

CACHE-A

ARCHIVE APPLIANCES

Cache-A FAQ

Last updated: August 26, 2011

How long does it take to archive with Cache-A appliances?

This really depends on a variety of factors, but mostly upon how your content is connected to the appliance. The following table will give you an idea of how long this may take in your specific environment. Typical data rates shown in the table can vary significantly in your environment (some users have archived as much as 80MB/s over GbE networks), so these are only rough guidelines. Another thing to keep in mind is that archiving big video files is very efficient and will be much faster than archiving lots of small files, all other things being equal.

Data Rate (MB/S)	Typical Of	100 GB Archive	LTO-4 (740 GB) Archive	LTO-5 (1425 GB) Archive
20	Slow network	85.3 Min	10.5 Hr	20.3 Hr
40	Average network	42.7 Min	5.3 Hr	10.1 Hr
60	Fast network	28.4 Min	3.5 Hr	6.8 Hr
75	Firewire drive	22.8 Min	2.8 Hr	5.4 Hr
110	eSATA RAID	15.5 Min	1.9 Hr	3.7 Hr

What kind of tapes can I use in my Cache-A archive appliance?

The LTO Consortium has dreamed that all LTO tape drives will read and write the current generation and the previous generation and read back two generations. Thus an LTO-5 drive can fully utilize (read and write) both LTO-4 and LTO-5 tapes and read-only LTO-3 tapes.

The LTO Consortium has also provided extremely tight specifications and qualification tests for any media claiming to be LTO Ultrium tape. So we expect tape from any qualified LTO manufacturer to work well in our equipment. That said, we have seen forgeries of LTO media that in fact did not work properly, so we recommend you obtain your tape from a reputable reseller you can trust.

It is also worth noting that Cache-A appliances can only read tapes written in Cache-A appliances unless using command-line tools - to learn how [[see this tech brief](#)]

Cache-A is sometimes asked if we support WORM media. WORM stands for Write Once Read Many; in other words, once you archive to this media, your data can never be erased.

Because we overwrite our previous TOCs (Table Of Contents) at the beginning of new sessions, WORM media is **not** compatible with Cache-A devices.

Can I archive from my P2 or SxS camera memory card directly to LTO Tape?

While this will work with most USB memory card adapters, Cache-A archive appliances do not yet have drivers to read these types of cards from our ExpressCard slot at this time.

For most workflows this has turned out to not be a problem as users need to have both archive and working copies of acquisition content. Thus many users are happily off-loading with their laptops or edit station using software like [ShotPut Pro](#) to make copies simultaneously to their local and or connected external drives and to a Cache-A at the same time.

What is the fastest way to make a duplicate tape in a Cache-A appliance?

The fastest workflow for making two copies of a tape is to start your archive via the VTAPE, thus creating a copy of all the content on the Cache-A local hard drive on the way to the tape. Then, simply select **Eject** (not **Erase**) - then insert a fresh tape, selecting **Copy** in the dialog that appears upon insertion. Note that this should only be done with a new or erased tape as any existing information will be replaced.

If you need to duplicate a previously written tape, you will have to do a **Recover-All** to get that content onto the VTAPE, and then proceed as described above.

How can I easily manage archiving my Avid projects?

Short of using Avid's full-up Interplay solution, they certainly don't make it easy to manage your content in a simple manner as they keep all their files using undecipherable file-names in a massive "Avid Media Files" directory.

If you are in a Unity or other shared storage environment, some users find that a good workflow is to set up workspaces on the shared storage on a project by project basis. This then makes it easy to archive off and delete those projects when completed.

In a non-workspace oriented user configuration, if you are simply looking to archive projects in toto with reasonable organization and project naming, a tool that some users

have found to be helpful is [Media Mover](#) which can organize your media into project specific folders that are then easily archived in an organized manner.

What's the best way to label my LTO tapes?

If you are shelf-storing your tapes and never plan on getting an automated library, then there is no reason to use anything but "human readable" labels like you would use on a video tape. If, on the other hand, you plan on migrating to a library then it might be a good idea to think about using barcodes to be ready for that eventuality.

If you are currently using a library, then we strongly recommend barcode labels as this greatly reduces the time it takes for the library to inventory its tapes. With barcodes, the robot can scan the tapes without moving them out of the magazine - without, each tape must be pulled and loaded into the drive, one at a time, to obtain the cartridge's media ID.

If you are going to use barcodes, they should conform to the LTO spec [[click here for IBM's spec](#)] and we recommend that they be printed professionally as problems have been reported with user-printed labels. That being said, here is a link to an interesting site that provides the ability for you to print your own labels:

<http://tapelabels.librelogiciel.com/>. We recommend that you use a well known tape label supplier like [Tri-Optic](#) or obtain pre-printed tapes from a reputable tape supplier like [Fujifilm](#).

Because there is potential that a label could come off inside the drive mechanism, manufacturers strongly recommend that only the outside spine or somewhere on tape box be used for labeling. If you have quality stickers and want to take the risk, the top of the tape would be the least dangerous area.

Can I erase a single file off of an LTO tape? Can I erase and reuse an old tape?

Our solutions do not allow you to erase any fraction of a tape - erasing is an all-or-nothing proposition. If you really, really need to remove one or more specific files from an archive, the only choice is to restore all the data on the tape, remove the offending files, and then re-archive everything left.

It is perfectly OK to erase and reuse a tape. However, you should keep an eye on the total number of insertions as reported by our **Tape Information** command. While LTO media is rated for something like a million passes, just archiving to fill a tape once uses 80 passes (on LTO-5) due to the fact that tracks are 'serpentine' and run many times up and back down the tape. Thus, if you have many archive sessions on a tape, these can add up quickly - we recommend users think about retiring tapes after 50-100 insertions.

Customers have also asked us if erasing overwrites the data so no one could see it; it does not. When we erase a tape, we simply write to the cartridge memory so that any LTO drive that looks at the tape will think it is empty. We also write an EOD mark (end of data) at the beginning of the tape. It requires very special low level software to read data on a tape past the EOD mark so even we have to go to outside labs to have them use their lab tools to recover such data. So, it is highly unlikely (but not impossible) that anyone other than a tape lab would be able to pull data from an erased tape.

How can I use Quantum A-Series tapes or other data tapes with Cache-A appliances?

Cache-A did license technology from Quantum to read A-Series LTO tapes, however the software effort to actually decode their proprietary format is so large that we may never be able to undertake it. If we ever do, at best, all we will support would be a recover-all.

In lieu of this ability, we offer a workflow developed for migrating data and is probably better and safer as long as your Quantum drive is still operational - please see the link under "Tech Brief Downloads" in the left column of this page - if your Quantum doesn't work any more let us know, Quantum is putting together a loaner program.

Cache-A appliances can in fact be used to restore data from any POSIX 'tar' written tape, although this will require a trip to the command line via Terminal (Mac), PuTTY or equivalent (PC), or a direct attached monitor and keyboard - again check out [[this Tech Brief](#)] for more information.

How often must LTO tape drives be cleaned?

Under normal circumstances, cleaning tapes should only be used when the drive so advises. When the drive needs cleaning, the orange **Clean** LED on the tape drive will flash. Only insert a cleaning cartridge into the tape drive when the LED flashes. This nominally will happen after about 200 full tape archives.

In circumstances where you see repeated tape drive errors, it may be advisable to run a cleaning session even without the LED warning, however this should be done judiciously as cleaning tapes do put additional wear and tear on the heads.

It is essential to use only LTO Ultrium Universal cleaning cartridges with our LTO Ultrium tape drives. Do not use swabs or other means of cleaning the heads. LTO Ultrium Universal cleaning cartridges can be used up to 50 times. The cleaning cycle can take up to 5 minutes, during which the orange **Clean** LED will be on and the green **Ready** LED will flash. When it has finished, the drive ejects the cartridge.

My archive session shows a different total number of bytes than my computer - is something wrong?

As long as your Transfer Summary shows no errors, everything should be fine. Any difference in reported sizes is caused by variations in the way different file systems and operating systems keep track of file sizes.

In the Mac environment, things changed in this regard substantially with the advent of OS 10.6 - according to MacWorld:

"There are, believe it or not, now two sets of entirely different terms for these two ways of thinking. For example, a gigabyte, or GB, is now defined as 1,000 bytes cubed, or 1,000,000,000 bytes. A gibibyte, or GiB, is equal to 1024 bytes cubed, or 1,073,741,824

bytes. (The same rules apply for megabyte and mebibyte, which are defined as 1000 bytes squared and 1,024 bytes squared, respectively.)

Wake up, you in back! Because here's where Snow Leopard comes in. In previous versions of Mac OS X, Apple used the 1024³ definition of GB. Rather than keep that math and start calling it GiB, Apple has started using the 1000³ definition." [[click here for more details](#)]

Even without this change in how sizes are calculated (i.e. a Mac running Tiger or a PC), there are still variations in how file sizes are calculated and you will most likely never see exactly the same number of bytes in both places.

How much data can I get on a tape and what about the compressed figures on the tape box.

LTO Tapes rate their capacity as both 'native' and 'compressed' and the compressed value is double the native capacity (LTO-4 is rated at 800 GB native / 1600 GB compressed while LTO-5 is 1500 GB Native / 3 TB compressed). LTO compression is lossless (every bit can be restored exactly) and is implemented inside the tape drive by hardware designed to the specifications for LTO.

Note that this compression is done transparently by dedicated hardware in the tape drive. While it is possible to turn this off, we never have seen a reason to do so. Thus Cache-A devices (like most every other LTO environment) always compresses data, always in the background, and always with no detriment to the user.

This compression is very much like zipping a file and some data compresses well, some does not. Typical data in the IT industry will compress at around a 2:1 ratio, so tape manufacturers quote double the native capacity for compressed values to reflect what an IT user might get on tape.

In the video world however, this lossless compression usually will not save much if any space, as most digital video formats already make very efficient use of every bit. However, this can vary considerably by video "essence" type. For instance, the compression you'll get with DPX raw tiffs of blue screen content may be quite a bit more than 2:1. DPX files containing RLE tiffs of forest footage may show very little compression. Mpeg file content will generally not compress at all. File headers and metadata do compress though, so lots of little mpegs may show savings whereas few large ones may not.

Any savings from compression shows up as a *negative* "Space Lost" value in a Cache-A Tape Info window. Actual space lost shows up as a positive value. You can end up actually losing space on tape for a number of reasons including making many sessions,

having an irregular data rate coming in, and actual wear and tear on the tape and/or drive.

It is important to understand that data tape drives are at their most basic level, analog systems and a number of factors can affect how much data will fit on a given tape. Keeping an uninterrupted supply of data to the drives will cause them to most efficiently utilize tape - any interruptions or competition for bandwidth on the network can cause loss of space when writing to tape, especially if your workflow is such that data is not mediated through our internal disk array. Even an exact copy of a data set is likely to have variations between tapes.

Normally, you should plan on a tape's capacity as slightly less than the stated native value. We recommend that you stay under 740 GB for LTO-4 and under 1425 GB for LTO-5. We also recommend that you keep an eye on the "Tape Info" dialog report to judge how much remaining space you have.