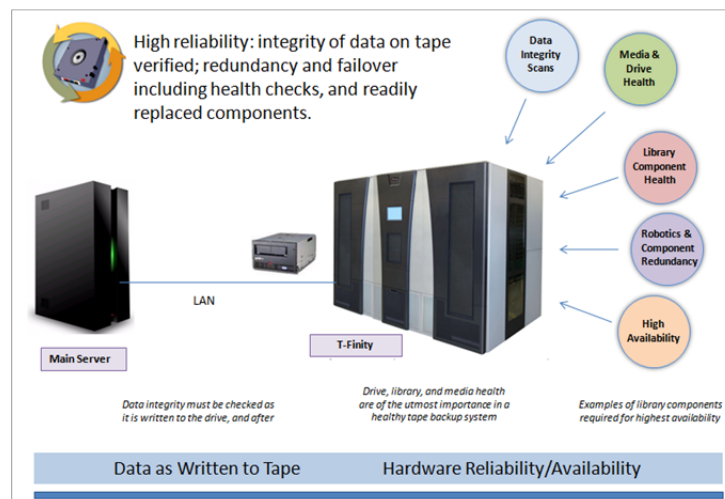


The New Reliability Requirements: Tape Libraries



Spectra Logic
6285 Lookout Road
Boulder, Colorado 80301

May 2011
(303) 449-6444
fax: (303) 939-8844

Contents

Abstract 3

What is Reliability?..... 3

Step by Step 3

 Reliability Features Assumed to be Available 4

Feature Comparison..... 4

Data As Written to Tape 5

 Data Integrity Check—no Intervention Required 6

Redundancy and Health Checks..... 6

 Drive, Media, and Library Component Health Checks 7

Availability..... 7

Conclusion..... 8

BlueScale, TranScale, Spectra, and the Spectra Logic are registered trademarks of Spectra Logic Corporation. All rights reserved worldwide. All other trademarks and registered trademarks are the property of their respective owners. All library features and specifications listed in this white paper are subject to change at any time without notice. Copyright © 2011 by Spectra Logic Corporation. All rights reserved.

Abstract

Documents of historical significance, like the Dead Sea Scrolls, the Rosetta stone, and the U.S. Constitution, only survive thanks to reliable media—stone and parchment. Digital efforts have yet to prove themselves over such a span of time, but right now, the closest equivalent storage media for digital data is magnetic tape, with a life of thirty or more years. In fact, data has been restored from tapes more than fifty years old¹. Tape libraries, such as Spectra tape libraries, have added features that enhance and strengthen tape reliability and support it with high-availability features. These features include verification that data written to tape can be read, redundant tape library components that can take over if one fails, and on-going library, tape drive, and tape media health checks.

What is Reliability?

Very often in the tape storage market, the word “reliability” appears, yet it is difficult to find a comprehensive definition of reliability. In some cases, it means redundant hardware: in others, read-after-write technology. While reliability is an important factor in tape library technology, if a company regards reliability across a limited spectrum, its reliability features will reflect the relative narrowness of this viewpoint. Spectra Logic assesses reliability from the moment data leaves the backup framework to 30 years after its creation, and so offers the widest range of reliability features available in both middle and enterprise markets. By breaking down the reliability process step by step, administrators can examine features that are important to all aspects of high reliability.

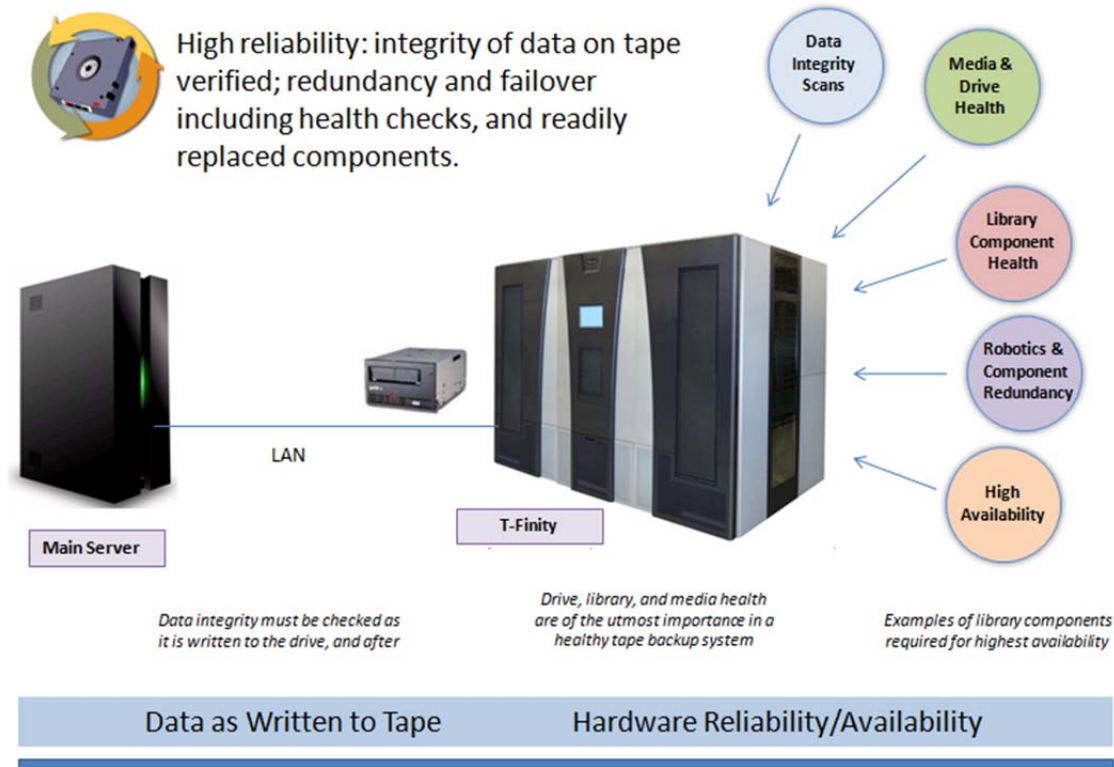
Step by Step

While the process can be deconstructed any number of ways, one method reviews reliability in three ways:

- **Data integrity support:** checking drive and media health, and checking data integrity after it is written to tape.
- **Hardware reliability:** the library depends on the working of each part so that the library may reliably perform data saving and retrieval functions over the life of the library. An unreliable or faulty hardware setup may prevent data from being written to tape.
- **Availability:** making sure that any library issues are resolved very rapidly and with minimal interruption, if any.

The diagram below illustrates the movement of data over time, and the features that Spectra libraries offer to address the intertwined phases of data, tape, and library health.

¹ “Fifty Years of Storage Innovation: Magnetic Tape and Beyond.” IBM Archives, IBM.com. Accessed May 12, 2011. http://www-03.ibm.com/ibm/history/exhibits/storage/storage_fifty.html



Reliability Features Assumed to be Available

All tape libraries offer some reliability features, especially for the enterprise market. You can expect dual robotics, dual power, dual port drives, call home features, and drive cleaning from every enterprise class library. Some even have redundant electronic controllers. These features are understood to be standard support/reliability/availability options.

Feature Comparison

Placing the features that Spectra offers in context will help compare the extent and breadth of Spectra's reliability features. The following data is taken from a wide range of vendor-supplied materials, and is accurate to best of the researcher's ability at the time of publication.

Feature/Option	T-Finity	Competing Enterprise Libraries
dual-robotics option	yes	yes
dual/redundant power supplies	yes	yes
automatic drive cleaning	yes	yes
dual-port drives	yes	yes

Feature/Option	T-Finity	Competing Enterprise Libraries
path failover	yes	yes
phone home auto-support	yes	yes
Separate redundant modules for interface, robotics, library controller	yes	one other library
global spare drive	yes	no
library configuration automatic backup	yes	no
automatic data verification of tapes that are also available for backup/restore	yes	no (some systems offer verification but it's not automatic)
User-installable components i/o blades, robotics, power supply, drive	yes	no

Data As Written to Tape

Drives used in all enterprise tape libraries support an extremely low error bit rate, with anywhere from 10^{-17} to 10^{-19} , depending on the drive type, with SATA drive error rate at 10^{-15} . Storage analyst Curtis Preston writes, “While 10^{-15} may look really close to 10^{-17} , it's not. When it's bits we're talking about, it's the difference between 113 TB and 11.1 PB! It means you are 100 times more likely to have bad data on disk than you are on an LTO-5 tape drive, and 10,000 times more likely than if the data is stored on a T1000C or TS1130 drive!²”

Spectra libraries have added features that enable the library to check data integrity once data is written to tape, and to check the health of the tape. These features are unique in that they can run automatically, as scheduled, without requiring a separate partition or separate attention. The Spectra library features ensure data integrity include Drive and Media Lifecycle Management, and the unique feature set (Data Integrity Verification) that includes of Pre-, Post- and QuickScan at no charge. Other systems may offer features like this, but the features cost extra and are not integrated into standard library operations.

² Preston, Curtis. “Have we put Tape out to Pasture too soon?” *Backup Central Blog*, April 21, 2011. www.backupcentral.com

Data Integrity Check—No Intervention Required

Data Integrity Verification builds on Media Lifecycle Management (MLM) by scanning some or all of the data that has been written to tape. It seamlessly integrates into existing environments and operates as a background library-managed process with no impact on existing backup and archive schedules. Spectra's Data Integrity Verification work proactively performs three levels of fully automated validation:



- **PreScan** - Checks newly imported tapes to ensure they are ready to use: for example, in good health, generationally compatible and not write-protected.
- **QuickScan** - Quickly and immediately verifies the backup and archive data was successfully written to tape with a rapid single pass from the beginning of the tape to the end of the first track. Along with PostScan, rapid or gradual media degradation is tracked in the library and reported to users.
- **PostScan** - Automatically verifies the physical tape cartridge and the integrity of all of the data stored upon it for the duration the data is retained. As with PreScan and QuickScan, the PostScan process is performed by the library and is independent of the backup application normally used to read and write data to the tape.


David Hill, principal of the analyst firm Mesabi Group, notes that it is significant that "...Data Integrity Verification's features [are] at no cost. Most users will run both Spectra's PreScan and QuickScan to validate that the tape is in good condition and that data has been properly written to and can be read back from tape. QuickScan's single pass data verification check takes less than one minute to complete, an impressive time frame."³

Redundancy and Health Checks

Highest end libraries include redundant robotics, so that even if one set of robotics fails, the other can take over while the failed set is being repaired. Additional component redundancy similarly permits data write/restore tasks to continue even if a component runs into trouble.



Spectra libraries are the only to provide a failover option for drives, so that if a drive fails, a global spare can be used in its place even if no one is on-site to replace the failed drive. To use the global spare drive option, install a spare drive in the library and configure the spare so that it is associated with some or all of the library partitions. If a drive in a designated partition fails, you can immediately switch to the global spare, so operations continue unimpeded.



³ <http://www.spectralogic.com/blog/index.cfm/2011/4/12/Spectra-Logic-Defines-Data-Integrity-Verification-Looks-Forward-to-Tape-Summit>

The Spectra T-Finity library has additional redundancy built throughout the library, including spare power brushes and fans, and multiple redundant electronics modules that can carry on in spite of failures of paired electronics and interface modules.

Drive, Media, and Library Component Health Checks

Drive Lifecycle Management provides data about drives that you can use to identify drives with issues before those drives affect on-going library operations. DLM gathers data each time a tape is unloaded from a drive. This data, stored by the library, includes the following information about the last fifty tapes unloaded:

- Read/write errors—including soft and hard errors
- Tape alerts
- Any flags generated while the tape was loaded in the drive
- Recent errors or recent information posted through the drive’s single character display

With this data, the library generates a report on both overall drive health status and health reports for individual drives. These reports, along with BlueScale drive clean/test/resets, help you determine where a problem rests: with a drive, a tape, or some other library component.

Using Spectra Certified Media, BlueScale Media Lifecycle Management (MLM) reports help you identify tapes that required retries to complete an operation, tapes with high error rates, and tapes with other problems. Further, you can track absolute values of tape usage from initial use/ tape’s manufacture date, to identify tapes that are nearing end-of-life as specified by the manufacturer. You can remove these tapes before issues arise—to make sure that the data written to the tape can be retrieved.

BlueScale Library Lifecycle Management (LLM) tracks the library and its components across its lifecycle. Through Hardware Health Monitoring (HHM), maintenance thresholds for key library components are tracked notifying you when maintenance tasks are required to keep your library in optimum operating condition. This utility generates a report that displays, and that you can save to a USB device, showing the current status of the HLM counters. It includes information about library serial number, key components and maintenance threshold, and more.

Availability

The most reliable libraries will provide options for the user to minimize downtime. You can expect most library vendors to provide some “hot-swappable” components, but only Spectra libraries offer on-site customer-replaceable components, along with automatic backup features for inventory and configuration data, reducing the time it takes to bring up a library after an issue may arise, along with an extensive package of parts replaceable in minutes on site.

The Assisted Self-Maintenance program addresses availability, because a library needs to be online in minutes, not hours or days. The T950’s Assisted Self-Maintenance (ASM) option stores components such as the T950 power supply on-site, so that staff can replace any failed component in

A circular icon with a light blue background and a darker blue border. The text "Media & Drive Health" is centered in a dark blue font.

Media &
Drive
Health

A circular icon with a light blue background and a darker blue border. The text "Library Component Health" is centered in a dark blue font.

Library
Component
Health

A circular icon with a light orange background and a darker orange border. The text "High Availability" is centered in a dark orange font.

High
Availability

minutes. Components including the drive and the power supply do not require any tools—the components are unlatched, removed, replaced, and latched in under a minute. This minimizes downtime – and the parts that can be stored on-site include drive, I/O controller, power supply, filters, compact flash, transporter/robotics, library and robotics controllers



Our Assisted Self Maintenance feature enables you to replace select parts yourself.

Most components can be replaced without tools, and in under three minutes. The T950 and some of the smaller libraries (the T200, T380, and T680, for example) provide an option to store a second set of library robotics on-site; this replacement procedure does require a screwdriver that is provided with the ASM kit.

With the parts on-site option, site administrators can readily replace components when issues arise and as directed by Spectra technical support. Instead of waiting four or eight hours for an engineer to arrive at the site, administrative staff can quickly install the component, and get the library issues resolved very quickly.

Conclusion

Spectra libraries extend the notion of reliability beyond an expected feature set. Users can expect the highest end libraries to provide at least an option of dual robotics, dual power, call-home features, and some path redundancy. Spectra libraries have new features that are very important to reliable library /tape automation: a method of checking library health, drive health, media health, the integrity of data stored on tape, and methods to increase availability through parts on-site.