

Quantum®

WHITE PAPER

CLOUD-READY QUANTUM CERTIFIED LTO MEDIA

Data Integrity With an Unthinkable Eleven 9s and More

CONTENTS

Introduction	3
Overview	3
Capacity.....	4
Data Durability	5
Number of 9s.....	6
Green Tape.....	6
Environmental Archiving Test	6
Quantum's LTO Media Database.....	6
Summary	7
References	7

INTRODUCTION

To maintain Linear Tape-Open (LTO) Ultrium trademark compliance, LTO media suppliers are required to pass vigorous compliance testing every year. This required process ensures the media adheres to the LTO Ultrium format specification, including interchange testing that validates that all media and LTO drives are interoperable. This well-defined and strict compliance testing is managed by the LTO consortium with the media supplier's full collaboration. See reference 1 for more details.



In-between yearly compliance testing, it's the media supplier's responsibility to ensure production cartridges continue to meet Ultrium specifications. The Quantum LTO certification is a value-added process focusing on five key user metrics: capacity, transfer rate, servo, and green media characteristics, with specific emphasis on testing the durability of the written data. This process is accomplished using the Quantum Certification System, an automated library test system with built-in calibrated test drives, reference media, and test cartridges based on a proprietary stratified sampling statistical process control. Quantum maintains and manages a live database with cartridge and drive test data logs, a key to support in-house drive and media studies by Quantum's media research team.

OVERVIEW

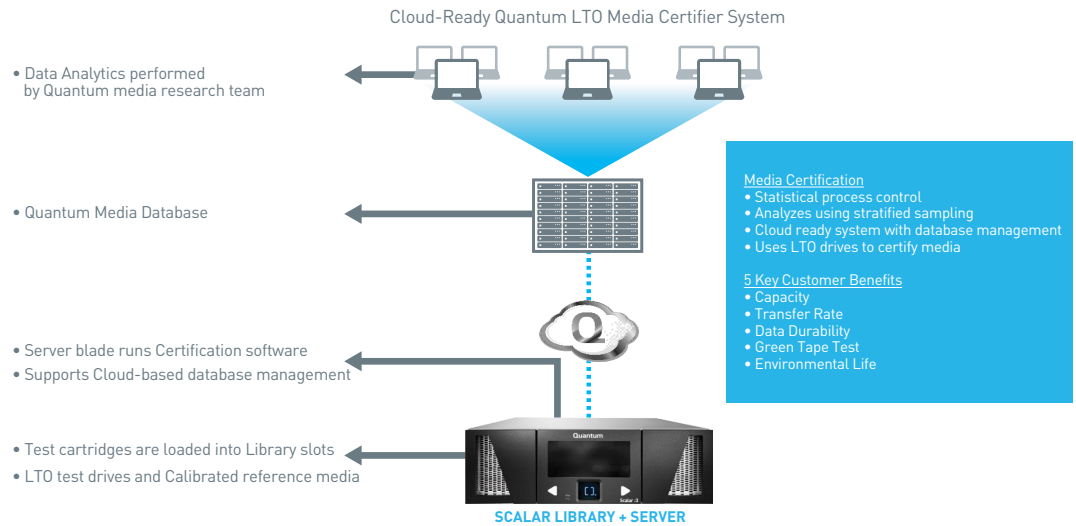
Quantum's certified media process tests production cartridge samples with calibrated drives and software. The architecture is based on statistical process control. The system is designed using Quantum Scalar® libraries with a dedicated local server (see Figure 1). Its purpose is to simulate customer usage conditions such that the five key metrics are checked and verified (see Table 1).

Table 1. *Quantum's Certified Media Key Metrics*

Metrics	Process
Capacity	Statistical process control for capacity during production
Transfer rate	Monitor for magnetic and physical media damage during production
Data durability	Monitor user data durability during production
Green tape test	Analyze data captured from certification for potential green tape signature
Environmental life	Accelerated aging testing using Quantum certified media to estimate user data durability as a function of media life and environment

During certification, sampled cartridges are tested using a special tape drive read-and-write process such that the reported performance data can be analyzed. The read-and-write process includes controlled underruns to ensure that testing includes back-hitch motions rather than benign full streaming modes. Special calibrated reference cartridges are used to support the analysis.

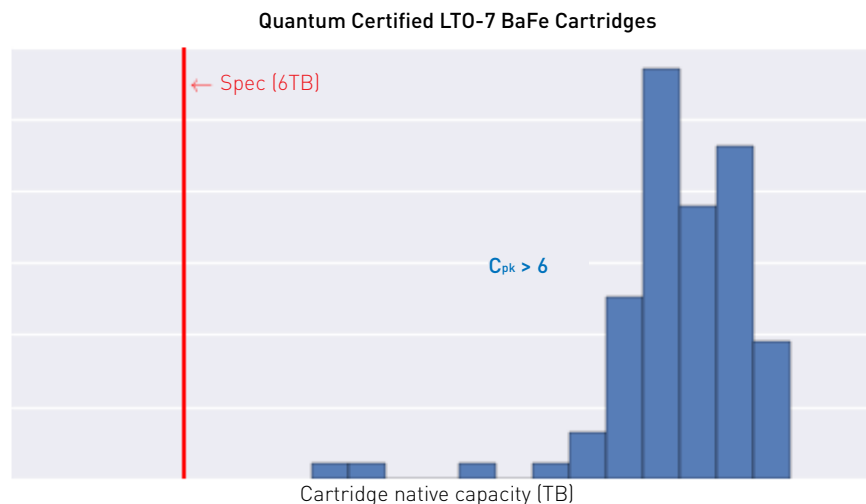
Figure 1. Quantum LTO Media Certifier



CAPACITY

Key process control variables are: minimum capacity, data transfer rates, tracking performance, and hard errors as attribute events. Figure 2 shows a sample plot for cartridge capacity distribution based on the amount of data sets that can be written to tape under test. Quantum monitors the distribution of the maximum capacity for the sample set and establishes a C_{pk} metric based on the distribution shape. The goal is to make sure the shape of the histogram is skewed to the right with a left tail resulting in a $C_{pk} > 6$. The data from each tested sample is stored in a database for further detailed analysis to help with key metrics, such as green media and drive cartridge interactions. This process ensures that when certified cartridges are used in constant capacity mode to generate copies for disaster recovery, especially for Tape Object Storage and archiving applications, the user application problems are minimized as the very large C_{pk} numbers indicate.

Figure 2. Quantum Certified Media Write Performance



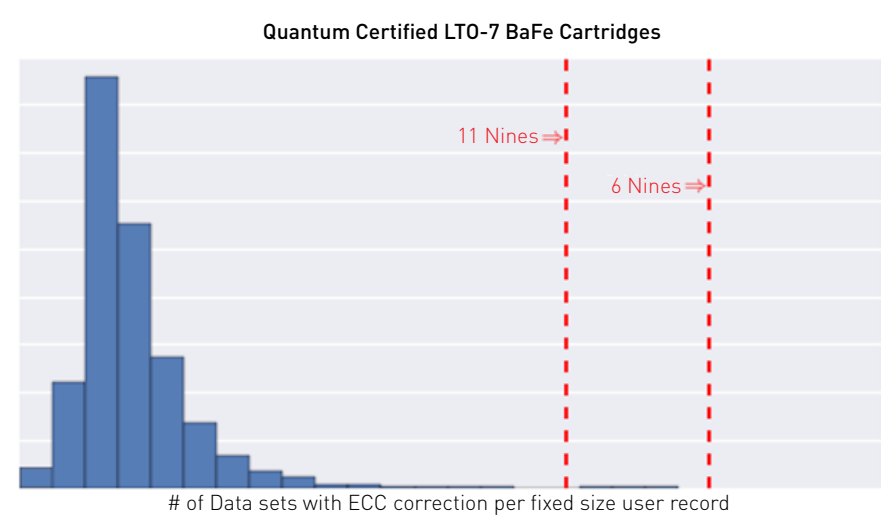
DATA DURABILITY

Durability is the uncorrectable bit errors that cannot be recovered from the underlining media even with the existing error correction code (ECC) system, such that one of those bytes after writing massive amounts of data might not be correct when it's read. The LTO Ultrium 7 specification for bit error rate (BER) is an amazing number with one error event in 10^{19} user bits; this means more than 200K cartridges or 10 exabytes (EB) of user data can be reliable.

With two orthogonal layers and highly interleaved multi-channel ECC code, the LTO BER specification is the outer Reed Solomon (RS) code error rate in erasure mode. This is the error rate of the C2 code where multiple channels of data are fed after two-dimensional interleaving. It's not feasible to measure such an error rate. For example, in order to measure a 10^{17} BER, the drive would need to process at least 1840×10^{17} bits which would be equal to reading nearly 4 million LTO-7 tapes. Therefore, the research community is inclined to devise approximate and reasonable analytical expressions, such as the binomial model, to estimate the actual BER. There are also data-based advanced estimation studies for accurately predicting the real life C2 failure numbers. The binomial model has been shown to be a very useful and accurate model under ideal conditions as long as errors are random and not correlated.

Quantum's Certification System utilizes a series of repeated writes and reads using data sets to estimate the durability of the archived data. By controlling actual locations on the tape, this process is also able to identify localized defects, potential suspect or damaged media sections down and across the tape. Figure 3 is a sample archival durability histogram plot for 3PB user data based on Quantum certified media. The x axis is the read quality of the data sets, number of data sets with C2 ECC error correction. The y axis is the number of occurrences sensed by the drive per recoding zone. The data durability limits are shown for reference only; these are based on model, simulation, and empirical test data estimates.

Figure 3. Quantum Certified Media Read Performance



Quantum certification monitors for the data sets requiring ECC correction to ensure that no hard errors, no read recovery process were needed and the sections of media that have large counts of data sets with C2 correction are mostly due to non-correlated random errors. See references 2-4 and 6-9 for more details.

NUMBER OF 9s

Traditionally durability (9s) is calculated from Mean Time To Data Loss (MTTDL) in hours, which is estimated by using the Markov Model and key metrics such as BER, annual failure rates (AFRs), transfer rates, repair rates, and capacity, assuming errors are random and uncorrelated. Using Quantum's Markov-Model-based Nines estimation, single LTO-7 cartridge durability would be approximately six 9s. With Quantum's certified media, properly managing environmental conditions and employing periodic data integrity checks, users can achieve greater than twelve 9s durability with a simple cost-effective two-copy replication-based archiving solution. See references 5 and 10-12 for more details.

GREEN TAPE

The term "green tape" refers to an unused data cartridge. Typically, green tape is more abrasive on the drive heads than cartridges that have been used several times. Drive performance can be affected if large numbers of green tapes are used for archiving (e.g., used once) versus using the same tapes repeatedly for backup operations. Quantum works with media vendors to establish an initial green tape performance prior to the start of manufacturing as defined in the Quantum media specification. Since this is a one-time initial test, the Quantum Certification System is used to check, monitor, and estimate the green tape performance as an ongoing real-time test process based on a special test algorithm. This algorithm utilizes calibrated reference tapes, multiple drives, and test sample cartridges. Periodically, the system estimates green tape performance and correlates to the initial test metrics to forecast if a potential green tape issue might be observed. The data and analysis results are archived in Quantum's certification database which is used for detailed background engineering analysis.

ENVIRONMENTAL ARCHIVING TEST

Archiving the life of LTO media is a well-studied topic focusing on media chemistry and error rates internal to the drive. However, no previously conducted research has actually focused on the durability of user data as a function of the environmental conditions in which the media is archived. The Quantum Certified System addresses user data durability in terms of the number of 9s as a function of the archiving conditions, such as years stored, temperature, and humidity. The Quantum LTO research team manages a detailed and extensive LTO archiving study using sampled cartridges from the Certification Test System by utilizing an industry standard guide for accelerated aging testing using the Arrhenius equation and data durability estimations. Sampled media with prewritten and calibrated user data are stored in extreme conditions over long periods to determine the archiving life of user data per Quantum library operating specifications; for example, the number of 9s durability for written user data when stored at 35°C for a minimum of 10 years or 25°C for 30 years. This is the first time any company has established this type of controlled testing and data collection process as part of a certification practice.

QUANTUM'S LTO MEDIA DATABASE

In addition to the certification process, the Quantum LTO media research team maintains and manages an industry-first comprehensive Quantum LTO media database, which is a collection of test data from every cartridge tested. This database contains key metrics, such as cartridge and pancake ID, drive internal temperature, serial numbers, history of media, environmental

conditions, tape alerts, underrun and overrun events, various error rates, capacity data, servo information, and transfer rates. This database also includes special test data from calibrated reference tapes which are tested periodically by Quantum's certification test system as part of the overall certification process. This database is instrumental in the ongoing research performed by Quantum's LTO media research team to explore various media performance metrics, such as green tape, head cleaning, and floating debris.

SUMMARY

As tape is taking an ever-increasing role in long-term cost-effective data storage, it is imperative that the drives and media are up to the task. As there is no visible or physical way to guarantee the quality and reliability of media, it falls to the LTO developers and suppliers to ensure this level of quality to give customers confidence their data is safe. Quantum tests LTO media to ensure they meet the required specifications and standards at multiple points during the manufacturing process. Quantum believes it is critical customers know that when they purchase a Quantum branded LTO media product it has been tested to a higher standard. It is this additional quality assurance process that enables Quantum to offer the most comprehensive quality and warranty in the industry. Customers using Quantum tape automation and branded LTO media can be confident they are storing the data on the best available products.

REFERENCES

- [1] Ultrium LTO website, <http://www.lto.org>.
- [2] Ultrium LTO tape reliability web page, http://www.lto.org/wp-content/uploads/2014/07/LTO_Get_The_Facts_10_15.pdf.
- [3] S. S. Arslan, Jaewook Lee, Jerry Hodges, James Peng, Hoa Le, and Turguy Goker, "MDS Product Code Performance Estimations Under Header CRC Check Failures and Missing Syncs," *IEEE Transactions on Device and Materials Reliability* Vol. 14, No. 3, pp. 921–930, Sept. 2014.
- [4] S.S. Arslan, "Post-ECC Error Analysis and Erasure Rate Estimation for the Current Tape Drive System," unpublished, 2012. Email: Suayb.Arslan@quantum.com.
- [5] S. S. Arslan, "Durability and Availability of Erasure-Coded Systems with Concurrent Maintenance," unpublished, 2014. Available online: http://www.suaybarslan.com/Reliability_Systems_14.pdf.
- [6] S. S. Arslan, "Redundancy and Aging of Efficient MDS-Parity Protected Distributed Storage Systems," *IEEE Transactions on Device and Materials Reliability*, Vol. 14, No. 1, pp. 275–285, Mar. 2014.
- [7] S. S. Arslan, J. Lee, and T. Goker, "Error Event Corrections Using List-NPMLD Decoding and Error Detection Codes," *IEEE Transactions on Magnetics*, Vol. 49, No. 7, pp. 3,775–3,778, July 2013.
- [8] S. S. Arslan, J. Lee, and T. Goker, "Bit Error Detection and Correction with Error Detection Code and List-NPMLD," US 20140173381 A1.
- [9] S. S. Arslan, J. Lee, and T. Goker, "Cycle Slip Detection and Correction Through Classification of Run Length Limited Code Failures," *IEEE Transactions on Magnetics*, Vol. 49, No. 9, pp. 4,988–4,998, Sept. 2013.
- [10] Shawn Brume, "The 50 Year Total Cost of Ownership: Real considerations for long term storage," <https://www.ibm.com/blogs/data-protection-retention/the-50-year-total-cost-of-ownership-real-considerations-for-long-term-storage>.
- [11] Amazon Simple Storage Service (S3) FAQs, AWS Data Protection [eleven 9s durability metric], <https://aws.amazon.com/s3/faqs>.
- [12] Dan Iacono and Carol Sliwa, "Durability for Cloud Applications," [audio podcast], <http://searchcloudstorage.techtarget.com/podcast/Data-durability-guarantees-shouldnt-be-big-concern-with-cloud-storage>.



ABOUT QUANTUM

Quantum is a leading expert in scale-out storage, archive, and data protection, providing solutions for capturing, sharing, and preserving digital assets over the entire data lifecycle. From small businesses to major enterprises, more than 100,000 customers have trusted Quantum to address their most demanding data workflow challenges. Quantum's end-to-end, tiered storage foundation enables customers to maximize the value of their data by making it accessible whenever and wherever needed, retaining it indefinitely and reducing total cost and complexity. See how at www.quantum.com/customerstories.

www.quantum.com • 800-677-6268

Quantum[®]