.

When you have finished a task or sequence of tasks, this is clearly noted in the document:

Stop here.

To minimize the time that you spend switching between publications or between major sections of the document, we have made an effort to avoid cross references to external information wherever possible. If you need to have a figure, a table, or a procedure, it should always be, at worst, on a neighboring page.

Taking advantage of this book’s hypertext features

If you choose to view this book online, rather than in printed form, you can jump quickly to any part of the book by clicking on the corresponding entry under the `Bookmarks` tab on the left side of the Adobe Acrobat interface. In addition, clicking on entries in the table of contents, cross references, or references to subsequent tasks will take you directly to the indicated part of the document. You can then use the back arrow on the Adobe Acrobat Reader to return, if desired, to the point you left. In addition, clicking on most Uniform Resource Locators (URLs) and on most references to online resources will open your default web browser to the

corresponding web page. You can, if necessary, obtain a required download immediately (be aware, however, URLs to specific pages change frequently and may not always be accurate).

Understanding the conventions used in this book

The table below illustrates the conventions that represent literal and variable values, commands, and property names in this book.

Convention	Meaning	Examples
AaBbCc123	Fixed-width text is used for literal values, including names of commands, files, directories, literal computer inputs/outputs, and Uniform Resource Locators (URLs)	Edit your <code>.login</code> file. Use <code>ls -a</code> to list files. % You have mail.
<i>AaBbCc123</i> <i>AaBbCc123</i>	Oblique text is used for variables that stand for real names or values and for book titles.	To delete a file, type: <code>rm filename</code> .
ABCDEF	Fixed-width uppercase text represents electrical and system states, Boolean values, and constants.	ON, OFF TRUE, FALSE ONLINE, DOWN
ABCD	Bold, san-serif, uppercase text indicates callouts in illustrations.	Click Submit (A below).
1.	Numbered paragraphs indicate steps in a process that should be executed in sequential order.	
■	Bulleted paragraphs indicate lists of alternatives or components.	
[iasd01-nol]root#	A commandline prompt	
Tab, Shift, A	Keys on the keyboard	
Ctrl-Alt-Del, Alt-0169	A combination of keys on the keyboard, pressed at the same time.	

Obtaining the latest information and supporting resources

Obtain the latest Sun documentation from `<docs.sun.com>`. Always check the documentation portal for updates to this document before proceeding. Documents distributed by other means, such as CDROM, may not reflect the latest changes to appliance hardware, software, and services.

This manual is part of a set of documents that describes the operations of the Sun StorageTek SAM, Sun StorageTek QFS, and Sun Cluster products. The table below shows the complete documentation set for these products at the time the current document was published, but be aware that titles change from time to time.

Title

Sun StorageTek QFS File System Configuration and Administration Guide

Sun StorageTek Storage Archive Manager (SAM) File System Configuration and Administration Guide

Sun StorageTek Storage Archive Manager (SAM) Archive Configuration and Administration Guide

Sun StorageTek Storage Archive Manager (SAM) Installation and Upgrade Guide

Sun StorageTek QFS Installation and Upgrade Guide

Sun StorageTek Storage Archive Manager (SAM) Troubleshooting Guide

Sun StorageTek QFS and Sun StorageTek SAM Release Notes

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

Commenting on this book

Sun welcomes your comments and suggestions for improving this book. Contact us at dsc-customer-feedback@sun.com. Please include the title, part number, issue date, and revision: *Sun Customer Ready Infinite Archive System User Operations Guide for Version 3.6*, part number 316196801 (Oct 2009 Revision AD).

Contents

About this book	iii
Understanding the intended scope of this book	iii
How this book is organized	iv
Taking advantage of this book's hypertext features	iv
Understanding the conventions used in this book	v
Obtaining the latest information and supporting resources	v
Commenting on this book	vi
1. Understanding IAS	1
Benefits of appliance architecture	1
Faster solution deployment	2
Manageability, reliability, and availability	2
Proven components	3
SAM for information lifecycle management	3
QFS for fast, reliable access to files	4
Solaris Cluster for high availability	5
IAS appliance best practices	5
Use subdirectories and network shares as virtual filesystems	5
Use subdirectories and file sharing for performance tuning	6
Always consult Sun services before altering IAS internals	6

Always use the IAS Java Web Console when making changes	7
Let IAS software manage the storage	7
2. Changing root passwords	9
Connecting the service host	10
▼ Attaching a host to the ILOM of a dual-node appliance	10
▼ Attaching a host to the ILOM of a single-node appliance	12
Changing the ILOM root password	15
▼ Changing ILOM passwords	16
Changing the Solaris root password	18
▼ Logging in to the Solaris commandline using ssh	18
▼ Logging in to the Solaris Java Desktop via the ILOM	19
▼ Changing the root password from the commandline	21
3. Configuring network name services	23
Name Service Switch (NSS)	23
Accessing Solaris on IAS appliances	24
▼ Logging in to the Solaris commandline using ssh	24
▼ Logging in to the ILOM	25
▼ Logging in to the Solaris Java Desktop via the ILOM	26
Setting name services on IAS appliances	28
Configuring Domain Name System (DNS)	28
▼ Setting up DNS services	28
Configuring Network Information Services (NIS)	29
▼ Configuring /etc/inet/hosts files on IAS server nodes	30
▼ Configuring /etc/inet/hosts files on NIS servers	31
▼ Configuring nsswitch.conf for use with NIS	31
▼ Setting up NIS and NIS+	32
4. Using the IAS Java Web Console GUI	35

Supported web browsers	35
Navigating the IAS Java Web Console	36
The status banner and dashboard	36
The dashboard	37
The server information area	37
The button bar	37
The role-based task tree	38
The content pane	39
Tables	39
Forms	40
Pop-up browser windows	40
▼ Enabling pop-up windows in Firefox browsers	40
Understanding IAS operational roles	41
Standard operating roles	42
Special roles	43
Understanding role-based user accounts	44
Using the IAS help system	44
5. Managing IAS user accounts	47
Managing access to an archival system	47
Role-based access controls	47
Auditable, individual accounts	48
Multiple accounts for auditable, role-based access	48
Working with user accounts	49
▼ Adding user accounts	49
▼ Modifying user account properties	50
▼ Deleting user accounts	52
6. Monitoring the system	53

Monitoring appliance health	54
Assessing basic system availability	54
Checking the system configuration	55
▼ Checking IAS software version information	55
▼ Checking basic system health	56
▼ Checking the archive-server node configuration	57
▼ Displaying software packages installed on a server node	58
▼ Displaying log and trace files	58
▼ Displaying configuration files	59
Monitoring archive system faults	61
▼ Monitoring file system alerts	61
▼ Examining associated archive system log and trace files	62
▼ Monitoring disk-storage hardware alarms	64
Investigating file system problems	66
▼ Investigating file-system, directory, or file access problems	66
Monitoring appliance metrics	70
Assessing storage utilization with Tier Summary	70
▼ Creating Tier Summary reports	71
▼ Viewing Tier Summary reports	72
▼ Deleting a Tier Summary report	73
Assessing media utilization and status	74
▼ Accessing the Media Utilization and Status page	74
▼ Creating a media utilization report	76
▼ Viewing a media utilization report	77
▼ Deleting a media utilization report	78
Assessing the distribution of files in the archive	78
▼ Generating a graphical file-distribution report	79
Monitoring archiving processes with the File System Manager Monitoring Console	82

▼ Accessing the Monitoring Console	82
Using IAS Explorer	85
▼ Running the IAS Explorer	85
Using SAM Explorer	86
▼ Running SAM Explorer	87
7. Sharing directories with IAS clients	89
▼ Preparing to share a directory	89
▼ Creating a Network File System (NFS) share	90
▼ Creating a SAMBA share	93
▼ Setting Windows Networking Options	96
8. Managing archival media	99
How IAS organizes storage media	99
Media pools	100
Hierarchical storage tiers	100
Organizing media with VSN pools	101
▼ Creating a VSN Pool	102
▼ Editing a VSN Pool	102
▼ Deleting a VSN Pool	103
Managing disk volumes	104
▼ Creating a disk VSN	104
▼ Editing the media attributes of a disk-based VSN	105
Managing removable tape media	105
Tape media in IAS solutions	106
Managing robotic tape libraries	106
▼ Adding a robotic tape library	107
▼ Changing the state of a library	108
▼ Unloading a robotic library	108

▼ Deleting a robotic library	109
Managing tape drives	109
▼ Changing the state of a tape drive	110
▼ Idling a drive	110
▼ Restarting archiving after idling a drive	111
▼ Unloading a drive	111
▼ Cleaning a drive	112
Managing tape volumes	112
Searching for tape volumes	113
▼ Searching for a volume by VSN	113
Importing tape cartridges	113
▼ Importing volumes into the tape library	113
Exporting tape cartridges	114
▼ Exporting a tape volume from the library	115
Loading and unloading tape drives	115
▼ Loading a tape volume into a drive in the library	116
▼ Unloading a volume from a drive in the library	116
Relabeling tape volumes	117
▼ Relabeling a tape volume	117
Reserving tape volumes	118
▼ Reserving a tape volume	119
▼ Cancelling a tape volume reservation	120
Auditing slots and updating the library catalog	120
▼ Auditing a library slot	121
Changing tape media attributes	121
▼ Editing the media attributes of a tape volume	121
Managing exported tapes with the file system Historian	122
9. Maintaining the archive file systems	125

Locating and repairing file corruption	126
▼ Checking file systems	126
Mounting and unmounting file systems	126
▼ Mounting or unmounting a file system	127
▼ Editing mount options	127
Saving a point-in-time file system image	128
▼ Creating a recovery point on demand	129
10. Restoring a file system	131
Protecting file systems and their contents	131
Protecting both data and metadata	131
Automating administration with SAMPM	132
Restoring a file system	134
▼ Indexing recovery point files for filesystem recovery	134
▼ Restoring a file system	136
Restoring directories and files	140
▼ Indexing recovery point files for directory and file recovery	140
▼ Restoring directories and files	143
Conserving disk space	148
Deleting indices as part of disaster recovery	149
Deleting indices as part of routine file recovery	149
Deleting recovery point indices	149
▼ Deleting the index for a recovery point	149
11. Archiving data	151
Archive management and IAS	152
Avoid complexity	152
Know what you are doing	153
Know your documentation	153

Storage Archive Manager (SAM)	153
The archiver	154
The releaser	154
The stager	154
The recycler	155
Managing archiving policies	155
Writing practical archiving policies	156
Adding a policy to a file system	156
▼ Starting the <code>New Archive Policy</code> wizard	157
▼ Defining the files to which the policy will apply	157
▼ Naming the policy and defining basic archiving behavior	158
▼ Defining copy parameters	159
▼ Setting work-queue and recycler parameters	160
▼ Finishing the new archive policy	161
Deleting a policy	162
Editing a policy	162
▼ Accessing policies for editing	162
▼ Defining additional file-match criteria	163
▼ Defining additional archive copies	165
▼ Modifying file-match criteria	167
▼ Modifying copy parameters	168
Controlling the archiver	169
▼ Stopping or idling archiving	169
▼ Running archiving	169
12. Managing archive server nodes	171
Changing the operating state of a node	172
Quiescing the archiver	172
▼ Stopping archiver activity	172

Shutting down server nodes	173
▼ Shutting down server processes from the IAS console	174
▼ Turning system power OFF	175
▼ Remotely controlling server nodes with the ILOM	177
▼ Powering down the server node using the ILOM	179
Starting up server nodes	179
▼ Turning system power ON	180
▼ Remotely controlling server nodes with the ILOM	181
▼ Powering up the server node using the ILOM	183
Rebooting	183
▼ Rebooting a node or appliance from the IAS console	183
Switchover (failover and failback)	185
▼ Switching over to the other server node	186
Maintaining network connectivity	187
Setting IP addresses on high-availability systems	187
▼ Setting IP addresses on high-availability systems	188
Setting IP addresses on single-node systems	189
▼ Setting IP addresses on single-node systems	189
Maintaining network routes	190
▼ Viewing a network route	190
▼ Adding a network route	190
▼ Deleting a network route	191
Managing email notifications	192
▼ Creating an Email Alert	193
Maintaining time synchronization	194
▼ Using external NTP servers to maintain the time and date	194
▼ Using internal NTP servers to maintain the time and date	195
13. Maintaining the IAS disk arrays	197

Scheduling repairs promptly if needed	197
Monitoring disk hardware regularly	198
14. Maintaining registration information	199
▼ Entering Sun Online Account Information	200
▼ Entering Contact Information	201
▼ Entering Site Information	202
▼ Specifying Internet Connection Settings	203
▼ Reading the privacy statement and saving changes	203
15. Troubleshooting	205
Using IAS diagnostic tools	205
Error conditions	205
No Remaining Space for a VSN	205
Damaged or Stale Files in a Recovery Point	206
Existing Fault Conditions Not Displayed	206
A. Registering for the ASR feature	207
▼ Starting the ASR registration wizard	207
▼ Declining ASR support	208
▼ Accepting ASR support	210
Glossary	213

Understanding IAS

The Infinite Archive System (IAS) is an integrated, fully assembled, ready-to-use appliance that is optimized for secure, highly accessible, long-term storage of large volumes of archival data. The appliance design integrates carefully matched, well-proven hardware and software components that minimize cost while maximizing performance and reliability. Extensive automation and user-friendly interfaces keep maintenance and administration to a minimum, so that you can concentrate on archiving and the business requirements that drive it, rather than on the implementation.

This chapter contains two main sections. The first, “[Benefits of appliance architecture](#)” (see page 1), explains how the IAS appliance works and why. This basic overview of the appliance design lays the groundwork for the second section, “[IAS appliance best practices](#)” (see page 5). This second section outlines the principles that will guide your approach to all of the specific procedures that we cover later in this book.

Benefits of appliance architecture

Packaging an archiving solution as an appliance separates routine but complex, low-level implementation tasks from the unique, high-level configuration tasks that make a technical solution work in your environment. The former are handled at the factory, in a tested, repeatable, standard process that insures consistent, baseline results. The latter are handled on site by Sun service professionals and by your own administrators. This approach ensures “[Faster solution deployment](#)” (see page 2), improves “[Manageability, reliability, and availability](#)” (see page 2), and makes optimal use of “[Proven components](#)” (see page 3).

Faster solution deployment

The Sun Customer Ready Program delivers factory matched, tested, fully integrated storage components in standardized configurations. The hardware and software are fully installed and work together out of the box. The complex installation and setup work that would otherwise consume an inordinate amount of time at your site is already complete when the system arrives. With functioning hardware and software as a given, the implementation team can focus on the business requirement—on archiving—rather than on making individual pieces work.

During initial system installation, your network and storage administrators work with Sun services personnel to define archive requirements—redundancy, retention, and availability—for the different types of data that you will be archiving. They jointly determine the numbers of files and capacity required for each data type, and the directory organization and user/client access appropriate to each. Once the requirements have been established, the deployment team configures archive policies to suit using dedicated IAS interfaces. The interfaces then adjust all of the underlying component settings to match the policies. Complex, repetitive configuration changes are made consistently, safely, and with minimal opportunity for human error.

Connecting the fully configured IAS appliance to users, applications, and client hosts is equally straightforward. The appliance design hides the internal configuration of the archive behind a standard, network file system interface that users, clients, and applications already use. To clients, the IAS archive looks like one or more standard SAMBA or NFS shares.

Manageability, reliability, and availability

The same factors that streamline solution deployment—proven components arranged in a standardized configuration, extensive automation, and user access via industry standard network file sharing—are the key to the ease of operation, reliability, and high availability of an IAS archive.

Once the appliance is installed, the internal operation of the archive is largely automatic. You can read from and write to files and directories, control access, and administer resources using normal, network-administration tools, just as you would for any other NFS or SAMBA share. The system implements your policies behind the scenes, automatically creating and maintaining the specified number of copies of each file in specified repositories, retaining each file for the required interval, migrating files between repositories, and managing space in each repository. You need only monitor the IAS internals for problems, using the browser-based, IAS graphical administration console and IAS email notifications.

Automation is the key factor in insuring the long-term stability and availability of an IAS archive. Human errors—subtle, inadvertent misconfigurations—are the most common causes of data loss, poor performance, and poor system availability in archiving systems. In a complex system, in the absence of standardization and automation, such errors are easy to make and hard to diagnose and correct. Heavy administrative workloads increase the incidence of problems. By design, IAS appliances limit the ways in which operators, administrators, and archive users can interact with the system, so that you have to make fewer changes, have less to do, and have fewer opportunities for making costly mistakes.

When you do have to make changes—either to data or, in the rare instances, to the underlying configuration—automated IAS interfaces insure that changes are tracked and applied correctly. Users, administrators, and service personnel each interact with a defined set of control points—exposed via NFS or SAMBA shares or via the IAS graphical administration interface—and not with the specially tuned operating system, optimized internal file systems, and preconfigured hardware interfaces. The defined control points are designed to be flexible and feature-rich enough to provide any useful degree of customization, while insuring that highly integrated archive configuration remains intact.

Proven components

The IAS appliance’s specialized architecture rests on proven Sun foundations:

- [SAM for information lifecycle management](#)
- [QFS for fast, reliable access to files](#)
- [Solaris Cluster for high availability](#)

SAM for information lifecycle management

Sun StorageTek SAM is a hierarchical storage manager that implements fully automated, policy-driven, information lifecycle management behind a familiar, straightforward file system interface. When a user or application saves a file to a SAM file system, software automatically works in the background to insure that the data remains available throughout the retention period specified. The software makes multiple copies of the data and distributes the copies across multiple, tiered storage systems, so that the accessibility and responsiveness of the storage provide a quality and cost of service commensurate with the value and importance of the data.

The tiers of a SAM storage subsystem match the input/output (I/O) characteristics of storage devices and media to the usage patterns of the data and the I/O characteristics of the user and application processes that create, edit, and consume data. The first tier, *primary cache*, holds working copies of the data files that users and applications are actively accessing and modifying. Primary storage media must

thus reliably sustain many random I/O operations per second and many concurrent I/O streams from multiple host systems. This high performance comes at a proportionately high price that effectively limits the capacity of the primary tier, and, in any case, only a subset of a typical organization's data is in active use at any given time. For these reasons, the *secondary* and, if present, *tertiary storage* tiers of a SAM hierarchy are normally implemented using higher-capacity, lower-cost media that can be economically provided in the required quantities.

In enterprise-class archive systems and in systems that must handle a highly transactional workloads or many small, rapidly changing files—such as many databases and most email systems—Sun engineers specify a primary cache consisting of high-performance, low-density, iSCSI or Fibre Channel disk arrays. In smaller, two-tier archive systems and in systems that handle large, for the most part sequentially accessed files—such as medical imagery, geological data, and engineering drawings—they select moderately priced, high-density SATA disk arrays that are more cost-effective for these kinds of workload. In the enterprise and I/O-intensive archive solutions, SATA disk arrays form the secondary storage tier and are backed by a third tier of inexpensive, high-density, sequential-access tape media stored in a robotic library. Systems that use SATA disk as primary cache can use either a tape or SATA disk for secondary storage, depending upon the requirement, and may or may not include a third tier.

The SAM software migrates copies of each data file from tier to tier based on file usage and user-defined policies. Actively accessed, frequently changed data is kept in the primary cache and backed by multiple, redundant copies in the other storage tiers. As the file ages and is less frequently accessed, SAM expires the primary copy in order to free primary cache for more active data, effectively migrating the file to the secondary or tertiary tier. If a user or process at some point needs to start accessing a migrated file, SAM stages the file back to primary cache and retains it there until it is no longer being used.

The complex, information lifecycle management processes are, however, invisible to users, applications, and hosts. The archive presents one copy of each file, resident in an ordinary file system, a file system that just supplies what is needed when it is needed.

QFS for fast, reliable access to files

QFS is a high-performance file system that facilitates rapid location of and access to files across multiple media types, while at the same time enhancing the redundancy and resilience of the file system itself. QFS file systems store file data and filesystem metadata (the information necessary for locating and accessing files) on separate storage devices for optimum security, scalability, and performance.

Solaris Cluster for high availability

When assured access to data is essential, an IAS appliance can combine multiple SAM-QFS server nodes in a high-availability cluster managed by proven Sun Solaris Cluster software. An active *Meta Data Server (MDS)* node manages QFS metadata devices and supplies filesystem metadata to the other, *QFS client* nodes. All of the QFS file servers—both MDS and client—serve files to user application clients. But, if the MDS becomes unavailable, the system automatically fails over to a client node, which then becomes the active MDS. Moreover, if maintenance or repairs necessitate a shutdown of the active MDS, you can manually fail over to a client before proceeding and thus keep the archive continuously available .

IAS appliance best practices

An IAS appliance is at its best when you leave it alone and think of it as a sealed box, like a toaster. You do not open up your toaster and alter the wiring when you want toast: you simply drop the bread in and push down the handle provided. You may adjust the timer control in order to get lighter or darker toast. But, once you have the setting the way you like it, you generally leave it alone, in order to get consistently satisfactory results. When you use an IAS appliance, you should likewise use the external controls provided and limit the adjustments that you make as much as possible. This section outlines some best practices:

- [“Use subdirectories and network shares as virtual filesystems”](#) (see page 5)
- [“Use subdirectories and file sharing for performance tuning”](#) (see page 6)
- [“Always consult Sun services before altering IAS internals”](#) (see page 6)
- [“Always use the IAS Java Web Console when making changes”](#) (see page 7)
- [“Let IAS software manage the storage”](#) (see page 7).

Use subdirectories and network shares as virtual filesystems

Most of us are used to organizing and controlling access to data by placing different data on different servers, different disks, or different file systems. You can—and should—follow a similar rationale when you administer an IAS archive. However, you should never alter the preconfigured file systems that come with the IAS appliance without explicit instructions from a Sun support representative. Instead, you should create subdirectories within these standard file systems and then share them with NFS and/or SAMBA.

The automated processes that give an IAS appliance many of its special properties will only work with the standard file systems and standard filesystem names (Tx_y, where *x* and *y* are digits). So you should consider these file systems part of the internal, physical implementation of the system. Leave them alone.

On the other hand, your users, applications, and client systems require multiple file systems, ideally with informative, meaningful names and varying access requirements. To accommodate these requirements, create user-facing, *virtual* file systems on top of the physical IAS implementation. Using the IAS Java Web Console, create a subdirectory for each filesystem that you need within the IAS filesystem (Tx_y). Name each *subdirectory* in a way that makes sense *to you and your fellow administrators*. Then share each subdirectory using either NFS or SAMBA. Set up access control for each share just as you would when sharing any other networked storage resource. Name each *share* in a way that makes sense *to your users or applications*. You then have an archive structure that is both easy for you and your colleagues to manage and easy for your users to understand.

By virtualizing your user- and client-facing file systems in this way, you increase the usability of the archive while simultaneously hardening it against both inadvertent and intentional interference. Both the end users and the system administrators can work with share names that are convenient and meaningful. Yet the IAS standard file systems have not been altered. As a bonus, the virtual, client-side names mask the actual, physical directory structure—a modest but useful extra protection against intruders.

Use subdirectories and file sharing for performance tuning

As with any filesystem, performance may seem to fall over time, as you add files to an IAS archive. To restore or maintain, you simply modify your subdirectory structure and network shares (do not alter the internal configuration of your appliance!). Performance can fall off when a SAM-QFS subdirectory contains 10,000 or more files. So distribute user and client data across multiple subdirectories, each with fewer files. Use access-controlled NFS and/or SAMBA shares to balance loads by channelling files to the separate subdirectories.

Always consult Sun services before altering IAS internals

If you have a requirement that you cannot address using subdirectories and network file sharing or if you must correct an internal fault, contact Sun services unless you are absolutely sure of what you are doing, have the requisite training, and have mastered the full range of SAM-QFS and (where appropriate) Sun Solaris Cluster documentation.

Always use the IAS Java Web Console when making changes

If you must alter the internal configuration of the IAS appliance, always do so using the browser-based IAS Java Web Console administrative interface, unless directed by a Sun services representative. Never use the IAS Solaris operating system, Integrated Lights Out Management interface, server- or disk-management utilities, or individual hardware interfaces as substitutes for the IAS administrative console, even if you are intimately familiar with general-purpose Sun Solaris servers and storage systems. IAS components are specially configured and tuned for use in an integrated appliance, and may not function correctly when individually reconfigured.

Note – while you can view native CAMS and FSM interfaces by pressing the graphical APPLICATIONS button at the top of the IAS Java Web Console banner, you should never use these applications for appliance configuration unless specifically instructed to do so by a Sun support representative. In particular, CAMS and FSM are not cluster-aware applications and cannot manage the clustered components of IAS appliances.

Let IAS software manage the storage

In general, give the IAS software processes free access to all of the storage and let IAS decide how storage and media are apportioned and utilized. Do not try to micromanage or second-guess the archive. Archiving always works best when you keep things simple and let the software choose from the largest possible selection of storage locations when processing any given job.

In most circumstances, you should thus avoid reserving file systems and/or Volume Serial Number (VSN) pools for particular users, groups, or purposes. When you divide up the space statically and reserve each piece for a specific use, unused space in one subdivision cannot be combined with space in the others. Overall capacity is under-utilized, yet archiving may stop for lack of media. There is more than enough aggregate storage tied up in the different subdivisions, but too little is available in the designated pool or file system.

Changing root passwords

Before you begin using the Infinite Archive System (IAS) appliance, secure the server nodes by changing the default `root` passwords. Replace the defaults with secure passwords of your choosing.

Secure passwords contain eight characters, of which at least two are alphabetic characters and at least one is a numeric or special character. Each new password differs significantly from the corresponding account name (changing the case of letters is not sufficiently different) and differs from the old/default password by at least three characters.

There are two `root` passwords that you must change on each server node:

- the `root` password for the Integrated Lights Out Manager (ILOM)
- the `root` password for the Sun Solaris operating system.

The Sun Integrated Lights Out Manager (ILOM) is a service processor that provides management access to IAS server nodes even when the server CPU and Solaris operating system are not running. It consists of an independent processor called the Graphics Redirect and Service Processor (GRASP), firmware (software stored in non-volatile memory), and serial and Ethernet communications interfaces. The ILOM firmware includes a full set of basic communications and management applications, including a server-side Secure Shell (ssh) implementation for encrypted remote login, a commandline interface (CLI), an integral web server, and a downloadable, browser-based graphical user interface (GUI).

You can change both the ILOM and Solaris `root` passwords using the ILOM, as described in the remainder of this chapter. There are three basic tasks:

- [“Connecting the service host” on page 10](#)
- [“Changing the ILOM root password” on page 15](#)

Connecting the service host

To use the ILOM, you must connect a host to the IAS appliance and log in. This service host can be a laptop or workstation running any operating system that supports Java and a compatible web browser.

Sun recommends the free, open-source Firefox browser for accessing the Java Web Console graphical user interfaces. Firefox is available from the Mozilla Foundation at <http://www.mozilla.com/>. Other standards-compliant browsers may also work with the IAS console, but have not been tested.

The Microsoft Internet Explorer browser has been tested and is *not* compatible with the IAS console application.

Note that you must enable popup windows when using Java Web Console applications. Consult your browser documentation for instructions.

Proceed as follows.

▼ Attaching a host to the ILOM of a dual-node appliance

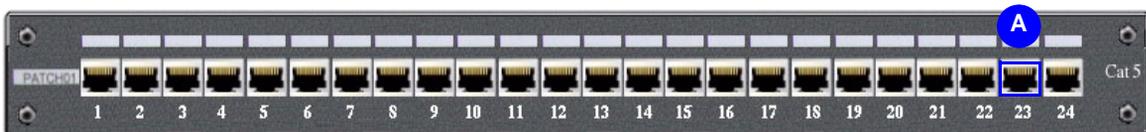
The ILOM service processors of dual-node, high-availability IAS appliances are accessible via port 23 on the IAS patch panel. To attach a management host to this port, proceed as follows.

1. **Arrange the required physical Ethernet connections between the host and the IAS appliance.**
2. **On the host, configure the IP address of the Ethernet port that connects to the IAS appliance. Set it to 192.168.51.251.**

This is IAS private-network address that is reserved for customer use.

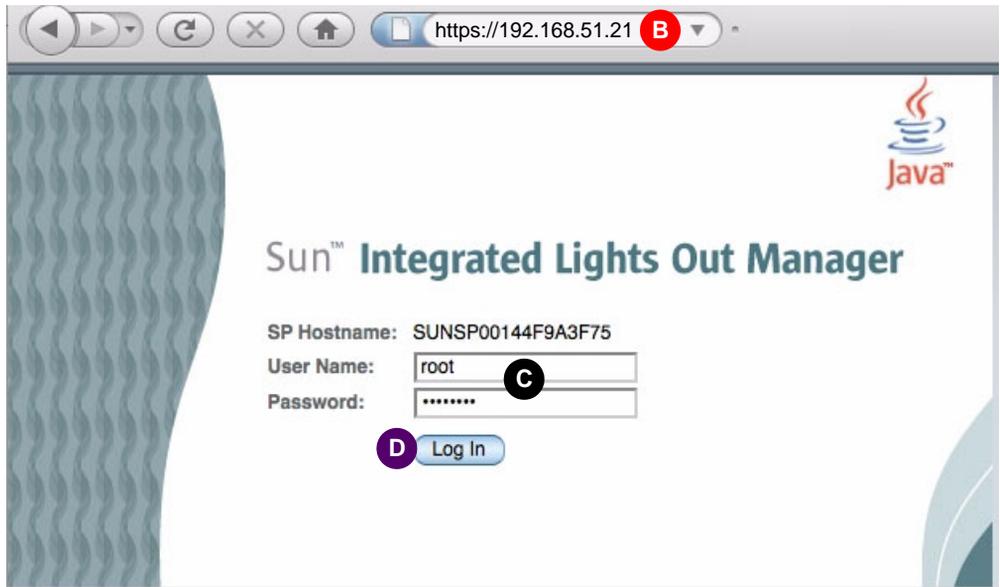
3. **On the IAS appliance, connect the Ethernet cable from the host to port 23 on the IAS patch panel (A below).**

Port 23 is reserved for customer-supplied hosts.



4. **Test the connection. On the host, open a compatible web browser, such as Mozilla Firefox.**

5. In the navigation menu of the browser (B below), enter 192.168.51.21, the IP address of the ILOM interface on the SAM-QFS Meta Data Server (MDS) node, IAS3N-n01:



6. When the login page appears, log in as root, and enter the ILOM root password (C above).

Remember that the ILOM root password is not the same as the Solaris root password.

7. Press **Log In (D above)** to open the **ILOM interface (below)**.

ABOUT REFRESH LOG OUT

Role (User): Administrator (root) SP Hostname : SUNSP00144F9A3F75

Sun™ Integrated Lights Out Manager

Java™

- System Information
- System Monitoring
- Configuration
- User Management
- Remote Control
- Maintenance

Versions Session Time-Out Components

Versions

View the version of ILOM firmware currently in use.

Version Information	
Property	Value
SP Firmware Version	1.1.8
SP Firmware Build Number	18244
SP Firmware Date	Sat Apr 21 11:25:11 PDT 2007
SP Filesystem Version	0.1.14

Next step: [“Changing the ILOM root password” on page 15.](#)

▼ **Attaching a host to the ILOM of a single-node appliance**

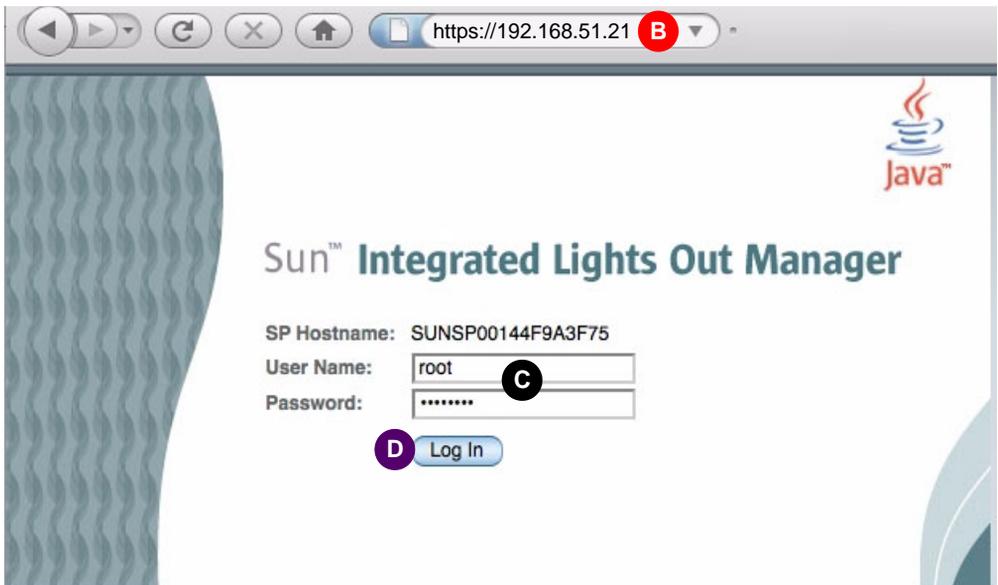
The ILOM service processor of single-node IAS appliances are accessible via the **Net Mgt** port on the lower edge of the **System Controller** at the rear of the server chassis. To attach a management host to this port, proceed as follows.

- 1. Arrange the required physical Ethernet connections between the designated host and the IAS appliance.**
- 2. If the host and the ILOM will be attached to the local area network (LAN) using static addressing, obtain IP addresses, net masks, and gateway addresses for both.**
- 3. If the host will be will directly attached to the ILOM using a private connection, configure an Ethernet port on the host to 192.168.51.251, the default IAS private address used with customer equipment.**

4. On the IAS appliance, connect the Ethernet cable coming from the host or network switch to the Net Mgt port on the lower edge of the System Controller at the rear of the server chassis (A below).



5. On the host, open a compatible web browser, such as Mozilla Firefox.
6. In the navigation menu of the browser (B below), enter 192.168.51.21, the default IP address of the ILOM interface on the IAS server node:



7. When the login page appears, log in as `root`, and enter the ILOM `root` password (C above).

Remember that the ILOM `root` password is not the same as the Solaris `root` password.

8. Press Log In (D above) to open the ILOM interface (below).

ABOUT REFRESH LOG OUT

Role (User): Administrator (root) SP Hostname : SUNSP00144F9A3F75

Sun™ Integrated Lights Out Manager

System Information System Monitoring Configuration User Management Remote Control Maintenance

Versions Session Time-Out Components

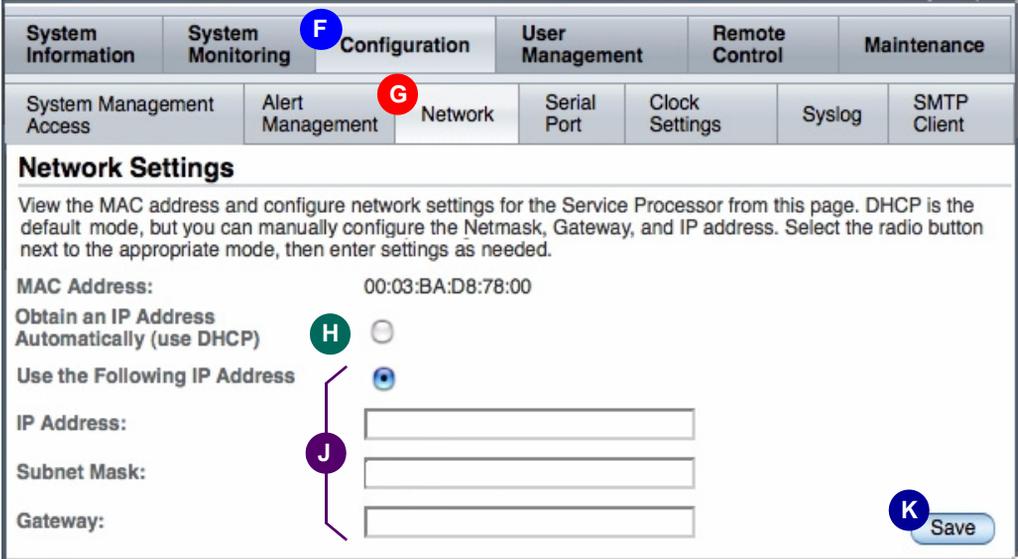
Versions

View the version of ILOM firmware currently in use.

Version Information	
Property	Value
SP Firmware Version	1.1.8
SP Firmware Build Number	18244
SP Firmware Date	Sat Apr 21 11:25:11 PDT 2007
SP Filesystem Version	0.1.14

9. If the host is directly attached to the ILOM using a private connection or if the ILOM network port has already been configured for the LAN, you have finished configuring the ILOM connection. Stop here.
10. If the host and the ILOM are attached to the local area network (LAN) and if the , ILOM network interface is not currently configured, proceed as follows.

11. Select the **Configuration** tab from the system properties (F below).



The screenshot shows a web-based configuration interface. At the top, there are several tabs: System Information, System Monitoring, Configuration (marked with a blue 'F'), User Management, Remote Control, and Maintenance. Below these, there are more specific tabs: System Management Access, Alert Management, Network (marked with a red 'G'), Serial Port, Clock Settings, Syslog, and SMTP Client. The main content area is titled 'Network Settings'. It contains a paragraph of text explaining the purpose of the page. Below the text, there are several fields: 'MAC Address' with the value '00:03:BA:D8:78:00', 'Obtain an IP Address Automatically (use DHCP)' with a selected radio button (marked with a green 'H'), 'Use the Following IP Address' with an unselected radio button, and three text input fields for 'IP Address', 'Subnet Mask', and 'Gateway' (marked with a purple 'J'). A blue 'Save' button (marked with a blue 'K') is located at the bottom right of the form.

12. When configuration property tabs appear, select **Network** (G above).

The **Network Settings** properties appear.

13. To use a dynamically assigned IP address, click the radio button labeled **Obtain an IP Address Automatically (use DHCP)** (H above).

14. Otherwise, click the radio button labeled **Use the Following IP Address** and enter the **IP Address, Subnet Mask, and Gateway** address in the text boxes provided (J above).

15. Press the on-screen **Save** button to make your changes (K above).

Next step: [“Changing the ILOM root password” on page 15.](#)

Changing the ILOM root password

ILOM service processors have user accounts much like the servers they support. Privileges associated with these accounts govern the degree to which users can control the server remotely. To change the default ILOM root password, proceed as follows.

▼ Changing ILOM passwords

1. If you have not already done so, log in to the ILOM graphical user interface.
2. Click **User Management** (A below) on the ILOM main menu.

ABOUT REFRESH LOG OUT

Role (User): Administrator (root) SP Hostname : SUNSP00144F9A3F75

Sun™ Integrated Lights Out Manager

Java™

System Information System Monitoring Configuration **User Management** (A) Remote Control Maintenance

User Accounts Active Sessions LDAP Settings Radius

User Settings

Add, delete, or modify local ILOM user accounts from this page. ILOM offers 10 local user accounts. The system uses *root*, which you cannot delete. The other 9 are available for your use. Single Sign On enables an ILOM administrator, logged in to the Chassis Monitoring Module, to access any blade Service Processor without the need to log in again. Single Sign On must be enabled on each Service Processor you want to access with this feature, which is enabled by default.

Enable Single Sign On

Save

Users

Add Edit Delete

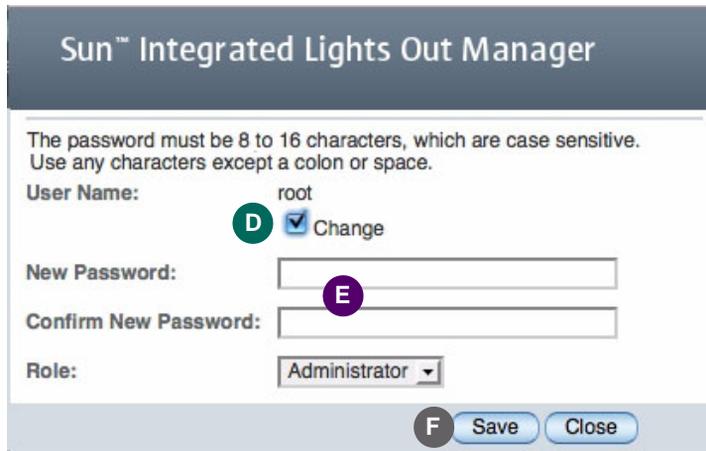
Name	Role
root	Administrator

(B) (C)

3. In the **Users** table, click the radio button corresponding to the account that you want to change (B above), and press the on-screen **Edit** button (C).

A password-maintenance dialog appears.

4. In the dialog, check the `Change` checkbox (D below).



The dialog box is titled "Sun™ Integrated Lights Out Manager". It contains the following text and controls:

- Instruction: "The password must be 8 to 16 characters, which are case sensitive. Use any characters except a colon or space."
- User Name: "root" (with a green circle 'D' next to it)
- Change checkbox: checked (with a blue circle 'D' next to it)
- New Password: empty text field (with a purple circle 'E' next to it)
- Confirm New Password: empty text field (with a purple circle 'E' next to it)
- Role: "Administrator" (dropdown menu)
- Buttons: "Save" and "Close" (with a grey circle 'F' next to the Save button)

5. Enter and confirm the new password in the fields provided (E above).
6. Press `Save` to make the change, `Close` to abort (F above).

Next step: ["Changing the Solaris root password" on page 18](#)

Changing the Solaris root password

To change the default Solaris root password on a server node, carry out the following tasks:

- [“Logging in to the Solaris commandline using ssh” on page 18](#) or, if you have just changed the ILOM password, [“Logging in to the Solaris Java Desktop via the ILOM” on page 19](#)
- [“Changing the root password from the commandline” on page 21.](#)

▼ Logging in to the Solaris commandline using ssh

1. Open a terminal window, and run ssh with the root account and the IP address of the IAS server:

```
# ssh root@IAS_server_IP
```

2. When you are challenged for a password, enter the root password for the IAS server:

```
# ssh root@IAS_server_IP
Password:
Last login: Fri Nov 21 09:57:43 2008 from 129.150.36.18
Sun Microsystems Inc. SunOS 5.10 Generic January 2005

                System BUILT: Thu Sep 18 14:27:01 MDT 2008

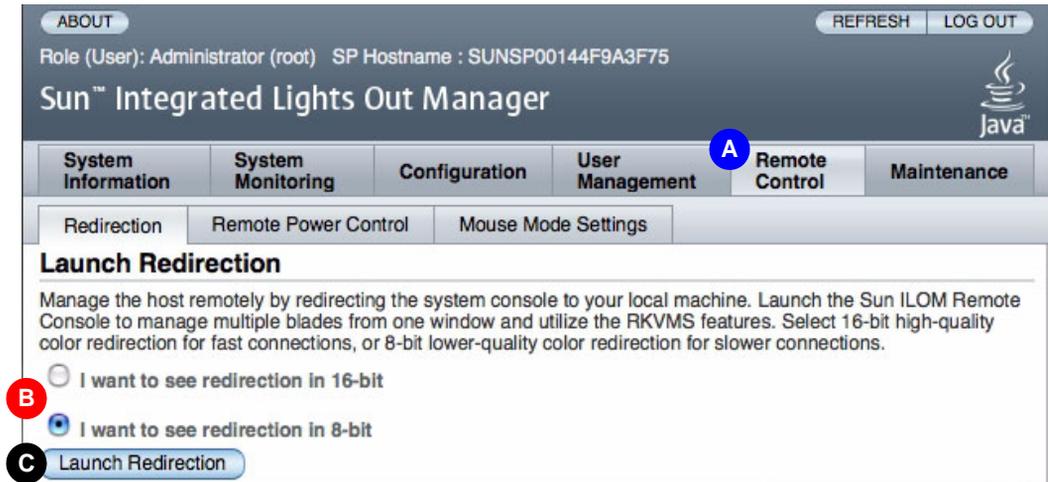
SAM-QFS Release:      4.6.50
Flash Archive Image:  3.5_1
SUNWcistk Version:    3.5.0.52

                CIS Model: CIS3-V-SAS-1-SATA-1-T1
You have new mail.
IAS3N-n0x#
```

Next task: [“Changing the root password from the commandline” on page 21.](#)

▼ Logging in to the Solaris Java Desktop via the ILOM

1. On the Integrated Lights Out Manager (ILOM) page, press the Remote Control button (A below).



2. On the Launch Redirection panel, click the radio button for 8- or 16-bit resolution (B above), and press Launch Redirection (C).

Choose 16-bit redirection for better resolution, 8-bit for better performance on a slow network.

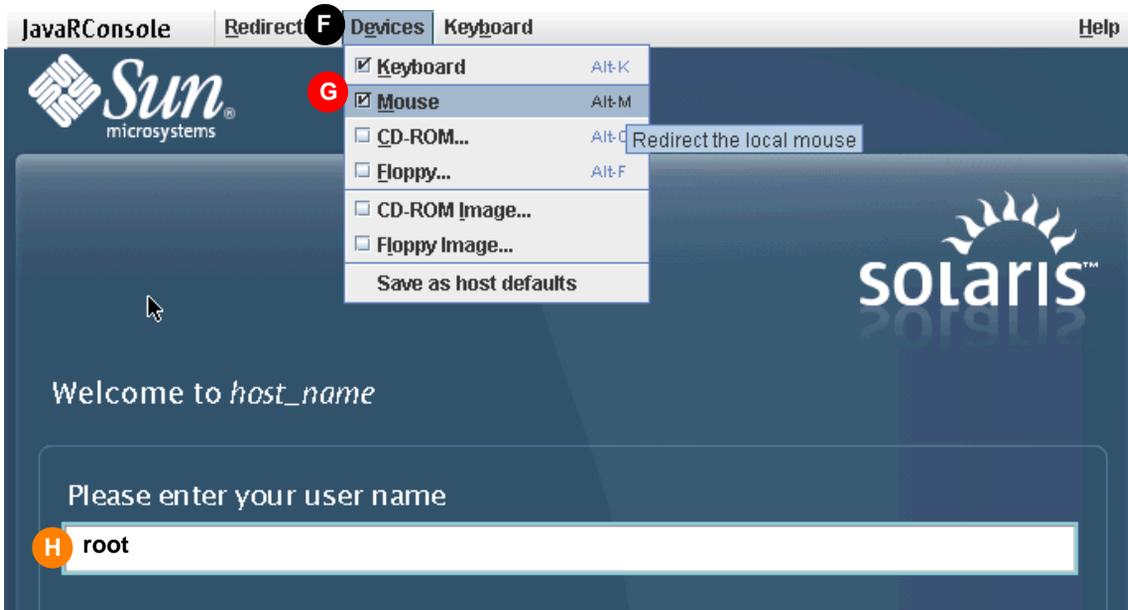
3. When the download dialog appears, click the Open with radio button, and select Java(TM) Web Start Launcher from the list (D below). Then press OK (E).



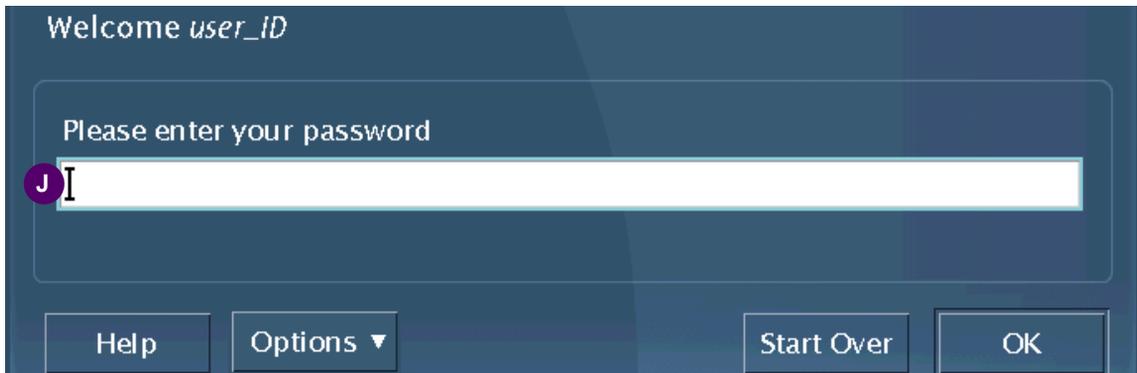
The file launches a Remote Console window holding a Solaris login panel.

4. When the Remote Console application interface appears, select **Devices** from the main menu (F below), and check the **Mouse** check box in the submenu (G).

You check the Mouse check box to enable mouse support for the Remote Console.



5. In the Please enter your user name field (H above) of the login panel, enter root.
6. When the Welcome root appears, enter the password for the Solaris root account in the Please enter your password field (J below).



When you enter your password, the Solaris desktop appears.

From the Sun Java Desktop, you can use Solaris graphical user interface applications or open Solaris terminal windows.

Next step: “Changing the root password from the commandline” on page 21.

▼ Changing the `root` **password** from the commandline

1. If you are working from the Solaris desktop, use the main menu to open a Solaris terminal window.
2. At the command prompt, enter the `passwd` command, as shown below:

```
[iasd03-n0x]# passwd
passwd: Changing password for root
New Password:
Re-enter new Password:
passwd: password successfully changed for root
```

3. When prompted, enter the `New Password` and then re-enter it to confirm your choice.

The `passwd` command confirms that the password has been changed.

Stop here.

Configuring network name services

Name services provide a directory that translates between human-readable host/domain names and machine-readable network addresses. Without this information, network services, such as email notifications and filesystem sharing, cannot function. Sun service personnel set up network name services during initial installation and configuration. But, as with any file server, you will need to manage and update name-services on the IAS appliance as changes are made on your network. This chapter explains how to use the Solaris [Name Service Switch \(NSS\)](#) when configuring an IAS appliance with two of the most common name-service protocols—Domain Name System (DNS) and Network Information Service (NIS).

For more information on Solaris 10 and name services, see [System Administration Guide: Naming and Directory Services \(DNS, NIS, and LDAP\)](#).

Name Service Switch (NSS)

The Sun Solaris *Name Service Switch (NSS)* is a configuration file, `/etc/nsswitch.conf`, that tells Solaris to redirect requests for name information from local configuration files on the host to specified, centralized configuration databases on the network. These databases supply a wide range of network configuration information, including aliases, boot parameters, Ethernet hardware addresses (MAC addresses), and information on groups, hosts, networks, netmasks, passwords, protocols, services, and email.

As a starting point, IAS comes with the following, service-specific templates:

- `nsswitch.dns` directs name queries to Domain Name System (DNS) databases.
- `nsswitch.nis` directs name queries to Network Information Service (NIS) databases.
- `nsswitch.nisplus` directs name queries to Network Information Service Plus (NIS+) databases.

- `nsswitch.ldap` directs name queries to Lightweight Directory Access Protocol (LDAP) databases.
- `nsswitch.files` directs name queries to local files.

To set up NSS, carry out the following tasks:

- [“Accessing Solaris on IAS appliances” on page 24](#)
- [“Setting name services on IAS appliances” on page 28.](#)

Accessing Solaris on IAS appliances

You access the Solaris operating system on the IAS server nodes remotely, either via a Secure Shell (`ssh`) commandline or via a web browser, using the Sun Integrated Lights Out Manager service processor on the node:

- [“Logging in to the Solaris commandline using `ssh`” on page 24](#)
- [“Logging in to the ILOM” on page 25](#) and [“Logging in to the Solaris Java Desktop via the ILOM” on page 26.](#)

▼ Logging in to the Solaris commandline using `ssh`

1. Open a terminal window, and run `ssh` with the root account and the IP address of the IAS server:

```
# ssh root@IAS_server_IP
```

2. When you are challenged for a password, enter the `root` password for the IAS server:

```
# ssh root@IAS_server_IP
Password:
Last login: Fri Nov 21 09:57:43 2008 from 129.150.36.18
Sun Microsystems Inc. SunOS 5.10 Generic January 2005

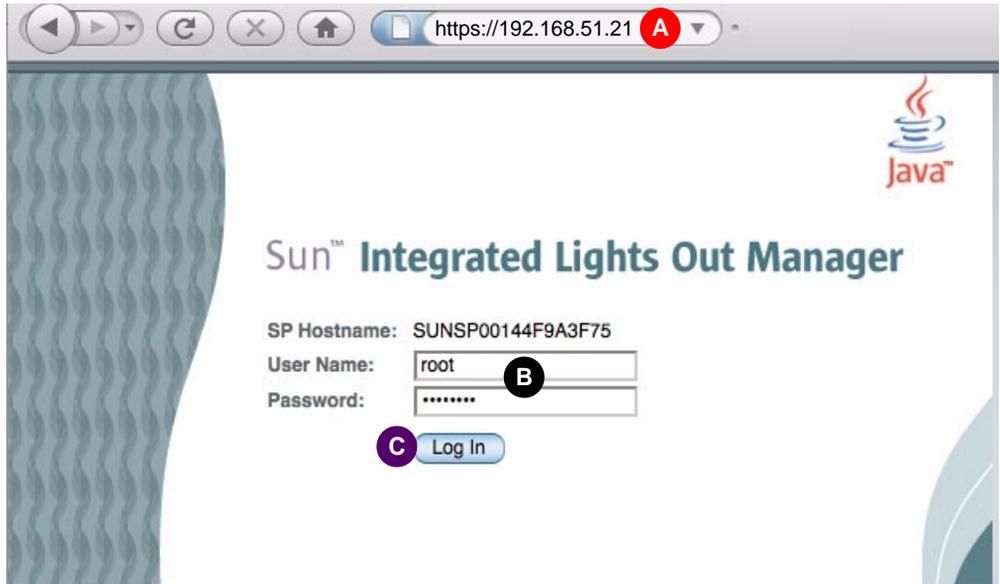
                System BUILT: Thu Sep 18 14:27:01 MDT 2008

SAM-QFS Release:      4.6.50
Flash Archive Image:  3.5_1
SUNWcistk Version:    3.5.0.52

                CIS Model: CIS3-V-SAS-1-SATA-1-T1
You have new mail.
IAS3N-n0x#
```

▼ Logging in to the ILOM

1. On the IAS management host, open a web browser session.
2. In the navigation menu of the web browser (A below), enter the IP address of the server node: 192.168.51.21 for node 01, 192.168.51.22 for node 02.

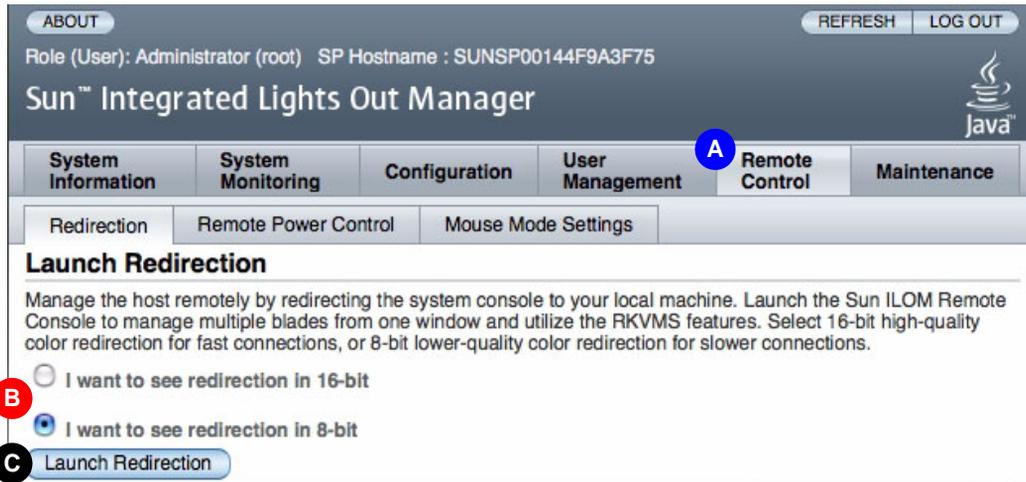


3. When the login page appears, log in as `root`, and enter the ILOM `root` password (B above). Press Log In (C) to open the ILOM interface.

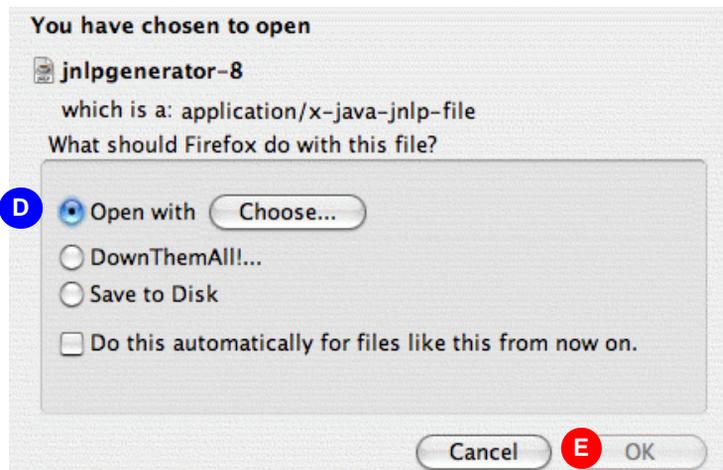
Next task: [“Logging in to the Solaris Java Desktop via the ILOM” on page 26.](#)

▼ Logging in to the Solaris Java Desktop via the ILOM

1. On the Integrated Lights Out Manager (ILOM) page, press the Remote Control button (A below).



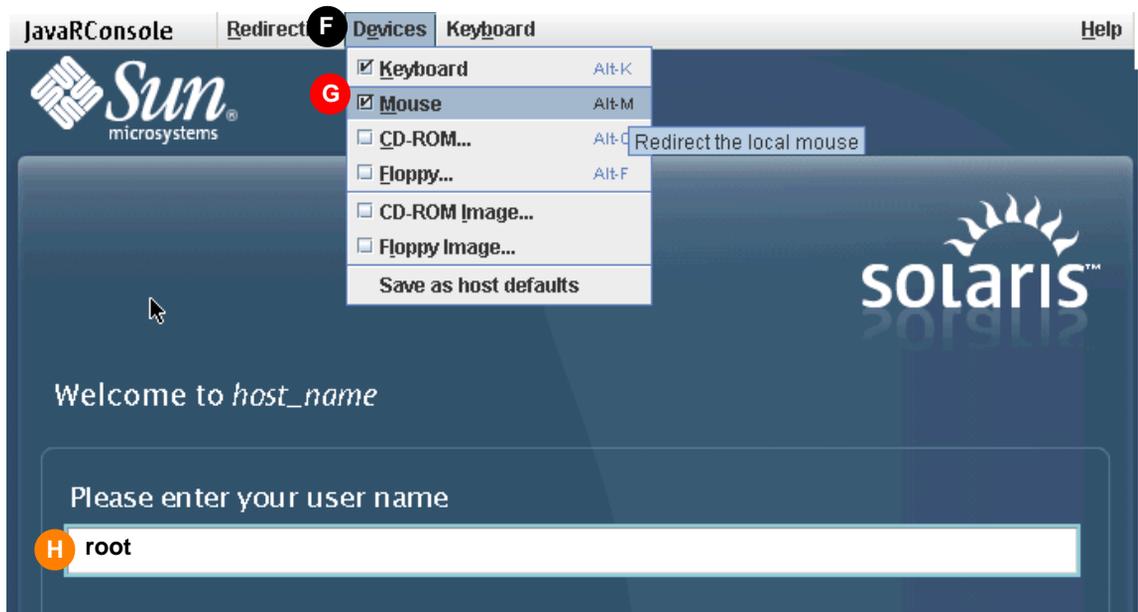
2. On the Launch Redirection panel, click the radio button for 8- or 16-bit resolution (B above), and press Launch Redirection (C).
Choose 16-bit redirection for better resolution, 8-bit for better performance on a slow network.
3. When the download dialog appears, click the Open with radio button, and select Java(TM) Web Start Launcher from the list (D below). Then press OK (E).



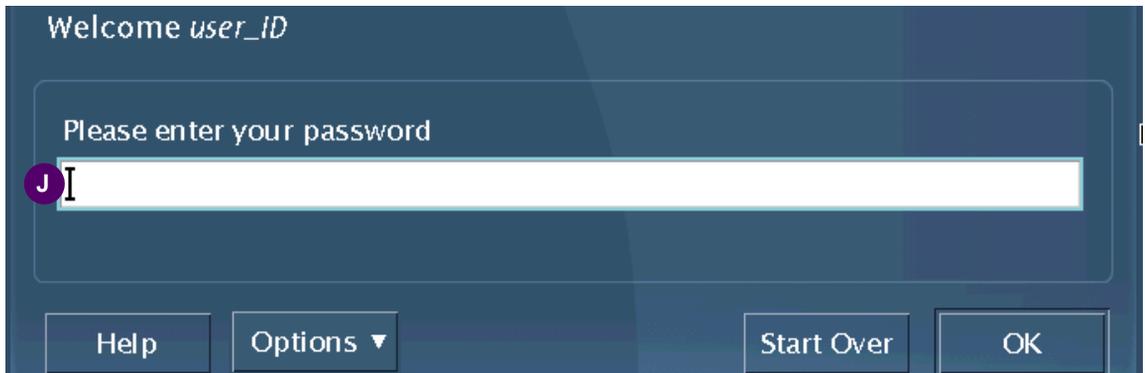
The file launches a Remote Console window holding a Solaris login panel.

4. When the Remote Console application interface appears, select Devices from the main menu (F below), and check the Mouse check box in the submenu (G).

You check the Mouse check box to enable mouse support for the Remote Console.



5. In the Please enter your user name field (H above) of the login panel, enter root.
6. When the Welcome root appears, enter the password for the Solaris root account in the Please enter your password field (J below).



When you enter your password, the Solaris desktop appears.

From the Sun Java Desktop, you can use Solaris graphical user interface applications or open Solaris terminal windows.

Next task: [“Setting name services on IAS appliances” on page 28.](#)

Setting name services on IAS appliances

Select the name service that you need to set up from the list below, and then go to the section indicated:

- [“Configuring Domain Name System \(DNS\)” on page 28](#)
- [“Configuring Network Information Services \(NIS\)” on page 29.](#)

Configuring Domain Name System (DNS)

To set up DNS, proceed as follows.

Caution – Be careful when editing IAS configuration files. On IAS servers nodes, the SAM-QFS default location for configuration files, `/etc/opt/SUNWsamfs/`, is a symbolic link to the IAS subdirectory `/global/etc/opt/SUNWsamfs`. IAS will not work properly unless the symbolic links are intact. So do not edit files in any other location! Make sure that you know what you are doing, and make back up copies using `cp -p` before editing files. Never use commands that will damage links, such as `mv`. If, at any point you are at all uncertain, please contact Sun support services for assistance.

▼ Setting up DNS services

1. In a terminal window, `ssh` to the IP address of the public-network management interface on the SAM-QFS Meta Data Server (MDS) node.

2. Change to the `/etc` directory:

```
IAS3N-n0x# cd /etc
IAS3N-n0x#
```

3. Back up the current `nsswitch.conf` file using a naming convention that makes sense (by adding a suffix representing today's date, for example):

```
IAS3N-n0x# cp -p nsswitch.conf nsswitch.conf.yyyymmdd
IAS3N-n0x#
```

4. Copy the DNS template file over the current `nsswitch.conf` file:

```
IAS3N-n0x# cp -p nsswitch.dns nsswitch.conf
IAS3N-n0x#
```

5. Specify the IP address for the default network router by echoing the default router address to the `/etc/defaultrouter` file:

```
IAS3N-n0x# echo defaultrouter_ip > /etc/defaultrouter
IAS3N-n0x#
```

6. In a text editor, such as `vi`, create a `resolv.conf` file for the network.

The file should look something like the example below

```
domain subdomain.example.com
nameserver best_name-server_ip-address
nameserver optional_2nd_name-server_ip-address
nameserver optional_3rd_name-server_ip-address
search subdomain.example.com othersubdomain.example.com
```

where:

- `domain` identifies the default domain that should be used when a hostname cannot be resolved on the first query.
- `nameserver` identifies the IP address of a name server.

The `resolv.conf` file can specify from one to three addresses. The node will try to connect to first on the list. If the name server cannot be reached, the node tries the remaining entries, if any, in the order given. So list the name servers in order of reliability, starting with the most reliable.

- `search` introduces a space- or tab-delimited list of alternate domains that should be searched if the hostname cannot be resolved.

7. Start the DNS client:

```
IAS3N-n0x# svcadm enable svc:/network/dns/client:default
```

8. From a terminal window, `ssh` to the IP address of the public-network management interface on the SAM-QFS client node.
9. Repeat steps 2-7 on the client.

Next task: If you have configured both IAS server nodes and need to set up DNS as a way of looking up NIS servers, go to [“Configuring Network Information Services \(NIS\)” on page 29](#). Otherwise, stop here.

Configuring Network Information Services (NIS)

NIS and NIS+ provide a simple network lookup service based on databases and client/server processes. The databases are stored on hosts that run an NIS server process. An NIS domain is the set of hosts that share a common set of databases called *maps*. During booting, clients join a domain by sending a request, obtaining a domain name, and then *binding* the domain name to the first NIS server host that replies.

To set up NIS or NIS+, you must carry out the following tasks:

- [“Setting up DNS services” on page 28](#) or [“Configuring /etc/inet/hosts files on NIS servers” on page 31](#)
- [“Configuring nsswitch.conf for use with NIS” on page 31](#)
- [“Setting up NIS and NIS+” on page 32.](#)

Caution – Be careful when editing IAS configuration files. On IAS servers nodes, the SAM-QFS default location for configuration files, `/etc/opt/SUNWsamfs/`, is a symbolic link to the IAS subdirectory `/global/etc/opt/SUNWsamfs`. IAS will not work properly unless the symbolic links are intact. So do not edit files in any other location! Make sure that you know what you are doing, and make back up copies using `cp -p` before editing files. Never use commands that will damage links, such as `mv`. If, at any point you are at all uncertain, please contact Sun support services for assistance.

▼ Configuring `/etc/inet/hosts` files on IAS server nodes

The hosts file is a local database that associates the names of hosts with their Internet Protocol (IP) addresses. It serves as a local alternative to name services such as DNS and LDAP.

1. If you plan to use DNS for looking up NIS servers and have not already done so, configure DNS before proceeding.
2. Otherwise, from a terminal window, `ssh` to the IP address of the SAM-QFS Meta Data Server (MDS) node.
3. Change to the `/etc/inet/` directory, and back up the current `hosts` file using a naming convention that makes sense (such as adding a suffix representing today’s date to the filename):

```
IAS3N-n0x# cd /etc/inet
IAS3N-n0x# cp -p hosts hosts.yyyymmdd
```

4. Open the `hosts` file in a text editor, such as `vi`, and append an entry for the NIS server at the end of the file.

An entry consists of an IP address in the form `w.x.y.z` and a *hostname*, separated by one or more spaces or tabs:

```
nis-server_ip-address      nis-server_hostname
```

5. Save the edited `hosts` file, and exit the editor application.
6. From a new terminal window, `ssh` to the IP address of the public-network management interface on the SAM-QFS client node.
7. Repeat steps 2-7 on the SAM-QFS client node.

Next task: [“Configuring /etc/inet/hosts files on NIS servers” on page 31.](#)

▼ Configuring /etc/inet/hosts files on NIS servers

The hosts file is a local database that associates the names of hosts with their Internet Protocol (IP) addresses. It serves as a local alternative to name services such as DNS and LDAP.

1. Log in to each NIS server, either directly, from the NIS server’s console, or using `ssh` from a host that is on the same network.
2. On each NIS server, change to the `/etc/inet/` directory, and back up the current hosts file using a naming convention that makes sense (such as adding a suffix representing today’s date to the filename):

```
[nis-server_hostname]root# cd /etc/inet
[nis-server_hostname]root# cp -p hosts hosts.yyyymmdd
```

3. On each NIS server, open the `hosts` file in a text editor, such as `vi`, and, at the end of the file, append entries for the IAS public network interfaces.

An `/etc/inet/hosts` entry consists of an IP address and a hostname, separated by one or more spaces or tabs. Here, each entry must include one of the customer-supplied IP addresses that you entered when you ran the IAS configuration wizard and the NIC name of the corresponding IAS virtual network interface.

For each IAS data network, you thus need to create entries for two virtual network interfaces. For the management network, you need an entry for one virtual interface:

```
public_ip-address_1    iast01-net1    # DATA01 virtual-IP 1
public_ip-address_2    iast01-net3    # DATA01 virtual-IP 2
public_ip-address_3    iast01-net2    # DATA02 virtual-IP 1
public_ip-address_4    iast01-net4    # DATA02 virtual-IP 2
public_ip-address_5    iast01         # MGMTNET virtual-IP 2
```

4. On each NIS server, save the edited `hosts` file, and exit the editor application.
5. Repeat steps 1-4 until you have edited the host files on all the NIS servers to which the IAS nodes may bind.

Next task: [“Configuring nsswitch.conf for use with NIS” on page 31.](#)

▼ Configuring nsswitch.conf for use with NIS

1. In a terminal window, `ssh` to the IP address of the SAM-QFS Meta Data Server (MDS) node.
2. Change to the `/etc` directory:

```
IAS3N-n0x# cd /etc
```

3. **Back up the current `nsswitch.conf` file using a naming convention that makes sense (such as adding a suffix representing today's date to the file name):**

```
IAS3N-n0x# cp -p nsswitch.conf nsswitch.conf.yyyyymmdd
```

4. **If you are setting up NIS, copy the NIS template file over the current `nsswitch.conf` file:**

```
IAS3N-n0x# cp -p nsswitch.nis nsswitch.conf
```

5. **If you are setting up NIS+, copy the NIS+ template file over the current `nsswitch.conf` file:**

```
IAS3N-n0x# cp -p nsswitch.nisplus nsswitch.conf
```

6. **In a new terminal window `ssh` to the IP address of the public-network management interface on the SAM-QFS client node.**

7. **Repeat steps 2-5 above on the SAM-QFS client node.**

Next step: If you have only configured `nsswitch.conf` on the first IAS server node, repeat [“Configuring `nsswitch.conf` for use with NIS” on page 31](#) on the second node, `iasd0x-n02`. Otherwise, go to [“Setting up NIS and NIS+” on page 32](#).

▼ Setting up NIS and NIS+

1. **In a terminal window, `ssh` to the IP address of the of the SAM-QFS Meta Data Server (MDS) node.**

2. **Change to the `/etc/` directory.**

```
IAS3N-n0x# cd /etc
```

3. **Set the Meta Data Server (MDS) node's NIS domain name using the `domainname` command with the argument `nis.example.domain`**

where `nis.example.domain` is the name of the correct NIS domain:

```
IAS3N-n0x# domainname nis.example.domain
```

4. **Verify the change by using the `domainname` command without an argument.**

The `domainname` command displays the name of the domain:

```
IAS3N-n0x# domainname
nis.example.domain
```

5. **Redirect the output of the `domainname` command to the `defaultdomain` configuration file:**

```
IAS3N-n0x# domainname > defaultdomain
```

6. **Start the Solaris NIS client software using the command `ypinit` with the `-c` parameter:**

```
IAS3N-n0x# ypinit -c
```

When it starts, `ypinit` prompts for the list of NIS servers to which the client should bind. The prefix `yp` derives from the original name for NIS, *Yellow Pages*. Trademark issues forced the abandonment of the name itself, but the abbreviation lives on in the names of NIS commands. `-c` stands for *client*.

7. Add the names of the NIS servers in order of preference, one per line, until all servers have been listed.

Generally, you should enter NIS servers starting with the closest server (in network terms) and working down to the farthest server.

```
IAS3N-n0x# ypinit -c
next host to add: nis-server_hostname0
...
next host to add: nis-server_hostnameN
next host to add:
```

8. When all NIS servers have been listed, enter `Ctrl-D` or enter a blank line.

The `ypinit -c` command lists the current set of NIS (`yp`) servers, and asks you to confirm that it is correct:

```
IAS3N-n0x# ypinit -c
next host to add: nis-server_hostname0
...
next host to add: nis-server_hostnameN
next host to add:
The current list of yp servers looks like this:
nis-server_hostname0
...
nis-server_hostnameN
Is this correct? [y/n: y] y
```

9. At the `Is this correct? [y/n: y]` prompt, enter `y` (for yes) if all is well.

10. Enable the NIS `svc:/network/nis/client:default` service using the `svcadm enable` command:

```
IAS3N-n0x# svcadm enable svc:/network/nis/client:default
```

11. If your network requires any of the optional NIS services, enable them using the `svcadm enable` command:

```
IAS3N-n0x# svcadm enable svc:/network/nis/passwd:default
IAS3N-n0x# svcadm enable svc:/network/nis/update:default
```

12. In a terminal window, `ssh` to the IP address of the SAM-QFS client node, and repeat steps 2-11.

Next step: If you have only configured NIS on the first IAS server node, repeat [“Setting up NIS and NIS+” on page 32](#) on the second node, `iasd0x-n02`. Otherwise, stop here.

Using the IAS Java Web Console GUI

The Sun Customer Ready Infinite Archive System (IAS) graphical user interface is a standard Java Web Console that runs inside a standard web browser, so that you control remote systems using pages and links that are much like those served up by ordinary web sites. This chapter introduces the characteristics and use of the IAS console:

- “Supported web browsers” on page 35
- “Navigating the IAS Java Web Console” on page 36
- “Understanding IAS operational roles” on page 41
- “Understanding role-based user accounts” on page 44
- “Using the IAS help system” on page 44.

Supported web browsers

Sun recommends the free, open-source Firefox browser for accessing the Java Web Console graphical user interfaces. Firefox is available from the Mozilla Foundation at <http://www.mozilla.com/>. Other standards-compliant browsers may also work with the IAS console, but have not been tested.

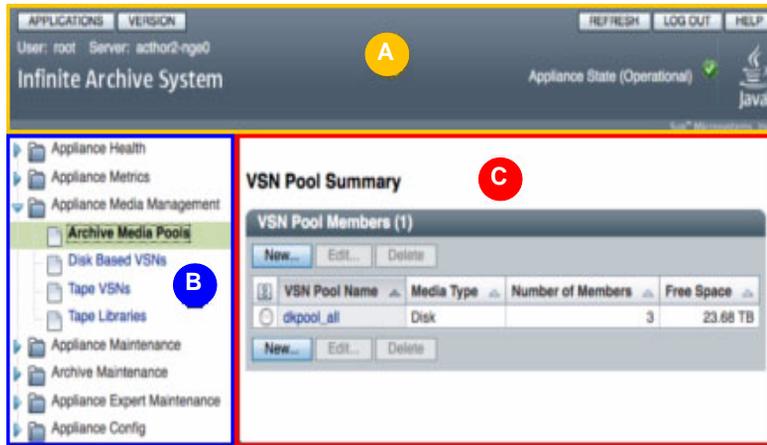
The Microsoft Internet Explorer browser has been tested and is *not* compatible with the IAS console application.

Note that you must enable popup windows when using Java Web Console applications. Consult your browser documentation for instructions.

Navigating the IAS Java Web Console

The interface divides each browser window into three panes:

- a *banner* that displays system status information (**A** below)



- a *task navigation tree* that lets you locate controls and displays (**B** above)
- a *content pane* that displays system information and hosts control panels (**C**).

The task navigation tree and content pages are organized around a set of standard user roles, each with an associated privilege level. This approach ensures that users always have the capabilities needed for any given job, at the same time, isolating potentially dangerous functionality that is not normally required.

The functions and features of the three console panes are discussed in more detail below.

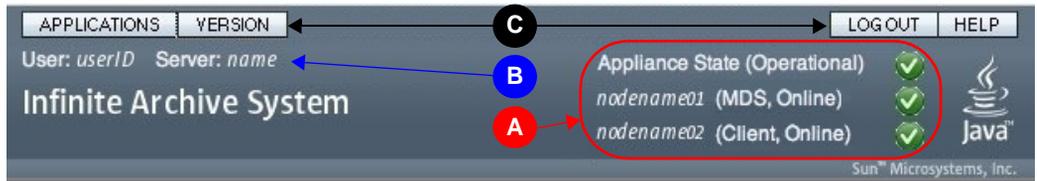
The status banner and dashboard

The banner is the area that spans the top of the page. It displays the name of the application, along with summary status information and controls for the interface. Each of the following parts of the banner displays a specific kind of information:

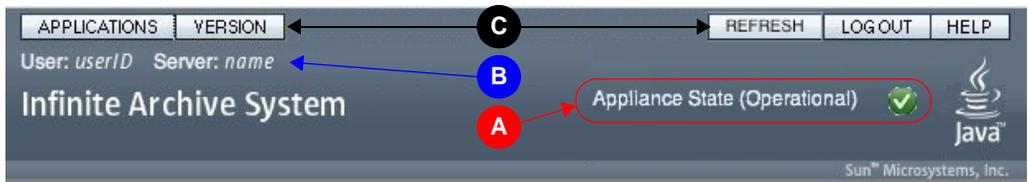
- [The dashboard](#)
- [The server information area](#)
- [The button bar](#)

The dashboard

The *dashboard* part of the banner displays the `Appliance State` (A below). On a clustered, multinode server, the dashboard displays the status of the `Appliance` as a whole, the status of the `MDS` (MetaData Server) node, and the status of the `Client` (standby) node:



On an unclustered, single-node system, the dashboard displays only the `Appliance State`, since the state of the server node is the state of the appliance:



The following text and graphics describe appliance and component states:

-  means that the appliance or node is operational or online;
-  warns you of a problem;
-  alerts you to a critical error.

The server information area

The upper left side of the banner displays the `Server` where you are currently logged in and your current `User` identity (B above).

The button bar

Control buttons line the top of the banner (C above):

- The `APPLICATIONS` button returns you to the parent `Java Web Console`, so that you can use one of the other applications (`CAMS` or `FSM`) for monitoring.
Note that these applications should not be used for IAS configuration. Use only the IAS interface when administering the archive.

- The `VERSION` button displays revision-level and copyright information for the Sun Infinite Archive System software.
- The `LOGOUT` button ends your working session and logs you out of the Sun Infinite Archive System console interface.
- The `REFRESH` button (when present) updates the dashboard information.
- The `HELP` button opens the IAS help browser.

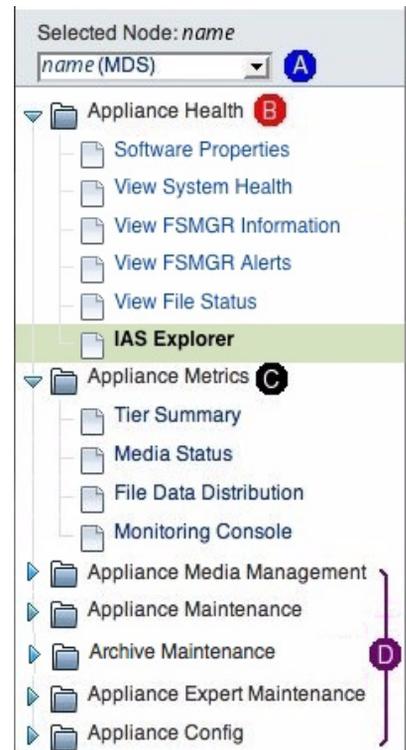
The role-based task tree

The lefthand pane of the IAS interface is the task tree that you use when moving around between the pages that hold the various displays and controls (do not use the browser's forward and back buttons for navigating, because they may display a stale, cached view of the page instead of retrieving fresh information from the archive appliance).

On clustered, multinode systems, a list control (**A** at right) lets you select the archive node that you want to work with: either the `MDS` (Meta Data Server) or the `Client` (the standby node that is configured to take over for the MDS following failover).

The remainder of the tree is a hierarchical menu of tasks that is organized around the user roles that are defined below in [“Understanding IAS operational roles” on page 41](#). When you log in to perform a given role, the menu tree gives you access to all the tools that you need while hiding functionality that might be destructive if invoked unintentionally.

As a rule, users log in under the most restrictive, least privileged role that allows them to perform a given task, because routine use of unnecessarily privileged roles exposes the system to unintended applications of powerful tools and may result in damage to irreplaceable, business-critical, archival data. The appliance comes fully configured and needs little attention thereafter. So, most of the time, you log in under a user role that lets you monitor the system and check metrics (**B** and **C**). You neither need nor have access to system configuration and maintenance tools (**D**).



Note that only users created with the IAS interface can access IAS functionality.

The content pane

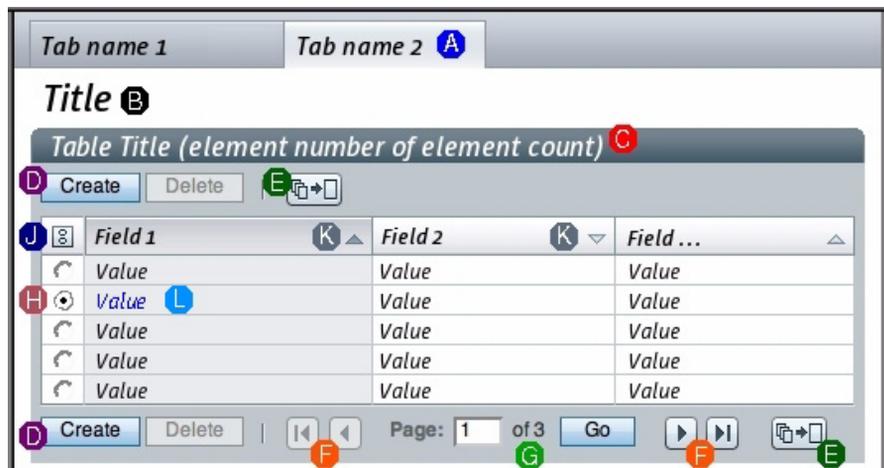
The content pane is the area on the right side of the page that displays system data and controls. When a page holds more information than a screen can comfortably hold, it may be divided into multiple tabs, each of which contains a single screen of information. For the most part, content pages and tabs use two basic layout elements:

- [Tables](#)
- [Forms](#)

Tables

Tables typically display existing data. The figure below shows a typical tabbed page with a table. You activate a tab and display its contents by selecting it (**A** below). A title is prominently displayed at the top of the display area (**B**).

In the display area, each table has a colored title bar that displays the name of the table and the number of records it displays out of the total available (**C** below). Under the title bar and immediately above the bottom of the table, there is a control bar that lets you add and delete records (**D**) and navigate through records. A paging button lets you expand the table to show all records on a single page (**E**). You can also browse move pages by pressing an arrow button (**F**) or by entering a page number and pressing **Go** (**G**).



Select a record by clicking the corresponding radio button (H above). Pressing the button at the top of the radio button column clears the selection (J). Clicking on the up and down arrows in the column headers sorts the records on the selected column (K). In some records, you can also click a hyperlink to view details for a single record (L).

Forms

Forms generally create/add or edit records. The figure below shows a typical, untabbed form. The page title is displayed prominently at the top (A below). You make changes using a set of familiar graphical user interface controls, including checkboxes, text boxes, list controls, and radio buttons (B). An asterisk, *, indicates a required field (C). Save and Reset buttons at the top and bottom of the page let you commit your changes or discard them (D).

The image shows a web form with the following elements:

- Title** (A) at the top left.
- Save** and **Reset** buttons (D) at the top right.
- A legend: **C** * Indicates required field.
- Checkbox**: (B).
- * Text box**: (B).
- * List control**: (B).
- Text/List**: (B).
- Radio button 1** and **Radio button 2** (B).
- Save** and **Reset** buttons (D) at the bottom right.

Pop-up browser windows

Infinite Archive System (IAS) console requires pop-up window support. So, if your web browser currently blocks pop-up windows, you must either enable them or make an exception for IAS. Different browsers have different capabilities and interfaces, so an exhaustive explanation of how to enable pop-ups is not possible. However, the following procedure works with the recommended Mozilla Firefox 2 browser.

▼ Enabling pop-up windows in Firefox browsers

1. **Launch a web browser session.**

2. From the main menu, execute one of the following command sequences (the exact selection depends on your operating system):
 - Firefox > Preferences
 - Tools > Options
 - Edit > Preferences
3. When the tabbed properties sheet appears, select the `Content` tab.
4. If you want unblock all pop-up windows, uncheck the `Block pop-up windows` check box, press the `OK` button at the bottom of the properties sheet, and stop here.

This allows all applications and web sites to launch popup windows, but this should not be a problem on a host that is not used for routine web browsing.
5. If you want unblock only pop-up windows that are part of the IAS console, press the `Exceptions` button.
6. When the `Allowed Sites-Pop-ups` dialog appears, enter the URL of the IAS console application in the `Address of Web Site` text field, and press `Allow`.
7. Close the dialog.

Stop here.

Understanding IAS operational roles

The Infinite Archive System (IAS) integrates a powerful set of tools that can handle almost every conceivable archiving situation, from the most routine to the most unlikely. Since many tools are seldom used and powerful enough to cause damage to systems or data if used improperly, IAS limits their use to those who perform a specific archiving role. Each role is defined by a set of situations and tasks that require specific tools and capabilities.

IAS distinguishes between two types of user roles:

- **Standard operating roles** encompass routine monitoring and basic, non-routine administration and management.
- **Special roles** encompass decidedly non-routine fault-management, repair, and major reconfiguration tasks that, if mishandled, can result in significant data loss and damage to the archive.

Standard operating roles

The Infinite Archive System (IAS) distributes normal user responsibilities between the following operational roles:

- `guest` grants users read-only access to the storage subsystem.

The `guest` role is defined by the Common Array Manager Software (CAMS) and a CAMS `guest` user has the same Infinite Archive System (IAS) user privileges as an IAS `monitor`, and an IAS `monitor` accesses CAMS interfaces as a CAMS `guest`.

- `monitor` is responsible for most of the routine administrative work associated with the Infinite Archive System (IAS).

The `monitor` role is defined by IAS and grants the user the extensive status and problem reporting capabilities that routine archive administration requires. IAS appliances are highly automated and need little of the administrative intervention associated with free-standing, general-purpose servers and storage subsystems. Instead, administrators monitor the system for abnormalities.

- `operator` is responsible for non-routine file-system administration.

The `operator` role is defined by the SAM-QFS File System Manager (FSM) software and grants the user the ability to correct SAM-QFS faults, modify jobs, and generate SAM reports. An `operator` can manually start, stop, and idle archiving. In other respects, an `operator` has the same access to the archive as a `monitor`.

- `storage` is responsible for non-routine administration of the disk storage subsystem.

The `storage` role is defined by the Common Array Manager Software (CAMS) that administers Sun disk arrays and array controllers. An IAS `storage` user has all of the capabilities of a CAMS `storage` administrator. He can modify the array configuration, change array attributes, acknowledge and address hardware alarms, and add or delete CAMS login accounts. In addition, an IAS `storage` user can monitor utilization and availability of all disk and tape archival storage media via the Infinite Archive System (IAS) Appliance Media Management interface.

Caution – When logging in with `storage` privileges, be careful! Know what you are doing. Work in the `storage` role only when necessary. If you just need to examine the storage configuration, use a CAMS `guest` account to minimize the risk of inadvertent configuration changes. While storage users cannot alter the overall archive system configuration, careless use of CAMS can cause loss of data and loss of storage subsystem functionality.

Special roles

The remaining IAS roles grant users access to powerful tools that, if misapplied, may cause significant data loss and disruption of core archive functions. For this reason, these special roles should be granted only to specially qualified individuals and should be used only when and as necessary. Two such roles are defined by the IAS system:

- `admin` can correct faults detected during normal operations, such as routine monitoring, and may carry out occasional housekeeping tasks and minor reconfigurations.

An `admin` user can access all of the displays and controls that are available to the `monitor` that reports a problem, plus an extensive set of maintenance and repair tools, stored under the `Appliance Maintenance` and `Appliance Config` branches of the task tree.

- `expert` can perform non-routine, disruptive tasks that become necessary only in unusual circumstances.

In the event of a major component failure, an `expert` user can configure a replacement. If requirements change significantly enough to invalidate the design of the archive itself, an expert can re-architect the archiving solution to meet the new need. The `expert` user has access to all of the displays and controls that are available to other users, and a special set of tools kept under the `Appliance Expert Maintenance` branch of the task tree.

Only experienced, fully trained system administrators are qualified to log on as an `admin` or `expert`. These individuals should be well versed in Sun Solaris server operations, disk- and tape-storage subsystems, local area networks (LANs), storage area networks (SANs), SAMBA and/or NFS shared file systems, SAM-QFS concepts and operations, and Sun Solaris Cluster concepts and operations.

Caution – If you do not have the required level of expertise available in your organization, do not make use of the special roles. You do not want to risk the security and stability of your archive! Instead, contact Sun support for advice and assistance.

Caution – Since powerful tools are liable to accidental misuse in even the best-trained hands, users should not routinely log on in special roles. Users should normally log in under a less privileged role, such as `monitor`. Then, when circumstances warrant, they should log on in their special role, carry out necessary duties, and then log out.

Understanding role-based user accounts

Roles and role-based user login accounts are not synonymous! Multiple users can fulfill the same role while using individual login accounts that uniquely identify each user. Conversely, a single user can perform multiple roles by having a separate, individual login account for each. For example:

- Users that log in as `User1`, `User2`, and `User17` may share the `IAS monitor` role without sharing user IDs and passwords.
- A single user may login in as `User19m` to perform the `IAS monitor` role and as `User19a` to function as an `IAS admin`.

By matching the tools to the work at hand and limiting access to those authorized to do the work, role-based access control provides the best balance between operational flexibility and system security.

For a full discussion of user and login account management, see [“Managing IAS user accounts” on page 47](#).

Using the IAS help system

A comprehensive online help system supplements the Infinite Archive System (IAS) documentation. To access the system, press the `Help` button at the upper right corner of the IAS banner. This launches the help browser popup window. The help browser has a navigation pane at the left and a topic pane at the right.

You can access context-sensitive help topics for most of the panels in the IAS Java Web Console interface at any time by pressing the `Help` button. Help for the current IAS application context appears automatically in the topic pane of the help browser.

You can also browse through the help system yourself. Press the `Help` button just as you would for context-sensitive help. But use the tabs in the navigation pane at left to browse through all the available topics. The navigation pane includes a

Contents tab, an Index tab, and a Search tab backed by a full-text search database. You can also step to the preceding or following topics using arrow buttons. These options are summarized in the table below.

Help control	Description
Help button	Press the graphical Help button to display the help topic for the current page.
Contents tab	In the help browser, select the Contents tab to view a hierarchical list of help topics. Within the list, click on triangle icons to expand the subtopics beneath a given topic.
Index tab	In the help browser, select the Index tab to display an alphabetized list of subject keywords. Click an index entry to display the help page for that topic.
Search tab	Click the Search tab to locate a word or phrase wherever it occurs in the help system. Type the text in the Search field, and press the graphical Search button. The Navigation pane displays a list of links to topics that match your search criteria in order of relevancy.
 Back arrow button	Press the back arrow button to return to the previous help topic.
 Forward arrow button	Press the forward arrow button to go to the next help topic. This button is enabled only if you have clicked the back arrow button..
 Printer button	Click to printer button to print the help topic that is currently displayed.

Managing IAS user accounts

Managing user accounts—and thus access to the IAS archive appliance—is critical both to the security of archival records and to their efficient use. Records must be protected from inadvertent damage and unauthorized use. At the same time, authorized users must have unfettered access the records and tools that they need when doing their jobs. A well thought-out, properly maintained approach to user accounts addresses both requirements. This chapter describes best practices for creating and maintaining accounts and provides detailed, step-by-step instructions for the most common account-management tasks.

Managing access to an archival system

You should apply two basic principles when managing access to the Infinite Archive System (IAS) appliance:

- Access to data and system-management functions should be limited to those who actually need the access, an approach commonly called *role-based access controls* (sometimes abbreviated RBAC).
- Access to data and system-management functions should be auditable, so that each user's actions can be traced back to a specific individual.

Role-based access controls

IAS implements a predefined role-based access-control scheme that associates sets of duties—monitor, operator, storage, admin, expert—with sets of tools and access privileges required for carrying out those specific duties (see [“Understanding IAS operational roles” on page 41](#) for details). By associating each user login account with a specific role, you provide the necessary access without unnecessarily exposing the archive.

While user accounts are associated with user roles, the two are not synonymous. IAS appliances come preconfigured with a generic, default login account for each user role. But these accounts are intended primarily for convenience during setup rather than for routine access by groups of users sharing a common user name. As we shall see, each individual user should have a unique, individual account, user name, and password for each IAS role that he or she has to perform.

Auditable, individual accounts

Generally accepted security practices mandate individually identifiable login accounts for all users. Individual accounts let system administrators log user access and audit user activities. Audit trails expose security breaches and document them in support of subsequent investigation and recovery efforts. For an archiving appliance, audit trails are also crucial for documenting the integrity of the repository and, consequently, the reliability and security of individual records. Increasingly, government and industry regulations require audit trails whenever an organization has custody of individual medical and financial records and whenever records pertain to the governance and financial transactions of publicly traded companies.

Multiple accounts for auditable, role-based access

Typically, users who perform roles that require privileged access—like IAS `admin` and `expert` or Solaris `root`—do not need elevated access privileges all of the time. An individual will, in most cases, do actual `admin`, `expert`, or `root` work relatively rarely. Much of the time, an IAS `expert` user's duties may be little different from those of a `monitor`. For this reason, using the least privileged account consistent with the tasks that need to be performed is likewise central to role-based access control.

Each user should have a separate login account for each role performed, rather than a single, highly privileged account for both routine daily activities and rarer, more specialized duties. For example, a user who has IAS `expert` privileges might normally log in as `userID-M` with IAS `monitor` privileges, reserving his IAS `expert` account, `userID-X`, for rare situations that require this level of intervention. This approach gives the user the powerful tools necessary when circumstances so require while protecting the archive against inadvertent misuse of those tools at other times.

Working with user accounts

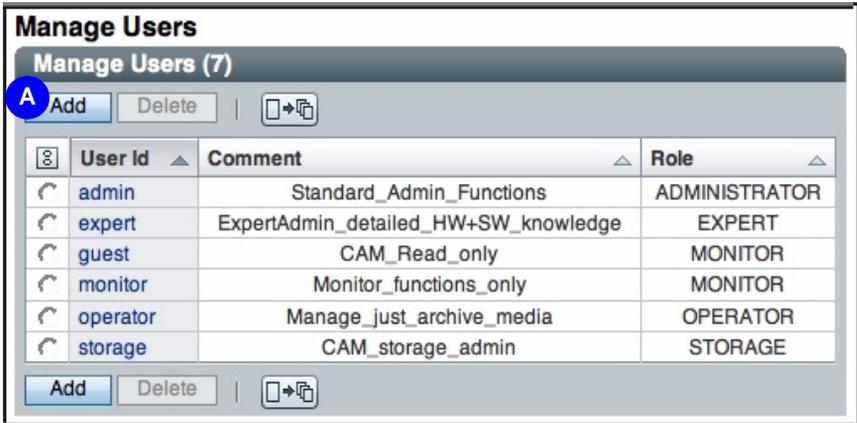
You can carry out the following account-management tasks using the IAS Java Web Console user interface:

- “Adding user accounts” on page 49
- “Modifying user account properties” on page 50
- “Deleting user accounts” on page 52.

▼ Adding user accounts

1. If you are monitoring a high-availability, multinode appliance, use the Selected Node list control at the top of the IAS task tree to select the MDS node.
2. In the IAS task tree, expand the Appliance Expert Maintenance branch and select Manage Users.
3. In the Manage Users table, press the Add button (A below).

Note that the Storage and Guest roles shown below are found only on high-availability IAS appliances that use Common Array Manager Software (CAMS) to manage disk arrays.



The screenshot shows the 'Manage Users' interface. At the top, there is a title 'Manage Users (7)' and a toolbar with 'Add', 'Delete', and a refresh icon. Below this is a table with columns for 'User Id', 'Comment', and 'Role'. The table contains six rows of user accounts. At the bottom, there is another toolbar with 'Add', 'Delete', and a refresh icon. A blue circle with the letter 'A' is positioned over the 'Add' button in the top toolbar.

User Id	Comment	Role
admin	Standard_Admin_Functions	ADMINISTRATOR
expert	ExpertAdmin_detailed_HW+SW_knowledge	EXPERT
guest	CAM_Read_only	MONITOR
monitor	Monitor_functions_only	MONITOR
operator	Manage_just_archive_media	OPERATOR
storage	CAM_storage_admin	STORAGE

4. When the Add User form appears, enter a new User Id and an optional Comment, if desired (B below).

The screenshot shows the 'Add User' form with the following fields and controls:

- * User Id:** A text input field with a red circle 'B' next to it.
- Comment:** A text input field.
- * Role:** A dropdown menu with a purple circle 'C' next to it. The dropdown is open, showing options: ADMINISTRATOR, EXPERT, MONITOR, OPERATOR, and STORAGE.
- * Password:** A text input field with a black circle 'D' next to it.
- * Confirm Password:** A text input field.
- Buttons:** 'Save' and 'Reset' buttons are located at the top right and bottom right. A blue circle 'E' is next to each 'Save' button.
- Legend:** A note '* Indicates required field' is located at the top right.

5. Using the spinner control (C above), select a user Role.
6. Enter a Password, and Confirm it (D above).
7. Press Save to finalize your changes, Reset to clear them (E above).
8. To back out of the screen, re-select Manage Users from the tree menu at left.

Stop here.

▼ Modifying user account properties

1. If you are monitoring a high-availability, multinode appliance, use the Selected Node list control at the top of the IAS task tree to select the MDS node.
2. In the IAS task tree, expand the Appliance Expert Maintenance branch and select Manage Users.

3. In the **Manage Users** table, click the hyperlinked account name for the account that you want to modify (A below).

Manage Users

Manage Users (1 - 5 of 7)

Add Delete |

	User Id	Comment	Role
<input type="radio"/>	user_ID A		STORAGE
<input type="radio"/>	storage	CAM_storage_admin	STORAGE
<input type="radio"/>	operator	Manage_just_archive_media	OPERATOR
<input type="radio"/>	monitor	Monitor_functions_only	MONITOR
<input type="radio"/>	guest	CAM_Read_only	MONITOR

Add Delete | Page: 1 of 2 **Go**

Note that the `Storage` and `Guest` roles shown above are found only on high-availability IAS appliances that use Common Array Manager Software (CAMS) to manage disk arrays.

4. When the **Edit User** form appears, change the `User Id` and/or the optional `Comment`, as desired (B below).

Edit User **E** Save Reset

User Id: **B**

Comment:

Current Role: CURRENT_ROLE

New Role: **C**

New Password: **D**

Confirm Password:

—User Roles—

—User Roles—

ADMINISTRATOR

EXPERT

MONITOR

OPERATOR

STORAGE

E Save Reset

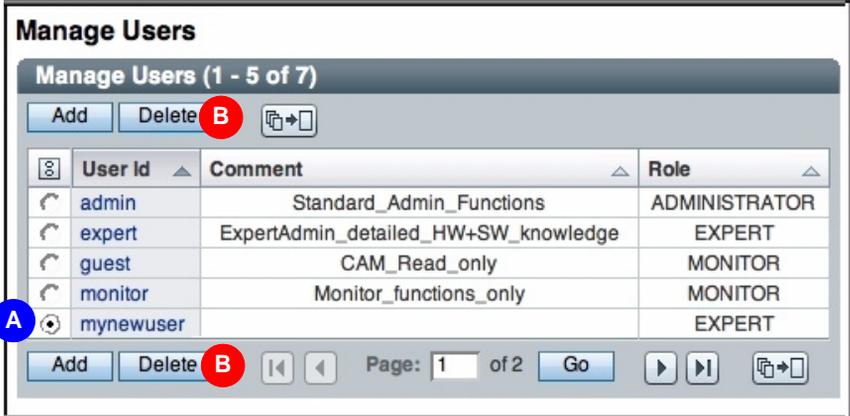
5. If you want to change the `Current Role`, select the new user Role using the spinner control (C above).
6. To change the current password, enter a `New Password` and `Confirm` it (D above).
7. Press **Save** to finalize your changes, **Reset** to clear them (E above).
8. To back out of the screen, re-select `Manage Users` from the tree menu at left.

Stop here.

▼ Deleting user accounts

1. If you are monitoring a high-availability, multinode appliance, use the Selected Node list control at the top of the IAS task tree to select the MDS node.
2. In the IAS task tree, expand the Appliance Expert Maintenance branch and select Manage Users.
3. In the Manage Users table, click the radio button beside the account that you need to delete (A below).

Note that the Storage and Guest roles shown below are found only on high-availability IAS appliances that use Common Array Manager Software (CAMS) to manage disk arrays.



Manage Users

Manage Users (1 - 5 of 7)

Add Delete **B** 

	User Id ▲	Comment ▲	Role ▲
	admin	Standard_Admin_Functions	ADMINISTRATOR
	expert	ExpertAdmin_detailed_HW+SW_knowledge	EXPERT
	guest	CAM_Read_only	MONITOR
	monitor	Monitor_functions_only	MONITOR
	mynewuser		EXPERT

A Add Delete **B**   Page: 1 of 2 Go   

4. Press Delete (**B** above).
5. When you are prompted, confirm or cancel the deletion.

Stop here.

Monitoring the system

Infinite Archive System (IAS) is a fully integrated appliance that needs little of the routine administrative intervention associated with free-standing, general-purpose servers and storage subsystems. You merely need to monitor the system for signs of abnormalities using the wide range of tools and reports that the IAS Java Web Console user interface provides. In general, you monitor two things: *appliance health*, which has to do with the availability of the archive, and *appliance metrics*, which are measures of archive utilization.

Availability is conceptually straightforward and easy to monitor. When key components such as a server node, a network interface, an archive file system, or a storage subsystem become unavailable, core functionality is abruptly lost or degraded, and alerts are posted on the IAS console and/or in email directed to an administrator.

Utilization issues are more subtle and require more judgment on your part. When usage is excessive, lack of resources, such as storage media, can stop the archiving process as effectively as a component failure, even though the system is functioning normally. Utilization levels that presage imminent trouble in one organization may be perfectly consistent with years of trouble-free operation in another. So recognizing trends and rates is crucial when you are monitoring utilization. A resource that is 80% used is fine if utilization grows 1% per year but a crisis if growth exceeds 1% every week.

In the sections listed below, this chapter explains the procedures for monitoring appliance health and examining appliance metrics:

- [“Monitoring appliance health” on page 54](#)
- [“Monitoring appliance metrics” on page 70.](#)

Monitoring appliance health

As we have seen, appliance health is a measure of the availability of the Infinite Archive System appliance and its subcomponents and, hence, of the level of service that the appliance provides at any given time. When the appliance is *operational*, it provides the full range of functionality for which it was designed and sold (although, as we will see, it may be overutilized if your requirements have grown in the interim). When the appliance is *degraded*, it is still functioning but at reduced performance or capacity, following a component failure of some kind. Backup components and systems have taken on additional loads in order to prevent any interruption in service. But there is little or no redundancy remaining, so you need to take action before anything further happens. An appliance that has completely failed and is down is no longer available at all.

The following discussion explains how you use the various IAS Appliance Health monitoring tools to detect problems that degrade system availability.

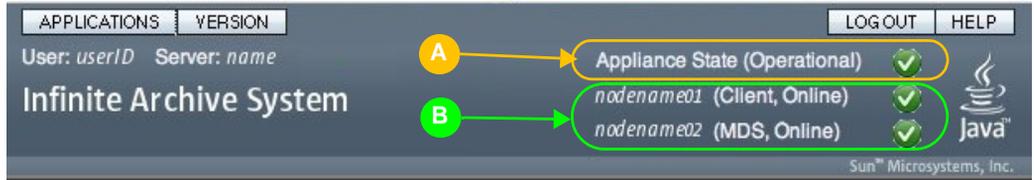
- [“Assessing basic system availability” on page 54](#)
- [“Checking the system configuration” on page 55](#)
- [“Monitoring archive system faults” on page 61](#)
- [“Investigating file system problems” on page 66.](#)

Assessing basic system availability

Whenever you log in to the Infinite Archive System (IAS) console, check the IAS banner for the Appliance State (**A** below).



On a high-availability system, you can also see the status of the two server nodes, the Client (standby) and the MDS (MetaData Server) nodes (**B** below).



The following text and graphics describe appliance and component states:

- ✓ means that the appliance or node is operational and online;
- ⚠ warns you of a problem;
- ❗ alerts you to a critical error.

If you note a warning or error, you should investigate using the tools and procedures listed in the remainder of this chapter.

Checking the system configuration

If you detect a problem, you should start by gathering basic system configuration information. If you need to engage Sun support services or if you wish to make the best use of on-line customer-support resources, you will, at a minimum, need to start by “[Checking IAS software version information](#)” (see page 55). For instructions on gathering additional configuration information that may prove useful, either to you or to a Sun support representative, see the following:

- “[Checking basic system health](#)” on page 56
- “[Displaying software packages installed on a server node](#)” on page 58
- “[Displaying log and trace files](#)” on page 58
- “[Displaying configuration files](#)” on page 59.

▼ Checking IAS software version information

Problems and solutions are often specific to the installed version of a software product. So, before initiating a support call, open the `Software Properties` sheet and note the version numbers of each core IAS software component.

1. **In the IAS task tree, expand the Appliance Health branch.**
2. **Under the Appliance Health branch, click Software Properties.**
The property sheet displays version information for the core IAS software components:

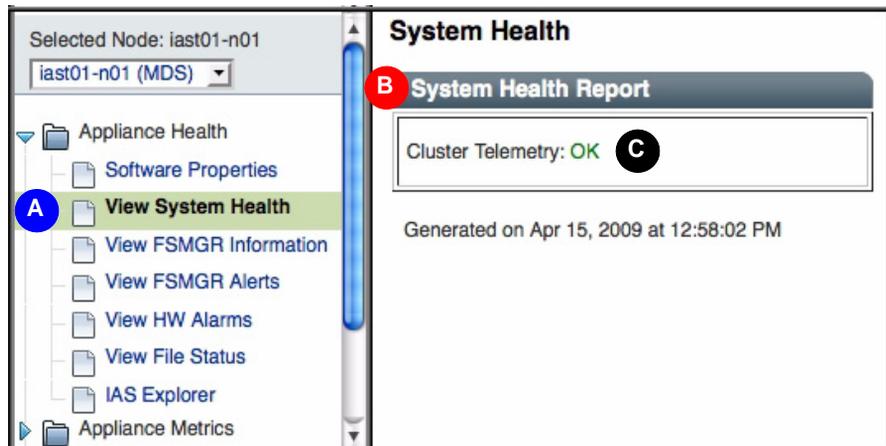
- the SAMS-QFS archive-management and file system software
- the Flash Archive Image from which the software was installed
- the SUNwcistk console utilities package.

Next task: If you are actively monitoring or troubleshooting, you probably want to go to [“Monitoring archive system faults” on page 61](#). But if you want to drill down further into the appliance configuration, including versioning information for installed component software, start with [“Checking basic system health” on page 56](#).

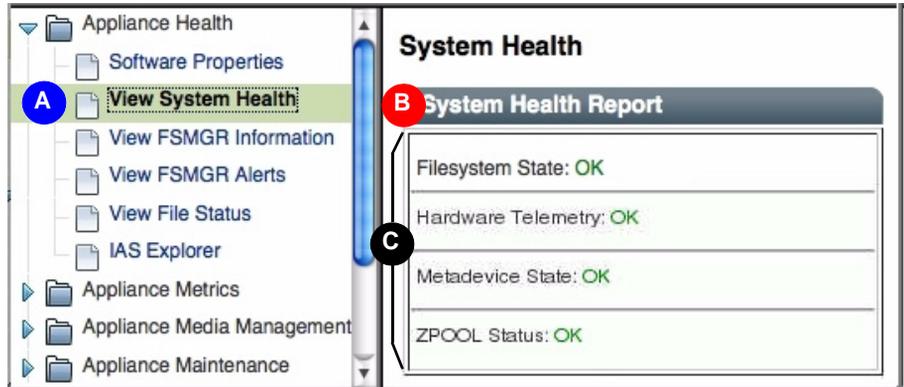
▼ Checking basic system health

1. Open the **Appliance Health** node of the IAS system tree.
2. Select **View System Health (A below)** to open the **System Health** summary panel (B).

High-availability, multi-node system:



Single-node system:



Stop here.

▼ Checking the archive-server node configuration

Changes in the operation or configuration of a system can help to explain many types of problem. Each archive-server node on an IAS appliance records these types of changes on an ongoing basis in log and trace files and in lists of installed software packages. The IAS console application provides convenient tools for examining this information safely. To use them, proceed as follows:

1. **If you are monitoring a high-availability, multinode appliance, use the Selected Node list control at the top of the IAS task tree to select the name of the server node that you wish to examine.**
2. **In the IAS task tree, expand the Appliance Health branch.**
3. **Under the Appliance Health branch, click View FSMGR Information.**
This displays the Server Configuration page.

Next task: this depends on what you are looking for. See the following:

- [“Displaying software packages installed on a server node” on page 58](#)
- [“Displaying log and trace files” on page 58](#)
- [“Displaying configuration files” on page 59.](#)

▼ Displaying software packages installed on a server node

To display the software packages that are installed locally on a server, proceed as follows:

1. If you have not already done so, start by [“Checking basic system health” on page 56](#).
2. Then scroll down to the **Packages** section, and view the **Application Packages** table.

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

Next task: to continue examining the server node configuration, go to [“Displaying log and trace files” on page 58](#)

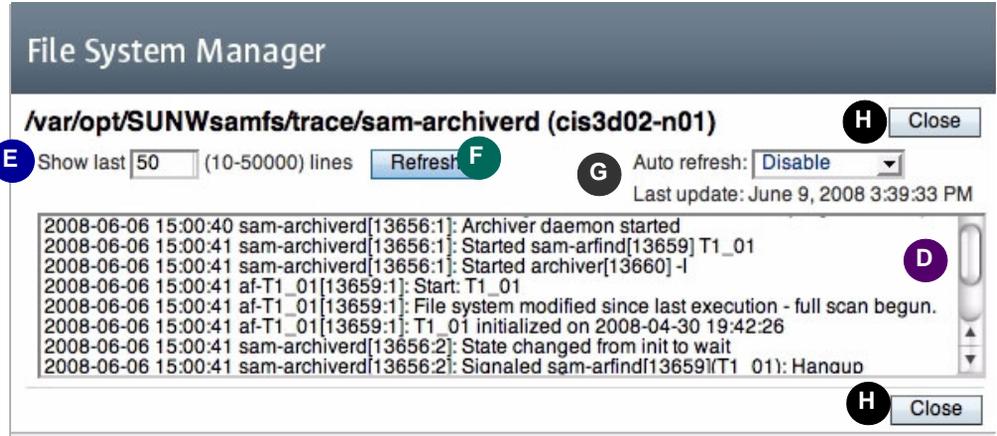
▼ Displaying log and trace files

Log and trace files record the history of a system and reveal outages, errors, and configuration changes that can help to diagnose faults and performance problems. These files can also reveal performance trends that give advance warning of future issues. To view the log and trace files for two core IAS components, the Sun StorageTek SAM archiving software and the Sun StorageTek QFS file systems, proceed as follows:

1. If you have not already done so, start by [“Checking basic system health” on page 56](#).
2. Scroll down to the **Log and Trace** section (A below).
3. In the **Path** column of the **Log and Trace Files** table (B below), click on the hyperlinked path specification of the log or trace file that you want to view (C).

Log and Trace A					
Log and Trace Files (20)					
Name	Type	Status	Path B	Flags	Size
Master automated library daemon	Trace	ON	/var/opt/SUNWsamfs/trace/sam-amld		60 B
Archiver daemon	Trace	ON C	/var/opt/SUNWsamfs/trace/sam-archiverd		59.24 kB
Media catalog daemon	Trace	ON	/var/opt/SUNWsamfs/trace/sam-catserverd		66 B
File system daemon	Trace	ON	/var/opt/SUNWsamfs/trace/sam-fsd		119.45 kB
File transfer daemon	Trace	ON	/var/opt/SUNWsamfs/trace/sam-rftd		60 B
Recycler	Trace	ON	/var/opt/SUNWsamfs/trace/sam-recycler		64 B
Alternate Recycler	Trace	ON	/var/opt/SUNWsamfs/trace/sam-nrecycler		65 B
Shared file system daemon	Trace	ON	/var/opt/SUNWsamfs/trace/sam-sharefsd		196.08 kB

4. When the file-viewer pop-up window appears, examine the file by scrolling through the text in the view pane (D below).



5. To limit the number of lines displayed, enter a number between 10 and 50000 in the Show last text field (E above).
6. To update the display, press the Refresh button for a manual update (F above) or set the Auto refresh control to Enable (G).
7. Press Close to finish (H above).

Next task: to continue examining the server node configuration, go to [“Displaying configuration files” on page 59](#).

▼ Displaying configuration files

The configuration files for the Sun StorageTek SAM archiving software and Sun StorageTek QFS file systems control the way that the archive performs and the way that it manages archival data. To view the configuration files, proceed as follows:

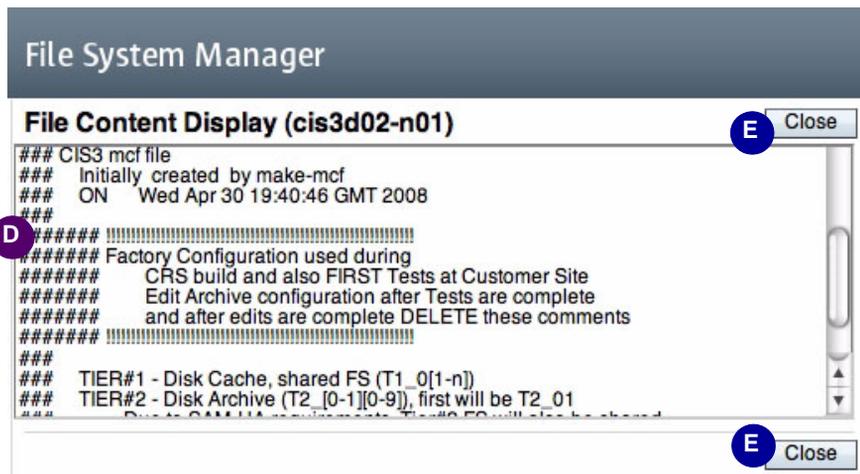
1. If you have not already done so, start by [“Checking basic system health” on page 56](#).
2. Scroll down to the Configuration File Status section (A below).

- In the **Path** column of the Application Configuration File Status table (B below), click on the hyperlinked path to the file that you want to view (C).

Configuration File Status A

Application Configuration File Status (8)	
Name	Status
/etc/opt/SUNWsamfs/mcf	OK
/etc/opt/SUNWsamfs/defaults.conf	OK
/etc/opt/SUNWsamfs/diskvols.conf	OK
/etc/opt/SUNWsamfs/samfs.cmd	OK
/etc/opt/SUNWsamfs/archiver.cmd	OK
/etc/opt/SUNWsamfs/releaser.cmd	OK
/etc/opt/SUNWsamfs/recycler.cmd	OK
/etc/opt/SUNWsamfs/stager.cmd	OK

- When the File System Manager file-viewer pop-up window appears, examine the file by scrolling through the text in the view pane (D below).



- Press the **Close** button when you finish (E above).

Stop here. If you wish to check system alerts, see [“Monitoring archive system faults”](#) on page 61.

Monitoring archive system faults

The Infinite Archive System (IAS) continually monitors the file systems and storage hardware that jointly store and protect archival information. When IAS detects a problem, it warns you by displaying a file-system alert or a hardware alarm in the IAS console. This section covers the monitoring tasks:

- “Monitoring file system alerts” on page 61
- “Monitoring disk-storage hardware alarms” on page 64

▼ Monitoring file system alerts

1. If you are monitoring a high-availability, multinode appliance, use the Selected Node list control at the top of the IAS task tree to select the MDS node.
2. In the IAS task tree, expand the Appliance Health branch.
3. Under the Appliance Health branch, click View FSMGR Alerts.
4. When the Fault Summary page appears (A below), view fault alerts in the records of the Faults table (B).

A Fault Summary

Log and Trace Files: Select a Log/Trace file to be viewed View

B Faults (2)

Acknowledge Delete Filter: All Items Critical Major Minor State Active message

ID	Severity	Time Generated	State	Description
ID_#	Critical	date-time	Active	message
ID_#	Minor	date-time	Active	message

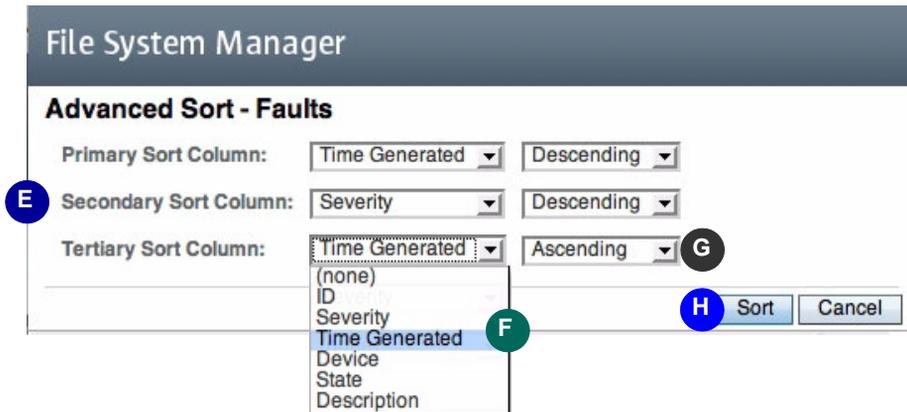
Acknowledge Delete

C Filter: All Items Critical Major Minor

D Filter: All Items Critical Major Minor

5. If there are many faults, you can limit the display to a selected severity using the Filter list control (C above).
6. You can also sort the alerts by clicking on the up and down arrows in the column headers (see “Tables” on page 39 for details).

- Alternatively, you can perform an **Advanced Sort** by pressing the large up arrow in the control-button area for the table (D above). When the **Advanced Sorts-Faults** pop-up window appears, specify a sort. For each **Sort Column** (E below), select a criterion (F) and a sort order (G) using the list controls. Then press **Sort** to sort the alerts or **Cancel** to abort (H).



Next task:

- If you wish to drill down to examine the details behind an alert, go to [“Examining associated archive system log and trace files”](#) on page 62.
- Otherwise, proceed with [“Monitoring disk-storage hardware alarms”](#) on page 64.

▼ Examining associated archive system log and trace files

- If you are already on the **Fault Summary** page (A below), go to step 5.
- Otherwise, if you are monitoring a high-availability, multinode appliance, use the **Selected Node** list control at the top of the IAS task tree to select the MDS node.
- In the IAS task tree, expand the **Appliance Health** branch.
- Under the **Appliance Health** branch, click **View FSMGR Alerts**.

- To view trace logs for information connected with alerts, use the Log and Trace Files list control to select a file (B below).

The control lists the files in the `/var/opt/SUNWsamfs/trace`, `/var/adm/`, and `/global/CIS/var/opt/SUNWsamfs/` system subdirectories.

A **Fault Summary**

Log and Trace Files: Select a Log/Trace file to be viewed View **B**

Faults (2)

Acknowledge

ID	Severity
	Critical
	Minor

Acknowledge

/var/opt/SUNWsamfs/trace/sam-amld
 /var/opt/SUNWsamfs/trace/sam-archiverd
 /var/opt/SUNWsamfs/trace/sam-catserverd
 /var/opt/SUNWsamfs/trace/sam-fsd
 /var/opt/SUNWsamfs/trace/sam-rftd
 /var/opt/SUNWsamfs/trace/sam-recycler
 /var/opt/SUNWsamfs/trace/sam-nrecycler
 /var/opt/SUNWsamfs/trace/sam-sharefsd
 /var/opt/SUNWsamfs/trace/sam-stagerd
 /var/opt/SUNWsamfs/trace/sam-serverd
 /var/opt/SUNWsamfs/trace/sam-clientd
 /var/opt/SUNWsamfs/trace/fsmgmt
 /var/adm/sam-log
 /var/adm/messages
 /var/adm/messages
 /global/CIS/var/opt/SUNWsamfs/recycler.log
 /global/CIS/var/opt/SUNWsamfs/archiver.log
 /global/CIS/var/opt/SUNWsamfs/releaser.log
 /global/CIS/var/opt/SUNWsamfs/stager.log

- Once you have selected a file, press the View button (C below).

The console opens a file viewer in a pop-up window.

Fault Summary

Log and Trace Files: /var/opt/SUNWsamfs/trace/fsmgmt View **C**

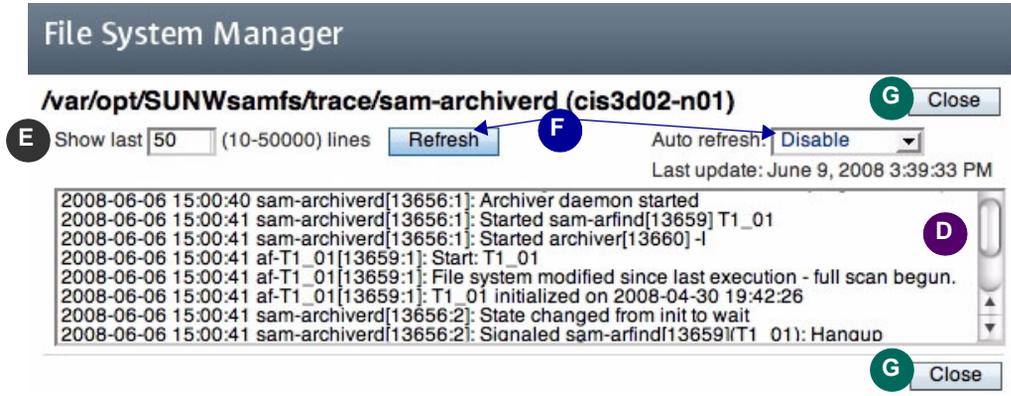
Faults (2)

Acknowledge Delete | Filter: All Items

ID	Severity	Time Generated	Device	State	Description
ID_#	Critical	date-time		Active	message
ID_#	Minor	date-time		Active	message

Acknowledge Delete

7. When the File System Manager file viewer pop-up screen appears, scroll through the contents of the file in the text window (D below).



8. To limit or expand the number of lines displayed (starting from the end of the file) by entering a number between 10 and 50000 in the Show last text box (E above).
9. To update the display with the latest version of the file, press the Refresh button each time that you want an update or set the Auto refresh list control to Enable (F above).
10. Press the Close button when you are finished (G above).

Next task: [“Monitoring disk-storage hardware alarms” on page 64.](#)

▼ Monitoring disk-storage hardware alarms

On high-availability IAS appliances that use Common Array Manager Software (CAMS) to manage disk arrays, you can view disk-array alarms in the IAS Java Web Console. When CAMS monitoring software detects a disk-array problem that requires intervention, the IAS interface displays an alarm at one of four severity levels:

- **Down** alarms tell you that the alarmed device or component has failed and that you must arrange for immediate service.
- **Critical** alarms tell you that a device or component has made a serious error and that you must arrange for immediate service.
- **Major** alarms tell you that a device or component has made a significant error and that you may or may not have to arrange for service.
- **Minor** alarms tell you that a device or component error has occurred, but that you probably do not need to arrange for service.

To view hardware alarms, proceed as follows:

1. Use the Selected Node list control at the top of the IAS task tree to select the MDS node.
2. In the IAS task tree, expand the Appliance Health branch.
3. Under the Appliance Health branch, click View HW Alarms.

If the initial system installation is not complete, you will see the following Automated Service Request registration form and will either have to complete the registration process or decline the service before proceeding (see “Registering for the ASR feature” on page 207):

Site Information
Complete the site information form first, before you start to use this application.

Site Information Contact Information * Indicates required field

Site Information
Enter information describing where arrays are installed. This information may be used to ensure that parts are shipped to the correct location.

* Company Name:
Company name must match that associated with a Sun Online Account.

* Site Name:

Otherwise, an Alarms table appears in the content pane of the interface (A below). Severity is indicated by the icons (B), where the grey icon with the down arrow means Down, the red icon with the vertical bar means Critical, the orange icon with the diagonal bar means Major, and the yellow icon with the horizontal bar means Minor.

Alarms
Alarms (1)

Acknowledge... Reopen... Delete | Filter: All Items

All Items
All Items
Open
Acknowledged

Severity	Alarm Details	Device	Component	Type	Date	State	Auto Clear
	link_to_details	name (type)	part_affected	of_alarm	date-time	status	Yes No

Down Critical Major Minor

4. To filter alarms by state, select All Items | Open | Acknowledged from the Filter list control (C above).
5. To view Alarm Details, click on the hyperlinked message (D above).

Next task:

- If you want to drill down and examine detailed status information for individual file systems, go to [“Investigating file-system, directory, or file access problems”](#) on page 66.

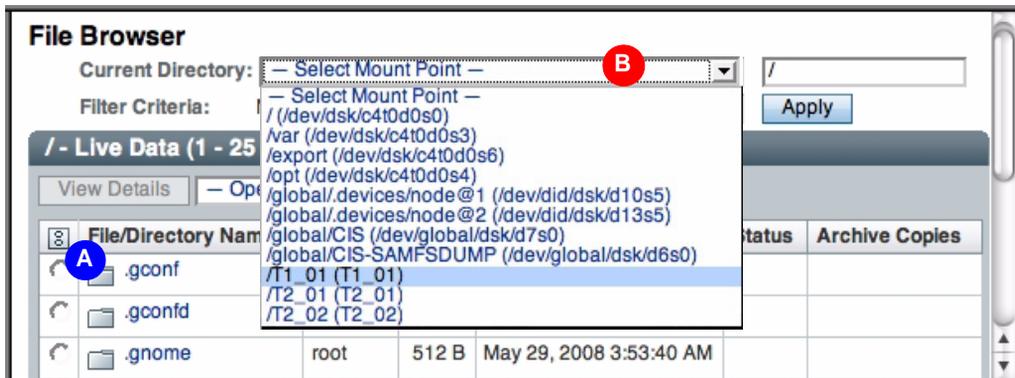
Investigating file system problems

Sometimes, a user will be the first to detect an archive problem: a file or directory cannot be found or accessed, generally because a user moved or inadvertently deleted it. In such cases, you can use the IAS File Browser to locate the missing item in either the active data repository or in the recovery history. If the item is indeed missing, you can then arrange to have an administrator restore the file from the last good IAS SAM-QFS *recovery point*. Proceed as follows.

▼ Investigating file-system, directory, or file access problems

1. If you are monitoring a high-availability, multinode appliance, use the `Selected Node` list control at the top of the IAS task tree to select the MDS node.
2. In the IAS task tree, expand the `Appliance Health` branch.
3. Under the `Appliance Health` branch, click `View File Status`.

The File Browser appears.



4. To drill down to a specific file system, you can click on the hyperlinked name of the mount-point directory in the `File/Directory` column of the table (A above), or you can select a mount point from the `Current Directory` list control (B).

5. You can then examine the status of the selected file system (C below).

File Browser

Current Directory: /T1_01

Filter Criteria: No Filter **E** Maximum Entries:

Live Data Recovery Point:

/T1_01 - Live Data (7)

	File/Directory Name	User	Size	Last Modified	Status C	Archive Copies
	Up One Level D					
	.archive	root	4.00 kB	June 5, 2008 3:47:59 PM	Online	
	.inodes	root	1.25 MB	June 5, 2008 6:35:05 PM	Online	
	.stage	root	4.00 kB	June 5, 2008 3:48:36 PM	Online	
	user_data	root	4.00 kB	June 5, 2008 3:48:31 PM	Online	
	lost+found	root	16.00 kB	June 5, 2008 3:47:59 PM	Online	
	samppm	root	4.00 kB	June 12, 2008 12:00:01 AM	Online	
	temp	root	4.00 kB	June 5, 2008 3:48:31 PM	Online	

6. If you do not see what you are looking for, you can search up and down in the directory hierarchy. Click on a hyperlinked subdirectory name in the File/Directory Name column to descend a level or click on the Up One Level link to move to the directory above (D above).

7. To reduce the visual clutter and simplify searches of large file systems, set Filter Criteria by pressing the associated Change button (E above).

The Filter Files popup window appears (see below).

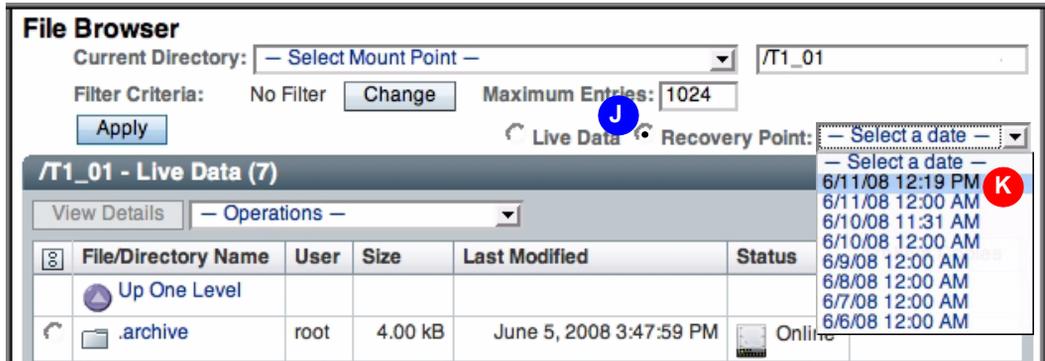
8. When the **Filter Files** popup window appears, enter your criteria using the fields and controls provided (**F** below), and press **OK** to apply the filters, **Cancel** to discard them (**G**):

The **File Name Pattern** field (**H** above) uses the pattern-matching syntax outlined in the following table:

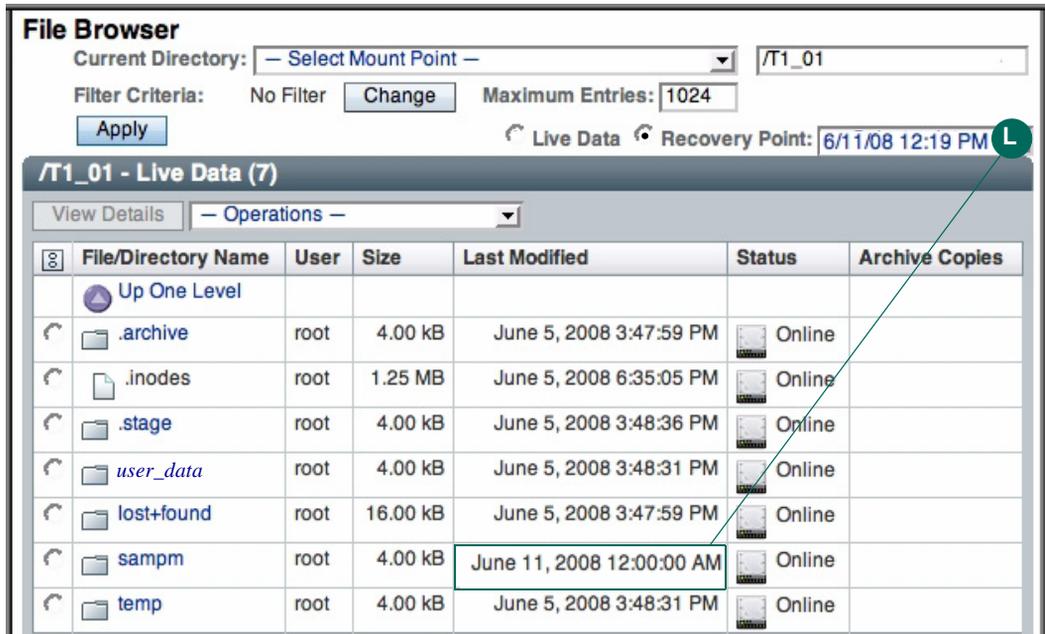
Wildcard	Description	Sample pattern	Sample match
/	matches a directory path only and cannot be matched by any wildcard character	/var /abc/ *	/abc /abc/abc, /abc/def, etc.
?	matches any single character	a?d	aad, abd, acd, etc.
*	matches zero or more characters	a*d	ad, abcd, acbd, aaaad, etc.
[]	matches files any one of the characters enclosed by the brackets	a[bc]	ab or ac
""	matches only the literal string within the quotation marks	"abc"	only abc
\	escapes the wildcard character and makes it into a literal value	a\?b a\\b a\\?	a?b (NOT acb) a\b a\b, a\c, or a\?

- If you have indexed recovery points for the file system and wish to trace deleted files, click the **Recovery Point** radio button (J below), and select a date from the corresponding list control (K).

If you have not indexed at least one recovery point, the radio buttons (J below) do not appear. To index recovery points and recover missing files, see “Restoring a file system” on page 131.



The File Browser displays the file system as it was at the recovery point (L below).



Next task: If you discovered anything that requires administrative intervention, such as a file system that needs to be recovered, notify a user that has IAS admin authority.

Monitoring appliance metrics

IAS Appliance Metrics measure archive utilization, storage capacity, status and availability of media, data distribution, and job queues. Appliance Metrics help with routine management and planning and when trouble shooting availability issues. If, for instance, you find that archive growth is rapidly depleting your remaining storage capacity, then the archive administrators will need to either reconsider your organization's archiving policies or purchase additional storage from Sun. If users report problems accessing the archive or archive functions, checking utilization of key resources can help you to diagnose the problem.

The following topics explain how you use the various IAS Appliance Metrics monitoring tools:

- [“Assessing storage utilization with Tier Summary” on page 70](#)
- [“Assessing media utilization and status” on page 74](#)
- [“Assessing the distribution of files in the archive” on page 78](#)
- [“Monitoring archiving processes with the File System Manager Monitoring Console” on page 82.](#)

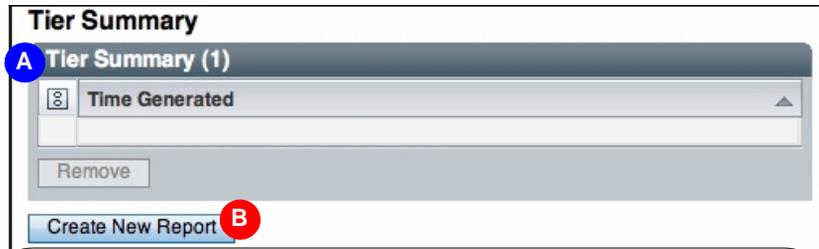
Assessing storage utilization with Tier Summary

The Infinite Archive System uses SAM-QFS archival file systems that organize storage media into hierarchical *tiers* for efficient information lifecycle management. When users and applications are actively reading and modifying data, IAS stores the data files on high-speed, random-access disk media that can deliver the expected responsiveness. But fast disk storage is an expensive way to store relatively static archival data. So, once data is no longer being regularly accessed, IAS migrates the files to less expensive media: higher capacity, lower speed disk storage and/or high-capacity sequential-access tape. Migration is completely transparent to users, client applications, and operating systems, all of which see only a standard, disk-based file system. If a user or application subsequently re-accesses a migrated file, IAS can automatically *stage* it back to the fast-access disk tier so that users experience no loss of responsiveness.

The Tier Summary displays storage utilization reports for each tier in the archive. This section explains [“Creating Tier Summary reports” on page 71](#), [“Viewing Tier Summary reports” on page 72](#), and [“Deleting a Tier Summary report” on page 73](#).

▼ Creating Tier Summary reports

1. If you are monitoring a high-availability, multinode appliance, use the Selected Node list control at the top of the IAS task tree to select the MDS node.
2. In the IAS task tree, expand the Appliance Metrics branch.
3. Under the Appliance Metrics branch, click the Tier Summary hyperlink.
A Tier Summary table appears (A below) that lists the available tier reports.



4. To generate the new report, press the Create New Report button (B above).
The new report appears in the summary table (C below), listed by time generated.



5. To select a pre-existing or newly generated report, locate it in the Tier Summary table and click on the hyperlinked date of the report (C above).

Next task: [“Viewing Tier Summary reports” on page 72.](#)

▼ Viewing Tier Summary reports

1. In list shown in the **Tier Summary** table (A below), click on the hyperlinked `Time Generated` for the report that you want to view (B).

The selected report appears in the `Current Report` area (C).

Tier Summary

A Tier Summary (1)

B Time Generated
Wed May 07 18:24:38 GMT+00:00 2008

Remove

Create New Report

C Tier Usage Summary

D Tier #1 Tier #2 Tier #3

Used
Free

2.3G of 4.0T Used 28M of 12.0T Used

E Tier #1 Disk Cache (2)

Name ▲	Size ▲	Percent Free ▲	Available ▲	HWM ▲	LWM ▲	Mount Point ▲	Used ▲
T1_01	100G	1	100G	85	80	/T1_01	808K
Total	0.098T		0.098T				0T

F Tier #2 Disk Archive (3)

Name ▲	Size ▲	Percent Free ▲	Available ▲	HWM ▲	LWM ▲	Mount Point ▲	Used ▲
T2_01	100G	1	100G	95	90	/T2_01	2.0M
T2_02	100G	1	100G	95	90	/T2_02	2.0M
Total	0.195T		0.195T				0T

G Tier #3 Tape Archive (0)

Name	Size	Percent Free	Available	HWM	LWM	Mount Point	Used
Tier #3 Tape Archive							

The report contains up to three tables, depending on the number of tiers in the solution. In the example above:

- Tier Usage Summary pie-charts (D above) provide a high-level, graphical representation of space utilization in each tier of the archive.
- Tier #1 Disk Cache (E above) lists utilization statistics for the fast, low-capacity, random-access Fibre Channel or SCSI disk that holds active files.

- Tier #2 Disk Archive (**F**) lists utilization statistics for the slower, higher-capacity, random-access SATA disk that hold less frequently accessed files.
- Tier #3 Disk Archive (**G**) lists utilization statistics for the high-capacity, sequential access tape media that provides bulk storage of extremely large or very infrequently accessed files.

Each table in the report displays both the total space available in the corresponding tier of the archive and the following detail information for each file system stored on the tier:

- the `Name` of each file-system
- the `Size` of each file system
- `Percent Free`, the percentage of space in each file system that is available for storing new files
- the amount of space in each file system that is `Available` for storing new files
- the `HWM` (high water mark) at which IAS begins to release each file system's files from a storage tier in order to reclaim space (expressed as a percentage of capacity in use)
- the `LWM` (low water mark) at which IAS stops releasing each file system's files from a storage tier (also expressed as a percentage of capacity in use)
- the `Mount Point` for each file system
- the amount of space `Used` in each file system.

Stop here.

▼ Deleting a Tier Summary report

1. **At the top of the Tier Summary page, click the radio button corresponding to the report that you wish to delete (A below).**



2. **Press the Remove button (B above).**

Stop here.

Assessing media utilization and status

The Media Utilization & Status page provides an overview of archival storage capacity and summaries of the status and utilization of tape cartridges.

▼ Accessing the Media Utilization and Status page

1. If you are monitoring a high-availability, multinode appliance, use the Selected Node list control at the top of the IAS task tree to select the MDS node.
2. In the IAS task tree, expand the Appliance Metrics branch.
3. Under the Appliance Metrics branch, click the Media Status hyperlink.

The Media Utilization and Status page appears. The Media Utilization Historic Summary table lists the available reports and provides controls for creating new reports and deleting old ones (A below).

	Time Generated	Size	Description
↻	June 12, 2008 10:36:19 AM	686 B	Archive storage capacity utilization and tape
↻	May 29, 2008 6:12:22 AM	712 B	Archive storage capacity utilization and tape

Reports include the following information:

- The VSNS Requiring Operator Attention table (B below) lists the media volumes that have suffered errors that you specify when you create a report (a VSN is the *Volume Serial Number* that uniquely identifies a tape storage cartridge).
- The Archive Media Utilization table (C) lists the total size, space used, and percentage space used for all archival media in an archiving file system.

- The `VSN is blank` table (D) lists the media volumes that are empty at the time the report is created.

Archive storage capacity utilization and tape VSN error status (June 12, 2008 6:14:14 PM)				
B VSNs Requiring Operator Attention				
Name	Media Type	Status		
AC6613	IBM 3580 LTO	Damaged Media		
C Archive Media Utilization				
Name	Media Type	Total Space	Free Space	Space Consumed (%)
samfs2.1	IBM 3580 LTO	97.87 GB	93.65 GB	4
D VSN is blank				
Name	Media Type			
AC6616	IBM 3580 LTO			

- The `VSN is in use in the last 24 hours` table (E below) lists all VSNs accessed in the 24 hours prior to creation of the report.
- The `Archiver Pool Levels` table (F) lists media pools and displays the number of volumes and free space available in each.
- The `VSN Summary` table (G) lists all of the media types in the file system and displays the number of volumes and the total size and free space available on each type.

E VSN is in use in the last 24 hours				
Name	Media Type			
No items found				
F Archiver Pool Levels				
Name	Media Type	VSN Count	Free Space	
myPool	IBM 3580 LTO	4	689.32 GB	
G Tape VSN Summary				
Media Type	VSN Count	Total Space	Free Space	
Disk	2	0 B	0 B	
IBM 3580 LTO	6	791.33 GB	400.00 GB	

If one or more Sun StorageTek Automated Cartridge System Library Software (ACSL) libraries are attached to the server node, you can also generate an ACSL report. ACSL reports include the following additional information:

- The `Drives` table lists the drives that are attached to the storage module, with the physical location in the library and status of each.

- The `Locks` table lists the lock status of the drives.
- The `Scratch Pools` table lists the pools of empty media volumes that are available.
- The `Accessed or Imported Volumes` table lists the volumes that have been accessed or imported within the 24 hours preceding creation of the report.

The remainder of this section explains the procedures for the following tasks:

- [“Creating a media utilization report” on page 76](#)
- [“Viewing a media utilization report” on page 77](#)
- [“Deleting a media utilization report” on page 78.](#)

▼ Creating a media utilization report

1. To create a new report, go to the `Media Utilization Historic Summary` table and press the `New` button (A below).

Media Utilization & Status (Historic)
Click the New button to save the current media utilization and status.

Media Utilization Historic Summary (2)

A

	Time Generated	Size	Description
↻	June 12, 2008 10:36:19 AM	686 B	Archive storage capacity utilization and tape
↻	May 29, 2008 6:12:22 AM	712 B	Archive storage capacity utilization and tape

The `New Metric` popup window appears (see below).

2. On the **New Metric** popup page, check the **Archival storage utilization and tape VSN error status** radio button if you wish to include error information in the output (B below).

Note: *VSN* is the abbreviation for *Volume Serial Number*, the barcode that uniquely identifies a tape storage cartridge.

3. Next, specify the type of **VSN Status** information that you want to include by checking the corresponding check boxes (D above). Or accept the defaults.
4. Specify the type of **Utilization** information that you want to include by checking the corresponding check boxes (E above). Or accept the defaults.
5. Press the **Submit** button to generate the report or **Cancel** to abort (F above).

Next task: [“Viewing a media utilization report” on page 77.](#)

▼ Viewing a media utilization report

1. To view a report, go to the **Media Utilization Historic Summary** table and click on the **Time Generated** hyperlink for the report (A below).

Time Generated	Size	Description
June 12, 2008 10:36:19 AM	686 B	Archive storage capacity utilization and tape
May 29, 2008 6:12:22 AM	712 B	Archive storage capacity utilization and tape

2. When the **Time Generated** for the report is greyed out and the corresponding radio button is selected (B below), scroll down to view the report (C).

Media Utilization & Status (Historic)
Click the New button to save the current media utilization and status.

Media Utilization Historic Summary (2)

New Remove

 Time Generated	Size	Description
<input type="radio"/> June 12, 2008 10:36:19 AM	686 B	Archive storage capacity utilization and tape
<input checked="" type="radio"/> May 29, 2008 6:12:22 AM	712 B	Archive storage capacity utilization and tape

Archive storage capacity utilization and tape VSN error status (June 12, 2008 6:14:14 PM)

VSNs Requiring Operator Attention

Name	Media Type	Status
AC6613	IBM 3580 LTO	Damaged Media

Stop here.

▼ Deleting a media utilization report

1. To delete a report, go to the **Media Utilization Historic Summary table** and click on the radio button corresponding to the **Time Generated for the report** (A below).

Media Utilization & Status (Historic)
Click the New button to save the current media utilization and status.

Media Utilization Historic Summary (2)

New Remove

 Time Generated	Size	Description
<input checked="" type="radio"/> June 12, 2008 10:36:19 AM	686 B	Archive storage capacity utilization and tape
<input type="radio"/> May 29, 2008 6:12:22 AM	712 B	Archive storage capacity utilization and tape

2. Press the **Remove** button (B above).

Stop here.

Assessing the distribution of files in the archive

The distribution of data files across the tiers of the archival file system is a measure of the efficiency of the policies that migrate data from the Tier-1 primary disk cache to secondary and tertiary storage. This information is valuable when you are

troubleshooting apparent capacity problems or when users complain about file access performance. Depending on the results, your archive administrators may want to alter the relative sizing of the storage tiers to better suit usage patterns.

The Infinite Archive System can graph file distributions for you. Just remember that graphs are most accurate when you collect the largest number of data points possible over the largest practical time period. Proceed as follows.

▼ Generating a graphical file-distribution report

1. If you are monitoring a high-availability, multinode appliance, use the Selected Node list control at the top of the IAS task tree to select the MDS node.
2. In the IAS task tree, expand the Appliance Metrics branch, and click the File Data Distribution hyperlink.
3. If you do not have indexed recovery points, press the Schedule Collection from File System button (A below).

Indexed recovery points are discussed in Chapter 10, “Indexing recovery point files for directory and file recovery” on page 140.

File Data Distribution

The data for metric generation comes from the indexed recovery points and (optionally) from the live file system. No data is available to generate the file data distribution metric. You can create a schedule for data collection from the live file system by clicking on the 'Schedule Collection from File System' button.

File system name: T1_01

Storage Tiers: T1_01, T2_01, T2_02

Start Date: 05/15/2008

End Date: 06/14/2008

go

Storage Tiers

Distribution of files on different tiers

No data is available to generate the metric

If you have recovery points, you can index them to collect data

You can also create a schedule to collect data from the live system

Collection Time

■ Tape Archive ■ Disk Archive ■ STK5800 Archive ■ Disk Cache

Description of legend

- If you are scheduling data collection, when the Schedule Data Collection from file system popup window appears, set the Start Date. Press the calendar button (C below).

Schedule Data Collection from file system T1_01 for Metric Generation

Schedule

Start Date: 06/14/2008 mm/dd/yyyy

Start Time: 17 : 55

Repeat Interval: Daily

Calendar

June 2008

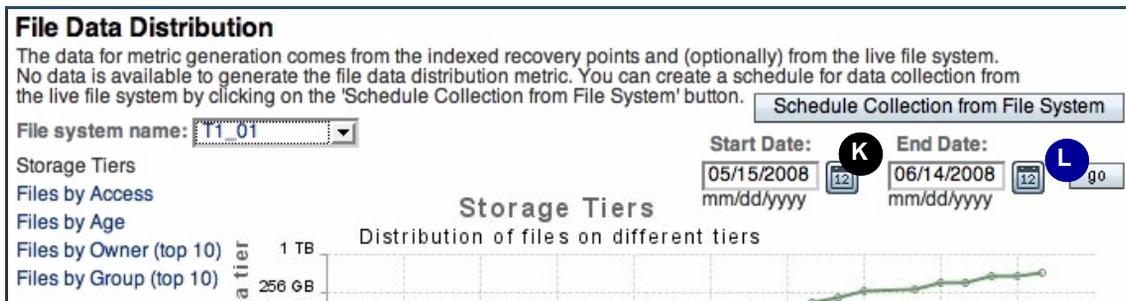
* Indicates required field

* Start Date: 06/14/2008 mm/dd/yyyy

S	M	T	W	T	F	S
1	2	3	4	5	6	7
8	9	10	11	12	13	14 <input type="button" value="E"/>
15	16	17	18	19	20	21
22	23	24	25	26	27	28
29	30	1	2	3	4	5

- If you are scheduling data collection, select a month and year from the list controls at the top of the Calendar (D above), click on a day in the calendar display (E), and press the OK button to set the date, Cancel to abort (F).
- If you are scheduling data collection, specify the hours and minutes of the Start Time (G above) and the Repeat Interval (H) using the list controls provided. Then press the Submit button to confirm the schedule, Cancel to abort (J).

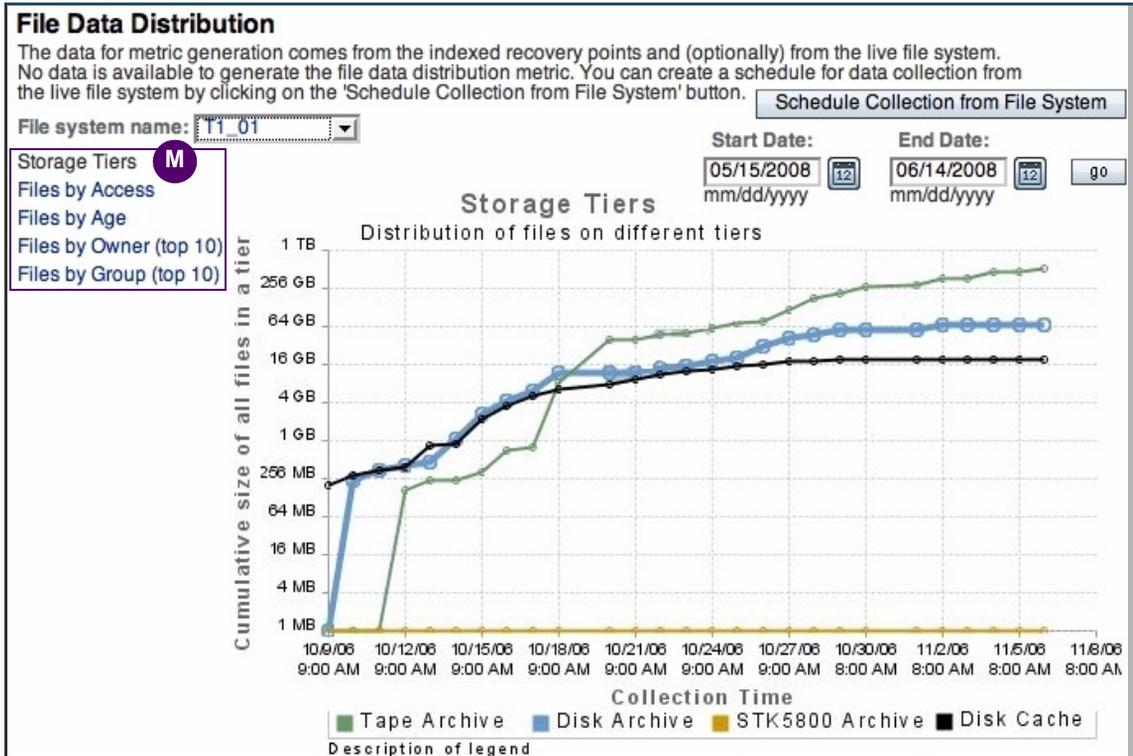
The new Start Date and End Date appear on the File Data Distribution page.



- To contract or expand the time interval displayed on the horizontal axis of the graph, use the Start Date and End Date calendar controls (K above).

8. Press the **go** button to display the output (L above).

The output graph appears as shown below. Move your mouse over points on the graph to view additional information, such as the size, number of files, files online, or files offline.

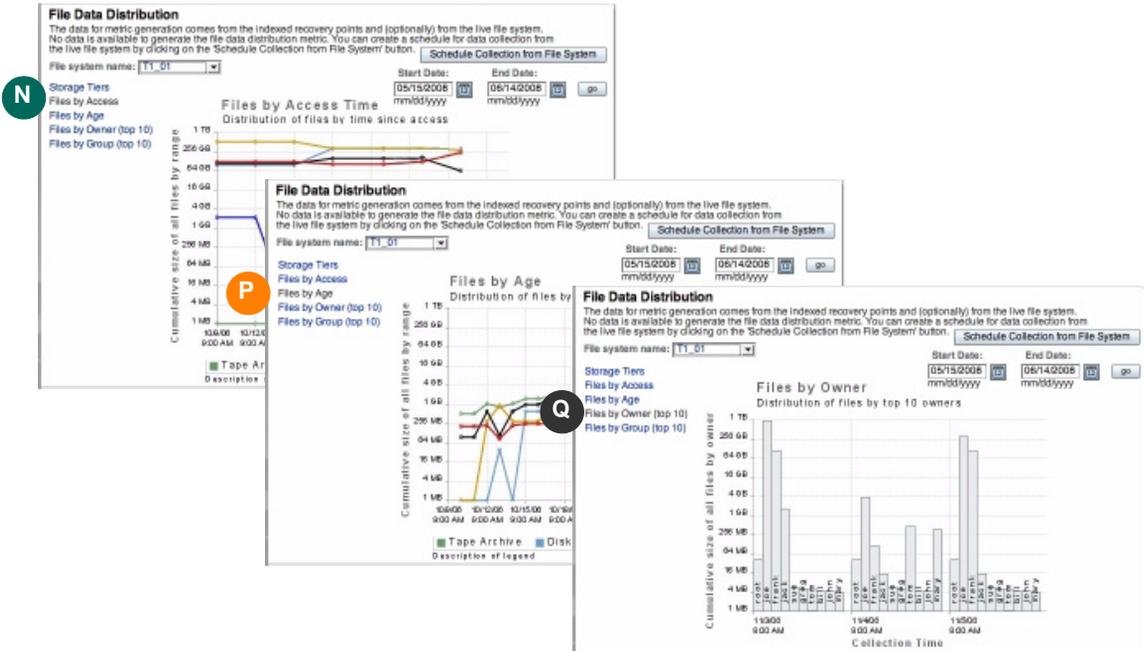


9. To see alternative data views, click the hyperlinks at the left of the graph (M above).

The following data views are available:

- Storage Tiers (above) displays a time-series plot of the total amount of storage that the file system uses in each tier of storage.
- Files by Access (N below) displays time-series plots that show the cumulative size of all files during each time interval. This graph shows how much of the file system 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 (P) displays time-series plots that show 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 (Q)** displays a bar chart showing the percentage of files owned by each of the top ten users.
- **Files by Group** displays a bar chart showing the percentage of files owned by the top ten groups.



Stop here.

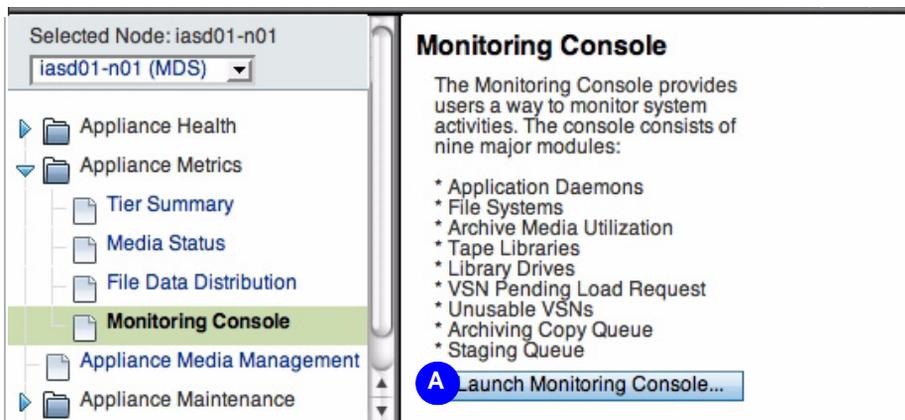
Monitoring archiving processes with the File System Manager Monitoring Console

The File System Manager Monitoring Console displays state information for core archiving system processes and services, file-system statistics, and media utilization statistics.

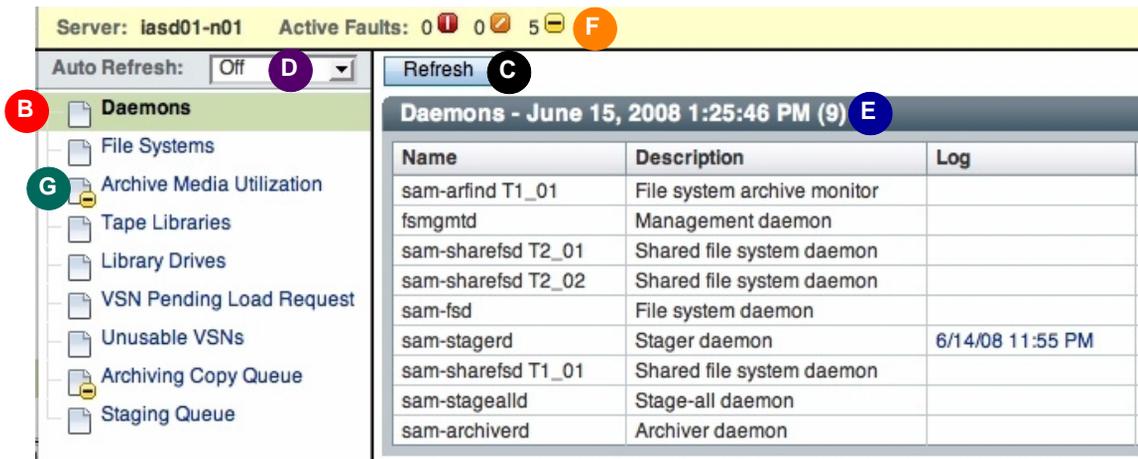
▼ Accessing the Monitoring Console

1. If you are monitoring a high-availability, multinode appliance, use the **Selected Node** list control at the top of the IAS task tree to select the MDS node.
2. In the IAS task tree, expand the **Appliance Metrics** branch, and click the **Monitoring Console** hyperlink.

- When the Monitoring Console page appears in the content pane, press the Launch Monitoring Console button (A below).



- When the Monitoring Console appears, locate the type of information that you need in the lefthand navigation tree and click on the corresponding hyperlink (B below).



Using the navigation tree, you can access monitoring panels for core archive system components (see the documentation for Sun StorageTek SAM-QFS file systems for descriptions of these components)

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

- Staging Queue

5. To refresh the display, press the **Refresh** button in the content pane (C above) or set **Auto-Refresh** to **On** using the list control at the top of the navigation tree (D). The time stamp in each table title (E) shows the last time the data was retrieved.

6. To monitor for problems, check the fault counts in the banner area of the Monitoring Console (F above) and then check for warning icons in the navigation tree (G).

In the example, four file systems are 99% full. In general, you want to keep file systems no more than 85-90% full, so you would want to bring this to the attention of someone who has the authority and training to rectify the situation and/or contact Sun support services.

Auto Refresh:

Refresh

Policy	Copy #	Type	Capacity	Usage
TESTD12	1	Disk	199.98 GB	99%
TESTD1	1	Disk	199.98 GB	99%
sampm	1	Disk	99.99 GB	99%
allfs1	1	Disk	199.98 GB	99%

7. When you are finished viewing the data, press the **Close** button at the upper right corner (J below).

Close

Description	Log	Start Time
File system archive monitor		June 5, 2008 4:09:42 PM
Management daemon		May 29, 2008 8:16:08 AM

Stop here.

Using IAS Explorer

While you can diagnose many Infinite Archive System (IAS) faults easily using information provided by the Java Web Console user interface, you may need the assistance of Sun support engineers in more complex cases. In such cases, you generally need to provide some fairly extensive information, including all of the following:

- installed software packages, revision levels, and licensing information
- system configuration and log files
- SAM Explorer output.

The IAS Explorer application makes gathering this information straightforward. The IAS Explorer automatically gathers the required files and state information and saves the result as a compressed archive file at a location that you specify.

In general, a user with `admin` authority should run IAS Explorer as soon as possible if a server node panics, crashes, core dumps, hangs, or becomes unresponsive and whenever a problem arises that requires the assistance of Sun support services. To run the IAS Explorer, proceed as follows:

▼ Running the IAS Explorer

1. In the `Appliance Health` branch of the IAS task tree, select `IAS Explorer` (A below).

IAS Explorer

You may now backup this server's state to another server. You must enter the other server's name (or IP address), a valid user name and password, and a preexistent target directory. The transfer may be accomplished using either the FTP or SCP protocol.

Host Name/IP Address: (B)

User Name: (C)

Password:

Target Directory: (D)

Use FTP or SCP: FTP (E) SCP

(F)

2. **When the IAS Explorer panel appears at right, enter the Host Name/IP Address of the host to which you want to send the output (B above).**

Note that any firewalls between the IAS appliance and the specified host address must allow communication via `ftp` (ports 21, 22) and/or `ssh` (port 23).

3. **In the User Name and Password fields (C above), enter the user ID and password that the IAS Explorer will use when accessing the server specified in the previous step.**

Note that the specified user account must have read, write, and execute (+rwx) permissions for the Target Directory specified below.

4. **In the Target Directory field (D above), enter the path to the directory where the IAS Explorer output should be saved.**

5. **Click a radio button from the Use FTP or SCP control (E above) to specify either FTP (File Transfer Protocol) or SCP (Secure Copy Protocol) for the data transmission.**

Secure Copy Protocol encrypts data for transfer via Secure Shell (SSH). File Transfer Protocol transfers data in clear text.

6. **Press Send to run IAS Explorer and transfer the data (F above).**

7. **Contact a Sun support representative, and forward the file for analysis.**

Stop here.

Using SAM Explorer

The SAM Explorer application is to archival file systems what the IAS Explorer is to the archiving system as a whole (see [“Using IAS Explorer” on page 85](#)). IAS Explorer runs SAM Explorer and includes the output—the *SAMreport* file—with the rest of its own output. So, normally, you should not need to run SAM Explorer separately. But, if it ever becomes necessary, you can do so.

A SAMreport typically includes the following:

- installed software packages, revision levels, and licensing information for the SAM software
- copies of archive and filesystem configuration files, including `mcf`, `archiver.cmd`, `recycler.cmd`, `inquiry.conf`, and `defaults.conf`

- copies of the log files that maintain a history of the state of the system, including the operating system's messages file and the archive log files `sam-log`, `archiver.log`, `recycler.log`, and `releaser.log`.
- copies of trace files (the log files that are created by individual system processes)
- core files (the contents of selected system memory dumped to disk for fault analysis).

This information is written to disk in the `/var/tmp/` or `/tmp` subdirectory, under the name `SAMreport.host.version.date`, where `host` is the name of the IAS server node, `version` is the revision level of the SAM-QFS archive file system, and `date` is the date on which SAM Explorer ran.

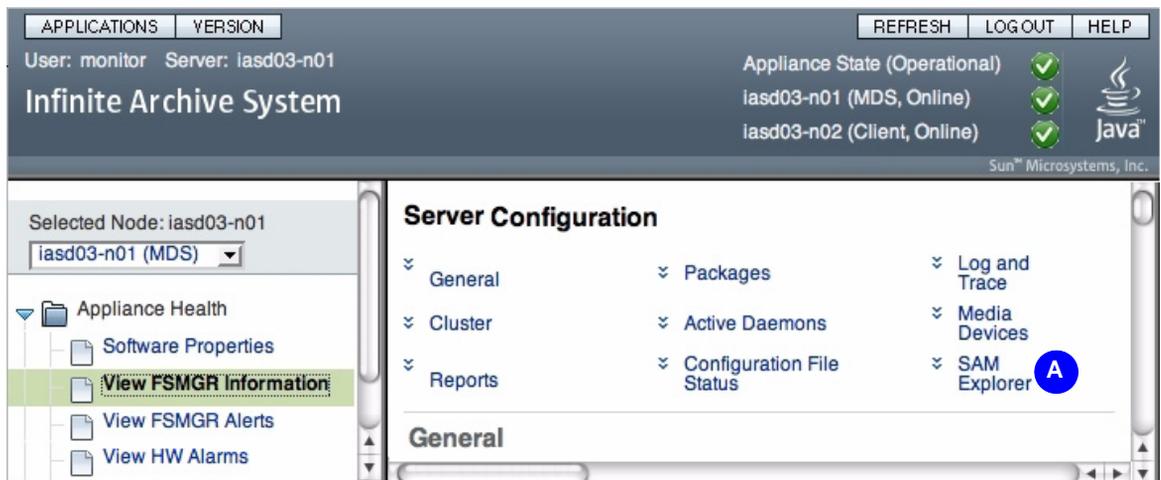
Note – Do not change the log configuration of the IAS appliance. If logs are not routinely maintained as specified by the factory configuration, critical diagnostic information will be lost.

To run SAM Explorer, proceed as follows.

▼ Running SAM Explorer

1. Log in to the IAS console as `admin`.
2. At the top of the IAS task tree, use the `Selected Node` list control to select the name of the server node that you wish to examine.
3. In the IAS task tree, expand the `Appliance Health` branch.
4. Under the `Appliance Health` branch, click `View FSMGR Information`.

This displays the `Server Configuration` page.



5. On the Server Configuration page, click on the SAM Explorer hyperlink (A above).

You jump down the page to the Sam Explorer Output table.

SAM Explorer

The screenshot shows the SAM Explorer interface. At the top, there is a 'Generate Report...' button labeled 'B'. Below it is a table titled 'SAM Explorer Output (1)' with columns: Path, Name, Size, Time Created, and Last Modified. The table contains one entry: Path: /var/tmp, Name: SAMreport.13-n01.4_6_50.Jun162008_1913 (labeled 'F'), Size: 20.96 kB, Time Created: June 16, 2008 7:14:04 PM, Last Modified: June 16, 2008 7:14:04 PM. Below the table is a 'Back to top' link. A 'Run SAM Explorer' dialog box is open, showing 'Generate Report in:' with two radio buttons: /tmp (labeled 'C') and /var/tmp. Below this is a text input field for 'Number of lines include from each log file (1-50000):' with the value '500' (labeled 'D'). At the bottom of the dialog are 'Run' (labeled 'E') and 'Cancel' buttons.

6. In the Sam Explorer Output table, click on the Generate Report Button (B above)
7. When the Run SAM Explorer popup dialog appears, click the radio button that corresponds to the directory where the report should be saved (C above)
8. In the popup dialog, enter the number of log lines you want to include in the output (D above).
9. When you are finished with the popup dialog, press the Run button to generate the report or Cancel to abort (E above). Confirm that you want to run the report when prompted.
10. To view the report, return to the Sam Explorer Output table and click the hyperlinked name of the new report (F above).
11. Contact a Sun support representative, and forward the file for analysis.

Stop here.

Sharing directories with IAS clients

Sharing directories with Network File System (NFS) or SAMBA makes Infinite Archive System (IAS) archives available to clients—the users and applications that must access or store archival information—while limiting direct access to the archive appliance itself. Shared directories are, moreover, easy to use and easy to integrate with applications. UNIX and Linux users access the directories as if they were mounted on the local filesystem, while users of Microsoft Windows access them as if they were local disk volumes.

This chapter outlines the tasks that you must carry out when setting up directory sharing:

- [“Preparing to share a directory” on page 89](#)
- [“Creating a Network File System \(NFS\) share” on page 90](#) or [“Creating a SAMBA share” on page 93](#).

Caution – File sharing has security implications that have to be considered by archive administrators, particularly when stored records are sensitive, private, or required for regulatory reasons.

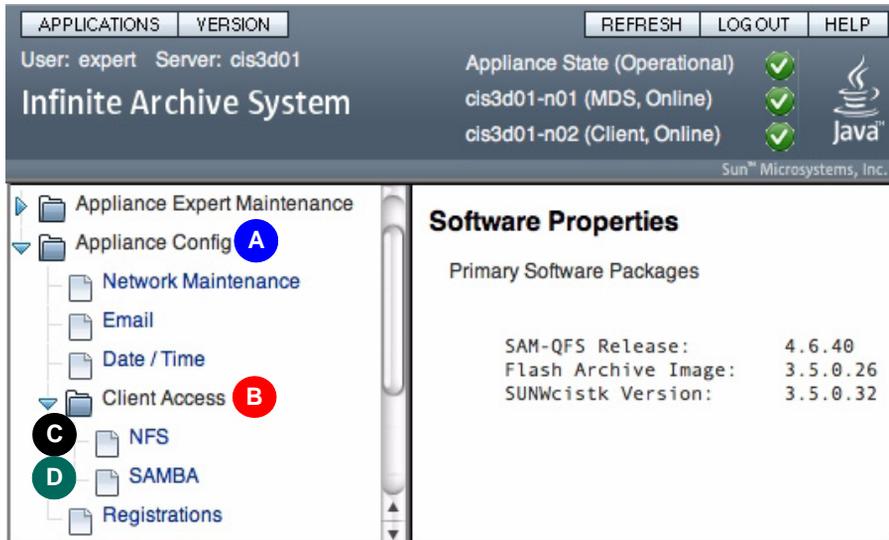
▼ Preparing to share a directory

1. **Use the IAS Java Web Console administrative interface when sharing archive directories, not SAM File System Manager (FSM) or similar tools.**

The IAS Java Web Console is fully aware of all aspects of the archive configuration. Other tools are not. For example, on high-availability, multinode IAS appliances, the IAS Java Web Console automatically configures multiple, redundant paths for each share, so that a single-point failure does not render the archive inaccessible to clients. Other tools will not.

2. **If you are working with a high-availability, multinode appliance, use the Selected Node list control at the top of the IAS task tree to select the MDS node.**

- In the IAS task tree, expand the **Appliance Config** branch (A below) and the **Client Access** branch (B).



- In the IAS task tree, click on the hyperlink for the type of file sharing you want to use: **NFS** (C above) or **SAMBA** (D).

Next task: “Creating a Network File System (NFS) share” on page 90 or “Creating a SAMBA share” on page 93.

▼ Creating a Network File System (NFS) share

- When the **NFS Exports** page appears in the IAS content pane, go to the **NFS Exports** table, and press the **Create** button (A below).



The Add **NFS** page appears.

2. When the Add NFS page appears, enter a Description that will help users to understand the purpose of the directory (B below).

The screenshot shows the 'Add NFS' configuration page. At the top right, there are 'Save' and 'Reset' buttons, with a green circle 'M' next to the 'Save' button. Below this, a note says '* Indicates required field'. The main form contains several fields and options:

- Description:** A text input field with a red circle 'B' next to it.
- * Export Directory:** A text input field with a black circle 'C' next to it. Below it, a note says '(Only Tier 1 directories can be exported)'. A red bracket 'L' is positioned to the left of this field.
- nosub:** A checkbox with a blue circle 'D' next to it. The text says 'Prevents clients from mounting subdirectories of shared directories.'
- nosuid:** A checkbox with a green circle 'E' next to it. The text says 'Checking this causes the Server File System to ignore any attempt to enable the setuid or setgid mode bits.'
- ro:** A checkbox with a purple circle 'F' next to it. The text says 'Read-Only Access.'
- Read-Only Host List:** A text input field with a pink circle 'G' next to it. The text says '(colon-separated)'.
- rw:** A checkbox with an orange circle 'H' next to it. The text says 'Read-Write Access.'
- Read-Write Host List:** A text input field with a black circle 'J' next to it. The text says '(colon-separated)'.
- root:** A checkbox with a blue circle 'K' next to it. The text says 'Root Access.'
- Root Host List:** A text input field with a blue circle 'K' next to it. The text says '(colon-separated)'.
- Is Auto Directory Created:** A checkbox with a red circle 'L' next to it.
- Auto Create Owner:** A text input field with a red circle 'L' next to it.

At the bottom right, there are 'Save' and 'Reset' buttons, with a green circle 'M' next to the 'Save' button.

3. In the text field provided, enter the Export Directory that you are going to share via NFS (C above).
4. To keep remote users from mounting subdirectories that lie beneath the exported directory, check the nosub checkbox (D above).

In general, you should be as specific as possible about what you choose to share, particularly when you are administering an archive. For this reason, most NFS administrators choose to specify nosub on exports.
5. To keep remote users from changing setuid and getgid file attribute bits, check the nosuid checkbox (E above).

In most situations, nosuid should be checked. Allowing use of the UNIX setgid and getgid commands is a potential security risk. Unscrupulous users can exploit these commands to gain privileged access to the system.
6. To grant users read-only access to an exported directory, check the ro checkbox (F above).

7. To limit a grant of read-only access to users on specified hosts, enter a colon-delimited list of hosts in the `Read-Only Host List` text box (G above).
8. To grant users read/write access to an exported directory, check the `rw` checkbox (H above).
9. To limit a grant of read/write access to users on specified hosts, enter a colon-delimited list of hosts in the `Read-Write Host List` text box (J above).
10. To allow users to log in remotely as `root` from a limited number of specified hosts, check the `root` checkbox and enter a colon-delimited list of hosts in the `Root Host List` text box (K above).

Be careful about allowing remote root logins, because they can expose the archive to security threats. Do not check this option unless necessary. If possible, allow root logins only from a limited number of secured hosts, such as servers within your data center.

11. To automatically create a share directory that does not currently exist, check the `Is Auto Directory Created` checkbox and enter the user ID of the desired directory owner in the `Auto Create Owner` text box (L above).

These options go together. If you supply one without the other, the system returns an error.

12. Press the `Save` button to make your changes, `Reset` to discard them (M above).

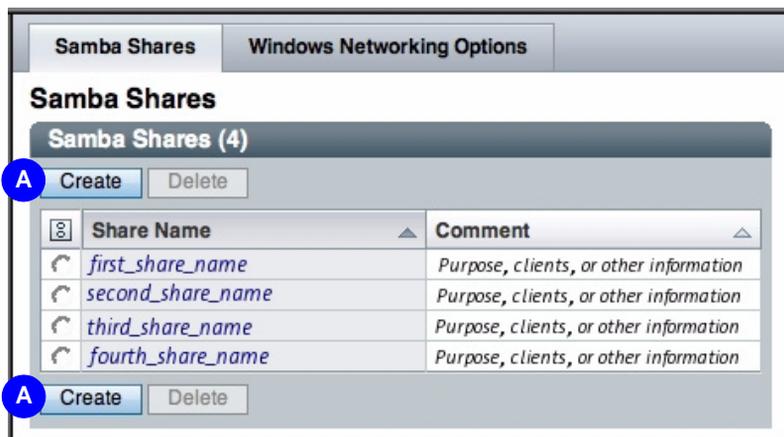
Stop here.

▼ Creating a SAMBA share

SAMBA supports directory and file sharing between UNIX-type systems (such as Sun Solaris and Linux hosts) and Microsoft Windows machines. Before configuring SAMBA file system shares, you should be thoroughly familiar with SAMBA configuration, options, and security considerations. If you are not, please consult a Sun services representative before proceeding.

Note that the information in this section is not intended to substitute for SAMBA documentation. For authoritative information on SAMBA configuration options, including those described below, consult samba.org and third-party publishers.

1. When the Samba Shares page appears in the IAS content pane, go to the Samba Shares table, and press the Create button (A below).



The Create Samba Share page appears.

2. When the `Create Samba Share` page appears, enter a `Share Name` (B below) and specify the path to the `Shared Directory` (C).

The screenshot shows a web form titled "Create Samba Share". At the top right, there are "Save" and "Reset" buttons, with a circled 'M' next to the "Save" button. A note "* Indicates required field" is located to the right of the "Is Home Directory Share" checkbox. The form contains the following fields and checkboxes, each with a letter label:

- `Is Home Directory Share:` checkbox (D)
- `* Share Name:` text input field (B)
- `* Shared Directory:` text input field (C)
- `Comment:` text input field (E)
- `Valid Users:` text input field (F)
- `Is Available:` checkbox (G)
- `Is Browseable:` checkbox (H)
- `Is Guest Ok:` checkbox (J)
- `Is Writeable:` checkbox (K)
- `Is Auto Directory Created:` checkbox (L)
- `Auto Create Owner:` text input field (L)

At the bottom right of the form, there are "Save" and "Reset" buttons, with a circled 'M' next to the "Save" button.

3. If you want SAMBA to automatically create home directories in the share when users log in for the first time, check the `Home Directory Share` (D above).

When a SAMBA user tries to access a share that does not exist, the server checks to see if its SAMBA configuration file includes a `[homes]` section. If it does, the server checks to see if the share name that the user submitted is a valid user login ID on the server. If the share name is a user ID, SAMBA creates a shared `/home/user_name` directory for the user.

4. Add a SAMBA `Comment`, if desired (E above).

The comment option can be used to provide details about the share to users that are browsing a file system. For example, you might identify a share as such, by supplying a comment string like `Shared Network Drive`, or you might want users to know that the `R:` drive share is the `Engineering Drawings Repository (EDR)`.

5. To share the directory only with specified users and/or groups, enter a space- or comma-delimited list of user and/or group IDs in the `Valid Users` text box (F above).

6. To make the share accessible to clients, check the `Is Available` checkbox (G above).

7. To let clients see the share in listings, check the `Is Browseable` checkbox (H above).

8. To create a public share that does not require a password, check the `Guest Ok` checkbox (J above).

Caution – the `Guest Ok` option has security implications that have to be considered by archive administrators, particularly when stored records are sensitive, private, or required for regulatory reasons.

9. To give users read/write access to the share, check the `Is Writeable` checkbox (K above).

Caution – by default, SAMBA shares allow read-only access. Archive administrators should take care when allowing read/write access from a network, especially if records are sensitive, private, or required for regulatory compliance.

10. To automatically create a share directory that does not currently exist, check the `Is Auto Directory Created` checkbox and enter the user ID of the desired directory owner in the `Auto Create Owner` text box (L above).

These options go together. If you supply one without the other, the system returns an error.

11. Press the `Save` button to make your changes, `Reset` to discard them (M above).

Next task: [“Setting Windows Networking Options” on page 96](#) .

▼ Setting Windows Networking Options

Note that the information in this section is not intended to substitute for SAMBA or Microsoft networking documentation. For authoritative information, consult samba.org, the Microsoft Corporation and third-party publishers.

1. Select the **Windows Networking Options** tab of the SAMBA configuration screen (A below).

Samba Shares **A** Windows Networking Options

Windows Networking Options **M** Save Reset

* Indicates required field

Workgroup: **B**

WINS Server: **C**

Server Description: **D**

Server Name: **E**

Server Aliases: **F**

Default Service: **G**

Always Show Services: **H**

* Max Reported Disk Size: **J** mB

Security: **K**

Password Server: **L**

M Save Reset

2. If desired, enter the name of a Windows Workgroup in the text box provided (B above).
3. If desired, enter the hostname or IP address of a Windows Internet Naming Service (WINS) server in the WINS Server text box (C above).
4. If desired, enter a Server Description in the text box provided (D above).
5. If desired, enter a Server Name (E above) and Server Alias (F) in the corresponding text boxes.
6. If desired, select a Default Service using the list control provided (G above).
7. If desired, tell Windows to Always Show Services that you select using the corresponding list control (H above).

8. Enter a maximum number of megabytes in the `Max Reported Disk Size` text box (J above).

This is a required parameter.

9. If desired, select a `Security` type using the corresponding list control (K above).

10. If desired, enter a hostname or IP address in the `Password Server` text box (L above).

11. Press `Save` to make your changes, `Reset` to discard them (M above).

Stop here.

Managing archival media

This chapter presents overviews of the Infinite Archive System (IAS) storage tiers and outlines the basic procedures for managing them:

- [“How IAS organizes storage media” on page 99](#)
- [“Organizing media with VSN pools” on page 101](#)
- [“Managing disk volumes” on page 104](#)
- [“Managing removable tape media” on page 105.](#)

The discussion is by no means comprehensive and is not intended as a substitute for full SAM-QFS and File System Manager (FSM) documentation. If you are not familiar with SAM-QFS and FSM, please see [“Obtaining the latest information and supporting resources” on page v](#) of this volume.

Remember, too, that IAS appliances are factory configured and optimized for the purpose. They are almost entirely self-throttling and self-tuning. So make changes only when necessary and only if you know what you are doing.

Caution – ill-considered alterations to the IAS tiered storage configuration may interfere with the functioning of the archive. Do not hesitate to call Sun support if you need assistance.

How IAS organizes storage media

The Infinite Archive System (IAS) appliance manages media using two key organizational structures:

- [“Media pools”](#) (see page 100)
- [“Hierarchical storage tiers”](#) (see page 100).

Media pools

You manage IAS storage media using *archive media pools*, sets of volumes that are administered as a group. The members of a media pool are all of the same type (random-access SAS disk, random-access SATA disk, or sequential-access tape). They are all reserved for the same purposes. So the members of the pool can be used interchangeably. This approach greatly simplifies the task of administering a large archive. You assign designated media pools to your archiving jobs instead of individually selecting volumes. The archiving software then selects an appropriate volume from the designated pool.

The IAS appliance comes with preconfigured media pools that are configured for maximum flexibility and full use of the available capacity. All of the media in the default pools are available on a first-come, first-served basis to all archive processes that require storage. This insures that the archive capacity is fully utilized and that no archiving job fails to finish for lack of media while media remain in the archive. If necessary, however, you can create and reserve media pools for specific archiving jobs, so that media are apportioned according to overall organizational priorities and policies.

Hierarchical storage tiers

IAS organizes media pools into a tiered storage hierarchy that matches the performance characteristics, capacity, and cost of each type of media to the characteristics of your data and the ways in which you use it. *Tier 1 (primary storage)* holds online data that clients and applications are actively reading and writing. *Tier 2 (secondary storage)* holds online data that clients and applications access infrequently or only in bulk. *Tier 3*—long-term archival storage—holds still less frequently accessed nearline material that is kept mainly for historical reference or data-protection reasons.

Because primary data is in active use, users may need to read or write any part of it at any time. For this reason, Tier 1 resides on random-access disk media. The exact type depends on the way that you plan to use the archive. If many users and/or application processes must read and write many small files at more or less the same time, Tier 1 resides on Serial-Attached SCSI (SAS) disks that are optimized for responsiveness and high seeking performance rather than high capacity. If a small number of users and/or processes read and write relatively small numbers of relatively large, static files—medical images, satellite photographs, and the like—Tier 1 can reside on SATA disks. The lower seeking performance and higher capacity of SATA disks better match the requirements of this type of solution.

Tier 2, secondary storage, also stores data that may be accessed at any time. But files are written and read by a small number of archive system processes—users and applications access Tier-2 data indirectly, via Tier 1 (Tier 2 stages data to Tier 1 when

requested). Tier 2 may hold backup copies of Tier-1 data files, in which case files are staged back to Tier 1 after a restore request. Or Tier-2 may store the bulk of the primary data, as in a medical imaging or surveillance solution. In this case, files are staged back to Tier-1 when a user requests the file. When Tier 2 stores backup data, high throughput, random-access SATA disks are ideal for rapidly locating and restoring files. When Tier 2 stores large, primary data files, high-throughput, sequential-access tape media may be a better choice, because, once the file is located and opened, the bulk of the data is read sequentially.

Tier-3 storage is generally included in IAS solutions that backup data from a Tier-2 SATA disk archive, either for added security (disaster recovery) or because large, primary data files are stored in Tier 2. Since Tier 3 usually stores data that is written once and seldom if ever read thereafter, efficient bulk writes and high-capacity at low-cost storage are the main requirements. So Tier-3 resides on sequential-access tape media.

The remainder of this chapter outlines basic tasks associated with maintenance and tuning of the IAS storage hierarchy:

- [“Organizing media with VSN pools” on page 101](#)
- [“Managing disk volumes” on page 104](#)
- [“Managing removable tape media” on page 105.](#)

Organizing media with VSN pools

In general, Sun recommends that you leave media management to the IAS archive system. Attempting to micromanage media allocation can reduce overall media utilization and limit effective capacity. When you reserve tapes for specific uses, the archiving software cannot utilize free space as efficiently. Some space always ends up wasted on at least some volumes. In extreme cases, archiving can stop because there is not enough space remaining on assigned volumes, even though the aggregate space remaining in the library is ample. So, for best utilization, you want to let the system use all available media for all archiving jobs.

That said, when you need to reserve media for a specified user or purpose, you can define a corresponding archive policy and media pools. You could, for example, assign one pool to the policy that archives your billing department’s files and another to the policy that archives files from human resources. You would then assign the number of volumes that each user group requires to its media pool.

You define a VSN pool by assigning it a name, a media type, and a specified set of one or more member VSNs. A *scratch pool* is a special pool of unassigned, free media that the archiver can draw upon when all of the volumes assigned to a purpose-

specific pool have been used up. If you reserve a volume in a pool, it is no longer available to the pool in which it originated. So the number of VSNs within a named pool changes as VSNs are used.

This section outlines the following procedures:

- [“Creating a VSN Pool” on page 102](#)
- [“Editing a VSN Pool” on page 102](#)
- [“Deleting a VSN Pool” on page 103.](#)

▼ Creating a VSN Pool

1. **If you are working with a high-availability, multinode appliance, use the Selected Node list control at the top of the IAS task tree to select the MDS node.**
2. **In the IAS task tree, expand the Appliance Media Management branch, and select Archive Media Pools.**

The VSN Pool Summary page appears.

3. **On the VSN Pool Summary page, in the VSN Pool Members table, press the New button.**

The New VSN Pool window pops up.

4. **In the New VSN Pool window, enter a Name for the new pool.**
5. **Select the type of media that the pool will contain from the Media Type list.**
6. **Specify VSNs: select the Start radio button and enter start and End values in the corresponding fields or select the Range of VSNs radio button and enter a comma-delimited list of ranges and/or individual VSNs (for example: EW4276-EW4288,KQ7086,FV8434).**
7. **Press the Submit button to commit your changes, Cancel to abort.**

Stop here.

▼ Editing a VSN Pool

1. **If you are working with a high-availability, multinode appliance, use the Selected Node list control at the top of the IAS task tree to select the MDS node.**
2. **In the IAS task tree, expand the Appliance Media Management branch, and select Archive Media Pools.**

The VSN Pool Summary page appears.

3. **On the VSN Pool Summary page, in the VSN Pool Members table, select the radio button corresponding to the VSN pool that you want to edit.**
4. **Press the Edit button.**
The the Edit VSN Pool window pops up.
5. **If desired, select the type of media that the pool will contain from the Media Type list.**
6. **If desired, Specify VSNS: select the Start radio button and enter start and End values in the corresponding fields or select the Range of VSNS radio button and enter a comma-delimited list of ranges and/or individual VSNS (for example: EW4276-EW4288 ,KQ7086 ,FV8434).**
7. **Press the Submit button to commit your changes, Cancel to abort.**

Stop here.

▼ Deleting a VSN Pool

You can delete a volume serial number (VSN) pool as long as none of its member VSNS are being used by a policy. To delete a VSN pool, proceed as follows.

1. **If you are working with a high-availability, multinode appliance, use the Selected Node list control at the top of the IAS task tree to select the MDS node.**
2. **In the IAS task tree, expand the Appliance Media Management branch, and select Archive Media Pools.**
The VSN Pool Summary page appears.
3. **On the VSN Pool Summary page, in the VSN Pool Members table, select the radio button corresponding to the VSN pool that you want to delete.**
4. **Press the Delete button.**
If the selected VSN pool has a member that is being used by a policy, the Delete button is unavailable.
5. **If you are sure that you want to delete the pool, confirm your decision when prompted.**

Stop here.

Managing disk volumes

In an IAS archive policy, you can archive copies of your files to disk-based volume serial numbers (VSNs) located in a tier-2 archival disk storage array. Typically, you make the entire file system into 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. This makes planning simpler.

If you do decide to create multiple VSNs within a file system, exercise caution: these are logical volumes and are not physically independent. The SAM-QFS file systems that are used on IAS appliances do not restrict a policy's use of the VSNs within a file system. So you need to be sure that you do not inadvertently create backup copies of your source files on VSNs that reside on the same physical media. You do not want a failure to cause the loss of both source files and archived copies. While the RAID architecture used in IAS disk archives offers considerable protection, you should always back up primary data to physically distinct media. In fact, it is best practice for policies to create at least *two* additional copies of primary data files on independent media, such as archival disk and/or tape.

When possible, each backup copy should itself reside on a separate type of media. This provides the maximum possible redundancy and the best protection for your data. To optimize performance, specify disk for the first copy of the source data and tape for additional copies. The SAM-QFS stager stages preferentially from copy 1. Using disk VSNs thus minimizes tape loading, improves responsiveness, increases reliability, and reduces wear-and-tear on tape drives and media.

This section addresses the following topics:

- [“Creating a disk VSN” on page 104](#)
- [“Editing the media attributes of a disk-based VSN” on page 105.](#)

▼ Creating a disk VSN

1. **If you are working with a high-availability, multinode appliance, use the Selected Node list control at the top of the IAS task tree to select the MDS node.**
2. **In the IAS task tree, expand the Appliance Media Management branch, and select Disk-Based VSNs.**
The Disk-Based VSNs Summary page appears in the content pane.
3. **Press the New button.**
The New Disk-Based VSN window pops up.
4. **Select Disk from the Media Type list.**

5. In the text field provided, enter a `pool Name` of from one to 31 characters selected from the set {A-Z0-9!"%&'()*+,-./:;<>=?_}.
6. Select the `Host server node` from the list provided.
7. If you know the fully qualified path to the directory in which you want to store archive copies, enter it in the `Path` text field. If you want to create a new directory, check the `Create Path if does not exist` checkbox.
8. If you do not know the fully qualified path to the directory in which you want to store archive copies, use the `Browse` button to locate and select a directory.
9. Press the `Submit` button to commit your changes, `Cancel` to abort.

Stop here.

▼ Editing the media attributes of a disk-based VSN

1. At the top of the IAS task tree, use the `Selected Node` list control to select the name of the server node that you wish to examine.
2. In the IAS task tree, expand the `Appliance Media Management` branch, and select `Disk-Based VSNs`.
The `Disk-Based VSNs Summary` page appears in the content pane.
3. In the `Disk-Based VSNs` table, select the radio button corresponding to the disk VSN that you want to edit, and press the `Edit Media Flags` button.
The `Edit Disk Media Flags` window pops up.
4. Check the checkboxes corresponding to the attributes that you want to modify
5. Press the `Submit` button to commit your changes, `Cancel` to abort.

Stop here.

Managing removable tape media

Infinite Archive System (IAS) appliances can include a tier of tape storage media behind Tier-1 primary disk storage or behind Tier-2 archival disk. Tape is ideal for economically storing infrequently accessed files and unusually large files, such as medical and cartographic imagery. It also provides redundancy for critical data. IAS tape media are kept in a tape library that is integrated into the appliance or maintained in a customer-supplied, external tape library.

This section starts with an overview of the role of [“Tape media in IAS solutions” on page 106](#). The remainder of the chapter then details the procedures associated with managing tape:

- [“Managing robotic tape libraries” on page 106](#)
- [“Managing tape drives” on page 109](#)
- [“Managing tape volumes” on page 112](#)
- [“Searching for tape volumes” on page 113](#).

Tape media in IAS solutions

Traditionally, tape media are used when making backup copies of files stored in traditional file systems. Tapes are periodically exported from the library and transported to a remote site—the vault—where they are safely stored for use in disaster recovery. Alternatively, tape has been frequently used for long-term off-line storage of archival data that will not be accessed again, in most cases. Tape serves a somewhat different function in an IAS solution.

In an IAS appliance, tape media are themselves part of the working file system. The SAM-QFS file systems used in IAS appliances can incorporate disk and tape resources seamlessly. Users and applications simply access the file, as they would when reading a file from disk. But, depending on the configuration, the IAS system may supply the data from disk storage, from tape, or, in some cases, from some combination of both. The tape copies are not *merely* backups, and none of the data is offline. Instead, files stored on tape are *nearline*, available for use with only slight delay. The redundancy provided by backup tapes and external software in traditional environments is an integral function of the SAM-QFS file system s used in the IAS. The file systems maintain multiple copies of file data and filesystem metadata using disparate media and synchronized to predefined recovery points.

For the above reasons, some traditional tape-management practices are not best practice in a correctly sized and provisioned IAS environment. In particular, you will not, as a rule, import or export media to any great degree, because there are no backup or vault copies: all of the data is active, either online or nearline.

Managing robotic tape libraries

Libraries are unattended, robotically controlled devices that store, load, and unload the physical tape cartridges that host data volumes (VSNs). A robotic arm transports tape cartridges between the slots that store them and one or more drives that read and write data stored on them.

IAS libraries are preconfigured and provisioned at the factory or during a Sun Professional Services engagement. They should normally function with little or no subsequent reconfiguration.

However, should it become necessary, you can use the File System Manager component of IAS to manage and modify your tape-storage subsystem. This chapter describes the following tasks:

- [“Adding a robotic tape library” on page 107](#)
- [“Changing the state of a library” on page 108](#)
- [“Unloading a robotic library” on page 108](#)
- [“Deleting a robotic library” on page 109](#)

▼ Adding a robotic tape library

Before adding a library to the browser interface, you must connect compatible equipment to the IAS appliance. Consult Sun customer services for advice and assistance.

1. **At the top of the IAS task tree, use the Selected Node list control to select the name of the server node that you wish to examine.**
2. **In the IAS task tree, expand the Appliance Media Management branch, and select Tape Libraries.**
The Tape Library Summary page appears.
3. **On the Tape Library Summary page, press the Add button.**
The Add Library wizard appears.
4. **Specify the way in which the library will attach to the IAS appliance, and press the Next button.**
 - Select `Direct Attached` if the library is connected to the appliance rather than to an intermediate storage network switch.
 - Select `STK ACSLS` if the library is connected to the appliance via ACSLS software.
 - Select `Network Attached` if the library is attached to the appliance via a storage area network (SAN).
5. **Enter any library information that the wizard requires.**
The information requested varies depending on which type of library you are adding.
6. **Review the specified information for the new library.**
7. **If you are satisfied, click Finish. Otherwise, click Previous to make changes.**

Stop here.

▼ Changing the state of a library

A library can be **ON**, **OFF**, or **UNAVAILABLE**.

When a library is **ON**, the IAS software controls the library. IAS can discover the internal state of the library, such as where tapes are located or whether barcodes are used, and can update the library catalog and other internal structures.

When a library is **OFF**, IAS does not control the library. I/O operations and automatic cartridge movement stop. But the drives in the library remain **ON**.

When a library is **UNAVAILABLE**, you can reconfigure the library and tape cartridges.

To change the state of a library, proceed as follows.

1. **If you are working with a high-availability, multinode appliance, use the Selected Node list control at the top of the IAS task tree to select the MDS node.**
2. **In the IAS task tree, expand the Appliance Media Management branch, and select Tape Libraries.**
The Tape Library Summary page appears.
3. **Select the radio button that corresponds to the library whose state you want to change.**
4. **Choose Change State from the Operations list control**
5. **When the Change State dialog pops up, use the Change State to list control to select the new state.**
6. **Press the Submit button to make the change, Cancel to abort.**

Stop here.

▼ Unloading a robotic library

When you unload a library, all the volume serial numbers (VSNs) in the library are moved to the `Historian` catalog, the library goes **OFF**, and all archiving activity stops. You can reimport VSNs from the `Historian` catalog later, if required.

Caution – be aware of the consequences before you unload a library. Unloading the library stops all archiving activity for that library.

To unload a library, proceed as follows.

1. **If you are working with a high-availability, multinode appliance, use the Selected Node list control at the top of the IAS task tree to select the MDS node.**
2. **In the IAS task tree, expand the Appliance Media Management branch, and select Tape Libraries.**
The Tape Library Summary page appears.
3. **Select the radio button that corresponds to the library that you want to unload.**
4. **Choose Unload from the Operations list control**
A message box prompts you to confirm the unload operation.
5. **Press OK to confirm the unload.**

Stop here.

▼ Deleting a robotic 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.

To delete a library, proceed as follows.

1. **If you are working with a high-availability, multinode appliance, use the Selected Node list control at the top of the IAS task tree to select the MDS node.**
2. **In the IAS task tree, expand the Appliance Media Management branch, and select Tape Libraries.**
The Tape Library Summary page appears.
3. **Select the radio button that corresponds to the library you want to delete.**
4. **Choose Delete from the Operations list control**
A message box prompts you to confirm the deletion.
5. **Press OK.**

Managing tape drives

This section provides instructions for the following tape-drive management tasks:

- [“Changing the state of a tape drive” on page 110](#)

- “Idling a drive” on page 110
- “Restarting archiving after idling a drive” on page 111
- “Unloading a drive” on page 111
- “Cleaning a drive” on page 112.

▼ Changing the state of a tape drive

A tape drive can be **ON**, **OFF**, or **UNAVAILABLE**. To change the state of a library, proceed as follows.

1. **If you are working with a high-availability, multinode appliance, use the Selected Node list control at the top of the IAS task tree to select the MDS node.**
2. **In the IAS task tree, expand the Appliance Media Management branch, and select Tape Libraries.**
The Tape Library Summary page appears.
3. **Click the radio button that corresponds to the library that holds the drive, and press the View Drives button.**
The Library and Drive Details page appears.
4. **Click the radio button that corresponds to the drive, and press the Change State button.**
5. **When the Change State dialog pops up, use the Change State to list control to select the new state.**
6. **Press the Submit button to make the change, Cancel to abort.**

Stop here.

▼ Idling a drive

Idling a drive stops any archiving or staging activity that is in process so that you can manipulate cartridges in a library without disturbing archiving or staging operations.

To idle a drive, proceed as follows.

1. **At the top of the IAS task tree, use the Selected Node list control to select the name of the server node that you wish to examine.**
2. **In the IAS task tree, expand the Appliance Media Management branch, and select Tape Libraries.**
The Tape Library Summary page appears.

3. Click the radio button that corresponds to the library that holds the drive, and press the `View Drives` button.

The `Library and Drive Details` page appears.

4. Click the radio button that corresponds to the drive, and press the `Idle` button.

Next task: [“Restarting archiving after idling a drive” on page 111.](#)

▼ Restarting archiving after idling a drive

To restart archiving after idling a drive, proceed as follows:

1. If you are working with a high-availability, multinode appliance, use the `Selected Node` list control at the top of the IAS task tree to select the MDS node.
2. Expand the `Appliance Maintenance` branch of the IAS task tree.
3. Expand the `Archive Maintenance` node of the `Appliance Maintenance` branch.
4. Select `Archiver Activity`.
5. When the `Archive Activity` page appears, click the `Restart` radio button under `Perform the following archiving action`, and press the `Submit` button.

Stop here.

▼ Unloading a drive

If a volume is no longer needed, it is automatically unloaded. However, you can manually unload a drive if necessary. When you unload a drive, the mounted cartridge is removed from the drive and placed in its storage slot.

To unload a drive, proceed as follows.

1. If you are working with a high-availability, multinode appliance, use the `Selected Node` list control at the top of the IAS task tree to select the MDS node.
2. In the IAS task tree, expand the `Appliance Media Management` branch, and select `Tape Libraries`.
The `Tape Library Summary` page appears.
3. Click the radio button that corresponds to the library that holds the drive, and press the `View Drives` button.

The `Library and Drive Details` page appears.

4. Click the radio button that corresponds to the drive, and press the `Idle` button.
You must idle the drive before unloading it.
5. Re-click the radio button that corresponds to the drive, and select `Unload from the Operations` pull-down menu.

Stop here.

▼ 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, proceed as follows.

1. If you are working with a high-availability, multinode appliance, use the `Selected Node list control` at the top of the IAS task tree to select the MDS node.
2. In the IAS task tree, expand the `Appliance Media Management` branch, and select `Tape Libraries`.
The `Tape Library Summary` page appears.
3. Click the radio button that corresponds to the library that holds the drive, and press the `View Drives` button.
The `Library and Drive Details` page appears.
4. Click the radio button that corresponds to the drive, and press the `Clean` button.

Stop here.

Managing tape volumes

This section provides instructions for the following tape-management tasks:

- [“Searching for tape volumes” on page 113](#)
- [“Importing tape cartridges” on page 113](#)
- [“Exporting tape cartridges” on page 114](#)
- [“Loading and unloading tape drives” on page 115](#)
- [“Relabeling tape volumes” on page 117](#)
- [“Reserving tape volumes” on page 118](#)
- [“Auditing slots and updating the library catalog” on page 120](#)
- [“Changing tape media attributes” on page 121](#)
- [“Managing exported tapes with the file system Historian” on page 122.](#)

Searching for tape volumes

Tape volumes are identified by volume serial numbers (VSNs) that you search from the IAS console interface. Proceed as follows.

▼ Searching for a volume by VSN

- 1. If you are working with a high-availability, multinode appliance, use the Selected Node list control at the top of the IAS task tree to select the MDS node.**
- 2. In the IAS task tree, expand the Appliance Media Management branch, and select Tape VSNs.**
The Tape VSN Summary page appears.
- 3. Press the Search for VSN button.**
The Search Tape VSNs page appears.
- 4. On the Search Tape VSNs page, enter the volume serial number for the volume that you are trying to locate in the VSN text box, and press the Search VSN button.**
If an exact match is found, the VSN Details page displays the particulars of the volume. If several matches are found, a Search Results page displays.

Stop here.

Importing tape cartridges

When a volume is imported, a robotic arm moves the cartridge from the import/export slot to the first available slot in the library.

Generally, a properly sized and provisioned Infinite Archive System (IAS) appliance should not require additional tapes under normal circumstances and, when new tapes are necessary, the IAS File System Manager software can automatically import the volumes from the mail slots of a compatible direct-attached library. If volumes cannot be automatically imported, however, you can also import them manually, using the procedure below.

▼ Importing volumes into the tape library

- 1. Place the cartridge(s) into the cartridge access port (CAP) on the library.**

2. **If you do not have the IAS Java Web Console graphical user interface running, start a compatible web browser, such as Mozilla Firefox, and enter the URL `https://ip_address:6789`, where `ip_address` is the address of the private management interface on the IAS server node and `6789` is the port number.**

On single-node appliances, use the IP address `192.168.51.31`.

On multi-node, high-availability appliances, use the IP address of the SAM-QFS Meta Data Server (MDS) node. The MDS node is usually host `IAS3N-n01` (`192.168.51.31`) on a new installation, where *N* is the appliance number. The client node is usually `IAS3N-n02` (`192.168.51.32`).

3. **If you are configuring a high-availability appliance, at the top of the IAS task tree, use the `Selected Node` list control to select the Meta Data Server (MDS).**
4. **In the IAS task tree, expand the `Appliance Media Management` branch, and select `Tape Libraries` hyperlink.**

The `Tape Library Summary` page appears.

5. **Select the radio button next to the name of the library into which you want to import volumes, and press the `Import` button.**
 - If the `SAMST` library driver is configured the appropriate volumes are automatically imported.
 - If `ACSL` is configured, the `Specify VSN Name` window appears.
6. **If `ACSL` is configured, when the `Specify VSN Name` window appears, specify either a starting VSN and an ending VSN or a single VSN.**
7. **Press the `Submit` button.**

Stop here.

Exporting tape cartridges

Exporting a tape volume removes the cartridge from the library and records the change in the SAM-QFS Historian database. For direct attached libraries, at least one mail slot is needed when exporting a tape volume.

Generally, you should not need to export tapes from an Infinite Archive System (IAS) library appliance. Sun does not recommend routinely removing media from the IAS tape library. The system is designed to keep all archival data managed and available at all times. The tape media are part of the active file system, not simply static backup copies. Maintaining tapes in the library greatly reduces operator intervention, improves access to stored data, and minimizes opportunities for loss or damage to tape cartridges.

However, when necessary, you can export a volume using the procedure below. If you later require data stored on an exported tape volume, you can locate it using information stored by the SAM-QFS Historian.

▼ Exporting a tape volume from the library

- 1. If you are working with a high-availability, multinode appliance, use the Selected Node list control at the top of the IAS task tree to select the MDS node.**
The Selected Node list control appears.
- 2. In the IAS task tree, expand the Appliance Media Management branch, and select Tape Libraries.**
The Tape Library Summary page appears.
- 3. Click the radio button corresponding to the library from which you want to export the volume, and press the View VSNs button.**
The Tape VSN Summary page appears.
- 4. Click the radio button corresponding to the slot that holds the volume that you want to export.**
- 5. Select Export from the Operations list control.**
The cartridge is ejected from the library.

Stop here.

Loading and unloading tape drives

Normally, Sun recommends that you leave tape movement to the IAS software. When a tape volume is needed, either for archiving or for staging data requested by a filesystem user, the SAM-QFS file system automatically loads the volume into a drive in the library. Direct operator intervention complicates operation of the archive and introduces opportunities for error.

Nonetheless, you can initiate this operation manually, if necessary. The software then determines the drive into which the volume is loaded. Proceed as follows.

- [“Loading a tape volume into a drive in the library” on page 116](#)
- [“Unloading a volume from a drive in the library” on page 116.](#)

▼ Loading a tape volume into a drive in the library

1. **If you are working with a high-availability, multinode appliance, use the Selected Node list control at the top of the IAS task tree to select the MDS node.**
2. **In the IAS task tree, expand the Appliance Media Management branch, and select Tape Libraries.**
The Tape Library Summary page appears.
3. **Select the radio button corresponding to the library into which you want to load the volume, and press the View VSNS button.**
The VSN Summary page appears.
4. **Click the radio button corresponding to the slot that holds the volume that you want to load.**
5. **Select Load from the Operations list control.**
The cartridge is drawn into the drive unit.

Stop here.

▼ Unloading a volume from a drive in the library

1. **If you are working with a high-availability, multinode appliance, use the Selected Node list control at the top of the IAS task tree to select the MDS node.**
2. **In the IAS task tree, expand the Appliance Media Management branch, and select Tape Libraries.**
The Tape Library Summary page appears.
3. **Select the radio button corresponding to the library that holds the drive, and press the View Drives button.**
The Library and Drive Details page is displayed.
4. **Select the radio button corresponding to the drive that holds the volume that you want to unload, and press the Unload button.**
A message tells you whether the unload operation has been successfully issued.

Stop here.

Relabeling tape volumes

Sun recommends against manually relabeling media. Labeling a volume reinitializes it and renders any data stored inaccessible. This may make sense when you are reusing an unmanaged backup tape. But the tape media in an IAS library are part of the active file system, not simply static backup copies. If you manually reinitialize part of this file system, you may disrupt the normal data-protection mechanisms built in to a SAM-QFS file system and may lose data.

You can, of course, preserve the data on a volume that will be relabeled by making a copy before proceeding. But you must be able to do so without the help of the IAS console: this is a purely external process. The copy will also be of limited use. The data that it contains will no longer be synchronized with other copies in the file system and SAM-QFS will not recognize the volume serial number of the external copy.

Nonetheless, if special circumstances make it necessary, you can relabel a tape using the procedure below.

▼ Relabeling a tape volume

- 1. Outside of the browser interface, make a copy of the volume if you need to preserve data.**

Be aware that externally created copies are of limited use. The data are no longer synchronized with other copies in the file system and SAM-QFS cannot recognize the volume serial numbers of the external copies.

- 2. If you are working with a high-availability, multinode appliance, use the Selected Node list control at the top of the IAS task tree to select the MDS node.**
- 3. In the IAS task tree, expand the Appliance Media Management branch, and select Tape Libraries.**

The Tape Library Summary page appears.

- 4. To locate the volume by volume serial number (VSN), press the Search for Tape VSN button. When the Search Tape VSNs page appears, enter the VSN in the text box provided, and press the Search Tape VSN button. When the VSN Details page appears, go to step 7.**
- 5. To select the volume by library slot, select the radio button corresponding to the library that holds the volume, and press the View VSNs button.**

The VSN Summary page appears.
- 6. From the VSN Summary pages, select the radio button corresponding to the slot that holds the volume that you want to relabel.**

7. Press the Label button.

The Label Tape window appears.

8. When the Label Tape popup window appears, choose a Label Type by selecting the Relabel radio button unless the library lacks a barcode reader, in which case you should select the Label radio button.

9. Enter a Label Name in the text box provided.

Label names must be from one to six characters, each selected from the set A-Z0-9!"%&'()*+,-./:;<=>?_.

10. Select a Block Size using the list control provided.

The block size determines the basic unit of online storage for the volume.

11. Press the Submit Job button to commit your changes, Reset to clear your entries, or Cancel to abort.

12. If you intend to proceed, click OK when you are prompted for confirmation.

A message tells you whether the label operation was issued successfully. If it was, a tape label job is created. After the job completes, the VSN is updated with the new label.

Stop here.

Reserving tape volumes

In general, Sun does not recommend reserving volumes in an IAS archive. Typically, backup copy/vault applications reserve tape volumes for data sets that will be exported from the library and vaulted at a remote site. Reserving tape volumes for specific jobs keeps related data on contiguous volumes and minimizes the number of tapes that are shuttled between the data center and the vault. IAS is not a backup application, however. IAS tape volumes are part of an active file system that should remain under automated IAS management at all times. If media are not exported from the library, reserving them offers no real benefit.

But reserving tape volumes can cause problems. It can reduce overall media utilization and limit effective capacity. When tapes are reserved, some space is always wasted on at least some volumes, because the archiving software cannot use all of the free space equally. In extreme cases, archiving may stop because there is not enough space remaining on assigned volumes, even though the aggregate space remaining in the library is ample.

That said, if special circumstances dictate, you can use the procedures below to reserve volumes for specific purposes:

- [“Reserving a tape volume” on page 119](#)
- [“Cancelling a tape volume reservation” on page 120.](#)

▼ Reserving a tape volume

- 1. If you are working with a high-availability, multinode appliance, use the Selected Node list control at the top of the IAS task tree to select the MDS node.**
- 2. In the IAS task tree, expand the Appliance Media Management branch, and select Tape Libraries.**
The Tape Library Summary page appears.
- 3. To locate the volume by volume serial number (VSN), press the Search for Tape VSN button. When the Search Tape VSNs page appears, enter the VSN in the text box provided, and press the Search Tape VSN button. When the VSN Details page appears, go to step 6.**
- 4. To select the volume by library slot, select the radio button corresponding to the library that holds the volume, and press the View VSNs button.**
The VSN Summary page appears.
- 5. From the VSN Summary pages, select the radio button corresponding to the slot that holds the volume.**
- 6. Press the Reserve VSN button.**
The Reserve VSN wizard starts.
- 7. Check the checkboxes corresponding to one or more of the reservation methods described below, and press the Next button.**
 - File System reserves the volume for archiving files from a specific file system.
 - Archive Policy reserves the volume for archiving files that are associated with a specified archive policy.
 - Owner reserves the volume 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.
- 8. When the wizard prompts you, provide all information required, and press the Next button.**
Depending on your selections in the previous step, the wizard prompts you for one or more of the following:
 - a File System Name (select from the list provided).
 - an Archive Policy (select from the list provided).
 - an Owner, Group, or Directory (select and supply the user name, owner name, or directory path in the corresponding text field).

9. Press the `Finish` button to reserve the volume, `Previous` to return to the preceding step in the wizard, or `Cancel` to abort.

Stop here.

▼ Cancelling a tape volume reservation

To remove a reservation and return a tape volume to unrestricted general use, proceed as follows.

1. **If you are working with a high-availability, multinode appliance, use the `Selected Node` list control at the top of the IAS task tree to select the MDS node.**
2. **In the IAS task tree, expand the `Appliance Media Management` branch, and select `Tape Libraries`.**

The `Tape Library Summary` page appears.

3. **To locate the volume by volume serial number (VSN), press the `Search for Tape VSN` button. When the `Search Tape VSNS` page appears, enter the VSN in the text box provided, and press the `Search Tape VSN` button. When the `VSN Details` page appears, go to step 6.**
4. **To select the volume by library slot, select the radio button corresponding to the library that holds the volume, and press the `View VSNS` button.**

The `VSN Summary` page appears.
5. **From the `VSN Summary` pages, select the radio button corresponding to the slot that holds the volume.**
6. **Press the `Unreserve VSN` button.**

Stop here.

Auditing slots and updating the library catalog

Occasionally, the space actually remaining on a tape differs from that reported in the library catalog. Auditing the slot corrects this variance. When you audit a slot, the library loads the volume specified, reads the label, and updates the library catalog to reflect the actual capacity.

▼ Auditing a library slot

To audit the slot for a volume in a library, proceed as follows.

- 1. If you are working with a high-availability, multinode appliance, use the Selected Node list control at the top of the IAS task tree to select the MDS node.**
- 2. In the IAS task tree, expand the Appliance Media Management branch, and select Tape Libraries.**
The Tape Library Summary page appears.
- 3. Select the radio button corresponding to the library that holds the volume, and press the View VSNs button.**
The VSN Summary page appears.
- 4. From the VSN Summary pages, select the radio button corresponding to the slot that holds the volume.**
- 5. Using the Operations list control, select Audit.**
The library catalog is updated with the slot capacity information.

Stop here.

Changing tape media attributes

In general, you should leave management of media attributes to the automated IAS archiving software. You should seldom if ever need to change media attributes, and regular use of read-only and write-protected attributes will reduce the capacity of the archival file system and may interfere with regular archiving.

Nonetheless, if it becomes necessary, you can edit the media attributes for a volume using the procedure below:

▼ Editing the media attributes of a tape volume

- 1. If you are working with a high-availability, multinode appliance, use the Selected Node list control at the top of the IAS task tree to select the MDS node.**
- 2. In the IAS task tree, expand the Appliance Media Management branch, and select Tape Libraries.**
The Tape Library Summary page appears.

3. **To locate the volume by volume serial number (VSN), press the Search for Tape VSN button. When the Search Tape VSNs page appears, enter the VSN in the text box provided, and press the Search Tape VSN button. When the VSN Details page appears, go to step 6.**
4. **To select the volume by library slot, select the radio button corresponding to the library that holds the volume, and press the View VSNs button.**
The VSN Summary page appears.
5. **From the VSN Summary pages, select the radio button corresponding to the slot that holds the volume.**
6. **Press the Edit Media Attributes button.**
The Edit Media Attributes window pops up.
7. **Check the check boxes corresponding to the attributes that you want to change.**
The following attributes can be modified:
 - Damaged Media flags a suspect volume.
 - Duplicate VSN indicates that there is another volume with the same Volume Serial Number.
 - Read-only makes the volume read-only.
 - Write-protected makes the volume unwritable.
 - Foreign media flags the volume as a non-SAM-QFS volume created by some other application.
 - Recycle flags the volume for recycling.
 - Volume is Full indicates that the volume has no space available for storage.
 - Unavailable indicates that the volume cannot be used for storage at present.
 - Need Audit indicates that the library catalog needs to be updated with the actual slot capacity for this volume.
8. **Press the OK button to make your changes, Reset to clear your selections and start over, or Cancel to abort.**

Stop here.

Managing exported tapes with the file system Historian

Sun does not recommend routinely exporting media from the IAS tape library. The tape media are part of the active SAM-QFS file system, not simply static copies, and should thus be under IAS management and available for use at all times. However, should circumstances justify exporting tape volumes, the SAM-QFS file system *Historian* can help you to later identify and retrieve cartridges if needed.

The Historian is a catalog that maintains entries for tape cartridges that have been exported from the IAS robotic library. Logically, the SAM-QFS software represents the Historian as a virtual library and lists it with the physical libraries on the `Tape Library Summary` page of the IAS console. This virtual library stands in for the external, physical repositories (offsite vaults, other libraries, etc.) that hold cartridges that have been exported from the library.

If the file system ever receives a request for data stored uniquely on an exported tape cartridge, the Historian supplies the minimum information that you need in order to identify the required physical volume. You can then request the volume from your external repository and reimport it into the IAS physical library.

Maintaining the archive file systems

On a properly installed Infinite Archive System (IAS) appliance, all normal interactions between clients (users and applications) and archived files are handled indirectly, using network file sharing (see [“IAS appliance best practices” on page 5](#)). The IAS internal file systems are never directly exposed to clients and should thus never need modification. In fact, alterations to these file systems may damage the appliance and may cause data loss. So you should not, in general, use the SAM-QFS File System Manager (FSM) to modify IAS file systems.

However, in rare cases, you may need to correct file system corruption, mount/unmount internal file systems, or create a one-time backup of a filesystem configuration. This chapter outlines these topics in the following sections:

- [“Locating and repairing file corruption” on page 126](#)
- [“Mounting and unmounting file systems” on page 126](#)
- [“Saving a point-in-time file system image” on page 128](#).

At a minimum, you should fill in the bare outlines presented here by consulting the *Sun StorEdge SAM-FS Software Configuration and Administration Guide* and the *Sun StorEdge QFS Software Configuration and Administration Guide*. See [“Obtaining the latest information and supporting resources” on page v](#) for instructions on obtaining these documents.

Caution – This chapter is not intended to be comprehensive and cannot substitute for specialized Sun StorageTek SAM-QFS training, documentation, and expertise. If you lack these prerequisites, if you encounter filesystem errors, or if filesystem maintenance other than that described here seems necessary, engage Sun support and/or Sun Professional Services. Do not attempt the work on your own.

Locating and repairing file corruption

SAM-QFS file systems should not normally require repairs. But, if they do, the SAM-QFS file-checking process identifies and, if desired, repairs file corruption in disk partitions that belong to the specified file system. Proceed as follows.

▼ Checking file systems

1. **If you are working with a high-availability, multinode appliance, use the Selected Node list control at the top of the IAS task tree to select the MDS node.**
2. **In the IAS task tree, expand the Appliance Expert Maintenance branch.**
3. **In the Appliance Expert Maintenance branch, select File System.**
The File Systems Summary page appears in the content pane.
4. **On the File Systems Summary page, select the radio button that corresponds to file system that you want to check.**
5. **Select Check File System from the Operations list control.**
The Check File System window pops up.
6. **Select a file-check option**
You can either check the file system or, if the file system is unmounted, check the file system and automatically make repairs.
7. **Specify a location to which the file system check report will be sent.**
8. **Press the Submit button to commit your changes, or Cancel to abort.**

Stop here.

Mounting and unmounting file systems

Mounting a file system attaches it to the file system hierarchy at a designated directory called the *mount point*. Mounting a file system makes its contents accessible to users and applications, while unmounting it makes the contents completely inaccessible. *Mount options* control how a file system is accessed and how it will behave when mounted or unmounted.

This section documents the following procedures:

- “Mounting or unmounting a file system” on page 127
- “Editing mount options” on page 127

▼ Mounting or unmounting a file system

To mount a file system, proceed as follows

1. **If you are working with a high-availability, multinode appliance, use the Selected Node list control at the top of the IAS task tree to select the MDS node.**
2. **In the IAS task tree, expand the Appliance Expert Maintenance branch.**
3. **In the Appliance Expert Maintenance branch, select File System.**
The File Systems Summary page appears in the content pane.
4. **On the File Systems Summary page, select the radio button that corresponds to file system that you want to mount or unmount.**
5. **Select Mount or Unmount from the Operations list control.**

Stop here.

▼ Editing mount options

Caution – do not proceed unless you have IAS expert authority, UNIX systems administration experience, and extensive training in and experience with SAM-QFS.

Exercise care when editing mount options for IAS file systems. IAS appliances come fully preconfigured and optimized for your archiving needs. Changes should not normally be necessary. If, however, changes prove necessary, you can edit filesystem mount options using the IAS console.

Before proceeding, note the following general points:

- Some changes cannot take effect until you unmount and then remount the target file system.
- When the file system is mounted, some mount options are unavailable.
- If you want to change a mount option value to the default value, delete the existing value and save the changes, leaving the modified fields empty.
- 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 server node's `/etc/vfstab` configuration file.

For more information, see the *Sun StorageTek QFS Software Installation and Upgrade Guide* or the *Sun StorageTek Storage Archive Manager (SAM) Installation and Upgrade Guide* (for information on obtaining these documents, see [“Obtaining the latest information and supporting resources” on page v](#)).

Then edit mount options using the procedure below.

1. **If you are working with a high-availability, multinode appliance, use the Selected Node list control at the top of the IAS task tree to select the MDS node.**
2. **In the IAS task tree, expand the Appliance Expert Maintenance branch.**
3. **In the Appliance Expert Maintenance branch, select File System.**
The File Systems Summary page appears in the content pane.
4. **On the File Systems Summary page, select the radio button that corresponds to file system for which you want to edit mount options.**
5. **Using the Operations list control, select Edit Mount Options.**
The file system Edit Mount Options page pops up.
6. **Edit the fields that you want to change.**
7. **Press the Save button to commit your changes, Cancel to abort.**

Stop here.

Saving a point-in-time file system image

All IAS file systems are automatically protected by system software that periodically saves the filesystem metadata to a *recovery point* file (see [Chapter 10, “Maintaining the archive file systems” on page 125](#)). But you may at times wish to create unscheduled recovery points of your own—at the close of a quarter or fiscal year, for instance. If so, you can manually create a one-time recovery point using the following procedure.

▼ Creating a recovery point on demand

Please note that the creation of a recovery point can be a resource-intensive, high-overhead operation. So, if possible, carry out this procedure at a time when the archive is not heavily used.

1. If you are working with a high-availability, multinode appliance, use the Selected Node list control at the top of the IAS task tree to select the MDS node.

2. In the IAS task tree, expand the Appliance Expert Maintenance branch.

3. In the Appliance Expert Maintenance branch, select Metadata Maintenance.

The Recovery Points page appears in the content pane.

4. Select the Current File System from the list control supplied, and press the Create Recovery Point Now button located at the top of the Recovery Points for File System *filesystem_name* table.

The Take Recovery Point page pops up.

5. In the Fully Qualified Recovery Point File Name field, type a full path and name for the recovery point file that you want to create or accept the default shown in the field.

6. Press the Submit button to commit your changes, Cancel to abort.

Creating a recovery file takes about one hour for every 5-10 gigabytes of data, and the size of the resulting file will be about 666 multiplied by the number of files in the file system.

Stop here.

Restoring a file system

Note – the procedures in this chapter require IAS `expert` authority. Users should also have UNIX systems administration experience and a sound working knowledge of SAM-FS/QFS. If you lack this degree of expertise, please consider engaging Sun professional Services.

Infinite Archive System appliances automatically protect your data against loss and corruption. In most cases, you can restore a lost or damaged file system, directory, or file as it was prior to the loss. This chapter briefly explains the IAS mechanisms that are involved in [“Protecting file systems and their contents”](#), then moves on to describe the actual procedures:

- [“Restoring a file system”](#) on page 134
- [“Restoring directories and files”](#) on page 140.

Protecting file systems and their contents

This section provides an overview of the mechanisms that protect the file systems in an IAS archive. It starts with a brief overview of the problem, [“Protecting both data and metadata”](#). It then explains that IAS appliances address this problem by [“Automating administration with SAMPM”](#) on page 132.

Protecting both data and metadata

To protect a file system, you have to do two things: you must protect the data stored in the files and you must protect the filesystem metadata that makes the files and their contents accessible. The preconfigured IAS archiving process protects data by continuously duplicating changed files to physically independent media (a separate

disk array and or tape). But you cannot reconstitute individual files, directories, and/or file systems with data alone. You also need the file paths, inodes, directories, access control lists (ACLs), symbolic links, and pointers that tell you where data are stored on media and how they are organized into meaningful units. If you lose this information, you lose the filesystem, because any surviving data remains permanently inaccessible.

In traditional implementations of Sun StorageTek Storage Archive Manager (SAM) and SAM-FS/SAM-QFS file systems, administrators protect the metadata in two ways: by periodically running `samfsdump` and by enabling *archive logging*.

The `samfsdump` filesystem utility saves a snapshot of the metadata to a compressed archive file called a *recovery point* or *dump file*. Together, recovery point files and archive copies give the administrator everything needed to recover a single file or an entire file system. If he saves recovery points and archives file data regularly, he can, in effect, roll back a file system or file to a moment when it had not yet been deleted, damaged, or modified.

Archive logs record the particulars of the file metadata and its storage locations as files are archived. In principle, the administrator can recreate the files using the information in the log. But in practice, the process is so time consuming that it is seldom used unless recovery point files are not available. For example, if files were lost shortly before `samfsdump` was next scheduled to run, the administrator might use the logs to recover the items. But he would not waste time using the logs to recover files that were included in the last recovery point.

Direct, manual administration of a SAM system and associated archive can produce excellent results. But archive administration is a demanding task that requires specialized training, considerable expertise, and time. The administrator has to schedule both recovery point creation and archiving in such a way that both parts of a restorable file system—data and metadata—will be available and in synch when they are needed. Individuals with the required skills are not always available within a given organization, and, when they can be found, they are always overtaxed. Their time is valuable, and routine archive management may not be their most pressing assignment.

Automating administration with SAMPM

IAS appliances automate the potentially complex, labor-intensive administrative tasks associated with SAM and SAM-FS/QFS file systems. The SAM archiver is preconfigured to archive continuously, so that copies are made as soon as a significant number of data files have changed. To protect metadata, IAS includes special utility software: SAM Preventative Maintenance (SAMPM).

For each Tier-1 file system configured on an IAS appliance, SAMPM runs `samfsdump` nightly, so that a new recovery point file is created every 24 hours, and automatically manages the space consumed by recovery points. This schedule protects filesystem metadata without adversely affecting active data operations (more frequent creation of recovery points can impose significant overhead and may consume excessive amounts of disk space). If there are multiple Tier-1 file systems, SAMPM runs `samfsdump` in parallel by default. Sun Professional Services may also configure SAMPM to run in parallel on a single file system when this would be advantageous.

SAMPM saves the compressed recovery point file to a subdirectory of the appliance's `root` file system, `/global/CIS-SAMFSDUMP/` (the default). This lets you recover following the loss of an entire Tier-1 file system. Space in the `root` file system is limited, so only a few of the most recent recovery points are retained in this directory. The older recovery points are deleted.

Deleting the older recovery points makes sense for disaster recovery, when you generally want only the most recent files. But for routine recovery of individual directories or files, you need a wider recovery window than a few recovery points can provide. For this reason, SAMPM copies every recovery-point file to a subdirectory of the mount point for each Tier-1 file system. Within this subdirectory, SAMPM stores recovery-point files in a directory structure organized by year, month, and server-node name. By default, the path would look like this:

```
/T1_01/sampm/year/month/server-node_name/
```

Space is not as critical in the Tier-1 file system, so you can retain enough recovery point files to roll back a file or directory to almost any point in its prior history.

Since there is a 24-hr window between the completion of one recovery point and the creation of the next, you can lose some data under some circumstances. Data that is created and lost within this window cannot be recovered from the recovery point file. How much data is lost thus depends on how rapidly data changes relative to the intervals between recovery points. If you anticipate sufficiently rapid changes in our data to make data loss a concern, please consult Sun Professional Services for assistance.

Data that cannot be recovered from a recovery point file can sometimes be retrieved using archiver logs. The procedures are beyond the scope of this document, however, so consult the SAM-FS/QFS documentation (see [“Obtaining the latest information and supporting resources” on page v](#)) or Sun Professional Services.

Caution – SAM-FS/QFS documentation includes procedures for configuring recovery points. Do not change the default IAS recovery point configuration unless you are a highly skilled SAM-FS/QFS administrator. Consult Sun Professional Services instead.

Restoring a file system

When you are recovering an entire SAM-FS/QFS file system, you cannot access the recovery point files (SAM-FS dump files) that are stored on the file system itself, as you would when recovering files and directories. So the appliance's SAM Preventative Maintenance (SAMPM) software automatically saves copies of a few of the most recent recovery point files on an independent UFS file system. By default these files are stored in the `/global/CIS-SAMFSDUMP` subdirectory. To access these disaster recovery point files and restore an entire file system, carry out the tasks listed below:

- [“Indexing recovery point files for filesystem recovery”](#)
- [“Restoring a file system”](#).

▼ Indexing recovery point files for filesystem recovery

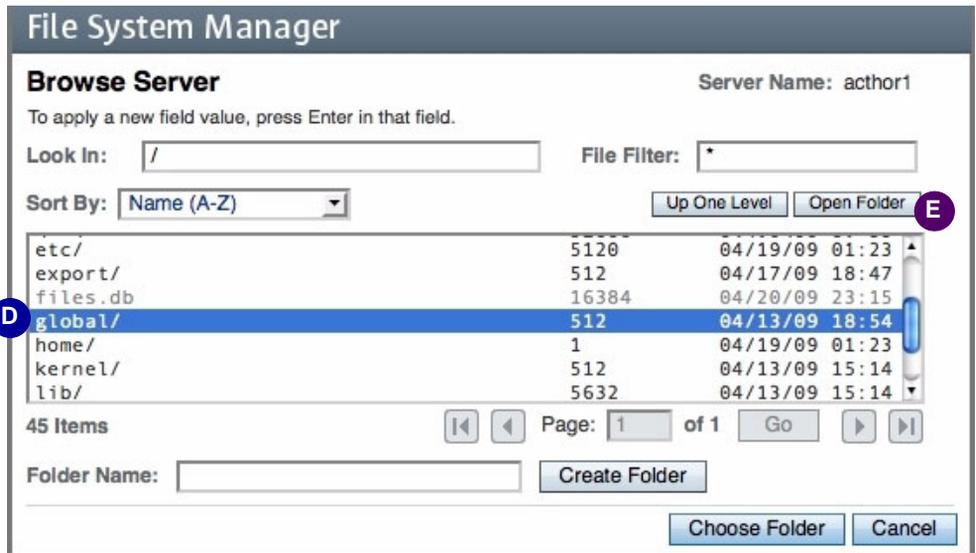
Before you can restore a file or directory, you need to select the correct recovery point file and create an index of its contents. Proceed as follows.

1. **If you are working with a high-availability, multinode appliance, use the Selected Node list control at the top of the IAS task tree to select the MDS node.**
2. **In the IAS task tree, expand the Appliance Expert Maintenance branch.**
3. **In the Appliance Expert Maintenance branch, select Metadata Maintenance (A below).**

The Recovery Points page appears in the content pane.

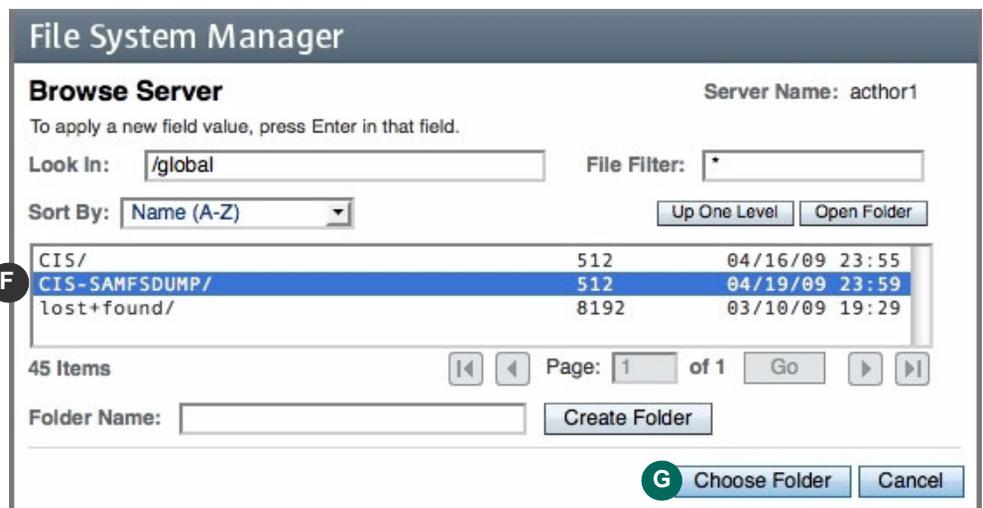
The screenshot shows the 'Recovery Points' page in the IAS task tree. The left sidebar contains the navigation menu, with 'Metadata Maintenance' selected and marked with a blue circle 'A'. The main content area is titled 'Recovery Points' and contains a 'Current File System:' dropdown menu (marked with a red circle 'B') and a 'Recovery Points found in:' browse button (marked with a black circle 'C'). Below these is a summary bar showing 'Choose a file system in the above drop down menu (0)'. A table with columns for File Name, Date, Size, Number of Entries, Indexed, Compressed, and Retain Permanently is shown, with a message below it stating 'No recovery points were found. Select a file system and a directory that contains recovery points'.

4. Using the Current File System list control, select the file system that you intend to restore (B above).
5. Press the Browse button in the Recovery Points found in field (C above).
6. In the directories list of the Browse Server dialog, scroll down, select the global/ directory (D below), and press the Open Folder button (E).



7. When the /global folder opens in the directory list, select CIS-SAMFSDUMP (F below), and press the Choose Folder button (G).

By default, disaster recovery point files are stored in /global/CIS-SAMFSDUMP.



8. In the table on the `Recovery Points` page, click the radio button corresponding to the recovery point file that you want to index (H below).

The SAMPM software names recovery point files `T1_nn-yyyy-Mon-dd`, where `T1_nn` is the name of the file system (`T1_01` by default), `yyyy` is the year (for example `2009`), `Mon` is the abbreviated name of the month (for example `Apr`), and `dd` is the day of the month.

Recovery Points

Current File System: `/T1_01 (T1_01)`

Recovery Points found in: `/global/CIS-SAMFSDUMP`

Recovery Points for File System T1_01 (1 - 25 of 43)

<input type="checkbox"/>	File Name	Date	Size	Number of Entries	Indexed	Compressed	Retain Permanently
<input checked="" type="radio"/> (H)	T1_01-2009-Apr-22	April 23, 2009 12:08:25 AM	1.26 GB		No (K)	Yes	<input type="checkbox"/>
<input type="radio"/>	T1_01-2009-Apr-21	April 22, 2009 12:08:26 AM	1.26 GB		No	Yes	<input type="checkbox"/>
<input type="radio"/>	T1_01-2009-Apr-20	April 21, 2009 12:08:23 AM	1.26 GB		No	Yes	<input type="checkbox"/>

9. Press the `Create Index` button, and wait for indexing to finish.

Be patient. A large recovery point file can take some time to index. When indexing is complete, the value of the `Indexed` field will change from `No` (K) to `Yes`.

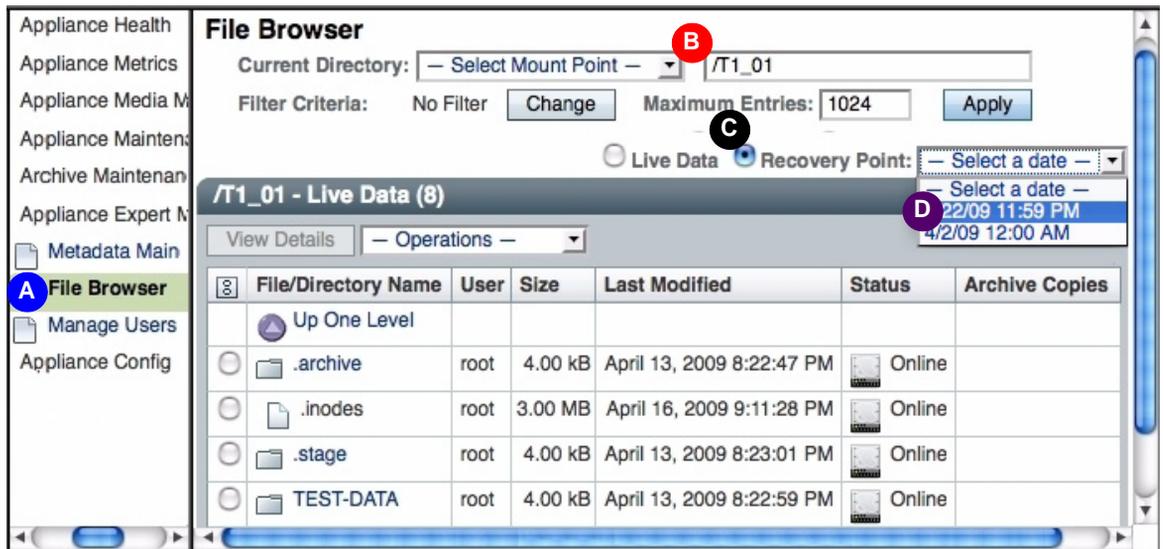
Next task: “Restoring a file system” on page 136.

▼ Restoring a file system

1. If you are working with a high-availability, multinode appliance, use the `Selected Node` list control at the top of the IAS task tree to select the MDS node.
2. In the IAS task tree, expand the `Appliance Expert Maintenance` branch.

3. In the Appliance Expert Maintenance branch, select File Browser (A below).

The File Browser page appears in the content pane.



4. Select the file system mount point using the Current Directory list control (B above).

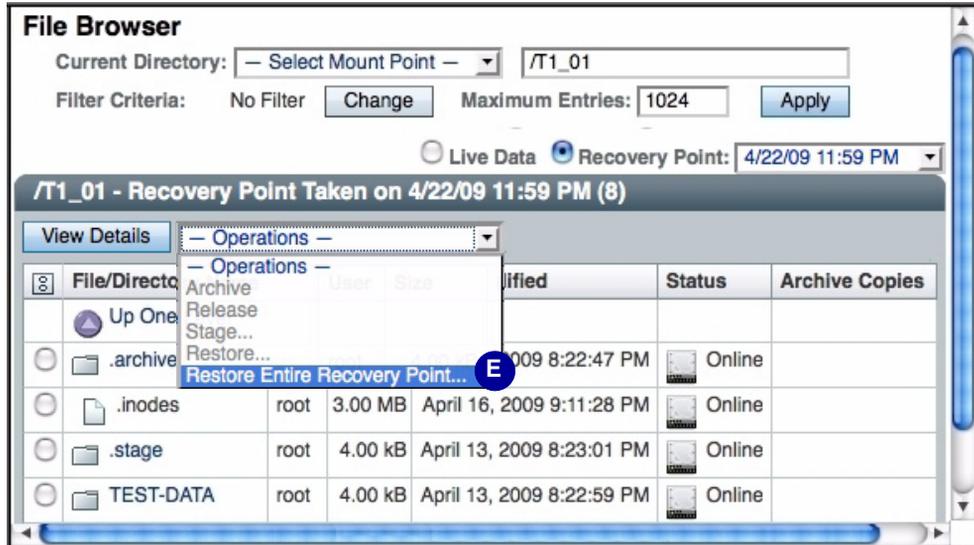
The file system files and directories are displayed in the *filesystem_name* - Live Data table.

5. Click the Recovery Point radio button at the top of the table (C above).
6. Expand the Recovery Point list control, and select the recovery point date for the file that you indexed (D above).

The *filesystem_name* - Live Data table is replaced by the *filesystem_name* - Recovery Point Taken on date time table. You can now browse the recovery point as if you were browsing the live file system.

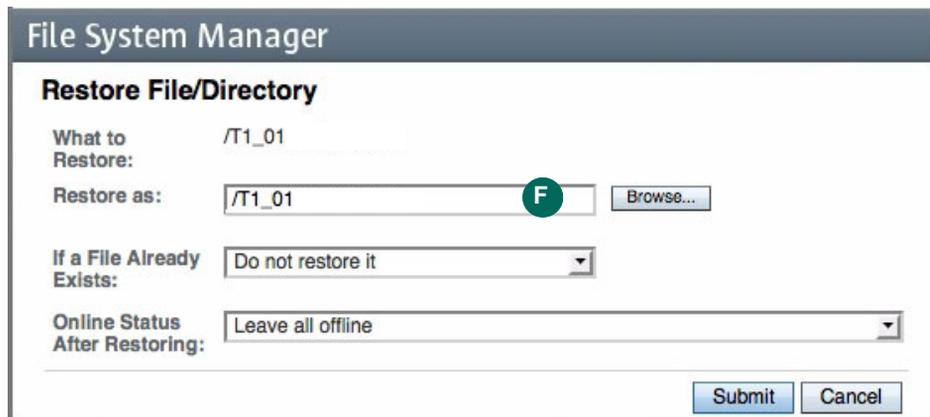
Only indexed recovery point files are listed. If you cannot find the recovery point in the list, stop here and go to ["Indexing recovery point files for filesystem recovery"](#) on page 134.

7. Select **Restore Entire Recovery Point** from the **Operations** list control (E below).

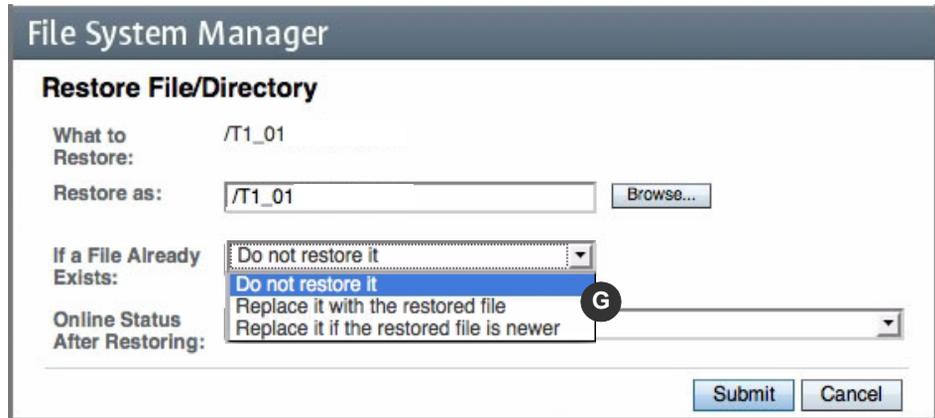


8. When the **Restore File/Directory** popup appears, enter the mount-point to which you want to restore the file system in the **Restore as** text box or press the **Browse** button to search for a mount point (F below).

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.



9. On the same page, use the **If a File Already Exists** list control to specify the action that you want the system to take if the file already exists in the location to which you are restoring files (**G** below).



File System Manager

Restore File/Directory

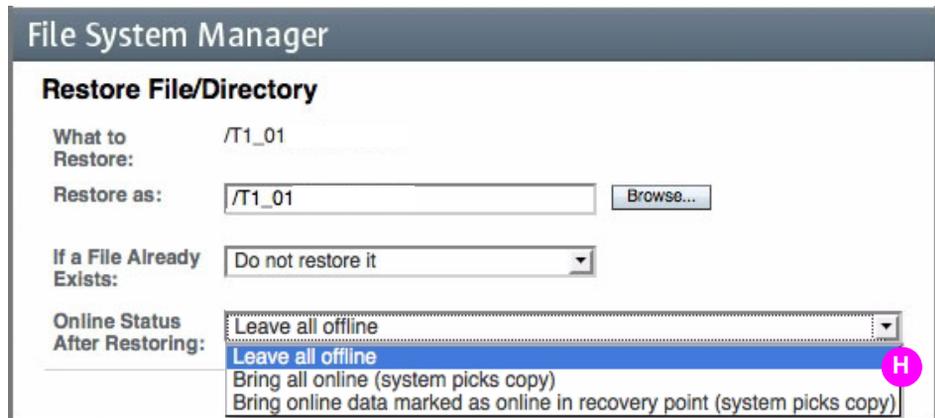
What to Restore: /T1_01

Restore as: /T1_01

If a File Already Exists: **Do not restore it** **Do not restore it**

Online Status After Restoring:

10. From the **Online Status After Restoring** list control, choose the actions that you want the file system to take after completing the restore process (**R** below).



File System Manager

Restore File/Directory

What to Restore: /T1_01

Restore as: /T1_01

If a File Already Exists:

Online Status After Restoring: **Leave all offline** **Leave all offline**

11. Press the **Submit** button to commit your changes, **Cancel** to abort (**J** below).



File System Manager

Restore File/Directory

Online Status After Restoring:

Next task: "Deleting the index for a recovery point" on page 149.

Restoring directories and files

Restoring directories and files is similar to restoring an entire file system, except that you need to perform an additional step: you have to locate the metadata for the desired files and directories within the recovery point file, so that the IAS system can restore exactly the items you need. To locate individual file and directory metadata, you must *index* the chosen recovery point file. A recovery point index lets you go through the directories and files listed in a recovery point file and select items using the IAS File Browser graphical user interface, just as you would when browsing files in the live file system.

To facilitate indexing, the IAS appliance's SAM Preventative Maintenance (SAMPM) software automatically copies daily recovery point files from the directory where recovery point files are created, `/global/CIS-SAMFSDUMP`, to the `sampm/` subdirectory of each IAS Tier-1 file system. The `/global/CIS-SAMFSDUMP` directory has limited space—just enough for a handful of the most recently generated recovery points—and indices take up substantial amounts of space. Copying recovery points to the `sampm/` subdirectory thus gives you the space you need, plus room to store a larger number of recovery points.

This section describes the processes involved in recovering files and file systems:

- [“Indexing recovery point files for directory and file recovery” on page 140](#)
- [“Restoring directories and files” on page 143.](#)

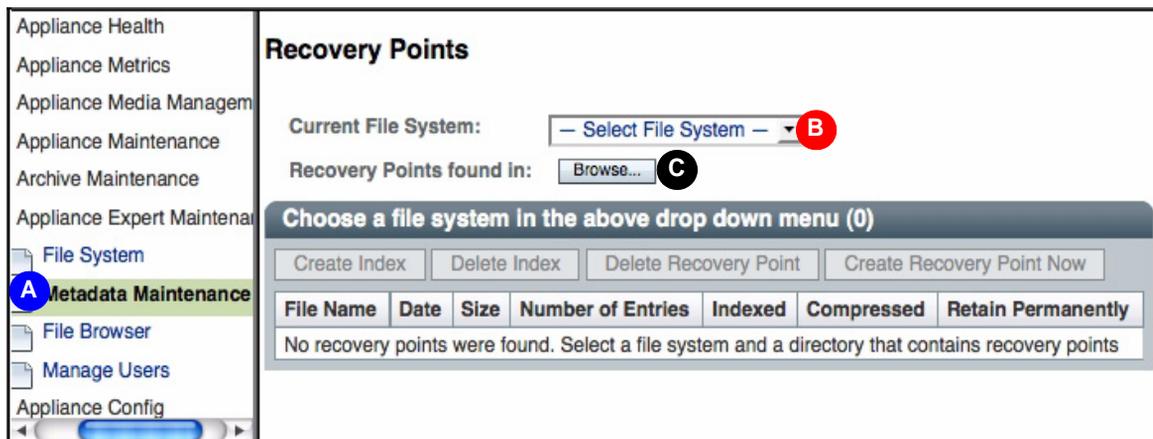
▼ Indexing recovery point files for directory and file recovery

Before you can restore a file or directory, you need to select the correct recovery point file and create an index of its contents. To locate and index a recovery point file, proceed as follows.

1. **If you are working with a high-availability, multinode appliance, use the Selected Node list control at the top of the IAS task tree to select the MDS node.**
2. **In the IAS task tree, expand the Appliance Expert Maintenance branch.**

- In the **Appliance Expert Maintenance** branch, select **Metadata Maintenance** (A below).

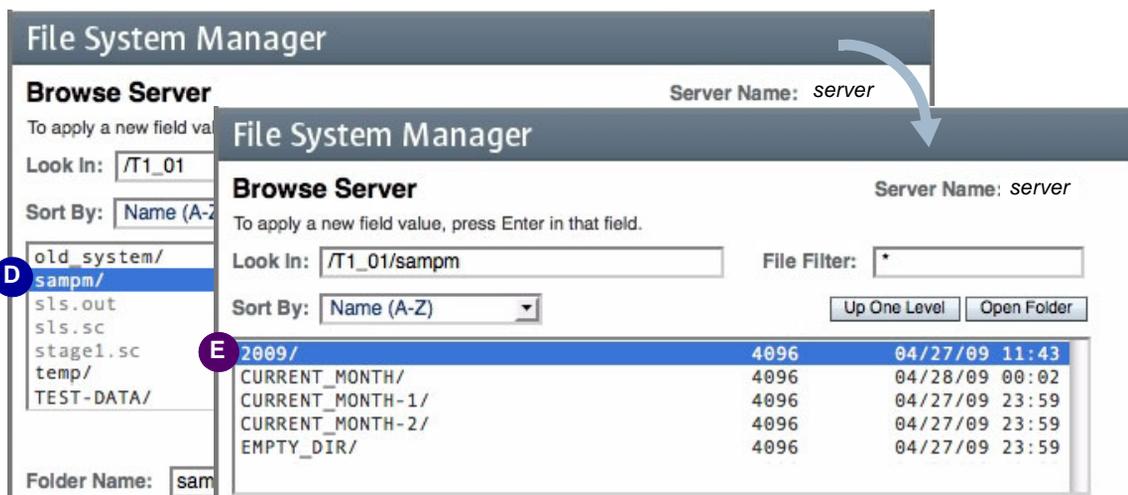
The **Recovery Points** page appears in the content pane.



- Using the **Current File System** list control, select the **Tier-1** file system that you intend to restore (B above), and press the **Browse** button (C) to search for recovery point files.

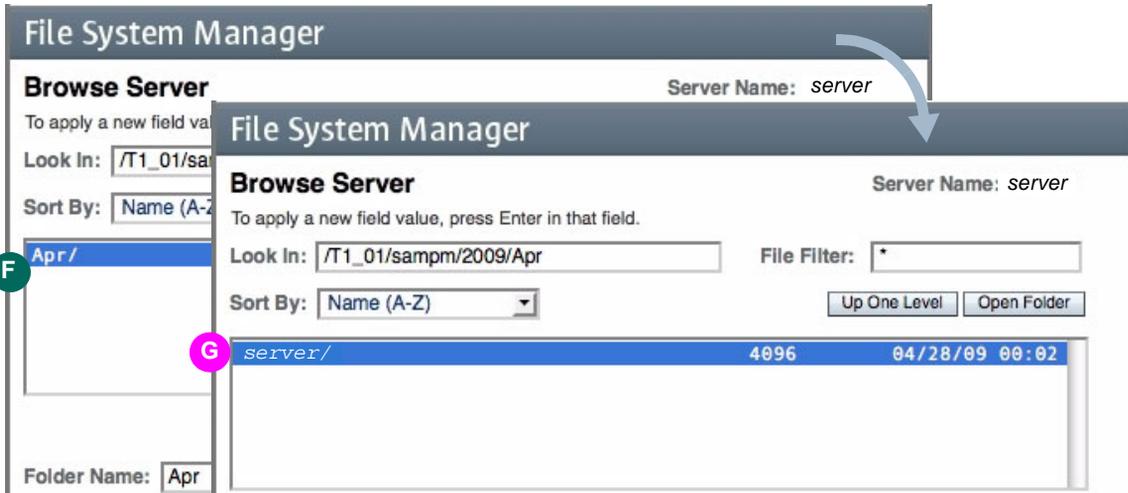
The **File System Manager** server browser opens in the mount-point directory of the selected file system (`/T1_01` in the example below).

- In the **FSM** server browser, scroll down and double click on the `sampm/` directory (D below).



- When the directory browser refreshes, double click on the entry for the year of the recovery point that you are looking for (E above).

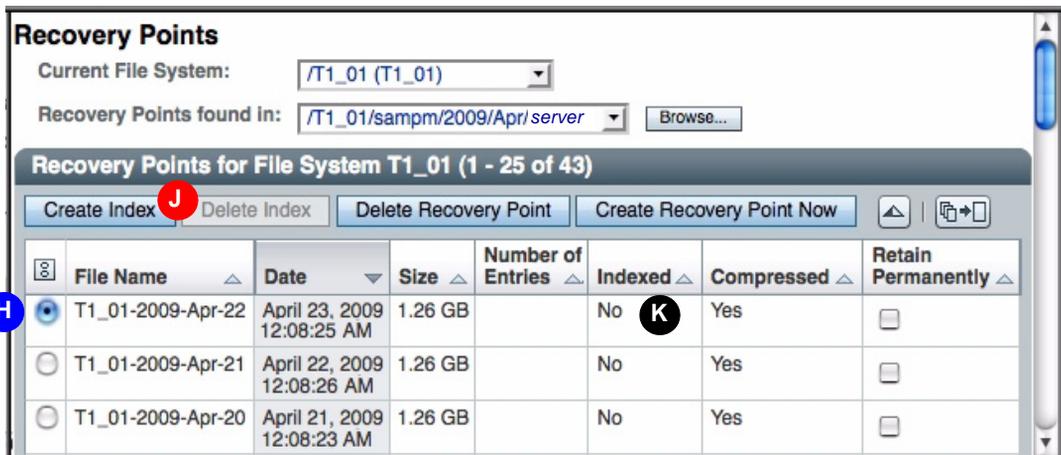
- When the directory browser refreshes, double click on the entry for the month of the recovery point that you are looking for (F below).



- When the directory browser refreshes, double click on the entry for the server of the recovery point that you are looking for (G above).

- Click the radio button beside the recovery point file that you want to index (H below).

The SAMPM software names recovery point files *T1_nn-yyyy-Mon-dd*, where *T1_nn* is the name of the file system (T1_01 by default), *yyyy* is the year (for example 2009), *Mon* is the abbreviated name of the month (for example Apr), and *dd* is the day of the month.



10. Press the `Create Index` button (J above), and wait for indexing to finish.

Be patient. A large recovery point file can take some time to index. When indexing is complete, the value of the `Indexed` field will change from `No` (K) to `Yes`.

Next task: “Restoring directories and files” on page 143.

▼ Restoring directories and files

To restore a file system, proceed as follows:

1. If you are working with a high-availability, multinode appliance, use the `Selected Node` list control at the top of the IAS task tree to select the MDS node.
2. In the IAS task tree, expand the `Appliance Expert Maintenance` branch.
3. In the `Appliance Expert Maintenance` branch, select `File Browser` (A below).

The `File Browser` page appears in the content pane.

The screenshot shows the File Browser interface. On the left is a navigation pane with 'File Browser' selected (A). The main area has a 'Current Directory' dropdown set to '/T1_01' (B). Below it are 'Filter Criteria' (No Filter), 'Change', 'Maximum Entries' (1024), and 'Apply' buttons. There are radio buttons for 'Live Data' (C) and 'Recovery Point' (D). A date selection dropdown is also visible. Below these controls is a table titled '/T1_01 - Live Data (8)'. The table has columns for File/Directory Name, User, Size, Last Modified, Status, and Archive Copies. The table contains entries for '.archive', '.inodes', '.stage', and 'TEST-DATA', all with status 'Online'. There is also an 'Up One Level' button.

File/Directory Name	User	Size	Last Modified	Status	Archive Copies
Up One Level					
.archive	root	4.00 kB	April 13, 2009 8:22:47 PM	Online	
.inodes	root	3.00 MB	April 16, 2009 9:11:28 PM	Online	
.stage	root	4.00 kB	April 13, 2009 8:23:01 PM	Online	
TEST-DATA	root	4.00 kB	April 13, 2009 8:22:59 PM	Online	

4. Select the file system mount point using the `Current Directory` list control (B above).

The file system files and directories are displayed in the `filesystem_name - Live Data` table.

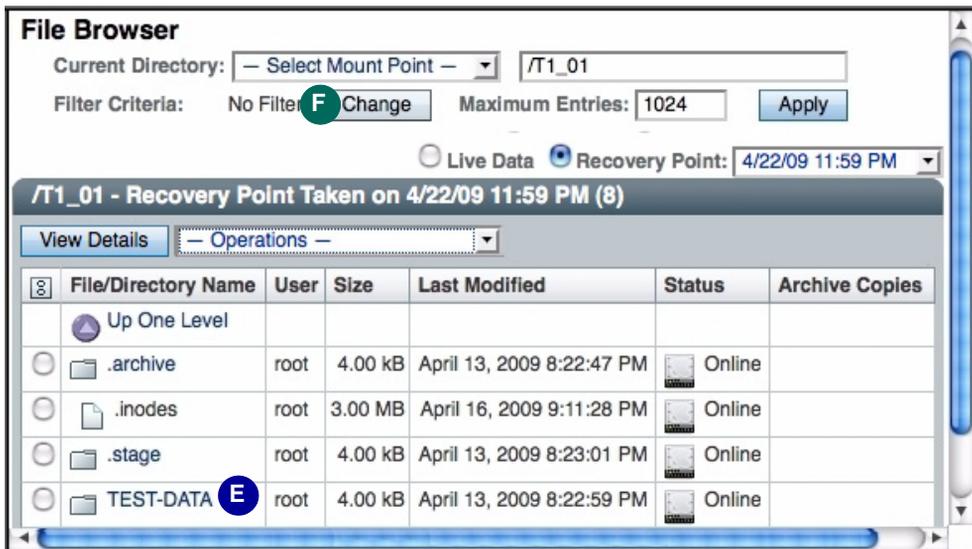
5. Click the `Recovery Point` radio button at the top of the table (C above).

- Expand the Recovery Point list control, and select the recovery point date to which you want to restore (D above).

The *filesystem_name* - Live Data table is replaced by the *filesystem_name* - Recovery Point Taken on *date time* table. You can now browse the recovery point as if you were browsing the live file system.

If you cannot find the recovery point in the list, stop here and go to [“Indexing recovery point files for directory and file recovery” on page 140.](#)

- Browse through the recovery point file to locate the directories and/or files that you want to restore. To step down through the directories under the current mount point, click on the hyperlinks in the File/Directory Name field of the *filesystem_name* - Recovery Point Taken on *date time* table (E below).



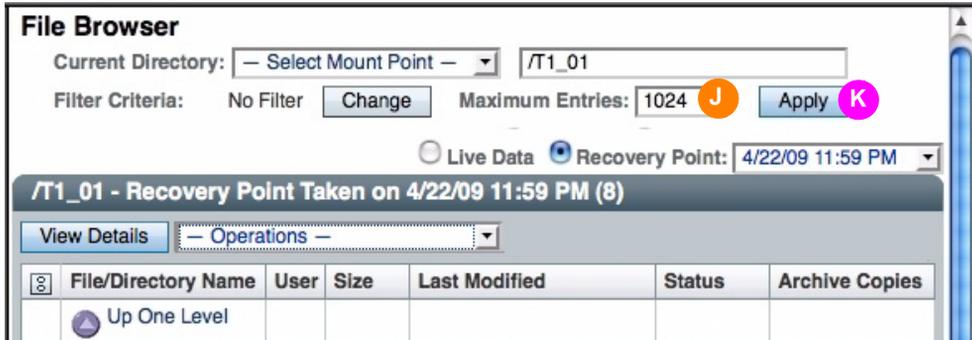
- To reduce the visual clutter and simplify a search in a large file system, set a Filter Criteria by pressing the associated Change button at the top of the page (F above).

9. If you chose to set a filter, enter your criteria in the **Filter Files** popup using the fields and controls provided (G below). Press **OK** to define the filter, **Cancel** to discard changes (H).

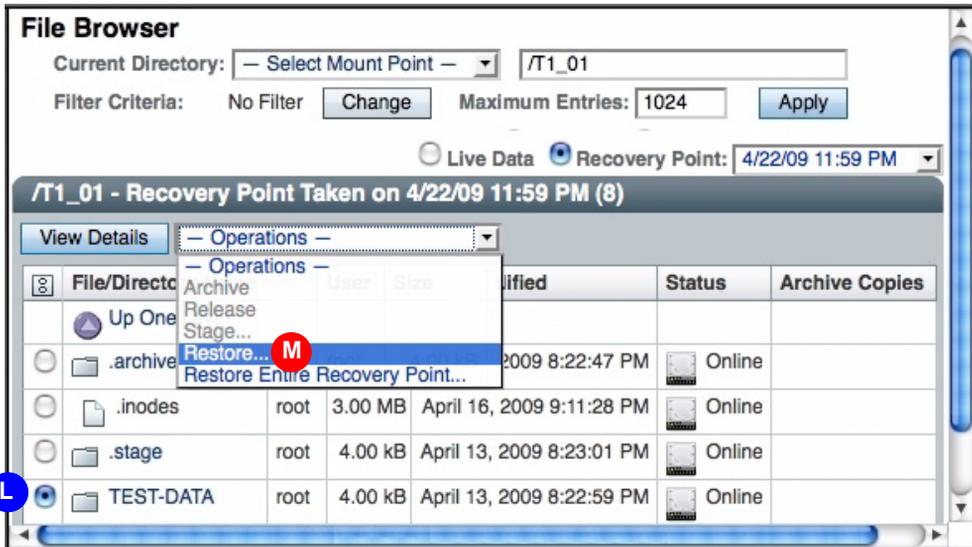
The **File Name Pattern** field uses the pattern-matching syntax outlined in the following table:

Description	Sample pattern	Sample match
/	matches a directory path only and cannot be matched by any wildcard character	/var /abc /abc/* /abc/abc, /abc/def, etc.
?	matches any single character	a?d aad, abd, acd, etc.
*	matches zero or more characters	a*d ad, abcd, acbd, aaaad, etc.
[]	matches files any one of the characters enclosed by the brackets	a[bc] ab or ac
" "	matches only the literal string within the quotation marks	"abc" only abc
\	escapes the wildcard character and makes it into a literal value	a\?b a?b (NOT acb) a\\b a\b a\\? a\b, a\c, or a\?

- To limit the number of directory and file entries that are read from the recovery point file and displayed on screen, enter a number in the `Maximum Entries` field (J below).



- If you defined a filter or set a maximum number of entries displayed, press the `Apply` button (K above) to make your limits take effect.
- Once you have located the file or directory that you want to restore in the `File/Directory Name` field of the `filesystem_name - Recovery Point Taken on date time` table, click the corresponding radio button (L below).



- Select `Restore` from the `Operations` list control (M above).

14. When the **Restore File/Directory** popup appears, enter the location to which you want to restore the file or directory in the **Restore as** text box or press the **Browse** button to search for a destination (**N** below).

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.

File System Manager

Restore File/Directory

What to Restore: /T1_01/TEST-DATA

Restore as: /T1_01/TEST-DATA **N** Browse...

If a File Already Exists: Do not restore it

Online Status After Restoring: Leave all offline

Submit Cancel

15. On the same page, use the **If a File Already Exists** list control to specify the action that you want the system to take if the file already exists in the location to which you are restoring files (**P** below).

File System Manager

Restore File/Directory

What to Restore: /T1_01/TEST-DATA

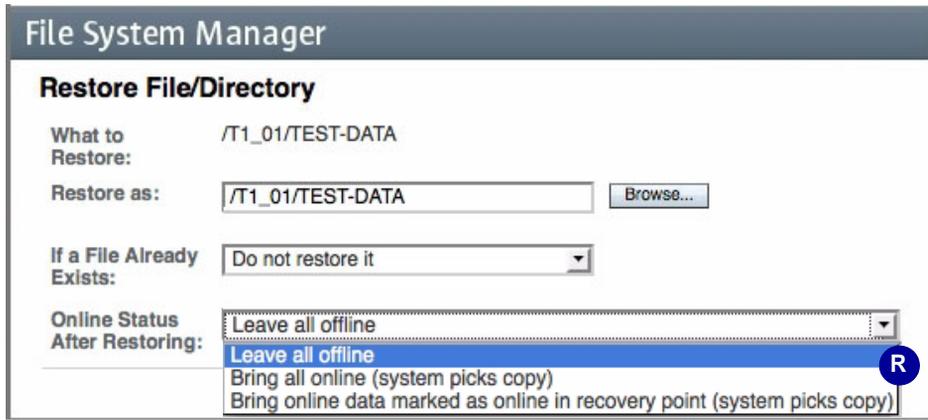
Restore as: /T1_01/TEST-DATA Browse...

If a File Already Exists: Do not restore it

Online Status After Restoring: Replace it with the restored file **P**

Submit Cancel

16. From the **Online Status After Restoring** list control, choose the actions that you want the file system to take after completing the restore process (R below).



File System Manager

Restore File/Directory

What to Restore: /T1_01/TEST-DATA

Restore as: /T1_01/TEST-DATA

If a File Already Exists: Do not restore it

Online Status After Restoring:
Leave all offline
Bring all online (system picks copy) **R**
Bring online data marked as online in recovery point (system picks copy)

17. Press the **Submit** button to commit your changes, **Cancel** to abort (S below).



Online Status After Restoring: Leave all offline

Next task: [“Deleting indices as part of routine file recovery” on page 149.](#)

Conserving disk space

The indices to recovery point files can consume a great deal of disk space over time. For this reason, you should always delete indices that are created during restoration of entire file systems and may want to periodically delete indices that are used for directory and file recovery as well. This section begins by explaining the reasons for [“Deleting indices as part of disaster recovery”](#) and for either retaining or [“Deleting indices as part of routine file recovery” on page 149.](#) It concludes with the detailed procedure for [“Deleting the index for a recovery point” on page 149.](#)

Deleting indices as part of disaster recovery

In disaster recovery situations, where entire file systems are recovered, the need for an index vanishes once disaster recovery is complete. Since space in the default `/global/CIS-SAMFSDUMP` directory is particularly limited, you need to delete the index and any associated files as soon as you have restored the file system. See [“Deleting the index for a recovery point”](#) for instructions.

Deleting indices as part of routine file recovery

Deleting indices that are used for day-to-day directory and file recovery is optional. Since the recovery point files and indices reside, by default, in the subdirectories of the `sampm/` subdirectory of each file system, space is not an immediate problem. Having at least some indices available can be convenient if you frequently restore directories and files in response to user requests. When you have a ready-made index, you can start restores immediately, without waiting for an index to generate for each request.

Nonetheless, space is not unlimited in IAS file systems. Index files compete for storage space with recovery point files and data. So you should think carefully about keeping indices. Try to keep just enough on hand to address the bulk of user requests within an acceptable response time. Delete indices that outlive their immediate usefulness. If you delete an index that you later need, you can always regenerate it by re-indexing the recovery point file. For instructions, see [“Deleting the index for a recovery point”](#).

Deleting recovery point indices

To delete an index and its associated file(s), proceed as follows.

▼ Deleting the index for a recovery point

1. **If you are working with a high-availability, multinode appliance, use the Selected Node list control at the top of the IAS task tree to select the MDS node.**
2. **In the IAS task tree, expand the Appliance Expert Maintenance branch.**
3. **In the Appliance Expert Maintenance branch, select Metadata Maintenance.**
The Recovery Points page appears.
4. **Using the list control provided, select the Current File System.**

5. **Using the list control provided, select the Recovery Points found in directory, or press the Browse button to search for other recovery point directories.**
6. **Select the radio button corresponding to the recovery point file for which you want to delete an index.**
7. **Press the Delete Index button and then confirm your choice when prompted.**
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.

Stop here.

Archiving data

Archiving has two primary functions.

- It implements an information lifecycle management policy, so that the characteristics of storage media are closely matched with the characteristics and usage of the data
- It creates redundant copies of data on independent media, so that a damaged or lost file system can be restored using the filesystem metadata in recovery point file.

IAS appliances come preconfigured with archive policies that archive data continuously. The archiver determines how and when to schedule archiving by balancing the elapsed time since files were last archived against the amount of data that needs archiving. When the elapsed time exceeds a `Start Age` threshold or the amount of data pending exceeds a `Start Size` threshold, the archiver makes the required number of copies.

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

- [“Archive management and IAS” on page 152](#)
- [“Storage Archive Manager \(SAM\)” on page 153](#)
- [“Managing archiving policies” on page 155](#)
- [“Controlling the archiver” on page 169.](#)

Archive management and IAS

Sun Customer Ready Infinite Archive System (IAS) appliances come with preconfigured files systems and archiving policies that automatically administer the files. Accordingly, you should not have to do any significant archive management under normal circumstances. Indeed, needless operator intervention can significantly reduce archive performance, capacity, and/or reliability.

Nonetheless, operational requirements do sometimes justify changes to the archive configuration. When this is the case, this chapter can guide you through the necessary tasks. But before it can do so, you need to be aware of both the general management principles that guide IAS file system design and the considerable potential for unanticipated side effects that accompanies any modifications to a working archive. There are three fundamental rules for working with IAS file systems:

- [“Avoid complexity” on page 152](#)
- [“Know what you are doing” on page 153](#)
- [“Know your documentation” on page 153.](#)

Avoid complexity

The first principle of IAS filesystem management is to keep things simple. The SAM-QFS file systems manage data migration and media utilization automatically, under the guidance of the archiving policies that are part of the core IAS system configuration. If you create, modify, or redefine any of these policies, do so with care and restraint. Do not over-specify, and do not try to micromanage every conceivable situation. Doing so reduces the efficiency of the archive, wastes storage space, makes the configuration difficult to troubleshoot, and adds overhead that may reduce performance and compromise reliability.

In particular, satisfying the multiple conditions in a complex archiving policy generally requires more loading and unloading of tape media than a simpler policy would require. This is a significant problem in a SAM-QFS file system that integrates tape and disk media. Loading, mounting, and positioning a tape is a multistep mechanical operation that is relatively time-consuming compared to disk I/O. Each additional mount adds overhead that subtracts from overall filesystem performance. In addition, each additional mount exposes both the tape media and the moving parts in drives and robots to added wear and tear. The additional loads shorten the life of media and components, increase the risk of mechanical failure, and increase maintenance requirements. Keeping policies simple minimizes tape movement and avoids these issues.

Know what you are doing

The second principle of IAS filesystem management is to always know what you are doing before you proceed. The SAM-QFS file systems and File System Manager (FSM) software that are integrated into the IAS appliance are powerful, multifunctional tools that support a wide range of configurations and user requirements. Such powerful, highly configurable tools are highly unforgiving—too much so for unskilled users. IAS comes preconfigured to reliably meet most requirements. If you choose to alter this configuration, you must have the skills to do so successfully.

No one should undertake SAM-QFS filesystem maintenance or reconfiguration without adequate prior training and considerable expertise. If you do not have the necessary training and experience, defer to someone who does. If the requisite expertise is simply not available in your organization, please do not hesitate: contact Sun for assistance.

Know your documentation

Finally, no one should undertake SAM-QFS filesystem maintenance or reconfiguration without access to and familiarity with the comprehensive SAM-QFS and FSM document sets. While this document outlines many archiving procedures, it is not, on its own, an adequate substitute. SAM-QFS file systems and FSM software are highly flexible tools that present the user with many options and consequently many potential pitfalls—far more than can be fully documented in a guide to a preconfigured appliance. So consult the full range of SAM-QFS documentation that is available online from Sun. See [“Obtaining the latest information and supporting resources” on page v](#) of this guide for details.

Storage Archive Manager (SAM)

The Sun StorageTek Storage Archive Manager (SAM) component of the Infinite Archive System archives QFS filesystem data and manages multiple copies of files on the IAS hierarchical storage system. SAM functionality is distributed across four logical components: the *archiver*, the *releaser*, the *stager*, and the *recycler*.

The archiver

The SAM archiver makes up to four copies of each modified file that it finds in tier-1 disk storage and stores them on different volumes in tier-2 archival disk storage and/or tier-3 archival tape storage. The archiver follows the instructions contained in *policy files*. Each policy is a set of directives that specify the following system characteristics:

- the file systems that are archived
- the *archive age* (interval since last modification) that makes a file eligible for archiving
- the numbers of copies that are made
- the properties of each copy,
- the media on which copies reside
- the kinds of files that should be kept together.

System-defined default policies, such as `allfs1` and `allsets`, determine the base behavior of the archiver. Additional, custom policies can then fine-tune behavior by overriding individual parameters in the default policies. A special `no-archive` policy contains no copy directives, only file-match criteria, so that the archiver can skip specified directories (for instructions on using the `no-archive` policy, see [“Defining additional file-match criteria” on page 163](#)).

The releaser

The releaser maintains enough free space in tier-1 disk storage so that there is always enough room for the files that benefit most from fast, responsive storage: files that are being actively written and read. When tier-1 utilization exceeds a threshold value called the *high water mark*, the releaser begins releasing the disk space occupied by files that have been archived to tier-2 or tier-3 media and have not been accessed recently. The releaser then continues to release files until tier-1 utilization drops below a second threshold value, the *low water mark*.

The stager

The stager restores files to tier-1 disk storage when a user or process requests file data that has been released. The stager copies the file from tier-2 or tier-3 storage back to tier-1. 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. If an error occurs, the system attempts to find the next

available archive copy of the file. Media errors, unavailability of media, and unavailability of a library can often be handled automatically, without interrupting the process.

The recycler

The recycler maintains enough space in tier-2 and tier-3 archival storage to insure that there is always room for newly made archival copies. The recycler scans for expired file copies that have exceeded their retention periods and deletes them.

Managing archiving policies

Archiving policies supply the core logic for Infinite Archive System (IAS) solutions. They determine what data is archived and when. They specify the number of copies that are kept in the archive and the kind of storage media that are used. They control how the files migrate within the archives and how users access them.

You build archiving policies from two basic types of component: *file-matching criteria* and *copy parameters*. File-matching criteria identify the files that are governed by the policy. The criteria can be as broad as a simple path from the filesystem mount point to a starting directory or as specific as a list of individual file names. Copy parameters specify the number of copies, the media used, sorting (if any), file aging, device and queuing behavior, and recycling criteria. Copy parameters can also be broad, high-level directives or extremely fine-grained sets of commands.

Most customers should not alter the preconfigured, IAS factory default policies. They have been selected to meet the needs of the vast majority of users. However, the IAS Java Web Console interface provides all of the tools needed should special circumstances require. If you need to alter the default policy set, review the introductory material in [“Archive management and IAS” on page 152](#), exercise care, and consider all the possible ramifications before proceeding.

This section starts with some general rules for modifying the IAS default configuration, [“Writing practical archiving policies” on page 156](#). It then proceeds with specific instructions for [“Adding a policy to a file system” on page 156](#), [“Deleting a policy” on page 162](#), [“Editing a policy” on page 162](#).

Writing practical archiving policies

If you choose to alter the default IAS archiving configuration, you should keep some basic guidelines in mind.

Choose a sensible archive interval. The archive interval is the time between file system scans. IAS appliances come pre-configured for continuous archiving, with scanning disabled (the `No Scan` method). If you choose to override this arrangement, base archive intervals on how often files are created and modified and on whether you want to save all modification copies. A very short archive interval keeps the archiver scanning almost continuously.

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.

Finally, keep complexity to a minimum. Complexity increases your maintenance burden and may needlessly reduce archive performance and capacity. So do not over-specify when you set criteria and parameters:

- specify the highest level file-matching criteria that you can: do not provide a file list or a matching pattern if a simple start directory will do
- set the least restrictive media-usage parameters that you can: use pools rather than ranges and ranges rather than particular volume serial numbers whenever possible.

Adding a policy to a file system

When you add a policy, you change the archiving behavior of a file system by defining new file-match criteria and copy directives. In most cases, IAS comes preconfigured with tested, optimized values for most of the tunable parameters that you can adjust. Since many parameters are interdependent, ill-considered changes can have unexpected, deleterious effects. So, if you need to make changes, be careful, and know what you are doing.

[“Starting the New Archive Policy wizard” on page 157](#)

[“Defining the files to which the policy will apply” on page 157](#)

[“Naming the policy and defining basic archiving behavior” on page 158](#)

[“Defining copy parameters” on page 159](#)

▼ Starting the New Archive Policy wizard

1. **Login to the IAS Java Web Console under the expert or admin role.**
2. **If you are working with a high-availability, multinode appliance, use the Selected Node list control at the top of the IAS task tree to select the MDS node.**
3. **In the IAS task tree, expand the Appliance Maintenance branch.**
4. **On the Appliance Maintenance branch, expand Archive Maintenance.**
5. **Under Archive Maintenance, select Archiver Policies.**
The Archive Policies Summary page appears in the content pane.
6. **Press the New button.**
The New Archive Policy wizard starts.

Next task: [“Defining the files to which the policy will apply” on page 157.](#)

▼ Defining the files to which the policy will apply

1. **In the first step of the New Archive Policy wizard, Define File Match Criteria, enter the relative path from the target filesystem mount point to the Starting Directory.**
2. **If desired, enter a regular expression that matches the names of the files that you want to archive in the Name Pattern text field and use the fields and controls provided to specify an Owner, a Minimum Size, a Group, a Maximum Size, and/or an Access Age.**
Be judicious when using these optional parameters, because they complicate the archiving task and increase the likelihood that archiving will produce unanticipated results. Remember: simple policies are usually best. Simple policies produce more predictable results, run faster, and are easier to debug if problems arise.
3. **Select the target file system from the Apply to File System list control.**
4. **Press the Next button to continue, Previous to return to the last step, or Cancel to abort.**

Next task: [“Naming the policy and defining basic archiving behavior” on page 158.](#)

▼ Naming the policy and defining basic archiving behavior

1. **In the next step**, Define Policy Name and Number of Copies, **enter a New Archive Policy Name in the text field provided.**
2. **Specify the Number of Archive Copies using the list control provided.**
Specify at least two.
3. **Select a Staging Behavior using the list control provided.**

You have the following choices:

- Stage these files automatically when they are accessed provides the best responsiveness under most conditions and is the best choice for most purposes.

Files are *staged*—copied back from tape to the disk cache—when they are called for and before they are read. This approach lets the user browse files on random-access disk media, which, for most files, is more responsive than searching sequentially through tape. The sequential-access tape drive operates more efficiently as well, because it can write the whole file out to disk in one go instead of constantly repositioning media in response to user requests.

- Mark these files for associative staging automatically stages all files within the target directory when any one of the associated files is accessed.

This method is useful when the contents of a directory are likely to be used together, so that a call for one file is likely to be followed by calls for the others. Examples might include multifile software-development or CAD projects and collections of medical images associated with a single patient's medical records.

- Never stage is useful when you are storing very large files, such as medical imagery or surveillance video, that can be *streamed* from tape to primary cache.

When files are large enough to be streamed from tape at the optimum transfer speed for the tape drive, sequential access is as fast as or faster than random access from disk. So large files can be read from tape to primary disk cache while the user is actually accessing the file. This just-in-time approach to transferring data saves the disk cache for the smaller files that are best accessed from a random I/O device. The approach also greatly simplifies cache management.

- Reset the stager behavior of these files to the default.

4. **Select a Releaser Behavior using the list control provided.**

You have the following choices:

- Stage these files automatically when they are accessed provides the best responsiveness under most conditions and is the best choice for most purposes.

Files are *staged*—copied back from tape to the disk cache—when they are called for and before they are read. This approach lets the user browse files on random-access disk media, which, for most files, is more responsive than searching sequentially through tape. The sequential-access tape drive operates more efficiently as well, because it can write the whole file out to disk in one go instead of constantly repositioning media in response to user requests.

- Mark these files for associative staging automatically stages all files within the target directory when any one of the associated files is accessed.

This method is useful when the contents of a directory are likely to be used together, so that a call for one file is likely to be followed by calls for the others. Examples might include multifile software-development or CAD projects and collections of medical images associated with a single patient's medical records. If used with files that are not, in fact, used together, this method makes poor use of the disk cache and may cause excessive mounting and positioning of tape media.

Note that you cannot use associative archiving if the files do not fit on a single tape volume. In this case, use either the `-sort path` or `-rsort path` parameters to keep the files contiguous.

- Never stage is useful when you are storing very large files, such as medical imagery or surveillance video, that can be *streamed* from tape to primary cache.

When files are large enough to be streamed from tape at the optimum transfer speed for the tape drive, sequential access is as fast as or faster than random access from disk. So large files can be read from tape to primary disk cache while the user is actually accessing the file. This just-in-time approach to transferring data saves the disk cache for the smaller files that are best accessed from a random I/O device. The approach also greatly simplifies cache management.

- Reset the stager behavior of these files to the default.

Next task: [“Defining copy parameters” on page 159.](#)

▼ Defining copy parameters

1. **In the next step, Copy 1 - Media Parameters, specify an Archive Age by entering a number in the text field provided and then selecting units using the accompanying list control.**

The `archive age` is the amount of time since a file's last modification.

2. **Specify a Media Type and, if desired, a VSN Pool Name (media pool name) using the list controls provided.**

3. **If desired, Specify VSNs (tape volumes) by entering a range of volume serial numbers (VSNs) in the two accompanying text fields and/or by entering a regular expression matching the desired VSNs or a comma-delimited list of VSNs.**

Be careful about over specifying. Use pools rather than ranges and ranges rather than particular volume serial numbers whenever possible. If you restrict the archiver's choice of media, you may reduce media utilization and archiving performance, while increasing manipulation of tape media. An archiving operation might block waiting for media.

4. **If desired, Reserve the VSNs used for copies from the same Policy and/or File System by checking the corresponding check boxes. You can also use the accompanying list control to further restrict use of the VSN to files from the same Directory, User, or Group.**

Again, be careful about over specifying. If you restrict the archivers choice of media, you may reduce media utilization and archiving performance and may increase tape manipulation by increasing the need for mounting and positioning.

5. **Optionally, in the next step, Copy 1 - Advanced Disk Copy Option, use the list control to specify archiver behavior** If the files to be archived are offline.

This option lets you specify how subsequent copies are made once the first copy is complete. Your choices are:

- Stage files back as they are needed copies files from the offline tape media to the disk cache whenever the archiver is ready to copy them to tape media.
- Copy them directly from offline media bypasses the disk cache and copies files directly from tape to tape.
- Stage one file ahead of the current copy pre-stages the next file from offline tape to disk while the current file is being copied to tape.
- Stage them all back first pre-stages all files from offline tape to disk before copying them to tape.

Next task: [“Setting work-queue and recycler parameters” on page 160.](#)

▼ Setting work-queue and recycler parameters

1. **If desired, you can adjust the Archiver's Work Queue. Adjust the archiving Start Age, Start Count, and Start Size using the controls provided.**

Be judicious about using this option. IAS comes preconfigured with Start Age, Start Count, and Start Size values that are optimal for most users. Changing these values complicates the solution, so be sure that you know what you are doing.

2. **If desired, you can Disable Recycling for this policy by checking the corresponding checkbox.**

Be careful with this option. If recycling is disabled, you cannot reclaim tape media for reuse after copies expire.

3. **If desired, you can set a Recycler High Water Mark by entering a percentage in the field provided.**

A recycler high water mark sets the percentage of media usage below which recycling cannot occur. This percentage is the ratio of the used space in the library to its total capacity.

4. **If desired, you can enter an email address in the Mail Recycler Information to field.**

5. **If desired, enter a percentage in the Only recycle VSNs that would gain at least (%) field.**

Recycling a tape is a comparatively high-overhead process that may require repeated mounting and repositioning of tapes. Entering a value in this field lets you balance the costs of this overhead against the benefit of additional library capacity. If a recycling a tape will reclaim too little space to make recycling worthwhile, the recycler skips it.

6. **Press the Next button to continue, Previous to return to the last step, or Cancel to abort.**

Next task: [“Finishing the new archive policy” on page 161.](#)

▼ Finishing the new archive policy

1. **In the next step, the Summary, double-check your entries.**
2. **If everything is correct, press the Finish button to commit the changes.**
3. **Otherwise, press Previous to return to the last step, or Cancel to abort.**

The New Archive Policy wizard displays the results of your submission.

Stop here.

Deleting a policy

Deleting a policy removes the policy from the `archiver.cmd` file that controls the archiver. To delete a policy, proceed as follows.

1. **Login to the IAS Java Web Console under the `expert` or `admin` role.**
2. **If you are working with a high-availability, multinode appliance, use the Selected Node list control at the top of the IAS task tree to select the MDS node.**
3. **In the IAS task tree, expand the `Appliance Maintenance` branch.**
4. **On the `Appliance Maintenance` branch, expand `Archive Maintenance`.**
5. **Under `Archive Maintenance`, select `Archiver Policies`.**
The `Archive Policies Summary` page appears in the content pane.
6. **In the `Archive Policies` table, select the radio button that corresponds to the policy that you want to delete, and press the `Delete` button.**
7. **If you are sure, confirm the deletion when prompted.**

Stop here.

Editing a policy

You edit policies from the `Archive Policies` leaf of the `Archive Maintenance` and `Appliance Maintenance` branches of the IAS task tree. Start with [“Accessing policies for editing”](#) on page 162, below:

▼ Accessing policies for editing

1. **Login to the IAS Java Web Console under the `expert` or `admin` role.**
2. **If you are working with a high-availability, multinode appliance, use the Selected Node list control at the top of the IAS task tree to select the MDS node.**
3. **In the IAS task tree, expand the `Appliance Maintenance` branch.**
4. **On the `Appliance Maintenance` branch, expand `Archive Maintenance`.**
5. **Under `Archive Maintenance`, select `Archiver Policies`.**
The `Archive Policies Summary` page appears in the content pane.

Next task: from the `policy-name` `Details` page, you can do any of the following:

- [“Defining additional file-match criteria” on page 163](#)
- [“Defining additional archive copies” on page 165](#)
- [“Modifying file-match criteria” on page 167](#)
- [“Modifying copy parameters” on page 168.](#)

▼ Defining additional file-match criteria

Adding file-match criteria to an existing policy applies the actions defined by the policy to additional files. Adding file-match criteria to the `no-archive` policy, for instance, excludes additional files and directories from archiving.

Once you have carried out the steps in [“Accessing policies for editing” on page 162](#), proceed as follows.

1. **In the File Match Criteria table, press the Add button.**

The New Policy Criteria wizard starts.

2. **In the first step, Define File Match Criteria, enter the relative path from the target filesystem mount point to the Starting Directory.**
3. **If desired, enter a regular expression that matches the names of the files that you want the policy to govern in the Name Pattern text field. Then use the fields and controls provided to specify an Owner, a Minimum Size, a Group, a Maximum Size, and/or an Access Age.**

Be judicious when using these optional parameters, because they complicate the archiving task and increase the likelihood that archiving will produce unanticipated results. Remember: simple policies are usually best. Simple policies produce more predictable results, run faster, and are easier to debug if problems arise.

4. **After you add file-match criteria, select the desired Staging Behavior using the list control provided.**

You have the following choices:

- Stage these files automatically when they are accessed provides the best responsiveness under most conditions and is the best choice for most purposes.

Files are *staged*—copied back from tape to the disk cache—when they are called for and before they are read. This approach lets the user browse files on random-access disk media, which, for most files, is more responsive than searching sequentially through tape. The sequential-access tape drive operates more efficiently as well, because it can write the whole file out to disk in one go instead of constantly repositioning media in response to user requests.

- Mark these files for associative staging automatically stages all files within the target directory when any one of the associated files is accessed.

This method is useful when the contents of a directory are likely to be used together, so that a call for one file is likely to be followed by calls for the others. Examples might include multifile software-development or CAD projects and collections of medical images associated with a single patient's medical records.

- `Never stage` is useful when you are storing very large files, such as medical imagery or surveillance video, that can be *streamed* from tape to primary cache.

When files are large enough to be streamed from tape at the optimum transfer speed for the tape drive, sequential access is as fast as or faster than random access from disk. So large files can be read from tape to primary disk cache while the user is actually accessing the file. This just-in-time approach to transferring data saves the disk cache for the smaller files that are best accessed from a random I/O device. The approach also greatly simplifies cache management.

- Reset the stager behavior of these files to the default.

5. Select a Releaser Behavior using the list control provided.

You have the following choices:

- Stage these files automatically when they are accessed provides the best responsiveness under most conditions and is the best choice for most purposes.

Files are *staged*—copied back from tape to the disk cache—when they are called for and before they are read. This approach lets the user browse files on random-access disk media, which, for most files, is more responsive than searching sequentially through tape. The sequential-access tape drive operates more efficiently as well, because it can write the whole file out to disk in one go instead of constantly repositioning media in response to user requests.

- Mark these files for associative staging automatically stages all files within the target directory when any one of the associated files is accessed.

This method is useful when the contents of a directory are likely to be used together, so that a call for one file is likely to be followed by calls for the others. Examples might include multifile software-development or CAD projects and collections of medical images associated with a single patient's medical records. If used with files that are not, in fact, used together, this method makes poor use of the disk cache and may cause excessive mounting and positioning of tape media.

Note that you cannot use associative archiving if the files do not fit on a single tape volume. In this case, use either the `-sort path` or `-rsort path` parameters to keep the files contiguous.

- `Never stage` is useful when you are storing very large files, such as medical imagery or surveillance video, that can be *streamed* from tape to primary cache.

When files are large enough to be streamed from tape at the optimum transfer speed for the tape drive, sequential access is as fast as or faster than random access from disk. So large files can be read from tape to primary disk cache while

the user is actually accessing the file. This just-in-time approach to transferring data saves the disk cache for the smaller files that are best accessed from a random I/O device. The approach also greatly simplifies cache management.

- Reset the stager behavior of these files to the default.

6. **Select the target file system from the Apply to File System list control.**
7. **Press the Next button to proceed.**
8. **In the Summary page, review your entries.**
9. **Press the Finish button to commit your changes, Previous to make changes, or Cancel to abort.**

Stop here.

▼ Defining additional archive copies

1. **In the Copy Information table, press the Add button.**

The New Policy Copy wizard starts.

2. **In the next step, Copy Media Parameters, specify an Archive Age by entering a number in the text field provided and then selecting units—Week(s), Day(s), Hour(s), Minute(s), or Second(s)—using the accompanying list control.**

The archive age is the amount of time since a file's last modification.

3. **Specify a Media Type using the list controls provided.**
4. **Select a VSN Pool Name (media pool name) with the list control provided or Specify VSNs (tape volumes) by entering a range of volume serial numbers in the two accompanying text fields and/or by entering a regular expression or a comma-delimited list in the third text box.**

Be careful about over specifying. Use pools rather than ranges and ranges rather than particular volume serial numbers whenever possible. If you restrict the archivers choice of media, you may reduce media utilization and archiving performance, while increasing manipulation of tape media. An archiving operation might block waiting for media.

5. **If desired, Reserve the VSNs used for copies from the same Policy and/or File System by checking the corresponding check boxes. You can also use the accompanying list control to further restrict use of the VSN to files from the same Directory, User, or Group.**

Again, be careful about over specifying. If you restrict the archivers choice of media, you may reduce media utilization and archiving performance and may increase tape manipulation by increasing the need for mounting and positioning.

6. Press the Next button.

The Advanced Disk Copy Option panel appears.

7. In the Tuning Archive Media Writes section, you can use the list control to specify archiver behavior If the files to be archived are offline.

This option lets you specify how subsequent copies are made once the first copy is complete. Your choices are:

- Stage files back as they are needed copies files from the offline tape media to the disk cache whenever the archiver is ready to copy them to tape media.
- Copy them directly from offline media bypasses the disk cache and copies files directly from tape to tape.
- Stage one file ahead of the current copy pre-stages the next file from offline tape to disk while the current file is being copied to tape.
- Stage them all back first pre-stages all files from offline tape to disk before copying them to tape.

8. In the Managing the Archiver's Work Queue section, you can adjust the archiving Start Age, Start Count, and/or Start Size using the controls provided.

Be judicious about using this option. IAS comes preconfigured with Start Age, Start Count, and Start Size values that are optimal for most users. Changing these values complicates the solution, so be sure that you know what you are doing.

9. In the Archive Policy Based Recycling section, you can Disable Recycling for this policy by checking the corresponding checkbox.

Be careful with this option. If recycling is disabled, you cannot reclaim tape media for reuse after copies expire.

10. In the Archive Policy Based Recycling section, you can enter an email notification address in the Mail Recycler Information to field.

11. In the Archive Policy Based Recycling section, you can also set a Recycler High Water Mark by entering a percentage in the field provided.

A recycler high water mark sets the percentage of media usage below which recycling cannot occur. This percentage is the ratio of the used space in the library to its total capacity.

- 12. In the Archive Policy Based Recycling section, you can also enter a percentage in the Only recycle VSNS that would gain at least (%) field.**

Recycling a tape is a comparatively high-overhead process potentially involving multiple mounting and repositioning of tapes. Entering a value in this field lets you balance the costs of this overhead against the benefit of additional library capacity. If a recycling a tape will reclaim too little space to make recycling worthwhile, the recycler skips it.

- 13. Press the Next button to continue, Previous to return to the last step, or Cancel to abort.**
- 14. In the Summary page, review your entries.**
- 15. Press the Finish button to commit your changes, Previous to make changes, or Cancel to abort.**

Stop here.

▼ Modifying file-match criteria

Once you have carried out the steps in [“Accessing policies for editing” on page 162](#), proceed as follows.

- 1. In the Policy Name field of the Archive Policies table, click on the hyperlinked name of the policy that you want to edit.**
The *policy-name* Details page appears in the content pane.
- 2. In the File Match Criteria table, select the radio button that corresponds to the criterion that you want to modify. Press the Edit button.**
The *policy-name* Criteria Details page appears in the content pane.
- 3. In the first step, Define File Match Criteria, enter the relative path from the target filesystem mount point to the Starting Directory.**
- 4. If desired, enter a regular expression that matches the names of the files that you want the policy to govern in the Name Pattern text field. Then use the fields and controls provided to specify an Owner, a Minimum Size, a Group, a Maximum Size, and/or an Access Age.**
Be judicious when using these optional parameters, because they complicate the archiving task and increase the likelihood that archiving will produce unanticipated results. Remember: simple policies are usually best. Simple policies produce more predictable results, run faster, and are easier to debug if problems arise.
- 5. In the upper and lower right corners of the panel, press the Save button to commit your changes, Cancel to abort.**

Stop here.

▼ Modifying copy parameters

Once you have carried out the steps in [“Accessing policies for editing”](#) on page 162, proceed as follows.

1. In the `Copy Information` table, select the radio button that corresponds to the copy that you want to modify.
2. To modify the media type, tape media pool, or tape volumes used for the copy, press the `Modify VSN Assignments` button.
The `policy-name - Copy number Copy VSNs` page appears.
3. To modify the file-sorting method, unarchiving criteria, media writes and buffering, archive queuing, and recycling parameters, press the `Edit Advanced Options` button.
The `policy-name - Copy number Details` page appears.
4. If you are on the `policy-name - Copy number Copy VSNs` page, use the controls provided to modify the media-type, tape media-pool, and/or tape-volume specifications for the copy.
5. If you are on the `policy-name - Copy number Details` page, use the controls provided to modify the file-sorting method, unarchiving criteria, media writes and buffering, archive queuing, and/or recycling parameters.
6. If you want to enable associative archiving, go to the `Archive Organization` section and select `Force files in a directory to be archived together`.

Keeping files together on one piece of tape media can be useful when the files in a directory are related in some way and are thus likely to be used together. In such cases, keeping files together improves staging performance and reduces mounting and positioning of tape media.

Note that you cannot use associative archiving if the files do not fit on a single tape volume. In this case, use either the `-sort path` or `-rsort path` parameters to keep the files contiguous.

Remember, however, that forcing the archive to use contiguous media makes capacity utilization less efficient and management more complex. So use this technique only as necessary.

7. Press the `Save` button to commit your changes, `Cancel` to abort.

Stop here.

Controlling the archiver

This section explains how you control the state of the archiver, including:

- [“Stopping or idling archiving” on page 169](#)
- [“Running archiving” on page 169](#).

▼ Stopping or idling archiving

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. Proceed as follows.

1. **Login to the IAS Java Web Console under the `expert` or `admin` role.**
2. **If you are working with a high-availability, multinode appliance, use the `Selected Node` list control at the top of the IAS task tree to select the `MDS` node.**
3. **In the IAS task tree, expand the `Appliance Maintenance` branch.**
4. **On the `Appliance Maintenance` branch, expand `Archive Maintenance`.**
5. **Under `Archive Maintenance`, select `Archiver Activity`.**
The `Archive Activity` page appears in the content pane.
6. **To immediately stop archiving, select the `Stop` radio button.**
7. **To stop archiving at a logical point in the process, select the `Idle` radio button.**
8. **Specify the target file system by selecting it from the `For File System` list control.**
9. **Press the `Submit` button to commit your changes.**

Stop here.

▼ Running archiving

1. **Login to the IAS Java Web Console under the `expert` or `admin` role.**
2. **If you are working with a high-availability, multinode appliance, use the `Selected Node` list control at the top of the IAS task tree to select the `MDS` node.**
3. **In the IAS task tree, expand the `Appliance Maintenance` branch.**

4. **On the** Appliance Maintenance **branch, expand** Archive Maintenance.
5. **Under** Archive Maintenance, **select** Archiver Activity.
The Archive Activity page appears in the content pane.
6. **Select the** Run Now **radio button.**
7. **Specify the target file system by selecting it from the** For File System **list control.**
8. **Press the** Submit **button to commit your changes.**

Stop here.

Managing archive server nodes

The server nodes of Infinite Archive System (IAS) appliances are specially equipped and configured Sun server platforms running specially configured Sun Solaris operating systems and software. IAS server nodes generally require little human intervention, because system software automates most normal administrative tasks. However, when necessary, you can manage nodes directly. IAS appliances provide two interfaces:

- The Infinite Archive System (IAS) console software lets archive administrators reboot, place server nodes in maintenance mode, and, on high-availability systems, switchover (failover/failback) between nodes, all without direct physical access to the machine.
- The Integrated Lights Out Manager (ILOM) is a service processor that lets administrators access server node hardware, even when the operating system is not running and main system power is **OFF**. Archive administrators can use the ILOM to remotely power up a system that is physically connected to power but not switched **ON**.

This chapter outlines the procedures for using these tools to carry out the following tasks:

- [“Changing the operating state of a node” on page 172](#)
- [“Maintaining network connectivity” on page 187](#)
- [“Managing email notifications” on page 192](#)
- [“Maintaining time synchronization” on page 194.](#)

Changing the operating state of a node

When you reboot, shut down, place a server in maintenance mode, or, on high-availability systems, switchover (failover/failback), you change the operating state of the node and must therefore take steps to insure that data and system availability are properly maintained during the transition. This section documents the necessary procedures.

In particular, state changes always have some potential for disrupting filesystem operations. IAS software tries to insure that pending archive transactions are fully committed and consistent before proceeding. But, nonetheless, it is still best practice to manually stop archiving before you initiate a reboot, shutdown, failover/failback, or a switch to maintenance mode. Accordingly, this section begins with [“Quiescing the archiver” on page 172](#).

Quiescing the archiver

While IAS software tries to insure that pending file-system transactions are fully committed and consistent before proceeding, it is nonetheless best practice to manually stop archiving before you reboot, shut down, switchover (failover/failback), or place a server node in maintenance mode. Under some circumstances, when a server changes state, the SAM archiver process fails to write an end-of-file (**EOF**) marker on the media. SAM interprets the resulting inconsistency as a media error and marks the catalog entry for the volume with the `E` flag.

Clearing spurious media errors is outside the scope of this volume. But the relevant procedures can be found under *“To Clear Media Errors”*, *Sun StorageTek SAM Archive Configuration and Administration Guide*
<<http://docs.sun.com/source/819-7931-10/chap02.html#28098>>.

▼ Stopping archiver activity

1. **At the top of the IAS task tree, use the Selected Node list control to select the name of the MDS (metadata server) node.**
2. **In the IAS task tree, expand the Archive Maintenance branch, and select Archiver Activity.**

3. When the Archive Activity page appears in the content pane, go to Perform the following archiving action (A below) and click the Stop radio button (B).

The screenshot shows the 'Archive Activity' page with two columns of radio buttons. The left column is titled 'Perform the following archiving action:' and contains options: Restart, Idle, Run Now, Soft Restart (Rerun), and Stop. The right column is titled 'Perform the following staging action:' and contains options: Idle and Run Now. Below these columns are two 'For File System' dropdown menus, both set to 'All'. At the bottom are two 'Submit' buttons. Callouts A, B, C, and D point to the 'Perform the following archiving action:' header, the 'Stop' radio button, the 'All' dropdown menu, and the left 'Submit' button, respectively.

4. Accept the default in the For File System list control, All (C above)
5. Press the Submit button (D above).

When the console displays the message `ArchiveActivity.arstop` was issued successfully, archiving has stopped.

Next task: carry out the required state change, as described below:

- [“Shutting down server nodes” on page 173](#)
- [“Starting up server nodes” on page 179](#)
- [“Rebooting” on page 183](#)
- [“Switchover \(failover and failback\)” on page 185.](#)

Shutting down server nodes

When you power down IAS server nodes, you stop the software processes, stop the operating systems and turn off power to processors, disks, and peripherals. Whenever possible, you thus want to power down *gracefully*, so that the software and operating system have time to stop running processes and write unsaved data out to disk. If you do not shut the system down in an orderly fashion, you risk corrupting system configuration or file systems. This section explains the proper procedure.

▼ Shutting down server processes from the IAS console

1. At the top of the IAS task tree, use the Selected Node list control to select the name of the MDS (metadata server) node.
2. In the IAS task tree, expand the Appliance Maintenance branch.
3. Under the Appliance Maintenance branch, click Appliance State Operations.
The Appliance State Operations table appears in the content pane at right.
4. If you are working on a high-availability, multinode system, click the radio button beside the server node that you want to shut down or select Appliance to select both (A below).

Appliance State

Appliance State Operations (3)

-- Operations --

<input type="checkbox"/>	Node Name	Node Role	Current State
<input checked="" type="radio"/>	Appliance	Appliance	Operational
<input type="radio"/>	IASserial#n01	MDS	Online
<input type="radio"/>	IASserial#n02	Client	Online

Disable Preventive Maintenance Reports

Time Interval: Hour

If you are working on a single-node system, as shown below, you do not need to select a server node.

Appliance State

Appliance State Operations (3)

-- Operations --

Node Name	Node Role	Current State
Appliance	Appliance	Operational

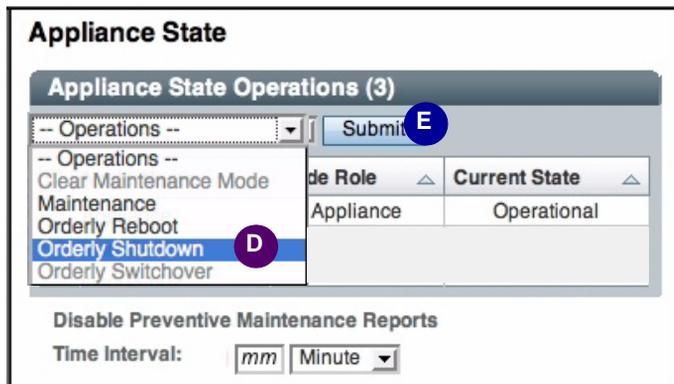
Disable Preventive Maintenance Reports

Time Interval: Minute

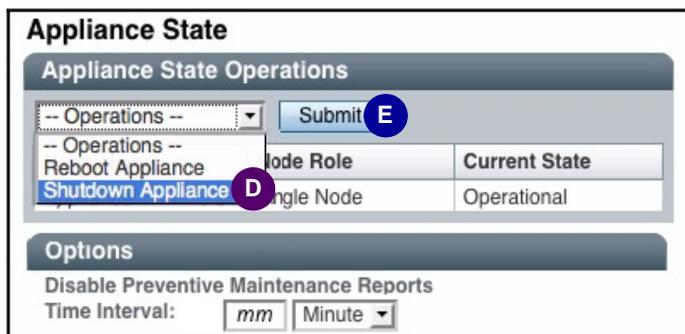
5. To turn off alert reporting during scheduled down period, enter an interval in the Time Interval text box and select units (Minute, Hour, or Day) using the Time Interval list control (B above).

- Expand the **Operations** list control (C above), and select **Orderly Shutdown** (D below).

The state controls for a dual-node high-availability appliance are shown immediately below:



The simpler controls for the single-node appliance shown below:



- Press **Submit** (E above) to execute the operation.

The IAS application stops server processes, shuts down the operating system, and, on Sun SPARC-based IAS models, shuts off system power (powering off Sun x86-based models requires additional steps).

Be patient. It can take some time to prepare an active server node for shutdown.

Next task: [“Turning system power OFF” on page 175.](#)

▼ Turning system power OFF

Caution – turning the *system power* **OFF** does not cut off all power to the enclosure! An orderly shutdown turns enclosure main power **OFF** and powers down the server. But *standby power* remains **ON**, and current continues to flow to the GRASP service

processor board and the power supply fans, as shown by the blinking power-status LEDs on the chassis. If you must disconnect *all* power (as you must for some service tasks), you also need to physically disconnect the power cords from the back of the server node. The blinking power-status LEDs then go dark.

1. If you are powering down an IAS server node remotely, stop here.

The IAS Java Web Console user interface automatically shuts off the system power after it stops the operating system. Standby power is, of course, still **ON**, as described in the caution above.

2. If you are powering down an IAS model that incorporates Sun x86 servers, such as the Sun Fire X4540, and if you have physical access to the server, you can safely press the power button located on the front of the IAS server node (A below).



When system power is OFF, the power LED (**B** above) blinks to indicate that the chassis and ILOM GRASP board are on standby power. Note that you must physically disconnect the power cords from the server-node power supplies when you need to shut off all power from the server node.

3. If you shut off the system power using the power button, stop here.

4. If you are powering down an IAS model x86 server node and *do not* have physical access to the server, you can turn system power OFF remotely using the ILOM service processor. See the **Next task listed at the end of this procedure.**

Next task: If you need to power down an x86 server node remotely, go to [“Remotely controlling server nodes with the ILOM” on page 177](#). Otherwise, stop here.

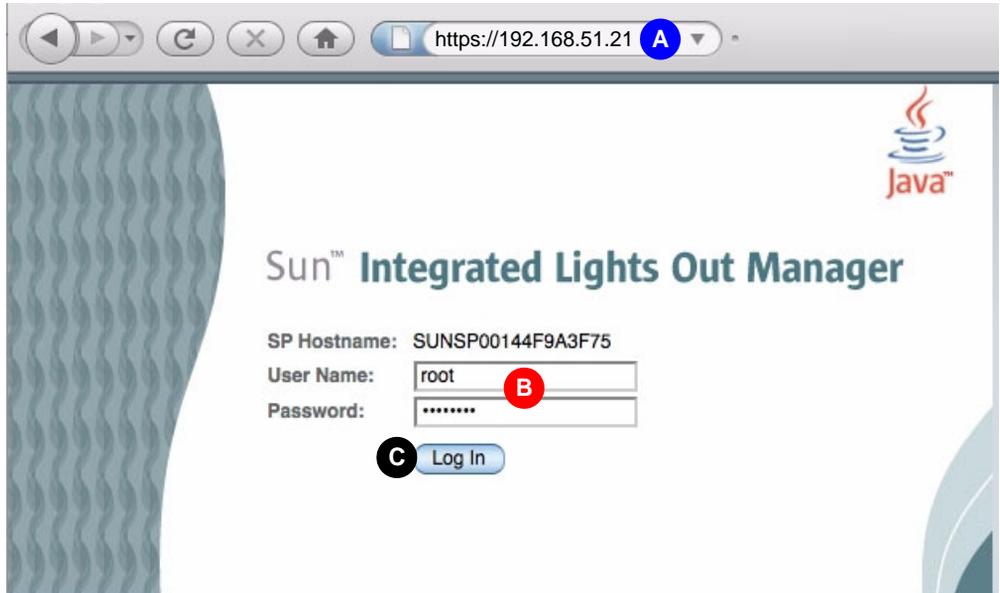
▼ Remotely controlling server nodes with the ILOM

The Sun Integrated Lights Out Manager is a service processor that can control the server node hardware independently of the operating system. A webserver running on the ILOM gives you access to a browser-based graphical management interface.

1. **Log on to a management host computer that is attached to the Customer IAS Mgt Connection, port 23, of the IAS appliance patch panel and configured with the private IP address of the Customer IAS Mgt Connection: 192.168.51.250.**

If you so request, Sun services personnel can connect a customer-supplied management host to the IAS appliance during installation.

2. **On the management host, open a compatible browser and enter the management URL for the desired node into the navigation bar of the browser (A below):**
 - <https://192.168.51.21> for node 01
 - <https://192.168.51.31> for node 02 (high-availability servers only).



Mozilla Firefox is recommended. Other standards-compliant browsers may work well. But Microsoft Internet Explorer is generally not compatible with IAS interfaces.

3. When the login page appears, log in as `root`, and enter the ILOM `root` password (B above). Press Log In (C) to open the ILOM interface (D below).

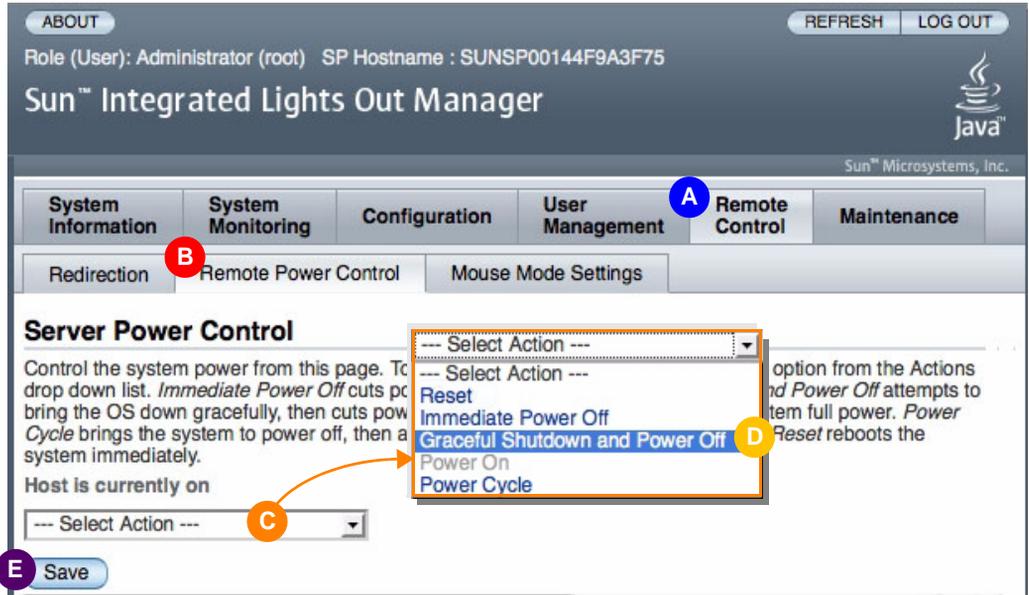
The screenshot displays the Sun Integrated Lights Out Manager (ILOM) web interface. At the top, there is a navigation bar with 'ABOUT', 'REFRESH', and 'LOG OUT' buttons. Below this, the user role is identified as 'Administrator (root)' and the SP Hostname is 'SUNSP00144F9A3F75'. The main title is 'Sun™ Integrated Lights Out Manager' with the Java logo on the right. A menu bar contains 'System Information', 'System Monitoring', 'Configuration', 'User Management', 'Remote Control', and 'Maintenance'. Under 'System Information', there are sub-tabs for 'Versions', 'Session Time-Out', and 'Components'. The 'Versions' page is active, showing the title 'Versions' and a subtitle 'View the version of ILOM firmware currently in use.' Below this is a 'Version Information' table with the following data:

Property	Value
SP Firmware Version	1.1.8
SP Firmware Build Number	18244
SP Firmware Date	Sat Apr 21 11:25:11 PDT 2007
SP Filesystem Version	0.1.14

Next task: [“Powering down the server node using the ILOM” on page 179.](#)

▼ Powering down the server node using the ILOM

1. In the ILOM graphical user interface, click on `Remote Control` (A below) and select the `Remote Power Control` tab (B).



2. Under `Server Power Control`, open the `Select Action` list control (C above) and select `Graceful Shutdown and Power Off` (D).
3. Press the `Save` button (E above).

The server node powers down to standby mode. System power is `OFF`. Standby power remains `ON` until power cables are disconnected from the power supplies.

Stop here.

Starting up server nodes

This section describes the procedures for restoring system power to IAS server nodes that are currently on standby power. If the power cables are currently disconnected from the power supplies and/or wall power, you must reconnect them and restore standby power before proceeding.

▼ Turning system power ON

1. If you have physical access to the server, you can safely press the power button located on the front of the IAS server node (A below).



When system power is **ON** the power LED (**B** above) stops blinking and shines steadily. Note that you must physically reconnect the power cords from the server-node power supplies before you can restore system power to the server node.

2. If you *do not* have physical access to the server, you can turn system power on remotely using the ILOM service processor. See the **Next** task listed at the end of this procedure.

Next task: [“Remotely controlling server nodes with the ILOM” on page 181.](#)

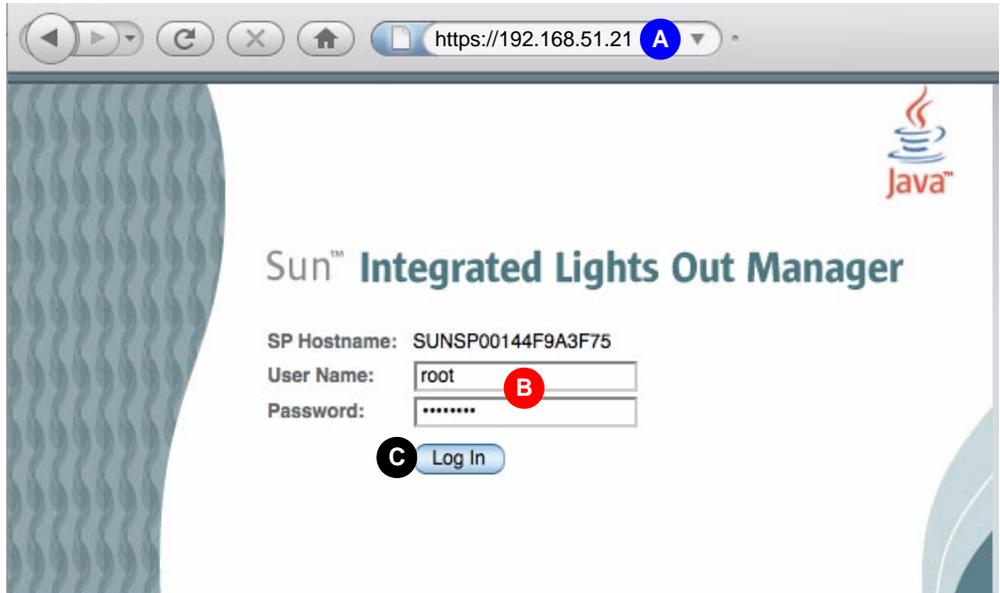
▼ Remotely controlling server nodes with the ILOM

The Sun Integrated Lights Out Manager is a service processor that can control the server node hardware independently of the operating system. A webserver running on the ILOM gives you access to a browser-based graphical management interface.

1. **Log on to a management host computer that is attached to the Customer IAS Mgt Connection, port 23, of the IAS appliance patch panel and configured with the private IP address of the Customer IAS Mgt Connection: 192.168.51.250.**

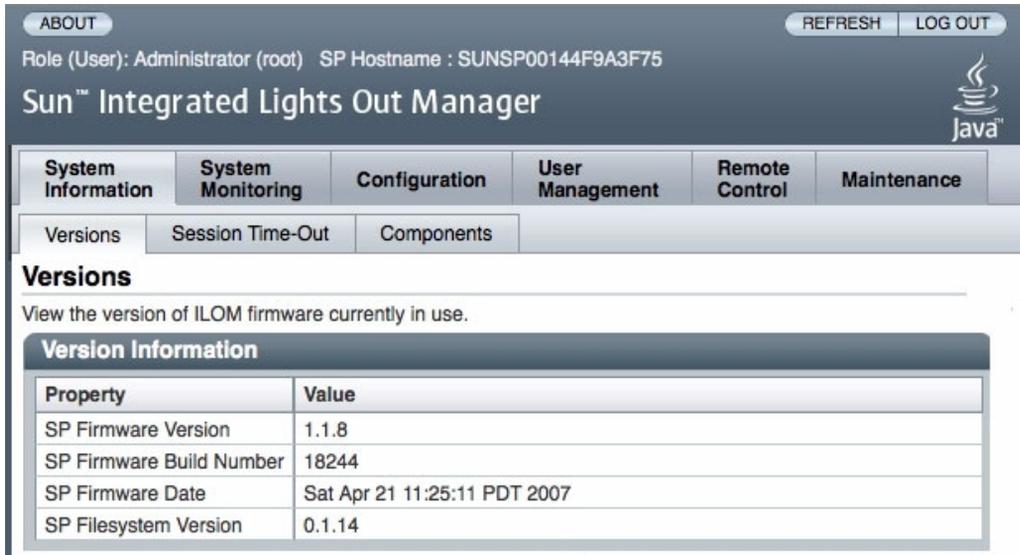
If you so request, Sun services personnel can connect a customer-supplied management host to the IAS appliance during installation.

2. **On the management host, open a compatible browser and enter the management URL for the desired node into the navigation bar of the browser (A below):**
 - <https://192.168.51.21> for node 01
 - <https://192.168.51.31> for node 02 (high-availability servers only).



Mozilla Firefox is recommended. Other standards-compliant browsers may work well. But Microsoft Internet Explorer is generally not compatible with IAS interfaces.

3. When the login page appears, log in as `root`, and enter the ILOM `root` password (B above). Press Log In (C) to open the ILOM interface (D below).



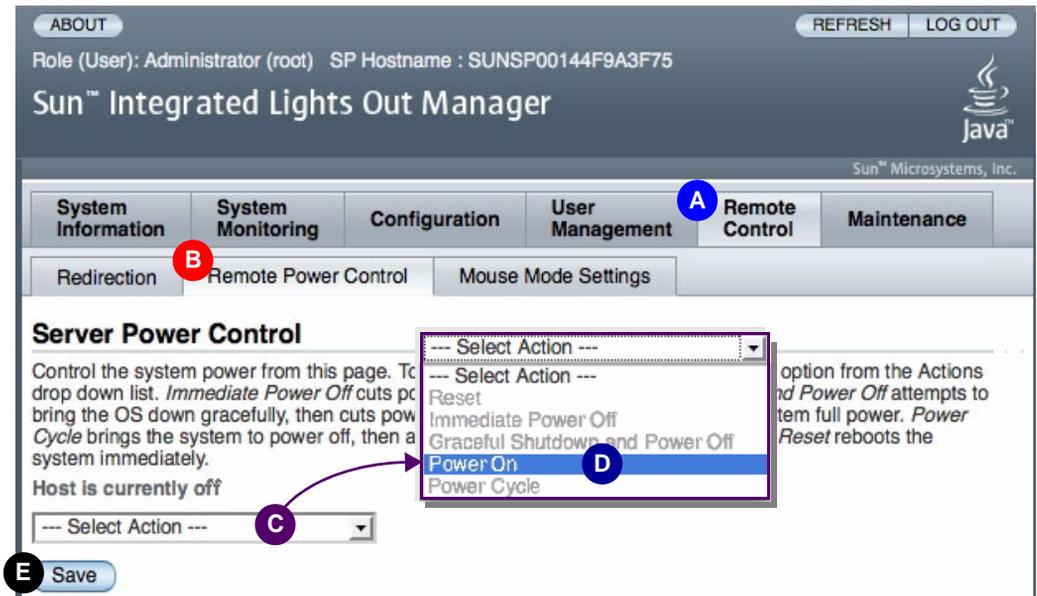
The screenshot displays the Sun Integrated Lights Out Manager (ILOM) web interface. At the top, there is a navigation bar with 'ABOUT', 'REFRESH', and 'LOG OUT' buttons. Below this, the user role is identified as 'Administrator (root)' and the SP Hostname is 'SUNSP00144F9A3F75'. The main title is 'Sun™ Integrated Lights Out Manager' with the Java logo on the right. A menu bar contains 'System Information', 'System Monitoring', 'Configuration', 'User Management', 'Remote Control', and 'Maintenance'. Under 'System Information', there are sub-tabs for 'Versions', 'Session Time-Out', and 'Components'. The 'Versions' section is active, showing the text 'View the version of ILOM firmware currently in use.' Below this is a 'Version Information' table.

Property	Value
SP Firmware Version	1.1.8
SP Firmware Build Number	18244
SP Firmware Date	Sat Apr 21 11:25:11 PDT 2007
SP Filesystem Version	0.1.14

Next task: [“Powering up the server node using the ILOM” on page 183.](#)

▼ Powering up the server node using the ILOM

1. In the ILOM graphical user interface, click on `Remote Control` (A below) and select the `Remote Power Control` tab (B).



2. Under `Server Power Control`, open the `Select Action` list control (C above) and select `Power On` (D).
3. Press the `Save` button (E above).

The server node powers up.

Stop here.

Rebooting

When you reboot an IAS server node, you restart the operating system and reinitialize the system hardware. This section explains the proper procedures.

▼ Rebooting a node or appliance from the IAS console

The `Appliance State Operations` page of the IAS console lets you reboot nodes.

1. At the top of the IAS task tree, use the Selected Node list control to select the name of the MDS (metadata server) node.
2. In the IAS task tree, expand the Appliance Maintenance branch.
3. Under the Appliance Maintenance branch, click Appliance State Operations.
4. If you are working on a high-availability, multinode system, click the radio button beside the server node that you want to reboot (A below).

You cannot reboot both nodes at once, so do not select Appliance.

Appliance State

Appliance State Operations (3)

-- Operations -- **C** Submit

	Node Name	Node Role	Current State
A	Appliance	Appliance	Operational
<input checked="" type="radio"/>	IASserial#n01	MDS	Online
<input type="radio"/>	IASserial#n02	Client	Online

Disable Preventive Maintenance Reports

Time Interval: **B** mm Hour

If you are working on a single-node system, as shown below, you do not need to select a server node.

Appliance State

Appliance State Operations (3)

-- Operations -- **C** Submit

<input type="radio"/>	Node Name	Node Role	Current State
<input type="radio"/>	Appliance	Appliance	Operational

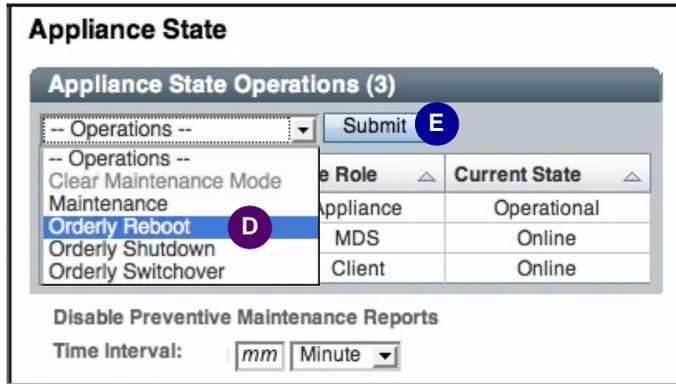
Disable Preventive Maintenance Reports

Time Interval: **B** mm Minute

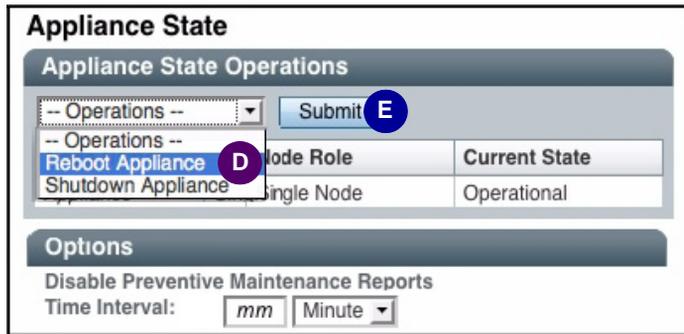
5. To turn off alert reporting during scheduled down period, enter an interval in the Time Interval text box and select units (Minute, Hour, or Day) using the Time Interval list control (B above).

- Expand the **Operations** list control (C above), and select **Orderly Reboot** from the list (D below).

The state controls for a dual-node high-availability appliance are shown immediately below:



The simpler controls for the single-node appliance shown below:



- Press **Submit** (E above) to execute the operation.

Be patient. The reboot process can take 3-5 minutes (or longer on an active archive).

Stop here.

Switchover (failover and failback)

On high-availability systems, switching over is the process of manually initiating a failover from the `mds` (the active metadata server) node to the `client` (the standby metadata server) node. This is useful when you need to perform disruptive maintenance on the MDS and need to keep the archive fully available. .

▼ Switching over to the other server node

1. At the top of the IAS task tree, use the `Selected Node` list control to select the name of the MDS (metadata server) node.
2. In the IAS task tree, expand the `Appliance Maintenance` branch.
3. Under the `Appliance Maintenance` branch, click `Appliance State Operations`.
4. When the `Appliance State Operations` table appears in the content pane of the IAS console, click the radio button beside the name of the MDS (metadata server) node (A below).

You can only initiate switchover from the active metadata data server node.

Appliance State

Appliance State Operations (3)

-- Operations -- C

	Node Name	Node Role	Current State
<input type="radio"/>	Appliance	Appliance	Operational
<input checked="" type="radio"/>	IASserial#n01	MDS	Online
<input type="radio"/>	IASserial#n02	Client	Online

Disable Preventive Maintenance Reports

Time Interval: B

5. To turn off alert reporting during scheduled down period, enter an interval in the `Time Interval` text box and select units (Minute, Hour, or Day) using the `Time Interval` list control (B above).
6. Expand the `Operations` list control (C above), and select `Orderly Switchover` from the list (D below).

Appliance State

Appliance State Operations (3)

-- Operations -- E

-- Operations --

Clear Maintenance Mode

Maintenance

Orderly Reboot

Orderly Shutdown

D Orderly Switchover

	Node Name	Node Role	Current State
<input type="radio"/>	Appliance	Appliance	Operational
<input type="radio"/>	MDS	MDS	Online
<input type="radio"/>	Client	Client	Online

Disable Preventive Maintenance Reports

Time Interval:

7. **Press** `Submit` (**E** above) to execute the operation.

Be patient. The shutdown process can take some time on an active archive.

Stop here. If you wish to failback to the starting MDS node (when maintenance is completed, for example), simply repeat the process above.

Maintaining network connectivity

This section covers maintenance of the Ethernet local area networks (LANs) that connect archive clients to the archive and IAS Java Web Console user interfaces to the archive appliance. Sun service personnel configure IAS LAN connections during initial IAS installation. If you need to change them subsequently, use one of the following procedures:

- [“Setting IP addresses on high-availability systems” on page 187](#)
- [“Setting IP addresses on single-node systems” on page 189](#)
- [“Maintaining network routes” on page 190.](#)

Setting IP addresses on high-availability systems

High-availability IAS appliances can have up to five local area network connections:

- One to four `Data Network` connections support high-availability network file sharing between the archive and archive clients, such as application hosts and user workstations.
- One `Management Network` supports high-availability HTTPS communication between the IAS server nodes and management stations that access the IAS Java Web Console graphical user interface.

Each network includes one or two virtual IP addresses mapped to two static, physical IP addresses. Application hosts and user workstations connect to the appliance using the virtual addresses. Normally, each virtual address points to a physical port address on one of the two IAS server nodes. But, if an IAS Ethernet port or node fails, the affected virtual IP address automatically fails over and shares the physical address on the other node. The IAS Java Web Console user interface, web server, and shared NFS and SAMBA file systems remain available, albeit at reduced bandwidth.

▼ Setting IP addresses on high-availability systems

You can configure from one to four data networks and a management network on a clustered, high-availability IAS appliance. To modify your network configuration, proceed as follows.

1. Use the `Selected Node` list control at the top of the IAS task tree to select the MDS node.
2. In the IAS task tree, expand the `Appliance Config` branch and select `Network Maintenance`.
3. Select the `Network Maintenance` tab (A below).

The tab displays address configuration tables for one to four `Data Networks` (B) and one `Management Network` (C).

A Network Maintenance | Network Routes

B Network Maintenance * Indicates required field

Data Network 1	
D * Static IP 1:	<input type="text" value="IP-address-01"/>
* Static IP 2:	<input type="text" value="IP-address-02"/>
E * Virtual IP 1:	<input type="text" value="IP-address-03"/>
* Virtual IP 2:	<input type="text" value="IP-address-04"/>
F * Netmask:	<input type="text" value="netmask"/>

Data Network 2	
Static IP 1:	<input type="text" value="IP-address-11"/>
Static IP 2:	<input type="text" value="IP-address-12"/>
Virtual IP 1:	<input type="text" value="IP-address-13"/>
Virtual IP 2:	<input type="text" value="IP-address-14"/>
Netmask:	<input type="text" value="netmask"/>

Data Network 3	
Static IP 1:	<input type="text" value="IP-address-21"/>
Static IP 2:	<input type="text" value="IP-address-22"/>
Virtual IP 1:	<input type="text" value="IP-address-23"/>
Virtual IP 2:	<input type="text" value="IP-address-24"/>
Netmask:	<input type="text" value="netmask"/>

Data Network 4	
Static IP 1:	<input type="text" value="IP-address-31"/>
Static IP 2:	<input type="text" value="IP-address-32"/>
Virtual IP 1:	<input type="text" value="IP-address-33"/>
Virtual IP 2:	<input type="text" value="IP-address-34"/>
Netmask:	<input type="text" value="netmask"/>

C Management Network

D * Static IP 1:	<input type="text" value="IP-address-41"/>
* Static IP 2:	<input type="text" value="IP-address-42"/>
E * Virtual IP:	<input type="text" value="IP-address-43"/>
F * Netmask:	<input type="text" value="netmask"/>

G Save Reset

4. For each network connection that you wish to modify, enter `Static IP` addresses (D), `Virtual IP` addresses (E), and a `Netmask` (F).

5. Press **Save** to make your changes, **Reset** to discard them (G).

Stop here.

Setting IP addresses on single-node systems

An unclustered, single-node IAS appliance has eight Ethernet interfaces that can be connected to a local area network. There are four onboard Gig-E ports and four additional Gig-E ports on a network interface card. Each port is identified by a Solaris device name (in the examples below, `e1000g0-3` and `nge0-3`).

▼ Setting IP addresses on single-node systems

1. In the IAS task tree, expand the **Appliance Config** branch and select **Network Maintenance**.
2. Select the **Network Maintenance** tab (A below).

The tab displays address configuration tables for each of the **Network Interfaces** (B).

The screenshot shows the 'Network Maintenance' configuration page. At the top, there are two tabs: 'Network Maintenance' (labeled A) and 'Network Routes'. Below the tabs, the page is titled 'Network Maintenance'. There are three sections for network interfaces, each labeled B: 'Network Interface #1 (e1000g0)', 'Network Interface #2 (e1000g1)', and 'Network Interface #8 (nge3)'. Each interface section has two input fields: 'IP Address' (labeled C) and 'Netmask' (labeled D). The 'IP Address' field for the first interface contains the text 'IP-address-01' and the 'Netmask' field contains 'netmask'. At the bottom right of the page, there are two buttons: 'Save' (labeled E) and 'Reset'.

3. For each network connection that you wish to modify, enter an **IP Address** (C above) and a **Netmask** (D).
4. Press **Save** to make your changes, **Reset** to discard them (E).

Stop here.

Maintaining network routes

Network route maintenance consists of the following tasks:

- [“Viewing a network route” on page 190](#)
- [“Adding a network route” on page 190](#)
- [“Deleting a network route” on page 191.](#)

▼ Viewing a network route

1. Use the `Selected Node` list control at the top of the IAS task tree to select the **MDS node**.
2. In the IAS task tree, expand the `Appliance Config` branch and select `Network Maintenance`.
3. To view network routing information, select the `Network Routes` tab of the content window.

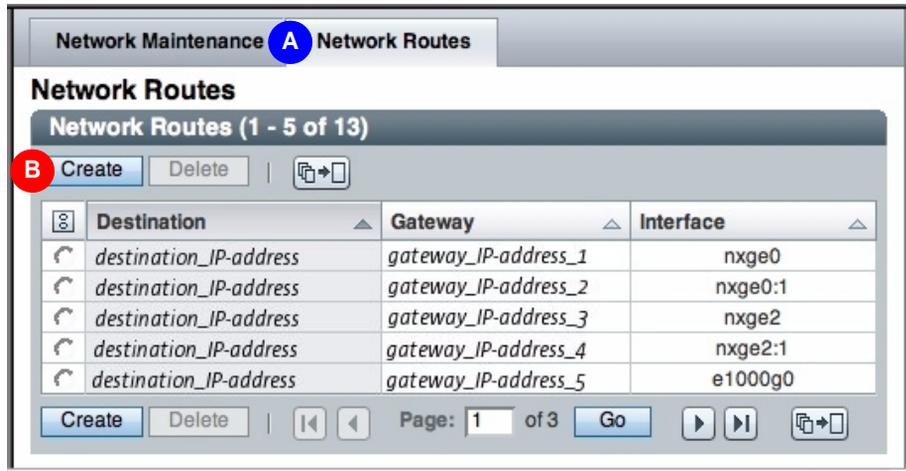
For instructions on displaying, browsing, and sorting network routes, see [“Tables” on page 39](#).

Stop here.

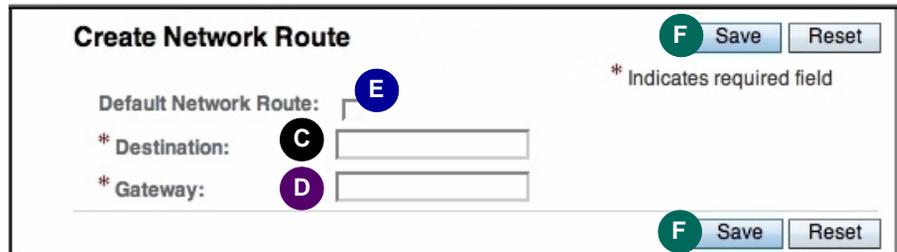
▼ Adding a network route

1. Use the `Selected Node` list control at the top of the IAS task tree to select the **MDS node**.
2. In the IAS task tree, expand the `Appliance Config` branch and select `Network Maintenance`.

3. Select the `Network Routes` tab of the content window (A below).



4. Press `Create` (B above).
5. When the `Create Network Route` dialog appears, enter a `Destination IP` address (C below) and a `Gateway IP` address (D).



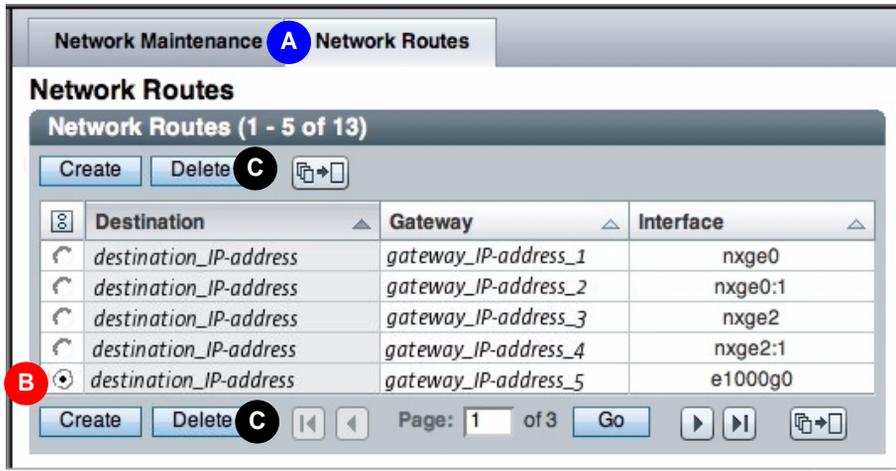
6. If you are adding a `Default Network Route`, check the checkbox (E above).
7. Press `Save` to make your changes, `Reset` to discard them (F above).

Stop here.

▼ Deleting a network route

1. Use the `Selected Node` list control at the top of the IAS task tree to select the `MDS` node.
2. In the IAS task tree, expand the `Appliance Config` branch and select `Network Maintenance`.

3. Select the `Network Routes` tab of the content window (A below).



4. Click the radio button corresponding to the route that you wish to delete (B above).
5. Press `Delete` (C above).
6. When the browser prompts you, confirm or cancel the deletion.

Stop here.

Managing email notifications

Infinite Archive System (IAS) email alerts inform a user of specified events and conditions by generating a message and sending it to a designated email address at specified intervals (while the error condition persists). The recipient is the `root` account.

▼ Creating an Email Alert

1. Use the `Selected Node` list control at the top of the IAS task tree to select the `MDS` node.
2. In the IAS task tree, expand the `Appliance Config` branch and select `Email`. The `Email Notifications` page appears in the content pane.

Email Notifications E Save Reset

A Notification Level 1
Email Addresses: **B**
Time Interval: **C** **D** HOUR

A Notification Level 2
Email Addresses: **B**
Time Interval: **C** **D** HOUR

A Notification Level 3
Email Addresses: **B**
Time Interval: **C** **D** HOUR

A Notification Level 4
Email Addresses: **B**
Time Interval: **C** **D** HOUR

E Save Reset

3. For each of the four `Notification Levels` displayed (A above), enter a comma-delimited list of one or more `Email Addresses` in the corresponding text field (B) and specify a `Time Interval` value (C) and units (D) using the controls provided.

The `Notification Levels` let you configure an automatic escalation chain within your organization.

4. Press the `Submit` button (E above) to commit your changes, `Cancel` to abort.

Stop here.

Maintaining time synchronization

IAS appliances maintain consistent timing with Network Time Protocol (NTP). NTP keeps all of the system clocks on a network synchronized, so that the time and date are the same for all devices. To configure NTP, use one of the following methods:

- “Using external NTP servers to maintain the time and date” on page 194
- “Using internal NTP servers to maintain the time and date” on page 195.

▼ Using external NTP servers to maintain the time and date

To set and maintain the date and time automatically, synchronize the internal Network Time Protocol (NTP) servers on the IAS server nodes with the network-wide timing values established by NTP servers on your network. Proceed as follows.

1. If you are working on a high-availability, multinode system, use the Selected Node list control at the top of the IAS task tree to select the MDS node.
2. In the IAS task tree, expand the Appliance Config branch and select Time Zone / NTP.
3. Select the Network Time Protocol tab of the content window (A below).

The screenshot shows the 'Network Time Protocol' configuration window. It features a tabbed interface with 'Time Zone' and 'Network Time Protocol' tabs. The 'Network Time Protocol' tab is active. The main content area is titled 'Set up to three External Network Time Protocol (NTP) Servers.' and contains two radio buttons: 'Enable NTP Servers' (selected) and 'Disable NTP Servers'. Below these are three rows for NTP servers, each with a label, a text box for the IP address, and a 'Test Access to NTP Server' button. At the bottom right are 'Save' and 'Reset' buttons.

4. Click the Enable NTP Servers radio button (B above).
5. Enter an IP address for each NTP server that you want to configure using the NTP Server ... text box (C above), and validate your typing by pressing the corresponding Test Access to NTP Server ... button (D).
6. Press Save to make your changes, Reset to discard them (E above).

Stop here.

▼ Using internal NTP servers to maintain the time and date

If Network Time Protocol (NTP) servers are not available on the network or if local timing is desirable, configure the appliance to synchronize its clocks with its own internal NTP servers. Proceed as follows.

1. If you are working on a high-availability, multinode system, use the `Selected Node` list control at the top of the IAS task tree to select the node that you want to configure.

2. In the IAS task tree, expand the `Appliance Config` branch and select `Time Zone / Zone / NTP`.

3. If you have just set the date and/or time, wait 5-10 minutes before proceeding, so that the previous change can complete.

If you attempt to update the date and time too soon, the appliance nodes may still be adjusting to the previous change, in which case you may see an NTP error.

4. If an external/remote NTP server currently handles timing for the appliance, start by changing to local NTP control. Select the `Network Time Protocol` tab of the content window (A below), click the `Disable NTP Servers` radio button (B), and press `Save` (C).

The screenshot shows the 'Network Time Protocol' configuration window. At the top, there are two tabs: 'Time Zone' and 'Network Time Protocol'. The 'Network Time Protocol' tab is selected, indicated by a blue circle with the letter 'A'. Below the tabs, the text reads 'Set up to three External Network Time Protocol (NTP) Servers.' There are two radio buttons: 'Enable NTP Servers' and 'Disable NTP Servers'. The 'Disable NTP Servers' radio button is selected, indicated by a red circle with the letter 'B'. Below the radio buttons, there are three rows for NTP servers. Each row has a label 'NTP Server 1:', 'NTP Server 2:', and 'NTP Server 3:' followed by a text input field containing 'IP-address-ntp1', 'IP-address-ntp2', and 'IP-address-ntp3' respectively. To the right of each input field is a 'Test Access to NTP Server' button. At the bottom right of the window, there are two buttons: 'Save' and 'Reset'. The 'Save' button is highlighted with a black circle and the letter 'C'.

5. Select the `Time Zone` tab of the content window (D below).

The screenshot shows the 'Time Zone' configuration window. At the top, there are two tabs: 'Time Zone' and 'Network Time Protocol'. The 'Time Zone' tab is selected, indicated by a black circle with the letter 'D'. Below the tabs, the text reads 'Set the time zone for the Appliance.' There is a label '* Time Zone:' followed by a dropdown menu. The dropdown menu is open, showing 'Greenwich' selected, indicated by a green circle with the letter 'E'. Below the dropdown menu, the text reads 'In order for the Appliance Time Zone change to take effect, the Appliance must be re-started.' At the bottom right of the window, there are two buttons: 'Save' and 'Reset'. The 'Save' button is highlighted with a blue circle and the letter 'F'.

6. **Select a time zone using the Time Zone list controls (E above).**
7. **Press Save to make your changes, Reset to discard them (F above).**
8. **If you see an Exception Saving Date Time followed by a Network Time Protocol (NTP) error and if you have attempted to reset the time twice in succession after only a short interval, retry the change later.**

Note that you must restart the appliance (both nodes) for your changes to take effect.

Next task: [“Rebooting” on page 183.](#)

Maintaining the IAS disk arrays

Sun preconfigures the disk storage arrays in Infinite Archive System (IAS) appliances to fit the needs of specific archiving solutions. So, under normal circumstances, the arrays should require no user configuration or maintenance as such. To insure the health of your disk storage subsystem and protect your data, you need only attend to the following tasks:

- [“Scheduling repairs promptly if needed” on page 197](#)
- [“Monitoring disk hardware regularly” on page 198.](#)

Scheduling repairs promptly if needed

The RAID storage architecture used in IAS appliances automatically protects data against storage hardware failures. RAID storage is highly reliable and self-healing. Data is stored redundantly on sets of multiple disks, so that no single disk failure can cause loss of data. For additional security, the IAS factory configuration allocates two additional disks as hot spares, so that up to two RAID sets can automatically replace a failed disk and restore full redundancy within a day or so.

Nonetheless, whenever a disk fails, you need to schedule replacement as soon as possible. Once the two hot spare disks have been used, the array cannot rebuild a fully redundant RAID set. The resulting *degraded* RAID set functions normally—no data is lost yet—but it cannot tolerate subsequent disk failures. If you let things get to this point, you use up the protections built into the IAS storage arrays and make data loss significantly more likely.

Monitoring disk hardware regularly

Disk failures can be difficult to detect, so be diligent about [“Monitoring disk-storage hardware alarms”](#) (see page 64). IAS storage arrays are so well protected that you may not notice any change in the functionality of the archive following disk problems—until it is too late. Monitoring the health of the array lets you correct problems long before they have any effect on your data.

You should also register your IAS appliance if it has not already been registered. By [“Accepting ASR support”](#) (see page 210) during the registration process, you allow IAS disk array-management software to submit an automatic, online Auto Service Request to Sun services whenever a disk fails. Normally, a Sun technician fills out and submits the online registration forms during initial installation and configuration. But if, for whatever reason, this was not done or if you elected to decline ASR support during installation, you can register or reregister using the procedure outlined in [Appendix A, “Registering for the ASR feature”](#) on page 207.

Maintaining registration information

The Infinite Archive System appliance software includes a registrations management page that lets you maintain your appliance and component registration information using a single, centralized interface. On the `Registrations` page you can carry out the following maintenance tasks:

- [“Entering Sun Online Account Information” on page 200](#)
- [“Entering Contact Information” on page 201](#)
- [“Entering Site Information” on page 202](#)
- [“Specifying Internet Connection Settings” on page 203](#)
- [“Reading the privacy statement and saving changes” on page 203.](#)

▼ Entering Sun Online Account Information

1. If you are working with a high-availability, multinode appliance, use the Selected Node list control at the top of the IAS task tree to select the MDS node.
2. In the IAS task tree, expand the Appliance Config branch and select Registrations.

The Registrations form appears. Note that the hypertext links at the top of the page let you quickly jump to any part of the form (A below), while Back to the top hypertext links let you return (B).

The screenshot shows a web form titled "Registrations" with "Save" and "Reset" buttons. Below the title is a sub-header "Sun Online Account Information" and a paragraph of instructions. A navigation menu contains links for "Sun Online Account Information", "Site Information", "Privacy Statement", "Contact Information", and "Internet Connection Settings". A legend indicates that an asterisk (*) denotes a required field. The form includes two input fields: "Sun Online Account Name" (marked with C) and "Password" (marked with D). A "Back to top" link (marked with B) is at the bottom left. A link for "Sun Online Account Registration" (marked with E) is within the instructions text. Callout A points to the navigation menu.

3. If you already have a Sun Online account, enter the Account Name (C above) and Password (D) in the fields provided.
4. If you do not have a Sun Online account, click the Sun Online Account Registration hypertext link (E above) to create one now.

Next task: To continue, go to “[Entering Contact Information](#)” on page 201. To save or clear changes, go to “[Reading the privacy statement and saving changes](#)” on page 203.

▼ Entering Contact Information

1. Scroll down to the **Contact Information** section of the Registrations page (A below).

A **Contact Information**

Enter the name and email address of the person who will receive subscription summary information from Sun. The contact must be the same person who has a Sun Online Account.

B * First Name:

C * Last Name:

D Telephone Number:

E * Email Address:

2. In the text box provided, enter the **First Name** of the person who has the Sun Online account (B above).
3. In the text box provided, enter the **Last Name** of the person who has the Sun Online account (B above).
4. If desired, enter the contact person's **Telephone Number** in the text box provided (C above).
5. In the text box provided, enter the contact person's **Email Address** (D above).

Next task: To continue, go to [“Entering Site Information”](#) on page 202. To save or clear changes, go to [“Reading the privacy statement and saving changes”](#) on page 203.

▼ Entering Site Information

1. Scroll down to the **Site Information** section of the Registrations page (A below).

The screenshot shows a registration form titled "Site Information" (A). The form contains the following fields and options:

- B** * Company Name: [Text box] Company name must match that associated with a Sun Online Account.
- C** * Site Name: [Text box]
- D** Address: [Text box]
- D** Address 2: [Text box]
- E** * City: [Text box]
- F** State/Province: [Text box]
- G** Postal Code: [Text box]
- H** * Country: [Dropdown menu] United States
- J** * Contract ID: [Text box]
- Registration externally with Sun: **K**

2. In the text box provided, enter the **Company Name** (B above).
3. In the text box provided, enter the **Site Name** (C above).
4. In the two text boxes provided, enter the optional **Address** information (D above).
5. In the text box provided, enter the **City** where the company site is located (E above).
6. In the text box provided, enter the **State/Province** where the company site is located (F above).
7. In the text box provided, enter the **Postal Code** for the company site (G above).
8. In the text box provided, enter the **Country** for the company site (H above).
9. In the text box provided, enter the **Contract ID** for the company site (J above).
10. To register externally with Sun, check the **Registration externally with Sun** check box (K above).

Next task: To continue, go to [“Specifying Internet Connection Settings”](#) on page 203. To save or clear changes, go to [“Reading the privacy statement and saving changes”](#) on page 203.

▼ Specifying Internet Connection Settings

1. Scroll down to the **Internet Connection Settings** section of the **Registrations** page (A below).

The screenshot shows the 'Internet Connection Settings' form. Callout A points to the section title. Callout B points to the 'Connection Type' section, which has two radio buttons: 'Direct Connection to the Internet' (selected) and 'Use Proxy Server'. Callout C points to the 'Proxy Host Name' text input field. Callout D points to the 'Proxy Port' text input field, which contains the value '8080'. Callout E points to the 'Enable Proxy Authentication' checkbox, which is unchecked. Callout F points to the 'User Name' and 'Password' text input fields.

2. If you connect to the Internet directly, without an intervening proxy server, specify your **Connection Type (B above)** by clicking the **Direct Connection to the Internet** radio button. Then stop here, and go to the **Next** task listed at the end of this procedure.
3. If you connect to the Internet via a proxy server, specify your **Connection Type (B above)** by clicking the **Use Proxy Server** radio button.
4. In the two text boxes provided, enter the **Proxy Host Name (C above)** and **Proxy Port number (D)**.
5. If you must provide credentials before accessing the proxy server, check the **Enable Proxy Authentication** check box (E above) and enter your **User Name** and **Password** in the two text boxes provided (F).

Next task: [“Reading the privacy statement and saving changes” on page 203.](#)

▼ Reading the privacy statement and saving changes

1. Scroll down to the **Privacy Statement** section of the **Registrations** page (A below).

The screenshot shows the 'Privacy Statement' section. Callout A points to the section title. Callout B points to the first paragraph of text. Callout C points to the second paragraph of text. Callout D points to the 'Save' and 'Reset' buttons at the bottom right of the section.

2. Use the controls at the right side of the text area (B above) to scroll through the statement (C).
3. Once you have read the privacy statement and reviewed all entries, press `Save` to save your registration information or `Reset` to discard the information and start over (D above).

Stop here. This completes the registration process.

Troubleshooting

This chapter provides information about troubleshooting the Sun Customer Ready Infinite Archive System software. It contains the following sections:

- [“Using IAS diagnostic tools” on page 205](#)
- [“Error conditions” on page 205.](#)

Using IAS diagnostic tools

Always start by [“Running the IAS Explorer” on page 85](#). The IAS Explorer gathers a wide range of diagnostic and support information that is essential when solving problems. [“Running SAM Explorer” on page 87](#) can also be a useful source of supplemental information in some circumstances.

Error conditions

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. Contact Sun support for further information.
- Recycle the expired archive copies on the VSN by running the recycler.

You can use the Monitoring Console to track VSN usage. One of the available reports lists archive copies with the highest usage.

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.

To address this issue, reboot the server and log in to the server as `root`. Then enter the command `pkill -HUP syseventd` at the system prompt. The Solaris `sysevent` daemon will then restart and find the required `SUNwsamfs` binaries.

Registering for the ASR feature

Auto Service Request (ASR) is a feature of the Common Array Manager Software (CAMS) that manages the disk-storage subsystems integrated into Infinite Archive System appliances. By registering with Sun, you setup the ASR feature to automatically initiate a service request—with supporting telemetry—as soon as CAMS detects a problem.

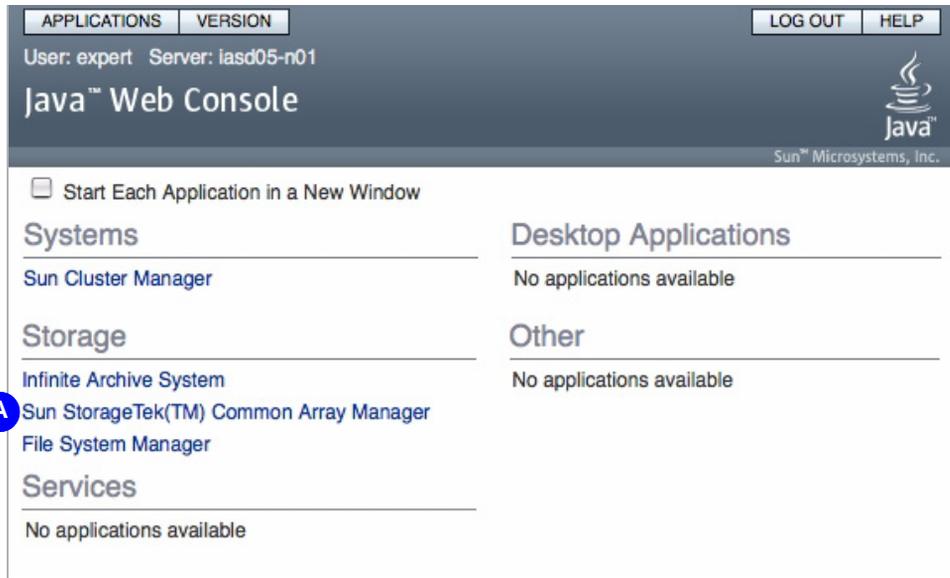
You must accept or decline ASR registration before you can use the CAMS software or the IAS disk-array monitoring features. Normally, Sun Services personnel perform ASR registration during initial configuration of an IAS appliance. But if necessary, you can take care of this step yourself using the procedures in this appendix. Start with [“Starting the ASR registration wizard” on page 207](#). Then continue with one of the following procedures:

- [“Declining ASR support” on page 208](#)
- [“Accepting ASR support” on page 210](#).

▼ Starting the ASR registration wizard

1. **If you have not already done so, open a compatible web browser, such as Mozilla Firefox, enter the IP address of the management interface on the SAM-QFS Meta Data Server (MDS) node, and log in as an administrative user, such as `expert`.**

2. When the Java Web Console interface appears, click on the Sun StorageTek(TM) Common Array Manager hyperlink (A below).

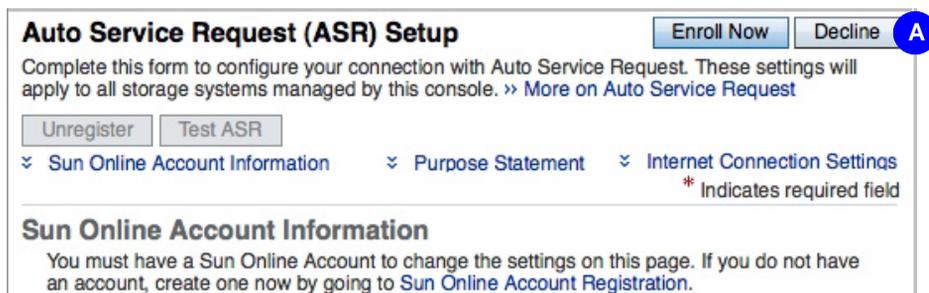


If you have not already completed ASR registration, the Auto Service Request (ASR) registration wizard starts automatically.

Next task: “Declining ASR support” on page 208 or “Accepting ASR support” on page 210.

▼ Declining ASR support

1. When the Auto Service Request (ASR) Setup screen appears, press Decline (A below).



2. If ASR is desired at a later time, the customer can register from the CAMS interface.

CAMS can be accessed via the `Applications` button at the top left corner of the IAS graphical user interface and the `Sun StorageTek(TM) Common Array Manager` hypertext link.

Stop here.

▼ Accepting ASR support

1. If the customer has a Sun Online Account, enter the Sun Online Account Name and Password in the fields provided in the Sun Online Account Information section of the Auto Service Request (ASR) Setup screen (A below).

Auto Service Request (ASR) Setup E

Complete this form to configure your connection with Auto Service Request. These settings will apply to all storage systems managed by this console. » [More on Auto Service Request](#)

☾ Sun Online Account Information ☾ Purpose Statement ☾ Internet Connection Settings

* Indicates required field

Sun Online Account Information

You must have a Sun Online Account to change the settings on this page. If you do not have an account, create one now by going to [Sun Online Account Registration](#) B

Registration Status: Not Registered

A * Sun Online Account Name:

* Password:

[Back to top](#)

Internet Connection Settings

Connection Type: Direct Connection to the Internet

Use Proxy Server C

Proxy Host Name:

Proxy Port:

Enable Proxy Authentication

User Name:

Password:

[Back to top](#)

Purpose Statement D

Sun Microsystems Inc. respects your desire for privacy. Personal information collected from this form will not be shared with organizations external to Sun without your consent, except to process data on Sun's behalf in connection with this transaction.

We will use the personal information collected for communications and management of your Sun Online Account, the services and applications you access using your Sun Online Account, and the products you register with your Sun Online Account.

If you are providing an alternative contact, you represent and warrant to Sun that you have

[Back to top](#) E

2. Otherwise, if you do not have a Sun Online Account, click the Sun Online Account Registration hypertext link in the Sun Online Account Information section of the Auto Service Request (ASR) Setup screen (B above), follow the instructions on screen, and return here.

3. **In the Internet Connections Settings section (C above), click the radio button corresponding to your Internet Connection Type, and enter proxy information (if any) in the text fields provided.**
4. **Read the Purpose Statement (D above).**
5. **Press Enroll Now to complete registration or Decline to abort (E above).**

Stop here.

Glossary

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 mount points 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 (mcf).

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 *disk allocation unit*.

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 there is one Historian for all automated libraries at a site.

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 `/opt/SUNWsamfs/mibs/SUN-SAM-MIB.mib` 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 *disk allocation unit*.

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 *striping*.

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 `/opt/SUNWsamfs/mibs/SUN-SAM-MIB.mib` 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.

K

L

last update

The most recent time that the browser interface was updated with data.

library

See *automated library*.

library catalog

See *catalog*.

low water mark

A specified percentage of the total disk space that is used to control recycling. The releaser continues to release files as long as the file system remains above the configured Low Water Mark.

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 `/opt/SUNWsamfs/mibs/SUN-SAM-MIB.mib` 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 *master configuration file*.

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 that 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 `.inodes` 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.

recovery point

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.

recovery point schedule

The automatic creation of recovery point files for archiving file systems. It is good practice to create a recovery point 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 `/opt/SUNWsamfs/mibs/SUN-SAM-MIB.mib` 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 file 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 `mcF` file as one or more (usually two) `gXXX` 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 *World Wide Name*.