An Oracle White Paper
June 2011

StorageTek T10000C Tape Drive Enterprise-Class Design for Maximum Reliability

ORACLE

Executive Overview

When you store your valuable data on tape, you expect it to be there when you need it. 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 24x7 like your business?

Oracle's StorageTek T10000C tape drive builds on proven technology to bring you unprecedented durability, reliability, and, ultimately, peace of mind. This paper outlines several of the features and technologies that make the StorageTek T10000C tape drive the best choice for demanding data environments.

Introduction

Your company's data is your most valuable asset. Without access to data, no company can survive. This reliability requirement means you cannot compromise on the technology that you use 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% 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 the StorageTek T10000C tape drive, a high-duty-cycle enterprise-class tape drive with superior reliability. Oracle designed this drive 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 this enterprise-class drive is based on the customer requirement that data, once written to a StorageTek cartridge, must be available. Reliability of the drive and media is paramount.

Background

Before delving into the StorageTek T10000C tape drive's advanced enterprise features, a brief review of tape drive components might be helpful.

There are two parts to a tape drive system: the drive and the cartridge. The StorageTek T10000 T2 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

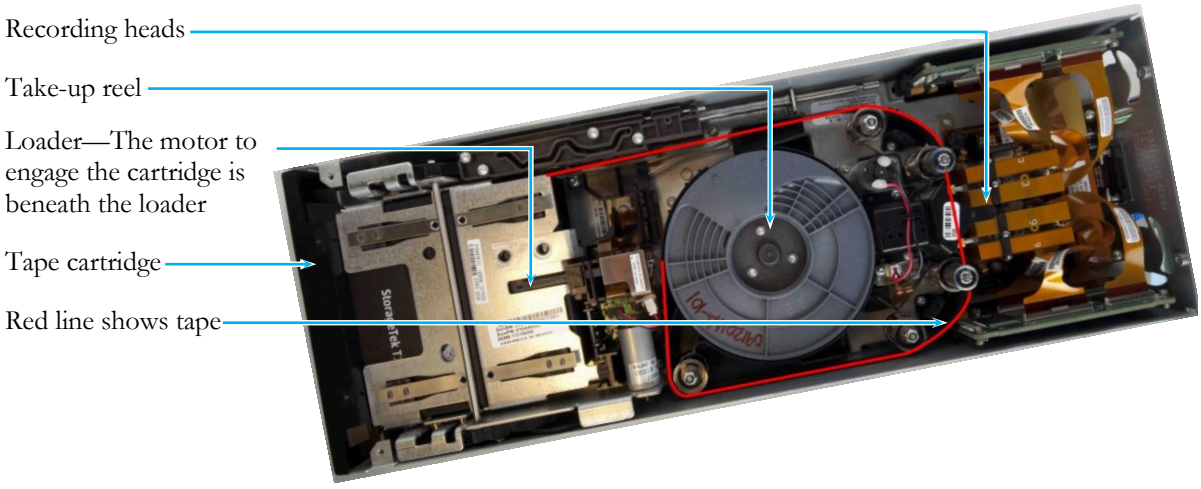

Figure 1. Tape drive components. Some elements cannot be seen, such as the electronics cards beneath the mechanical assembly.

The StorageTek T10000C tape drive writes data in tracks approximately 3 microns wide down the entire two-thirds-mile length of the tape. Each millimeter of track length contains about 1,273 bytes of data. When multiplied by the 3,584 tracks of data on the tape, each cartridge yields a capacity of 5 terabytes of user data.

As you will learn in more detail shortly, the StorageTek T10000C tape drive uses two heads, each writing and reading 32 tracks of data simultaneously. To keep the heads on the correct track while reading and writing, the StorageTek T10000C tape drive uses nine special prewritten tracks on the tape called *servo bands*. These bands are read by servo readers on each head, allowing the head to be positioned 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

Keeping in mind a general understanding of the physical drive and tape path, here are some of the engineering considerations employed in the StorageTek T10000C tape drive 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. The StorageTek T10000C tape drive uses two independently servo-controlled heads to lay down 32 tracks of data quickly and reliably. As shown 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 might 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, the StorageTek T10000C tape drive foregoes the chase and eliminates the head azimuth problem altogether.

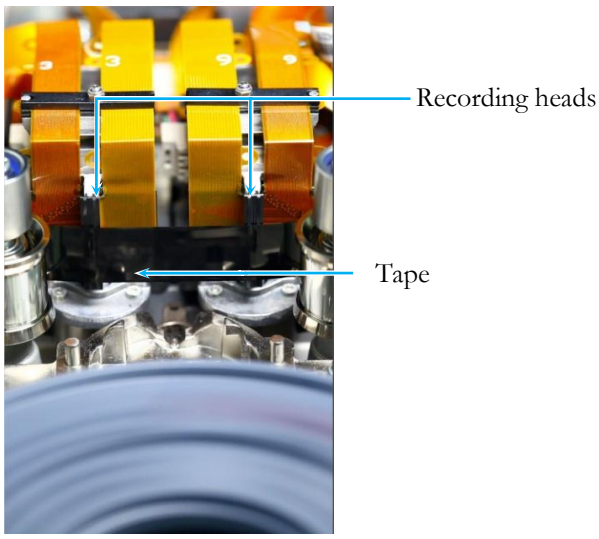

Figure 2. Dual heads provide 32 parallel read/write channels.

32-Channel Operation

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

- Reduce tape speed as much as possible while still achieving the required data rate.
- Reduce the number of tape passes required to achieve the required cartridge capacity.
- 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.

Tape Speed

Lower tape speed minimizes media and drive wear. Using 32 channels simultaneously, the StorageTek T10000C tape drive achieves a data rate of 240 megabytes per second (MB/sec) native (uncompressed) with a tape speed of only 5.62 meters per second (m/s). By comparison, mid-range drives support only 16-channel operation, achieving lower data rates with higher tape speeds. For example, Oracle's StorageTek LTO5 tape drive provides a data rate of 140 MB/sec while operating at a tape speed of 6.1 m/s.

Tape Passes

Fewer end-to-end passes of tape results in less media and drive wear. The StorageTek T10000C tape drive's 32 channels enable a 5-TB cartridge to be written with half as many passes as would be required with 16-channel operation. For example, the StorageTek T10000C tape drive writes 3,584 tracks on the media. The 32 parallel read/write elements require 112 passes to fill the tape, or 22 passes per terabyte. By comparison, the StorageTek LTO5 tape drive utilizes a 16-channel head, requiring the drive to make 53 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 T10000C 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. Mid-range drives distribute data across only 16 channels and are, therefore, unable to correct as many simultaneous errors. There will be more on error correction later in this paper.

Improved Tape Guiding

The SafeGuide Tape Path

Several of the following features comprise the StorageTek T10000C tape drive's unique tape guidance system designed to ensure superior media reliability. These features combine to make the SafeGuide Tape Path. See Figure 3.

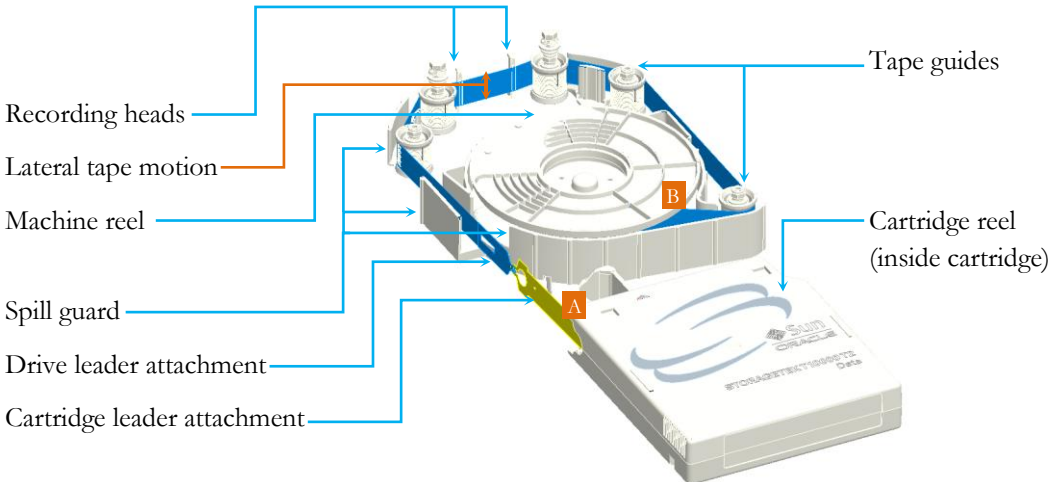

Figure 3. A complete tape path system is the length of tape from position A to position B.

Improved Tape Guides

The elements of the tape path that control the lateral position of the tape are called the *tape guides*. The StorageTek T10000C tape drive includes 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 T10000C 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 (see 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 your 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 T10000C tape path is more than two times the length of competitors' tape paths. This length allows the StorageTek T10000C tape drive to guide the tape using 18 times less force, greatly reducing edge or surface wear and thereby reducing the chance of damage to the tape. If the edge gets damaged, you might not be able to read back your 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

The StorageTek T10000C tape drive uses 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.

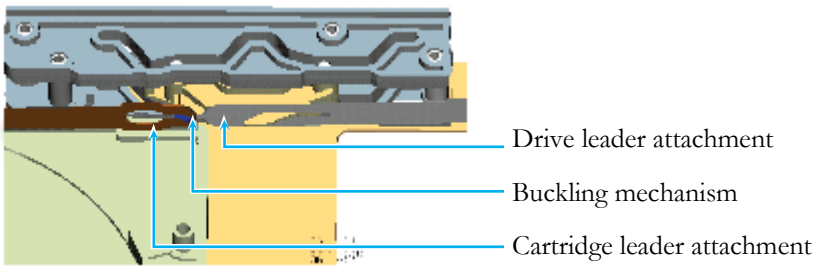

Figure 5. Two-leader design and buckling mechanism.

Other cartridge technologies offer exposed leader pins, which can dislodge from the leader during handling, causing cartridge load failure. The StorageTek T10000C 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 your data, the StorageTek T10000 T2 cartridge contains a locking mechanism that prevents the flanges from hitting the inside of the cartridge case during handling (see Figure 6).

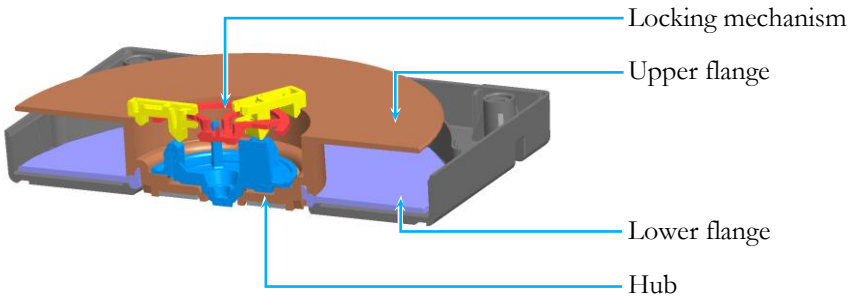

Figure 6. A locking mechanism prevents tape damage in the event of a fall.

Media Stability

When data is saved to tape, you want to be confident that the 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, the dimensions and magnetic properties of magnetic tape can change as a function of environmental changes. Minimizing any tape dimensional changes and magnetic degradation ensures the robustness of data, even after long-term storage in different environments.

Oracle's selection of aramid as the StorageTek T10000 T2 substrate ensures superior tape dimensional stability performance and long-term archival life. For further detail on the importance of tape dimensional stability on the long-term accessibility of your archive, please refer to the white paper Protecting Your Archival Data with Improved Tape Dimensional Stability.

Oracle's selection of barium ferrite (BaFe) as the StorageTek T10000 T2 magnetic particle ensures superior chemical stability and storage performance compared to other metal particles in extreme environmental conditions. For the details behind this comparison, please refer to 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

The StorageTek T10000C tape drive improves data integrity by combining a sophisticated dual error correction code (ECC) and a cyclic redundancy check (CRC) to provide one of the highest levels 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 1).

TABLE 1. A COMPARISON OF UNCORRECTABLE BIT ERROR RATES

DRIVE	BIT ERROR RATE
Oracle StorageTek T10000C	1×10^{-19}
IBM TS1140	Unavailable
Oracle StorageTek LTO5	1×10^{-17}
Hitachi Ultrastar 7K3000 Disk Drive	1×10^{-15}

The overlapping CRCs in the data path throughout the drive provide the best protection from rare data corruption events due to buffer corruption or hardware failure. In fact, the uncorrectable bit error rate of the StorageTek T10000C tape drive is two orders of magnitude better than that of the StorageTek LTO5 tape drive and four orders of magnitude better than Hitachi's latest enterprise disk offering.

StorageTek Data Integrity Validation

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 data is on the tape medium. However, there is still a possibility for data corruption as data migrates across various pieces of equipment between the host and the tape drive (for example, switches, backplanes, buses, adapters and various memory buffers).

Oracle knows how important your data is and has developed StorageTek Data Integrity Validation (DIV) for the StorageTek T10000C tape drive 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 it allows the application to check end-to-end data integrity when the record is read at a later time. For further detail on Oracle's end-to-end protection of your data, please refer to the white paper StorageTek Data Integrity Validation for the StorageTek T10000C Tape Drive.

Intelligent Data Transfer

StorageTek File Sync Accelerator

StorageTek File Sync Accelerator significantly improves throughput performance when the drive detects a high frequency of file syncs coming from the host. A high frequency of file syncs between files can be a problem, because each requires the drive to empty the contents of the buffer to tape, provide confirmation to the host that the data was written successfully to tape, and then back-hitch or reposition the tape in preparation for receiving the next file in the stream. Not only is there a drop in data throughput in this situation, there is unnecessary wear on the media due to back hitches.

StorageTek File Sync Accelerator enables the drive to continue streaming even when it receives a high frequency of file syncs from the host. Instead of stopping after each file sync is written to tape, the tape drive writes all the incoming data in segments on one track of the tape, providing confirmation to the host after each file is successfully written to tape. Once the host transfer is complete or the data buffer is filled to the appropriate level, the data is written contiguously to another track on the tape. The 2-GB buffer in the StorageTek T10000C tape drive, twice as large as any other drive's buffer, enables twice the amount of data to be received from the host before data is rewritten to the tape. While the primary purpose of the 2-GB buffer is to ensure no loss in performance due to back-hitch operations, it also enables StorageTek File Sync Accelerator to match host performance more effectively while reducing media wear due to repositioning.

StorageTek Tape Application Accelerator

A number of Multiple Virtual Storage (MVS) applications do not support the new 3590 streaming-enable commands or lack the optimizations required to fully realize the benefits of the commands (for example, DFS-HSM). The purpose of these optimizations is to minimize undesirable tape drive performance hits by buffering write-tape-marks and sync-marks between files that are much smaller than the relatively large data buffers within the drive. Applications that are not optimized or that do not support streaming-enable commands prevent you from getting the highest possible performance out of your tape drive investment and expose your drives and media to unnecessary wear.

The StorageTek Tape Application Accelerator allows the StorageTek T10000C tape drive to stream to its full potential even when it is being controlled by MVS applications not tuned to maximize streaming. By reducing back-hitches in environments that use these types of applications, the StorageTek Tape Application Accelerator can dramatically improve performance while minimizing wear on the drive and the media.

Note: Applications that support streaming-enable commands typically recommend *duplexing* output type jobs to two tape drives to minimize exposure to buffered data loss due to a catastrophic event, such as an individual drive power failure. Likewise, Oracle recommends duplexing output-type jobs when the StorageTek Tape Application Accelerator is enabled. By default, this feature is disabled and may be enabled only on FICON-configured drives using the Virtual Operator Panel (VOP).

Progressive Write Technology

You can't read what hasn't been written. With that in mind, the StorageTek T10000C tape drive improves data read reliability by utilizing progressive write technology. This technology allows the drive to skip down the media if it encounters difficulty in writing, remapping the data farther down the tape rather than trying to rewrite again in the same location. This technique has two distinct advantages:

- It reduces back hitches.
- It writes the data in a better location on the media, giving more read-back reliability.

Rigorous Testing

How much is 20,000 terabytes of data? That's the number of bytes that were recorded on StorageTek T10000 T2 cartridges without a single data-loss event before we declared the StorageTek T10000C tape drive worthy of your demanding datacenter environment. Spooled out end to end, this amount of written media would span San Francisco to Boston and be guided through StorageTek T10000C hardware a total of 112 times, the equivalent of 11 trips around the globe.

Conclusion

There is really only one legitimate criterion for an enterprise tape system and that is reliability. If your data isn't there when you need it, there is no reason to spend your money. The StorageTek T10000C tape drive includes several features that support the intense, high-duty-cycle applications of your enterprise datacenter. From thoughtful hardware design to sophisticated error detection and correction, the proven reliability of the StorageTek T10000C tape drive from Oracle provides more than secure data. It provides peace of mind.

StorageTek T10000C Tape Drive
Enterprise-Class Design for Maximum Reliability
June 2011

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

Hardware and Software, Engineered to Work Together