AMASS for UNIX Version 5.6 Addendum

This addendum documents features that are new or enhanced in Quantum AMASS for UNIX version 5.6. Features that were new or enhanced in AMASS 5.4 and 5.5 are also included, as well as updated information for changes in previously documented functionality.

> **Note:** The features described in this addendum are not documented in the most recent release of the complete AMASS for UNIX documentation set, which was last updated for the 5.3 release.

This addendum includes the following sections:

New and Enhanced Features in AMASS 5.6

The following features are either new or have been enhanced in AMASS for UNIX version 5.6:

Daemon Updates

[Update to *Managing the AMASS File System* (6-00028-01), page 4-11.]

The libio daemon performs read/write operations by sending SCSI commands to drives. There are two types of libio daemon:

- The libio_tape daemon performs read/write operations on tape drives.
- The libio_optical daemon performs read/write operations on optical drives.

On occasion (for example, when a drive is repaired replaced), the libio daemon associated with a drive may need to be restarted. In previous versions of AMASS for UNIX, restarting a libio daemon required stopping and then starting the entire system, causing interruption of service.

In AMASS for UNIX version 5.6, libio daemons (libio_tape and libio_optical) can now be restarted independently. You can restart the libio daemon for any drive without affecting other daemons or restarting the system.

Restarting the libio Daemon

To restart a libio daemon, first make the affected drive inactive, and then restart the daemon associated with the drive:

1. Make the affected drive inactive. At the command prompt, type:

   drivestat -i *drive_number* [*juke_number*]

   where *drive_number* is the number of the drive and *juke_number* is the number of the juke containing the drive. (If you do not specify a juke number, the system assumes juke 1.)

2. Determine the process ID of the libio daemon. At the command prompt, type:

   ps -elf | grep libio

   The system displays process IDs for all currently running libio_tape and libio_optical daemons. Each daemon is followed by the juke and drive number, for example:

   libio_tape 1 2

   In the output, the first number is the juke number and the second number is the drive number. Use these numbers to identify the daemon associated with the affected drive, and then note the process ID for that daemon.

3 Stop the libio daemon for the affected drive. At the command prompt, type:

kill -15 *process_id*

where *process_id* is the process ID you identified in step 2.

4 When you are ready to start the libio daemon again, change to the **/usr/amass/daemons** directory.

5 Start the libio daemon for the affected drive. At the command prompt, type:

./libio_tape *juke_number drive_number*

where *juke_number* is the number of the juke containing the drive and *drive_number* is the number of the drive. (For an optical drive, type ./libio_optical *juke_number drive_number* instead.)

6 Repeat step 2 to verify that libio daemon for the affected drive has successfully started.

7 Make the affected drive active. At the command prompt, type:

drivestat -a *drive_number* [*juke_number*]

where *drive_number* is the number of the drive and *juke_number* is the number of the juke containing the drive. (If you do not specify a juke number, the system assumes juke 1.)

> **Note:** Before restarting the libio daemon, you must reset any exported environment variables specified in the **amass_start** script. If you do not do this, when the libio daemon starts, it will use the default value for all environment variables. To export a variable, at the command prompt, type:
>
> *variable_name=value*
> export *variable_name*
>
> where *variable name* is the name of the environment variable and *value* is the value of the variable.

Database Checking Performance

[Update to *Managing the AMASS File System* (6-00028-01), page 4-55.]

In AMASS for UNIX version 5.6, the dbcheck command invokes a new preprocessor when validating the AMASS file system database for data integrity. This preprocessor is called FADS, or Fast/Accurate Data Scanner.

FADS identifies the database records that require validation, and dbcheck validates only those records rather than the entire database. As a result, when dbcheck is run regularly, the FADS preprocessor significantly increases the speed of database validation without affecting accuracy or data integrity.

Running dbcheck with FADS

In AMASS 5.6, the FADS preprocessor is used by default when running the dbcheck command. The dbcheck -a option is assumed even if it is not specified, and you must include the full database path and name in the arguments. Also, if the AMASS file system is active, use the dbcheck -R option.

For example, if the file system is not active:

dbcheck -a /usr/filesysdb/filesv4

Or, if the file system is active:

```
dbcheck -a -R /usr/filesysdb/filesv4
```

The first time dbcheck is run, FADS creates companion files that are used to track the database records that require validation. During the initial run, all database records are validated. As a result, the first time dbcheck is run, the validation process will take as long as it did in previous versions of AMASS.

The next time dbcheck is run, the companion files are used to determine which records require validation. Only those records that are new or have changed since the last dbcheck will be validated. This results in a significantly shorter time for the validation process.

Setting up a dbcheck Cron Job

To realize the full speed benefits of the FADS preprocessor, run dbcheck frequently. The more often dbcheck is run, the fewer the number of records that require validation each time, and the faster the validation process.

To automatically run dbcheck at regular intervals, create a cron job. For example, to run dbcheck at 2:10 a.m. each day, add the following line to the **amass crontab** file:

```
10 2 * * * /usr/amass/utils/dbcheck –a –R /usr/filesysdb/filesv4
```

Running dbcheck in Legacy Mode

If desired, you can run dbcheck in legacy mode without using the FADS preprocessor. (When dbcheck is run in legacy mode, database validation takes significantly longer.)

To run dbcheck in legacy mode, use the following syntax:

```
dbcheck -a -L database_path
```

where *database_path* is the full path and name of the AMASS database to be validated.

Support for Running Multiple Admin Requests

[Update to *Managing the AMASS File System* (6-00028-01), pages 3-9, 3-61, and 3-158.]

With AMASS 5.6, administrators can run multiple admin commands simultaneously. For example, with the enhanced support, an administrator can run any of the following command combinations:

- driveclean + amassbackup
- driveclean + driveclean
- driveclean + volformat
- volformat + volformat

> **Note:** If an admin command uses a particular resource (such as a cleaning volume), a second command that requires the same resource will not start until the resource becomes available. While waiting for the required resource, the command prompt is unavailable.

New and Enhanced Features in AMASS 5.5

The following features are either new or have been enhanced in AMASS for UNIX version 5.5:

AMASS Infinite File Life

[Update to the AMASS *Infinite File Life* manual (6-00323-01).]

AMASS Infinite File Life (IFL) support is provided for LTO1, LTO2, and LTO3. AMASS IFL tracks the quality of AMASS volumes over time. Volumes are checked regularly and if a problem is determined, a replacement is scheduled. There are a number of criteria that determine when a volume is verified or replaced, for example: age, errors, Mbytes read, and Mbytes written. All of this functionality is scheduled automatically and will even verify media that is not actively being read or written to.

IFL support for LTO utilizes a generic drive type. This generic drive type relies on Log Sense Tape Alert page functionality to identify media failures. The user specifies which Tape Alert flags are used for verification and replacement of media. The selection of these flags are drive dependent and can be determined by referring to the SCSI manual for a given drive. Certification of other drive types with the generic IFL drive type will be provided in future releases of AMASS.

Commands

drivechk
Command modified to support LTO1, LTO2, and LTO3. Turn on debug (see config file below) to see new functionality.

logsense
This is a new command that tests the Log Sense tape alert functionality for a drive. The logsense command loads the specified volume into a drive and performs a series of tape alert commands. The tape alert values returned are displayed. Turn on debug (see config file below) to see new functionality.

logsense -u
Usage:

logsense [-u | -y] [-d *drive*] -v *volume*

Where:

-u Display this usage message (only).

-y Suppress confirmation messages.

[-v *volume*]...

-d Drive number.

-v Verify specific volume.

mqverify

Command modified to support LTO1, LTO2, and LTO3. Turn on debug (see config file below) to see new functionality.

volreplace

The volreplace command has been modified to move media being replaced to the library mailbox. At this point the suspect media can be removed from the library. Note that a message is printed in the AMASS **tac** log and on the command line output.

Config File

The following changes have been made in the AMASS IFL config file located at

/usr/amass/ifl/

Config

Debug

Existing variable used to provide debug output:

```
debug all
```

Generic drive designator

This configures IFL to utilize a generic drive type which uses Log Sense Tape Alert functionality. This will also override the behavior of the currently supported DTF and AIT drives. To enable, define the following line in the config file:

```
drive generic 1
```

Tape Alert Threshold Parameters

Specify which tape alert flags are used to determine when a volume needs to be verified or replaced. The flags available are from 1–64 and are based on a given drive type. Consult the drive SCSI manual for more details. For LTO1, LTO2, and LTO3 the following flag settings are recommended:

```
verify GEN TpAparam 3
verify GEN TpAparam 5
verify GEN TpAparam 6
verify GEN TpAparam 15
verify GEN TpAparam 18
verify GEN TpAparam 52
verify GEN TpAparam 53
replace GEN TpAparam 4
replace GEN TpAparam 19
```

Archiving of ifldaily Logs

The ability to better manage the logs created by the nightly ifldaily script has been added. Users can specify the number of days in which to retain the current set of logs and specify a directory in which to archive the older logs.

Two variables have been added to the AMASS IFL config file for this purpose:

- Number of copies to be retained in the logs directory:

  maximum no_of_logs 7

- Archive these logs if needed in this directory:

  archive path /usr/amass/ifl/logs/archive

Bypassing AMASS Startup Tests

[Update to *Managing the AMASS File System* (6-00028-01), page 4-10.]

When AMASS starts, a number of checks are run to validate the AMASS configuration. These tests are located in:

/usr/amass/tools/TESTS

If a given test takes too long to run, it can be excluded. To do so, add the script name to this file:

/usr/amass/tools/TESTS/exclude_tests

Specify the path relative to:

/usr/amass/tools/TESTS

For example, to exclude the **avolume.t** test case in the **/usr/amass/tools/TESTS/BACKUP** directory, add the following line to the **/usr/amass/tools/TESTS/exclude_tests** file:

BACKUP/avolume.t

When running the **amass_start** script, you will see this message:

Bypassing BACKUP/avolume.t test specified in
/usr/amass/tools/TESTS/exclude_tests

> **Note:** Users should understand the impact of not running a given test.

Disk Based Backup/ Restore of the AMASS Database

[Update to *Managing the AMASS File System* (6-00028-01), pages 3-8 and 3-19.]

This allows backup/restore of the AMASS database to/from disk in addition to tape/optical. There are several significant advantages to this support:

- Improved backup times thereby reducing the time in which AMASS is inactive.
- Ability to verify the alternate database on disk while mainline AMASS is still running. Therefore, users can check their database at any time without having to bring AMASS down. Users can schedule the backup to disk through cron, followed by database checks. The integrity of the database is then known even when AMASS runs for long periods of time.
- Ability to handle AMASS databases sizes that have exceeded the space available on media. This is typically an issue with optical media.
- Easy to navigate directory structure on disk allowing users to view complete set of full and partial backups.
- If errors were detected in a database, we can rehearse the changes required to the database by manipulating the database on disk.

- The amassbackup -x *dir* can now be run even if AMASS is not running.

Commands

Commands changed for this feature include:

- amassbackup and amassrestore have new -x option

> **Note:** When using the new -x option, it is important to use the same directory name for each backup. Otherwise, an accurate restore is not possible.

amassbackup -u
Usage:

amassbackup [-fvu] [-d *tapedev*] [-x *dir*]

Where:

-f	Perform full backup.
-v	Verbose.
-u	Display this usage message.
-d *tapedev*	Backup to specified tape device.
-x *backup_dir*	Backup to specified directory.

amassrestore -u
Usage:

amassrestore [-tuv] [-x *directory*] [-D *drivenumber*] [-d *tapedevice*] [-e *date*] [-s *date*] [-J *jukeboxnum*] [-L *label*] [-P *position*]

Where:

-t	List files on backup volume.
-u	Display this usage message.
-v	Verbose.
-x *directory*	Restore from supplied directory on disk.
-D *drivenumber*	Drive number, [1 - max drives].
-d *tapedevice*	Use specified tape device for restore.
-e *date*	Ending date of restore. Date format: MM/DD/YYYY MM/DD/YYYY hh MM/DD/YYYY hh:mm MM/DD/YYYY hh:mm:ss
-s *date*	Date of partial backup to skip. Date format: MM/DD/YYYY MM/DD/YYYY hh MM/DD/YYYY hh:mm MM/DD/YYYY hh:mm:ss
-J *jukeboxnum*	Find backup vol in this jukebox.
-L *label*	Find backup volume with this media label.
-P *position*	Find backup vol in this position.

- sysdbchk and inittaf have new -a argument to specify alternate database path.
- checkjournal, printjournal, dbcheck, dbedit, dchain, and keybuild were modified to accept path to alternate database.

Improved AMASS Shutdown

[Update to *Managing the AMASS File System* (6-00028-01), page 3-24.]

AMASS has been modified to allow user processes accessing AMASS to be interruptible. This allows AMASS to be shut down cleaner, avoiding costly reboots and database checks. Additionally, user processes which hang can now be terminated.

Improved Multi-Threaded Support

[Update to *AMASS Overview* (6-00026-01), page 1-6.]

AMASS has been enhanced to better handle multi-threaded processes and a larger number of single threaded processes in parallel. This provides improved throughput in multi-processor host environments.

Mediaread Able to Read Multiple Blocks

[Update to *Managing the AMASS File System* (6-00028-01), page 4-39.]

The mediaread utility has been modified to read more than one block. The -n option has been added for this purpose. Usage is as follows:

mediaread [-p] [-o *ofile*] [-n *nblks*] *drive block jukeno*

-p	Print amass header information on standard output.
-o *ofile*	Write requested block to ofile (default **/tmp/mediaread**).
-n *nblks*	Number of blocks to read (default 1).

Mixed Media

[Update to *Installing AMASS* (6-00027-01), page C-19.]

With this feature, one AMASS instance can utilize a single library with different drive/media types. You are provided with greater control over how drive resources are used. Drives can be dedicated to selected volume(s) or volume group(s).

Mixed media allows you to create an association between volumes and a group of drives. Volumes are identified by slot, volume group, volume number, or barcode. The association is defined in:

/usr/amass/etc/amass.rc

Additionally, you can change the contents of the **amass.rc** file and initiate a "refresh" by send a "SIGHUP" to libsched "kill -s HUP *pid_libsched*." The following message will be seen in the **tac_00** file:

```
LIBSCHED1[5595]: E5109(16)<00000>:scsilib10856: DEBUG:
Caught signal: [1] will re-read amass.rc
```

The next time libsched is asked to schedule a volume with a drive, the **amass.rc** file will be re-read.

Solaris sgen Driver Support

[Update to *Installing AMASS* (6-00027-01), page 2-36.]

AMASS 5.5 provides support for the Solaris **sgen** driver in addition to the AMASS juke driver. Support of the **sgen** driver simplifies the AMASS install process in fibre environments. In addition, the **sgen** driver provides support for a wider variety of fiber devices and environments. During the install process a user can select from the existing AMASS or juke driver or the **sgen** driver.

> **Note:** For Solaris 8, **sgen** driver version 1.2 or later is required. The 1.2 version is available from Sun as Patch Id: 110901-01.

New and Enhanced Features in AMASS 5.4

The following features are either new or have been enhanced in AMASS for UNIX version 5.4:

AMASS Read Only

[Update to *Managing the AMASS File System* (6-00028-01), page 3-24.]

AMASS read only allows users to set AMASS to a read only state. You can read from AMASS, but not write or delete. To change the state to read only we use the amasstat command. The options are as follows:

amasstat [-aicfswuy] [-t *sec*]

-a	Changes filesystem to active status.
-i	Changes filesystem to inactive status (reads and writes from/to open files continue).
-r	Changes filesystem status to read only.
-R	Changes filesystem status to read/write.
-c	Checks current status, but do not change.

-f Forces status change even if timeout occurs.

-s Suppresses messages from kernel.

-t Time, in seconds, to wait for AMASS to go inactive.

-u Displays this usage message.

-v Displays the current AMASS version.

-w Disables writes to AMASS when inactive (reads from open files continue).

-y Suppresses confirmation and informational messages.

If -a, -i, and -c are not specified, then the active status of the AMASS filesystem is toggled.

Cache

[Update to *AMASS Overview* (6-00026-01), page 1-9.]

AMASS supports up to 256 partitions and has been certified with up to 4 TB cache. The 4 TB cache budget may be spent on a single partition or dispersed among multiple partitions. In theory, each partition could use 4 TB cache, or up to 64 PB total; however, the total amount of cache among all partitions has only been certified up to 4 TB.

Enhanced Cache

There are enhancements to the AMASS feature Enhanced Cache Control. The core of AMASS is a cache in which data is held prior to migrating to/from media. The cache is broke up into cache blocks whose size is configurable. Without the Enhanced Cache Control feature turned on, AMASS reuses cache blocks which are the oldest on the cache.

With Enhanced Cache Control the cache blocks are allocated to cache block lists. Each of these lists can then be weighted to control from where cache blocks are acquired. The cache block lists are then associated with AMASS volume groups. AMASS volume groups allow users to group data in a AMASS directory with a set of volumes in the tape/optical library. When associating a volume group to a cache block list, the user can do so for read or writes. So once a file is written to cache, it is then migrated to tape, and could be made to persist in cache by having the file's volume group be associated with a highly weighted cache block list.

- Eight cache block priority levels are maintained within the AMASS files system.
- Cache block reallocation priorities are established on a volume group basis. No additional controls are provided for individual directories or files.
- For each volume group, two cache list priorities (one through eight) are assigned. One list is assigned for reads and the other list is assigned for writes.
- For a particular volume group, the priority can be changed dynamically for the read/write priority, but the new priority takes effect for cache blocks accessed or modified after the change.
- Each cache list has an additional weighting factor that is used to compute a score for the cache block list. The score for a particular list is the weighting factor multiplied by the access time of the oldest cache block on the list.

- A special weighting factor can be assigned specifying that all other cache lists should be searched before this list is searched. This stops the oldest cache block age from being used to create a score for a given cache list.
- The weighting factor can be changed dynamically and the new settings immediately take effect.
- If the weighting factor is assigned to multiple cache lists, an implied priority is established based on the cache list number (in other words, search the cache list with the lower priority level first).
- **The special weight factor for list 1 CAN NO LONGER be set to zero.** Initially, all of the cache blocks available to the system live in the first list and setting the weight factor on this list causes undesired effects. When setting the weight to zero, the system will only use this list if no other cache blocks in the system are available. The moment a file has been written to media and the cache blocks are put back on the free list, the blocks will be re-used immediately, even if there are thousands of unused cache blocks available on list one (with a weight of zero).
- A cache block is reallocated from the list with the highest score for all eight cache priority lists.
- All cache blocks exist on one list. Cache block reallocation occurs on a first-in-first-out basis.

Check Drive

[Update to *Managing the AMASS File System* (6-00028-01), page 4-10.]

AMASS logic handles media and communication with libraries. If drives are found obstructed, drives are cleared and AMASS continues. If the library is offline, a message is written to the log and AMASS waits until the library is back online.

This functionality can be disabled by setting the environmental variable, AMASS_DISABLE_CHKDRV.

Drive Mounting and Spinup Changes

[Update to *Installing AMASS* (6-00027-01), page A-3.]

Two environmental variables control how drives mount and spin up volumes. These environmental variables, like all AMASS environmental variables, must be set before you execute the amass_start command or add the variables to the **amass_start** script itself. Please contact Quantum Technical Assistance Center for further assistance on setting AMASS environmental variables.

The DO_ASYNC_STK_MOUNT and AMASS_SPINUP_DELAY variables are described as follows:

- DO_ASYNC_STK_MOUNT

  This environmental variable enables the functionality for StorageTek (STK) libraries to connect to AMASS using Quantum's eXtensible Device Interface (XDI) utility. When setting this variable, multiple drive loads can be in progress simultaneously because mounts are performed asynchronously. Any customer running with an XDI/STK configuration benefits from this option.

The environmental variable does not have to be set to any value, but only needs to be defined. For example:

```
DO_ASYNC_STK_MOUNT= ;
export DO_ASYNC_STK_MOUNT
```

- AMASS_SPINUP_DELAY

  Prior to AMASS 5.3, *DO_ASYNC_STK_MOUNT* was an environmental variable that served two functions. First, it enabled the XDI/STK asynchronous mounting. Second, it overrode the default spinup delay for a drive. However, starting with AMASS 5.3, the second function of DO_ASYNC_STK_MOUNT was moved to another environment variable called AMASS_SPINUP_DELAY.

  This environmental variable overrides the default spinup delay for any drive. The environmental variable should be set to the number of seconds needed to spin up a drive. To determine the recommended and default settings for your particular drive, contact the Quantum Technical Assistance Center for assistance. For example:

```
AMASS_SPINUP_DELAY=600;
export AMASS_SPINUP_DELAY
```

> **Note:** This variable only needs to be set if you see timeout messages while a drive is spinning up a volume.

drivestat

[Update to *Managing the AMASS File System* (6-00028-01), page 3-73.]

The drive status can be viewed or modified using the drivestat command. For example:

drivestat [-airju] [-y] [-e #] *device_number juke_number*

No arguments lists all drives.

-a Changes device to active status.

-i Changes device to inactive status.

-e # Changes idle_eject timeout to # for this device.

-r Resets device runtime state flags to initial condition.

-j Resets device eject state flags to initial condition.

-u Displays this usage message.

-y Suppress confirmation and informational messages.

If only device number is specified, then the active status of the specified device will be toggled.

Three command line options have been added to the drivestat tool:

-r Resets device runtime state flags to initial condition.

-j Resets device eject state flags to initial condition.

The -rj option address a drive in an active state, but has no work scheduled for the drive. This situation arises after a volume, stuck in a drive, has been cleared offline. After restoring the drive to an active state, the drive remains idle. The -rj option will remove the drive state

at the time of failure and restore the drive to its initial conditions. AMASS will assume and trust the operator to have cleared the drive and put the media away.

The -rj option cannot be run on an active drive.

-D Displays detailed AMASS debug logging.

The -D option has been a hidden option for several release cycles and one AMASS support person has instructed users how to use it. It is now part of the standard -u usage output.

Usage:

drivestat -D *device_number juke_number debug_option*

Where *debug_option* is defined as the arithmetic sum of the following:

0 = Resets all debug flags.

1 = Sets scsi_debug flag.

2 = Sets drvfd_debug flag.

4 = Sets scsi_verbose flag.

8 = Sets select_debug flag (in libsched select_drive function).

15 = ALL: 1+2+4+8

medialist

[Update to *Managing the AMASS File System* (6-00028-01), page 4-33.]

Additional command line arguments added to the medialist command allow users to list specific information rather than all media information. These additions extend to both SCSI and network attached libraries.

Usage:

medialist [-v] [-j *jukebox_number*] [-D] [-V] [-T] [-M] [-x] [-j]

Where:

-D Shows Drive information only.

-v Shows barcode information (Valid only for SCSI attached libraries).

-V Shows volume information only.

-T Shows Transport information only (not valid for network attached library).

-M Shows Mailbox information only (not valid for network attached library).

-x Turns on SCSI debugging.

-j Directs the medialist command to jukebox ID rather than the default of 1.

mediaread

[Update to *Managing the AMASS File System* (6-00028-01), page 4-39.]

The mediaread command provides options for displaying the AMASS header and for redirecting the output to a file other than the default file **/tmp/mediaread**. This affects all operating systems.

Usage:

mediaread [-p] [-o *ofile*] *drive block jukeno*

Where:

-p Prints the AMASS header information on standard output.

-o *ofile* Writes the requested block to **ofile** (default **/tmp/mediaread**).

Multiple volformat

[Update to *Managing the AMASS File System* (6-00028-01), page 3-158.]

AMASS does not allow multiple instances of volformat, volcopy, and volclean admin requests. The multiple volformat feature is the first step toward relaxing this restriction, thereby allowing AMASS administrators to get more done in less time. Multiple volformat means the AMASS administrator can execute simultaneous volformat requests. This feature is targeted at the AMASS administrator; the typical AMASS user will see no change in day-to-day operations.

The number of simultaneous volformat requests is controlled by ADMIN DRIVES, a number dictated by the AMASS administrator during installation. An ADMIN DRIVE is a drive selected by AMASS to satisfy an admin request. All active drives are eligible for selection as an admin drive. The administrator can set the number of drives to be equal to the number of physical drives in the library; in this case, user requests will not be serviced until the number of administrative requests is less than the number of available drives.

The AMASS administrator will set the number of admin drives to some value that allows optimal drive sharing between admin requests and user requests. The prompt for number admin drives is added to config_prod; since the install procedure (installamass) calls config_prod, this new prompt appears in all AMASS installations since the 5.3 release.

The volformat AMASS utility uses more than one admin drive. Option -d *number_of_drives* directs volformat to use the specified number of admin drives to format the media. In addition, the new range volformat specification allows for a more convenient way to specify several volumes to be formatted. Volume range can be specified as *n-m*, where the value of *m* is greater than *n*.

For example:

volformat -d 4 8-12

This formats volumes 8, 9, 10, 11, and 12 using not more than four drives at the same time.

Usage:

volformat [-b *blocksize*] [-c on | off] [-pquyd] [*volumelist*]

Where:

-p Forces physical format.

-q Quick formats faster, but writes are slower.

-u Displays this usage message.

-d Uses multiple (*num*) drives.

-b Specifies device blocksize.

-c on | off Specifies device compression mode.

-y Suppresses confirmation and informational messages.

Network Attached Libraries

[Update to *Managing the AMASS File System* (6-00028-01), page 4-10.]

Enhancements to the network attached library interface include:

- Reducing the number of "query drive" commands thus softening the load on the XDI interface. A function to test the need to inquire about library drive status is added. If there are no drive candidates to service a pending IO request, the library query drive command is not called. This will reduce the load on the network library interface.
- Releasing the shared memory semaphore across the library query drive command. Failure to do this will cause the entire AMASS system to lock up should a query drive command fail to complete (is hung).
- Adding a feature that will time out a hung XDI command. This happens frequently in the presence of library errors.
- Adding XDI failed mount/dismount retry option. XDI failed mounts and dismounts are retried three times before a drive/volume are taken out of service.

The combination of these four fixes will prevent AMASS from locking up and losing drive/volume resource in the presence of library errors. This new functionality is controlled by using the following environmental variables:

- AMASS_DISABLE_TEST_DRIVE - should the reduction in queries to the drive cause undesired effects, this environmental variable will return AMASS to pre-fix default state that continually pings the library interface.
- AMASS_DRIVESTAT_TIMEOUT=*new_val_in_seconds* - changes query drive time out from the default of 10 seconds to the value specified here.
- AMASS_STORE_TIMEOUT=*new_val_in_seconds* - changes the library mount/dismount timeout from a default 360 seconds to the value specified here.
- AMASS_DISABLE_LIBTIMEOUT - should the library timeout feature cause undesired effects, this environment variable will disable the feature.

> **Note:** These environmental variables can be set in the shell in which AMASS is started or put directly in the **amass_start** script.

vgimport

[Update to *Managing the AMASS File System* (6-00028-01), page 3-111.]

The vgimport function accepts a volume list from a file in addition to the command line.

The vgimport function's hidden -q option reduces the time it takes to import large volume groups.

Usage:

vgimport [-yun] [-p *path*] [-v *volumenumber(s)* | -f *volume_list_file*] *filename*

Where:

-q *quick*	Do not verify NSR label against volume.
-y	Suppress questions and informational messages.
-u	Prints this usage message.

-n	Parses the metadata file, but does not update the database.
-p *path*	Specifies volume group root directory name.
-v *volumenumber*	Specifies import volume number(s) (# #,# #-#).
-f *volume_list_file*	Obtains volumes from input file (one volume per file line).

volclean / voldelete

[Update to *Managing the AMASS File System* (6-00028-01), pages 3-132 and 3-146.]

The behavior of AMASS commands volclean and voldelete have been changed to require the UNIX user who is running the command to have permissions to remove the files associated with the volume being cleaned or deleted. The volclean/voldelete files that were not removed are listed. The removal of files in the AMASS system follows standard UNIX conventions.

> **Note:** With this change, volclean and voldelete are only functional while AMASS is running. Previously both commands were functional even when AMASS was not running.

volstat

[Update to *Managing the AMASS File System* (6-00028-01), page 3-206.]

Three command line options are added to the volstat command.

- -f Adds back the default size, vlk_maxxfr (as libio_tape does) and sets the volume to a status of not full. The volume will now eligible for write selection.
- -F Marks a volume as full. This is primarily for testing purposes but has a similar effect as marking a volume as read only.
- -s Enables setting the explicit volume capacity. The value input represents the desired volume size and is measured in megabytes.

For example:

```
o300m1# volstat -u
```

Usage:

volstat [-aifFsuy] *volume_number*

Where:

-a	Changes volume status to active.
-i	Changes volume status to inactive.
-f	Sets a tape volume to NOT FULL. Adds the value vlk_maxxfr found in **amass/scripts/amassconfig2** to the volume capacity.
-F	Sets a tape volume to FULL.
-s	Sets volume available (AVAIL) size (MB). volstat -s *available_size_in_MB vol_num*

-u Displays a usage message.

-y Suppresses confirmation and informational messages.

If -a or -i are not specified, then the active status of the specified volume is toggled.

Updated Information for AMASS 5.6

The following updates consist of changes to existing functionality that were previously documented in AMASS 5.3:

fileincache / fileonmedia

[Update to *Managing the AMASS File System* (6-00028-01), pages 3-79 and 4-57.]

For the -v option, the fileincache command now prints the resulting status to **stdout**, not **stderr**.

For the -v option, the fileonmedia command now prints the resulting status to **stdout**, not **stderr**.

healthcheck

[Update to *Managing the AMASS File System* (6-00028-01), page 3-81.]

The healthcheck command tests several AMASS components. The command does not correct any malfunction that it encounters, and only advises you that a test failed.

The -f *path* option tests the transfer procedure. The next available drive is used for this test.

- Enter the absolute path (beginning with the mount point) for the file you want AMASS to use for this test. For example: /archive/test/jasonfile
- The volume where the file resides must be online and active in the database.
- AMASS must be able to read the volume within 10 minutes or the test fails.

Resolving Check Drive Problems with Optical Media

[Update to *Managing the AMASS File System* (6-00028-01), page 4-28.]

The mediaeject command can fail to eject media in the throat of the drive. The resulting failure can cause a debris flow that may inactivate resources. This problem can occur if the media in the throat is not a member of the supported media types under which the test for media in the drive is executed.

To resolve this problem, test the drive again for media not in the drive after the failed mediaeject command.

voldelete / volslot / volclean

[Update to *Managing the AMASS File System* (6-00028-01), pages 3-132, 3-146, and 3-198.]

The AMASS commands voldelete, volslot. and volclean have been enhanced to support the -f (force) flag. This support allows you to delete a volume from a non-existent juke.

> **Caution:** Because the force (-f) flag bypasses safety checks to allow access to the database, this option should only be used by trained personnel or with the guidance of the Quantum Technical Assistance Center.

Removing Block Size and Compression Restrictions for Volume Groups

[Update to *Managing the AMASS File System* (6-00028-01), pages 1-13, 3-117, and 3-158.]

Volume groups can now contain volumes with different block sizes and compression settings. Previously, AMASS required volumes in a volume group to use the same block size and have the same compression flag. You can now choose to remove one or both of these requirements.

To change volume group compatibility settings, set the volMvBlkIgnore and volMvCmpIgnore variables in the **/user/amass/etc/amass.rc** file.

For example:

```
volMvBlkIgnore=TRUE
volMvCmpIgnore=FALSE
```

Usage:

volMvBlkIgnore=[FALSE | TRUE]

FALSE	(Default) Requires all volumes assigned to a block group to have the same block size.
TRUE	Allows volumes with different block sizes to be assigned to the same volume group.

volMvCmpIgnore=[FALSE | TRUE]

FALSE	(Default) Requires all volumes assigned to a block group to have the same compression flag.
TRUE	Allows volumes with different compression flags to be assigned to the same volume group.

> **Caution:** If you assign volumes with different block size and compression settings to the same group, then at a later time set volMvBlkIgnore or volMvCmpIgnore to FALSE, AMASS does not automatically remove any volumes from the group. If you want the volume group to contain only volumes with the same block size and compression settings, you must manually move volumes out of the group. All volumes in the group are still available for I/O, but if you add a new volume to the group, its block size or compression flag must match the first volume in the group.

Using the vglist Command

The vglist command displays information only for the first volume in a volume group. This is the case even if volumes in the group have different block sizes or compression settings.

When the vglist -a option is used, asterisks appear next to the blksiz column if the volMvBlkIgnore variable is set to TRUE. Also, asterisks appear next to the cmpr column if the volMvCmpIgnore variable is set to TRUE.

The asterisks alert you that volumes in the group might have different block sizes or compression settings. For example:

```
# vglist -a 800
```

VOLGRP	Juke	.....	**blksiz	cmpr**	vols	used	avail
800	1	.....	32768	On	4	9235	516344

Using the volformat Command

The volformat command can now automatically determine the largest block size allowed by a device, and then format a volume using that block size. To enable or disable this feature, set the amass_use_maxblks variable in the **/user/amass/etc/amass.rc** file.

For example:

```
amass_use_maxblks=TRUE
```

Usage:

amass_use_maxblks=[FALSE | TRUE]

FALSE	(Default) A volume in a volume group is formatted using the block size used by the first volume in the volume group. (If the volume group is empty, a predefined block size based on device type is used.)
TRUE	A volume in a volume group is formatted using the largest block size allowed by the device (but not exceeding the block size specified by the MAXIOSIZE variable).

> **Note:** Using the volformat -b option allows you to specify the block size to use when formatting regardless of what amass_use_maxblks is set to.

Updated Clean Request Process

[Update to *Managing the AMASS File System* (6-00028-01), page 3-61.]

In previous versions of AMASS, a drive cleaning request would fail if another admin request was already being processed or if no cleaning volume was available (for example, because another drive was being cleaned). When a drive cleaning request failed, the affected drive was taken offline.

In AMASS 5.6, multiple admin requests are supported. This means that multiple drive cleaning requests are queued and are completed without failing or causing the drive to go offline.

Using the AMASS_CLN_INTERVAL Variable

The AMASS_CLN_INTERVAL environmental variable is new in AMASS 5.6. Use this variable to specify how often tape drives are automatically cleaned. This variable is set in the **/usr/amass/tool/amass_start** script.

Use the AMASS_CLN_INTERVAL variable to set the cleaning interval for the following drive types:

- LTO family
- DLT
- STK 3590
- IBM 3592
- STK9940

> **Note:** For DTF or AIT drives, use the DRV_CLEAN_COUNT variable to specify the number of tape mounts between cleanings, or use the DRV_CLEAN_OFF variable to turn off cleaning.

For example:

```
AMASS_CLN_INTERVAL=300
```

Usage:

AMASS_CLN_INTERVAL=<*tape mounts*>

<*tape mounts*> The number of tape mounts after which the drive is scheduled to be cleaned. A value of 0 turns automatic drive cleaning off. (If there are I/O requests already in queue for the drive, the number of actual tape mounts before cleaning occurs may be higher than this value.)

> **Note:** To use the AMASS_CLN_INTERVAL variable, there must be a cleaning volume group that contains at least one cleaning volume.

Updated Information for AMASS 5.5

The following updates consist of changes to existing functionality that were previously documented in AMASS 5.3:

- Return Codes for fileonmedia Command

Return Codes for fileonmedia Command

[Update to *Managing the AMASS File System* (6-00028-01), page 3-79.]

Three return codes are added to the fileonmedia command:

```
fileonmedia -u
```

Usage:

fileonmedia [-uv] *fullpath_filename*

-u Display this usage message.

-v Verbose.

Return codes:

0 The file is NON-RESIDENT on media.

1 The file is RESIDENT on media.

255 No file given or file does not exist.

Updated Information for AMASS 5.4

The following updates consist of changes to existing functionality that were previously documented in AMASS 5.3:

Automatic Drive Cleaning

[Update to *Managing the AMASS File System* (6-00028-01), page 3-61.]

If a 3590 OR 3592 drive is **installed in an IBM 3494 library**, automatic drive cleaning by AMASS is **not supported** because the library has its own internal drive cleaning capabilities.

If a DLT or StorageTek 9840 or 9940 drive is installed in a **StorageTek 97xx or Lxx library**, automatic drive cleaning by AMASS is **not supported** because the library has its own internal drive cleaning capabilities.

Logging Mechanism

[Update to *Managing the AMASS File System* (6-00028-01), page 4-49.]

The AMASS logging mechanism has been changed to use different file and directory names.

The **/usr/amass/logs** directory is used instead of the **/usr/amass/emasslogs** directory.

The **etac** directory name has been changed to **tac**.

Log file names have been changed from **el_ETAC_XX.XXXXXX** to **tac_XX.XXXXXX**.

MAXLTIME / MINLTIME

[Update to *Installing AMASS* (6-00027-01), page A-9.]

The MAXLTIME value is the maximum amount of time, in seconds, that a volume can stay in a drive with other I/O pending. The interval starts at the first I/O for the volume. The MINLTIME value is the minimum amount of time, in seconds, that

a volume remains in a drive—as long as no requests are pending for the current volume—before AMASS ejects the volume. A volume is only ejected if the drive needs to load another volume. The previous logic for MINLTIME did not correctly determine that there was still an I/O pending on the loaded volume. Now the logic adds a function to analyze the volume I/O queue. The queue returns a logical yes/no answer.

Yes The I/O is pending on the volume.

No The volume is idle.

With this change in place, full MAXLTIME allocations are observed and I/O throughput is greatly improved.

sysperf -f

[Update to *Managing the AMASS File System* (6-00028-01), page 3-97.]

The sysperf -f command provides a way to direct output to a file in an unrestricted and suitable format to import into a spreadsheet or other data analysis tool. Table 1 lists the data and examples of the raw output of the sysperf command.

Table 1 sysperf -f Command Data

	Data Description	Example Data
1	Date	09/13/2004
2	Time	15:08:14
3	Read requests queued	0
4	Number of volumes in read queue	0
5	Write requests in queue	13
6	Number of volumes in write queue	4
7	Number of fnodes	544
8	Used fnodes	23
9	Free cache blocks	72
10	Dirty cache blocks	2115
11	Cache blocks queued	44
12	Cache blocks pending	1
13	Cache blocks done	3
Drive 1	Drive 1: I/O perf of drive 1 [bytes/interval]	25165824
	Drive 1: Read/write mode	W
	Drive 1: Volume number	3933
Drive 2	Drive 2: I/O perf of drive 2[bytes/interval]	23068672
	Drive 2: Read/write mode	W
	Drive 2: Volume number	549

Note: If -c option is not selected, cache blocks related to fields 11–13 are set to zero.
If -k option is not selected, -fnode related fields (7, 8) and cache fields (9–13) are set to zero.
[bytes/interval]: Number of bytes in the interval (given at sysperf -ck -f *logfile interval*).

List of Acronyms

The acronyms listed in Table 2 are used throughout the AMASS documentation set.

Table 2 List of Acronyms

ACSLS	Automatic Cartridge System Library Software
AIO	Asynchronous Input/Output
AMASS	Archival Management and Storage System
AML	Automated Media Library
API	Application Programming Interface
CR	Change Request
DAS	Distributed AML Server
FC	Fibre Channel
FCR	Fibre Channel Router
FTIF	Foreign Tape Import Feature
GA	General Availability
HBA	Host Bus Adapter
IFL	InfiniteFileLife
LTO	Linear Tape Option
LUN	Logical Unit Number
NFS	Network File System
SIO	Streaming Input/Output
Scalar DLC	Scalar Distributed Library Controller
SDLT	Scalar Digital Linear Tape
SNC	Storage Network Controller

STK	StorageTek
XDI	eXtensible Device Interface

Contacting Quantum

More information about this product is available on the Customer Service Center website at www.quantum.com/csc. The Customer Service Center contains a collection of information, including answers to frequently asked questions (FAQs). You can also access software, firmware, and drivers through this site.

For further assistance, or if training is desired, contact the Quantum Technical Assistance Center:

North America	+1 800-284-5101
UK, France, and Germany	00800 4 QUANTUM
EMEA	+44 1256 848 766
World Wide Web	www.quantum.com/support