Oracle® Hierarchical Storage Manager and StorageTek QFS Software

`samu` User Interface Guide

Release 6.1.2

E48431-07

February 2018

ORACLE

Oracle Hierarchical Storage Manager and StorageTek QFS Software `samu` User Interface Guide, Release 6.1.2

E48431-07

Contents

3 Operator Commands

4 Service and Support Commands

Preface

This document addresses the needs of system administrators, and network administrators, and service engineers who are tasked with configuring and operating file systems and archiving solutions using Oracle Hierarchical Storage Manager (formerly StorageTek Storage Archive Manager) and Oracle StorageTek QFS Software.

Documentation Accessibility

For information about Oracle's commitment to accessibility, visit the Oracle Accessibility Program website at
http://www.oracle.com/pls/topic/lookup?ctx=acc&id=docacc.

Access to Oracle Support

Oracle customers that have purchased support have access to electronic support through My Oracle Support. For information, visit
http://www.oracle.com/pls/topic/lookup?ctx=acc&id=info or visit
http://www.oracle.com/pls/topic/lookup?ctx=acc&id=trs if you are hearing impaired.

Prerequisites for Using this Document

This document assumes that you are already familiar with the Oracle Solaris operating system, storage, and network administration. Please refer to the Solaris documentation and man pages and to storage hardware documentation for information on relevant tasks, commands, and procedures.

Conventions

The following textual conventions are used in this document:

- *Italic* type represents book titles and emphasis.
- `Monospace` type represents commands and text displayed in a terminal window and the contents of configuration files, shell scripts, and source code files.
- **`Monospace bold`** type represents user inputs and significant changes to commandline output, terminal displays, or file contents. It may also be used to emphasize particularly relevant parts of a file or display.
- ***`Monospace bold oblique`*** type represents variable inputs and outputs in a terminal display or file.
- *`Monospace oblique`* type represents other variables in a terminal display or file.

- **...** (three-dot ellipsis marks) represent file contents or command output that is not relevant to the example and has thus been omitted for brevity or clarity.
- `[ - ]` (brackets surrounding values separated by a hyphen) delimit value ranges.
- `[ ]` (brackets) in command syntax descriptions indicate optional parameters.
- `root@`*`solaris-host`*`:~#` represents aSolaris command shell prompt.
- `[root@`*`linux-host`* `~]#` represents a Linux command shell prompt.

Available Documentation

The *Oracle® Hierarchical Storage Manager and StorageTek QFS Software* `samu` *User Interface Guide* is part of the multivolume Oracle HSM *Customer Documentation Library*, available from `docs.oracle.com`.

For Solaris operating system documentation, see the *Oracle Solaris 11.1 Information Library*, also available from `docs.oracle.com`.

For information on system requirements, new features, and bug fixes, consult the release notes, `README.txt`, in the download ZIP file or on your Oracle HSM file-system server at `/opt/SUNWsamfs/doc/README.txt`.

1

Introducing samu

This document describes the Oracle HSM samu operator utility, a menu-driven, commandline user interface that greatly simplifies configuration and management tasks. With samu, you can monitor Oracle HSM devices and file system activity, control access to devices and file systems. and tune running file systems.

The samu utility is in some respects similar to the UNIX vi editor. You select displays, set display options, navigate within and between displays, enter commands, refresh displays, and quit the utility using similar control key sequences. The last line of each display window displays error messages. Displays refresh automatically unless an error occurs, in which case the display halts until the operator takes further action. When desired, you can take snapshots of display windows for later reference.

The remainder of this introduction describes prerequisites for using samu and some common tasks. Subsequent chapters provide detailed descriptions of the available commands.

Prerequisites

The samu operator utility requires a display terminal that is at least 24 lines long by 80 characters wide.

Like the vi editor, samu is based on the UNIX curses(3CURSES) library. If you have trouble invoking samu, make sure that your terminal type is defined correctly.

Note that what you see in the sample screens below may differ from what you see online. The samples are representative rather than exact representations of any given installation. Actual displays depend on your terminal characteristics and on your file-system configuration. Many samu screens relate to storage and archive management, for example, and will not appear if you are using QFS file systems without archiving.

Common Tasks

The following sections describe how to start and stop samu, interact with the utility, access the help windows, view operator displays, and bring up drives.

- Start samu
- Stop samu
- Display a Screen
- Change a Drive State From down to on
- Get Help

Start samu

You can start the samu operator utility using the utility's defaults or you can specify display and/or command options:

- To start the operator utility using the default display options, enter the samu command at the system command prompt:

  ```
  root@solaris:~# samu
  ```

  The samu interface starts and displays the default initial screen, the help:

  ```
  Help information           page 1/15   samu  version time date

  Display control commands:
  Displays:

      a  Archiver status                x  Migration status
      c  Device configuration           y  Migration vsn list
      d  Daemon trace controls          C  Memory
      f  File systems                   D  Disk volume dictionary
      g  Shared clients                 F  Optical  disk label
      h  Help information               I  Inode
      l  Usage information              J  Preview  shared memory
      m  Mass storage status            K  Kernel statistics
      n  Staging  status                L  Shared memory tables
      o  Optical  disk status           M  Shared memory
      p  Removable media load requests  N  File system parameters
      r  Removable media                P  Active Services
      s  Device status                  R  SAM-Remote
      t  Tape drive status              S  Sector data
      u  Staging  queue                 T  SCSI sense data
      v  Robot catalog                  U  Device table
      w  Pending  stage queue

      more (ctrl-f)                              samu on host
  ```

- To start the operator utility with a specified display, enter the samu command with the -d switch and the name of the command:

  ```
  # samu -d commandname
  ```

  The samu interface starts and shows the specified screen display. For example, to start samu with the configuration screen, you would enter the following:

  ```
  # samu -d c
  ```

  The utility starts with the device configuration screen:

  ```
  Device configuration:              samu               version time date

  ty  eq  state   device_name                             fs  family_set
  sk  100 on      /etc/opt/SUNWsamfs/libraryA_T10K_parms  100 libraryA
  ti  101 on      /dev/rmt/0cbn                           100 libraryA
  ti  102 on      /dev/rmt/1cbn                           100 libraryA
  ti  103 on      /dev/rmt/2cbn                           100 libraryA
  hy  104 on      historian                               104
  ```

- To start the operator utility and change the default display refresh interval, enter the samu command with the -r switch and the length of the interval in seconds:

  ```
  # samu -r seconds
  ```

- To start the operator utility and immediately run a specified command, enter the `samu` command with the `-c` switch and the name of the command:

  ```
  # samu -c commandname
  ```

- To start the operator utility and immediately run a series of commands from a file containing one command per line, enter the `samu` command with the `-f` switch and the name of the command file:

  ```
  # samu -c filename
  ```

Stop samu

To exit `samu`, enter `:q` or `q`.

```
Device configuration:              samu               version time date

ty  eq  state   device_name                              fs  family_set
sk  100 on      /etc/opt/SUNWsamfs/libraryA_T10K_parms   100 libraryA
ti  101 on      /dev/rmt/0cbn                            100 libraryA
ti  102 on      /dev/rmt/1cbn                            100 libraryA
ti  103 on      /dev/rmt/2cbn                            100 libraryA
hy  104 on      historian                                104

q
```

The `samu` operator utility exits and returns you to the command shell.

Display a Screen

To display a particular `samu` screen, proceed as follows.

1. If you do not know the display control command that corresponds to the desired screen, enter `h` to call up the help screen. Lookup the required command in the list displayed.
2. Enter the display control command that corresponds to the desired screen.

   For example, to display the library catalog and list the volume serial numbers (VSNs) for the tapes in the robotic tape library, enter `v`.

   ```
   Device configuration:              samu               version time date

   ty   eq state   device_name                              fs  family_set
   sk  100 on      /etc/opt/SUNWsamfs/libraryA_T10K_parms   100 libraryA
   ti  101 on      /dev/rmt/0cbn                            100 libraryA
   ti  102 on      /dev/rmt/1cbn                            100 libraryA
   ti  103 on      /dev/rmt/2cbn                            100 libraryA
   hy  104 on      historian                                104

   v
   ```

   The utility displays the `Robot VSN catalog`:

   ```
   Robot VSN catalog by slot : eq 900  samu               version time date
                                                        count 32
   slot       access time  count  use flags         ty vsn
      0   2013/05/01 13:36     2  86% -il---b-----  li S00001
      1   2013/05/03 14:14     0  98% -il---b-----  li S00002
      2   2013/05/03 14:17     0  86% -il---b-----  li S00003
      3   2013/05/03 14:20     0  95% -il---b-----  li S00004
      4   none                 0   0% -il-o-b-----  li S00005 ...
   ```

3. If the display is longer than the available screen area, the utility breaks the information into pages and displays the prompt `more` near the bottom of the screen. To navigate between pages, press the Ctrl-F (`^f`) key combination to move forward to the next page, or Ctrl-B (`^b`) to move back to the previous page.

Enter a `samu` Command

You enter commands from the display screens using the procedure below:

1. On any `samu` screen, enter a **`:`** (colon).
2. When the `command:` prompt appears, enter the `samu` command.

```
Device configuration:              samu               version time date

ty   eq state   device_name                           fs   family_set
sk  100 on      /etc/opt/SUNWsamfs/libraryA_T10K_parms 100  libraryA
ti  101 on      /dev/rmt/0cbn                         100  libraryA
ti  102 on      /dev/rmt/1cbn                         100  libraryA
ti  103 down    /dev/rmt/2cbn                         100  libraryA
hy  104 on      historian                             104
:
command:q
```

Run a Shell Command and Return

You run operating system shell commands from `samu` as follows:

1. On any `samu` screen, enter a **`:`** (colon).
2. When the `command:` prompt appears, enter `!` (exclamation point) followed by the shell command.

```
Device configuration:            samu               version time date

ty   eq state   device_name                           fs   family_set
sk  100 on      /etc/opt/SUNWsamfs/libraryA_T10K_parms 100  libraryA
ti  101 on      /dev/rmt/0cbn                         100  libraryA
ti  102 on      /dev/rmt/1cbn                         100  libraryA
ti  103 down    /dev/rmt/2cbn                         100  libraryA
hy  104 on      historian                             104
:
command:!ls
cores          faults         log            reports        sendtrap.log
tmpfiles
Press Return to continue
```

3. When you are ready to return to the `samu` interface, press Return.

Change a Drive State From `down` to `on`

1. Open a `samu` display that shows drive and automated library device states. Enter `c`, `m`, `o`, `r`, `s`, or `t`.
2. In the `eq` field of the screen, note the equipment ordinal number for the `down` device.

```
Device configuration:              samu               version time date
```

```
ty  eq  state   device_name                              fs   family_set
sk  100 on      /etc/opt/SUNWsamfs/libraryA_T10K_parms  100  libraryA
ti  101 on      /dev/rmt/0cbn                           100  libraryA
ti  102 on      /dev/rmt/1cbn                           100  libraryA
ti  103 down    /dev/rmt/2cbn                           100  libraryA
hy  104 on      historian                               104
```

3. Stop all device activity. Enter the command character: and, when the command prompt appears, enter the `off` command with the equipment ordinal number for the `down` device:

```
Device configuration:             samu                version time date

ty   eq state   device_name                              fs   family_set
sk  100 on      /etc/opt/SUNWsamfs/libraryA_T10K_parms  100  libraryA
ti  101 on      /dev/rmt/0cbn                           100  libraryA
ti  102 on      /dev/rmt/1cbn                           100  libraryA
ti  103 down    /dev/rmt/2cbn                           100  libraryA
hy  104 on      historian                               104
:
command:off 103
```

4. Enter the command character: and, when the command prompt appears, enter the `on` command with the equipment ordinal number for the `off` device:

   For example:

```
Device configuration:             samu                version time date

ty   eq state   device_name                              fs   family_set
sk  100 on      /etc/opt/SUNWsamfs/libraryA_T10K_parms  100  libraryA
ti  101 on      /dev/rmt/0cbn                           100  libraryA
ti  102 on      /dev/rmt/1cbn                           100  libraryA
ti  103 off     /dev/rmt/2cbn                           100  libraryA
hy  104 on      historian                               104
:
command:on 103
```

Get Help

You can reach the help pages from any `samu` display by entering `:h` or `h`.

2

Commands Listed by Function

This chapter groups the samu commands by task and function. Locate the task that you need to carry out and then see the alphabetical command listings in Chapter 3, "Operator Commands", and Chapter 4, "Service and Support Commands", for more information:

Controlling Operator Displays

This section lists the samu operator display commands. Note that each command is a lower-case letter of the alphabet (uppercase commands are described in Chapter 4, "Service and Support Commands".

- a (Display Archiver Status Screen)
- c (Display Device Configuration Screen)
- d (Display Daemon Trace Controls Screen)
- f (Display File Systems Screen)
- h (Display the Help Screen)
- l (Display Utilization Information)
- m (Display the Mass Storage Status Screen)
- n (Display the Staging Status Screen)
- o (Display the Optical Disk Status Screen)

- p (Display the Removable Media Load Requests Screen)
- r (Display the Removable Media Status Screen
- refresh (Set the Automatic Refresh Interval for the samu Display)
- s (Display Device Status Screen)
- snap (Save a Snapshot of the Current Screen to a File)
- t (Display Tape Drive Status Screen)
- u (Display Staging Queue Screen)
- v (Display Robot VSN Catalog Screen)
- w (Display the Pending Stage Queue Screen)
- x (Display Current Migration Status)
- y (List Migrating Volumes by VSN)

Tuning I/O Performance

This section starts with a brief summary of the factors that affect file-system performance, the objectives of performance tuning, and the Oracle HSM features that support tuning efforts. It then links to the commands that control the input/output methods that Oracle HSM supports: paged (buffered) I/O, direct I/O, and switched I/O.

Understanding I/O Performance Tuning Objectives and Options

Disk I/O (input/output) involves mechanical processes that are vastly more time-consuming than other file-system operations. So I/O performance tuning focuses on keeping the mechanical work to the absolute minimum necessary for transferring a given amount of data. This means reducing both the number of separate I/Os per data transfer and the number of *seeks* operations required to service each I/O.

So the basic objectives of I/O tuning are as follows:

- Read and write large blocks of data.
- Write blocks in units that align with the sector boundaries of the underlying media, so that the disk controller does not have to read and modify existing data before writing the new data.
- Queue up small I/Os in cache and write larger, combined I/Os to disk.

To achieve these objectives, Oracle HSM file systems support three kinds of I/O, each of which addresses a different set of file-system usage patterns.

Paged (buffered) I/O is the default method because it works well with the usage characteristics of common, general-purpose file-systems. Most file-systems are used by multiple users and applications. Each reads and writes small amounts of data in blocks that often depend on application design rather than on disk sector boundaries. File access tends to be random rather than sequential. When properly configured, paged I/O is ideal for this kind of use. User data is cached in virtual memory pages until the Oracle Solaris kernel writes the data to disk. So users and applications write to and read from cache memory, where random access and small block sizes are less of a problem. Performance-critical, physical reads and writes to and from disk are made in larger, better aligned, more nearly sequential chunks.

Direct I/O addresses limitations of paged I/O that become apparent in file systems are dedicated to particular types of data, such as medical imagery, geological information bases, and real-time surveillance imagery. Generally, users and applications read and write large files sequentially in blocks that align with disk sector boundaries. So paging has no performance advantages. It does, however, impose system overhead and delays the point at which data is securely written from memory to disk—the latter potentially critical in realtime applications. In this situation, direct I/O can significantly improve performance. Data is transferred directly between the local buffer on the application host and the Oracle HSM disk device, with no intermediate delays or additional complication.

Switched I/O is a combination of the other two that works by setting an I/O size threshold for the file system. Paged I/O is used until this threshold is reached. Thereafter, the system automatically switches to direct I/O until the request has been completely addressed. Most small requests are handled before the threshold is reached, so they are read from and written to cache memory and transferred to or from disk asynchronously, in larger, better aligned units. But larger requests are handled directly once the threshold is exceeded, without extra caching overhead.

Paged (Buffered) I/O

The following commands let you tune the paged I/O characteristics of Oracle HSM file systems based on file size and usage.

- `flush_behind` (Set the Flush-Behind Parameter)
- `force_nfs_async` (Force Asynchronous NFS)
- `noforce_nfs_async` (Do Not Force Asynchronous NFS)
- `readahead` (Set The Maximum Read-Ahead Parameter)
- `wr_throttle` (Set the Maximum Size of Outstanding Writes)
- `writebehind` (Set the Write-Behind Parameter)

Direct I/O

The following `samu` commands let you change the default I/O behavior of a specified family set device from paged I/O to direct I/O and vice versa:

- `forcedirectio` (Use Direct I/O by Default)
- `noforcedirectio` (Do Not Use Direct I/O by Default)

The commands below control how sparse files are initialized during direct I/O:

- `dio_szero` (Zero Uninitialized Areas of Sparse Files During Direct I/O)
- `nodio_szero` (Do Not Zero Uninitialized Areas of Sparse Files with Direct I/O)

Switched I/O

As described above, Oracle HSM can be configured to switch from paged I/O to direct I/O when file reads and/or writes exceed specified thresholds. The commands below enable this feature and define the thresholds:

- `dio_rd_consec` (Limit the Number of Consecutive Direct Reads)
- `dio_rd_form_min` (Set the Size of Well-Aligned Direct I/O Reads)
- `dio_rd_ill_min` (Set the Size of Misaligned Direct I/O Reads)

- dio_wr_consec (Limit the Number of Consecutive Direct Writes)
- dio_wr_form_min (Set the Size of Well-Formed Direct I/O Writes)
- dio_wr_ill_min (Set the Size of Misaligned Direct I/O Writes)

Configuring File Systems

- add (Add Equipment to a Mounted File System)
- atime (Configure Access Time Updates for a File System)
- def_retention (Set Default WORM Retention Time)
- mm_stripe (Set the Metadata Stripe Width)
- noatime (Disable Access Time Updates for a File System)
- noqwrite (Disable Simultaneous Reads/Writes from Different Threads)
- nosuid (Mount a File System with setuid Execution Disabled)
- nosw_raid (Turn Off Software RAID Mode)
- notrace (Disable File System Tracing)
- qwrite (Enable Simultaneous Reads/Writes to a File from Different Threads)
- remove (Move Data and Remove Equipment from a File System)
- stripe (Set the Stripe Width for a File System)
- suid (Mount a File System with setuid Execution Enabled)
- sw_raid (Turn On Software RAID Mode)
- sync_meta (Synchronize File System Metadata)
- trace (Enable File System Tracing)

Configuring Archival File Systems

The commands below control the archiving and staging behavior of Oracle HSM file systems:

- maxpartial (Set the Maximum Partial Release Size for the File System)
- nohwm_archive (Do Not Start Archiving at the High Water Mark)
- partial (Set the Default Partial Release Size for the File System)
- partial_stage (Set Partial Stage-Ahead Point for the File System)
- release (Release Archived Data Files and Disable Allocation on a Device)
- stage_flush_behind (Set the Stage Flush Behind Size)
- stage_n_window (Size the Buffer for Reading Directly from Archival Media)
- thresh (Set High and Low Watermarks)

Configuring Shared File Systems

The following commands control how Oracle HSM coordinates interactions between hosts that share a file system.

- aplease (Set Append Lease Expiration Time for a Shared File System)

- `lease_timeo` (Set the Unused Lease Timeout for a Shared File System)
- `maxallocsz` (Set the Maximum Allocation Size for a Shared File System)
- `mh_write` (Turn on Multi-Host Read/Write)
- `minallocsz` (Set the Minimum Allocation Size for a Shared File System)
- `migconfig` (Configure Migration to New Media)
- `nomh_write` (Turn Off Multi-Host Read/Write)
- `rdlease` (Set the Read Lease Time for a Shared File System)
- `wrlease` (Set the Write Lease Time for a Shared File System)

Configuring Single-Writer/Multi-Reader File Systems

The commands below control how and when a read-only host obtains metadata updates from the metadata server.

- `invalid` (Set the Invalidate-Cache Delay for a Multi-Reader File System)
- `norefresh_at_eof` (Do Not Refresh at EOF in a Multi-Reader File System)
- `refresh_at_eof` (Refresh at EOF in a Multi-Reader File System)

Configuring File Systems for Oracle Real Application Clusters (RAC)

The following commands are specific to implementations of Oracle Real Application Clusters (RAC) that store data on Oracle HSM file systems using asynchronous I/O and Solaris Volume Manager mirrored volumes.

- `abr` (Enable Application Based Recovery)
- `dmr` (Enable Directed Mirror Reads of Software Mirrors)
- `noabr` (Disable Application Based Recovery)
- `nodmr` (Disable Directed Mirror Reads of Software Mirrors)

Managing Storage Devices

- `alloc` (Enable Storage Allocation on a Device)
- `clear` (Clear a Load Request for a Removable Storage Volume)
- `devlog` (Set Device Logging Options)
- `idle` (Idle Device)
- `noalloc` (Disable Storage Allocation on a Device)
- `off` (Turn Off A Device)
- `on` (Turn On A Device)
- `readonly | ro` (Make Device Read Only)
- `unavail` (Make a Device Unavailable to Oracle HSM)

Managing Robotic Storage Libraries

- `audit` (Audit Library or Library Slot)

- export (Move Removable Media to the Robotic Library's Mailbox)
- import (Move Removable Media from the Mailbox into the Robotic Library)
- load (Load a Cartridge into a Drive)
- priority (Set Priority in Load-Request Preview Queue)
- unload (Unload a Drive)

Archiving

- aridle (Stop Archiving Gracefully)
- arrerun (Restart the Archiver Non-Disruptively)
- arrestart (Restart the Archiver Disruptively)
- arrmarchreq (Remove Archive Requests)
- arrun (Start Archiving)
- arscan (Scan a File System for Unarchived Files)
- arstop (Stop All Archiving Immediately)
- artrace (Perform Archiver Tracing)

Staging

- stclear (Clear a Stage Request)
- stidle (Stop Staging as Soon as Possible)
- strun (Start Staging)

Migrating Volumes to New Media

- migconfig (Configure Migration to New Media)
- migidle (Configure Migration to New Media)
- migstart (Start the Media Migration Service)
- migstop (Stop the Media Migration Service)

Working with Service and Support Representatives

The samu interface includes a number of commands that are designed to help you assist Oracle technical support and service representatives. These handle the following tasks:

- Controlling Diagnostic Displays
- Gathering Diagnostic Information

Controlling Diagnostic Displays

The samu service and support display commands are designated by uppercase letters: C, D, F, I, J, K, L, M, N, P, R, S, T, and U. They provide debugging and internal state information that is not generally useful without the assistance of a member of the Oracle technical support staff.

- `:C` (Display the Contents of Shared Memory at a Specified Address)
- `D` (Display Disk Volume Dictionary Screen)
- `F` (Display Optical Disk Label Screen)
- `I` (Display Inode Information)
- `J` (Display the Shared Memory Segment for the Load-Request Preview Queue)
- `K` (Display Kernel Statistics)
- `L` (Display Shared Memory Tables and System Defaults)
- `M` (Display Raw Shared Memory)
- `N` (Display File System Parameters)
- `P` (Display Active Services)
- `R` (Display SAM-Remote Configuration Information)
- `S` (Display Raw Sector Data)
- `T` (Display SCSI Sense Data for a Current or Specified Device)
- `U` (Display Device Table in Human-Readable Form)

Gathering Diagnostic Information

- `diskvols` (Set/Clear Flags on Disk Volumes Used for Archiving)
- `dtrace` (Enable/Configure Tracing for Oracle HSM Daemons)
- `open` (Open a Disk Storage Device for Reading)
- `read` (Read Sectors On a Disk Storage Device)

3

Operator Commands

This section lists detailed summaries of `samu` operator commands in alphabetical order.

a (Display Archiver Status Screen)

The `a` command opens the `Archiver status` screen and displays current archiver activity.

When issued without a parameter, the `a` command displays archiver status for all file systems. When the command is issued with the name of a file system, the command displays status for the specified file system only.

Command Synopsis

```
a
:a filesystem
```

Controls

Ctrl-B (^b)
Move back one file system.

Ctrl-F (^f)
Move forward one file system.

Ctrl-D (^d)
Move down one archive copy.

Ctrl-U (^u)
Move up one archive copy.

Ctrl-J (^j)
Toggle between binary and decimal units of size.

Examples

The following example shows a typical archiver status screen. The archiver is idle:

```
Archiver Status                   samu              version time date

sam-archiverd: Idle

sam-arfind:  samqfs1 mounted at /export/samqfs1
```

```
Files waiting to start 32         schedule 0         archiving 0
Monitoring file system activity.
```

The next example shows the same file system, but the archiver is actively archiving:

```
Archiver Status                   samu               version time date

sam-archiverd: Archiving

sam-arfind:  samqfs1 mounted at /export/samqfs1
Files waiting to start 149        schedule 0         archiving 0
Monitoring file system activity.

sam-arcopy:  samqfs1.archset1.1.83 li.VOL119
Copying file /export/samqfs1/2014/data/A00876.dat
```

Field Descriptions

The screen lists status information for all Oracle HSM file systems by file-system name and mount point.The fields in the detail display are explained below.

sam-archiverd

The `sam-archiverd` field displays the current status of the Oracle HSM archiver daemon, the software component that copies modified files to backup media:

- `Idle` means that Oracle HSM is waiting for archive-ready files
- `Archiving` means that Oracle HSM is copying files to archival storage.
- `Waiting for resources` means that Oracle HSM is ready to copy files to archival storage, but is waiting for a drive to become available or for a volume to mount.
- Various other messages indicate errors.

sam-arfind

The `sam-arfind` field displays the most recent results reported by the software component that monitors a file system and locates candidates for archiving. The results include:

- the name that you assign to the Oracle HSM file system (`samqfs1` in the example)
- the mount-point directory for the Oracle HSM file system (`/export/samqfs1` in the example)
- the number of archive-ready `Files waiting to start` archiving (`32` in the example)
- the number of archive-ready files currently in the archiving `schedule` (`0` in the example)
- the number of archive-ready files that are currently `archiving` (`0` in the example).

sam-arcopy

The `sam-arcopy` field displays the current activity reported by the software component that copies archive-ready files to archival media. It reports the following:

- the name of the file system (`samqfs1` in the example)

  The files in an archive request file are members of the same *archive set* and thus share the same archiving characteristics and requirements. The filename consists of the name of the *copy directive* (the archiving rule that sets the criteria for copying and retaining a file on specified media) and a sequence number set by the archiver.

- the name of the archive set (archset1 in the example)
- the copy number of the current copy (1 in the example)
- a sequence number (83 in the example)
- the volume serial number of the media volume that is being used to store the current copy (VOL119 in the example).
- the path and file name of the file that is currently being copied (/export/samqfs1/2014/data/A00876.dat in the example).

abr (Enable Application Based Recovery)

When issued with the Oracle HSM equipment ordinal number of a disk family-set device, the abr command lets Oracle Real Application Clusters (RAC) use the Application Based Recovery feature of Solaris Volume Manager to handle recovery of databases stored in Oracle HSM file systems that use asynchronous I/O with Solaris Volume Manager mirrored volumes.

By default, ABR is enabled.

Command Synopsis

```
:abr family-set-equipment-number
```

Examples

The following example enables ABR on family-set device 20:

```
File systems                          samu                version time date

ty  eq state   device_name   status        high   low   mountpoint   server
ma  20 on      sam2          -----2---r-   80%    70%   /sam2
 mm 21 on      /dev/dsk/c0t600A0B800026B6F000001F0F4CCFF833d0s0
 mr 22 on      /dev/dsk/c0t600A0B800026E584000021864C4943AAd0s0
 mr 23 on      /dev/dsk/c0t600A0B800026E5840000218B4CBD9C4Ed0s0
ma  10 on      sam1          m----2---r-   80%    70%   /sam1
 mm 11 on      /dev/dsk/c0t600A0B8000118ACA000009594E1565E8d0s0
 mr 12 on      /dev/dsk/c0t600A0B8000114D5400000CA64E156793d0s0
 mr 13 on      /dev/dsk/c0t600A0B8000114D5400000CAA4E156897d0s0
 mr 14 on      /dev/dsk/c0t600A0B8000118ACA0000095B4E1568C8d0s0
 mr 15 on      /dev/dsk/c0t600A0B8000118ACA0000095D4E156948d0s0
:
command:abr 20
```

add (Add Equipment to a Mounted File System)

When issued with an Oracle HSM equipment ordinal number, the add command updates the corresponding file system so that it can use disk devices that have been added to the family set definition in the mcf (Oracle HSM Master Configuration File) and set up with the samd config command.

See "remove (Move Data and Remove Equipment from a File System)" on page 3-80.

Command Synopsis

```
:add equipment-number
```

Examples

For example, assume that we need to add an additional disk slice to file system family set qfs1. We start by adding a line to the /etc/opt/SUNWsamfs/mcf file that identifies the slice, /dev/dsk/c5t8d0s3, as Oracle HSM equipment ordinal number 13 (the selected equipment number must be unique within the Oracle HSM system):

```
root@solaris:~# vi /etc/opt/SUNWsamfs/mcf
...
# Section 1: File Systems
#
# Equipment          Equipment  Equipment Family          Device  Additional
# Identifier         Ordinal    Type      Set             State   Parameters
# -----------------  ---------  --------- --------------  -----   ----------
qfs1                 10         ma        qfs1            on
/dev/dsk/c5t8d0s0    11         mm        qfs1            on
/dev/dsk/c5t8d0s1    12         md        qfs1            on
/dev/dsk/c5t8d0s3    13         md        qfs1            on
:wq
root@solaris:~#
```

Then we reconfigure the running Oracle HSM daemon to incorporate the changes in the file:

```
root@solaris:~# samd config
```

At this point, the File systems screen shows that the equipment is off and thus not yet available. So we enter the add command with the user-assigned Oracle HSM equipment number of the new slice:

```
File systems                          samu                version time date

ty  eq state   device_name   status        high   low   mountpoint   server
ma  10 on      qfs1          -----2---r-   80%    70%
 mm 11 on      /dev/dsk/c5t8d0s0
 md 12 on      /dev/dsk/c5t8d0s1
 md 13 off     /dev/dsk/c5t8d0s3
:
command:add 13
```

The device state changes to on:

```
File systems                          samu                version time date

ty  eq state   device_name   status        high   low   mountpoint   server
ma  10 on      qfs1          -----2---r-   80%    70%
 mm 11 on      /dev/dsk/c5t8d0s0
 md 12 on      /dev/dsk/c5t8d0s1
 md 13 on      /dev/dsk/c5t8d0s3
```

alloc (Enable Storage Allocation on a Device)

When issued with an Oracle HSM equipment ordinal number, the alloc command starts storage allocation on the specified data device. See "noalloc (Disable Storage Allocation on a Device)" on page 3-58.

Command Synopsis

```
:alloc  equipment-number
```

Examples

The following example starts storage allocation on disk device 14:

```
File systems                          samu                version time date

ty  eq state   device_name   status        high   low   mountpoint   server
ma  20 on      sam2          -----2---r-   80%    70%
 mm 21 on      /dev/dsk/c0t600A0B800026B6F000001F0F4CCFF833d0s0
 mr 22 on      /dev/dsk/c0t600A0B800026E584000021864C4943AAd0s0
 mr 23 on      /dev/dsk/c0t600A0B800026E5840000218B4CBD9C4Ed0s0
ma  10 on      sam1          m----2---r-   80%    70%   /sam1
 mm 11 on      /dev/dsk/c0t600A0B8000118ACA000009594E1565E8d0s0
 mr 12 on      /dev/dsk/c0t600A0B8000114D5400000CA64E156793d0s0
 mr 13 on      /dev/dsk/c0t600A0B8000114D5400000CAA4E156897d0s0
 mr 14 noalloc /dev/dsk/c0t600A0B8000118ACA0000095B4E1568C8d0s0
 mr 15 on      /dev/dsk/c0t600A0B8000118ACA0000095D4E156948d0s0
:
command:alloc 14
```

The device state changes to on:

```
File systems                          samu                version time date

ty  eq state   device_name   status        high   low   mountpoint   server
ma  20 on      sam2          -----2---r-   80%    70%   /sam2
 mm 21 on      /dev/dsk/c0t600A0B800026B6F000001F0F4CCFF833d0s0
 mr 22 on      /dev/dsk/c0t600A0B800026E584000021864C4943AAd0s0
 mr 23 on      /dev/dsk/c0t600A0B800026E5840000218B4CBD9C4Ed0s0
ma  10 on      sam1          m----2---r-   80%    70%   /sam1
 mm 11 on      /dev/dsk/c0t600A0B8000118ACA000009594E1565E8d0s0
 mr 12 on      /dev/dsk/c0t600A0B8000114D5400000CA64E156793d0s0
 mr 13 on      /dev/dsk/c0t600A0B8000114D5400000CAA4E156897d0s0
 mr 14 on      /dev/dsk/c0t600A0B8000118ACA0000095B4E1568C8d0s0
 mr 15 on      /dev/dsk/c0t600A0B8000118ACA0000095D4E156948d0s0
```

aplease (Set Append Lease Expiration Time for a Shared File System)

Before appending data to a shared file, a host must request and obtain an *append lease* that grants it exclusive access to the file. The aplease command defines the number of seconds during which this lease remains valid on the specified device. If the specified time runs out before the append operation is complete, the host must reacquire the lease before proceeding.

The specified number of seconds must fall in the range [15-600]. The default is 30 seconds.

For related information, see "rdlease (Set the Read Lease Time for a Shared File System)" on page 3-75, "wrlease (Set the Write Lease Time for a Shared File System)" on page 3-102, and "lease_timeo (Set the Unused Lease Timeout for a Shared File System)" on page 3-47.

Command Synopsis

```
:aplease family-set-equipment-number interval_in_seconds
```

Examples

The following example sets the append lease time to 2 minutes (120 seconds) on family set device 800:

```
File systems                    samu                    version time date

ty   eq   state device_name           status      high low mountpoint server
ma   800  on    shareqfs1             m----2c--r-  80% 70% /shareqfs1
 mm   801  on    /dev/dsk/c6t50020F2300004655d0s0
 mr   802  on    /dev/dsk/c6t50020F2300004655d0s1
ma    810  on   shareqfs2             m----2c--r-  80% 70% /shareqfs2
 mm   811  on    /dev/dsk/c6t50020F2300004655d0s6
 mr   812  on    /dev/dsk/c6t50020F2300004655d0s7
:
command:aplease 800 120
```

aridle (Stop Archiving Gracefully)

The aridle command stops archiving at the next logical point (such as the end of the current archive tar file). The command is typically used prior to disruptive maintenance, such as unmounting, growing, or shrinking file systems. Archiving remains inactive until the arrun command is entered.

When entered with no parameters, aridle stops all archiver activity.

When entered with the dk parameter, aridle stops all archiving to disk media.

When entered with the rm parameter, aridle stops all archiving to removable media, such as tape.

When entered with the fs. parameter and a file-system name, aridle stops all archiver activity on the specified file system.

Command Synopsis

```
:aridle
:aridle dk
:aridle rm
:aridle fs.file-system-identifier
```

Examples

In the first example, aridle stops all archiving:

```
Archiver Status                  samu              version time date

sam-archiverd: Archiving

sam-arfind: sam1 mounted at /sam1
Files waiting to start 0        schedule 482          archiving 0
Monitoring file system activity.

sam-arcopy: sam1.copy-directive.sequence-number vol-ser
```

```
Copying file path/filename
:
command:aridle
```

After the aridle command is entered, the archiver daemon is Idle until the arrun command is entered:

```
Archiver Status                  samu               version time date

sam-archiverd: Idle

sam-arfind:  sam1 mounted at /sam1
Files waiting to start 0        schedule 482          archiving 0
Waiting for :arrun
```

In the second example. aridle stops archiving on the sam1 file system only:

```
Archiver Status                   samu               version time date

sam-archiverd: Archiving

sam-arfind: sam1 mounted at /sam1
Files waiting to start 0        schedule 482          archiving 0
Monitoring file system activity.

sam-arcopy: sam1.copy-directive.sequence-number vol-ser
Copying file path/filename
:
command:aridle fs.sam1
```

After the aridle command is entered, the archiver daemon is Waiting... until the arrun command is entered with the file-system directive fs.sam1:

```
Archiver Status                   samu               version time date

sam-archiverd: Waiting for :arrun fs.sam1

sam-arfind:  sam1 mounted at /sam1
Files waiting to start 0        schedule 482          archiving 0
Waiting for :arrun fs.sam1
```

In the third example. aridle stops archiving to disk:

```
Archiver Status                   samu               version time date

sam-archiverd: Archiving

sam-arfind: sam1 mounted at /sam1
Files waiting to start 0        schedule 482          archiving 0
Monitoring file system activity.

sam-arcopy: sam1.copy-directive.sequence-number vol-ser
Copying file path/filename
:
command:aridle dk
```

After the aridle command is entered, the archiver daemon is now Waiting... until the arrun command is entered with the dk (disk) directive:

```
Archiver Status                   samu               version time date

sam-archiverd:  Waiting for :arrun dk
```

```
sam-arfind:  sam1 mounted at /sam1
Files waiting to start 0        schedule 482          archiving 0
Waiting for :arrun dk
```

arrerun (Restart the Archiver Non-Disruptively)

The arrerun command restarts archiving without disrupting ongoing operations. The archiver daemons are restarted, and all work in progress is recovered. Compare arrerun to arrestart (Restart the Archiver Disruptively).

Command Synopsis

```
:arrerun
```

Examples

In the example, arrerun restarts archiving and recovers work in progress:

```
Archiver Status                   samu               version time date

sam-archiverd: Archiving

sam-arfind: sam1 mounted at /sam1
Files waiting to start 0        schedule 482          archiving 0
Monitoring file system activity.

sam-arcopy: sam1.copy-directive.sequence-number vol-ser
Copying file path/filename
:
command:arrerun
```

arrestart (Restart the Archiver Disruptively)

The arrestart command stops and restarts the archiver regardless of the state of the archiver. Compare the command with "arrerun (Restart the Archiver Non-Disruptively)".

Use arrestart with caution, because it stops operations immediately, and any copy operations that fail to complete will need to be repeated. This waste space on archival media.

Command Synopsis

```
:arrestart
```

Examples

```
Archiver Status                   samu               version time date

sam-archiverd: Archiving

sam-arfind: sam1 mounted at /sam1
Files waiting to start 0        schedule 482          archiving 0
Monitoring file system activity.
```

```
sam-arcopy: sam1.copy-directive.sequence-number vol-ser
Copying file path/filename
:
command:arrestart
```

arrmarchreq (Remove Archive Requests)

When issued with a file-system name and either the name of an archive-request file or a wildcard, the arrmarchreq command deletes the specified archive requests.

An archive-request filename consists of the name of the *copy directive* (the archiving rule that sets the criteria for copying and retaining a file on specified media) and a sequence number set by the archiver.

Command Synopsis

```
:arrmarchreq file-system-name.archive-request-file-name
:arrmarchreq file-system-name.*
```

Examples

In the first example, arrmarchreq deletes archive request file data0.1.83 from the sam1 file system:

```
Archiver Status                   samu               version time date

sam-archiverd: Idle

sam-arfind: sam1 mounted at /sam1
Files waiting to start 0        schedule 0            archiving 0
Monitoring file system activity.
:
command:arrmarchreq sam1.data0.1.83
```

In the second example, arrmarchreq deletes all archive request files from the sam1 file system:

```
Archiver Status                   samu               version time date

sam-archiverd: Idle

sam-arfind: sam2 mounted at /sam2
Files waiting to start 0        schedule 0            archiving 0
Monitoring file system activity.
:
command:arrmarchreq sam2.*
```

arrun (Start Archiving)

The arrun command starts archiving.

When entered with no parameters, arrun starts all archiver activities.

When entered with the dk parameter, arrun starts archiving to disk media.

When entered with the rm parameter, arrun starts archiving to removable media.

When entered with the fs. parameter and a file-system name, arrun starts archiver activity on the specified file system.

Command Synopsis

```
:arrun
:arrun dk
:arrun rm
:arrun fs.file-system-identifier
```

Examples

In the first example, arrun starts all archiving:

```
Archiver Status                   samu               version time date

sam-archiverd: Idle

sam-arfind:  sam1 mounted at /sam1
Files waiting to start 0        schedule 482          archiving 0

Waiting for :arrun
:
command:arrun
```

After the arrun command is entered, the archiver daemon resumes Archiving:

```
Archiver Status                   samu               version time date

sam-archiverd: Archiving

sam-arfind: sam1 mounted at /sam1
Files waiting to start 0        schedule 482          archiving 0
Monitoring file system activity.

sam-arcopy: sam1.copy-directive.sequence-number vol-ser
Copying file path/filename
```

In the second example. arrun starts archiving on the sam1 file system:

```
Archiver Status                   samu               version time date

sam-archiverd: Waiting for :arrun fs.sam1

sam-arfind:  sam1 mounted at /sam1
Files waiting to start 0        schedule 482          archiving 0
Waiting for :arrun fs.sam1
:
command:arrun fs.sam1
```

After the arrun command is entered, the archiver daemon resumes archiving file system sam1:

```
Archiver Status                   samu               version time date

sam-archiverd: Archiving

sam-arfind: sam1 mounted at /sam1
Files waiting to start 0        schedule 482          archiving 0
Monitoring file system activity.
```

```
sam-arcopy: sam1.copy-directive.sequence-number vol-ser
Copying file path/filename
```

In the third example. arrun starts archiving to disk media:

```
Archiver Status                   samu              version time date

sam-archiverd:  Waiting for :arrun dk

sam-arfind:  sam1 mounted at /sam1
Files waiting to start 0        schedule 482          archiving 0
:
command:arrun dk
```

After the arrun command is entered, the resumes archiving to disk:

```
Archiver Status                   samu              version time date

sam-archiverd: Archiving

sam-arfind: sam1 mounted at /sam1
Files waiting to start 0        schedule 482          archiving 0
Monitoring file system activity.

sam-arcopy: sam1.copy-directive.sequence-number vol-ser
Copying file path/filename
```

arscan (Scan a File System for Unarchived Files)

When issued with the name of an Oracle HSM file system, the arscan command scans the specified file system for unarchived files.

When entered with no parameters other than the file-system name, arscan recursively scans all directories in the specified file system.

When entered with the file-system name and a directory, arscan recursively scans all directories in the specified file system starting at the specified directory.

When entered with the file-system name and the .inodes key word, arscan scans the .inodes file. If most files in a file system are already archived, this is faster than scanning directories.

Entering an integer representing a number of seconds with any of the other parameters delays the scan for the specified number of seconds.

Command Synopsis

```
:arscan file-system-name
:arscan file-system-name delay-in-seconds
:arscan file-system-name.starting-directory
:arscan file-system-name.starting-directory delay-in-seconds
:arscan file-system-name..inodes
:arscan file-system-name..inodes delay-in-seconds
```

Examples

In the example, arscan:

```
Archiver Status                   samu              version time date

sam-archiverd: Idle
```

```
sam-arfind:  sam1 mounted at /sam1
Files waiting to start 0        schedule 482          archiving 0
:
command:arscan
```

arstop (Stop All Archiving Immediately)

The arstop command stops archiving immediately.

When entered with no parameters, arstop stops all archiver activities.

When entered with the dk parameter, arstop stops archiving to disk media.

When entered with the rm parameter, arstop stops archiving to removable media.

When entered with the fs. parameter and a file-system name, arstop stops archiver activity on the specified file system.

Command Synopsis

```
:arstop
:arstop dk
:arstop rm
:arstop fs.file-system-identifier
```

Examples

In the first example, arstop stops all archiving:

```
Archiver Status                   samu               version time date

sam-archiverd: Archiving

sam-arfind: sam1 mounted at /sam1
Files waiting to start 0        schedule 482          archiving 0
Monitoring file system activity.
sam-arcopy: sam1.copy-directive.sequence-number vol-ser
Copying file path/filename
:
command:arstop
```

In the second example, the arstop command stops archiving on file system sam1:

```
Archiver Status                   samu               version time date

sam-archiverd: Archiving

sam-arfind: sam1 mounted at /sam1
Files waiting to start 0        schedule 482          archiving 0
Monitoring file system activity.

sam-arcopy: sam1.copy-directive.sequence-number vol-ser
Copying file path/filename
:
command:arstop fs.sam1
```

artrace (Perform Archiver Tracing)

The artrace command starts archive tracing.

Tracing writes a detailed log of archiver activity to the sam-archiverd file in the /var/opt/SUNWsamfs/trace/ subdirectory.

When entered with no parameters, artrace traces archiver activities on all file systems.

When entered with the fs. parameter and a file-system name, artrace traces archiver activity on the specified file system.

Command Synopsis

```
:artrace
:artrace fs.file-system-identifier
```

Examples

In the first example, artrace starts archiver tracing for all file systems:

```
Archiver Status                   samu               version time date

sam-archiverd: Archiving

sam-arfind: sam1 mounted at /sam1
Files waiting to start 0        schedule 482          archiving 0
Monitoring file system activity.
sam-arcopy: sam1.copy-directive.sequence-number vol-ser
Copying file path/filename
:
command:artrace
```

In the second example, the artrace command starts archiver tracing for file system sam1:

```
Archiver Status                   samu               version time date

sam-archiverd: Archiving

sam-arfind: sam1 mounted at /sam1
Files waiting to start 0        schedule 482          archiving 0
Monitoring file system activity.

sam-arcopy: sam1.copy-directive.sequence-number vol-ser
Copying file path/filename
:
command:artrace fs.sam1
```

atime (Configure Access Time Updates for a File System)

When issued with the Oracle HSM equipment ordinal number of a disk family-set device -1, 0, or 1, the atime command controls how and when the access time attribute is updated for files in the corresponding file system.

Setting a 0 (zero) value enables caching and delayed writes up to one minute for access time updates. Access time updates are only written to disk immediately when utilization of an Oracle HSM file system is above the low water mark, when the access

time change coincides with updates to the creation time (ctime) or modification time (mtime), or when the file system is unmounted. This setting reduces disk I/O and thus improves performance.

Setting a 1 value immediately updates the access time on disk whenever a file is accessed. This increases disk I/O and reduces performance accordingly.

Setting a -1 value or executing the noatime command disables access time updates. If access times are unimportant, this can significantly reduce I/O and thus improve performance. However, the POSIX standard requires that access times be marked on files. Never use this setting with an archiving, Oracle HSM file system.

The default is value is 0 (caching and delayed writes enabled).

For more information, see the mount_samfs and stat man pages for further information.

Command Synopsis

```
:atime family-set-equipment-number -1|0|1
```

Examples

The following example disables caching and forces immediate access time updates to disk on family set device 100:

```
File systems               samu                   version time date

ty    eq  state device_name           status       high low   mountpoint server
ma    100  on   qfs1                  m----2----d   90% 70%   /qfs1
 mm    101  on   /dev/dsk/c5t10d0s0
 md    102  on   /dev/dsk/c5t10d0s1
 md    103  on   /dev/dsk/c5t11d0s0
 md    104  on   /dev/dsk/c5t11d0s1
:
command:atime 100 1
```

audit (Audit Library or Library Slot)

When issued with the Oracle HSM equipment ordinal number of a robotic library and, optionally, a slot identifier, the audit command causes the specified automated library to mount volumes, read the volume serial number (VSN) of each volume, and rebuild the library catalog. If specified, the slot must be in use and physically occupied by media.

If the specified slot contains a tape cartridge, the -e parameter tells the audit to skip to the end of data (EOD) and determine the space available. Note, however, that this process is not interruptible and may take hours to complete.

If the specified slot contains a two-sided optical cartridge, identify the desired side as 1 or 2. If a side is not specified, both sides are audited.

For more information, see the auditslot(1M), mount_samfs and stat man pages for further information.

Command Synopsis

```
:audit equipment-number
```

```
:audit equipment-number:slot
:audit -e equipment-number:slot
:audit equipment-number:slot:side
:audit -e equipment-number:slot:side
```

Examples

The following example audits slot 17 in the robotic library with Oracle HSM equipment ordinal number 900:

```
Robot VSN catalog by slot : eq 900  samu       version time date
                                                        count 32
slot     access     time count  use  flags         ty vsn
   0     2013/07/02 12:53    3  96%  -il-o-b----f  li VOL001
   1     2013/07/02 12:53    2  98%  -il-o-b----f  li VOL002
   2     2013/07/02 12:52    1  86%  -il-o-b----f  li VOL003
   3     2013/07/02 12:52    1  95%  -il-o-b----f  li VOL004
   4     2013/07/02 12:51    1  79%  -il-o-b-----  li VOL005
   5     2013/07/02 12:45    0  14%  -ilEo-b----f  li VOL006 MEDIA ERROR
   6     2013/07/02 12:46    0   0%  -il-o-b-----  li VOL007
   7     2013/07/02 12:46    0   0%  -il-o-b-----  li VOL008
   8     2013/07/02 12:47    0   0%  -il-o-b-----  li VOL009
   9     2013/07/02 12:47    0   0%  -il-o-b-----  li VOL010
  10     2013/07/02 12:48    0   0%  -il-o-b-----  li VOL011
  11     2013/07/02 12:48    0   0%  -il-o-b-----  li VOL012
  12     2013/07/02 12:48    0   0%  -il-o-b-----  li VOL013
  13     2013/07/02 12:49    0   0%  -il-o-b-----  li VOL014
  14     2013/07/02 12:49    0   0%  -il-o-b-----  li VOL015
  15     2013/07/02 12:50    0   0%  -il-o-b-----  li VOL016
  16     2013/07/02 12:50    0   0%  -il-o-b-----  li VOL017
  17     2013/07/02 12:51    0   0%  -il-o-b-----  li VOL018
  18     2013/07/02 12:51    0   0%  -il-o-b-----  li VOL019
  19     none               50   0%  -il-oCb-----  li CLN020
:
command: audit 900:17
```

c (Display Device Configuration Screen)

The c command opens the Device configuration screen, which lists all device names and equipment ordinal numbers.

Command Synopsis

```
c
:c
```

Controls

Ctrl-B (^b)
Move back one page.

Ctrl-F (^f)
Move forward one page.

Ctrl-D (^d)
Move down the page.

Ctrl-U (^u)
Move up the page.

Example

The following example shows the device configuration screen:

```
Device configuration:            samu         version time date

ty   eq state   device_name                           fs   family_set
sk  100 on      /etc/opt/SUNWsamfs/SL8500_T10K_parms  100  SL8500
ti  101 down    /dev/rmt/0cbn                         100  SL8500
ti  102 down    /dev/rmt/1cbn                         100  SL8500
ti  103 down    /dev/rmt/2cbn                         100  SL8500
hy  104 on      historian                             104
```

Field Descriptions

ty
The ty field displays the Oracle HSM device type. Device-type identifiers are listed in Appendix A and on the mcf man page.

eq
The eq field displays the user-assigned Oracle HSM equipment ordinal number of the device.

state
The state field displays the current operating state of the device, which may include any of the following:

- on means that the device is available.
- ro means that the device is available read-only.
- off means that the device is not available.
- down means that the device is available for maintenance only.
- idle means that the device is completing operations that are already in progress but is not accepting new connections.
- noalloc means that no more storage can be allocated on this device. An administrator has run the noalloc command (see "noalloc (Disable Storage Allocation on a Device)" on page 3-58). The noalloc state persists until an administrator issues an alloc command (see "alloc (Enable Storage Allocation on a Device)" on page 3-4).

device_name
The device_name field displays the path and file name of the device.

In this instance, the family set equipment type, sk, indicates an Oracle StorageTek ACSLS interface to a network-attached library. So the device name is the path to the parameters file for the library, SL8500_T10K_parms. For detailed information on network-attached libraries and parameter files, see the mcf man page.

fs
The fs field displays the user-assigned Oracle HSM equipment ordinal number for the family-set that contains a device.

family_set
The family_set field displays the user-assigned name of the Oracle HSM family set that contains the device. Family sets represent the equipment associated with a Oracle HSM component, such as the group of disk devices associated with an Oracle HSM file system or the tape library associated with archiving.

clear (Clear a Load Request for a Removable Storage Volume)

When issued with a volume serial number (VSN) and, optionally, an index value, the clear command cancels pending load requests for the specified VSN, clears any reference to the VSN from the Removable media load requests screen, and aborts any processes that are waiting for the volume to mount.

Optionally, you may specify the value of the index column of the row that corresponds to the VSN on the screen.

See "p (Display the Removable Media Load Requests Screen)" on page 3-69 and "load (Load a Cartridge into a Drive)" on page 3-48 for further information.

Command Synopsis

```
:clear vsn
:clear vsn index
```

Examples

The following example clears a load request for VSN VOL005:

```
Removable media load requests all both     samu               version time date
                                                              count: 0
index type pid    user        rb   flags     wait count  vsn
    0   li   0    root        900  -b-f---   0:00        VOL005
:
command: clear VOL005
```

d (Display Daemon Trace Controls Screen)

The d command opens the Daemon trace controls screen, which displays the events that Oracle HSM is currently tracing, as specified in the defaults.conf file. For more information about enabling trace files, see the defaults.conf man page.

Command Synopsis

```
d
:d
```

Controls

Ctrl-B (^b)
Move back one page.

Ctrl-F (^f)
Move forward one page.

Example

The following example shows trace file information. It includes information about the daemons being traced, the paths to the trace files, the events being traced, and information about the size and age of the trace files.

```
Daemon trace controls            samu               version time date

sam-amld    /var/opt/SUNWsamfs/trace/sam-amld
        cust err fatal misc proc debug date
        size  10  age 0
sam-archiverd /var/opt/SUNWsamfs/trace/sam-archiverd
        cust err fatal misc proc debug date
        size  10  age 0
sam-catserverd /var/opt/SUNWsamfs/trace/sam-catserverd
        cust err fatal misc proc debug date
        size  10  age 0
sam-fsd    /var/opt/SUNWsamfs/trace/sam-fsd
        cust err fatal misc proc debug date
        size  10  age 0
sam-rftd    /var/opt/SUNWsamfs/trace/sam-rftd
        cust err fatal misc proc debug date
        size  10  age 0
sam-recycler  /var/opt/SUNWsamfs/trace/sam-recycler
        cust err fatal misc proc debug date
        size  10  age 0
```

def_retention (Set Default WORM Retention Time)

When issued with the Oracle HSM equipment ordinal number of a disk family-set device and a time interval, the def_retention command changes the default retention period for WORM-enabled files that do not have a specified retention period to the specified time interval.

The data and path of a WORM (*Write Once Read Many*) file cannot be changed or deleted until the retention period expires. Specify the retention period in any of three ways:

- To indicate that the file should never be deleted or changed, enter the string permanent or 0.
- To specify an interval in years, days, hours, and/or minutes, enter a string of the form

  [*YEARS*y][*DAYS*d][*HOURS*h][*MINUTES*m]

  where *YEARS* is the specified number of years (if any), *MONTHS* is the specified number of months (if any), *DAYS* is the specified number of days (if any), *HOURS* is the specified number of hours (if any), and *MINUTES* is the specified number of minutes (if any). For example, 5y3d1h4m specifies 5 years, days, 1 hour, and 4 minutes, 10y specifies 10 years, and 60d12h specifies 60 days and 12 hours.
- To specify retention as a number of minutes, enter an integer in the range [1-2147483647].

If left unchanged, the default retention period for WORM files is 43,200 minutes (30 days).

For additional information on WORM options, see the mount_samfs and sam_worm man page.

Command Synopsis

```
:def_retention family-set-equipment-number retention-interval
```

Examples

The following example sets the default retention period for WORM files stored in equipment ordinal number 10 (file system samfs1) to 7 years:

```
File systems                samu                    version time date

ty    eq  state device_name           status       high low   mountpoint server
ma  10    on    samfs1                m----2----d   90% 70%   /samfs1
 mm  11    on    /dev/dsk/c5t10d0s0
 md  12    on    /dev/dsk/c5t10d0s1
 md  13    on    /dev/dsk/c5t11d0s0
 md  14    on    /dev/dsk/c5t11d0s1
:
command:def_retention 10 7y
```

devlog (Set Device Logging Options)

When issued with the Oracle HSM equipment ordinal number of a device and a space-delimited list of event-logging options, the devlog command configures logging for the specified device.

Specify event-logging options using one or more of the following key words:

- all logs all events.
- none disables logging.
- default restores the default logging settings.
- detail logs events that help track the progress of operations.
- err logs error messages.
- label logs labeling operations.
- mig logs migration toolkit messages.
- msg logs thread and process communications.
- retry logs retries of device operations.
- syserr logs system library errors.
- time logs time device operations.
- module includes module names and source lines in log messages.
- event includes event names in the message in log messages.
- date includes event dates in the message.
- - *keyword* (a minus sign preceding one of the keywords above) clears removes the keyword from the current logging specification.

The default logging specification is `err retry syserr date`. For additional information, see the `samset` man page.

Command Synopsis

```
:devlog equipment-number option [next-option]...
:devlog all option [next-option]...
:devlog equipment-number -option [-next-option]...
:devlog all -option [-next-option]...
```

Examples

The following example logs the event name in addition to the standard options for device `101`:

```
Removable media status: all             samu                    version time date

ty  eq   status      act  use  state    vsn
ti  101  ---------p    0   0%  notrdy
        empty
ti  102  ---------p    0   0%  notrdy
        empty
:
command:devlog 101 event err retry syserr date
```

dio_rd_consec (Limit the Number of Consecutive Direct Reads)

The `dio_rd_consec` command specifies the number of consecutive reads larger than `dio_rd_form_min` or `dio_rd_ill_min` that Oracle HSM should perform using direct I/O.

The default value is `0`, which disables direct reads based on I/O size and causes Oracle HSM to ignore `dio_rd_form_min` or `dio_rd_ill_min`.

For more information, see "`dio_rd_form_min` (Set the Size of Well-Aligned Direct I/O Reads)" on page 3-21 and "`dio_szero` (Zero Uninitialized Areas of Sparse Files During Direct I/O)" on page 3-22.

Command Synopsis

```
:dio_rd_consec family-set-equipment-number  number_of_consecutive_IOs
```

Examples

The following example sets minimum size of well-formed direct reads for family set device `100` to `16`:

```
Mass storage status                 samu                 version time date

ty  eq   status      use  state  ord   capacity   free       ra  part high low
ma  100  m----2---r- 1%   on            7.270T     7.270T   1M  16    80% 70%
 mm 110              1%   on       0   465.253G   465.231G [975661056 inodes]
 mr 120              1%   on       1     1.817T     1.817T
 mr 130              1%   on       2     1.817T     1.817T
 mr 140              1%   on       3     1.817T     1.817T
 mr 150              1%   on       4     1.817T     1.817T
```

```
:
command:dio_rd_consec 100 16
```

dio_rd_form_min (Set the Size of Well-Aligned Direct I/O Reads)

The dio_rd_form_min command makes a specified number of kilobytes the minimum size at which Oracle HSM will automatically switch from paged I/O to direct I/O when reading well-aligned data on a specified family-set disk device (see "Switched I/O" on page 2-3).

The default value is 256 kilobytes. Setting the value to 0 disables switched I/O during reads of well-aligned data.

Command Synopsis

```
:dio_rd_form_min family-set-equipment-number  number_kilobytes
```

Examples

The following example sets minimum size of well-formed direct reads for family set device 100 to 1024 kilobytes:

```
Mass storage status                    samu                version time date

ty  eq  status       use  state  ord   capacity   free       ra  part high low
ma  100  m----2---r- 1%   on            7.270T     7.270T   1M  16    80% 70%
 mm 110              1%   on       0   465.253G   465.231G [975661056 inodes]
 mr 120              1%   on       1     1.817T     1.817T
 mr 130              1%   on       2     1.817T     1.817T
 mr 140              1%   on       3     1.817T     1.817T
 mr 150              1%   on       4     1.817T     1.817T
:
command:dio_rd_form_min 100 1024
```

dio_rd_ill_min (Set the Size of Misaligned Direct I/O Reads)

The dio_rd_ill_min command makes a specified number of kilobytes the minimum size at which Oracle HSM will automatically switch from paged I/O to direct I/O when reading misaligned data on a specified family-set disk device (see "Switched I/O" on page 2-3).

The default value is 0 kilobytes which disables switched I/O during reads of misaligned data.

Command Synopsis

```
:dio_rd_ill_min family-set-equipment-number  number_kilobytes
```

Examples

The following example sets minimum size of malformed direct reads for family set device 100 to 4096 kilobytes:

```
Mass storage status                samu                version time date
```

```
ty  eq  status       use  state  ord   capacity   free       ra  part high low
ma  100  m----2---r- 1%   on            7.270T     7.270T   1M  16    80% 70%
 mm 110              1%   on       0   465.253G   465.231G [975661056 inodes]
 mr 120              1%   on       1     1.817T     1.817T
 mr 130              1%   on       2     1.817T     1.817T
 mr 140              1%   on       3     1.817T     1.817T
 mr 150              1%   on       4     1.817T     1.817T
:
command:dio_rd_form_min 100 4096
```

dio_szero (Zero Uninitialized Areas of Sparse Files During Direct I/O)

When issued with the Oracle HSM equipment ordinal number of a disk family-set device, the dio_szero command makes direct I/O to the specified device handle sparse files in the same way that paged I/O handles them. Uninitialized areas are zeroed as the areas are accessed. Sparse file behavior the same as that for paged I/O. This reduces performance when writing sparse files directly.

The default is nodio_szero (Do Not Zero Uninitialized Areas of Sparse Files with Direct I/O).

Command Synopsis

```
:dio_szero family-set-equipment-number
```

Examples

The following example zeroes sparse files on family set device 100:

```
Mass storage status                 samu                version time date

ty  eq  status       use  state  ord   capacity   free       ra  part high low
ma  100  m----2---r- 1%   on            7.270T     7.270T   1M  16    80% 70%
 mm 110              1%   on       0   465.253G   465.231G [975661056 inodes]
 mr 120              1%   on       1     1.817T     1.817T
 mr 130              1%   on       2     1.817T     1.817T
 mr 140              1%   on       3     1.817T     1.817T
 mr 150              1%   on       4     1.817T     1.817T
:
command:dio_szero 100
```

dio_wr_consec (Limit the Number of Consecutive Direct Writes)

The dio_wr_consec command specifies the number of consecutive writes larger than dio_wr_form_min or dio_wr_ill_min that Oracle HSM should perform using direct I/O.

The default value is 0, which disables direct writes based on I/O size and causes Oracle HSM to ignore dio_wr_form_min or dio_wr_ill_min.

For further details, see "dio_wr_form_min (Set the Size of Well-Formed Direct I/O Writes)" on page 3-23 and "dio_wr_ill_min (Set the Size of Misaligned Direct I/O Writes)" on page 3-23.

Command Synopsis

```
:dio_wr_consec family-set-equipment-number  number_of_consecutive_IOs
```

Examples

The following example sets minimum size of well-formed direct reads for family set device 100 to 16:

```
Mass storage status                samu               version time date

ty  eq  status       use  state  ord   capacity   free       ra  part high low
ma  100  m----2---r- 1%   on            7.270T     7.270T   1M  16    80% 70%
 mm 110              1%   on       0   465.253G   465.231G [975661056 inodes]
 mr 120              1%   on       1     1.817T     1.817T
 mr 130              1%   on       2     1.817T     1.817T
 mr 140              1%   on       3     1.817T     1.817T
 mr 150              1%   on       4     1.817T     1.817T
:
command:dio_wr_consec 100 16
```

dio_wr_form_min (Set the Size of Well-Formed Direct I/O Writes)

The dio_wr_form_min command makes a specified number of kilobytes the minimum size at which Oracle HSM will automatically switch from paged I/O to direct I/O when writing well-aligned data to a specified family-set disk device (see "Switched I/O" on page 2-3).

The default value is 256 kilobytes. Setting the value to 0 disables switched I/O during writes of well-aligned data.

Command Synopsis

```
:dio_wr_form_min  family-set-equipment-number  number_kilobytes
```

Examples

The following example sets minimum size of well-formed direct writes for family set device 100 to 1024 kilobytes:

```
Mass storage status              samu               version time date

ty  eq  status       use  state  ord   capacity   free       ra  part high low
ma  100  m----2---r- 1%   on            7.270T     7.270T   1M  16    80% 70%
 mm 110              1%   on       0   465.253G   465.231G [975661056 inodes]
 mr 120              1%   on       1     1.817T     1.817T
 mr 130              1%   on       2     1.817T     1.817T
 mr 140              1%   on       3     1.817T     1.817T
 mr 150              1%   on       4     1.817T     1.817T
:
command:dio_wr_form_min 100 1024
```

dio_wr_ill_min (Set the Size of Misaligned Direct I/O Writes)

The dio_wr_ill_min command makes a specified number of kilobytes the minimum size at which Oracle HSM will automatically switch from paged I/O to direct I/O

when writing misaligned data to a specified family-set disk device (see "Switched I/O" on page 2-3).

The default value is 0 kilobytes which disables switched I/O during reads of misaligned data.

Command Synopsis

```
:dio_wr_ill_min family-set-equipment-number  number_kilobytes
```

Examples

The following example sets minimum size of malformed direct writes for family set device 100 to 4096 kilobytes:

```
Mass storage status                samu                version time date

ty  eq  status       use  state  ord   capacity   free       ra  part high low
ma  100  m----2---r- 1%   on            7.270T     7.270T   1M  16    80% 70%
 mm 110              1%   on       0   465.253G   465.231G [975661056 inodes]
 mr 120              1%   on       1     1.817T     1.817T
 mr 130              1%   on       2     1.817T     1.817T
 mr 140              1%   on       3     1.817T     1.817T
 mr 150              1%   on       4     1.817T     1.817T
:
command:dio_wr_form_min 100 4096
```

diskvols (Set/Clear Flags on Disk Volumes Used for Archiving)

When issued with the name of an archival disk volume, a plus or a minus, and one of the flags listed below, the diskvols command sets or clears the specified flag on the specified volume. When troubleshooting hardware issues, you may want to change the values of the U and E flags. But leave the remainder alone:

- l means that the volume is labeled; a seqnum file has already been created.
- r means that the volume is defined on a remote host.
- U means that the volume is unavailable.
- R means that the volume is read only.
- E means that a media error occurred when writing to the disk archive directory.
- A means that the volume needs to be audited.
- F means that the volume is full.
- c means that the volume is ready for recycling.

Command Synopsis

```
:diskvols volume +flag
:diskvols volume -flag
```

Examples

The following example sets the U (unavailable) flag on the archival storage volume disk01:

```
Disk volume dictionary         samu           version time date
header
version 460
volumes
magic 340322 version 9 nkeys 2 ndata 2
index  space        capacity     used      flags  volume
    0  12882411520  12887785472  10291200  -----  disk01
    1  6443827200   6443892736      70656  -----  disk02
clients
magic 340322 version 9 nkeys 1 ndata 1
:
command:diskvols disk01 +U
```

The flag is set:

```
Disk volume dictionary         samu           version time date
header
version 460
volumes
magic 340322 version 9 nkeys 2 ndata 2
index  space        capacity     used      flags  volume
    0  12882411520  12887785472  10291200  --U--  disk01
    1  6443827200   6443892736      70656  -----  disk02
clients
magic 340322 version 9 nkeys 1 ndata 1
```

dmr (Enable Directed Mirror Reads of Software Mirrors)

When issued with the Oracle HSM equipment ordinal number of a disk family-set device, the dmr command lets Oracle Real Application Clusters (RAC) use the Directed Mirror Reads feature of Solaris Volume Manager when databases are stored on Oracle HSM file systems that use asynchronous I/O with Solaris Volume Manager mirrored volumes. See "nodmr (Disable Directed Mirror Reads of Software Mirrors)" on page 3-60.

By default, DMR is enabled.

Command Synopsis

```
:dmr family-set-equipment-number
```

Examples

The following example enables DMR on family-set device 20:

```
File systems                          samu                version time date

ty  eq state   device_name   status        high   low   mountpoint   server
ma  20 on      sam2          -----2---r-   80%    70%   /sam2
 mm 21 on      /dev/dsk/c0t600A0B800026B6F000001F0F4CCFF833d0s0
 mr 22 on      /dev/dsk/c0t600A0B800026E584000021864C4943AAd0s0
 mr 23 on      /dev/dsk/c0t600A0B800026E5840000218B4CBD9C4Ed0s0
ma  10 on      sam1          m----2---r-   80%    70%   /sam1
```

```
 mm 11 on      /dev/dsk/c0t600A0B8000118ACA000009594E1565E8d0s0
 mr 12 on      /dev/dsk/c0t600A0B8000114D5400000CA64E156793d0s0
 mr 13 on      /dev/dsk/c0t600A0B8000114D5400000CAA4E156897d0s0
 mr 14 on      /dev/dsk/c0t600A0B8000118ACA0000095B4E1568C8d0s0
 mr 15 on      /dev/dsk/c0t600A0B8000118ACA0000095D4E156948d0s0
:
command:dmr 10
```

dtrace (Enable/Configure Tracing for Oracle HSM Daemons)

When issued with an Oracle HSM daemon specifier plus the keyword `on` or `off` or a variable name and value, the `dtrace` command controls how trace information is written to trace files.

Specify daemons using the keyword `all` (for all Oracle HSM daemons) or the name of a particular daemon. Valid daemon names include `sam-archiverd`, `sam-catserverd`, `sam-fsd`, `sam-rftd`, `sam-recycler`, `sam-sharefsd`, `sam-stagerd`, `sam-serverd`, `sam-clientd`, and `fsmgmt`.

Use the `on` keyword to enable tracing for the specified daemon(s).

Use the `off` keyword to disable tracing for the specified daemon(s).

Use a variable name with a value to control the tracing process for the specified daemon. Set the variable using the syntax *`daemon-specifier.variable value`*.

Command Synopsis

```
:dtrace daemon-specifier on
:dtrace daemon-specifier off
:dtrace daemon-specifier.file path-filename
:dtrace daemon-specifier.options option_name1 option_name2 -option_name3 ...
```

Variables

`file`

The `file` variable takes a path or file-name value:

- `all.file` *`path-filename`* names the directory that holds the trace files *`path-filename`*.

  By default, trace files reside in `/var/opt/SUNWsamfs/trace`.

- *`daemon-name`*`.file`*`path-filename`* names the trace file *`path-filename`*.

  By default, trace files are named for the daemon they trace.

`options`

The `options` variable takes a space-separated list of *trace options*, events to be traced or elements to be included in trace output. To deselect an option, use a minus sign with the option (*`-option_name`*).

The pre-defined event types are `cust`, `err`, `fatal`, `misc`, `proc`, and `rft`. The message elements `program[pid]` and `time` are always included and can't be deselected.

Optional events include the following:

- `none` excludes all event types.
- `all` includes the most useful events: `cust`, `err`, `fatal`, `ipc`, `misc`, `proc`, and `rft`.

- alloc includes memory allocation events.
- cust includes customer-notification, syslog and notify file messages.
- err includes non-fatal program errors.
- fatal includes fatal syslog messages.
- files includes file actions.
- rft includes file transfer events.
- ipc includes inter-process communications.
- misc includes miscellaneous, otherwise uncategorized events.
- oprmsg includes operator messages.
- proc process initiation and completion.
- queue archiver queue contents when changed.

Optional message elements include:

- date includes the date in the message (the time is always included).
- module includes the source file name and line number in the message.
- type includes the event type in the message.

age
The age variable sets the time between trace file rotations to the specified value (rotating trace files keeps their size within manageable limits). Set the value of age to a number of seconds, minutes, hours, days, weeks, and/or years greater than 2 minutes.

A value of two minutes or less disables trace-file rotation.

size
The size variable sets the file size at which trace file rotations occur. Set the value of size as an integer with one of the following suffixes specifying the units: b for *bytes*, k for *kilobytes*, M for *megabytes*, G for *gigabytes*, or T for *terabytes*.

Examples

The following example excludes non-fatal errors from the sam-archiverd trace file and includes the event date and the source file and line number where the error occurred in each trace file entry:

```
Help information            page 1/15  samu                version time date

Displays:

    a   Archiver status                 w       Pending stage queue
    c   Device configuration            C       Memory
    d   Daemon trace controls           D       Disk volume dictionary
    f   File systems                    F       Optical disk label
    g   Shared clients                  I       Inode
    h   Help information                J       Preview shared memory
    l   Usage information               K       Kernel statistics
    m   Mass storage status             L       Shared memory tables
    n   Staging status                  M       Shared memory
    o   Optical disk status             N       File system parameters
    p   Removable media load requests   P       Active Services
    r   Removable media                 R       SAM-Remote
    s   Device status                   S       Sector data
```

```
    t   Tape drive status              T       SCSI sense data
    u   Staging queue                  U       Device table
    v   Robot catalog
:
command:dtrace sam-archiverd.options -err date module
```

export (Move Removable Media to the Robotic Library's Mailbox)

When issued with the Oracle HSM equipment ordinal number of a robotic library and a slot identifier, the export command moves the cartridge that resides in the specified slot to the mailbox. If specified, the slot must be in use and physically occupied by media. If the slot contains a two-sided optical cartridge and no side is specified, both sides are audited.

When issued with an Oracle HSM media-type identifier and the Volume Serial Number (VSN) of a cartridge, the export command moves the specified cartridge to the mailbox.

If the specified equipment ordinal number identifies a StorageTek network-attached library, the -f parameter exports the cartridge to the library's Cartridge Access Port (CAP) and updates the Oracle HSM catalog. The CAPID parameter must be defined in the library parameters file. See the stk man page for details.

If the specified equipment ordinal number identifies some other network-attached library, the export command updates the Oracle HSM catalog but leaves physical movement of the cartridge to the library's control software.

See also "import (Move Removable Media from the Mailbox into the Robotic Library)" on page 3-46.

Note that, by default, exported volumes are tracked by the Oracle HSM in the *historian*. The historian acts is a virtual library that can handle staging and archiving requests for volumes that are no longer resident in the library. An operator can then import the required cartridges to satisfy the load requests. Exporting cartridges from the historian itself deletes all information on the exported cartridge. The information about volumes on this cartridge will be lost. See the historian man page for further information.

Command Synopsis

```
:export equipment-number:slot
:export -f equipment-number:slot
:export media-type.vsn
:export -f media-type.vsn
```

Examples

In the following example, the first form of the command exports the contents of slot 1 in the robotic library with Oracle HSM equipment ordinal number 900. The second form exports the LTO cartridge with VSN VOL004.

```
Robot VSN catalog by slot : eq 900  samu       version time date
                                                        count 32
slot     access     time count  use  flags         ty vsn
   0     2013/07/02 12:53    3  96%  -il-o-b----f  li VOL001
   1     2013/07/02 12:53    2  98%  -il-o-b----f  li VOL002
   2     2013/07/02 12:52    1  86%  -il-o-b----f  li VOL003
```

```
   3     2013/07/02 12:52    1  95%  -il-o-b----f  li VOL004
...
:
command:export 900:1
...
command:export li.VOL004
```

f (Display File Systems Screen)

The f command opens the File systems screen and displays the components of Oracle HSM file systems.

Command Synopsis

```
f
:f
```

Controls

None.

Examples

The following example shows the file systems screen. Note that member drives are indented one space and appear directly below the file system to which they belong.

```
File systems                  samu                       version time date

ty   eq  state  device_name         status      high low  mountpoint server
ms   10  on     sam1                m----2----d  90%  70% /sam1
 md   11  on     /dev/dsk/c5t8d0s3
 md   12  on     /dev/dsk/c5t8d0s4
 md   13  on     /dev/dsk/c5t8d0s5
 md   14  on     /dev/dsk/c5t8d0s6
 md   15  on     /dev/dsk/c5t8d0s7
ms   20  on     sam2                m----2----d  90%  70%  /sam2
 md   21  on     /dev/dsk/c5t9d0s3
 md   22  on     /dev/dsk/c5t9d0s4
 md   23  on     /dev/dsk/c5t9d0s5
 md   24  on     /dev/dsk/c5t9d0s6
 md   25  on     /dev/dsk/c5t9d0s7
ma   30  on     qfs1                m----2----d  90%  70%  /qfs1
 mm   31  on     /dev/dsk/c5t10d0s0
 md   32  on     /dev/dsk/c5t10d0s1
ma   40  on     qfs2                m----2----d  90%  70%  /qfs2
 mm   41  on     /dev/dsk/c5t11d0s0
 md   42  on     /dev/dsk/c5t11d0s1
ma   50  on     qfs3                m----2---r-  90%  70%  /qfs3
 mm   51  on     /dev/dsk/c5t12d0s0
 mr   52  on     /dev/dsk/c5t12d0s1
ma   60  on     qfs4                m----2---r-  90%  70%  /qfs4
 mm   61  on     /dev/dsk/c5t13d0s0
 mr   62  on     /dev/dsk/c5t13d0s1
ma  100  on     shareqfs1           m----2c--r-  80%  70%  /shareqfs1 server
 mm  101  on     /dev/dsk/c6t50020F2300004655d0s0
 mr  102  on     /dev/dsk/c6t50020F2300004655d0s1
ma  110  on     shareqfs2           m----2c--r-  80%  70%  /shareqfs2 server
 mm  111  on     /dev/dsk/c6t50020F2300004655d0s6
```

```
mr  112  on     /dev/dsk/c6t50020F2300004655d0s7
```

Field Descriptions

ty
The ty field displays the Oracle HSM device type. Device-type identifiers are listed in Appendix A and on the mcf man page.

eq
The eq field displays the user-assigned Oracle HSM equipment number of the device.

state
The state field displays the current operating state of the device:

- on means that the device is available.
- ro means that the device is available read-only.
- off means that the device is not available.
- down means that the device is available for maintenance only.
- idle means that the device is completing operations that are already in progress but is not accepting new connections.
- noalloc means that no more storage can be allocated on this device. An administrator has run the noalloc (Disable Storage Allocation on a Device) command. The noalloc state persists until an administrator issues an alloc (Enable Storage Allocation on a Device) command.

device_name
The device_name field displays the path and file name of the device.

status
The status field displays an eleven-character string where each character position represents the value of a different file-system attribute. Character positions are numbered starting from the left, position 1, and ending at the right, position 11. The following table defines the possible state codes for each position in the status string.

Position	Value	Meaning
1	m----------	A lowercase m in position 1 means that the file system is mounted.
1	M----------	An uppercase M in position 1 means that the file system is being mounted.
2	-u---------	A lowercase u in position 2 means that the file system is being unmounted.
3	--A--------	An uppercase A in position 3 means that data is being *archived* (copied to nearline or offline storage).
4	---R-------	An uppercase R in position 4 means that previously archived data is being *released* from online disk storage to free space for new files.
5	----S------	An uppercase S in position 5 means that previously released data is being *staged* (copied from archival media to online disk storage).
6	-----1-----	The numeral 1 in position 6 means that this is a version-1 file system.
6	-----2-----	The numeral 2 in position 6 means that this is a version-2 file system.
7	------c----	A lowercase c in position 7 means that this is a shared file system.

Position	Value	Meaning
8	-------W---	An uppercase W in position 8 means that this is a single-writer file system.
9	--------R--	An uppercase R in position 9 means that this is a multi-reader file system.
10	---------r-	A lowercase r in position 10 means that the file system stores data on single-allocation mr devices.
11	----------d	A lowercase d in position 11 means that the file system stores data on dual-allocation md devices.

high
The high field displays the *high-water mark*, a storage utilization threshold expressed as the percentage of the file system's total capacity. When percent utilization exceeds the high-water mark, Oracle HSM starts deleting files that have been archived to backup media, starting with the largest and least recently used files. If users subsequently access the deleted files, Oracle HSM restores them to disk from an archival copy.

low
The *low-water mark*, a storage utilization threshold expressed as the percentage of the file system's total capacity. When percent utilization drops below the low-water mark, Oracle HSM stops deleting files from the disk. The low-water mark balances the need to maintain free disk space for new files against file-system performance by keeping files resident on disk and avoiding the need for excessive staging from archival media.

mountpoint
The directory where the file system attaches to the host's root file system.

server
The name of the metadata server (MDS) for a file system that is shared by multiple hosts.

flush_behind (Set the Flush-Behind Parameter)

When issued with the Oracle HSM equipment ordinal number of a disk family-set device and a number of kilobytes, the flush_behind command configures the corresponding file system to asynchronously write the specified number of kilobytes from dirty (modified) cache memory pages to disk.

As files change in memory, flush-behind insures that the changes are quickly saved to non-volatile storage media. When a process writes modified data to cache memory pages, the flushing follows close behind, copying the changes to disk even as the main write process continues. Unsaved, *dirty* cache pages are kept to a minimum, which helps to maintain the integrity of the file data.

On the other hand, flush-behind increases overhead and can reduce write performance. The process of flushing a data block may delay subsequent re-writes of that block. So flush-behind is disabled by default.

Set the specified value in the range [0-8192] kilobytes, where 0—the default—disables the flush-behind capability.

Command Synopsis

```
:stage_flush_behind family-set-equipment-number number_kilobytes
```

Examples

The following example sets staging flush-behind for family set device 100 to 1024 kilobytes:

```
Mass storage status                 samu                version time date

ty  eq  status       use  state  ord   capacity   free       ra  part high low
ma  100  m----2---r- 1%   on            7.270T     7.270T   1M  16    80% 70%
 mm 110              1%   on       0   465.253G   465.231G [975661056 inodes]
 mr 120              1%   on       1     1.817T     1.817T
 mr 130              1%   on       2     1.817T     1.817T
 mr 140              1%   on       3     1.817T     1.817T
 mr 150              1%   on       4     1.817T     1.817T
:
command:flush_behind 100 1024
```

forcedirectio (Use Direct I/O by Default)

When issued with the Oracle HSM equipment ordinal number of a disk family-set device, the forcedirectio command makes direct I/O the default for all input/output operations on the specified device.

Paged I/O is the default.

For more information, see "Understanding I/O Performance Tuning Objectives and Options" on page 2-2, the noforcedirectio (Do Not Use Direct I/O by Default) command, and the directio, setfa, sam_setfa, and sam_advise man pages.

Command Synopsis

```
:forcedirectio  family-set-equipment-number
```

Examples

The following example makes direct I/O the default for family set device 100:

```
Mass storage status                 samu                version time date

ty  eq  status       use  state  ord   capacity   free       ra  part high low
ma  100  m----2---r- 1%   on            7.270T     7.270T   1M  16    80% 70%
 mm 110              1%   on       0   465.253G   465.231G [975661056 inodes]
 mr 120              1%   on       1     1.817T     1.817T
 mr 130              1%   on       2     1.817T     1.817T
 mr 140              1%   on       3     1.817T     1.817T
 mr 150              1%   on       4     1.817T     1.817T
:
command:forcedirectio 100
```

force_nfs_async (Force Asynchronous NFS)

When issued with the Oracle HSM equipment ordinal number of a disk family-set device, the force_nfs_async command causes the specified device to override the Oracle HSM default and cache NFS (Network File System) data on the server, even if NFS has requested that the data be written through to disk.

The force_nfs_async command is effective only if the file system is mounted as an NFS server and only if the clients are mounted with the noac NFS mount option. For more information about mounting an NFS file system, see the mount_nfs man page.

Caution: Use this command with caution. In the event of a server interruption, data can be lost!

The force_nfs_async option violates NFS protocols. Data is cached on the NFS server and cannot be seen immediately by all the clients when multiple NFS servers are present (multiple NFS servers can be enabled within the QFS shared file system).

Command Synopsis

```
:force_nfs_async family-set-equipment-number
```

Examples

The following example sets force_nfs_async on family set 100:

```
Mass storage status              samu                version time date

ty  eq  status       use  state  ord   capacity   free       ra  part high low
ma  100  m----2---r- 1%   on            7.270T     7.270T   1M  16    80% 70%
 mm 110              1%   on       0   465.253G   465.231G [975661056 inodes]
 mr 120              1%   on       1     1.817T     1.817T
 mr 130              1%   on       2     1.817T     1.817T
 mr 140              1%   on       3     1.817T     1.817T
 mr 150              1%   on       4     1.817T     1.817T
:
command:force_nfs_async 100
```

fs (Specify a File System)

The fs command specifies the file system displayed by f (Display File Systems Screen) and the m (Display the Mass Storage Status Screen). When you have many file systems, the command lets you avoid paging through the display.

Command Synopsis

```
:fs file-system name
```

g (Display Shared File System Clients)

The g command opens the Shared clients screen and displays current configuration information.

When issued without a parameter, the g command displays archiver status for all file systems. When the command is issued with the name of a file system, the command displays status for the specified file system only.

Command Synopsis

```
g
:g
```

Controls

Ctrl-I (^I)
Toggle between the abbreviated and detailed views.

Ctrl-B
Move back one page.

Ctrl-B (^b)
Move back one page.

Ctrl-F (^f)
Move forward one page.

Ctrl-D (^d)
Move down the page.

Ctrl-U (^u)
Move up the page.

Field Descriptions

ord
The `ord` field displays the Oracle HSM server ordinal number (an internal identifier).

hostname
The `hostname` field displays the name of the server, potential server, or client host.

status
The `status` field displays the mount status, host type (server or client), file system type, and distributed I/O (datamover) status. In the abbreviated view, the field is displayed in bit-mask form. In the detailed view, they are expanded as shown in the example.

config and conf1
The `config` and `conf1` fields display file system configuration details. In the abbreviated view, the fields are displayed in bit-mask form. In the detailed view, they are expanded as shown in the example.

flags
The `flags` field displays abbreviated, textual status information such as mount status (`MNT`) and host type (`SVR` or `CLI`).

Examples

The following example shows the default, abbreviated view:

```
Shared clients              samu                 version time date
samsharefs is shared, server is samsharefs-mds, 2 clients 3 max

ord hostname             seqno nomsgs status  config   conf1  flags
  1 samsharefs-mds          14      0   8091  808540d   4051  0 MNT SVR
  2 samsharefs-client1      71      0   a0a1  808540d   4041  0 MNT CLI
```

This example shows the detailed view:

```
Shared clients              samu                 version time date
samsharefs is shared, server is samsharefs-mds, 2 clients 3 max
```

```
ord hostname             seqno nomsgs status  config   conf1  flags
  1 samsharefs-mds          14      0   8091  808540d   4051  0 MNT SVR

    config   :   CDEVID      ARCHIVE_SCAN    GFSID   OLD_ARCHIVE_FMT
    "        :   SYNC_META   TRACE   SAM_ENABLED     SHARED_MO
    config1  :   NFSV4_ACL   MD_DEVICES      SMALL_DAUS      SHARED_FS
    flags    :
    status   :   MOUNTED     SERVER  SAM     DATAMOVER
    last_msg :  Wed Jul  2 10:13:50 2014

  2 samsharefs-client1     127      0   a0a1  808540d   4041      0 MNT CLI

    config   :   CDEVID      ARCHIVE_SCAN    GFSID   OLD_ARCHIVE_FMT
    "        :   SYNC_META   TRACE   SAM_ENABLED     SHARED_MO
    config1  :   NFSV4_ACL   MD_DEVICES      SHARED_FS
    flags    :
    status   :   MOUNTED     CLIENT  SAM     SRVR_BYTEREV
    "        :   DATAMOVER
    last_msg :  Wed Jul  2 11:09:04 2014
```

h (Display the Help Screen)

The h command opens the Help information screen. The help screen summarizes samu displays and controls. By default, this is the first display that the system presents when you enter the samu command at the command line.

Command Synopsis

```
h
:h
```

Controls

Ctrl-B
Move back one page.

Ctrl-B (^b)
Move back one page.

Ctrl-F (^f)
Move forward one page.

Ctrl-D (^d)
Move down the page.

Ctrl-U (^u)
Move up the page.

Examples

There are fifteen Oracle HSM help screens in all, but archive-related displays are not shown if the Oracle Hierarchical Storage Manager software is not installed.

Help for Display Commands

The first example shows the first page that appears by default whenever you launch samu. Subsequent help screens show samu commands.

```
Help information       page 1/15       samu                    version time date
Displays:
  a  Archiver status                w  Pending stage queue
  c  Device configuration           C  Memory
  d  Daemon trace controls          D  Disk volume dictionary
  f  File systems                   F  Optical disk label
  h  Help information               I  Inode
  l  Usage information              J  Preview shared memory
  m  Mass storage status            K  Kernel statistics
  n  Staging status                 L  Shared memory tables
  o  Optical disk status            M  Shared memory
  p  Removable media load requests  N  File system Parameters
  r  Removable media                P  Active Services
  s  Device status                  R  SAM-Remote
  t  Tape drive status              S  Sector data
  u  Staging queue                  T  SCSI sense data
  v  Robot catalog                  U  Device table

  more (ctrl-f)
```

Help for samu Interface Controls

The second page of the help shows the hot keys that control the samu interface. Note that some controls are only applicable to certain displays:

```
Help information              page 2/15   samu                    version time date

Hot Keys:
    q   Quit
    :   Enter command
    sp  Refresh display
    ^f  Page display forward
    ^b  Page display backward
    ^d  Half-page display forward
    ^u  Half-page display backward
    ^i  Show details (selected displays)
    ^j  Toggle power for size, base 2 or 10 (selected displays)
    ^k  Advance display format
    ^l  Refresh display (clear)
    ^r  Toggle refresh
    /   Search for VSN (v display)
    %   Search for barcode (v display)
    $   Search for slot (v display)

    more (ctrl-f)
```

Help for Controlling Displays from the samu Command Prompt

The third page of the help shows display-control commands that can be entered at the same command prompt that control the displays. Note that some controls are only applicable to certain displays:

```
Help information            page 3/15   samu                    version time date

Display control commands:
    a [filesystem]        Archiver status
    n [media]             Staging status
```

```
    p [media]             Removable media load requests
    r [media]             Removable media
    u [media]             Staging queue
    v [eq]                Robot catalog
    w [media]             Pending stage queue
    C address             Memory
    I [inode]             Inode
    J [address]           Preview shared memory
    M [address]           Shared memory
    R [eq]                SAM-Remote
    S [address]           Sector data
    T [eq]                SCSI sense data
    U [eq]                Device table

    more (ctrl-f)
```

To view the archive display for a specified file system, you would enter the command-prompt hot key **:** (the colon). Then, at the samu command prompt, you would enter a *filesystem-equipment-identifier* as shown below:

```
    ...
    T [eq]                SCSI sense data
    U [eq]                Device table

    more (ctrl-f)
Command:a filesystem-equipment-identifier
```

The samu interface then displays a detailed archiving status report for the specified file system:

```
Archiver status                samu                   version time date

sam-archiverd: Waiting for resources

sam-arfind:  filesystem-equipment-identifier mounted at /mountpoint-directory
Files waiting to start 0  schedule 482  archiving  0
Monitoring file system activity.
Examine: noscan  Interval: 4m
Logfile: /var/opt/SUNWsamfs/archiver.log
events          12,138  syscalls 41,537  buffers  2
 archive             0
 change              3
 close           7,961
 create          3,927
 hwm                29
 modify              0
 rearchive           0
 rename              9
 remove            209
 unarchive           0
idstat          18,101 opendir  28 getdents  28 cached
```

Help for Miscellaneous File System Commands

The fourth page of the help shows various file-system commands that can be entered at the samu command prompt:

```
Help information            page 4/15   samu                   version time date

File System commands - miscellaneous:
```

```
    stripe eq value              Set stripe width
    suid eq                      Turn on setuid capability
    nosuid eq                    Turn off setuid capability
    sync_meta eq value           Set sync_meta mode
    atime eq value               Set access time (atime) update mode
    trace eq                     Turn on file system tracing
    notrace eq                   Turn off file system tracing
    add eq                       Add eq to mounted file system
    remove eq                    Remove eq; copy files to ON eqs
    release eq                   Release eq and mark files offline
    alloc eq                     Enable allocation on partition
    noalloc eq                   Disable allocation on partition
    def_retention eq interval    Set default WORM retention time

    more (ctrl-f)
```

For more information on individual commands, see the full descriptions listed below:

- stripe (Set the Stripe Width for a File System)
- suid (Mount a File System with setuid Execution Enabled)
- nosuid (Mount a File System with setuid Execution Disabled)
- sync_meta (Synchronize File System Metadata)
- atime (Configure Access Time Updates for a File System)
- trace (Enable File System Tracing)
- notrace (Disable File System Tracing)
- add (Add Equipment to a Mounted File System)
- remove (Move Data and Remove Equipment from a File System)
- release (Release Archived Data Files and Disable Allocation on a Device)
- alloc (Enable Storage Allocation on a Device)
- noalloc (Disable Storage Allocation on a Device)
- def_retention (Set Default WORM Retention Time)

Help for Oracle HSM File System Commands

The fifth page of the help shows Oracle HSM file-system commands that can be entered at the samu command prompt:

```
Help information            page 5/15   samu                  version time date

File System commands - SAM:
    hwm_archive eq               Turn on hwm archiver start
    nohwm_archive eq             Turn off hwm archiver start
    maxpartial eq value          Set maximum partial size in kilobytes
    partial eq value             Set size to remain online in kilobytes
    partial_stage eq value       Set partial stage-ahead point in kilobytes
    stage_flush_behind eq value  Set stage flush behind size in kilobytes
    stage_n_window eq value      Set direct stage size in kilobytes
    thresh eq high low           Set high and low release thresholds

    more (ctrl-f)
```

For more information on individual commands, see the full descriptions listed below:

- hwm_archive (Start Archiving When the High Water Mark is Reached)

- nohwm_archive (Do Not Start Archiving at the High Water Mark)
- maxpartial (Set the Maximum Partial Release Size for the File System)
- maxpartial (Set the Maximum Partial Release Size for the File System)
- partial (Set the Default Partial Release Size for the File System)
- partial_stage (Set Partial Stage-Ahead Point for the File System)
- stage_flush_behind (Set the Stage Flush Behind Size)
- stage_n_window (Size the Buffer for Reading Directly from Archival Media)
- thresh (Set High and Low Watermarks)

Help for File System I/O Commands

The sixth page of the help summarizes file-system input/output (I/O) commands that can be entered at the samu command prompt:

```
Help information            page 6/15   samu                  version time date

File System commands - I/O:
    dio_rd_consec eq value          Set number of consecutive dio reads
    dio_rd_form_min eq value        Set size of well-formed dio reads
    dio_rd_ill_min eq value         Set size of ill-formed dio reads
    dio_wr_consec eq value          Set number of consecutive dio writes
    dio_wr_form_min eq value        Set size of well-formed dio writes
    dio_wr_ill_min eq value         Set size of ill-formed dio writes
    flush_behind eq value           Set flush behind value in kilobytes
    forcedirectio eq                Turn on directio mode
    noforcedirectio eq              Turn off directio mode
    force_nfs_async eq              Turn on NFS async
    noforce_nfs_async eq            Turn off NFS async
    readahead eq value              Set maximum readahead in kilobytes
    writebehind eq value            Set maximum writebehind in kilobytes
    sw_raid eq                      Turn on software RAID mode
    nosw_raid eq                    Turn off software RAID mode
    wr_throttle eq value            Set outstanding write size in kilobytes
    abr eq                          Enable Application Based Recovery
    noabr eq                        Disable Application Based Recovery
    dmr eq                          Enable Directed Mirror Reads20a
    nodmr eq                        Disable Directed Mirror Reads
```

For more information on individual commands, see the full descriptions listed below:

- dio_rd_consec (Limit the Number of Consecutive Direct Reads)
- dio_rd_form_min (Set the Size of Well-Aligned Direct I/O Reads)
- dio_rd_ill_min (Set the Size of Misaligned Direct I/O Reads)
- dio_wr_consec (Limit the Number of Consecutive Direct Writes)
- dio_wr_form_min (Set the Size of Well-Formed Direct I/O Writes)
- dio_wr_ill_min (Set the Size of Misaligned Direct I/O Writes)
- flush_behind (Set the Flush-Behind Parameter)
- forcedirectio (Use Direct I/O by Default)
- noforcedirectio (Do Not Use Direct I/O by Default)
- force_nfs_async (Force Asynchronous NFS)

- noforce_nfs_async (Do Not Force Asynchronous NFS)
- readahead (Set The Maximum Read-Ahead Parameter)
- writebehind (Set the Write-Behind Parameter)
- sw_raid (Turn On Software RAID Mode)
- nosw_raid (Turn Off Software RAID Mode)wr_throttle (Set the Maximum Size of Outstanding Writes)
- abr (Enable Application Based Recovery)
- noabr (Disable Application Based Recovery)
- dmr (Enable Directed Mirror Reads of Software Mirrors)
- nodmr (Disable Directed Mirror Reads of Software Mirrors)

Help for QFS File System Commands

The seventh page of the help summarizes QFS file-system commands that can be entered at the samu command prompt:

```
Help information            page 7/15   samu                  version time date

File System commands - QFS:
    mm_stripe eq value             Set meta stripe width
    qwrite eq                      Turn on qwrite mode
    noqwrite eq                    Turn off qwrite mode

    more (ctrl-f)
```

For more information on individual commands, see the full descriptions listed below:

- mm_stripe (Set the Metadata Stripe Width)
- qwrite (Enable Simultaneous Reads/Writes to a File from Different Threads)
- noqwrite (Disable Simultaneous Reads/Writes from Different Threads)

Help for Multi-Reader File System Commands

The eighth page of the help summarizes multi-reader file-system commands that can be entered at the samu command prompt:

```
Help information            page 8/15   samu                  version time date

File System commands - multireader:
    invalid eq value               Set multireader invalidate cache delay
    refresh_at_eof eq              Turn on refresh at eof mode
    norefresh_at_eof eq            Turn off refresh at eof mode

    more (ctrl-f)
```

For more information on individual commands, see the full descriptions listed below:

- invalid (Set the Invalidate-Cache Delay for a Multi-Reader File System)
- refresh_at_eof (Refresh at EOF in a Multi-Reader File System)
- norefresh_at_eof (Do Not Refresh at EOF in a Multi-Reader File System)

Help for Shared File-System Commands

The ninth page of the help summarizes shared file-system commands that can be entered at the samu command prompt:

```
Help information            page 9/15   samu                    version time date

File System commands - shared fs:
    minoallocsz eq value           Set minimum allocation size
    maxallocsz eq value            Set maximum allocation size
    meta_timeo eq interval         Set shared fs meta cache timeout
    lease_timeo eq interval        Set shared fs lease relinquish timeout
    min_pool eq value              Set shared fs minimum threads count
    mh_write eq                    Turn on multihost read/write
    nomh_write eq                  Turn off multihost read/write
    aplease eq interval            Set append lease time
    rdlease eq interval            Set read lease time
    wrlease eq interval            Set write lease time

    more (ctrl-f)
```

For more information on individual commands, see the full descriptions listed below:

- minallocsz (Set the Minimum Allocation Size for a Shared File System)
- maxallocsz (Set the Maximum Allocation Size for a Shared File System)
- meta_timeo (Set the Metadata Cache Timeout for a Shared File System)
- lease_timeo (Set the Unused Lease Timeout for a Shared File System)
- migconfig (Configure Migration to New Media)
- mh_write (Turn on Multi-Host Read/Write)
- nomh_write (Turn Off Multi-Host Read/Write)
- aplease (Set Append Lease Expiration Time for a Shared File System)
- rdlease (Set the Read Lease Time for a Shared File System)
- wrlease (Set the Write Lease Time for a Shared File System)

Help for Device Commands

The tenth page of the help summarizes shared device commands that can be entered at the samu command prompt:

```
Help information            page 10/15   samu                    version time date

Device commands:
    devlog eq [option ...]         Set device logging options
    idle eq                        Idle device
    off eq                         Turn off device
    on eq                          Turn on device
    readonly eq                    Make device read only
    ro eq                          Make device read only
    unavail eq                     Make device unavailable
    unload eq                      Unload device

    more (ctrl-f)
```

For more information on individual commands, see the full descriptions listed below:

- devlog (Set Device Logging Options)

- idle (Idle Device)
- off (Turn Off A Device)
- on (Turn On A Device)
- readonly | ro (Make Device Read Only)
- unavail (Make a Device Unavailable to Oracle HSM)
- unload (Unload a Drive)

Help for Robot Commands

The eleventh page of the help summarizes robot commands that can be entered at the samu command prompt:

```
Help information            page 11/15  samu                    version time date

Robot commands:
    audit [-e] eq[:slot[:side]]    Audit slot or library
    import eq                      Import cartridge from mailbox
    export [-f] eq:slot            Export cartridge to mailbox
    export [-f] mt.vsn             Export cartridge to mailbox
    load eq:slot[:side]            Load cartridge in drive
    load [-f] mt.vsn               Load cartridge in drive
    priority pid priority          Set priority in preview queue

    more (ctrl-f)
```

For more information on individual commands, see the full descriptions listed below:

- audit (Audit Library or Library Slot)
- import (Move Removable Media from the Mailbox into the Robotic Library)
- export (Move Removable Media to the Robotic Library's Mailbox)
- load (Load a Cartridge into a Drive)
- priority (Set Priority in Load-Request Preview Queue)

Help for Archiver Commands

The twelfth page of the help summarizes archiver commands that can be entered at the samu command prompt:

```
Help information            page 12/15  samu                   version time date

Archiver commands:
    aridle [dk | rm | fs.fsname]         Idle archiving
    arrerun                              Soft restart archiver
    arrestart                            Restart archiver
    arrmarchreq fsname.[* | arname]      Remove ArchReq(s)
    arrun [dk | rm | fs.fsname]          Start archiving
    arscan fsname[.dir | ..inodes][int]  Scan filesystem
    arstop [dk | rm | fs.fsname]         Stop archiving
    artrace [fs.fsname]                  Trace archiver

    more (ctrl-f)
```

For more information on individual commands, see the full descriptions listed below:

- aridle (Stop Archiving Gracefully)

- arrerun (Restart the Archiver Non-Disruptively)
- arrmarchreq (Remove Archive Requests)
- arrun (Start Archiving)
- arscan (Scan a File System for Unarchived Files)
- arstop (Stop All Archiving Immediately)
- artrace (Perform Archiver Tracing)

Help for Stager Commands

The thirteenth page of the help summarizes stager commands that can be entered at the samu command prompt:

```
Help information            page 13/15  samu                  version time date

Stager commands:
    stclear mt.vsn                Clear stage request
    stidle                        Idle staging
    strun                         Start staging

    more (ctrl-f)
```

For more information on individual commands, see the full descriptions listed below:

- stclear (Clear a Stage Request)
- stidle (Stop Staging as Soon as Possible)
- strun (Start Staging)

Help for Media Migration Commands

The fourteenth page of the help summarizes stager commands that can be entered at the samu command prompt:

```
Help information            page 14/15  samu           version time date
Migration commands:
    migconfig                     Config migration
    migstart                      Start migration
    migidle                       Idle migration
    migstop                       Stopmigration



    more (ctrl-f)
```

Help for Miscellaneous Commands

The fifteenth page of the help summarizes miscellaneous commands that can be entered at the samu command prompt:

```
Help information            page 15/15  samu            version time date

Miscellaneous commands:
    clear vsn [index]             Clear load request
    dtrace daemon[.variable] value Daemon trace controls
    fs [filesystem]               Select a filesystem name (ex samfs1)
    mount mountpoint              Select a mount point
    open eq                       Open device
    q                             Exit from samu
```

```
    refresh [interval]             Set display refresh interval in seconds
    read address                   Read sector from device
    snap [filename]                Snapshot screen to file
    diskvols volume +flag | -flag  Set or clear disk volume flags
    !shell-command                 Run command in a shell

    more (ctrl-f)
```

See the full descriptions for more information on the commands listed below:

- clear (Clear a Load Request for a Removable Storage Volume)
- dtrace (Enable/Configure Tracing for Oracle HSM Daemons)
- fs (Specify a File System)
- mount (Select a Mount Point)
- open (Open a Disk Storage Device for Reading)
- refresh (Set the Automatic Refresh Interval for the samu Display)
- read (Read Sectors On a Disk Storage Device)
- snap (Save a Snapshot of the Current Screen to a File)
- diskvols (Set/Clear Flags on Disk Volumes Used for Archiving)

Help for Media Types

The fifteenth page of the help lists media type codes recognized by the samu interface:

```
Help information            page 15/15  samu                    version time date

Media types:
    all All media types                     tp  tape
    at  Sony AIT tape                       sa  Sony Super AIT tape
    d2  Ampex DST310 (D2) tape
    d3  STK SD-3 tape
    dt  4mm digital tape (DAT)
    fd  Fujitsu M8100 128track tape
    i7  IBM 3570 tape                       ib  IBM 3590 tape
    li  IBM 3580, Seagate Viper 200 (LTO    m2  IBM 3592 tape
    lt  digital linear tape (DLT)
    se  STK 9490 tape                       sf  STK T9940 tape
    sg  STK 9840 tape                       st  STK 3480 tape
    so  Sony DTF tape
    vt  Metrum VHS tape
    xm  Exabyte Mammoth-2 8mm tape          xt  Exabyte 8mm tape
    od  optical
    mo  5 1/4 in. erasable optical disk
    wo  5 1/4 in. WORM optical disk         o2  12 in. WORM optical disk
```

hwm_archive (Start Archiving When the High Water Mark is Reached)

When issued with an Oracle HSM family-set equipment ordinal number, the hwm_archive command configures the corresponding file system to start archiving automatically when the total size of the stored files passes the *high water mark*, the maximum space-utilization threshold specified for the file system. See nohwm_archive (Do Not Start Archiving at the High Water Mark).

Command Synopsis

```
:hwm_archive equipment-number
```

Examples

The following example configures disk family-set device 20 to automatically start archiving when the corresponding file-system is 80% full:

```
File systems                          samu                  version time date

ty  eq state   device_name   status        high   low   mountpoint   server
ma  20 on      sam2          -----2---r-   80%    70%   /sam/sam2
 mm 21 on      /dev/dsk/c0t600A0B800026B6F000001F0F4CCFF833d0s0
 mr 22 on      /dev/dsk/c0t600A0B800026E584000021864C4943AAd0s0
 mr 23 on      /dev/dsk/c0t600A0B800026E5840000218B4CBD9C4Ed0s0
:
command:hwm_archive 20
```

idle (Idle Device)

The idle command blocks new connections to the device specified by an Oracle HSM equipment ordinal number. Operations that are already in progress continue.

Command Synopsis

```
:idle equipment-number
```

Examples

The example below idles device 103:

```
Device configuration:            samu               version time date

ty  eq  state   device_name                          fs   family_set
sk  100 on      /etc/opt/SUNWsamfs/ libraryB_T10K_parms  100   libraryB
ti  101 on      /dev/rmt/0cbn                        100   libraryB
ti  102 on      /dev/rmt/1cbn                        100   libraryB
ti  103 on      /dev/rmt/2cbn                        100   libraryB
hy  104 on      historian                            104
:
command:idle 103
```

The device state changes to idle:

```
Device configuration:            samu               version time date

ty   eq state   device_name                          fs   family_set
sk  100 on      /etc/opt/SUNWsamfs/ libraryB_T10K_parms  100   libraryB
ti  101 on      /dev/rmt/0cbn                        100   libraryB
ti  102 on      /dev/rmt/1cbn                        100   libraryB
ti  103 idle    /dev/rmt/2cbn                        100   libraryB
hy  104 on      historian                            104
```

import (Move Removable Media from the Mailbox into the Robotic Library)

When issued with the Oracle HSM equipment ordinal number of a robotic library, the import command tells the library to move media from the library's mailbox to the first available slot in the library.

See also export (Move Removable Media to the Robotic Library's Mailbox).

Command Synopsis

```
:import equipment-number
```

Examples

The following example below imports the contents of the mailbox in the robotic library with Oracle HSM equipment ordinal number 900:

```
Robot VSN catalog by slot : eq 900  samu       version time date
                                                        count 32
slot     access     time count  use  flags         ty vsn
   0     2013/07/02 12:53    3  96%  -il-o-b----f  li VOL001
   1     2013/07/02 12:53    2  98%  -il-o-b----f  li VOL002
   2     2013/07/02 12:52    1  86%  -il-o-b----f  li VOL003
   3     2013/07/02 12:52    1  95%  -il-o-b----f  li VOL004
   4     2013/07/02 12:51    1  79%  -il-o-b-----  li VOL005
   5     2013/07/02 12:45    0  14%  -ilEo-b----f  li VOL006 MEDIA ERROR
...
:
command: import 900
```

invalid (Set the Invalidate-Cache Delay for a Multi-Reader File System)

When issued with the Oracle HSM equipment ordinal number of a disk family-set device that is mounted read-only in a single-writer/multi-reader file system, the invalid command tells the readers of a single-writer/multi-reader file system to delay invalidating their file-system metadata caches for a specified number of seconds. When files are modified, the host waits the prescribed interval before checking the metadata server for updates.

> **Caution:** Setting a non-zero delay value keeps stale file-system metadata cached, so the host may not have a consistent, current view of a read-only file system during the delay.

The specified number of seconds must fall in the range [0-60]. The default is 0 seconds.

Command Synopsis

```
:invalid family-set-equipment-number interval_in_seconds
```

Examples

The following example sets the cache invalidation delay to 30 seconds on family set device 900:

```
File systems                        samu                    version time date

ty    eq   state  device_name           status       high  low  mountpoint server
ma    900  on     qfs1                  m----2--R-d   90%  70%  /qfs1
 mm    901  on     /dev/dsk/c5t10d0s0
 md    902  on     /dev/dsk/c5t10d0s1
 md    903  on     /dev/dsk/c5t11d0s0
 md    904  on     /dev/dsk/c5t11d0s1
:
command:invalid 900 30
```

l (Display Utilization Information)

The l command opens the Usage information screen. The screen displays storage utilization statistics for the file systems and associated archival storage.

Command Synopsis

```
l
:l
```

Controls

Ctrl-J (^j)
Toggle between binary and decimal units of size.

Examples

The following example shows a sample usage screen.

```
Usage information                    samu                     version time date

hostid = 0xID      OS name: SunOS   Architecture: SPARC   CPUs: 2 (2 online)
library    40:     capacity 389.3G  bytes   space 291.1G  bytes, usage  25%
library    51:     capacity   9.5G  bytes   space   9.5G  bytes, usage   0%
library    55:     capacity   0     bytes   space   0     bytes, usage   0%
library    56:     capacity  10.7G  bytes   space  10.7G  bytes, usage   0%
library totals:    capacity 409.5G  bytes   space 311.3G  bytes, usage  24%

filesystem samfs3: capacity  54.5M  bytes   space  13.4M  bytes, usage  75%
filesystem samfs4: capacity 319.5M  bytes   space 298.0M  bytes, usage   7%
filesystem samfs7: capacity  96.6M  bytes   space  69.6M  bytes, usage  28%
filesystem samfs6: capacity   5.0G  bytes   space   4.9G  bytes, usage   3%
filesystem samfs8: capacity   5.0G  bytes   space   4.9G  bytes, usage   2%
filesystem totals: capacity  10.5G  bytes   space  10.2G  bytes, usage   3%
```

lease_timeo (Set the Unused Lease Timeout for a Shared File System)

When issued with the Oracle HSM equipment ordinal number of a shared disk family-set device and a number, the lease_timeo command configures the way in

which the corresponding shared file system manages unused read, write, and append leases.

Specify a number in the range [0-15] to force processes to relinquish leases that remain unused for the specified number of seconds.

Specify the number -1 to disable unused lease timeouts and let processes hold unused leases for the full terms specified by aplease (Set Append Lease Expiration Time for a Shared File System), rdlease (Set the Read Lease Time for a Shared File System), and "wrlease (Set the Write Lease Time for a Shared File System)". The default is 0.

Note that read and write leases are not relinquished if mh_write (Turn on Multi-Host Read/Write) is enabled.

Command Synopsis

```
:lease_timeo family-set-equipment-number 0-to-15-seconds
:lease_timeo family-set-equipment-number -1
```

Examples

The following example sets the unused lease timeout for shared file system shareqfs1 on family set device 800 to 5 seconds:

```
File systems                          samu                    version time date

ty  eq  state  device_name              status      high low  mountpoint server
ma 800  on     shareqfs1                m----2c--r-  80% 70%  /shareqfs1
mm 801  on     /dev/dsk/c6t50020F2300004655d0s0
mr 802  on     /dev/dsk/c6t50020F2300004655d0s1
ma 810  on     shareqfs2                m----2c--r-  80% 70%  /shareqfs2
mm 811  on     /dev/dsk/c6t50020F2300004655d0s6
mr 812  on     /dev/dsk/c6t50020F2300004655d0s7
:
command:lease_timeo 800 5
```

load (Load a Cartridge into a Drive)

The load command loads media into the drive specified by either an Oracle HSM equipment ordinal number and slot location or a specified Oracle HSM media type and volume serial number (VSN). Media-type identifiers are listed in Appendix A and on the mcf man page.

Command Synopsis

```
:load equipment-number:slot[:side]
:load media-type.vsn
```

Examples

The following example loads the LTO (li) with volume serial number VOL001:

```
Device configuration:                 samu                    version time date

ty   eq state   device_name                           fs   family_set
sk  100 on      /etc/opt/SUNWsamfs/SL8500_T10K_parms  100  SL8500
```

```
li  101 on      /dev/rmt/0cbn                        100  SL8500
li  102 on      /dev/rmt/1cbn                        100  SL8500
li  103 on      /dev/rmt/2cbn                        100  SL8500
hy  104 on      historian                            104
:
command:load li.VOL001
```

m (Display the Mass Storage Status Screen)

The m command opens the Mass storage status screen. The screen displays the status of mounted file systems and their member drives.

Command Synopsis

```
m
:m
```

Controls

Ctrl-J (^j)
Toggle between binary and decimal units of size.

Examples

In the following example, note that member drives are indented one space and appear directly below the file system to which they belong.

```
Mass storage status                 samu                    version time date

ty   eq  status       use state ord  capacity   free     ra part high low
ms   10  m----2----d   1% on          68.354G   68.343G  1M   16  90% 70%
 md   11               1% on      0   13.669G   13.666G
 md   12               1% on      1   13.669G   13.667G
 md   13               1% on      2   13.669G   13.667G
 md   14               1% on      3   13.674G   13.672G
 md   15               1% on      4   13.674G   13.672G
ms   20  m----2----d   1% on          68.354G   68.344G  1M   16  90% 70%
 md   21               1% on      0   13.669G   13.667G
 md   22               1% on      1   13.669G   13.667G
 md   23               1% on      2   13.669G   13.667G
 md   24               1% on      3   13.674G   13.672G
 md   25               1% on      4   13.674G   13.672G
ma   30  m----2----d   4% on          64.351G   61.917G  1M   16  90% 70%
 mm   31               1% on      0    4.003G   3.988G [8363840 inodes]
 md   32               4% on      1   64.351G  61.917G
ma   40  m----2----d   1% on          64.351G  64.333G   1M   16  90% 70%
 mm   41               1% on      0    4.003G   3.997G [8382784 inodes]
 md   42               1% on      1   64.351G  64.333G
```

Field Descriptions

ty
The ty field displays the Oracle HSM device type. Device-type identifiers are listed in Appendix A and on the mcf man page.

eq

The eq field displays the user-assigned Oracle HSM equipment ordinal number for the mass storage device.

status

The status field displays an eleven-character string in which each character position represents the value of a different file-system attribute. Character positions are numbered starting from the left, position 1, and ending at the right, position 11. The following table defines the possible state codes for each position in the status string.

Position	Value	Meaning
1	m----------	A lowercase m in position 1 means that the file system is mounted.
1	M----------	An uppercase M in position 1 means that the file system is being mounted.
2	-u---------	A lowercase u in position 2 means that the file system is being unmounted.
3	--A--------	An uppercase A in position 3 means that data is being *archived* (copied to nearline or offline storage).
4	---R-------	An uppercase R in position 4 means that previously archived data is being *released* from online disk storage to free space for new files.
5	----S------	An uppercase S in position 5 means that previously released data is being *staged* (copied from archival media to online disk storage).
6	-----1-----	The numeral 1 in position 6 means that this is a version-1 file system.
6	-----2-----	The numeral 2 in position 6 means that this is a version-2 file system.
7	------c----	A lowercase c in position 7 means that this is a shared file system.
8	-------W---	An uppercase W in position 8 means that this is a single-writer file system.
9	--------R--	An uppercase R in position 9 means that this is a multi-reader file system.
10	---------r-	A lowercase r in position 10 means that the file system stores data on single-allocation mr devices.
11	----------d	A lowercase d in position 11 means that the file system stores data on dual-allocation md devices.

use

The use field displays the percentage of disk space in use.

state

The state field displays the current operating state of the device:

- on means that the device is available.
- ro means that the device is available read-only.
- off means that the device is not available.
- down means that the device is available for maintenance only.
- idle means that the device is completing operations that are already in progress but is not accepting new connections.
- noalloc means that no more storage can be allocated on this device. An administrator has run the noalloc (Disable Storage Allocation on a Device) command. The noalloc state persists until an administrator issues an alloc (Enable Storage Allocation on a Device) command.

ord
The ord field displays the ordinal number of the device within the user-defined Oracle HSM family set (family sets represent the equipment associated with an Oracle HSM component, such as the group of disk devices associated with an Oracle HSM file system or the tape library associated with archiving).

capacity
The capacity field displays the number of 1024-byte blocks of usable space on the disk.

free
The free field displays the number of 1024-byte blocks of disk space available.

ra
The ra field displays the read-ahead size in kilobytes.

part
The part field displays the partial stage size in kilobytes.

high
The high field displays the high water mark for the file system. The high water mark is the percentage disk utilization at which Oracle HSM starts to delete files that have archive copies from disk, so that there is space for new files.

low
The low displays the low water mark for the file system. The low water mark is the percentage disk utilization at which Oracle HSM has freed enough space for new files and stops deleting archived files.

maxallocsz (Set the Maximum Allocation Size for a Shared File System)

When issued with the Oracle HSM equipment ordinal number of a shared disk family-set device and a kilobyte value, the maxallocsz command sets maximum number of kilobytes that are allocated ahead of a write to the specified shared file system. See minallocsz (Set the Minimum Allocation Size for a Shared File System).

The specified allocation size must be a multiple of 8 kilobytes in the range [16-4194304]. The default value is 8 times the device allocation unit (DAU). See the discussion of the -a option on the sammkfs man page for additional information on DAUs.

Command Synopsis

```
:maxallocsz family-set-equipment-number  number_kilobytes
```

Examples

The following example sets the maximum allocation to 1024 kilobytes on family set device 800:

```
File systems                          samu                    version time date

ty  eq   state  device_name              status      high low  mountpoint server
ma 800   on     shareqfs1                m----2c--r-  80% 70%  /shareqfs1
mm  801  on     /dev/dsk/c6t50020F2300004655d0s0
mr  802  on     /dev/dsk/c6t50020F2300004655d0s1
ma  810  on     shareqfs2                m----2c--r-  80%  70% /shareqfs2
```

```
mm  811  on     /dev/dsk/c6t50020F2300004655d0s6
mr  812  on     /dev/dsk/c6t50020F2300004655d0s7
:
command:maxallocsz 800 1024
```

maxpartial (Set the Maximum Partial Release Size for the File System)

When issued with an Oracle HSM family-set equipment ordinal number and a number of kilobytes, the maxpartial command sets the maximum number of bytes from the start of a file that can be retained in disk cache when the file is released.

A large file that has been designated for partial release can open more quickly when restaged from archival media because the beginning of the file is already on disk when staging begins. For further information, see partial (Set the Default Partial Release Size for the File System), partial_stage (Set Partial Stage-Ahead Point for the File System), and the release man page.

The specified kilobyte value must be an integer in the range [8-2097152]. The default is 16 kilobytes.

Command Synopsis

```
:maxpartial equipment-number number_of_kilobytes
```

Examples

The following example configures disk family-set device 20 to allow retention of a maximum of 512 kilobytes on disk for each file designated for partial release:

```
File systems                          samu                    version time date

ty  eq   state  device_name              status      high  low  mountpoint server
ma  20   on      sam2                    -----2---r-   80% 70%  /sam1
 mm 21   on       /dev/dsk/c0t600A0B800026B6F000001F0F4CCFF833d0s0
 mr 22   on       /dev/dsk/c0t600A0B800026E584000021864C4943AAd0s0
 mr 23   on       /dev/dsk/c0t600A0B800026E5840000218B4CBD9C4Ed0s0
:
command:maxpartial 20 512
```

meta_timeo (Set the Metadata Cache Timeout for a Shared File System)

When issued with the Oracle HSM equipment ordinal number of a shared disk family-set device and a number of seconds, the meta_timeo command sets the maximum metadata cache age for the specified device to the value supplied. Hosts use cached file-system attributes and directory data for the specified number of seconds, then refresh the cache by checking with the metadata server (MDS).

The default caching interval is 3 seconds. An interval of 0 disables metadata caching.

Command Synopsis

```
:meta_timeo family-set-equipment-number interval_in_seconds
```

Examples

The following example sets the cache interval on family set device 800 to 0 seconds, so that the host checks metadata consistency with the MDS prior to each use of the file system:

```
File systems                          samu                    version time date

ty  eq   state  device_name              status      high low  mountpoint server
ma  800  on     shareqfs1                m----2c--r-  80% 70%  /shareqfs1
mm  801  on     /dev/dsk/c6t50020F2300004655d0s0
mr  802  on     /dev/dsk/c6t50020F2300004655d0s1
ma  810  on     shareqfs2                m----2c--r-  80%  70% /shareqfs2
mm  811  on     /dev/dsk/c6t50020F2300004655d0s6
mr  812  on     /dev/dsk/c6t50020F2300004655d0s7
:
command:meta_timeo 800 0
```

mh_write (Turn on Multi-Host Read/Write)

When issued with the Oracle HSM equipment ordinal number of a shared disk family-set device, the mh_write command enables simultaneous reads and writes to the same file from multiple hosts.

By default, shared Oracle HSM file systems allow simultaneous reads from multiple hosts but not writes. To reverse the effects of the mh_write command, use nomh_write (Turn Off Multi-Host Read/Write).

Command Synopsis

```
:mh_write family-set-equipment-number
```

Examples

The following example enables multi-host reads on family set device 800:

```
File systems                          samu                    version time date

ty  eq   state  device_name              status      high low  mountpoint server
ma  800  on     shareqfs1                m----2c--r-  80% 70%  /shareqfs1
mm  801  on     /dev/dsk/c6t50020F2300004655d0s0
mr  802  on     /dev/dsk/c6t50020F2300004655d0s1
ma  810  on     shareqfs2                m----2c--r-  80%  70% /shareqfs2
mm  811  on     /dev/dsk/c6t50020F2300004655d0s6
mr  812  on     /dev/dsk/c6t50020F2300004655d0s7
:
command:mh_write 800
```

mh_write (Turn on Multi-Host Read/Write)

When issued with the Oracle HSM equipment ordinal number of a shared disk family-set device, the mh_write command enables simultaneous reads and writes to the same file from multiple hosts.

By default, shared Oracle HSM file systems allow simultaneous reads from multiple hosts but not writes. To reverse the effects of the mh_write command, use nomh_write (Turn Off Multi-Host Read/Write).

Command Synopsis

```
:mh_write family-set-equipment-number
```

Examples

The following example enables multi-host reads on family set device 800:

```
File systems                    samu                    version time date

ty  eq   state  device_name              status      high low  mountpoint server
ma  800  on     shareqfs1                m----2c--r-  80% 70%  /shareqfs1
mm  801  on     /dev/dsk/c6t50020F2300004655d0s0
mr  802  on     /dev/dsk/c6t50020F2300004655d0s1
ma  810  on     shareqfs2                m----2c--r-  80%  70% /shareqfs2
mm  811  on     /dev/dsk/c6t50020F2300004655d0s6
mr  812  on     /dev/dsk/c6t50020F2300004655d0s7
:
command:mh_write 800
```

migconfig (Configure Migration to New Media)

The migconfig command loads the migrationd.cmd file that provides the instructions governing migration from one set of media to a set of replacement media.

Command Synopsis

```
:migconfig
```

Examples

```
Migration status              samu                    version date time
Source Vsns - wait:  0 fsscan: 0 copy: 0 update ino: 0 log: 0 done:  0
Status: Stop: Waiting for :migstart
source  dest  cmod  perc  status

:
command:migconfig
```

migidle (Configure Migration to New Media)

The migidle command stops media migration at the next logical point (such as the end of the current archive tar file).

Command Synopsis

```
:migidle
```

Examples

```
Migration vsn list            samu                  version date time
Status:  Run  Vsns:2 src:1 dest:1 maxcopy:2
ord m ty vsn     start time  end time  status  Inodes done/tot    bytes
  0 S li VOL001  10/16 12:12 none      copy            0/35023   164.50M
  0 D li VOL012  10/16 12:12 none      copy                      148.75M


:
command:migidle
```

migstart (Start the Media Migration Service)

The migstart command starts the media migration daemon. The daemon reads the configuration information in the migrationd.cmd file and starts the process.

Command Synopsis

```
:migstart
```

Examples

```
Migration status               samu                  version date time
Status: Stop: Waiting for :migstart  Vsns:2 src:1 dest:1 maxcopy:0
ord m ty vsn    start time  end time    status  Inodes done/tot         bytes
  0 S li VOL002 none        10/16 15:42 error             0/0             0
  0 D li VOL011 none        none        avail                             0


:
command:migstart
```

migstop (Stop the Media Migration Service)

The migstop command stops the media migration daemon.

Command Synopsis

```
:migstop
```

Examples

```
Migration vsn list             samu                  version date time
Status:  Run  Vsns:2 src:1 dest:1 maxcopy:2
ord m ty vsn     start time  end   time  status  Inodes done/tot    bytes
  0 S li VOL001  10/16 12:12 10/16 12:14 complete    35023/35023   550.00M
  0 D li VOL012  10/16 12:12 10/16 12:14 avail                     550.00M


:
command:migstop
```

minallocsz (Set the Minimum Allocation Size for a Shared File System)

When issued with the Oracle HSM equipment ordinal number of a shared disk family-set device and a kilobyte value, the minallocsz command sets minimum number of kilobytes that are allocated ahead of a write to the specified shared file system. See maxallocsz (Set the Maximum Allocation Size for a Shared File System).

The specified allocation size must be a multiple of 8 kilobytes in the range [16-4194304]. The default value is 8 times the device allocation unit (DAU). See the discussion of the -a option on the sammkfs man page for additional information on DAUs.

Command Synopsis

```
:minallocsz family-set-equipment-number number_kilobytes
```

Examples

The following example sets the minimum allocation to 512 kilobytes on family set device 800:

```
File systems                    samu                    version time date

ty  eq   state  device_name              status      high low  mountpoint server
ma  800  on     shareqfs1                m----2c--r-  80% 70%  /shareqfs1
mm  801  on     /dev/dsk/c6t50020F2300004655d0s0
mr  802  on     /dev/dsk/c6t50020F2300004655d0s1
ma  810  on     shareqfs2                m----2c--r-  80%  70% /shareqfs2
mm  811  on     /dev/dsk/c6t50020F2300004655d0s6
mr  812  on     /dev/dsk/c6t50020F2300004655d0s7
:
command:minallocsz 800 512
```

mm_stripe (Set the Metadata Stripe Width)

When issued with the Oracle HSM equipment ordinal number of a disk family-set device and a number of disk allocation units (DAUs) per stripe, the mm_stripe command configures the corresponding file system to write the specified number of 16-kilobyte disk allocation units (DAUs) to each logical unit number (LUN) that is available for metadata storage.

The default value, 1 DAU, writes one DAU of metadata to one LUN before switching to another LUN. A 0 (zero) value disables striping and writes metadata DAUs to metadata LUNs round-robin.

See also stripe (Set the Stripe Width for a File System) and the mount_samfs man page.

Command Synopsis

```
:mm_stripe family-set-equipment-number number_DAUs
```

Examples

The following example disables metadata striping on family set device 10:

```
File systems                    samu                    version time date

ty  eq state  device_name              status      high low  mountpoint server
ma  10 on     sam2                     -----2---r-   80% 70% /sam2      /server3
 mm 11 on      /dev/dsk/c0t600A0B800026B6F000001F0F4CCFF833d0s0
 mr 12 on      /dev/dsk/c0t600A0B800026E584000021864C4943AAd0s0
 mr 13 on      /dev/dsk/c0t600A0B800026E5840000218B4CBD9C4Ed0s0
:
command:mm_stripe 20 0
```

n (Display the Staging Status Screen)

The n command opens the Staging status screen. The screen displays the status of the stager and a list of outstanding stage requests.

When issued without a parameter, the n command displays staging status for all media types. When the command is issued with an Oracle HSM media-type identifier, the command displays status for the specified media type only. Media-type identifiers are listed in Appendix A and on the mcf man page.

Command Synopsis

```
n
:n mediatype
```

Controls

Ctrl-B (^b)
Move back one page.

Ctrl-F (^f)
Move forward one page.

Ctrl-D (^d)
Move down the page.

Ctrl-U (^u)
Move up the page.

Example

The following example shows the staging status screen.

```
Staging status                   samu                    version time date

Log output to: /var/opt/SUNWsamfs/stager/log

Stage request: li.VOL001
Loading VSN VOL001

Staging queues
ty pid   user     status   wait files  vsn
li 16097 root     active   0:00    12  VOL001
```

noabr (Disable Application Based Recovery)

When issued with the Oracle HSM equipment ordinal number of a disk family-set device, the noabr command does not let Oracle Real Application Clusters (RAC) use the Application Based Recovery feature of Solaris Volume Manager to handle recovery of databases stored in Oracle HSM file systems that use asynchronous I/O with Solaris Volume Manager mirrored volumes.

By default, ABR is enabled. See abr (Enable Application Based Recovery).

Command Synopsis

```
:noabr family-set-equipment-number
```

Examples

The following example disables ABR on family-set device **20**:

```
File systems                      samu                    version time date

ty  eq state  device_name               status      high low  mountpoint server
ma  20 on      sam2          -----2---r-   80%    70%   /sam2
 mm 21 on      /dev/dsk/c0t600A0B800026B6F000001F0F4CCFF833d0s0
 mr 22 on      /dev/dsk/c0t600A0B800026E584000021864C4943AAd0s0
 mr 23 on      /dev/dsk/c0t600A0B800026E5840000218B4CBD9C4Ed0s0
ma  10 on      sam1          m----2---r-   80%    70%   /sam1
 mm 11 on      /dev/dsk/c0t600A0B8000118ACA000009594E1565E8d0s0
 mr 12 on      /dev/dsk/c0t600A0B8000114D5400000CA64E156793d0s0
 mr 13 on      /dev/dsk/c0t600A0B8000114D5400000CAA4E156897d0s0
 mr 14 on      /dev/dsk/c0t600A0B8000118ACA0000095B4E1568C8d0s0
 mr 15 on      /dev/dsk/c0t600A0B8000118ACA0000095D4E156948d0s0
:
command:noabr 20
```

noalloc (Disable Storage Allocation on a Device)

When issued with an Oracle HSM equipment ordinal number, the noalloc command stops storage allocation on the specified data device until an administrator issues an alloc (Enable Storage Allocation on a Device) command.

Command Synopsis

```
:noalloc  equipment-number
```

Examples

The following example stops storage allocation on disk device 14:

```
File systems                      samu                    version time date

ty  eq state  device_name            status      high low  mountpoint server
ma  20 on     sam2                   -----2---r-  80% 70%  /sam/sam2  server4
 mm 21 on      /dev/dsk/c0t600A0B800026B6F000001F0F4CCFF833d0s0
 mr 22 on      /dev/dsk/c0t600A0B800026E584000021864C4943AAd0s0
 mr 23 on      /dev/dsk/c0t600A0B800026E5840000218B4CBD9C4Ed0s0
ma  10 on     sam1                   m----2---r-  80% 70%  /sam/sam1  server4
 mm 11 on      /dev/dsk/c0t600A0B8000118ACA000009594E1565E8d0s0
 mr 12 on      /dev/dsk/c0t600A0B8000114D5400000CA64E156793d0s0
 mr 13 on      /dev/dsk/c0t600A0B8000114D5400000CAA4E156897d0s0
 mr 14 on      /dev/dsk/c0t600A0B8000118ACA0000095B4E1568C8d0s0
 mr 15 on      /dev/dsk/c0t600A0B8000118ACA0000095D4E156948d0s0
:
command:noalloc 14
```

The device state changes to noalloc:

```
File systems                      samu                    version time date
```

```
ty  eq state  device_name             status      high low  mountpoint server
ma  20 on     sam2                    -----2---r-  80% 70%  /sam/sam2  server4
 mm 21 on      /dev/dsk/c0t600A0B800026B6F000001F0F4CCFF833d0s0
 mr 22 on      /dev/dsk/c0t600A0B800026E584000021864C4943AAd0s0
 mr 23 on      /dev/dsk/c0t600A0B800026E5840000218B4CBD9C4Ed0s0
ma  10 on     sam1                    m----2---r-  80% 70%  /sam/sam1  server4
 mm 11 on      /dev/dsk/c0t600A0B8000118ACA000009594E1565E8d0s0
 mr 12 on      /dev/dsk/c0t600A0B8000114D5400000CA64E156793d0s0
 mr 13 on      /dev/dsk/c0t600A0B8000114D5400000CAA4E156897d0s0
 mr 14 noalloc /dev/dsk/c0t600A0B8000118ACA0000095B4E1568C8d0s0
 mr 15 on      /dev/dsk/c0t600A0B8000118ACA0000095D4E156948d0s0
```

noatime (Disable Access Time Updates for a File System)

When issued with the Oracle HSM equipment ordinal number of a disk family-set device, the noatime command disables updates to the access time attribute on the corresponding file system. It is equivalent to the command atime -1. See atime (Configure Access Time Updates for a File System).

Never use the noatime command with an archiving, Oracle HSM file system.

nodio_szero (Do Not Zero Uninitialized Areas of Sparse Files with Direct I/O)

When issued with the Oracle HSM equipment ordinal number of a disk family-set device, the nodio_szero command makes direct I/O to the specified device skip the zeroing of uninitialized areas. This improves performance when writing sparse files directly.

The default is nodio_szero.

For further information, see dio_szero (Zero Uninitialized Areas of Sparse Files During Direct I/O).

Command Synopsis

```
:nodio_szero family-set-equipment-number
```

Examples

The following example disables the zeroing of sparse files on family set device 100:

```
Mass storage status                 samu                  version time date

ty  eq   status      use  state  ord  capacity       free ra  part high low
ma  100  m----2---r- 1%   on           7.270T     7.270T 1M  16    80% 70%
 mm  110              1%   on      0  465.253G   465.231G [975661056 inodes]
 mr  120              1%   on      1    1.817T     1.817T
 mr  130              1%   on      2    1.817T     1.817T
 mr  140              1%   on      3    1.817T     1.817T
 mr  150              1%   on      4    1.817T     1.817T
:
command:nodio_szero 100
```

nodmr (Disable Directed Mirror Reads of Software Mirrors)

When issued with the Oracle HSM equipment ordinal number of a disk family-set device, the dmr command does not let Oracle Real Application Clusters (RAC) use the Directed Mirror Reads feature of Solaris Volume Manager when databases are stored on Oracle HSM file systems that use asynchronous I/O with Solaris Volume Manager mirrored volumes. See dmr (Enable Directed Mirror Reads of Software Mirrors).

By default, DMR is enabled.

Command Synopsis

```
:nodmr family-set-equipment-number
```

Examples

The following example disables DMR on family-set device 20:

```
File systems                    samu                    version time date

ty  eq state   device_name   status        high   low   mountpoint   server
ma  20 on      sam2          -----2---r-   80%    70%   /sam2
 mm 21 on      /dev/dsk/c0t600A0B800026B6F000001F0F4CCFF833d0s0
 mr 22 on      /dev/dsk/c0t600A0B800026E584000021864C4943AAd0s0
 mr 23 on      /dev/dsk/c0t600A0B800026E5840000218B4CBD9C4Ed0s0
ma  10 on      sam1          m----2---r-   80%    70%   /sam1
 mm 11 on      /dev/dsk/c0t600A0B8000118ACA000009594E1565E8d0s0
 mr 12 on      /dev/dsk/c0t600A0B8000114D5400000CA64E156793d0s0
 mr 13 on      /dev/dsk/c0t600A0B8000114D5400000CAA4E156897d0s0
 mr 14 on      /dev/dsk/c0t600A0B8000118ACA0000095B4E1568C8d0s0
 mr 15 on      /dev/dsk/c0t600A0B8000118ACA0000095D4E156948d0s0
:
command:nodmr 20
```

noforcedirectio (Do Not Use Direct I/O by Default)

When issued with the Oracle HSM equipment ordinal number of a disk family-set device, the noforcedirectio command restores the usual default I/O method, Direct I/O, on the specified device.

For more information, see "Understanding I/O Performance Tuning Objectives and Options" on page 2-2, the forcedirectio (Use Direct I/O by Default) command, and the directio, setfa, sam_setfa, and sam_advise man pages.

Command Synopsis

```
:noforcedirectio family-set-equipment-number
```

Examples

The following example restores the default, paged I/O, on family set device 100:

```
Mass storage status                samu                    version time date

ty  eq   status      use  state  ord   capacity      free   ra  part high low
ma  100  m----2---r- 1%   on            7.270T    7.270T   1M  16    80% 70%
```

```
 mm 110              1%   on       0   465.253G  465.231G [975661056 inodes]
 mr 120              1%   on       1     1.817T    1.817T
 mr 130              1%   on       2     1.817T    1.817T
 mr 140              1%   on       3     1.817T    1.817T
 mr 150              1%   on       4     1.817T    1.817T
:
command:noforcedirectio 100
```

noforce_nfs_async (Do Not Force Asynchronous NFS)

When issued with the Oracle HSM equipment ordinal number of a disk family-set device, the noforce_nfs_async command causes the specified device to write NFS (Network File System) data through to disk unless NFS has requested that the data be cached.

Command Synopsis

```
:noforce_nfs_async family-set-equipment-number
```

Examples

The following example sets noforce_nfs_async on family set 100:

```
Mass storage status                 samu                    version time date

ty  eq  status       use  state  ord   capacity   free       ra  part high low
ma  100 m----2---r-  1%   on           7.270T   7.270T     1M  16    80% 70%
 mm  110              1%   on       0  465.253G 465.231G [975661056 inodes]
 mr  120              1%   on       1    1.817T   1.817T
 mr  130              1%   on       2    1.817T   1.817T
 mr  140              1%   on       3    1.817T   1.817T
 mr  150              1%   on       4    1.817T   1.817T
:
command:noforce_nfs_async 100
```

nohwm_archive (Do Not Start Archiving at the High Water Mark)

When issued with an Oracle HSM family-set equipment ordinal number, the nohwm_archive command configures the corresponding file system so that it does not start archiving automatically when the total size of the stored files passes the *high water mark*, the maximum space-utilization threshold specified for the file system. See maxpartial (Set the Maximum Partial Release Size for the File System).

Command Synopsis

```
:nohwm_archive equipment-number
```

Examples

The following example configures disk family-set device 20 so that it does not automatically start archiving when the corresponding file-system is 80% full:

```
File systems                        samu                    version time date
```

```
ty  eq state   device_name   status        high   low   mountpoint   server
ma  20 on      sam2          -----2---r-   80%    70%   /sam2
 mm 21 on      /dev/dsk/c0t600A0B800026B6F000001F0F4CCFF833d0s0
 mr 22 on      /dev/dsk/c0t600A0B800026E584000021864C4943AAd0s0
 mr 23 on      /dev/dsk/c0t600A0B800026E5840000218B4CBD9C4Ed0s0
:
command:nohwm_archive 20
```

nomh_write (Turn Off Multi-Host Read/Write)

When issued with the Oracle HSM equipment ordinal number of a shared disk family-set device, the nomh_write command disables simultaneous reads and writes to the same file from multiple hosts and restores the default setting for the specified device.

By default, shared Oracle HSM file systems allow simultaneous reads from multiple hosts but not writes. Use nomh_write to restore the default after running mh_write (Turn on Multi-Host Read/Write).

Command Synopsis

```
:nomh_write  family-set-equipment-number
```

Examples

The following example disables multi-host reads on family set device 800:

```
File systems                          samu                    version time date

ty  eq   state  device_name           status      high low  mountpoint server
ma  800  on     shareqfs1             m----2c--r-  80%  70% /shareqfs1
 mm  801  on     /dev/dsk/c6t50020F2300004655d0s0
 mr  802  on     /dev/dsk/c6t50020F2300004655d0s1
ma  810  on     shareqfs2             m----2c--r-  80%  70% /shareqfs2
 mm  811  on     /dev/dsk/c6t50020F2300004655d0s6
 mr  812  on     /dev/dsk/c6t50020F2300004655d0s7
:
command:mh_write 800
```

norefresh_at_eof (Do Not Refresh at EOF in a Multi-Reader File System)

When issued with the Oracle HSM equipment ordinal number of a disk family-set device that is mounted read-only in a single-writer/multi-reader file system, the norefresh_at_eof command disables automatic updates of the file size attribute when the read buffer exceeds the end of file. See refresh_at_eof (Refresh at EOF in a Multi-Reader File System).

Command Synopsis

```
:norefresh_at_eof family-set-equipment-number interval_in_seconds
```

Examples

The following example disables automatic refresh of the file size attribute on family set device 900:

```
File systems                          samu                 version time date

ty  eq   state  device_name          status       high  low  mountpoint server
ma  900  on     qfs1                 m----2--R-d   90%  70%  /qfs1
 mm  901  on     /dev/dsk/c5t10d0s0
 md  902  on     /dev/dsk/c5t10d0s1
 md  903  on     /dev/dsk/c5t11d0s0
 md  904  on     /dev/dsk/c5t11d0s1
:
command: norefresh_at_eof 900
```

noqwrite (Disable Simultaneous Reads/Writes from Different Threads)

When issued with the Oracle HSM equipment ordinal number of a disk family-set device, the noqwrite command disables simultaneous read and write operations to a single file from different threads. One writer has exclusive access to the file and other writers and readers must wait for the file to be released, per the UNIX standard.

While simultaneous multi-thread reads and writes to a single file can be useful in database applications (see qwrite (Enable Simultaneous Reads/Writes to a File from Different Threads), this capability is not generally useful or desirable. So noqwrite is the default setting.

Command Synopsis

```
:noqwrite family-set-equipment-number
```

Examples

The following example disables simultaneous read and write operations to a single file on family set device 100:

```
File systems                          samu                 version time date

ty  eq   state  device_name          status       high  low  mountpoint server
ma  100  on     qfs1                 m----2--R-d   90%  70%  /qfs1
 mm  101  on     /dev/dsk/c5t10d0s0
 md  102  on     /dev/dsk/c5t10d0s1
 md  103  on     /dev/dsk/c5t11d0s0
 md  104  on     /dev/dsk/c5t11d0s1
:
command: noqwrite 100
```

nosuid (Mount a File System with setuid Execution Disabled)

When issued with the Oracle HSM equipment ordinal number of a disk family-set device, the nosuid command mounts the file system with the setuid file attribute disabled. The setuid attribute (*set user ID upon execution*) lets file system users run executables with the temporarily elevated permissions that may be required when using the files.

See suid (Mount a File System with setuid Execution Enabled) and the mount man page for further information.

The default is suid, mount the file system with the setuid attribute enabled.

Command Synopsis

```
:nosuid family-set-equipment-number
```

Examples

The following example mounts the file system on family set device 100 with the setuid attribute disabled:

```
File systems                          samu                version time date

ty  eq   state  device_name          status       high  low  mountpoint server
ma  100  on     qfs1                 m----2--R-d   90%  70%  /qfs1
 mm  101  on     /dev/dsk/c5t10d0s0
 md  102  on     /dev/dsk/c5t10d0s1
 md  103  on     /dev/dsk/c5t11d0s0
 md  104  on     /dev/dsk/c5t11d0s1
:
command: nosuid 100
```

nosw_raid (Turn Off Software RAID Mode)

When issued with the Oracle HSM equipment ordinal number of a disk family-set device, the nosw_raid command disables software RAID mode on the specified device and restores the Oracle HSM default.

Command Synopsis

```
:nosw_raid family-set-equipment-number
```

Examples

The following example disables software-RAID mode for family set device 100:

```
Mass storage status                   samu                 version time date

ty  eq  status       use  state  ord   capacity     free     ra  part high low
ma  100 m----2---r-  1%   on           7.270T     7.270T   1M  16    80% 70%
 mm 110              1%   on       0   465.253G   465.231G [975661056 inodes]
 mr 120              1%   on       1     1.817T     1.817T
 mr 130              1%   on       2     1.817T     1.817T
 mr 140              1%   on       3     1.817T     1.817T
 mr 150              1%   on       4     1.817T     1.817T
:
command:nosw_raid 100
```

notrace (Disable File System Tracing)

When issued with the Oracle HSM equipment ordinal number of a disk family-set device, the notrace command disables tracing for the corresponding file system.

The default is trace (enabled).

For more information about file system tracing, see "trace (Enable File System Tracing)" and the defaults.conf man page.

Command Synopsis

```
:no
trace family-set-equipment-number
```

Examples

The following example disables tracing for the file system on family set device 100:

```
File systems                          samu                    version time date

ty  eq   state  device_name          status       high  low  mountpoint server
ma  100  on     qfs1                 m----2--R-d   90%  70%  /qfs1
 mm  101  on     /dev/dsk/c5t10d0s0
 md  102  on     /dev/dsk/c5t10d0s1
 md  103  on     /dev/dsk/c5t11d0s0
 md  104  on     /dev/dsk/c5t11d0s1
:
command: notrace 100
```

o (Display the Optical Disk Status Screen)

The o command opens the Optical disk status screen. The screen displays the status of all optical disk drives configured within the environment.

Command Synopsis

```
o
:o
```

Controls

Ctrl-K (^k)
Select (manual, automated library, both, priority)

Ctrl-B (^b)
Move back one page.

Ctrl-F (^f)
Move forward one page.

Ctrl-D (^d)
Move down the page.

Ctrl-U (^u)
Move up the page.

Examples

The following example shows the optical disc drive status display.

```
Optical disk status            samu                  version time date

ty  eq  status      act use  state  vsn
mo  35  --l---wo-r  1   29%  ready  OPD004
```

Field Descriptions

The following table shows the field descriptions for this display.

ty
The `ty` field displays the Oracle HSM media type. Media-type identifiers are listed in Appendix A and on the `mcf` man page.

eq
The `eq` field displays the user-assigned Oracle HSM equipment ordinal number of the optical disc drive.

status
The `status` field displays a ten-character string in which each character position represents the value of a different removable-media device attribute. Character positions are numbered starting from the left, position 1, and ending at the right, position 10. The following table defines the possible state codes for each position in the status string.

Position	Value	Meaning
1	`s---------`	A lowercase `s` in position 1 means that media is being scanned.
1	`m---------`	A lowercase `m` in position 1 means that the automated library is operational.
1	`M---------`	An uppercase `M` in position 1 means that the device is in maintenance mode.
2	`-E--------`	An uppercase `E` in position 2 means that the device encountered an unrecoverable error during scanning.
2	`-a--------`	A lowercase `a` in position 2 means that the device is in audit mode.
3	`--l-------`	A lowercase `l` in position 3 means that media has a label.
3	`--L-------`	An uppercase `L` in position 3 means that the media is being labeled
3	`--N-------`	An uppercase `N` in position 3 means that the media is foreign (not part of the Oracle HSM archive).
4	`---I------`	An uppercase `I` in position 4 means that the device is idle and not available for new connections. But operations already in progress are continuing until completion.
4	`---A------`	An uppercase `A` in position 4 means that the device needs operator attention.
5	`----C-----`	An uppercase `C` in position 5 means that the device needs cleaning.
5	`----U-----`	An uppercase `U` in position 5 means that an unload has been requested.

Position	Value	Meaning
6	-----R----	An uppercase R in position 6 means that the device is reserved.
7	------w---	A lowercase w in position 7 means that a process is writing to the media.
8	-------o--	A lowercase o in position 8 means that the device is open.
9	--------P-	An uppercase P in position 9 means that the device is positioning a tape volume.
9	--------F-	For a tape library, an uppercase F in position 9 means that all storage slots are occupied. For a drive, an uppercase F in position 9 means that the currently mounted removable media is full.
10	---------R	An uppercase R in position 10 means that the device is ready and the media is read-only.
10	---------r	A lowercase r in position 10 means that the device is spun up and ready.
10	---------p	A lowercase p in position 10 means that the device is present.
10	---------W	An uppercase W in position 10 means that the device is write protected.

act
The act field displays the activity count.

use
The use field displays the percentage of cartridge space used.

state
The state field displays the current operating state of the optical disc drive:

- ready means that the device is on and media are loaded.
- notrdy means that the device is on, but media are not loaded.
- idle means that the device is completing operations that are already in progress but is not accepting new connections.
- off The device is not available.
- down means that the device is available for maintenance only.

vsn
The vsn field displays either the volume serial number assigned to the media volume or the keyword nolabel if the volume is not labeled.

off (Turn Off A Device)

The off command sets the device specified by an Oracle HSM equipment ordinal number offline. All operations stop.

Command Synopsis

```
:off equipment-number
```

Examples

```
Device configuration:            samu                   version time date

ty  eq  state  device_name                            fs   family_set
sk  100 on     /etc/opt/SUNWsamfs/SL8500_T10K_parms  100  libraryA
ti  101 on      /dev/rmt/0cbn                         100  libraryA
ti  102 on      /dev/rmt/1cbn                         100  libraryA
ti  103 on      /dev/rmt/2cbn                         100  libraryA
hy  104 on      historian                             104
:
command:off 103
```

The device state changes to `off`:

```
Device configuration:            samu                   version time date

ty   eq state   device_name                            fs   family_set
sk  100 on      /etc/opt/SUNWsamfs/SL8500_T10K_parms  100  libraryA
ti  101 on      /dev/rmt/0cbn                         100  libraryA
ti  102 on      /dev/rmt/1cbn                         100  libraryA
ti  103 off     /dev/rmt/2cbn                         100  libraryA
hy  104 on      historian                             104
```

on (Turn On A Device)

When issued with the Oracle HSM equipment ordinal number of a device, the `on` command brings the device specified by an Oracle HSM equipment ordinal number online.

Command Synopsis

```
:on equipment-number
```

Examples

```
Device configuration:            samu                   version time date

ty  eq  state  device_name                            fs   family_set
sk  100  on    /etc/opt/SUNWsamfs/SL8500_T10K_parms  100  libraryA
ti  101  on    /dev/rmt/0cbn                         100  libraryA
ti  102  on    /dev/rmt/1cbn                         100  libraryA
ti  103  off   /dev/rmt/2cbn                         100  libraryA
hy  104 on     historian                             104
:
command:on 103
```

The device state changes to `on`:

```
Device configuration:            samu                   version time date

ty  eq  state  device_name                            fs   family_set
sk  100  on    /etc/opt/SUNWsamfs/SL8500_T10K_parms  100  libraryA
ti  101  on    /dev/rmt/0cbn                         100  libraryA
ti  102  on    /dev/rmt/1cbn                         100  libraryA
ti  103  on    /dev/rmt/2cbn                         100  libraryA
hy  104 on     historian                             104
```

open (Open a Disk Storage Device for Reading)

When issued with the Oracle HSM equipment ordinal number of a disk device, the open command opens the device for the read (Read Sectors On a Disk Storage Device) command.

Command Synopsis

```
:open equipment-number
```

Examples

The following example opens device 23:

```
File systems                samu                    version time date

ty      eq   state  device_name     status      high low mountpoint server
ma      20      on  sam2            -----2---r-  80% 70% /sam/sam2
 mm      21      on  /dev/dsk/c0t600A0B800026B6F000001F0F4CCFF833d0s0
 mr      22      on  /dev/dsk/c0t600A0B800026E584000021864C4943AAd0s0
 mr      23      on  /dev/dsk/c0t600A0B800026E5840000218B4CBD9C4Ed0s0
:
command:open 23
```

The screen switches to displaying disk sectors, starting at sector 0:

```
Sector:   00000000 (0)    raw       samu                  version time date

00000000  00000000 00000000 00000000 00000000  ................
00000010  00000000 00000000 00000000 00000000  ................
00000020  00000000 00000000 00000000 00000000  ................
00000030  00000000 00000000 00000000 00000000  ................
00000040  00000000 00000000 00000000 00000000  ................
00000050  00000000 00000000 00000000 00000000  ................
00000060  00000000 00000000 00000000 00000000  ................
00000070  00000000 00000000 00000000 00000000  ................
 ...
```

p (Display the Removable Media Load Requests Screen)

The p command opens the Removable media load requests screen. The screen displays information about pending load requests for removable media. Four different screen formats are available:

- Both manual and automated library requests sorted by user
- Both manual and automated library requests sorted by priority.
- Manual requests only
- Automated library requests only.

When issued without a parameter, the p command displays load requests for all media. When the command is issued with an Oracle HSM media-type identifier, the command displays load requests for the specified media type only. Media-type identifiers are listed in Appendix A and on the mcf man page.

Command Synopsis

```
p
:p mediatype
```

Controls

You can use the following control keys on this screen:

Ctrl-K (^k)
Toggle between the different display formats.

Ctrl-B (^b)
Move back one page.

Ctrl-F (^f)
Move forward one page.

Ctrl-D (^d)
Move down the page.

Ctrl-U (^u)
Move up the page.

Examples

The following example shows the removable media load requests display.

```
Removable media load requests all both    samu                 version time date
  count: 1

index type pid    user     rb   flags    wait count vsn
    0 li   15533  root     150  W--f---  0:00       VOL002
```

Field Descriptions

index
The index field displays the corresponding request's position in the preview table.

type
The type field displays the media type code assigned to the removable media.

pid
The pid field displays the UNIX process identifier. A process identifier of 1 indicates NFS access.

user
The user field displays the name assigned to the user requesting the load.

priority
The priority field displays the degree to which the request may take precedence over other requests.

rb
The rb field displays the user-assigned equipment ordinal number of the automated library in which the requested VSN resides.

flags
The flags field displays a seven-character strings in which each character position represents the value of a property of the request. Character positions are numbered starting from the left, position 1, and ending at the right, position 7. The following table defines the possible flag codes for each position in the status string.

Position	Flag	Meaning
1	W------	An uppercase W in position 1 means that write access has been requested.
2	-b-----	A lowercase b in position 2 means that the request is busy.
3	--C----	An uppercase C in position 3 means that the request for the VSN has been cleared.
4	---f---	A lowercase f in position 4 means that file-system access has been requested.
5	----N--	An uppercase N in position 5 means that the requested media is foreign to the Oracle HSM file system.
6	-----S-	An uppercase S in position 6 means that the flip side of a two-sided optical-media volume is mounted.
7	------s	A lowercase s in position 7 means that a file is to be staged from removable media to the file system.

wait
The wait field displays the elapsed time since the mount request was received.

count
The count field displays the number of requests for this VSN, if a file is being staged.

vsn
The vsn field displays the volume serial name of the removable media volume.

partial (Set the Default Partial Release Size for the File System)

When issued with an Oracle HSM family-set equipment ordinal number and a number of kilobytes, the partial command sets the default number of bytes from the start of a file that can be retained in disk cache when the file is released.

A large file that has been designated for partial release can open more quickly when restaged from archival media because the beginning of the file is already on disk when staging begins. Users can override the default on a file-by-file basis by specifying a size when marking a file for partial release. For further information, see maxpartial (Set the Maximum Partial Release Size for the File System), partial_stage (Set Partial Stage-Ahead Point for the File System), and the release man page.

The specified kilobyte value must be an integer in the range [8-*maximumpartial*], where *maximumpartial* is the value set by the maxpartial command. The default is 16 kilobytes.

Command Synopsis

```
:partial equipment-number number_of_kilobytes
```

Examples

The following example configures disk family-set device 20 to retain 256 kilobytes on disk for each file designated for partial release:

```
File systems                          samu                  version time date

ty  eq state   device_name   status        high   low   mountpoint   server
ma  20 on      sam2          -----2---r-   80%    70%   /sam2
 mm 21 on      /dev/dsk/c0t600A0B800026B6F000001F0F4CCFF833d0s0
 mr 22 on      /dev/dsk/c0t600A0B800026E584000021864C4943AAd0s0
 mr 23 on      /dev/dsk/c0t600A0B800026E5840000218B4CBD9C4Ed0s0
:
command:partial 20 256
```

partial_stage (Set Partial Stage-Ahead Point for the File System)

When issued with an Oracle HSM family-set equipment ordinal number and a number of kilobytes, the partial_stage command defines a file offset beyond which a file designated for partial release is staged to disk in its entirety rather than in part.

Partially staging a large file can give users access to data more quickly because the file can open before it is fully staged from archival media to the disk cache. But, once enough data is in cache to start opening the file, there is no further advantage to partial staging. For further information, see maxpartial (Set the Maximum Partial Release Size for the File System), partial (Set the Default Partial Release Size for the File System), and the release man page.

The specified kilobyte value must be an integer in the range [8-*maximumpartial*], where *maximumpartial* is the value set by the maxpartial command. The default is the value set by the partial command.

Command Synopsis

```
:partial_stage equipment-number number_of_kilobytes
```

Examples

The following example sets the partial stage-ahead point for disk family-set device 20 to 5124 kilobytes (5 megabytes):

```
File systems                          samu                  version time date

ty  eq state   device_name   status        high   low   mountpoint   server
ma  20 on      sam2          -----2---r-   80%    70%   /sam2
 mm 21 on      /dev/dsk/c0t600A0B800026B6F000001F0F4CCFF833d0s0
 mr 22 on      /dev/dsk/c0t600A0B800026E584000021864C4943AAd0s0
 mr 23 on      /dev/dsk/c0t600A0B800026E5840000218B4CBD9C4Ed0s0
:
command:partial_stage 20 5124
```

priority (Set Priority in Load-Request Preview Queue)

When issued with the process identifier (pid) of a load request and an integer value, the priority command moves the load request to the specified position in the list of waiting jobs.

Specify the pid for one of the load requests listed by the p (Display the Removable Media Load Requests Screen) command. Specify the desired priority as an integer, with 0 (zero) being first in line (the highest priority).

Command Synopsis

```
:priority pid priority
```

Examples

The following example gives the load request with pid 17708 (for volume VOL202) a priority of 0:

```
Removable media load requests all both   samu                version time date
                                                            count: 4
index  type pid    user  rb   flags     wait count  vsn
    0  li   13850  root  900  W--f---   0:42        VOL002
    1  li   15533  root  900  ---f---   0:28        VOL001
    2  li   17708  root  900  ---f---   0:28        VOL202
    3  li   19937  root  900  ---f---   0:05        VOL015
:
command:priority 17708 0
```

With its new priority, pid 17708 moves from third in the queue (index 2) to first (index 0):

```
index  type pid    user  rb   flags     wait count  vsn
    0  li   17708  root  900  ---f---   0:28        VOL202
    1  li  15533   root  900  ---f---   0:28        VOL001
    2  li   15533  root  900  ---f---   0:28        VOL001
    3  li   19937  root  900  ---f---   0:05        VOL015
```

qwrite (Enable Simultaneous Reads/Writes to a File from Different Threads)

When issued with the Oracle HSM equipment ordinal number of a disk family-set device, the qwrite command enables simultaneous read and write operations to a single file from different threads.

In databases and other applications that handle multiple simultaneous transactions to the same file, simultaneous reads and writes improve I/O performance by queuing multiple requests at the drive level.

The default is noqwrite (Disable Simultaneous Reads/Writes from Different Threads).

Command Synopsis

```
:qwrite family-set-equipment-number
```

Examples

The following example enables simultaneous read and write operations to a single file on family set device 100:

```
File systems                    samu                    version time date
```

```
ty  eq   state  device_name           status       high low   mountpoint server
ma  100  on     qfs1                  m----2----d   90% 70%   /qfs1
 mm  101  on     /dev/dsk/c5t10d0s0
 md  102  on     /dev/dsk/c5t10d0s1
 md  103  on     /dev/dsk/c5t11d0s0
 md  104  on     /dev/dsk/c5t11d0s1
:
command: qwrite 100
```

r (Display the Removable Media Status Screen

The r command opens the Removable media status screen. The screen displays the activity on removable media devices such as tape drives. You can monitor either a specific type of device or a family of devices such as all tape devices.

When issued without a parameter, the r command displays the status for all removable media devices. When the command is issued with an Oracle HSM media type identifier, the command displays the status of the specified media type only. Media-type identifiers are listed in Appendix A and on the mcf man page.

Command Synopsis

```
r
:r mediatype
```

Examples

The following example shows the removable media status display.

```
Removable media status: all         samu                    version time date

ty  eq  status     act use state  vsn
li  150 --l------r   0 63% ready  VOL004
```

Field Descriptions

The following table shows the field descriptions for this display.

ty
The ty field displays Oracle HSM media type. Device-type identifiers are listed in Appendix A and on the mcf man page.

eq
The eq field displays the user-assigned Oracle HSM equipment ordinal number of the mass storage device.

status
The status field displays a ten-character string in which each character position represents the value of a different removable-media device attribute. Character positions are numbered starting from the left, position 1, and ending at the right, position 10. The following table defines the possible state codes for each position in the status string.

Position	Value	Meaning
1	s---------	A lowercase s in position 1 means that media is being scanned.
1	m---------	A lowercase m in position 1 means that the automated library is operational.
1	M---------	An uppercase M in position 1 means that the device is in maintenance mode.
2	-E--------	An uppercase E in position 2 means that the device encountered an unrecoverable error during scanning.
2	-a--------	A lowercase a in position 2 means that the device is in audit mode.
3	--l-------	A lowercase l in position 3 means that media has a label.
3	--L-------	An uppercase L in position 3 means that the media is being labeled
3	--N-------	An uppercase N in position 3 means that the media is foreign (not part of the Oracle HSM archive).
4	---I------	An uppercase I in position 4 means that the device is idle and not available for new connections. But operations already in progress are continuing until completion.
4	---A------	An uppercase A in position 4 means that the device needs operator attention.
5	----C-----	An uppercase C in position 5 means that the device needs cleaning.
5	----U-----	An uppercase U in position 5 means that an unload has been requested.
6	-----R----	An uppercase R in position 6 means that the device is reserved.
7	------w---	A lowercase w in position 7 means that a process is writing to the media.
8	-------o--	A lowercase o in position 8 means that the device is open.
9	--------P-	An uppercase P in position 9 means that the device is positioning a tape volume.
9	--------F-	For a tape library, an uppercase F in position 9 means that all storage slots are occupied. For a drive, an uppercase F in position 9 means that the currently mounted removable media is full.
10	---------R	An uppercase R in position 10 means that the device is ready and the media is read-only.
10	---------r	A lowercase r in position 10 means that the device is spun up and ready.
10	---------p	A lowercase p in position 10 means that the device is present.
10	---------W	An uppercase W in position 10 means that the device is write protected.

rdlease (Set the Read Lease Time for a Shared File System)

Before reading a shared file, a host must request and obtain an *read lease* that grants it exclusive access to the file. The rdlease command defines the number of seconds during which this lease remains valid on the specified device. If the specified time runs out before the read operation is complete, the host must reacquire the lease before proceeding.

The specified number of seconds must fall in the range [15-600]. The default is 30 seconds.

For related information, see aplease (Set Append Lease Expiration Time for a Shared File System), wrlease (Set the Write Lease Time for a Shared File System), and lease_timeo (Set the Unused Lease Timeout for a Shared File System).

Command Synopsis

```
:rdlease family-set-equipment-number interval_in_seconds
```

Examples

The following example sets the read lease time to 2 minutes (120 seconds) on family set device 800:

```
File systems                      samu                    version time date

ty  eq   state  device_name           status      high low mountpoint server
ma  800  on     shareqfs1             m----2c--r-  80% 70% /shareqfs1
 mm  801  on     /dev/dsk/c6t50020F2300004655d0s0
 mr  802  on     /dev/dsk/c6t50020F2300004655d0s1
ma  810  on     shareqfs2            m----2c--r-   80% 70% /shareqfs2
 mm  811  on     /dev/dsk/c6t50020F2300004655d0s6
 mr  812  on     /dev/dsk/c6t50020F2300004655d0s7
:
command:rdlease 800 120
```

read (Read Sectors On a Disk Storage Device)

When issued with a hexadecimal sector address on a previously opened disk device, the read command displays disk sectors starting at the specified address.

See also open (Open a Disk Storage Device for Reading).

Command Synopsis

```
:read hexadecimal-address
```

Examples

The following example reads device 23 starting at sector 1f4 (500):

```
Sector:   00000000 (0)        raw       samu       version time date

00000000  00000000 00000000 00000000 00000000  ................
00000010  00000000 00000000 00000000 00000000  ................
00000020  00000000 00000000 00000000 00000000  ................
00000030  00000000 00000000 00000000 00000000  ................
00000040  00000000 00000000 00000000 00000000  ................
00000050  00000000 00000000 00000000 00000000  ................
00000060  00000000 00000000 00000000 00000000  ................
00000070  00000000 00000000 00000000 00000000  ................
 ...
:
command:read 1f4
```

The screen displaying disk sectors starting at sector 1f4:

```
Sector:   000001f4 (500)      raw      samu       version time date
```

```
00000000  0a01c1c1 e39ce073 8e00e319 39e80195  ..AAc.`s..c.9h..
00000010  a40cd199 bfd5285f ba1509f9 41c03820  $.Q.?U(_:..yA@8
00000020  90780583 7f78825b 9fa0a58b 4a519c5d  .x...x.[. %.JQ.]
00000030  d2774afa bb256497 569bd2fb b775ddf2  RwJz;%d.V.R{7u]r
00000040  ba5cd1b7 2a56d93d 6e9dafab ba8b5e56  :\Q7*VY=n./+:.^V
00000050  d5686a5b 334d1f96 70060ec1 b46483c8  Uhj[3M..p..A4d.H
00000060  190339e7 3ce3e5e7 a702ff00 9394528c  ..9g<ceg'.....R.
00000070  4c9f3990 6dc05c1e b907e6dd 918e14e4  L.9.m@\.9.f]...d
...
```

readahead (Set The Maximum Read-Ahead Parameter)

When issued with the Oracle HSM equipment ordinal number of a disk family-set device, the readahead command specifies the minimum number of contiguous bytes that the device should read into memory. When correctly configured, read-ahead improves file-system performance by reducing disk I/O and making larger, contiguous blocks of information immediately available in memory.

Specify a read-ahead value in the range [1-8192] kilobytes. The default is 1024 kilobytes (1 megabyte).

You can also set a readahead directive in the samfs.cmd file. See the samfs.cmd man page for details.

Command Synopsis

```
:readahead family-set-equipment-number number_kilobytes
```

Examples

The following example sets the read-ahead value for family set device 100 to 1 megabyte (1024 kilobytes):

```
Mass storage status                samu                   version time date

ty  eq   status       use  state  ord   capacity     free   ra    part high low
ma  100  m----2---r-  1%   on           7.270T     7.270T 128K  16    80% 70%
 mm  110               1%   on      0   465.253G   465.231G [975661056 inodes]
 mr  120               1%   on      1     1.817T     1.817T
 mr  130               1%   on      2     1.817T     1.817T
 mr  140               1%   on      3     1.817T     1.817T
 mr  150               1%   on      4     1.817T     1.817T
:
command:readahead 100 1024
```

The read-ahead value changes to 1 megabyte.

```
Mass storage status                samu                   version time date

ty  eq   status       use  state  ord   capacity     free     ra  part high low
ma  100  m----2---r-  1%   on           7.270T     7.270T   1M  16    80% 70%
 mm  110               1%   on      0   465.253G   465.231G [975661056 inodes]
 mr  120               1%   on      1     1.817T     1.817T
 mr  130               1%   on      2     1.817T     1.817T
 mr  140               1%   on      3     1.817T     1.817T
 mr  150               1%   on      4     1.817T     1.817T
```

readonly | ro (Make Device Read Only)

The readonly or ro command applies the read-only attribute to the device specified by an Oracle HSM equipment ordinal number.

Command Synopsis

```
:readonly device-equipment-number
:ro device-equipment-number
```

Examples

The following example makes device 12 read-only:

```
Mass storage status                samu                version time date

ty      eq  status        use state  ord  capacity      free    ra  part high low
ma      10  m----2---r-    1% on           7.270T    7.270T    1M    16  80% 70%
 mm      11                1% on       0  465.253G  465.231G [975661056 inodes]
 mr      12                1% on       1    1.817T    1.817T
 mr      13                1% on       2    1.817T    1.817T
 mr      14                1% on       3    1.817T    1.817T
 mr      15                1% on       4    1.817T    1.817T
:
command:readonly 12
```

refresh (Set the Automatic Refresh Interval for the samu Display)

When issued with a number of seconds, the refresh command configures samu to refresh screens after the specified number of seconds has passed.

Command Synopsis

```
:refresh interval_in_seconds
```

Examples

The following example configures samu to refresh the display every 20 seconds:

```
File systems                       samu                   version time date

ty  eq   state device_name          status       high low   mountpoint server
ma  900  on    qfs1                 m----2--R-d   90%  70%  /qfs1
 mm  901  on    /dev/dsk/c5t10d0s0
 md  902  on    /dev/dsk/c5t10d0s1
 md  903  on    /dev/dsk/c5t11d0s0
 md  904  on    /dev/dsk/c5t11d0s1
:
command: refresh 20
```

refresh_at_eof (Refresh at EOF in a Multi-Reader File System)

When issued with the Oracle HSM equipment ordinal number of a disk family-set device that is mounted read-only in a single-writer/multi-reader file system, the

refresh_at_eof command forces an update of the file size attribute whenever the read buffer exceeds the end of file. See norefresh_at_eof (Do Not Refresh at EOF in a Multi-Reader File System).

Command Synopsis

```
:refresh_at_eof family-set-equipment-number interval_in_seconds
```

Examples

The following example enables automatic refresh of the file size attribute on family set device 900:

```
File systems                          samu                    version time date

ty  eq   state  device_name           status       high low   mountpoint server
ma  900  on     qfs1                  m----2--R-d   90%  70%  /qfs1
 mm  901  on    /dev/dsk/c5t10d0s0
 md  902  on    /dev/dsk/c5t10d0s1
 md  903  on    /dev/dsk/c5t11d0s0
 md  904  on    /dev/dsk/c5t11d0s1
:
command: refresh_at_eof 900
```

release (Release Archived Data Files and Disable Allocation on a Device)

When issued with an Oracle HSM equipment ordinal number, the release command releases all archived files that are stored on the device, including any space associated with files otherwise designated for partial release, and turns the specified device off.

The release command sets the specified device to the noalloc (Disable Storage Allocation on a Device) state and has the sam-fsd daemon run a sub-process, sam-shrink. The sam-shrink process places all archived files that reside on the specified device in the offline state, so that the data are available from archival media only. When the process completes successfully, the specified device is left in the off state. If the process encounters any unarchived files or files that are in the process of being archived or staged back from archive to disk, the process fails and the specified equipment remains in the noalloc state.

The release process can take a significant amount of time. Monitor its progress via samu and/or via the /var/opt/SUNWsamfs/trace/sam-shrink file and the log file configured in /etc/opt/SUNWsamfs/shrink.cmd (if any).

Note that both remove (Move Data and Remove Equipment from a File System) and release shrink a file system by taking disk devices from a family set without losing data. They differ because the remove command has to consolidate files on the remaining disk devices in the set before freeing the specified device. The release command does not: in an archiving file system, all data is already duplicated on removable media. See add (Add Equipment to a Mounted File System), remove (Move Data and Remove Equipment from a File System), and the sam-shrink man page for related information.

Command Synopsis

```
:release equipment-number
```

Examples

The following example, releases all files on disk slice /dev/dsk/c5t8d0s8 and shrinks family set samfs4:

```
File systems                          samu                    version time date

ty  eq  state   device_name   status        high   low   mountpoint   server
ma  20  on      samfs4        -----2---r-   80%    70%
 mm  21  on      /dev/dsk/c5t8d0s6
 md  22  on      /dev/dsk/c5t8d0s7
 md  23  on      /dev/dsk/c5t8d0s8
:
command:release 23
```

The device state changes to noalloc:

```
 md 23 noalloc /dev/dsk/c5t8d0s8
```

Finally, when the process completes successfully, the device state changes to off:

```
 md 23 off     /dev/dsk/c5t8d0s8
```

remove (Move Data and Remove Equipment from a File System)

When issued with an Oracle HSM equipment ordinal number, the remove command moves data from the specified device to the remaining devices in the file system and then turns the specified device off.

The remove command sets the specified device to the noalloc (Disable Storage Allocation on a Device) state and has the sam-fsd daemon run a sub-process, sam-shrink, to carry out the actual work. The sam-shrink process moves data from the specified equipment to the file system's other, on storage devices (if the specified equipment is a stripe group, there has to be another stripe group with the same number of devices). When the process completes successfully, the specified device is left in the off state. If data could not be moved for any reason, the equipment is not removed and remains in the noalloc state.

The removal process can take a significant amount of time. Monitor its progress via samu and/or via the /var/opt/SUNWsamfs/trace/sam-shrink file and the log file configured in /etc/opt/SUNWsamfs/shrink.cmd (if any).

Note that both release (Release Archived Data Files and Disable Allocation on a Device) and remove shrink a file system by taking disk devices from a family set without losing data. They differ because the remove command has to consolidate files on the remaining disk devices in the set before freeing the specified device. The release command does not: in an archiving file system, all data is already duplicated on removable media. See add (Add Equipment to a Mounted File System), and the sam-shrink man page for related information.

Command Synopsis

```
:remove equipment-number
```

Examples

The following example moves all files from disk slice /dev/dsk/c5t8d0s3 and shrinks family set qfs1:

```
File systems                          samu                  version time date

ty  eq state   device_name   status        high   low   mountpoint   server
ma  10 on      qfs1          -----2---r-   80%    70%
 mm  11 on      /dev/dsk/c5t8d0s0
 md  12 on      /dev/dsk/c5t8d0s1
 md  13 on      /dev/dsk/c5t8d0s3
:
command:remove 13
```

The device state changes to noalloc:

```
 md 13 noalloc /dev/dsk/c5t8d0s3
```

Finally, when the process completes successfully, the device state changes to off:

```
 md 13 off      /dev/dsk/c5t8d0s3
```

s (Display Device Status Screen)

The s command opens the Device status screen. The screen displays the status for all devices configured within the environment.

Command Synopsis

```
s
:s
```

Controls

You can use the following control keys on this screen:

Ctrl-B (^b)
Move back one page.

Ctrl-F (^f)
Move forward one page.

Ctrl-D (^d)
Move down the page.

Ctrl-U (^u)
Move up the page.

Examples

The following example shows the device status display.

```
Device status                   samu                   version time date

ty  eq    state  device_name                    fs  status
sk  100   on     /etc/opt/SUNWsamfs/dcstkconf    100 m--------r
        stk_dismount / (2275) 0, volser VOL073
```

```
sg  120   on     /dev/rmt/2cbn                  100 ---------p
sg  130   on     /dev/rmt/5cbn                  100 --l----o-r
        Ready for data transfer
sg  140   on     /dev/rmt/6cbn                  100 ---------p
sg  150   on     /dev/rmt/4cbn                  100 ---------p
hy  151   on     historian                      151 ----------
```

Field Descriptions

ty
The ty field displays Oracle HSM media type. Device-type identifiers are listed in Appendix A and on the mcf man page.

eq
The eq displays the user-assigned Oracle HSM equipment ordinal number of the mass storage device.

state
The state field displays the current operating state of the device.

device_name
The device_name field displays the path to the device. For file system devices, this is the file system name.

fs
The fs field displays the user-assigned Oracle HSM equipment number of the family set to which the device belongs.

status
The status field displays a ten-character string in which each character position represents the value of a different removable-media device attribute. Character positions are numbered starting from the left, position 1, and ending at the right, position 10. The following table defines the possible state codes for each position in the status string.

Position	Value	Meaning
1	s---------	A lowercase s in position 1 means that media is being scanned.
1	m---------	A lowercase m in position 1 means that the automated library is operational.
1	M---------	An uppercase M in position 1 means that the device is in maintenance mode.
2	-E--------	An uppercase E in position 2 means that the device encountered an unrecoverable error during scanning.
2	-a--------	A lowercase a in position 2 means that the device is in audit mode.
3	--l-------	A lowercase l in position 3 means that media has a label.
3	--L-------	An uppercase L in position 3 means that the media is being labeled
3	--N-------	An uppercase N in position 3 means that the media is foreign (not part of the Oracle HSM archive).
4	---I------	An uppercase I in position 4 means that the device is idle and not available for new connections. But operations already in progress are continuing until completion.

Position	Value	Meaning
4	---A------	An uppercase A in position 4 means that the device needs operator attention.
5	----C-----	An uppercase C in position 5 means that the device needs cleaning.
5	----U-----	An uppercase U in position 5 means that an unload has been requested.
6	-----R----	An uppercase R in position 6 means that the device is reserved.
7	------w---	A lowercase w in position 7 means that a process is writing to the media.
8	-------o--	A lowercase o in position 8 means that the device is open.
9	--------P-	An uppercase P in position 9 means that the device is positioning a tape volume.
9	--------F-	For a tape library, an uppercase F in position 9 means that all storage slots are occupied. For a drive, an uppercase F in position 9 means that the currently mounted removable media is full.
10	---------R	An uppercase R in position 10 means that the device is ready and the media is read-only.
10	---------r	A lowercase r in position 10 means that the device is spun up and ready.
10	---------p	A lowercase p in position 10 means that the device is present.
10	---------W	An uppercase W in position 10 means that the device is write protected.

snap (Save a Snapshot of the Current Screen to a File)

When issued with a file path and name, the snap command saves the current screen as a text file with the specified name.

Command Synopsis

```
:snap file-path-and-name
```

Examples

The following example saves a snap shot of the File systems screen to the file /export/home/samadmin/filsysscrn.txt:

```
File systems                          samu                    version time date

ty  eq  state   device_name   status        high   low   mountpoint   server
ma  20  on      sam2          -----2---r-   80%    70%
 mm  21  on      /dev/dsk/c0t600A0B800026B6F000001F0F4CCFF833d0s0
 mr  22  on      /dev/dsk/c0t600A0B800026E584000021864C4943AAd0s0
 mr  23  on      /dev/dsk/c0t600A0B800026E5840000218B4CBD9C4Ed0s0
:
command:snap /export/home/samadmin/filsysscrn.txt
```

stage_flush_behind (Set the Stage Flush Behind Size)

When issued with the Oracle HSM equipment ordinal number of a disk family-set device and a number of kilobytes, the stage_flush_behind command configures the

corresponding file system to asynchronously write the specified number of kilobytes from dirty (modified) stage memory pages to disk.

As staged files are copied from archival media into memory, flush-behind insures that the staged data are quickly saved to the file system's non-volatile primary disk cache. Unsaved, *dirty* cache pages are kept to a minimum during staging, which helps to maintain the integrity of the file data.

On the other hand, flush-behind increases overhead and can reduce I/O performance. The process of flushing a data block may delay subsequent re-writes of that block. So flush-behind is disabled by default.

Set the specified value in the range [0-8192] kilobytes, where 0—the default—disables the flush-behind capability.

Command Synopsis

```
:stage_flush_behind equipment-number number_of_kilobytes
```

Examples

The following example sets the stage_flush_behind for disk family-set device 20 to 5124 kilobytes (5 megabytes):

```
File systems                          samu                    version time date

ty  eq  state   device_name   status        high   low   mountpoint   server
ma  20  on      sam2          -----2---r-   80%    70%
 mm  21  on      /dev/dsk/c0t600A0B800026B6F000001F0F4CCFF833d0s0
 mr  22  on      /dev/dsk/c0t600A0B800026E584000021864C4943AAd0s0
 mr  23  on      /dev/dsk/c0t600A0B800026E5840000218B4CBD9C4Ed0s0
:
command:stage_flush_behind 20 5124
```

stage_n_window (Size the Buffer for Reading Directly from Archival Media)

When issued with the Oracle HSM equipment ordinal number of a disk family-set device and a number of kilobytes, the stage_n_window command sets the buffer size that Oracle HSM uses when reading files from the corresponding file system that are configured to never automatically stage to disk. Oracle HSM reads such files directly from the archival media to memory in chunks of the specified number of kilobytes.

Specify a number of kilobytes in the range [64-2097152]. The default is 8192 kilobytes.

Command Synopsis

```
:stage_n_window family-set-equipment-number number_of_kilobytes
```

Examples

The following example sets the stage_n_window for disk family-set device 20 to 16384 kilobytes (16 megabytes):

```
File systems                          samu                    version time date

ty  eq  state   device_name   status        high   low   mountpoint   server
```

```
ma  20  on      sam2          -----2---r-   80%    70%
 mm  21  on      /dev/dsk/c0t600A0B800026B6F000001F0F4CCFF833d0s0
 mr  22  on      /dev/dsk/c0t600A0B800026E584000021864C4943AAd0s0
 mr  23  on      /dev/dsk/c0t600A0B800026E5840000218B4CBD9C4Ed0s0
:
command:stage_n_window 20 16384
```

stclear (Clear a Stage Request)

When entered with an Oracle HSM media-type identifier and the volume serial number (VSN) of a removable media cartridge, the stclear command removes the corresponding staging request from the staging queue. Media-type identifiers are listed in Appendix A and on the mcf man page.

Command Synopsis

```
:stclear mediatype.volume-serial-number
```

Examples

In the example, stclear clears the staging request for the LTO (li) volume with VSN VOL001 from the staging queue:

```
Staging queue by media type: all       samu                    version time date

ty  length   fseq    ino    position   offset  vsn
li  451.611k   20    1030      207cc      473  VOL001
li  341.676k   20    1031      207cc       7fc  VOL002
li  419.861k   20    1032      207cc       aa9  VOL003
li  384.760k   20    1033      207cc       df2  VOL004
li  263.475k   20    1034      207cc     10f5  VOL005
li  452.901k   20    1035      207cc     1305  VOL006
li  404.598k   20    1036      207cc     1690  VOL007
:
command:stclear li.VOL001
```

The staging request for the LTO (li) volume VOL001 has been cleared:

```
Staging queue by media type: all       samu                    version time date

ty  length   fseq    ino    position   offset  vsn
li  341.676k   20    1031      207cc       7fc  VOL002
li  419.861k   20    1032      207cc       aa9  VOL003
li  384.760k   20    1033      207cc       df2  VOL004
li  263.475k   20    1034      207cc     10f5  VOL005
li  452.901k   20    1035      207cc     1305  VOL006
li  404.598k   20    1036      207cc     1690  VOL007
```

stidle (Stop Staging as Soon as Possible)

The stidle command idles the stager as soon as current staging operations are complete.

Command Synopsis

```
:stidle
```

Examples

In the example, stidle idles staging:

```
Staging queue by media type: all       samu                    version time date

ty  length   fseq    ino    position   offset vsn
li  451.611k   20    1030      207cc      473 VOL001
li  341.676k   20    1031      207cc      7fc VOL002
li  419.861k   20    1032      207cc      aa9 VOL003
li  384.760k   20    1033      207cc      df2 VOL004
li  263.475k   20    1034      207cc     10f5 VOL005
li  452.901k   20    1035      207cc     1305 VOL006
li  404.598k   20    1036      207cc     1690 VOL007
:
command:stidle
```

stripe (Set the Stripe Width for a File System)

When issued with the Oracle HSM equipment ordinal number of a disk family-set device and a number of disk allocation units (DAUs), the stripe command configures the corresponding file system to write the specified number of DAUs to each logical unit number (LUN) that is available for data storage.

Specify a number of DAUs in the range [0-255], where a 0 (zero) value disables striping and writes DAUs to LUNs round-robin.

For file systems of the ms equipment type and file systems of the ma type that do not include striped group equipment (g*XXX*), the default value is 1 DAU with a minimum DAU size of 128 kilobytes. For shared file systems and for ma file systems that include striped groups, the default is 0.

The system also sets stripe width to 0 automatically if a file system contains mismatched striped groups.

For more information about file system types, see the mcf man page and mm_stripe (Set the Metadata Stripe Width).

Command Synopsis

```
:mm_stripe family-set-equipment-number number_DAUs
```

Examples

The following example changes data striping on family set device 10 to two DAUs:

```
File systems           samu                    version time date

ty  eq  state   device_name   status        high   low   mountpoint   server
ma  10  on      sam2          -----2---r-   80%    70%   /sam/sam2
 mm  11  on      /dev/dsk/c0t600A0B800026B6F000001F0F4CCFF833d0s0
 mr  12  on      /dev/dsk/c0t600A0B800026E584000021864C4943AAd0s0
 mr  13  on      /dev/dsk/c0t600A0B800026E5840000218B4CBD9C4Ed0s0
:
command:stripe 20 2
```

strun (Start Staging)

The strun command starts the stager.

Command Synopsis

```
:strun
```

Examples

In the example, strun starts staging:

```
Staging queue by media type: all      samu                    version time date

ty  length   fseq    ino  position   offset  vsn
li  451.611k   20    1030    207cc      473  VOL001
li  341.676k   20    1031    207cc      7fc  VOL002
li  419.861k   20    1032    207cc      aa9  VOL003
li  384.760k   20    1033    207cc      df2  VOL004
li  263.475k   20    1034    207cc      10f5 VOL005
li  452.901k   20    1035    207cc      1305 VOL006
li  404.598k   20    1036    207cc      1690 VOL007
:
command:strun
```

suid (Mount a File System with setuid Execution Enabled)

When issued with the Oracle HSM equipment ordinal number of a disk family-set device, the suid command mounts the file system with the setuid attribute enabled. The setuid attribute (*set user ID upon execution*) lets file system users run executables with the temporarily elevated permissions that may be required when using the files.

See nosuid (Mount a File System with setuid Execution Disabled) and the mount man page for further information.

The default is suid, mount the file system with the setuid attribute enabled.

Command Synopsis

```
:suid family-set-equipment-number
```

Examples

The following example mounts the file system on family set device 100 with the setuid attribute enabled:

```
File systems                        samu                    version time date

ty  eq   state   device_name          status       high low   mountpoint server
ma  100  on      qfs1                 m----2----d   90% 70%   /qfs1
 mm  101  on      /dev/dsk/c5t10d0s0
 md  102  on      /dev/dsk/c5t10d0s1
 md  103  on      /dev/dsk/c5t11d0s0
 md  104  on      /dev/dsk/c5t11d0s1
:
```

```
command: suid 100
```

sw_raid (Turn On Software RAID Mode)

When issued with the Oracle HSM equipment ordinal number of a disk family-set device, the sw_raid command enables software RAID mode on the specified device.

Command Synopsis

```
:sw_raid family-set-equipment-number
```

Examples

The following example enables software-RAID mode for family set device 100:

```
Mass storage status                  samu                    version time date

ty  eq   status       use  state  ord   capacity     free     ra  part high low
ma  100  m----2---r-  1%   on            7.270T     7.270T   1M  16    80% 70%
 mm  110              1%   on       0   465.253G   465.231G [975661056 inodes]
 mr  120              1%   on       1     1.817T     1.817T
 mr  130              1%   on       2     1.817T     1.817T
 mr  140              1%   on       3     1.817T     1.817T
 mr  150              1%   on       4     1.817T     1.817T
:
command:sw_raid 100
```

sync_meta (Synchronize File System Metadata)

When issued with the Oracle HSM equipment ordinal number of a disk family-set device and a Boolean value (0 for *no* or 1 for *yes*), the sync_meta command determines whether metadata is written to disk every time it changes. A **0** (*no*) value configures the file system to cache metadata changes in a buffer prior to writing. This reduces the amount of physical I/O, improves I/O efficiency, and thus improves overall performance. A 1 (*yes*) value configures Oracle HSM to write metadata changes to disk immediately. This sacrifices performance in order to maximize data consistency.

For Oracle HSM file systems that are neither shared nor mounted multi-reader, the default is 0. For shared and multi-reader file systems, the default is 1. If a shared file system is configured for failover, the value must be 1.

Command Synopsis

```
:sync_meta family-set-equipment-number 0|1
```

Examples

The following example configures the file system on family set device 100 to synchronize metadata for maximal data consistency:

```
File systems                         samu                    version time date

ty  eq  state   device_name      status       high low   mountpoint server
ma  100 on      qfs1            m----2----d   90% 70%   /qfs1
```

```
 mm  101 on     /dev/dsk/c5t10d0s0
 md  102 on     /dev/dsk/c5t10d0s1
 md  103 on     /dev/dsk/c5t11d0s0
 md  104 on /    dev/dsk/c5t11d0s1
:
command: sync_meta 100 1
```

The second example configures the file system on family set device 200 to cache metadata changes for best performance:

```
File systems                    samu                    version time date

ty  eq  state   device_name      status       high low   mountpoint server
ma  200  on qfs2                 m----2----d   90% 70%   /qfs2
 mm  201  on /dev/dsk/c5t10d0s3
 md  202  on /dev/dsk/c5t10d0s4
 md  203  on /dev/dsk/c5t11d0s3
 md  204  on /dev/dsk/c5t11d0s4
:
command: sync_meta 200 0
```

t (Display Tape Drive Status Screen)

The t command opens the Tape drive status screen. The screen displays status information for all tape drives configured within the environment.

Command Synopsis

```
t
:t
```

Controls

You can use the following control keys on this screen:

Ctrl-B (^b)
Move back one page.

Ctrl-F (^f)
Move forward one page.

Examples

The following example shows the tape drive status display.

```
Tape drive status                   samu                    version time date

ty eq  status     act use  state   vsn
sg 120 ---------p   0  0%  notrdy
sg 130 ---------p   0  0%  notrdy
sg 140 ---------p   0  0%  down
         drive set down due to ACS reported state
sg 150 --l------r   0 41%  ready   VOL014
```

Field Descriptions

ty
The ty field displays Oracle HSM media type. Device-type identifiers are listed in Appendix A and on the mcf man page.

eq
The eq displays the user-assigned Oracle HSM equipment ordinal number of the mass storage device.

status
The status field displays a ten-character string in which each character position represents the value of a different removable-media device attribute. Character positions are numbered starting from the left, position 1, and ending at the right, position 10. The following table defines the possible state codes for each position in the status string.

Position	Value	Meaning
1	s---------	A lowercase s in position one means that media is being scanned.
1	m---------	A lowercase m in position one means that the automated library is operational.
1	M---------	An uppercase M in position one means that the device is in maintenance mode.
2	-E--------	An uppercase E in position two means that the device encountered an unrecoverable error during scanning.
2	-a--------	A lowercase a in position two means that the device is in audit mode.
3	--l-------	A lowercase l in position three means that media has a label.
3	--L-------	An uppercase L in position three means that the media is being labeled
3	--N-------	An uppercase N in position three means that the media is foreign (not part of the Oracle HSM archive).
4	---I------	An uppercase I in position four means that the device is idle and not available for new connections. But operations already in progress are continuing until completion.
4	---A------	An uppercase A in position four means that the device needs operator attention.
5	----C-----	An uppercase C in position five means that the device needs cleaning.
5	----U-----	An uppercase U in position five means that an unload has been requested.
6	-----R----	An uppercase R in position six means that the device is reserved.
7	------w---	A lowercase w in position seven means that a process is writing to the media.
8	-------o--	A lowercase o in position eight means that the device is open.
9	--------P-	An uppercase P in position nine means that the device is positioning a tape volume.
9	--------F-	For a tape library, an uppercase F in position nine means that all storage slots are occupied. For a drive, an uppercase F in position nine means that the currently mounted removable media is full.

Position	Value	Meaning
9	--------X-	For a drive, an uppercase X in position nine means that files are actively migrating from one Oracle StorageTek T10000D to another, without consuming server resources. Copies are being made in xcopy-eod (end-of-data) mode. Copying started with the first valid file on the source volume and will continue until the drive reaches the end-of-data (EOD) mark on the tape.
9	--------x-	For a drive, a lowercase x in position nine means that files are actively migrating from one Oracle StorageTek T10000D to another, without consuming server resources. Copies are being made in xcopy mode. Copying started with the first valid file on the source tape and will continue until all files are copied.
10	---------R	An uppercase R in position ten means that the device is ready and the media is read-only.
10	---------r	A lowercase r in position ten means that the device is spun up and ready.
10	---------p	A lowercase p in position ten means that the device is present.
10	---------W	An uppercase W in position ten means that the device is write protected.

act
The status field displays the activity count.

use
The use field displays the percentage of cartridge space used.

state
The state field displays the current operating state of the device. Valid device states include:

- on means that the device is available.
- ro means that the device is available read-only.
- off means that the device is not available.
- down means that the device is available for maintenance only.
- idle means that the device is completing operations that are already in progress but is not accepting new connections.
- noalloc means that no more storage can be allocated on this device. An administrator has run the noalloc *eq* command, where *eq* is the equipment ordinal number of the device. The noalloc state persists until an administrator issues an alloc *eq* command.

vsn
The vsn field displays the volume serial number assigned to a media volume, the keyword nolabel if a volume is not labeled, or a blank if a volume is not loaded or if the device is off.

thresh (Set High and Low Watermarks)

When issued with the Oracle HSM equipment ordinal number of a disk family-set device and two percentages, the thresh (*threshold*) command defines the *high* and *low water marks* for the corresponding file system.

The high water mark is the percentage disk utilization where the Oracle HSM releaser process starts to delete archived (backed-up) files from disk in order to make space for new files. The low-water mark is the percentage of disk utilization where the Oracle HSM releaser has freed enough space for new files and stops deleting archived files from the disk. When properly configured, these two parameters insure that the largest possible number of files are kept available on disk while maintaining sufficient space for new and enlarged files, thus minimizing excessive staging activity and maximizing performance and capacity.

For each watermark, specify an integer value in the range [0-100]. A highs of 80% and a low of 70% is a typical setting.

Command Synopsis

```
:thresh family-set-equipment-number high% low%
```

Examples

The following example sets the high and low water marks to 80% and 70% percent for the sam1 file system on family set device 100:

```
File systems                          samu                    version time date

ty  eq    state device_name           status       high low   mountpoint server
ma  100   on    sam1                  m----2----d   90% 60%   /sam1
 mm  101   on    /dev/dsk/c5t10d0s0
 md  102   on    /dev/dsk/c5t10d0s1
 md  103   on    /dev/dsk/c5t11d0s0
 md  104   on    /dev/dsk/c5t11d0s1
:
command: thresh 100 80 70

File systems                          samu                    version time date

ty  eq   state device_name      status       high low   mountpoint server
ma  100  on    sam1             m----2----d   80% 70%   /sam1
 mm  101  on    /dev/dsk/c5t10d0s0
 md  102  on    /dev/dsk/c5t10d0s1
 md  103  on    /dev/dsk/c5t11d0s0
 md  104  on    /dev/dsk/c5t11d0s1
```

trace (Enable File System Tracing)

When issued with the Oracle HSM equipment ordinal number of a disk family-set device, the trace command enables tracing for the corresponding file system.

The default is trace (enabled).

For more information about file system tracing, see "notrace (Disable File System Tracing)" on page 3-65 and the defaults.conf man page.

Command Synopsis

```
:trace family-set-equipment-number
```

Examples

The following example enables tracing for the file system on family set device 100:

```
File systems                          samu                    version time date

ty  eq   state  device_name           status       high low   mountpoint server
ma  100  on     qfs1                  m----2----d   90% 70%   /qfs1
 mm  101  on     /dev/dsk/c5t10d0s0
 md  102  on     /dev/dsk/c5t10d0s1
 md  103  on     /dev/dsk/c5t11d0s0
 md  104  on     /dev/dsk/c5t11d0s1
:
command: trace 100
```

u (Display Staging Queue Screen)

The u command opens the Staging queue screen. The screen lists all files in the staging queue.

When issued without a parameter, the u command displays the staging queue for all media. When the command is issued with an Oracle HSM media-type identifier, the command displays the staging queue for the specified media type only. Media-type identifiers are listed in Appendix A and on the mcf man page.

Command Synopsis

```
u
:u mediatype
```

Controls

You can use the following control keys on this screen:

Ctrl-B (^b)
Move back one page.

Ctrl-F (^f)
Move forward one page.

Ctrl-D (^d)
Move down the page.

Ctrl-U (^u)
Move up the page.

Ctrl-K (^k)
Display the path on the second line of each entry

Examples

The following example shows the staging queue display.

```
Staging queue by media type: all        samu                    version time date

ty  length   fseq    ino  position offset  vsn
li  451.611k   20    1030    207cc    473  VOL001
```

```
li  341.676k   20    1031    207cc    7fc  VOL002
li  419.861k   20    1032    207cc    aa9  VOL003
li  384.760k   20    1033    207cc    df2  VOL004
li  263.475k   20    1034    207cc   10f5  VOL005
li  452.901k   20    1035    207cc   1305  VOL006
li  404.598k   20    1036    207cc   1690  VOL007
```

Field Descriptions

ty
The ty field displays the Oracle HSM device type. Media-type identifiers are listed in Appendix A and on the mcf man page.

length
The length field displays the length of the file in kilobytes.

fseq
The fseq field displays the file system's user-assigned equipment ordinal number

ino
The ino field displays the inode number of the file that is being staged.

position
The position field displays the location of the archive file on the specified media volume.

offset
The offset field displays the position of the file that is being staged relative to the start of the archive file.

vsn
The vsn field displays the volume serial number assigned to the media volume, the keyword nolabel if the volume is not labeled, or a blank if a volume is not loaded or if the volume is loaded into a device that is off.

unavail (Make a Device Unavailable to Oracle HSM)

The unavail command keeps Oracle HSM from using the device specified by a given Oracle HSM equipment ordinal number.

The command is useful when you want to reserve a drive for other uses, such as loading media for file-system restores during disaster recovery.

Command Synopsis

```
:unavail  equipment-number
```

Examples

```
Device configuration:                samu                      version time date

ty  eq   state   device_name                              fs   family_set
sk  100  on      /etc/opt/SUNWsamfs/SL8500_T10K_parms   100   libraryA
 ti  101  on       /dev/rmt/0cbn                          100   libraryA
 ti  102  on       /dev/rmt/1cbn                          100   libraryA
```

```
 ti  103  on      /dev/rmt/2cbn                          100   libraryA
hy  104 on      historian                              104
:
command:unavail 103
```

The device state changes to unavail:

```
Device configuration:               samu                    version time date

ty  eq   state   device_name                            fs   family_set
sk  100  on      /etc/opt/SUNWsamfs/SL8500_T10K_parms   100   libraryA
 ti  101  on      /dev/rmt/0cbn                          100   libraryA
 ti  102  on      /dev/rmt/1cbn                          100   libraryA
 ti  103  unavail /dev/rmt/2cbn                          100   libraryA
hy  104  on      historian                             104
```

unload (Unload a Drive)

The unload command unloads media mounted on the device specified by an Oracle HSM equipment ordinal number. If the device has magazines, the command unloads the mounted cartridge and ejects the magazine. See "load (Load a Cartridge into a Drive)" on page 3-48.

Command Synopsis

```
:unload equipment-number
```

Examples

```
Device status                    samu                    version time date

ty  eq    state  device_name                fs  status
sn  100   on   /etc/opt/SUNWsamfs/dcstkconf  100 m--------r
        empty
li  101   on   /dev/rmt/2cbn                100 ---------p
        empty
li  102   on   /dev/rmt/5cbn                100 ---------p
         empty
li  103   on   /dev/rmt/6cbn                100 --l------r
        idle
li  104   on   /dev/rmt/4cbn                100 ---------p
         empty
hy  151   on   historian                    151 ----------
:
command:unload 103
```

v (Display Robot VSN Catalog Screen)

The v command opens the Robot VSN catalog screen. The screen lists catalog information for all media volumes stored in a given automated library.

When issued without a parameter, the v command displays the catalogs for all libraries. When the command is issued with a user-assigned Oracle HSM equipment ordinal number, the command displays the volume catalog for the specified device only.

Command Synopsis

```
v
:v equipment-number
:v historian
```

Controls

Ctrl-B (^b)
Move back one page.

Ctrl-F (^f)
Move forward one page.

Ctrl-D (^d)
Move down to the next catalog.

Ctrl-U (^u)
Move up to the previous catalog.

Ctrl-I (^i)
Display detailed information, two lines per entry. Press **Ctrl-I** once to display times and barcodes. Press Ctrl-I a second time to display volume reservations on the second line.

1
Sort entries by library slot.

2
Sort entries by count.

3
Sort entries by percent used.

4
Sort entries by Volume Serial Number (VSN).

5
Sort entries by access time.

6
Sort entries by barcode.

7
Sort entries by label time.

Ctrl-K (^k)
Toggle between the sorting methods enumerated above.

/
Search for a specified Volume Serial Number (VSN) and, if found, display the corresponding record at the top of the list.

%
Search for a specified barcode and, if found, display the corresponding record at the top of the list.

$
Search for a specified slot number and, if found, display the corresponding record at the top of the list.

Examples

The following example shows the automated library catalog display sorted by slot:

```
Robot VSN catalog by slot   : eq 900  samu                    version time date
                                                                    count 32
slot         access time count  use  flags         ty vsn

   0     2013/07/02 12:53    3  96%  -il-o-b----f  li VOL001
   1     2013/07/02 12:53    2  98%  -il-o-b----f  li VOL002
   2     2013/07/02 12:52    1  86%  -il-o-b----f  li VOL003
   3     2013/07/02 12:52    1  95%  -il-o-b----f  li VOL004
   4     2013/07/02 12:51    1  79%  -il-o-b-----  li VOL005
   5     2013/07/02 12:45    0  14%  -ilEo-b----f  li VOL006 MEDIA ERROR
   6     2013/07/02 12:46    0   0%  -il-o-b-----  li VOL007
   7     2013/07/02 12:46    0   0%  -il-o-b-----  li VOL008
   8     2013/07/02 12:47    0   0%  -il-o-b-----  li VOL009
   9     2013/07/02 12:47    0   0%  -il-o-b-----  li VOL010
  10     2013/07/02 12:48    0   0%  -il-o-b-----  li VOL011
  11     2013/07/02 12:48    0   0%  -il-o-b-----  li VOL012
  12     2013/07/02 12:48    0   0%  -il-o-b-----  li VOL013
  13     2013/07/02 12:49    0   0%  -il-o-b-----  li VOL014
  14     2013/07/02 12:49    0   0%  -il-o-b-----  li VOL015
  15     2013/07/02 12:50    0   0%  -il-o-b-----  li VOL016
  16     2013/07/02 12:50    0   0%  -il-o-b-----  li VOL017
  17     2013/07/02 12:51    0   0%  -il-o-b-----  li VOL018
  18     2013/07/02 12:51    0   0%  -il-o-b-----  li VOL019
  19     none               50   0%  -il-oCb-----  li CLN020
```

The next example shows the same screen sorted by the percentage of each volume that has been used:

```
Robot VSN catalog by % used  : eq 900   samu                    version time date
                                                                    count 32
slot         access time count  use  flags         ty vsn

   1     2013/07/02 12:53    2  98%  -il-o-b----f  li VOL002
   0     2013/07/02 12:53    3  96%  -il-o-b----f  li VOL001
   3     2013/07/02 12:52    1  95%  -il-o-b----f  li VOL004
   2     2013/07/02 12:52    1  86%  -il-o-b----f  li VOL003
   4     2013/07/02 12:51    1  79%  -il-o-b-----  li VOL005
   5     2013/07/02 12:45    0  14%  -ilEo-b----f  li VOL006 MEDIA ERROR
  12     2013/07/02 12:48    0   0%  -il-o-b-----  li VOL013
  13     2013/07/02 12:49    0   0%  -il-o-b-----  li VOL014
  14     2013/07/02 12:49    0   0%  -il-o-b-----  li VOL015
  15     2013/07/02 12:50    0   0%  -il-o-b-----  li VOL016
  16     2013/07/02 12:50    0   0%  -il-o-b-----  li VOL017
  17     2013/07/02 12:51    0   0%  -il-o-b-----  li VOL018
  19     none               50   0%  -il-oCb-----  li CLN020
   6     2013/07/02 12:46    0   0%  -il-o-b-----  li VOL007
   7     2013/07/02 12:46    0   0%  -il-o-b-----  li VOL008
   8     2013/07/02 12:47    0   0%  -il-o-b-----  li VOL009
   9     2013/07/02 12:47    0   0%  -il-o-b-----  li VOL010
  10     2013/07/02 12:48    0   0%  -il-o-b-----  li VOL011
  11     2013/07/02 12:48    0   0%  -il-o-b-----  li VOL012
  18     2013/07/02 12:51    0   0%  -il-o-b-----  li VOL019
```

Field Descriptions

ty
The ty field displays Oracle HSM media type. Device-type identifiers are listed in Appendix A and on the mcf man page.

eq
The eq displays the user-assigned Oracle HSM equipment ordinal number of the mass storage device.

count
The count field displays the number of accesses to this volume since the last audit operation.

slot
The slot field displays the slot number that identifies a storage cell in a robotic library.

access time
The access time field displays the date and time when the volume was last accessed.

use
The use field displays the percentage of cartridge space used.

flags
The **flags** field displays a twelve-character string in which each character position represents an attribute of the catalog entry. Character positions are numbered starting from the left, position 1, and ending at the right, position 12. In some cases, more than one flag can occur in a field, and one flag overrides the other.

The following table defines the possible flag codes for each position in the string:

Position	Flag	Meaning
1	A----------	An uppercase A in position 1 means that the volume in the corresponding library slot needs an audit.
2	-i----------	A lowercase i in position 2 means that the corresponding library slot is in use.
3	--l---------	A lowercase l in position 3 means that the volume in the corresponding library slot is labeled.
3	--N---------	An uppercase N in position 3 means that the volume in the corresponding library slot is unlabeled or foreign to the Oracle HSM file system.
4	---E--------	An uppercase E in position 4 means that there was a media error when accessing the corresponding volume.
5	----o-------	A lowercase o in position 5 means that the corresponding library slot is occupied by a volume.
6	-----C------	An uppercase C in position 6 means that the corresponding library slot is occupied by a cleaning tape.
6	-----p------	A lowercase p in position 6 means that the corresponding library slot is occupied by a priority volume.
7	------b-----	A lowercase b in position 7 means that the volume in the corresponding library slot has a barcode.

Position	Flag	Meaning
8	-------W----	An uppercase W in position 8 means that the volume in the corresponding library slot is write-protected by a physical switch on the cartridge.
9	--------R---	An uppercase R in position 9 means that the volume in the corresponding library slot is logically read-only.
10	---------c--	A lowercase c in position 10 means that the volume in the corresponding library slot has been selected for recycling.
11	----------d-	A lowercase d in position 11 means that the volume in the corresponding library slot has a duplicate volume serial number.
11	----------D-	An uppercase D in position 11 means that the volume in the corresponding library slot is a migration destination volume.
11	----------e-	A lowercase e in position 11 means that the volume in the corresponding library slot failed to migrate to new media due to an error.
11	----------M-	An uppercase M in position 11 means that the volume in the corresponding library slot has been partially migrated to new media.
11	----------m-	A lowercase m in position 11 means that the volume in the corresponding library slot has been migrated to new media.
11	----------S-	An uppercase S in position 11 means that the volume in the corresponding library slot is a migration source volume.
11	----------U-	An uppercase U in position 11 means that the volume that should reside in the corresponding library slot is unavailable.
12	-----------f	A lowercase f in position 11 means that the archiver found the volume in the corresponding library slot full.
12	-----------X	An uppercase X in position 11 means that the corresponding library slot is an export slot.

vsn

The vsn field displays the volume serial number assigned to the media, the keyword nolabel if the volume is not labeled, or a blank if a volume is not loaded or if the device is off.

LVTime

The LVTime field displays the date and time when the entire tape was last verified.

LVPos

The LVPos field displays the position where the last verification was canceled.

w (Display the Pending Stage Queue Screen)

The w command opens the Pending stage queue screen. The screen displays requests for files that are not currently available on disk and must be staged from removable media volumes that are not currently loaded.

When issued without a parameter, the w command displays queued stage requests for all media types. When the command is issued with an Oracle HSM media-type identifier, the command displays the pending stage requests for the specified media type only. Media-type identifiers are listed in Appendix A and on the mcf man page.

Command Synopsis

```
w
:w mediatype
```

Controls

You can use the following control keys on this screen:

Ctrl-B (^b)
Move back one page.

Ctrl-F (^f)
Move forward one page.

Ctrl-D (^d)
Move down the page.

Ctrl-U (^u)
Move up the page.

Ctrl-K (^k)
Display the path on the second line of each entry

Examples

The following example shows the pending stage queue screen for all media types. Thirteen requested files reside on one volume that is waiting to load:

```
Pending stage queue by media type: all   samu   version time date
                                                volumes 1 files 13
ty    length fseq ino position offset  vsn
li    1.383M    1  42    3a786   271b  VOL002
li    1.479M    1  56    3a786   5139  VOL002
li 1018.406k    1  60    3a786   6550  VOL002
li    1.VOLM    1  65    3a786   7475  VOL002
li    1.528M    1  80    3a786   99be  VOL002
li    1.763M    1  92    3a786   ce57  VOL002
li    1.749M    1 123    3a786   11ece VOL002
li  556.559k    1 157    3a786   1532f VOL002
li  658.970k    1 186    3a786   17705 VOL002
li  863.380k    1 251    3a786   1dd58 VOL002
li    1.268M    1 281    3a786   1f2b7 VOL002
li    1.797M    1 324    3a786   23dfa VOL002
li    1.144M    1 401    3a786   2bb6d VOL002
```

Field Descriptions

ty
The ty field displays the Oracle HSM device type. Media-type identifiers are listed in Appendix A and on the mcf man page.

length
The length field displays the length of the file in kilobytes.

fseq
The fseq field displays the file system's user-assigned equipment ordinal number

ino
The ino field displays the inode number of the requested data file.

position
The position field displays the location on the media where the required archive file starts.

offset
The offset field displays the starting location of the requested data file relative to the archive file position.

vsn
The vsn field displays the volume serial number assigned to the media, the keyword nolabel if the volume is not labeled, or a blank if a volume is not loaded or if the device is off.

wr_throttle (Set the Maximum Size of Outstanding Writes)

The wr_throttle command limits the amount of data that any one file can keep in memory before writing to the Oracle HSM disk family-set device specified by a supplied equipment ordinal number. Throttling keeps outstanding writes from consuming excessive memory.

The default data amount is 5% of the system memory, expressed in kilobytes.

```
:wr_throttle eq value
```

For eq, specify the user-assigned Oracle HSM equipment number for a file system.

For value, specify an integer number of kilobytes. If value=0, there is no limit.

Command Synopsis

```
:wr_throttle family-set-equipment-number number_kilobytes
```

Examples

The following example sets the wr_throttle value for family set device 100 to 838656 kilobytes (819 megabytes), which is 5% of the 16 gigabytes of memory available on the host system:

```
Mass storage status                    samu                    version time date

ty  eq   status      use  state  ord   capacity     free     ra  part high low
ma  100  m----2---r-  1%  on            7.270T     7.270T   1M  16    80% 70%
 mm  110              1%  on       0   465.253G   465.231G [975661056 inodes]
 mr  120              1%  on       1     1.817T     1.817T
 mr  130              1%  on       2     1.817T     1.817T
 mr  140              1%  on       3     1.817T     1.817T
 mr  150              1%  on       4     1.817T     1.817T
:
command:wr_throttle 100 838656
```

writebehind (Set the Write-Behind Parameter)

When issued with the Oracle HSM equipment ordinal number of a disk family-set device, the writebehind command specifies the maximum number of bytes that the device should cache in memory before writing data to disk. When correctly configured, write-behind improves file-system performance by writing information in larger, contiguous blocks.

Specify a write-behind value in the range [1-8192] kilobytes. The default is 512 kilobytes.

You can also set a writebehind directive in the samfs.cmd file. See the samfs.cmd man page for details.

Command Synopsis

```
:writebehind family-set-equipment-number number_kilobytes
```

Examples

The following example sets the write-behind value for family set device 100 to 1024 kilobytes (1 megabyte):

```
Mass storage status                  samu                    version time date

ty  eq  status       use  state   ord  capacity       free   ra  part high low
ma  100 m----2---r-  1%   on            7.270T     7.270T   1M  16    80% 70%
 mm  110              1%   on       0  465.253G   465.231G [975661056 inodes]
 mr  120              1%   on       1    1.817T     1.817T
 mr  130              1%   on       2    1.817T     1.817T
 mr  140              1%   on       3    1.817T     1.817T
 mr  150              1%   on       4    1.817T     1.817T
:
command:writebehind 100 1024
```

wrlease (Set the Write Lease Time for a Shared File System)

Before writing data to a shared file, a host must request and obtain an *write lease* that grants it exclusive access to the file. The wrlease command defines the number of seconds during which this lease remains valid on the specified device. If the specified time runs out before the write operation is complete, the host must reacquire the lease before proceeding.

The specified number of seconds must fall in the range [15-600]. The default is 30 seconds.

For related information, see "aplease (Set Append Lease Expiration Time for a Shared File System)" on page 3-5, "rdlease (Set the Read Lease Time for a Shared File System)" on page 3-75, and "lease_timeo (Set the Unused Lease Timeout for a Shared File System)" on page 3-47.

Command Synopsis

```
:wrlease family-set-equipment-number interval_in_seconds
```

Examples

The following example sets the write lease time to 2 minutes (120 seconds) on family set device 800:

```
File systems                          samu                    version time date

ty    eq  state device_name           status      high low mountpoint server
ma    800 on    shareqfs1             m----2c--r-  80% 70% /shareqfs1
 mm    801 on    /dev/dsk/c6t50020F2300004655d0s0
 mr    802 on    /dev/dsk/c6t50020F2300004655d0s1
ma    810 on    shareqfs2             m----2c--r-  80% 70% /shareqfs2
 mm    811 on    /dev/dsk/c6t50020F2300004655d0s6
 mr    812 on    /dev/dsk/c6t50020F2300004655d0s7
:
command:wrlease 800 120
```

x (Display Current Migration Status)

The x command displays the Migration status screen. This screen shows the current state of the migration daemon and provides a summary view of the progress of the currently active migration (if any).

Command Synopsis

```
x
:x
```

Controls

You can use the following control keys on this screen:

Ctrl-B (^b)
Move back one page.

Ctrl-F (^f)
Move forward one page.

Ctrl-D (^d)
Move down the page.

Ctrl-U (^u)
Move up the page.

Ctrl-I (^i)
Display detailed information, two lines per entry.

Examples

The following example shows the migration status screen while a migration is under way:

```
Migration status        samu          version time date
Source Vsns - wait:  0 fsscan: 0 copy: 1 update ino: 0 log: 0 done:  0
Status:Run
source    dest      cmod perc status
li VOL002 li VOL012 s     24% 0x00000230 blocks read
```

The example below shows the detailed information displayed by the ^i option, when the migration is nearly complete. The expanded display shows that all five archive copies are complete, the total size is 10GB, the migration copy process is checking the archive (tar) file headers, and one file header has been checked so far:

```
Migration status        samu          version time date
Source Vsns - wait:  0 fsscan: 0 copy: 1 update ino: 0 log: 0 done:  0
Status: Run
source    dest      cmod perc status
bytes copied/total     Copies:total     copied tarchecked    updated     logged
li VOL002 li VOL0012 -  100% Validating tar header on li.000040
      10.00G/10.00G               5          5          1          0          0
```

Field Descriptions

Source Vsns
The Source Vsns field lists statistics for the tape volumes from which files are migrating. The Source Vsns field has the following subfields:

wait
The number of volumes awaiting processing.

fsscan
The number of volumes for which archive file copies are still being matched to file system inodes.

copy
The number of volumes that are currently being copied to new media.

update ino
The number of volumes for which file system inodes have been updated to reflect the new locations of the migrated archive files.

log
The number of volume-migration logs created.

done
The number of volumes that have been completely copied to to replacement media.

Status
The status field displays the current status of the migration daemon: Stop, Idle, or Run.

source
The source field identifies each volume that is being copied. Volumes are identified by media type and Volume Serial Number (VSN). Media-type identifiers are listed in Appendix A and on the mcf man page.

dest
The dest field identifies each volume to which data is being copied. Volumes are identified by media type and Volume Serial Number (VSN).

cmod
The cmod field displays the copy mode, which can be one of the following:

-
Copying has not yet started.

X
The copy is being made using the xcopy-eod (end-of-data) mode of the Oracle StorageTek T10000D Extended Copy feature. Files are copied from drive to drive, without consuming server resources. Copying starts with the first valid file on the source volume and continues until the drive reaches the end-of-data (EOD) mark on the tape.

x
The copy is being made using the xcopy mode of the Oracle StorageTek T10000D Extended Copy feature. Files are copied from drive to drive, without consuming server resources. Copying starts with the first valid file on the source tape and continues until all files are copied.

s
The copy is being made in server-copy mode. Data is read from the source volume to a buffer on the file-system server and then written out to the destination volume.

c
The copy is being made in server-copy mode, with the server converting the source volume block size to the tape block size required by the destination volume.

perc
The perc field displays the percentage of the source volume that has been copied to new media.

status
The status field displays additional information on the progress of the copy, such as the number of tape blocks read from the source or the fact that the volume is being loaded into a drive.

y (List Migrating Volumes by VSN)

The y command displays the Migration vsn list screen. From this screen, you can track the progress of a migration in detail, volume by volume.

Command Synopsis

```
y
:y
```

Controls

You can use the following control keys on this screen:

Ctrl-B (^b)
Move back one page.

Ctrl-F (^f)
Move forward one page.

Ctrl-D (^d)
Move down the page.

Ctrl-U (^u)
Move up the page.

Ctrl-I (^i)
Display detailed information, two lines per entry.

Examples

The following example shows the migration volume serial number (VSN) screen. The contents of LTO (li) volume VOL001 have finished migrating to LTO volume VOL112:

```
Migration vsn list        samu          version time date
Status:  Run  Vsns:2 src:1 dest:1 maxcopy:2
ord m ty vsn     start time  end   time  status  Inodes done/tot    bytes
  0 S li VOL001  10/16 12:12 10/16 12:14 complete    35023/35023    12.00G
  0 D li VOL112  10/16 12:12 10/16 12:14 avail                      12.00G
```

The example below shows the detailed information displayed by the ^i option. Here the migration copy process is checking the archive file (tar) header on source volume VOL040 and positioning destination volume VOL142 prior to checking the archive file header.

```
Migration vsn list        samu          version time date
Status: Run  Vsns:2 src:1 dest:1 maxcopy:1
ord m ty vsn    start time  end time    status  Inodes done/tot         bytes
    detail
  0 S li VOL040 11/15 17:17 none        tar_check         0/5          10.00G
    Validating tar header on li.VOL040
  0 D li VOL142 11/15 17:17 none        tar_check                      10.00G
    Validating tar header, positioning to 0x61897
```

Field Descriptions

Status
The status field displays the current status of the migration daemon: Stop, Idle, or Run.

Vsns
The Vsns field lists number of tape volumes involved in the current migration.

src
The src field lists the total number of tape volumes that are to be copied to new media.

dest
The dest field lists the total number of tape volumes that are to be copied to new media.

maxcopy
The maxcopy field displays the maximum number of copies that can be made at any one time. You specify this value in the migrationd.cmd file.

ord
The ord (*ordinal*) field lists the position that the corresponding migration directive occupies in the migration.cmd file.

m
The m field indicates the role that the corresponding volume plays in the migration: either S (*source*) or D (*destination*).

ty
The ty field lists a two-character code that identifies the type of media on which the corresponding volume has been written. Media-type identifiers are listed in Appendix A and on the mcf man page.

vsn
The vsn field lists the *volume serial number* assigned to the corresponding volume. The VSN is six-character, alphanumeric code that uniquely identifies each piece of media in a robotic library.

start time
The start time field lists the month, day, hour, and minute when each copy operation began.

end time
The end time field lists the month, day, hour, and minute when each copy operation ended.

status
The status field lists the current copy state of the corresponding volume. Values may include:

- avail means that the destination volume has space available for additional copies.
- complete means that the source volume has been successfully copied.
- copy means that the source volume is being copied to the destination volume.
- copy_wait means that the source volume is waiting for copying to start.
- error means that the source volume could not be copied.
- fs_scan_wait means that the source volume has been scheduled and is waiting for copying to start.
- fs_scan means that the migration process is scanning the file system.
- log means that logging is in progress.
- sched_wait means that the source volume has been scheduled and is waiting for copying to start.
- update_inode means that the corresponding inode is being updated.

inodes done/tot
The inodes done/tot field lists the number of inodes that have been updated with the new location of corresponding file copies and the total number of inodes that must be updated to complete the copy operation.

bytes
The bytes field lists the amount of data transferred. When no units are specified, the listed value is in bytes. Otherwise, unit may be specified using k for kilobytes, M for megabytes, G for gigabytes, T for terabytes, P for petabytes, or E for exabytes.

z (Cloud Library Activity)

The z command displays the Cloud Library Activity screen.

Command Synopsis

```
z
```

```
:z
```

Controls

You can use the following control keys on this screen:

Ctrl-B (^b)
Move back one page.

Ctrl-F (^f)
Move forward one page.

Ctrl-D (^d)
Move down the page.

Ctrl-U (^u)
Move up the page.

Ctrl-I (^i)
Display detailed information, two lines per entry.

Examples

The following example shows:

```
Cloud Library Activity   samu           version time date
Cloud Library eq# 200    : x62l0chi
Connection status        : Connected/Ready
Encryption status        : PKCS#11 KMS
Cloud library message    : running
Active streams           : 1 reading, 0 writing, 3 idle
I/O rates                : 8.9 MB/sec reading <#---<#---<#--- from cloud
                         : 0.0 MB/sec writing (idle)
Blocks xfer/sec          : 0.5 (read), 0.0 (write)
Retry counts/sec         : 0.0 (read), 0.0 (write)
Drive Status VSN                MB/sec Retry/s
201   Idle
202   Read   x62l0chi_0000000   8.9    0.0
203   Idle
204   Idle
```

Field Descriptions

Cloud Library eq#
The `Cloud Library eq#` field displays the equipment number that the Oracle HSM master configuration file (`mcf`) assigns to the named cloud library.

Connection status
The `Connection status` field displays the current status of the connection between the file system server and the cloud library.

Encryption status
The `Encryption status` field displays the encryption API and key manager that is currently in use, if any.

Active streams
The `Active streams` field displays the current activity of the cloud drives.

I/O rates
The `I/O rates` field displays the current, aggregate read and write speeds in megabytes per second.

Blocks xfer/sec
The `Blocks xfer/sec` field displays the aggregate number of blocks written and read per second.

Retry counts/sec
The `Retry counts/sec` field displays the aggregate numbers of read and write retries per second.

Drive
The `Drive field` list the equipment numbers that the Oracle HSM master configuration file (`mcf`) assigns to the cloud drives.

Status
The `Status` field displays the current state of the corresponding Drive: Idle, Read, Write, etc.

VSN
The `VSN` field displays the volume serial number of the cloud media volume that is being accessed by the corresponding Drive.

MB/sec
The `MB/sec` field displays the current read or write speed attained by the corresponding Drive, in megabytes per second.

Retry/s
The `Retry/s` field displays the number of read or write retries per second for the corresponding Drive.

4

Service and Support Commands

This section lists detailed summaries of `samu` service and support commands in alphabetical order.

:C (Display the Contents of Shared Memory at a Specified Address)

The `C` command displays the contents of a specified shared memory address. The screen supplies debugging information that may be needed by Oracle service and support staff.

The command is always issued with a parameter, the specified memory address in hexadecimal notation.

Command Synopsis

```
:C 0xADDRESS
```

Examples

The following example shows part of the memory screen output:

```
Memory  base: 0x1234567   samu           version time date

00000000  05961000 1881c7e0 0891e800 1c9de3bf  ......G`..h...c?
00000010  40901020 90350419 207ffd0b d5921020  @.. .5.. .}.U..
00000020  00b61000 08a00220 08d07620 287fffff  .6... . .Pv (...
00000030  e5d05e20 10921020 00941020 06ba06e0  eP^ ... ... .:.`
00000040  4039004b 46d076e0 009006e0 48b80723  @9.KFPv`...`H8.#
00000050  d87ff93c 6f961020 009006e0 50921020  X.y<o.. ...`P..
00000060  00941020 0133004b 467ffb41 11961020  ... .3.KF.{A...
00000070  00901000 1d921020 00941020 022d004b  ....... ... .-.K

00000080  467ff959 8f961020 00ea06a2 a02f004b  F.yY... .j." /.K
00000090  46940663 64f873a8 af9805a3 749605e3  F..cdxs(/..#t..c
000000a0  6cc073a8 b7933d60 00b806a2 a0d05e20  l@s(7.=`.8." P^
000000b0  10ac1020 00b20720 c8400276 039a1020  .,. .2. H@.v...
000000c0  008b3da0 00d076e0 28af2970 03d05dc0  ..= .Pv`(/)p.P]@
000000d0  197ff96a fad25e20 10809000 08024000  ..yjzR^ ......@.
000000e0  08ac05a0 01973da0 0080a2e0 020a6fff  .,. ..= .."`..o.
000000f0  f78b3da0 00108000 091b041a b0150419  w.= ........0...
```

D (Display Disk Volume Dictionary Screen)

The D command displays the Disk volume dictionary screen. The screen supplies internal debugging information.

Command Synopsis

```
D
:D
```

Controls

You can use the following control keys on this screen:

Ctrl-I (^j)
Use detailed, two-line display format.

Ctrl-J (^j)
Toggle between binary and decimal units of size.

Examples

The following example shows the device configuration screen.

```
Disk volume dictionary    samu           version time date

header
version 460

volumes
magic 340322 version 9 nkeys 2 ndata 2
index  space     capacity    used    flags  volume
  0  12882411520  12887785472  10291200  -----  disk01
  1   6443827200   6443892736    70656  -----  disk02
clients
magic 340322 version 9 nkeys 1 ndata 1
```

Flags

Disk volume flags are five-character strings in which each character position represents an attribute of the disk volume. Character positions are numbered starting from the left, position 1, and ending at the right, position 5.

To set or clear a disk volume dictionary flag, use the diskvols (Set/Clear Flags on Disk Volumes Used for Archiving) operator command.

The following table defines the possible flag codes for each position in the string:

Position	Flag	Meaning
1	l----	The volume is labeled; a seqnum file has been created. An administrator set the flag to prevent the software from creating a new seqnum file.
2	-r---	The volume is defined on a remote host.
3	--U--	The volume is unavailable.
4	---R-	The volume is read only.

Position	Flag	Meaning
5	----E	A media error occurred when writing to the disk archive directory.

F (Display Optical Disk Label Screen)

The Optical disk label screen displays the label on an optical disc. Oracle support staff may request that you run this command in support of debugging.

fs (Specify a File System)

The fs command specifies the file system displayed by N (Display File System Parameters). When you have many file systems, the command lets you avoid paging through the display

Command Synopsis

```
:fs file-system name
```

I (Display Inode Information)

The I command displays the Inode screen for a given mount point. Oracle support staff may request that you run this command in support of debugging.

When issued without a parameter, the I command displays all inodes for the current file-system mount point, starting with Inode 0x1 (1). If a mount point directory is not currently specified, you are prompted for one.

When the command is issued with a decimal or hexadecimal inode number, the command displays the specified inode.

Command Synopsis

```
I
:I
:I inode-number
:I 0xinode-number
```

Controls

You can use the following control keys on this screen:

Ctrl-B (^b)
Move back one inode.

Ctrl-F (^f)
Move forward one inode.

Ctrl-K (^k)
Step through the following display formats:

- archive
- raw
- rawincore

- file, Extents (4k displayed as 1k)
- file, Extents (raw)

Example

The following example shows an abbreviated inode screen for Inode 0x1 (1) in file, Extents (4k displayed as 1k) format:

```
Inode   0x1 (1) format: file      samu           version time date
  incore: y

00008100 mode    -r--------     409cdf57 access_time
00000001 ino         (1)        1d32ea20
00000001 gen         (1)        4096b499 modify_time
00000002 parent.ino (2)         02588660
00000002 parent.gen (2)         4096b499 change_time
00000000 size_u                 02588660
000c0000 size_l     (786432)    4096b443 creation_time
01000000 rm:media/flags         409a8a7c attribute_time
00000000 rm:file_offset         409c0ce6 residence_time
00000000 rm:mau                 00000000 unit/cs/arch/flg
00000000 rm:position            00000000 ar_flags ...
```

J (Display the Shared Memory Segment for the Load-Request Preview Queue)

The J command displays the Preview shared memory screen displays the shared memory segment for the removable media load-request preview queue. Oracle support staff may request that you run this command in support of debugging.

Command Synopsis

```
J
:J
:J 0xMemoryAddress
```

Controls

Ctrl-B (^b)
Move back one page.

Ctrl-F (^f)
Move forward one page.

Ctrl-D (^d)
Move down the page.

Ctrl-U (^u)
Move up the page.

Sample Display

The following example shows the preview shared memory display. This sample output has been truncated.

```
Preview shared memory  size: 155648   samu           version time date

00000000 00040000 00014d58 00000000 00000000 ......MX........
00000010 00000000 00000000 73616d66 73202d20 ........samfs -
00000020 70726576 69657720 6d656d6f 72792073 preview memory s
00000030 65676d65 6e740000 00026000 00000000 egment..........
00000040 00025fff 00000000 00040000 00014d58 .._...........MX
00000050 00000000 00000000 00000000 00000000 ................
00000060 0000d9e0 00000064 00000000 000001b8 ..Y.....d.......8
00000070 3f800000 447a0000 0000d820 00000008 ?...Dz....X ....
```

K (Display Kernel Statistics)

The K command displays the Kernel statistics screen. It shows various statistics for the Oracle HSM kernel module. The screen supplies debugging information that may be needed by Oracle service and support staff.

Command Synopsis

```
K
:K
```

Controls

Ctrl-B (^b)
Move back one page.

Ctrl-F (^f)
Move forward one page.

Examples

The following example shows the kernel statistics display.

```
Kernel statistic                   samu              version time date

module: SAM-QFS name: general instance: 0 class: fs
version                    version date time
configured file systems    8
mounted file systems       8
nhino                      16384
ninodes                    129526
inocount                   129527
inofree                    128577
```

L (Display Shared Memory Tables and System Defaults)

The **L** command displays the Shared memory tables screen. It shows shared memory-table attributes and system defaults that are kept in shared memory. The screen supplies debugging information that may be needed by Oracle service and support staff.

Command Synopsis

```
L
:L
```

Controls

None.

Examples

```
Shared memory tables            samu           version time date

shm ptr           tbl:                              defaults:
size              f000 (61440)                      optical            mo
left              212a (8490)                       tape               li
scanner pid       12673                             timeout            600
fifo path         01a8 /var/opt/SUNWsamfs/previews                     100
dev_table         01c4                              stages             1000
first_dev         0e80                              log_facility       184
scan_mess         6260                              label barcode      TRUE
preview_shmid     83886122                          barcodes low       FALSE
flags             0x40000000                        export unavail     FALSE
preview stages    53780                             attended           TRUE
preview avail     100                               start rpc          FALSE
preview count     0                                 div                OFF
preview sequence  0                                 distio             OFF
age factor        1                                 vsn factor         1000
fs tbl ptr        0xd1dc                            fs count           1
fseq  10 samfs1   state  0      0      0      0      0
fseq  20 samfs2   state  0      0      0      0      0
fseq  30 qfs1     state  0      0      0      0      0
fseq  40 qfs2     state  0      0      0      0      0
fseq  50 qfs3     state  0      0      0      0      0
fseq  60 qfs4     state  0      0      0      0      0
fseq 100 shareqfs1 state  0      0      0      0      0
fseq 110 shareqfs2 state  0      0      0      0      0
```

M (Display Raw Shared Memory)

The M command displays the Shared memory screen. The screen shows the raw shared memory segment at the current address or a specified address. Oracle support staff may request that you run this command in support of debugging.

Command Synopsis

```
M
:M
:M 0xMemoryAddress
```

Controls

Ctrl-B (^b)
Move back one page.

Ctrl-F (^f)
Move forward one page.

Ctrl-D (^d)
Move down the page.

Ctrl-U (^u)
Move up the page.

Examples

The following example shows an abbreviated shared memory display.

```
Shared memory    size: 73728         samu           version time date

00000000 00040000 00014d58 00000000 00000000 ......MX........
00000010 00000000 00000000 73616d66 73202d20 ........samfs -
00000020 73686172 6564206d 656d6f72 79207365 shared memory se
00000030 676d656e 74000000 00012000 000044c8 gment..... ...DH
00000040 0000dd20 00000000 00000742 00000745 ..] .......B...E
00000050 00000001 00000000 00000000 c0000000 ............@...
00000060 00000001 0001534d 00000000 00000000 ......SM........
00000070 00000000 00000000 00000000 00000000 ................

00000080 00000000 00000000 00000000 00000000 ................
00000090 20000000 000001b0 000001cc 00000450  ......0...L...P
000000a0 0000cf50 00000001 00000001 4c696365 ..OP........Lice
000000b0 6e73653a 204c6963 656e7365 206e6576 nse: License nev
000000c0 65722065 78706972 65732e00 00000000 er expires......
000000d0 00000000 00000000 00000000 00000000 ................
000000e0 00000000 00000000 00000000 00000000 ................
000000f0 00000000 00000000 00000000 00000000 ................
```

mount (Select a Mount Point)

Selects a mount point for the I (Display Inode Information) service command.

Command Synopsis

```
mount mountpoint
```

N (Display File System Parameters)

The N command displays the File system parameters screen. The screen shows mount-point parameters, the superblock version, and other file system information. Oracle support staff may request that you run this command in support of debugging.

Command Synopsis

```
N
:N
```

Controls

Ctrl-B (^b)
Move back one stage request.

Ctrl-F (^f)
Move forward one stage request.

Ctrl-D (^d)
Move down one in the status queue.

Ctrl-U (^u)
Move up one in the status queue.

Ctrl-I (^i)
Display detailed status interpretations

Example

```
File system Parameters           samu           version time date

mount_point    : /samma1         partial        : 16k
server         :                 maxpartial     : 16k
filesystem name: samma1          partial_stage  : 16384
eq  type       : 100 ma          flush_behind   : 0
state version  :    0   2        stage_flush_beh: 0
(fs,mm)_count  :    3   1        stage_n_window : 8388608
sync_meta      : 0               stage timeout  : 0
atime          : default         dio_consec r,w :    0    0
stripe         : 2               dio_frm_min r,w:  256  256
mm_stripe      : 1               dio_ill_min r,w:    0    0
high  low      :  50%  30%       ext_bsize      : 4096
readahead      : 1048576         def_retention  : 43200
writebehind    : 524288
wr_throttle    : 214725632
rd_ino_buf_size: 16384
wr_ino_buf_size: 512
maxphys        : ccc7400
config         : 0x0c08500c      config1        : 0x00000050
status         : 0x00000001      mflag          : 0x00000044
                                                : WORM Emulation Lite mode

Device configuration:
ty  eq      state  device_name                    fs  family_set
mm  101     on     /dev/zvol/dsk/dpool/vol1      100 samma1
md  103     on     /dev/zvol/dsk/dpool/vol2      100 samma1
md  104     on     /dev/zvol/dsk/dpool/vol3      100 samma1
```

P (Display Active Services)

The P command displays the Active Services screen lists the services registered with the Oracle HSM single port multiplexer. Oracle support staff may request that you run this command in support of debugging.

Command Synopsis

```
P
```

```
:P
```

Controls

Ctrl-B (^b)
Move back one page.

Ctrl-F (^f)
Move forward one page.

Examples

The following example shows the active services display.

```
Active Services              samu           version time date

Registered services for host 'server1':
  sharedfs.qfs2
  sharedfs.qfs1
  2 services registered.
```

R (Display SAM-Remote Configuration Information)

The R command displays the SAM-Remote screen, which shows the status of SAM-Remote configurations. Oracle support staff may request that you run this command in support of debugging.

S (Display Raw Sector Data)

The S command Sector data screen displays raw device data. Oracle support staff may request that you run this command in support of debugging.

When issued without a parameter, the S command displays sector data for the currently open device or prompts you to issue an open command if no device is open. When the command is issued with a user-assigned Oracle HSM equipment ordinal number, the command displays sector data for the specified device.

Command Synopsis

```
S
:S
:S equipment-number
```

Controls

Ctrl-B (^b)
Move back one page.

Ctrl-F (^f)
Move forward one page.

Ctrl-D (^d)
Move down the page.

Ctrl-U (^u)
Move up the page.

Ctrl-K (^k)
Step through the following display formats:

- raw
- inode
- archive
- label
- sblk

Examples

The following example shows an abbreviated **raw** data view for equipment ordinal **12**:

```
Sector:   00000012 (18)       raw        samu           version time date

00000000  00000000 00000000 00000000 00000000  ................
00000010  00000000 00000000 00000000 00000000  ................
00000020  00000000 00000000 00000000 00000000  ................
...
```

T (Display SCSI Sense Data for a Current or Specified Device)

The T command displays SCSI sense data screen. The screen shows the SCSI status of a current or specified SCSI device. Oracle support staff may request that you run this command in support of debugging.

When issued without a parameter, the T command displays SCSI sense data for the currently open device or returns No sense data available if no device is open.

When the command is issued with a user-assigned Oracle HSM equipment ordinal number, the command displays SCSI sense data for the specified device, if available, or returns No sense data available.

Command Synopsis

```
T
:T
:T equipment-number
```

Controls

Ctrl-B (^b)
Move back one equipment ordinal number.

Ctrl-F (^f)
Move forward one equipment ordinal number.

Examples

This example shows some of the sense data returned by a drive with equipment ordinal number 101:

```
SCSI sense data eq: 101 addr: 0xADDRESS     samu              version time date

Sense data:                          Additional sense bytes:
00       valid/code: not VALID       [18] 00000000
00       segment                     [22] 00000000
00       fmk/eqm/ili/key:No Sense    [26] 00000000
00000000 info[0-3]                   [30] 00000000
      00 add_len ...
```

U (Display Device Table in Human-Readable Form)

The U command displays the Device table as a human-readable screen.

When issued without a parameter, the U command displays the device table for the device with the lowest, user-defined Oracle HSM equipment ordinal number or the most recently specified ordinal number.

When the command is issued with an equipment ordinal number, the command displays the device table for the specified device.

Command Synopsis

```
U
:U
:U equipment-number
```

Controls

Ctrl-B (^b)
Move back one equipment ordinal number.

Ctrl-F (^f)
Move forward one equipment ordinal number.

Examples

The following example shows part of the device table for equipment ordinal **101**, a tape drive:

```
Device table: eq: 803   addr: 000043e0      samu            version time date

message: empty
Inquiry: HP        Ultrium 4-SCSI   rev:     A3K6 serial: 10000033

4d58000100000004 0000000000000000            0000003c delay
0000000000000000 mutex                       0000000f unload_delay
00004b80 next                                00000000 mtime
386c7476 set:  vtl800                        0000000a scan_tid
00003030                                     ffffffff slot
00000000                                     00000000 eod_position
00000000                                     00000000 stage position
03230320 eq/fseq                             00000000 next_read
```

```
02110211 type/equ_type                  00040000 def blk-size
0000     state                          00000708 pos timeout
00000000 st_rdev                        00000000 max blksize
0000002e ord/model                      9184e72a dflt cap
0000c2e8 mode_sense                     00000000 drvblksz
0000c0c0 sense                          00000000 fsn
00000000 space                          ffffffff mask
00000000 capacity                       0000 drv indx
00000000 active                         00 medium
00000000 open
00004000 sector_size
00000000 label_address
00000000 vsn:
00000000  Wed Dec 31 17:00:00 1969
00000000
00000000
00020000 status: ---------p
00000000 samnm:
7665642f name: /dev/rmt/2cbn
```

A

Glossary of Equipment Types

The value of the `Equipment Type` field of the Master Configuration File (`mcf`) identifies devices and device configurations within the Oracle Hierarchical Storage Manager and StorageTek QFS Software. Equipment types are specified as two-character codes. This glossary lists the codes for quick reference when working with the samples or when interpreting an existing `mcf` (for full details see the `mcf(4)` man page).

For convenience, the codes are divided into two sections and then listed alphabetically:

Recommended Equipment and Media Types

This section describes all of the equipment codes that you normally need: the generic equipment codes (`rb`, `tp`, and `od`) and codes for identifying network-attached library interfaces, archival disk volumes, cloud resources, and the Oracle HSM historian.

The generic equipment codes `rb`, `tp`, and `od` are the preferred equipment type codes for all SCSI-attached libraries, tape drives, and optical disk devices. When you specify a generic equipment type, Oracle HSM can automatically set the correct type based on SCSI vendor codes.

cl

Cloud media, the abstract media type that organizes cloud resources into logical media volumes suitable for archiving. Oracle HSM labels each volume with a 31-character volume serial number (VSN) of the form *`nameNumber`*, where:

- *`name`* is the customer-defined string of 4 to 20 alphanumeric characters that identifies the logical library (type `cr`) that owns the media.
- *`Number`* is a randomly generated number in the range `[0000000-9999999]`.

For additional information, see the `cloud` (7) and `sam-cloudd` (1m) man pages.

cr

A cloud library, the abstract equipment type that manages cloud media (volumes of type `cl`) by emulating a network-attached tape library with a configurable number of drives.

A parameter file defines the characteristics of the equipment, including the `name` parameter value that prefixes and uniquely identifies the type `cl` volumes that belong to the library.

For additional information, see the `cloud` (7) and `sam-cloudd` (1m) man pages.

dk
A disk-based file system that the Oracle HSM software uses as an archival volume. UFS, ZFS, QFS, and NFS file systems can serve as disk archives.

Disk volumes and volume serial numbers (VSNs) are defined in the `/etc/opt/SUNWsamfs/diskvols.conf` file. See the `diskvols.conf` (4) man page.

g*XXX*
Where *XXX* is an integer in the range `[0-127]`, a striped group of disk devices that is part of an `ma` disk-cache family set.

hy
The Oracle HSM historian, an optional, virtual library that maintains a media catalog, but has no associated hardware. Used for tracking exported media.

ma
A high-performance QFS file system that maintains file-system metadata on one or more dedicated `mm` disk devices. File data resides on separate `md`, `mr`, or `g`*XXX* data devices.

md
A disk device that stores file data for an `ma` file system or data and metadata for an `ms` file system. `md` devices store file data in small, 4-kilobyte Disk Allocation Units (DAUs) and large, 16-, 32-, or 64-kilobyte DAUs. The default DAU is 64-kilobytes.

mm
A disk device that stores file-system metadata for a high-performance `ma` file system.

mr
A disk device that stores file data for an `ma` file system. `mr` devices store file data in large, fully adjustable Disk Allocation Units (DAUs) that are multiples of 8 kilobytes in the range 8-65528 kilobytes. The default DAU is 64 kilobytes.

ms
A Oracle HSM file system that maintains file-system metadata on the same devices that store file data.

od
Any SCSI-attached optical disk. Oracle HSM sets the appropriate equipment type automatically using the SCSI vendor code.

of
An abstract media type that distinguishes foreign media from Oracle HSM file system media. Used when migrating files from a foreign file system to an Oracle HSM file system.

The `of` media type does not identify a physical equipment type. Never specify it in a master configuration file (`mcf`).

rb
Any SCSI-attached tape library. Oracle HSM sets the appropriate equipment type automatically using the SCSI vendor code.

rd
The SAM-Remote pseudo-device. In the Master Configuration File (`mcf`), the corresponding `Equipment Identifier` field has to contain the path to the pseudo-device (such as `/dev/samrd/rd2`). The corresponding `Family Set` field has to contain the hostname of the SAM-Remote server.

sc
A SAM-Remote client system. In the Master Configuration File (`mcf`), the corresponding `Equipment Identifier` field has to contain the path the SAM-Remote client-configuration file for the client. The corresponding `Family Set` field has to contain the family set name of the server. The `Additional Parameters` field must contain the full path to the client's library catalog file.

sk
An Oracle StorageTek ACSLS interface to a network-attached library. In the Master Configuration File (`mcf`), the corresponding `Equipment Identifier` field has to contain the path to the parameters file for the ACSLS interface. For more information, see the `stk(7)` man page.

ss
A SAM-Remote server. In the Master Configuration File (`mcf`), the corresponding `Equipment Identifier` field has to contain the path to the SAM-Remote server-configuration file. The corresponding `Family Set` field has to contain the family set name of the server, which must match the name used in the `Family Set` field of the `mcf` on the client.

tf
An abstract media type that distinguishes Oracle StorageTek T10000 and LTO volumes that are in Linear Tape File System (LTFS) format from volumes that contain Oracle HSM archive files.

The `tf` media type does not identify a physical equipment type. Never specify it in a master configuration file (`mcf`).

tp
Any SCSI-attached tape drive. Oracle HSM sets the appropriate equipment type automatically using the SCSI vendor code. No, however, that if you do use more specific equipment codes such as `li` and `ti`, you must do so consistently. If you specify `li` (LTO) tape equipment in the `mcf` file, for example, you cannot refer to the same equipment as `tp` equipment in the `archiver.cmd` file

z*X* (where *X* is a character in the range [0-9a-z])
An abstract media type that distinguishes foreign media from media controlled by Oracle HSM software.

Other Equipment and Media Types

The equipment types listed in this section are also supported.

Note that, in most cases, Oracle recommends identifying SCSI-attached libraries, tape drives, and optical disk devices using the generic equipment types `rb`, `tp`, and `od`. The generic equipment types tell Oracle HSM to identify the hardware dynamically, using SCSI vendor IDs. The type codes below are essential when migrating from one media type to another and may sometimes be useful for management purposes. But using them in a Master Configuration File (`mcf`), for example, hard-codes a static equipment configuration that may, at some point, no longer match the actual hardware.

ac
A Sun 1800, 3500, or L11000 tape library.

at
A Sony AIT-4 or AIT-5 tape drive.

cy
A Cygnet optical disk library.

d3
A StorageTek D3 tape drive.

dm
A Sony DMF library.

ds
A DocuStore or Plasmon optical disk library.

dt
A DAT 4-mm tape drive.

e8
An Exabyte X80 library.

fd
A Fujitsu M8100 128-track tape drive.

h4
An HP SL48 or SL24 library.

hc
An Hewlett Packard L9-/L20-/L60-series library.

i7
An IBM 3570 tape drive.

ic
An IBM 3570 media changer.

il
An IBM 3584 tape library.

li
An LTO-3 or later tape drive.

lt
A Digital Linear Tape (DLT), Super DLT, or DLT-S4 tape drive.

me
A Metrum library.

mf
An IBM Multi Function optical drive.

mo
A 5.25-in erasable optical drive.

o2
A 12-in WORM drive.

ov
An Overland Data Inc. Neo Series tape library.

pd
A Plasmon D-Series DVD-RAM library.

q8
A Qualstar 42xx, 62xx, or 82xx library.

s3
A StorageTek SL3000 library.

s9
An Oracle StorageTek 97xx series library.

se
A StorageTek 9490 tape drive.

sf
A StorageTek T9940 tape drive.

sg
A StorageTek 9840C or later tape drive.

sl
A Spectra Logic or Qualstar tape library.

st
A StorageTek 3480 tape drive.

ti
A StorageTek T10000 (Titanium) tape drive.

vt
A Metrum VHS (RSP-2150) tape drive.

wo
A 5.25-in optical WORM drive.

xt
An Exabyte (850x) 8-mm tape drive.

B

Product Accessibility Features

Users with low vision, blindness, color blindness, or other visual impairments can access the Oracle Hierarchical Storage Manager and StorageTek QFS Software (Oracle HSM) via the commandline interface. This text-based interface is compatible with screen readers, and all functions are controlled using a keyboard.

Glossary

This glossary focuses on terms specific to Oracle Hierarchical Storage Manager and StorageTek QFS Software and file systems. For industry standard definitions, please refer to the dictionary maintained by the Storage Networking Industry Association at `http://www.snia.org/education/dictionary/`.

active metadata server

See **metadata server (MDS)**.

addressable storage

All storage space that is user-referenced through an Oracle HSM file system. See **online storage**, **nearline storage**, and **offsite storage**.

admin set

A set of user- and/or group-owned storage that administrators use. Admin sets are typically created to administer storage for projects that involve users from several groups and span multiple files and directories.

archival media

Media that stores copies of the files in an Oracle HSM file system. Archival media can include removable tape cartridges, magneto-optical cartridges, disk file systems configured as archival volumes, and cloud storage volumes.

archival storage

Data storage space created on archival media.

archive set

A collection of files that are copied to archival media together, using a common set of policies and parameters. Set membership determines the number of copies made, the parameters of the copying process, and the media used.

The `archiver.cmd` file defines archive sets by file system, directory location, size, and/or user and group ownership. See the `archiver.cmd` (4) man page for additional details.

archiver

The Oracle HSM program that manages the process of copying files to archival media. See the `archiver` (1m) man page for additional details.

associative staging

Copying a group of files that are no longer resident in the disk cache from archival media back to the Oracle HSM disk cache when a user or application accesses any one

member of the group. Associative staging insures that files that are used together are staged together. File owners can associate any files that reside in the same directory by setting the associative-staging attribute on the related files. See the `staging` (1) man page, **staging**, and **disk cache** for additional details.

audit (full)

The process of loading 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 device that stores removable media cartridges, loads them into drives, and unloads them without operator intervention. An automated library contains cartridge storage slots, one or more drives, a transport mechanism for the cartridges, and, often, a mechanism for ingesting and exporting cartridges. See **direct-attached automated library**, **network-attached automated library**, **robot**, **transport**.

backup

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

block allocation map

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

block size

The size of the smallest addressable data unit on a block device, such as a magnetic tape cartridge or hard disk. The block size for a Linear Tape Open (LTO) cartridge is 256 kilobytes. For an Oracle StorageTek T10000 tape cartridge, the block size is 2048 kilobytes. For disk devices, block size is equivalent to *sector size*, which is typically 512 bytes.

buffered I/O

Writing and reading data to storage media, such as magnetic tape or disk, via an intervening segment of host memory, called the *buffer*. When an application writes to the storage device, the host lets the required changes accumulate in memory before writing them out to the media in a single operation, a process called *flushing the buffer*. When an application reads from the media, the host reads more data from the media than the application requested and stores it all in memory, in case the application subsequently requests the additional data.

By consolidating a large number of application I/O requests into a smaller number of hardware I/O operations, buffering improves I/O performance and uses storage hardware more efficiently, even when applications send or request data in suboptimal amounts or at inconsistent rates. Compare **direct I/O**.

cartridge

A container for data-storage media, such as magnetic tape or optical media. Also called a *volume*, a *tape*, a *piece of media*, or, loosely, a *VSN*. See **volume**, **volume serial number (VSN)**.

catalog

The Oracle Hierarchical Storage Manager software's record of the removable media volumes in an automated library. Volumes are identified and tracked using a *volume*

serial number. See **volume serial number (VSN)**, **historian**.

client-server

Adjective describing a distributed computer application that divides work between *servers* that specialize in providing files or services and *clients* that request files and services when performing particular tasks.

cloud library

Cloud storage that is accessed and managed as if it were a network-attached tape library containing a set of labeled media volumes. For additional information, see the `cloud` (7) man page.

cloud storage

Storage provided as an abstract, network service, without reference to any particular physical implementation or location. Cloud storage supplies users with an agreed level of service rather than a set of defined physical resources. Users and applications store and access data by addressing logical containers rather than physical locations.

A **cloud library** is the Oracle Hierarchical Storage Manager interface to cloud storage.

Cryptographic Framework

The Oracle Solaris Cryptographic Framework, a common store of algorithms and libraries that handle cryptographic requirements for users and applications. The Cryptographic Framework implements the PKCS #11 Cryptographic Token Interface (Cryptoki). See **Public Key Cryptography Standards #11 (Cryptoki)**.

Cryptoki

The PKCS #11 Cryptographic Token Interface. See **Public Key Cryptography Standards #11 (Cryptoki)**.

data device

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

Data Integrity Validation (DIV)

Data Integrity Validation, a feature of Oracle StorageTek tape drives that works with the Oracle HSM software to calculate and compare checksums during I/O.

During write operations, Oracle HSM calculates a four-byte checksum for each data block and passes the checksum to the drive along with the data. The tape drive then recalculates the checksum and compares the result to the value supplied by Oracle HSM. If the values agree, the drive writes both the data block and the checksum to tape. Optionally, when the write operation is complete, Oracle HSM can ask the tape drive to rescan the data, recalculate checksums, and compare the results to the checksums stored on the tape.

During read operations, both the drive and Oracle HSM read a data block and its associated checksum from tape. Each recalculates the checksum from the data block and compares the result to the stored checksum. If checksums do not match at any point, the drive notifies Oracle HSM that an error has occurred.

data mover

In an Oracle HSM shared file system, a client that is connected to tape drives and performs tape I/O on behalf of the metadata server. See **distributed I/O**.

DAU

See **disk allocation unit (DAU)**.

device logging

A configurable feature of Oracle HSM that provides specific error information for the hardware devices that support file systems.

device scanner

Software that periodically monitors the presence of all manually mounted removable devices and detects the presence of mounted cartridges that can be requested by users or by other processes.

direct access

Access to files on archival media without preliminary staging to the disk cache. The `-n` (*stage never*) staging attribute marks files for direct access. See **removable media file** and the `stage` (1) man page for more information.

direct-attached automated library

An **automated library** that is connected directly to the host via a SCSI interface. Oracle HSM software can directly control SCSI-attached libraries.

direct I/O

Reading from and writing to a storage device without using memory buffers on the host. Direct I/O can improve performance when transferring large amounts of block-aligned, sequential data. But otherwise, **buffered I/O** generally provides the best results.

directory

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

disk allocation unit (DAU)

In QFS file systems, the minimum amount of contiguous space that each I/O operation consumes, regardless of the amount of data written. The disk allocation unit thus determines the minimum number of I/O operations needed when transferring a file of a given size. The DAU should always be a multiple of the block size of the disk device.

The size of the disk allocation unit varies depending upon the QFS device type selected and user requirements. The `md` device type uses dual-allocation units: the DAU is `4` kilobytes for the first eight writes to a file and then a user-specified `16`, `32`, or `64` kilobytes for any subsequent writes, so that small files are written in suitably small blocks, while larger files are written in larger blocks. The `mr` and striped group device types use a DAU that is adjustable in increments of `8` within the range `[8-65528]` kilobytes. Files are thus written in large, uniform blocks that can closely approximate the size of the large, uniformly sized files.

See **block size**.

disk buffer

In Oracle Hierarchical Storage Manager SAM-Remote configurations, the buffer on the SAM-Remote server host that is used for archiving data from the client to the server.

disk cache

The disk-resident portion of an Oracle HSM file system where files are written, modified, and read. New and modified files are copied from the disk cache to archive media and may eventually be released from the disk. When users subsequently request non-resident files, the files are staged (copied) from the archival media to the disk. Individual disk partitions or an entire disk can be used as disk cache.

disk space threshold

The maximum or minimum level of disk cache utilization, as defined by an administrator. The releaser controls disk cache utilization based on these predefined disk space thresholds. See **high-water mark**, **low-water mark**, and **releaser**.

disk striping

Writing a file across several disks, thereby improving access performance and increasing overall storage capacity. See also **striping**.

distributed I/O

A feature of Oracle Hierarchical Storage Manager that lets the metadata server of a shared QFS file system delegate tape I/O to file system clients that are connected to tape drives. This reduces loads on the server and makes more efficient use of drives and SAN bandwidth. See **data mover**.

DIV

See **Data Integrity Validation (DIV)**.

drive

1. A electromechanical mechanism for transferring data to and from a removable-media volume, such as a magnetic tape cartridge.
2. An electromechanical, magnetic hard disk drive.
3. A solid-state device that emulates a disk drive. See **solid-state device**.

Ethernet

A packet-switched, local area network technology.

extent array

An array within a file's inode that defines the disk location of each data block assigned to the file.

family set

In Oracle HSM and QFS file-system configurations, a group of physical devices that function as a single logical devices, such as a set of data and metadata disks or an automated library and its associated drives.

Fibre Channel

The ANSI standard that specifies high-speed serial communication between devices. Fibre Channel is used as one of the bus architectures in SCSI-3.

file system

A logical structure that organizes data into a hierarchy of directories and files.

file system directives

In the Oracle HSM `archiver.cmd` file, archiver and releaser directives that are specific to a particular file system. File system directives follow global directives and include all directives between an `fs = filesystem-specifier` directive and the next `fs =` directive or the end of the file. File system directives override any global directives that may also apply. See the `archiver.cmd` (4) man page for details.

ftp

File Transfer Protocol, a network protocol for transferring files between two hosts. For a more secure alternative, see **sftp**.

global directives

In the Oracle HSM `archiver.cmd` file, archiver and releaser directives that apply to all file systems. Global directives appear before the first `fs = filesystem-specifier` directive. See the `archiver.cmd` (4) man page for details.

grace period

In a disk **quota**, the amount of time that the file system allows the total size of files belonging to specified user, group, and/or **admin set**s to exceed the **soft limit** specified in the quota.

HA-COTC

High-Availability Clients Outside the Cluster, the failover configuration for the metadata servers of a shared QFS file system that includes clients.

In an HA_COTC configuration, the file system is shared between active and potential QFS metadata servers and file-system clients. The metadata servers are hosted on a two-node, failover cluster. Clients are not hosted on cluster nodes. Solaris Cluster thus software ensures that the metadata servers remain available so that clients can access metadata and obtain I/O licenses. But clients are not configured for failover and cannot therefore compromise the integrity of the file system following a failure.

HA-QFS

High Availability QFS , the failover configuration that insures that a QFS unshared, standalone file system remains accessible in the event of a host failure.

In an HA-QFS configuration, the file system is configured on both nodes of a two-node cluster managed by Solaris Cluster software. At any given time, only one node mounts the QFS file system. If the node that is mounting the file system fails, the clustering software automatically initiates fail over and re-mounts the file system on the remaining node.

HA-SAM

High-Availability Oracle Hierarchical Storage Manager, the failover configuration for an archiving QFS file system.

In an HA-SAM configuration, Oracle Solaris Cluster software maintains the availability of the file system by insuring that the QFS metadata server and the Oracle Hierarchical Storage Manager application continue to operate even if a server host fails. The file system is shared between active and potential QFS metadata servers hosted on a two-node cluster that is managed by Solaris Cluster software and Data Services.

hard limit

In a quota, the amount of time that the file system allows the total size of files belonging to specified user, group, and/or admin set IDs to exceed the soft limit specified in the quota. See **quota**, **admin set**, **soft limit**.

high-water mark

1. The percentage disk-cache utilization at which Oracle HSM starts the releaser process, deleting previously archived files from disk. A properly configured high-water mark insures that the file system always has enough space available for new and newly staged files. For more information, see the `sam-releaser` (1m) and `mount_samfs` (1m) man pages. Compare **low-water mark**.
2. In a removable media library that is part of an archiving file system, the percentage media-cache utilization at which Oracle HSM starts the recycler process. Recycling empties partially full volumes of current data so that they can replaced by new media or relabeled.

historian

The Oracle HSM historian is a catalog of volumes that have been exported from the automated media libraries defined in the `/etc/opt/SUNWsamfs/mcf` file. By default, it is located at `/var/opt/SUNWsamfs/catalog/historian` on the Oracle HSM file-system host. For additional information, see **catalog** and the `historian` (7) man page.

hosts file

The `hosts.`*`filesystem-name`* file that identifies the hosts that can mount a shared QFS file system. See the `hosts.fs` (4) man page for details.

identity domain

In the Oracle Cloud, the authorization and authentication domain for a single tenant of the multiple-tenancy cloud environment. Within the identity domain, the tenant can define role-based user accounts and make subscribed Cloud services available to the specified users. Each identity domain is identified by a unique Identity Domain ID.

For additional details, see "Oracle Cloud Terminology", in Chapter 1 of *Getting Started with Oracle Cloud*, available at `docs.oracle.com`.

indirect block

A disk block that contains a list of storage blocks. Oracle HSM and QFS file-system metadata can contain up to three levels of indirect blocks. A first-level indirect block contains a list of blocks used for data storage. A second-level indirect block contains a list of first-level indirect blocks. A third-level indirect block contains a list of second-level indirect blocks.

inode

An index node, a 512-byte metadata structure that defines a file for the file system. An inode describes all the attributes associated with a file other than the name. The attributes include ownership, access, permissions, size, and the location of the file.

inode file

In a QFS file system, a metadata file called `.inodes` that contains the inode structures for all files resident in the file system.

Java Cryptography Extension (JCE)

A provider-based application program interface (API) that provides a uniform framework for implementing security features in Java applications. It is part of the Java Developer's Kit (JDK).

kernel

The program that provides basic operating system facilities. The UNIX kernel creates and manages processes, provides functions to access the file system, provides general security, and supplies communication facilities.

key label

A name that uniquely identifies a cryptographic key that is contained within a keystore. See **keystore**.

Key Management Appliance (KMA)

In an **Oracle Key Manager (OKM)** environment, a security-hardened server that manages and provisions encryption keys and authentication of the storage devices in accordance with prescribed policies. KMAs are clustered for high-availability.

Key Management Framework (KMF)

The Oracle Solaris Key Management Framework, a set of tools and application programming interfaces (APIs) for securely creating, exchanging, storing, and using public-key encryption objects, such as X.509 certificates and public/private key pairs. KMF supports both **Key Management Interoperability Protocol (KMIP)** and **Public Key Cryptography Standards #11 (Cryptoki)**.

For full information, see the encryption- and certificate-management information in the *Oracle Solaris Information Library* and the Solaris `pktool` (1) and `kmscfg` (1) man pages.

Key Management Interoperability Protocol (KMIP)

An extensible protocol that defines message formats for communicating with cryptographic key-management servers. KMIP lets client applications maintain cryptographic keystores on a server and lets the clients encrypt and decrypt data indirectly, without local access to the key.

The **Oracle Key Vault (OKV)** product implements KMIP, and the Solaris **Key Management Framework (KMF)** can create and manage KMIP keystores. See the `pktool` (1) man page for details.

For information on a related standard, see **Public Key Cryptography Standards #11 (Cryptoki)**.

Key Management Service (KMS)

A software service that lets client applications remotely create and store cryptographic keys on dedicated cryptographic servers. Clients can encrypt and decrypt data without direct access to the key. The key servers can track and manage all keys under a single, uniform set of policies. See **Public Key Cryptography Standards #11 (Cryptoki)**, **Key Management Interoperability Protocol (KMIP)**, and the `kmscfg` (1) man page.

keystore

1. In the Oracle Solaris Cryptographic Framework, persistent storage for token objects that serve as cryptographic keys.

2. Generally, any repository that holds security certificates, public keys, and private keys for encrypting and decrypting information.

LAN

Local area network.

lease

In a shared QFS file system, a function that grants a client permission to perform an operation on a file for a specified period of time. The metadata server issues leases to each client. Leases can be renewed as necessary.

library

See **automated library**.

library catalog

See **catalog**.

libsam

An application programing interface (API) library that lets applications manipulate Oracle Hierarchical Storage Manager operations and files stored in StorageTek QFS file systems. With the `libsam` library, software applications that run on the file-system metadata server can access and manipulate file systems using local function calls. See the `intro_libsam` (3) and `rest_libsam` (3) man pages for details. Compare **libsamrpc**, **rest_libsam**.

libsamrpc

An application programing interface (API) library that lets applications manipulate Oracle Hierarchical Storage Manager operations and files stored in StorageTek QFS file systems. The `libsamrpc` library makes remote procedure calls, so calling applications can run on any host on the network. It supports a subset of the `libsam` functions. Compare **libsam**, **rest_libsam**.

Linear Tape File System (LTFS)

An open standard for file systems on magnetic tape media. LTFS provides directory and file metadata that let users and applications use data as if it were stored on magnetic or solid-state disk.

local file system

1. A QFS file system that is not shared with other hosts.
2. A file system that is installed on a server for use by the operating system.
3. A file system that is installed on one node of a Solaris Cluster system and is not made highly available.

low-water mark

In an archiving file system, the percentage disk-cache utilization at which Oracle HSM stops the releaser process and stops deleting previously archived files from disk. A properly configured low-water mark insures that the file system retains as many file in cache as possible, for best performance, while making space available for new and newly staged files. For more information, see the `sam-releaser` (1m) and `mount_samfs` (1m) man pages. Compare **high-water mark**.

LTFS

See **Linear Tape File System (LTFS)**.

LUN

A Logical Unit Number, a logical partition of a physical device that is used as if it were an independent device.

mcf

The Master Configuration File that defines QFS file systems, data and metadata devices, and Oracle HSM archival data devices.

media

Material that stores data. Common storage media include magnetic tape, magnetic disks, solid-state devices, cloud services, and optical disks.

media migration

1. Copying files from one type or generation of archival tape media to a different type or a newer generation of the same type.
2. Copying files from old, worn archival tape media to new, replacement media.
3. A feature of Oracle Hierarchical Storage Manager 6.1 and later that automates the above processes.

metadata

Literally, data about data. In a file system, metadata is information about files and directories. It includes the locations of each file's data on storage media, file attributes such as file type (directory, regular file, character special file, block special file, etc), modification times, ownership, access permissions, and checksums. See **inode**, **indirect block**.

For additional details, see the Oracle HSM `sls` (1) and Solaris `ls` (1) man pages.

metadata device

In an Oracle HSM high-performance (type `ma`) file system, a dedicated storage device type (type `mm`) that stores only file system metadata. See the `mcf` (4) man page.

You can use solid-state disk, electromechanical magnetic disk, or hardware or software mirrored devices as metadata devices.

metadata server (MDS)

The host that controls Oracle Hierarchical Storage Manager and StorageTek QFS file systems. The metadata server manages the file-system metadata, maintains configuration information for file systems and related processes, such as archiving, participates in file-system I/O, and, in shared configurations, makes the file system available to clients.

Only one metadata server can be *active* at a time. But, in shared configurations, you can configure some or all clients as standby, *potential* metadata servers that can be activated should the active server fail or require disruptive maintenance.

mount point

The directory on which a file system is mounted.

multireader file system

An Oracle HSM file system configuration in which all file system hosts can read files but only one host can write them. For more information, see the `mount_samfs` (1m) man page.

nearline storage

Removable media storage that requires robotic mounting before it can be accessed. Nearline storage is usually less expensive than online storage, but it takes somewhat longer to access. See **automated library**.

network-attached automated library

A library that is controlled by a software package supplied by the vendor. A parameter file identifies network-attached libraries to the Oracle HSM software, and a special Oracle HSM media changer daemon provides the interface to the vendor software. Oracle StorageTek ACSLS software controls Oracle StorageTek network-attached libraries. See **automated library**.

NFS

Network File System, a file system that provides transparent access to remote file systems on heterogeneous networks.

OKM

See **Oracle Key Manager (OKM)**.

OKV

See **Oracle Key Vault (OKV)**.

offsite storage

Storage that is remote from the server and is used for disaster recovery.

online storage

Storage that is immediately available. See **disk cache**.

Oracle HSM

1. A common abbreviation for Oracle Hierarchical Storage Manager.
2. An adjective describing a QFS file system that is configured for archiving and managed by Oracle HSM software.

Oracle Solaris Key Management Framework (KMF)

See **Key Management Framework (KMF)**.

Oracle Key Manager (OKM)

Oracle's scalable, device-independent system for securely creating, exchanging, storing, and using storage encryption keys. See also **Key Management Appliance (KMA)**.

Oracle Key Vault (OKV)

A pre-configured, secured software appliance that centrally manages Oracle Advanced Security Transparent Data Encryption (TDE) master keys, other encryption keys, Oracle Wallets, Java Keystores, and credential files for Oracle database and application servers and clients that implement the Key Management Interoperability Protocol.

See **Key Management Interoperability Protocol (KMIP)**.

parameters file

The file that defines a network attached or cloud library for Oracle HSM. In the Oracle HSM master configuration file, mcf, the path to the parameters file is the equipment

identifier for the library, `/etc/opt/SUNWsamfs/family-set-name`, where `family-set-name` is the value of the family set name field in the `mcf`.

partition

A portion of a device or a side of a magneto-optical cartridge.

PKCS #11

See **Public Key Cryptography Standards #11 (Cryptoki)**.

pkcs11_kms

An Oracle Solaris implementation of a PKCS #11 Key Management Service that lets Oracle Key Manager clients communicate with remote servers using a private protocol. See **Oracle Key Manager (OKM)**, **Key Management Service (KMS)**, and the `pkcs11_kms` (5) and `kmscfg` (1M) man pages.

plugin

See **provider**.

potential metadata server

See **metadata server (MDS)**.

preallocation

The process of reserving a contiguous amount of space on the Oracle HSM disk cache for writing a file. Preallocation can be specified only for a file that is size zero. For more information, see the `setfa` (1) man page. See **disk cache**.

provider

In the context of the Oracle Solaris Cryptographic Framework, a supplier of cryptographic services that are used by *consumers*, such Oracle HSM and other applications, end users, or kernel operations. Providers include PKCS #11 libraries, kernel cryptographic modules, and hardware accelerators. Providers are also called Cryptographic Framework *plugins*.

pseudo device

A software subsystem or driver with no associated hardware.

Public Key Cryptography Standards #11 (Cryptoki)

A generic application programming interface (API) for working with cryptographic tokens, such as hardware security modules (HSM) and smart cards, public-key cryptography, and storage encryption. The PKCS #11 standard is currently maintained by the OASIS PKCS 11 Technical Committee.

The Solaris **Key Management Framework (KMF)** can create and manage PKCS #11 keystores. See the `pktool` (1) man page for details.

For information on a related standard, see **Key Management Interoperability Protocol (KMIP)**.

QFS

1. A QFS file system, an Oracle UNIX file system that offers high performance and high capacity. QFS file systems can be used on their own or with Oracle Hierarchical Storage Manager.
2. The StorageTek QFS product, which includes the file system without the Oracle Hierarchical Storage Manager software.

qfsdump

See **samfsdump (qfsdump)**.

qfsrestore

See **samfsrestore (qfsrestore)**.

quota

The amount of storage resources that specified user, group, or **admin set**s are allowed to consume. See **hard limit** and **soft limit**.

RAID

Redundant Array of Independent Disks, a disk technology that uses several independent disks to reliably store files. Depending on the architecture, it can protect against data loss should one or more disks fail and can provide higher throughput than individual disks.

recovery point

A compressed file that stores a point-in-time backup copy of the metadata for a Oracle HSM file system.

In the event of a data loss—anything from accidental deletion of a user file to catastrophic loss of a whole file system—an administrator can recover to the last known-good state of the file or file system almost immediately by locating the last recovery point at which the file or file system remained intact. The administrator then restores the metadata recorded at that time and either stages the files indicated in the metadata to the disk cache from archival media or, preferably, lets the file system stage files on demand, as users and applications access them. See **samfsdump (qfsdump)**.

recycler

An Oracle HSM utility that reclaims space when an archival tape volume is largely filled with *stale* copies (copies that no longer reflect the current state of the file). The recycler moves any remaining current file copies to other media and then relabels the volume. The relabeled volume can then be over-written with new files.

recycling

The process of moving current files from an archival media volume and relabeling the media for re-use. For details, see the `sam-recycler` (1m), `recycler.cmd` (1m), and `recycler.sh` (1m) man pages.

regular expression

A string of characters in a standardized pattern-matching language that is designed for searching, selecting, and editing other character strings, such as file names and configuration files. For full details of the regular expression syntax used in Oracle HSM file-system operations, see the Solaris `regex` and `regcmp` man pages.

release priority

The priority according to which the Oracle HSM **releaser** deletes a file from the file **disk cache** once the file has been successfully archived. See the `sam-releaser` (1m) man page for details.

releaser

A Oracle HSM 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 according to high and low thresholds. See the `sam-releaser` (1m) and `releaser.cmd` (4) man page for details.

remote procedure call

See **RPC**.

removable media file

In an Oracle HSM file system, a special type of user file that can be accessed directly from where it resides on a removable media cartridge, such as magnetic tape or optical disk cartridge. See **direct access**.

REST

Representational state transfer, the style of software architecture pioneered by the World Wide Web. REST emphasizes the roles of components, their interactions, and the ways that they represent data, rather than the internal implementation of components. Typical REST applications communicate via Hypertext Transfer Protocol (HTTP). REST is an alternative to RPC. See **RPC**.

rest_libsam

An application programing interface (API) that lets applications control Oracle Hierarchical Storage Manager operations and access files stored in StorageTek QFS file systems. The `rest_libsam` library provides a lightweight **REST** interface that operates over an authenticated HTTPS connection. It is implemented on top of the existing `libsam` library and supports a subset of the `libsam` functions. Compare **libsam**, **libsamrpc**.

robot

An **automated library** component that moves cartridges between storage slots and drives. Also called a **transport**.

round-robin

A data access method in which entire files are written to logical disks in a sequential fashion. When a single file is written to disk, the entire file is written to the first logical disk. The second file is written to the next logical disk, and so on. The size of each file determines the size of the I/O. See also **disk striping** and **striping**.

RPC

Remote procedure call, a mechanism that lets a client application execute subroutines that run under an independent server application.

SAM

A common abbreviation for Storage Archive Manager, the former name of the Oracle Hierarchical Storage Manager product.

SAM-Remote

An Oracle HSM client-server configuration that lets an Oracle HSM metadata server access an automated tape library that is controlled by another Oracle HSM metadata server. The client is configured with pseudodevices that represent the devices that the server makes available and uses a specified subset of the archive media on the server.

SAM-QFS

1. A common abbreviation for older versions of the Oracle Hierarchical Storage Manager product.

2. An adjective describing a QFS file system that is configured for archiving and managed by Oracle HSM software.

samfsdump (qfsdump)

An Oracle HSM command that backs up file system metadata to a dump file. See **recovery point**.

If the Oracle Hierarchical Storage Manager packages are not installed, the command is called `qfsdump`.

samfsrestore (qfsrestore)

A program that restores inode and directory information from a recovery point. See **samfsdump (qfsdump)**, **recovery point**.

SAN

Storage Area Network.

SAS

Serial-Attached SCSI.

SC-RAC

Solaris Cluster-Oracle Real Application Cluster (SC-RAC), a high-availability Oracle Database solution that use QFS file systems.

In an SC-RAC solution, Oracle RAC software coordinates I/O requests, distributes workload, and maintains a single, consistent set of database files for multiple Oracle Database instances running on the nodes of a cluster. In the SC-RAC configuration, Oracle Database, Oracle Real Application Cluster (RAC), and QFS software run on two or more of the nodes of a cluster managed by Oracle Solaris Cluster software. One node is configured as the metadata server (MDS) of a QFS shared file system. The remaining nodes are configured as potential metadata servers that share the file system as clients. If the active metadata server node fails, Solaris Cluster software automatically activates a potential metadata server on a healthy node and initiates failover. Since the QFS file system is shared and already mounted on all nodes, access to the data remains uninterrupted.

SCSI

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

seeking

Moving the read/write heads of a disk device from one disk location to another during random-access I/O operations.

sftp

Secure File Transfer Protocol, a secure implementation of `ftp`. See **ssh**.

shared hosts file

When you create a shared file system, the system copies information from the hosts file to the shared hosts file on the metadata server. You update this information when you issue the `samsharefs -u` command

Small Computer System Interface

See **SCSI**.

soft limit

In a quota, the maximum amount of storage space that a specified user, group, and/or admin set IDs can fill for an indefinite period. Files can use more space than the soft limit allows, up to the hard limit, but only for a short grace period defined in the quota. See **grace period**, **hard limit**, **quota**, **admin set**.

solid-state device

A storage device that uses electronically rewritable, non-volatile, NAND flash memory as the storage medium, such as a SAS-attached, solid-state disk drive (SSD).

Solid-state drives can provide significantly higher inputs and outputs per second (IOPS) and significantly lower latency compared to traditional magnetic hard drives. They are thus particularly good choices for use as the metadata devices of Oracle Hierarchical Storage Manager and StorageTek QFS , high-performance, `ma` file systems.

ssh

Secure Shell, an encrypted network protocol that allows secure, remote command-line login and command execution.

staging

In Oracle HSM file systems, the process of copying an archived file that is no longer resident in the disk cache from archival storage back to the disk cache. See the `staging` (1) and `stager.cmd` (4) man pages, **disk cache**, and **associative staging** for additional information.

Storage Archive Manager

The former name of the Oracle Hierarchical Storage Manager product.

storage slots

In an automated library, the storage bays that hold media cartridges that are not mounted in drives.

stripe size

During striped device access, the number of disk allocation units (DAUs) that a QFS file system writes before moving to the next device in the stripe, as specified by the `stripe=` mount option.

striped group

In a QFS file system, a collection of devices configured as a single logical device of type `gXXX`. See the `mcf` (4) man page for additional information.

striping

Writing files to multiple devices in parallel, so that each file is spread across all the devices.

QFS file systems support two types of striping:

- *Hard striping* is a permanent feature of the file system that you enable when you specify striped group (type `gXXX`) devices in the Master Configuration File (`mcf`) entries that define the file system.
- *Soft striping* is an optional feature that enable or disable when you mount the file system with the `stripe=` mount parameter.

Compare **round-robin**.

SUNW.hasam

A Solaris Cluster Data Services resource type that supports failover for the Oracle Hierarchical Storage Manager application. `SUNW.hasam` is included with the Oracle HSM software. See HA-SAM.

SUNW.HAStoragePlus

A Solaris Cluster Data Services resource type that manages failover of a server's local storage, so that critical state and dynamic configuration information remains available. `SUNW.HAStoragePlus` is included in the Solaris Cluster software as a standard resource type. See HA-QFS, HA-SAM.

SUNW.qfs

A Solaris Cluster Data Services resource type that supports failover for the metadata servers of a high-availability, StorageTek QFS file system. `SUNW.qfs` is included with the Oracle Hierarchical Storage Manager and StorageTek QFS software. See HA-QFS, HA-SAM, and HA-COTC.

superblock

A data structure in the file system that defines the basic parameters of the file system. The superblock is written to all partitions in the storage family set and identifies the partition's membership in the set.

tar

Tape archive. A standard file and data recording format used for archive images.

TCP/IP

Transmission Control Protocol/Internet Protocol. The internet protocols responsible for host-to-host addressing and routing, packet delivery (IP), and reliable delivery of data between application points (TCP).

timer

Quota software that keeps track of the period starting when a user reaches a soft limit and ending when the hard limit is imposed on the user.

transport

See **robot**.

`vfstab` file

The `vfstab` file contains mount options for the file system. Mount options specified on the command line override those specified in the `/etc/vfstab` file, but mount options specified in the `/etc/vfstab` file override those specified in the `samfs.cmd` file.

volume

1. On storage media, a single, accessible, logical storage area, usually addressed by a volume serial number (VSN) and/or volume label. Storage disks and magnetic tape cartridges can hold one or more volumes. For use, volumes are *mounted* on a file system at a specified mount point. See **volume serial number (VSN)**, **mount point**.
2. A magnetic tape cartridge that holds a single logical volume. See **cartridge**.
3. On a random-access disk device, a file system, directory or file that is configured and used as if it were a sequential-access, removable-media cartridge, such as a tape.

volume overflow

A capability that enables the system to span a single file over multiple **volume**s. Volume overflow is useful for sites using very large files that exceed the capacity of their individual cartridges.

volume serial number (VSN)

1. A serial number assigned to a tape or disk storage volume. A volume serial number can consist of up to six uppercase, alphanumeric characters, must start with a letter, and must identify the volume uniquely within a given context, such a tape library or partition. The volume serial number is written on the volume label.
2. Loosely, a specific storage volume, especially a removable media cartridge. See **cartridge**, **volume**.

WORM

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