



Too old and too worn out for the demands of the data center?

It may be time to retire some of your tape cartridges.



Using damaged or worn out tape cartridges can be expensive. It could force you to recreate lost data, which can be extremely costly or even, in some situations, impossible. How do you know which cartridges should be retired? Happily, for users of LTO media¹ there is a very reliable historical data source. That source is a cartridge memory chip (CM) inside the cartridge, and it contains a wealth of information about the past life of the cartridge.

This article describes what information is available on the cartridge memory chip, what it means, how to access this chip, and how to make informed decisions based on these data.

Even cartridges that are used in a well-maintained environment will eventually wear out. Worn out or damaged cartridges can carry debris that may be transferred to tape path elements, such as roller guides and the read/write heads of a tape drive. This debris can scratch the tape surface as well as the read/write head and seriously damage the drive. Once the tape drive is contaminated the debris can be transferred to the next cartridge and the next, *ad infinitum*. A single bad cartridge can damage your tape drives and contaminate your (thus far) perfectly good cartridges. The longer it takes you discover a bad cartridge, the more extensive is the damage.

Like a virus, a bad cartridge ought to be quarantined before it has a chance to contaminate the rest of your system. "More and more tape systems these days only start reporting error activity after a certain threshold is reached," says Kevin Burton, Manager for Quality and Engineering for Graham Magnetics. "The problem is, the time between when the system starts reporting the error activity and when the tape crashes is way too close of a time frame."

Relevant data in the memory chip

The cartridge memory (CM) chip inside LTO cartridges contains a complete history of the cartridge. The most important statistical data in the CM include the number and type of records with errors, the number of bytes read/written, the number of loads, and the age of the cartridge.

1. Percentage of records with errors

Error statistics can be one of the most important indicators of cartridge health.

There are three types of errors: recovered data errors, unrecovered data errors, and servo errors. Of course, the more data that have been transferred to and from the drive, the more errors are to be expected. We therefore are concerned with the percentage of records that had errors rather than the absolute number of errors recorded.

¹ Although this article addresses LTO tapes specifically, it also applies to 3592 cartridges.



During a write process, *recovered write errors* occur when the read-back check of the data shows that the record had too many errors. The error correction circuitry of the tape drive may have been able to correct these errors and the data may have been read back perfectly, but a drive may decide to allow only a few errors in order to have sufficient margins for a later restore process. The drive will then rewrite the same record further down the tape. This is a fairly common operation caused mainly by media defects or by debris and is of little concern as long as the number of recovered write errors is not excessive.

Recovered read errors occur when the error correction logic is not able to correct errors during a read process. The drive will retry the operation, which involves stopping the tape, moving it back and reading the record again. The tape drive will repeat this operation until the read operation is successful, or until it finally gives up. A recovered read error means that this retry operation was eventually successful. However, the process of repeatedly stopping the tape, reversing, and again forwarding can take a long time. Meanwhile your system sits idle, waiting for the data to finally arrive. In addition, these start/stop operations have much higher impact on the life expectancy of the tape than do normal tape movements. Recovered read errors are therefore much more serious than recovered write errors, so it is advisable to incur very few read-retries.

The second type of error is the *unrecovered read and write error*. These are by far the most undesirable errors since they always cause a complete crash of the backup or restore process. If this happens during a data restore and if you do not have another copy of the data, you will soon be getting *those* phone calls from very unhappy users. Unrecovered read errors are especially unpleasant.

The third type of error statistics consists of *servo errors*. LTO tape is manufactured with five servo bands pre-written by the tape manufacturer². Tape drives use these bands to locate the data tracks. A *Suspended Writes* error means that during a write process, the information in a servo band could not correctly be decoded. The servo band instructs the tape drive as to the location of the write head. So if a servo error occurs, the drive may not know the location of the write head. The head could even be in an area that has previously been written. Continuing to write in this situation could have disastrous results. Tape drives justly tend to be a bit paranoid about servo errors while writing. A servo error does not necessarily mean that the servo track was unreadable. It just means that the drive found too many errors in the servo track to continue writing. The drive will write the

² The cartridge memory contains the name of the manufacturer who wrote the servo track as well as the name of the cartridge manufacturer. These two names are not always identical because several LTO media manufacturers sell their cartridges under multiple brand names. It is here that the actual manufacturer of the tape media can be found.



data further down the tape when the quality of the servo tracks improves. Those areas of tape where the servo bands were questionable are completely bypassed. The usable length of such a tape becomes slightly shorter, but this is of little concern as long as there are not too many suspended writes.

However, if the drive is unable to read the servo band at any location, the write operation is suspended. This is known as a *Fatal Suspended Writes* error since the drive is unable to recover from this failure. If the word 'fatal' is frightening, it is meant to be so. Because LTO tape drives cannot regenerate servo bands, once the servo information is destroyed the tape is completely unusable. A cartridge with this type of error is one that you do not want in your data center under any circumstance.

2. Number of bytes read and written

The *number of bytes read and written* refers to the total amount of data that has been transferred to and from the cartridge³. This tells us how many times the tape was moved through the drive's tape path. A well-maintained cartridge is quite resilient to the challenges of the tape path elements throughout its lifetime. At some point however, a cartridge is past its life expectancy and rapidly moves into its twilight years.

3. Number of cartridge loads

If the number of times the cartridge has been loaded is disproportionately high relative to the number of bytes written or read, this indicates that each time the cartridge was loaded, only a few records were accessed. Accessing random records on tape requires movement of the tape to locate the records. If this is the case, the actual amount of tape that moved through the drive's tape path is probably much higher than the statistic of *number of bytes transferred* suggests. It should ideally require only one load to write at least one wrap (i.e., one movement of tape between beginning and end of tape in either direction⁴). The number of cartridge loads should not exceed 50 times the number of bytes transferred divided by the total cartridge capacity.

4. Age of cartridge

The CM contains the date of manufacture of the cartridge which allows us to know its age. An old cartridge is not necessarily a bad cartridge. A cartridge should last for a long time without any negative effects on its

³ To be precise, the memory contains the number of "data sets" written or read. The number of bytes can be easily calculated from these numbers.

⁴ This is due to the serpentine nature of tape writing.



performance if it has been stored in a controlled environment. However, the older the cartridge is, the more natural it is that more opportunities would arise for its abuse. Given a choice, a newer cartridge is preferable to an old cartridge, so statistically, old age is considered a weak negative.

Pointing fingers

When an error occurs or when a defective tape cartridge is found, there are many directions in which to point a finger of blame. For instance, the cartridge memory chip is written and read by the tape drive. If a tape drive is malfunctioning, the cartridge could appear to be defective because the drive may falsely increment the error counters. Another possibility is that the error condition may only be temporary when, for example, debris is deposited on the head. Movement of tape across the head or other conditions can also remove that debris and the drive may function correctly thereafter. In this case the damage is limited. Also, a permanently defective drive could simply increment the error counters of the cartridge to the extent that a perfectly good cartridge could merely appear to be damaged. "In general, cartridge failure is very rare," says Rich D'Ambrise, Director of Technology at Maxell Corporation of America. "Most of the errors occur from lack of drive maintenance, drive failure or software issues - even when most of the prompts point towards media or cartridge errors."

Fortunately there are some indicators that can help to pinpoint the reasons for error statistics on the cartridge memory chip. In addition to the overall number of errors made during the life of a tape, the cartridge memory also contains details of each of the last four sessions. If a cartridge has a high error count but performed flawlessly the last four times the cartridge was used, we tend to blame the errors on the tape drive. This is because tape cartridges generally cannot repair themselves (bad sections are simply skipped over). However, if most of these errors occurred *during* the last four sessions, the cartridge is more likely to be faulty. The details of the last four sessions also contain the serial numbers of the tape drives which were used, so if the errors occurred on several different drives (with, of course, the same tape) this is a very strong indicator that the problem is with the cartridge.

How to extract the information


There are at least two ways to extract the CM information. First, it is available through the host interface of LTO tape drives. Special software on the host can receive the data from the drive and analyze the quality of the cartridge.



Second, the CM can also be read externally without having to load the cartridge into a drive. Through a wireless RF interface, an external device can access the CM. The photo on the left shows a device called VeriTape™ manufactured by MPTapes, Inc. It connects to a system through a USB interface. Application software that runs on most Windows operating systems reads and displays the data. Because there is no need to load the

cartridge into a drive, the operation is almost instantaneous and allows many cartridges to be evaluated in a very short period of time.

Putting the pieces together



Looking at all the data to evaluate a cartridge can be time consuming, tedious, and inconsistent. Not everyone assigns the same weight to each parameter. We therefore found it quite helpful to create a single score number that is automatically and consistently calculated from all available data. We call this score number the VeriScore™, with values ranging from 0 (very bad) to 100 (very good). We display VeriScore™ as well as the other CM statistics discussed in this article. It is fairly safe to use cartridges with a VeriScore™ of 90 and higher. Cartridges with a VeriScore™ between 70 and 89 maybe a bit risky, but if you need to stretch your media budget and you don't mind an occasional job crash, you may want to risk using them. Those with scores between 50 and 70 are recommended for only the adventurous user. In general, we discourage the use of tape cartridges with a VeriScore™ lower than 50.



LTO CM Reader - Files read from D:\Mountain_Engineering\LtoReader\CM-data\server-backups

File Action Help

Summary

Manufacturer	Min	Max	Avg	SD	FUJIFILM	VERBATIM	FUJIFILM	IMATION	IMATION	FUJIFILM	IMATION	IMATION
Serial Number					03DC111385	F3XPD5253	03DK101443	0816660460	080360884P	0348003256	0821617420	0821606930
Tape Score	60.83	97.75	86.84	12.81	61	61	68	74	86	86	89	92
Manufacture Date	2.75	4.35	2.91	0.40	2003.12.22	2003.10.27	2003.12.22	2003.09.25	2002.05.26	2003.05.09	2003.12.29	2003.12.29
Cartridge Type					LTO 2 cartridge	LTO 2 cartridge	LTO 2 cartridge	LTO 2 cartridge	LTO 1 cartridge	LTO 2 cartridge	LTO 2 cartridge	LTO 2 cartridge
Tape Alerts					none	none	none	none	none	none	none	none
Last Unload					Normal	Normal	Normal	Normal	Normal	Normal	Normal	Normal
Thread Count	1	27	6.70588	7.39734	27	13	16	9	3	16	7	1
Total Data Written	55.4 GB	1.42 TB	342 GB	401 GB	775 gigabytes	1.06 terabytes	1.42 terabytes	425 gigabytes	55.4 gigabytes	671 gigabytes	152 gigabytes	145 gigabytes
Total Data Read	0 B	32.4 GB	2.98 GB	7.88 GB	7.42 gigabytes	959 megabytes	965 megabytes	345 megabytes	1.21 megabytes	251 megabytes	32.4 gigabytes	404 kilobytes
Recent Data Written	0 B	487 GB	140 GB	126 GB	111 gigabytes	487 gigabytes	350 gigabytes	206 gigabytes	55.4 gigabytes	203 gigabytes	0 bytes	145 gigabytes
Recent Data Read	0 B	32.4 GB	2.1 GB	7.81 GB	815 megabytes	345 megabytes	256 megabytes	339 megabytes	1.21 megabytes	49.7 megabytes	32.4 gigabytes	404 kilobytes
Recent Write Retries	0	1.467e-03	1.302e-04	3.848e-04	0	1 (8.286e-07)	1 (1.153e-06)	11 (2.158e-05)	0	2 (3.978e-06)	0	526 (1.467e-03)
Recent Read Retries	0	1.366e-03	1.050e-04	3.447e-04	0	0	1 (1.366e-03)	0	0	0	39 (4.864e-04)	0
Recent Servo Errors	0	7.290e-06	4.775e-07	1.767e-06	0	1 (8.286e-07)	0	0	1 (7.290e-06)	0	0	0
Recent Fatal Write Errors	0	0	0	0	0	0	0	0	0	0	0	0
Recent Fatal Read Errors	0	0	0	0	0	0	0	0	0	0	0	0
Recent Fatal Servo Errors	0	0	0	0	0	0	0	0	0	0	0	0
Total Write Retries	0	1.467e-03	1.719e-04	3.951e-04	3 (1.564e-06)	2 (7.623e-07)	71 (2.023e-05)	14 (1.332e-05)	0	20 (1.204e-05)	216 (5.756e-04)	526 (1.467e-03)
Total Read Retries	0	1.684e-03	1.843e-04	4.306e-04	10 (5.446e-04)	4 (1.684e-03)	1 (4.188e-04)	0	0	0	39 (4.864e-04)	0
Total Servo Errors	0	1.236e-05	1.179e-06	3.378e-06	0	1 (3.812e-07)	0	13 (1.236e-05)	1 (7.290e-06)	0	0	0
Total Fatal Write Errors	0	0	0	0	0	0	0	0	0	0	0	0
Total Fatal Read Errors	0	0	0	0	0	0	0	0	0	0	0	0
Total Fatal Servo Errors	0	5.446e-05	3.204e-06	1.321e-05	1 (5.446e-05)	0	0	0	0	0	0	0
Total Fatal Servo Errors	0	0	0	0	0	0	0	0	0	0	0	0

CM data read from D:\Mountain_Engineering\LtoReader\CM-data\server\stNoReader

VeriTape permits the user to quickly and consistently check if a cartridge is within his standards. "The VeriTape LTO cartridge reader is a 'must have' for all LTO owners" says Ken Cruden, Exec. V.P. and GM for Tandberg Data Corporation. "It's extremely easy to know if you have a good cartridge before archiving with a simple display of error rate performance, usage, age and a color coded scale to expedite the decision about usability."



Calculation of VeriScore™

How is VeriScore™ calculated? There are four steps to consider. First, a curve is defined for each parameter, e.g., the number of bytes transferred to and from the cartridge. Each data transfer means movement of tape through the drive's tape path. As long as the cartridge is still relatively new, this does not impact the quality of the cartridge, but once a cartridge reaches its end of life the quality rapidly decreases. So we define the curve accordingly to reflect this non-linear behavior.

The second step in calculating VeriScore™ is to use the number of data transfers to find the area in the previously defined curve where it impacts the score. We repeat this step for each parameter.

The third step is to assign a weight to each parameter. For example, unrecovered read errors that occurred on multiple drives within the last four sessions are weighted heavily and therefore have a significant impact on the score. The age of a cartridge is weighted very lightly. The weight of the other parameters is somewhere in between these two extremes.

The final step is to add the weighted results for each parameter, and at that point we assign it a VeriScore™ from between 0 and 100.

Conclusion

If you don't know the quality of your cartridge, you may be gambling with your company's data. As Kevin Burton says, "I am a big believer in knowing the condition of your media, and any tool that can give you that insight can make a data center run more efficiently."

"All data centers that use LTO cartridges should have the VeriTape LTO cartridge reader," adds Ken Cruden. "This simple to use product is extremely helpful for identifying the quality of a cartridge and can be beneficial in diagnosing the status of a cartridge or of the drive that used the cartridge."

The cartridge knows its past history. Don't you think you should know it too?

Peter Groel
President
MP Tapes, Inc.
www.mptapes.com