

ORACLE®

STORAGETEK

An Oracle White Paper
September 2013

StorageTek T10000 Tape Drives: Enterprise-Class Design for Maximum Reliability

ORACLE®

Executive Overview	2
Introduction	2
Background	3
High-Reliability Features	4
Dual Heads.....	4
32-Channel Operation	5
Recording Head Durability.....	5
Improved Tape Guiding	6
Rugged Cartridges and Media	9
Improved Data Integrity	10
Intelligent Data Transfer	11
Conclusion	11

Executive Overview

When valuable data is stored on tape, users expect to be there when it's needed. The success or failure of storing and recalling data depends on the reliability of both the tape drive that writes the data and the tape on which it's stored.

Reliability is affected by the design of the data storage system, usage, and environmental factors. Many low-cost tape drives can store data, but can they safely store the data in your high-usage, always-on business environment? And what are the potential consequences of using a tape drive that is not designed to operate 24/7 like many businesses?

Oracle's StorageTek T10000C and StorageTek T10000D tape drives build on proven technology to deliver unprecedented durability, reliability, and, ultimately, peace of mind. This paper outlines several of the features and technologies that make these StorageTek T10000 tape drives the best choice for demanding data environments.

Introduction

A company's data is its most valuable asset. Without access to data, no company can survive. This reliability requirement means enterprises cannot compromise on the technology that is used to maintain and secure that data.

Tape storage is the most trusted technology for data backup in long-term storage environments.

Many companies use tape storage for their data assets. While more than 80 percent of the world's digital data is stored on some kind of tape, it is enterprise tape that provides incomparable levels of data integrity, reliability, and robustness.

Enterprise tape is designed and constructed to perform in high-duty cycle robotic environments, to recover without failure, and to integrate into a heterogeneous environment (FICON for MVS or Fibre Channel for open systems).

This paper describes the design elements of StorageTek T10000 tape drives, which are high-duty cycle enterprise-class tape drives that provide superior reliability. Oracle designed these drives for heavy data retrieval in a robotic environment with excellent data availability, reliability, and performance. In many cases, customers will be reading and writing with the drive in excess of 20 hours per day, 365 days a year. Applications that manage and secure data typically require frequent starting, stopping, and seeking by the drive. The design technology in these enterprise-class drives is based on customers' requirement that data, once written to a StorageTek tape cartridge, must be available. Reliability of the drive and media is paramount.

Background

Before delving into StorageTek T10000 advanced enterprise features, a brief review of tape drive components may be helpful.

There are two parts to a tape drive system: the drive and the cartridge. The StorageTek T10000 T2 tape cartridge contains a reel wound with approximately two-thirds of a mile (~1000 meters) of tape. The tape is about half an inch wide and approximately 5 microns thick. These are among the components in the drive configuration shown in Figure 1:

- A loader to receive the cartridge and engage the reel onto a motor
- A tape path to guide the tape from the cartridge reel to the take-up reel inside the drive
- A head assembly to read and write data onto the tape
- Electronics cards to process data signals and control the drive's movements and functions

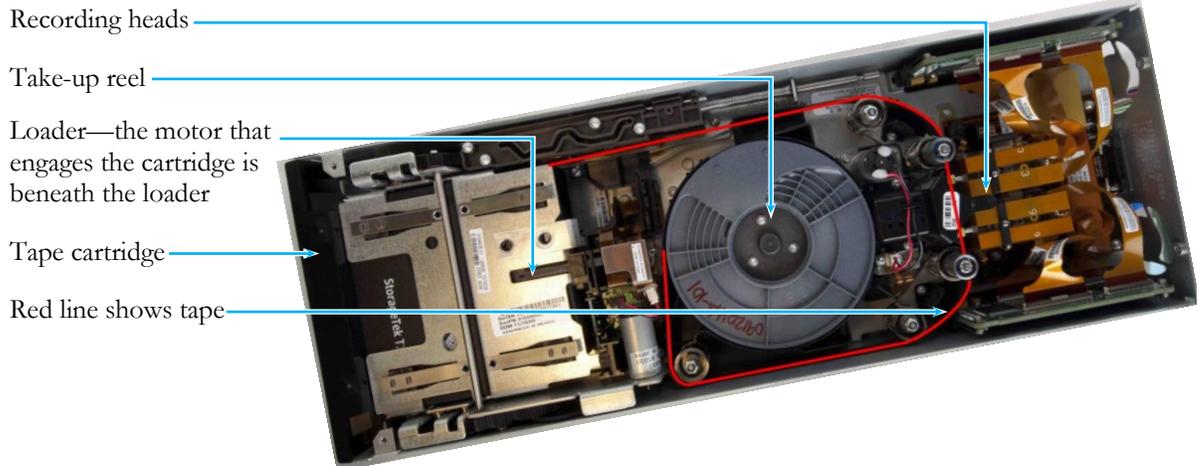


Figure 1. Tape drive components. Some elements, such as the electronics cards beneath the mechanical assembly, cannot be seen.

StorageTek T10000C and StorageTek T10000D tape drives write data in tracks down the entire two-thirds of a mile length of the tape. Each millimeter of track length contains more than 1,000 bytes of data. When multiplied by the number of tracks of data on the tape, the StorageTek T10000 T2 tape cartridge yields a capacity of up to 5.5 TB of user data using StorageTek T10000C tape drive, or a capacity of up to 8.5 TB using StorageTek T10000D tape drive.

Both the StorageTek T10000C and T10000D tape drives use two heads, each writing and reading 32 tracks of data simultaneously. To keep the heads on the correct track while reading and writing, the tape drive uses special prewritten tracks on the tape called servo bands. These bands are read by servo readers on each head, allowing the head to position to the desired tracks and stay on track during lateral tape motion (LTM) as it moves through the tape path.

The tape path's function is to guide the tape over the heads as the tape moves from one reel to the next, without damaging the tape or the data on the tape, while minimizing LTM.

High-Reliability Features

To add to the general picture of a physical drive and tape path, here are some of the engineering considerations employed in StorageTek T10000 tape drives to ensure high reliability in enterprise-class applications.

Dual Heads

Where the rubber meets the road, as it were, is the point at which the read/write head meets the tape. StorageTek T10000 tape drives use two independently servo-controlled heads to lay down 32 tracks of data quickly and reliably. In Figure 2, while the tape moves left to right, the head on the left writes 32 tracks simultaneously while the head on the right read-verifies those same 32 tracks. The opposite occurs when tape moves from right to left.

As the tape storage industry moves towards ever-shrinking track widths, single-head tape drives are forced to compensate for an inability to control the positioning of write elements and read elements independently relative to the tape as it speeds past. Single-head tape drives can have difficulty verifying written data during write operations, as the servo position of both the writers and readers becomes simultaneously critical. Due to slight variations in perpendicularity between the head and tape, the ideal position for the writers to put the data on tape may not be the ideal position for the readers to verify the integrity of what was written. The problem has driven some single-head tape drive manufacturers to incorporate a head azimuth control mechanism, which attempts to measure and adjust head tilt, effectively chasing the correct setting.

By writing and reading with independently servo-controlled heads, StorageTek T10000 tape drives forego the chase and eliminate the head azimuth problem altogether.

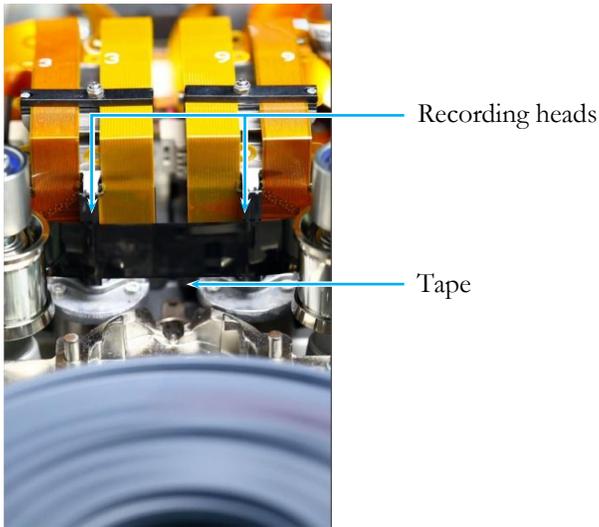


Figure 2. Dual heads provide 32 parallel read/write channels.

32-Channel Operation

In high-duty cycle enterprise applications, there are three goals that help to maximize reliability.

1. Reduce tape speed as much as possible while still achieving the required data rate.
2. Reduce the number of tape passes required to achieve the required cartridge capacity.
3. Distribute the data across as many channels as possible to maximize the error correction code's ability to correct simultaneous errors.

Thirty-two channels help improve data and drive reliability by addressing all three of the goals listed above.

Recording Head Durability

In high-duty cycle archive environments, the recording head must be designed to continuously record data on never-used (aka green) media, typically having higher abrasivity (that is, surface roughness) than used, burnished media. StorageTek T10000 tape drive recording heads are designed to handle this most demanding of tape storage environments. An advanced head contour design minimizes contact forces of the media over the magnetic recording elements. In addition to an advanced contour design, the StorageTek T10000D tape drive recording heads feature a three-layer coating of advanced, tribological-grade¹ materials that minimize wear due to higher media abrasivity:

¹ Tribology is the study of friction, wear, and interacting surfaces in relative motion.

1. The top layer consists of native-oxide and serves as an extremely durable and hard wearing protective barrier for the recording elements.
2. An unoxidized mid-layer adapts to changes that occur at the head/tape interface by shifting to its resilient native-oxide form, thereby refreshing the native-oxide layer as it wears.
3. An under-layer provides further protection by insulating the magnetic elements and providing a cohesive bond between the recording head and the upper layers.

Tape Speed

Lower tape speed minimizes media and drive wear. For example, using 32 channels simultaneously, the StorageTek T10000D tape drive achieves an industry-leading native data rate of 252 MB/sec while operating at a tape speed of only 4.75 m/sec. In comparison, LTO tape drives only support 16-channel operation, achieving lower data rates despite moving the tape through the tape path at higher speeds. For example, LTO-6 provides a data rate of 160 MB/sec while operating at a tape speed of 6.8 m/sec. With LTO-6, tape data speeds through the tape path 43 percent faster than with the StorageTek T10000D tape drive, and results in a 37 percent lower native data rate.

Tape Passes

Fewer end-to-end passes of tape results in less media and drive wear. With 32 channels, the StorageTek T10000C and T10000D tape drives are able to write a 5.5 TB (StorageTek T10000C tape drive) or 8.5 TB (StorageTek T10000D tape drive) tape cartridge with half as many passes as would be required with 16-channel operation. For example, the StorageTek T10000D tape drive writes 4,608 tracks on the media. The 32 parallel read/write elements require 144 passes to fill the tape, or 17 passes per terabyte. In comparison, LTO-6 utilizes a 16-channel head to write 2,176 tracks on the media, requiring the drive to make 54 passes per terabyte. Half as many channels mean lower drive and data reliability as well as throughput.

Data Distribution

The more channels used to put the data on tape, the higher the number of simultaneous errors may be corrected by the error correction code (ECC). For example, if the StorageTek T10000D tape drive detects that data read by 5 out of 32 channels is simultaneously in error, the ECC is able to correct the five errors using the “good” data from the other 27 channels. LTO drives distribute data across only 16 channels and are therefore unable to correct as many simultaneous errors. There is more on error correction later in this paper.

Improved Tape Guiding

The SafeGuide System Feature Tape Path

Several of the following features comprise the unique tape guidance system designed to ensure superior media reliability in StorageTek T10000 tape drives. These features are combined to make up the SafeGuide System, a feature of the StorageTek T10000 tape drives (Figure 3).

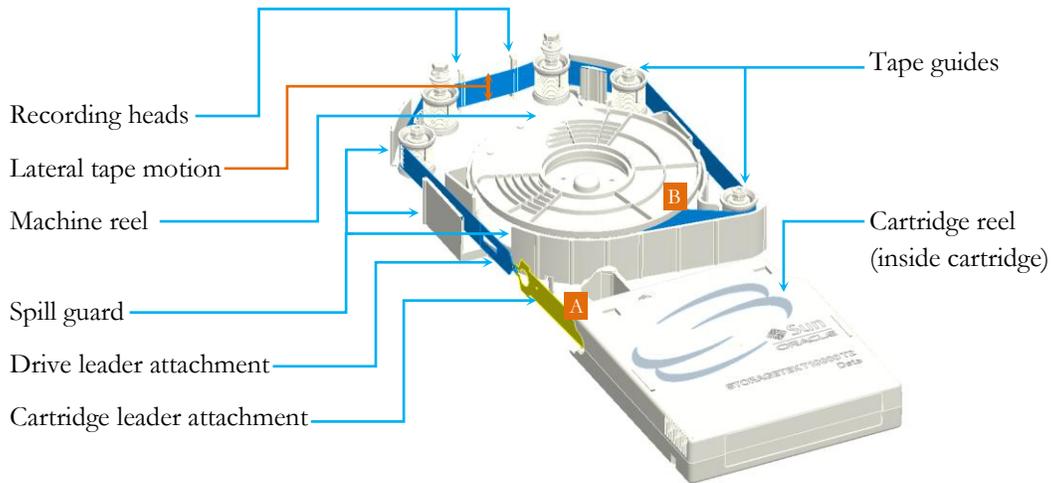


Figure 3. A complete tape path system; the length of tape from position A to

Improved Tape Guides

The elements of the tape path that control the lateral position of the tape are called the tape guides. StorageTek T10000 tape drives include five flanged rollers. The four closest to the head are specially designed to use surface guiding on the backside of the media in conjunction with traditional edge guiding.

The flanges at the top and bottom of the roller physically prevent large lateral tape movements. But that's only part of the StorageTek T10000 tape drive guidance strategy. Helical grooves on these rollers disperse the air film that might otherwise form between tape and roller, allowing rollers to gently grip the tape's backside surface to further reduce LTM (Figure 4). The helical design, its grooves spiraling around the roller like stripes on a barber pole, improves the reliability of the tape surface by preventing a wear pattern from forming on the media. At the same time, the grooves reduce tape edge wear by minimizing tape contact with the flanges.



Figure 4. One of the four grooved tape guides used in the StorageTek T10000C tape drive.

The key to Oracle's implementation of this grooved system is that it uses the backside of the tape rather than the data surface. Other implementations guide on the data surface where valuable data is stored. In high-duty cycle use, this can eventually damage the tape and degrade the data.

Length of Tape Path

For the highest media reliability, the tape path should be as long as possible. The longer the tape path, the less guiding force is required to keep the tape positioned properly relative to the head as it moves through the tape path. The StorageTek T10000 tape drive path is more than two times the length of competitors' tape paths. This length allows StorageTek T10000 tape drives to guide the tape using less force, greatly reducing edge or surface wear and thereby reducing the chance of damage to the tape. If the edge gets damaged, users may not be able to read back the data. One of the reasons enterprise drives are physically larger than low-end drives is to accommodate the longer tape path.

Leader and Take-Up Mechanism

StorageTek T10000 tape drives use an improved two-leader design and buckling mechanism (Figure 5) that Oracle believes provides the most reliable means available for connecting and pulling tape from a reel. This design is less sensitive to cartridge leader failure due to rough handling and provides a reliable positive engagement between the tape leader and cartridge leader.

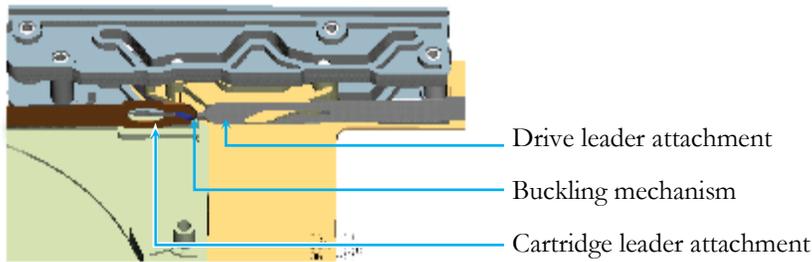


Figure 5. Two-leader design and buckling mechanism.

Other cartridge technologies offer exposed leader pins, which can dislodge from the leader during handling, causing cartridge load failure. The StorageTek T10000 tape drive design ensures reliable take-up while lessening the chances of media damage.

Rugged Cartridges and Media

Locking Hubs

A dropped cartridge can result in tape edge damage if a hub flange is pushed into the cartridge case. Edge damage can occur if the outer portion of the flange hits the inside of the cartridge case. To further protect data, the StorageTek T10000 T2 tape cartridge contains a locking mechanism that prevents the flanges from hitting the inside of the cartridge case during handling (Figure 6).

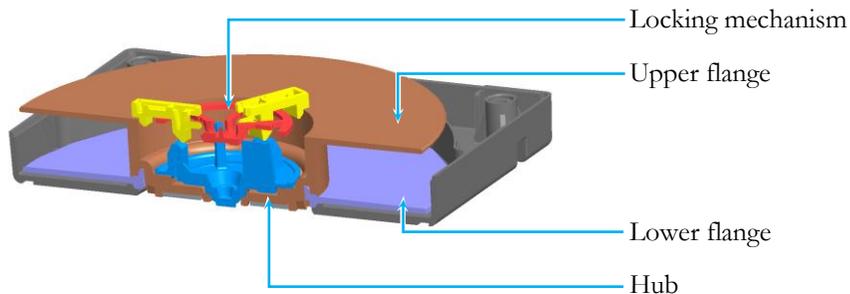


Figure 6. A locking mechanism prevents tape damage in the event of a fall.

Media Stability

When data is saved to tape, users want to be confident that data will be accessible now, as well as decades from now. Magnetic tape storage has one of the longest archive lifetimes, up to 30 years, of all storage solutions currently on the market. As with most materials, magnetic tape can change its dimensions and/or magnetic properties as a function of environmental changes. Minimizing any tape dimensional changes and magnetic degradation will ensure the robustness of data, even after long-term storage in different environments.

Oracle's selection of aramid as the StorageTek T10000 T2 tape cartridge substrate assures superior tape dimensional stability performance and long-term archival life. Further details about the importance of tape dimensional stability on the long term accessibility of your archive are available in the white paper, "[Protecting Archival Data with Improved Tape Dimensional Stability](#)."

Oracle's selection of barium ferrite (BaFe) as the StorageTek T10000 T2 tape cartridge magnetic particle assures superior chemical stability and storage performance compared to other metal particles in extreme environmental conditions. The details behind this comparison are available in the Fujifilm white paper, "[Long Term Archivability and Stability of Fujifilm Magnetic Tape Using Barium-Ferrite \(BaFe\) Particle.](#)"

Improved Data Integrity

Dual Error Correction Code

StorageTek T10000 tape drives improve data integrity by combining a sophisticated dual error correction code (ECC) and a cyclic redundancy check (CRC) to provide the highest level of error detection and correction available. The dual ECC, in conjunction with distributing data across 32 channels, creates a significant improvement over other ECC systems (Table 3).

TABLE 3. A COMPARISON OF UNCORRECTABLE BIT ERROR RATES

DRIVE	BIT ERROR RATE
StorageTek T10000D tape drive	1 per 10 ¹⁹
StorageTek T10000C tape drive	1 per 10 ¹⁹
IBM TS1140	Unavailable
LTO-6	1 per 10 ¹⁷
Seagate Terascale HDD	1 per 10 ¹⁴

The overlapping CRCs in the data path throughout the drive give the best protection from rare data corruption events due to buffer corruption or hardware failure. In fact, the uncorrectable bit error rate of StorageTek T10000 tape drives is 100 times better than LTO-6 and 100,000 times better than Seagate's latest capacity disk offering.

Data Integrity Validation Feature

CRCs ensure records received from the host do not get corrupted while moving between internal memories of the tape drive. And dual ECC ensures data recovery once it is on the tape medium. However, there is still a possibility for data corruption as it migrates across various pieces of equipment between the host and the tape drive (switches, backplanes, buses, adapters, and various memory buffers).

Oracle knows how important data is and has developed the Data Integrity Validation feature of the StorageTek T10000 tape drives to protect enterprise data from end to end. This important feature ensures the tape drive checks end-to-end data integrity when a record is written, and

allows the application to check end-to-end data integrity when the record is read at a later time. Additional details about Oracle's end-to-end protection of data are available in the white paper, "[Data Integrity Validation Feature of StorageTek T10000 Tape Drives](#)."

Intelligent Data Transfer

StorageTek T10000 tape drives from Oracle are designed to deliver the fastest performance in the industry for any tape storage environment. Whether a user has slower server technology or an application that writes discrete files rather than long data streams, T10000 tape drives have been engineered to store data faster than any other tape device. With fewer cartridges needed and faster performance the tape drive is able to complete jobs in less run time and with less drive wear and tear. Detailed explanations of the performance features available in the StorageTek T10000 tape drives are available in the white paper, "[Maximizing Tape Performance with StorageTek T10000 Tape Drives](#)".

Conclusion

There is really only one legitimate criterion for an enterprise tape system, and that is reliability. If the data isn't there when it's needed, there is no reason to spend the money on storage. Oracle's StorageTek T10000 tape drives include several features that support the intense, high-duty cycle applications of enterprise data centers. From thoughtful hardware design to sophisticated error detection and correction, the proven reliability of the StorageTek T10000C and StorageTek T10000D tape drives provide more than secure data. It provides peace of mind.



StorageTek T10000 Tape Drives:
Enterprise-Class Design for Maximum Reliability
September 2013

Oracle Corporation
World Headquarters
500 Oracle Parkway
Redwood Shores, CA 94065
U.S.A.

Worldwide Inquiries:
Phone: +1.650.506.7000
Fax: +1.650.506.7200
oracle.com



Oracle is committed to developing practices and products that help protect the environment

Copyright © 2013, Oracle and/or its affiliates. All rights reserved. This document is provided for information purposes only and the contents hereof are subject to change without notice. This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document and no contractual obligations are formed either directly or indirectly by this document. This document may not be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without our prior written permission.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. UNIX is a registered trademark licensed through X/Open Company, Ltd. 0913