



An Oracle White Paper
May 2011

Lowering Storage Costs with the World's First Five Terabyte Tape Drive

Executive Overview	1
Introduction	1
Unmatched Capacity	3
Lowering Tape Acquisition Costs.....	4
Reducing Management Costs	5
Reducing Datacenter Floorspace Costs	5
Minimizing Operational Risk by Reducing Cartridge Movement.....	5
Lowering Power and Cooling Costs	5
World's First Exabyte Tape Library	6
Conclusion	6

Executive Overview

Explosive data growth is causing storage costs to skyrocket, which is forcing IT organizations to look for more cost-effective archival and backup solutions. Oracle's StorageTek tape product family has maintained a market leadership position for many years and now includes the world's highest capacity tape drive. A single Oracle tape library unit can now house more than an exabyte of data¹ and can offer cost savings of up to 20% compared to other tape vendors².

Introduction

To cope with today's explosive data growth, many IT organizations are using a tiered storage approach that balances the cost of different types of storage media against application performance requirements. Tape solutions still offer the most cost-effective means to maintain long-term copies of infrequently used data and to backup large data sources. Since applications are generating and utilizing greater amounts of data, it is becoming challenging to backup hard disk drives to tape in a reasonable timeframe. Tape drive performance is therefore becoming a critical factor in completing backups within an allotted time window.

Similarly, as data sources keep growing, more and more data is being migrated or archived from hard disk to a near-line storage repository such as a tape library. While disk-to-disk backup solutions are useful for some situations, their capacity is much more limited — generally measured in hundreds of terabytes (TB), whereas tape libraries can store several orders of magnitude more data. Yet even tape library capacity can be stretched by today's storage and archival requirements. Today's IT managers need tape solutions that can keep

¹ Assumes 2:1 compression ratio.

² Based on list prices for IBM TS3500 and LTO-5 tape drives as of December 2010.

massive amounts of archived data available to applications while requiring minimal datacenter floor space.

Managing the vast amounts of data stored on tape can also be a big challenge for large enterprises. This is especially true when hard disk storage is nearly fully utilized and older data must be frequently archived to make room for new data. Data management becomes more complex and time consuming in this environment, and since more data is being archived, there are more tapes to manage as well. IT managers need tape solutions that can help reduce complexity and streamline data management processes to drive increased efficiency.

With the introduction of Oracle's StorageTek T10000C tape drive, Oracle has raised the bar again, delivering the world's highest capacity tape drive and the first tape solution to offer an exabyte (one million TBs) of storage capacity in a single tape library.

Unmatched Capacity

Oracle's StorageTek T10000C tape drive is the third generation of the StorageTek T10000 enterprise tape drive family. StorageTek T10000 tape drives are known for high capacity, high throughput, and enterprise quality to meet the requirements of the most demanding open systems and mainframe datacenters.

Typically, tape drive technology doubles capacity every two to three years. With the StorageTek T10000C, Oracle has achieved an unprecedented technological advancement, increasing capacity by five times that of Oracle's StorageTek T10000B. Native capacity on the StorageTek T10000C tape drive is a massive five terabytes. Five terabytes is 25% more capacity than any other enterprise-class drive and more than three times the capacity offered by any mid-range drive. To keep up with the 5 TB capacity, the StorageTek T10000C speeds data to and from tape with a native throughput of 240 megabytes (MB) per second.

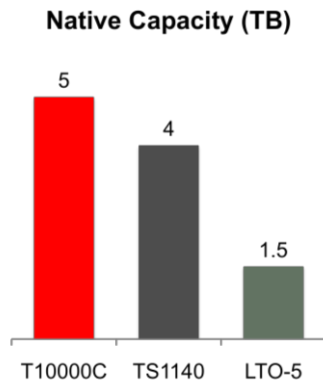


Figure 1. A comparison of StorageTek T10000C capacity with other tape drives.^{3,4}

Like its predecessors, the StorageTek T10000C supports both native mainframe and fibre channel attach. It is engineered to withstand the most demanding 24x7 enterprise data centers. In-drive encryption and write once read many (WORM) technology help ensure data security.

Compatibility with Oracle's StorageTek SL8500 and StorageTek SL3000 modular library systems protects previous tape library investments, while backward read compatibility with tape cartridges written by StorageTek T10000A and StorageTek T10000B tape drives also facilitates migration to the new technology.

³ HP LTO-5 Specifications, <http://h10010.www1.hp.com/wwpc/us/en/sm/WF06a/12169-304612-3446236-3446236-3446236-4150338.html>

⁴ IBM TS1140 Data Sheet, <http://www-03.ibm.com/systems/storage/tape/ts1140/index.html>

Lowering Tape Acquisition Costs

With the StorageTek T10000C tape drive's high capacity and fast throughput, fewer drives, libraries and media are needed to do the same job, resulting in acquisition costs that are 20% lower than other tape vendors.

A capacity up to three times greater than other tape drives, means fewer media cartridges are needed to store the same amount of data. As shown in Figure 2, a 10 PB archive, using 5 TB StorageTek T10000C tape drives, requires 2,000 media cartridges. That same 10 PB archive using 1.5 TB LTO-5 tape drives requires 6,667 media cartridges. Figure 2 highlights the savings in cartridges and libraries that can be achieved by using the StorageTek SL3000 tape library with the StorageTek T10000C tape drive. With fewer cartridges, the tape library can also be much smaller. The resulting savings in floorspace are discussed in the next section.

10 PB Example

Library	Drives	Libraries	Cartridges	Sq. Ft *
StorageTek SL3000	T10000C	1	2,000	45
StorageTek SL3000	LTO-5	2	6,667	116
Spectra Logic T950	LTO-5	1	6,667	127
IBM TS3500 (HD)	LTO-5	1	6,667	223
HP ESL	LTO-5	10	6,667	246
Quantum i6000	LTO-5	2	6,667	269
9310	T10000A	4	20,000	839

* Floorspace calculations include service area

Figure 2. Example of a 10 PB archive.^{5,6,7,8}

Another factor affecting tape acquisition costs is the number of drives required to achieve adequate throughput for completing a job. Backups, for example, often have a limited time window for completion. With throughput that is nearly two times more than any mid-range tape drive, a StorageTek T10000C tape drive can write the same amount of data as two mid-range drives. In an internal test, only six StorageTek T10000C tape drives were required to backup a 28 TB Exadata database in a four hour backup window, while 12 StorageTek T10000B tape drives were required to do

⁵ Spectra T-950 Library User Guide, <http://www.spectrallogic.com/index.cfm?fuseaction=home.displayFile&DocID=283>

⁶ IBM System Storage TS3500 Tape Library Introduction and Planning Guide IBM 3584 Tape Library, <http://www-01.ibm.com/support/docview.wss?uid=ssg1S7001738>

⁷ HP ESL QuickSpecs, http://h18006.www1.hp.com/products/quickspecs/11877_div/11877_div.pdf

⁸ Quantum Scalar i2000 Library Planning Guide, <http://downloads.quantum.com/scalar/S2K/6-00418-13A.pdf>

the same job. Likewise, LTO-5 tape drives, which can deliver 140 MB per second of throughput, or slightly more than half that of the StorageTek T10000C, also require nearly twice as many tape drives to do the same job.

Reducing Management Costs

As the data explosion continues, resources available to manage that data are not keeping up. Datacenter managers are often required to manage 20 to 40 percent annual increases in data with the same (or fewer) number of employees and capital expenditures. Oracle has designed the StorageTek tape product family to maximize scalability, simplify management, and provide the best value for the money.

As shown in Figure 2 above, a StorageTek T10000C, at 5 TB of capacity, requires three times fewer cartridges to store the same amount of data as a 1.5 TB LTO-5 tape drive. Fewer cartridges and tape drives mean smaller tape libraries and less complicated hardware to manage, ultimately leading to fewer people needed to run the datacenter.

Reducing Datacenter Floorspace Costs

An added benefit of minimizing storage equipment requirements is a reduction in data center floorspace needed. The last column in Figure 2 compares floorspace requirements for several tape storage systems. The StorageTek SL3000 modular library system with StorageTek T10000C tape drives requires one-fifth the floorspace of other comparable systems.

Minimizing Operational Risk by Reducing Cartridge Movement

The StorageTek T10000C's unprecedented ability to write 5 TB to a tape cartridge at one time means that the tape cartridge stays in the drive longer. Fewer exchanges are needed to write the same amount of data. Moving tape cartridges less leaves fewer opportunities for mechanical errors, further reducing the possible need for human intervention.

Lowering Power and Cooling Costs

Power and cooling are also an important consideration when evaluating datacenter expenses. The lightning-fast throughput speed of the StorageTek T10000C tape drive leads to reduced power consumption, as less power is ultimately required to write the same amount of data. Figure 3 shows a comparison of the energy required to record a 10 PB archive using existing storage technologies versus the StorageTek T10000C.




Figure 3. StorageTek T10000C requires less energy to record a 10 PB archive.^{9,10,11}

World's First Exabyte Tape Library

The StorageTek SL8500 modular library system has the ability to scale up to 100,000 cartridge slots in a single library instance. Used with the StorageTek T10000C tape drive, at 5 TB per tape cartridge, the StorageTek SL8500 can store 500,000 TB, or 500 PB. Assuming a 2:1 compression ratio, a total of 1,000 PB, or one exabyte, can be stored in a single library. The world's first exabyte tape storage system.

Conclusion

Acquisition costs and management costs are important to consider when evaluating any datacenter purchase. Acquisition costs are not only driven by unit prices, but also by the number of units required to do the same work. Management costs are affected by the number of people required to manage the equipment, the floorspace required, and the energy consumed.

Oracle's StorageTek T10000C tape drive with its unprecedented 5 TB capacity and 240 MB per second throughput helps customers achieve savings in all of these areas.

- Highest capacity means fewer tape cartridges, smaller tape libraries, smaller footprint.
- Faster throughput means fewer tape drives, faster back-up windows and lower energy consumption.

Each of these features contributes to a reduced total cost of ownership, often hard to achieve in this environment of massive data growth and limited IT management resources. The high performance

⁹ EMC Centerra Specification Sheet, <http://www.emc.com/collateral/hardware/specification-sheet/c932-centera-cas-ss.pdf>

¹⁰ IBM TS1140 Specifications, <http://www-03.ibm.com/systems/storage/tape/ts1140/specifications.html>

¹¹ HP LTO-5 Quickspecs, http://h18000.www1.hp.com/products/quickspecs/13572_na/13572_na.html

StorageTek T10000C is the world's highest capacity enterprise tape drive, and delivers the lowest total cost of ownership.

For more information about Oracle tape storage solutions, visit oracle.com, or call +1.800.ORACLE1 to speak to an Oracle representative.



Lowering Storage Costs with the World's First
Five Terabyte Tape Drive
May 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

Copyright © 2011, 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. 0511

Hardware and Software, Engineered to Work Together