

Sun StorageTek™ T10000

Fibre Channel Tape Drive

***Raising the Bar of Performance: A Performance Comparison Test by
Percept Technology Labs, Inc.
January, 2007***

The Sun StorageTek™ T10000 fibre channel tape drive is a high duty-cycle enterprise-class storage solution that achieves industry-leading data throughput rates - even in today's most demanding storage environments.


[1] – Why Raise the Bar?

In today's lightning-fast, information-driven business environment, data accessibility is becoming a 24x7 activity. With data transfer occurring constantly, and in ever-increasing amounts, far less time is being allocated to safeguard your data. In other words, data backup windows keep getting smaller and smaller.

But at the same time, the need for data storage is becoming more critical than ever. Rigorous compliance regulations recently enacted worldwide, for example, now mean that failure to recover data can do more than just compromise your company's competitive position; it can incur stiff legal penalties, as well.

The Sun StorageTek™ T10000 tape drive system was developed with all of this in mind. Specifically, the T10000 delivers industry-leading data throughput performance that enables fast-paced companies to cope with the challenges of doing business in today's world - by helping them to quickly backup their data stores while minimizing system "down-time." This technical paper demonstrates why the T10000 is the ideal choice for businesses that simply cannot afford to cut corners when it comes to data accessibility.

Top Close-Up View of the T10000



1 All testing for this technical paper was performed with a point-to-point configuration using Dell 1850 servers with a single 2.8 GHz dual-core Xeon processor, Q-Logic QLA2460 4GB HBAs, 2-3GB of RAM, Windows 2000 Server, SP4 Rollup1, and Redhat Enterprise Linux 4 Server.

[2] – Why the T10000?

Your company's data is arguably its most valuable asset. Many companies trust tape to store their data. In fact, tape is by far the most widely used technology for data backup in long-term storage environments. But while much 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 systems are designed for three purposes: to perform in high duty-cycle robotic environments (duty cycles between 70 and 100% are commonly expected); to recover without failure; and to integrate into a heterogeneous environment (FICON for z/OS or Fibre Channel for Open Systems).

However, not all enterprise tape systems are created equal. For example, modern applications that manage and secure data typically require that hardware perform frequent start, stop, and seek operations. Of course, such intense data transfer activity rarely ceases in today's fast-paced and always-on business environment - even after "regular" business hours. Unfortunately, not all tape drives are capable of delivering optimal performance when pushed to these extremes.

But Sun's T10000 enterprise-class tape drive is. Indeed, Sun's engineers and product developers recognize your company's need to minimize system down-time. They also understand that the performance of your tape drive - and tape cartridge - can profoundly impact your bottom line. The T10000 is Sun's answer to their customers' requirement that their data be written and read as fast as possible - without compromising reliability.

As this technical paper¹ will show, the T10000 is a high duty-cycle enterprise-class tape drive that significantly raises the bar with respect to data throughput in a robotic environment. The result? Lightning-fast drive performance combined with trustworthy data recall.

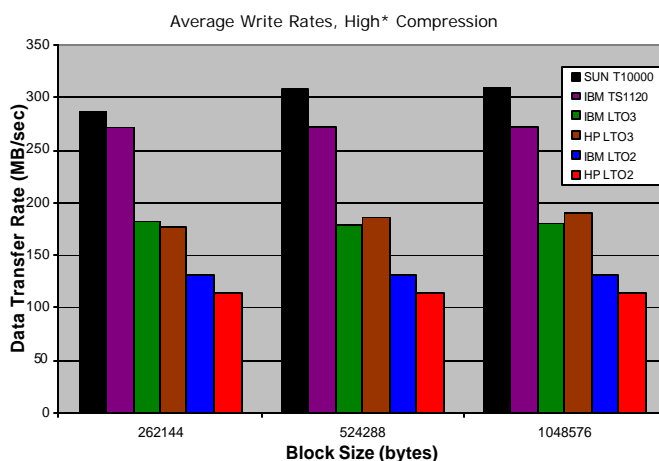
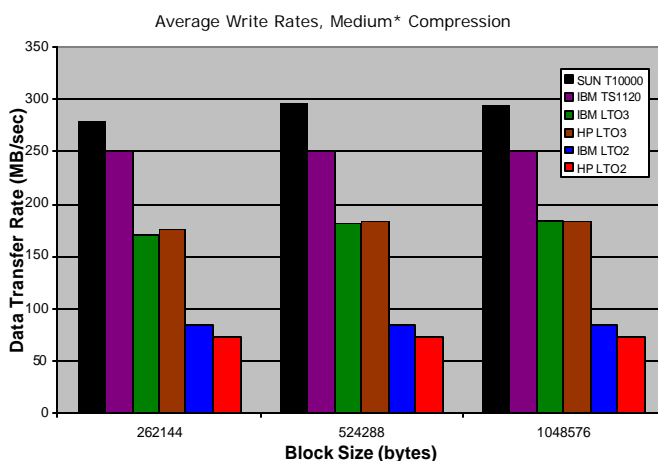

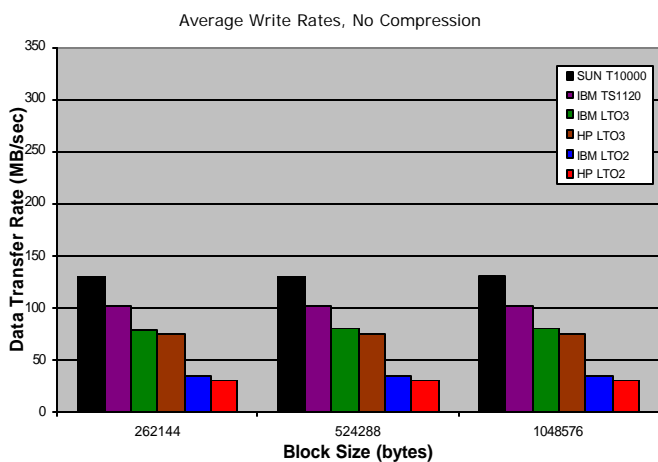
[3] – High-Performance Features of the T10000

The T10000 tape drive incorporates proven technology to bring you unprecedented data throughput performance. One of the major advantages the T10000 offers in this realm is industry-leading *capture and delivery*. Capture and delivery refers to the overall amount of work a tape drive is capable of doing. The T10000 maximizes capture and delivery by minimizing the duration of both basic drive operation and data transfer. The end result is the ability to **store and retrieve more data in a given amount of time**. This, of course, means your information system spends less time waiting for backup activities - which, in turn, can improve your company's performance.

[3.1] – Faster Data Transfer

The primary way to improve performance is to maximize data transfer rate, which is the same as minimizing the time the drive takes to store and/or retrieve a given amount of data. Independent laboratory read/write tests were performed to assess maximum data throughput by timing write and read operations. Each write operation transferred 10 GB of data in the form of a different compressible pattern. Write operation was then followed by rewind and read commands. Each write/read operation was performed with specific block sizes in order to measure the effect of data block size on transfer rate. These tests showed that, in the most common block size range (256-KB to 1-MB), the T10000 offers the fastest data transfer rate of any enterprise tape drive in the industry - whether transferring non-compressed, slightly-compressed, moderately compressed, or highly compressed data.²

Average write³ rates in 256-KB to 1-MB block size range



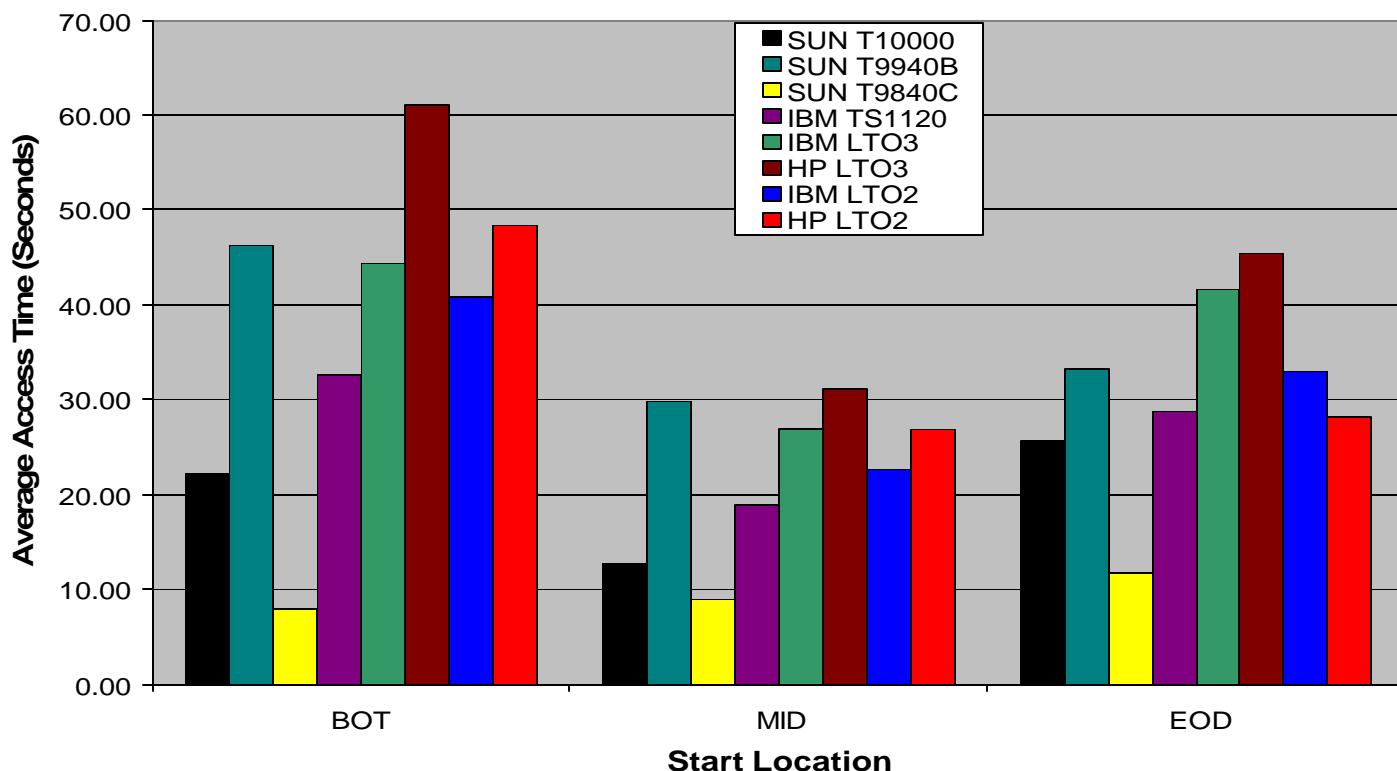
2 All data streams used during these test procedures were software-generated and homogenous. Data compressibility was quantified by passing the stream through a LZSS open source algorithm. Actual compression ratios achieved on the tape by the T10000 drive were 1:1 (no compression), 1.49:1 (low compression), 2.39:1 (medium compression), and 3.74:1 (high compression). These values were calculated by writing a full tape at each compression setting.

3 Throughput values for read operations were similar to those for write operations.

[3.2] – Faster Data Access

Another way to improve capture and delivery is to increase the data access rate, which refers to the speed at which the drive seeks and locates a specific set of data at a given location on the tape. This capability is especially important when recovering from incremental and multiplexed backups. To investigate drive access rates, a 10-GB data set was written by each drive in the selected test group. Next, each drive performed timed search operations from the beginning, middle, and end of the 10-GB data set to 15 different data blocks within that data set. This seek test showed that, aside from Sun's own T9840C access-centric enterprise tape drive, the T10000 delivers superior data access rates compared with the competition, in every class of data access operation.

Average time to locate 15 64-KB data blocks within a 10-GB data set at the beginning (BOT), middle (MID), and end (EOD) of data.⁴



[3.3] – Faster Cartridge Loading/Unloading

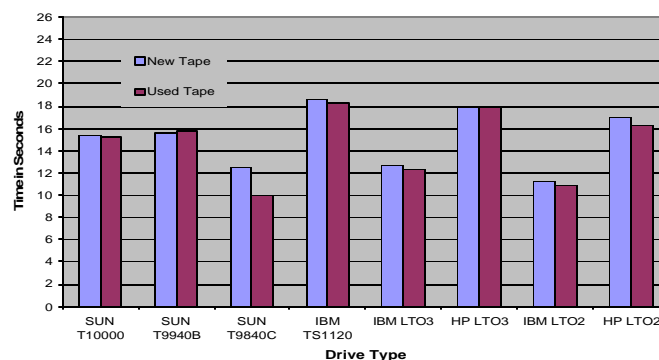
A third way to improve capture and delivery in a library system is to minimize tape handling time, i.e. the amount of time the cartridge spends being handled by the library robot, as well as the time the tape spends in the drive when data is not actually being transferred. Tests measuring the load and unload times of both new and used cartridges show that the T10000 achieves a total tape loading and unloading time lower than most other drives studied, resulting in a higher percentage of time spent performing data transfer operations. It is notable that the access-centric Sun T9840C tape drive is still the fastest in this category.

Average combined load and unload times (using both new and used cartridges)

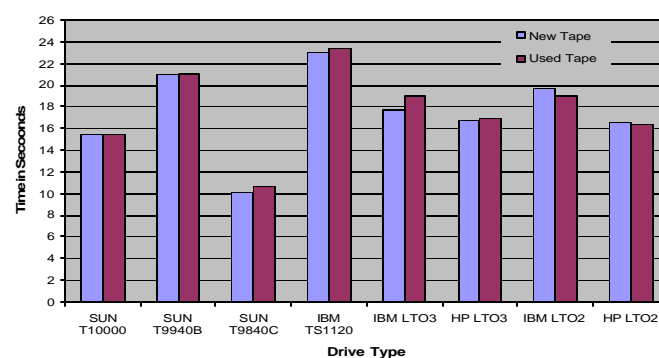
Drive	Average time to load and unload
T10000	31.3 seconds
T9940B	36.7 seconds
T9840C	21.2 seconds
TS1120	42.5 seconds
IBM LTO2	30.1 seconds
IBM LTO3	31.0 seconds
HP LTO2	33.0 seconds
HP LTO3	34.7 seconds

4 For the T10000, 10GB does not fill the full length of the tape. The 10GB dataset was chosen randomly for this test. Access times for all devices will differ depending on the tape capacity and the size of the dataset.

Average load times (using both new and used cartridges)



Average unload times (using both new and used cartridges)



[3.4] - Conclusion

The T10000 outperforms its closest rival the IBM TS1120 in Data Transfer Rate (at medium compression), 10.1% (256K block size), 15.3% (512K block size), 14.4% (1024K block size), in data access time, BOT 30.7%, MID 32.7%, EOD 10.5% and by 26% in total load and unload time.

[4] – Optimizing Your Data Backup System with the T10000


When assessing an enterprise-class tape drive, raw tape drive performance is really just part of the story. Other components of the backup system, such as the software, the server, and how these components interact with the drive, also impact ultimate performance. Thus, getting the whole system to work together in an optimal manner can be the key to obtaining the most out of your tape drive. In the interest of helping our customers achieve this ideal state of operation, the T10000 and a group of competing drives were evaluated with three common backup software applications, and on two common server platforms.⁵

5 Data compression ratios referenced in this section are: 1.62:1 (low-compressed data), and 1.91:1 (medium-compressed data). Also, additional hardware was used: A RAID-0 consisting of 4 Seagate 15,000 RPM Cheetah Drives, and an Adaptec 320 SCSI HBA, PN: AHA-29320.

[4.1] – Symantecä Veritas NetBackup with Windows Server

Veritas currently commands the largest share of the backup application market. When properly tuned, the Veritas NetBackup application, coupled with the T10000 tape drive, achieved the fastest throughput rate of all drives tested. Tuning proved to be an important step in maximizing NetBackup data backup rates, providing a six-fold increase in performance between tuned and non-tuned systems for all drives tested. Of all the tunable factors explored for Veritas, two had the largest influence on T10000 system performance: size and number of backup buffers. These two factors accounted for as much as 98% of the improvement in system performance compared to non-tuned systems (the other 2% was accounted for by physical and virtual memory settings).

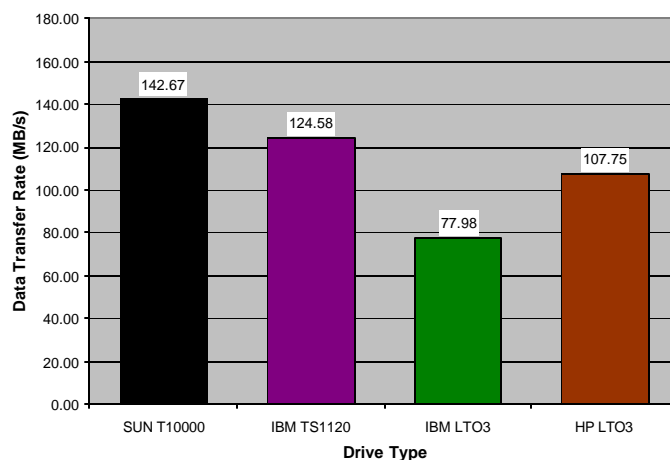
Average data throughput of a 200-GB data backup, using medium-compressed data at optimized NetBackup settings



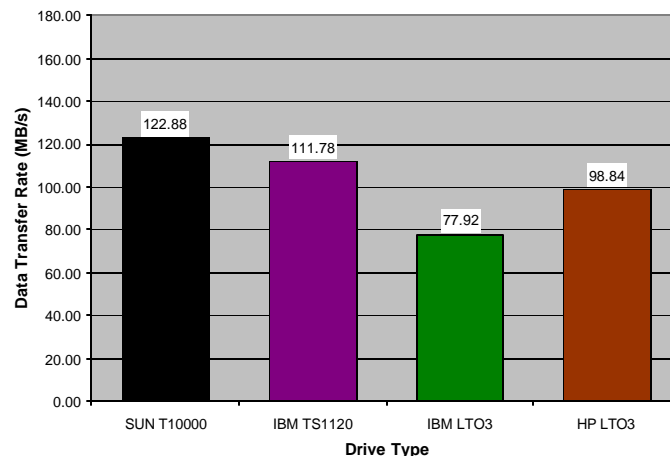
[4.2] – Veritas NetBackup with RedHat Enterprise Linux 4 Server

When coupled with the Veritas NetBackup application on a Linux server platform, the T10000 tape drive again achieved the fastest throughput rate of all drives tested. Tuning, which involved adjusting both the number and size of the buffers used, again proved important in maximizing performance. Specifically, proper tuning improved data throughput by 200 to 300% over default operating conditions.

Average data throughput of a 200-GB data backup, using medium-compressed data at optimized NetBackup settings



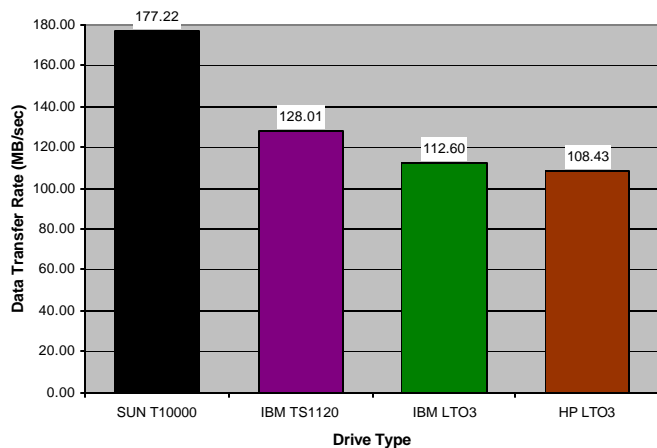
Average data throughput of a 200-GB data backup, using low-compressed data at optimized NetBackup settings




[4.3] – EMC Legato Networker 7.3 with Windows Server

EMC currently commands 31.8% of the backup application market. When properly tuned, the Legato Networker 7.3 application coupled with the T10000 drive achieved the fastest throughput rate of all the drives tested, outperforming the next-best system by as much as 40%. Proper tuning was critical in achieving this performance. Of all the tunable factors listed by Legato, four had the most influence on T10000 system performance: MaximumSGList OS parameter, size of physical memory on the server, amount of virtual memory available in the system, and device block size. These four factors accounted for as much as a 250% improvement in system performance compared with the non-tuned system.


Average data throughput of a 200-GB data backup, using medium-compressed data at optimized Networker 7.3 settings



Average data throughput of a 40-GB incremental data backup, using medium-compressed data at optimized Networker 7.3 settings



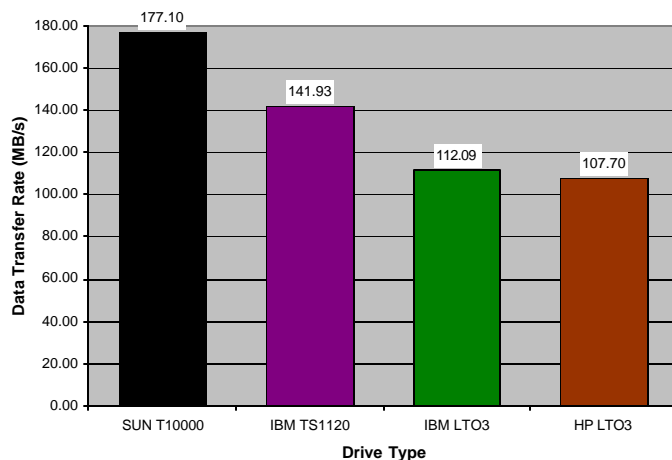
Average data throughput of a 200-GB data backup, using low-compressed data at optimized Networker 7.3 settings




[4.4] – EMC Legato Networker 7.3 with RedHat Enterprise Linux 4 Server

When properly tuned, the EMC Legato Networker 7.3 application coupled with the T10000 drive again far surpassed the throughput rates of all other drives tested. The tunable factors adjusted for the Linux platform, included increasing the amount of physical memory on the server and adjusting the device block size.


Average data throughput of a 200-GB data backup, using medium-compressed data at optimized Networker 7.3 settings



Average data throughput of a 40-GB incremental data backup, using medium-compressed data at optimized Networker 7.3 settings




Average data throughput of a 200-GB data backup, using low-compressed data at optimized Networker 7.3 settings



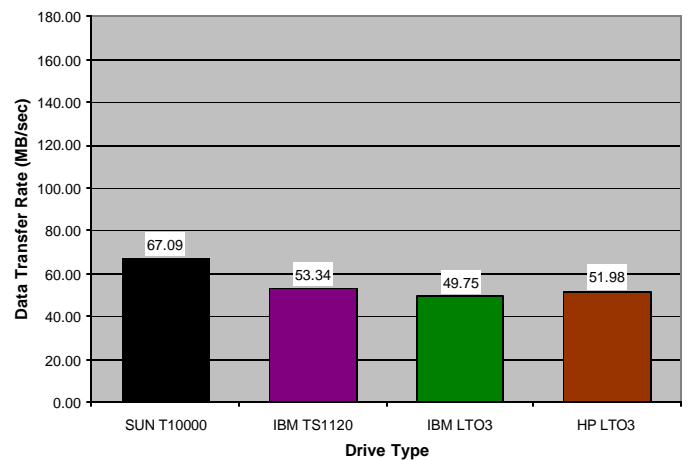
[4.5] – IBM Tivoli Storage Manager with Windows Server

When optimized, the Tivoli Storage Manager (TSM) application coupled with the T10000 drive achieved the fastest throughput rate of all drives tested. Proper tuning was required to achieve this performance. Specifically, system RAM was increased from 2 to 3 GB, MaximumSGList was optimized, Windows Virtual Memory was doubled, and one key TSM server parameter - TXNBYTELIMIT - was maximized. These actions increased data transfer rates by about 300% across the board for low-compressed data, and by anywhere from 200 to 400% for high-compressed data.

Average data throughput of a full 200-GB data backup, using medium-compressed data at optimized Tivoli settings



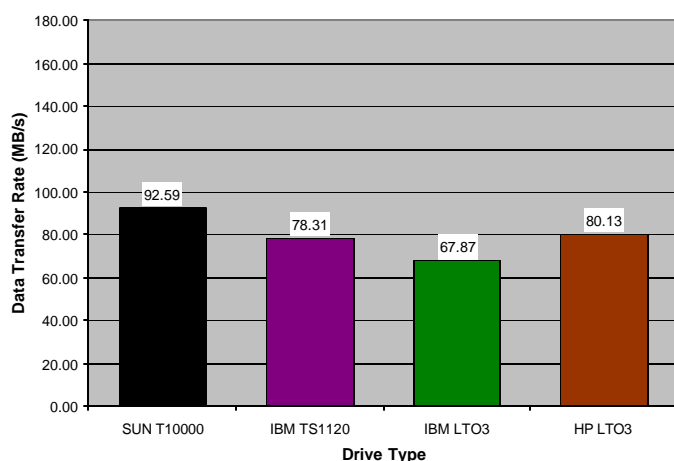
Average data throughput of a full 200-GB data backup, using low-compressed data at optimized Tivoli settings




[4.6] – IBM Tivoli Storage Manager with RedHat Enterprise Linux 4 Server

TSM coupled with the T10000 drive again achieved the fastest throughput rate of all drives tested. Tuning involved maximizing the TXNBYTELIMIT as it was on the Windows server. Physical Memory was also increased from 2GB to 3GB. Virtual Memory was not increased on the Linux server. These actions increased data transfer rates by about 200 to 400% across the board.

Average data throughput of a full 200-GB data backup, using medium-compressed data at optimized Tivoli settings



Average data throughput of a full 200-GB data backup, using low-compressed data at optimized Tivoli settings



ISV Performance Summary for the T-10000

Drive	ISV Package	OS	Optimized Transfer Rate
SUN T10000	NetBackup	Windows	160.44 MB/s
SUN T10000	NetBackup	Linux	142.67 MB/s
SUN T10000	Networker	Windows	177.22 MB/s
SUN T10000	Networker	Linux	177.10 MB/s
SUN T10000	TSM	Windows	99.14 MB/s
SUN T10000	TSM	Linux	92.59 MB/s

All testing for this technical paper was performed by Percept Technology Labs, Inc., located in Boulder, Colorado. Percept Technology Labs, Inc. is an established, independent product test and consulting company, with a ten-year history of high tech testing expertise and service excellence. For more information please visit our website at www.percept.com.

[5] – Summary

When choosing an enterprise tape system, speed and performance matter. This is especially true in today's 24x7 business environment. Simply put, the faster you can execute backup and recover commands to and from your tape, the faster your I-T system can return to the critical matter of managing your company's data. When both key measures of data throughput are considered - tape residency time and capture/delivery capability - the T10000 offers compelling performance advantages relative to the competition.