
An Oracle White Paper
January 2011

# Evaluating Tape Drive Performance

ORACLE

## Executive Overview

Performance and capacity of tape drives continues to increase. The StorageTek T10000B has a capacity of 1TB and throughput of 120 MB/s, LTO-5 has a capacity of 1.5TB and throughput of 140 MB/s, and the StorageTek T10000C has a capacity of 5TB and throughput of 240 MB/s. This paper discusses the importance of key specifications typically used to evaluate tape drive performance.

## Introduction

In determining which tape drives to buy, host interface speed is often a major gauge of tape drive performance. While interface speed is certainly important, it must be evaluated as part of a total system solution. Looking at host interface performance in isolation is equivalent to determining which automobile to buy based on the top speed shown on the speedometer. If the speed limit and road conditions require that a car be driven at 70 mph or less, the distinction between cars with a top speed of 160 mph versus 120 mph becomes less important.

Existing tape drives virtually all provide the 4Gb SCSI Fibre Channel Protocol (FCP) which supports a sustained throughput of between 350 and 400 Megabytes per second (MB/s). Next generation drives, like LTO-5, tout upgrading this interface to an 8Gb FCP, to support up to 800 MB/s.

While host interface speed is valid component of performance, more important factors should be considered. This paper will discuss the total performance of a tape drive, and how it should be evaluated to determine which drive provides the best performance solution.

## The Limits on Tape Drive Performance

Tape performance is always driven by three key components:

1. The speed that the storage application sends data to, or processes data from, the drive.
2. The speed of the host interface between the drive and the application.
3. The speed that the tape drive writes or reads data at the head/media interface.

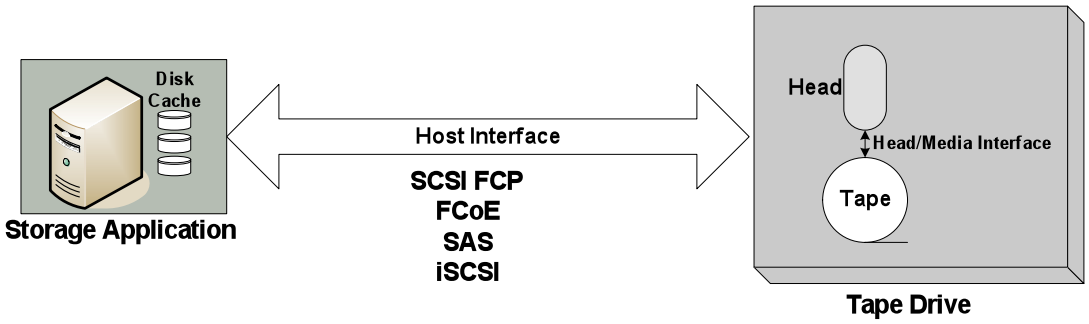


Figure 1: Performance Components

Application speed is typically the bottleneck in sustained throughput to a tape drive during write operations. This occurs because applications create metadata that is stored with each record and then package many records together before writing them to tape. This process is reversed for read operations with data being unpacked and translated from a serial tape format to a random access file format. Most tape drive applications (backup applications in particular) are based on code that has not been significantly changed for decades. These applications don't take advantage of the high performance characteristics of modern tape systems and often limit the amount of expensive disk cache used to hold data as it is being formatted for tape.

The host interface speeds for protocols like SCSI FCP, Serial Attached SCSI (SAS) or Fibre Channel over Ethernet (FCoE) are seldom the limiting factor in drive performance. With existing speeds of 200 MB/s to 400 MB/s, these interfaces are almost never the bottleneck in the performance pipe. Increasing this interface speed to 600-800 MB/s, has little, if any, impact on sustained drive performance.

In environments where applications have been optimized for performance, tape drive throughput at the head/media interface becomes the performance bottleneck. Existing enterprise and mid-range tape drives typically transfer data at about 100 - 150 MB/s. Compression ratios up to 2:1 are often achieved in these systems so that performance increases to 200 – 300 MB/s. Even with compression, these systems only require a 4Gb SCSI FCP interface to maintain sustained performance.

## Tape Drive Performance in Backup Environments

In a typical backup environment data is moved across a Local Area Network (LAN) to a media server. The media server controls the mounting of tape, catalogs the movement of records to tape volumes and manages tapes moved off site. The disk cache of the media server holds data until enough is buffered for movement to tape without causing a significant loss in streaming performance. The data coming across the LAN to the media server can be delivered at various speeds. These speeds range from 10MB/s to 100 MB/s for individual work stations, up to 1 GB/s (or higher), between servers. As this data is gathered in the disk cache, metadata is created and the data is aggregated into tape records. The media server then manages the loading of the correct tapes and transferring data to those tapes.

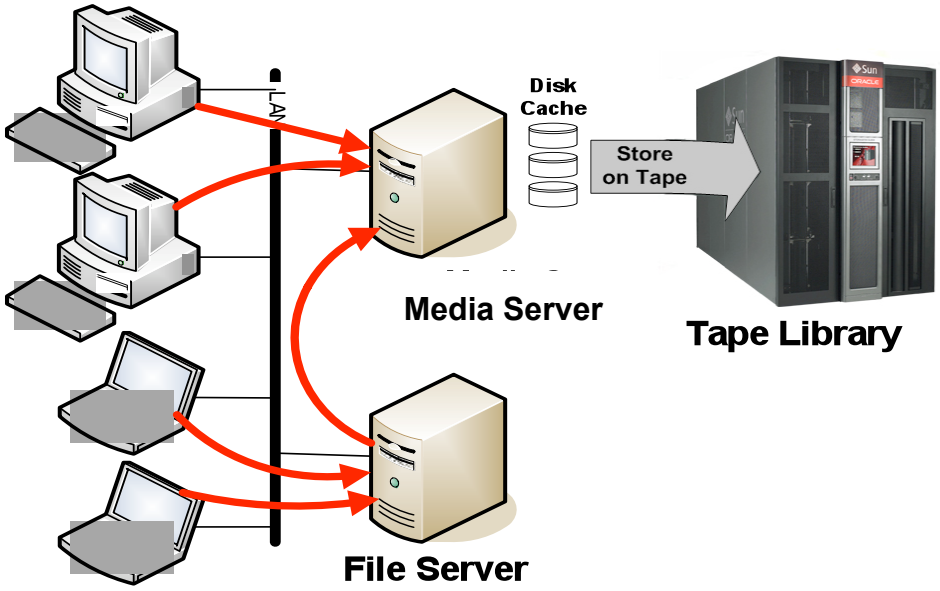


Figure 2: Backup using a Media Server

The above configuration results in a bottleneck at the media server where data is accumulated and processed for tape storage.

To avoid this bottleneck the Network Data Management Protocol (NDMP) was created to allow servers and workstations with connectivity to back data up directly to tape libraries. In this configuration a Data Management Application (DMA) server controls the tape library, and creates catalogs, while data is backed up directly from the workstation or file servers. Creation of metadata and coordinating tape access by the DMA takes as much time, or more than, the typical media server configuration shown in Figure 1. The main benefit of this configuration is a reduction in the disk cache used by the media server.

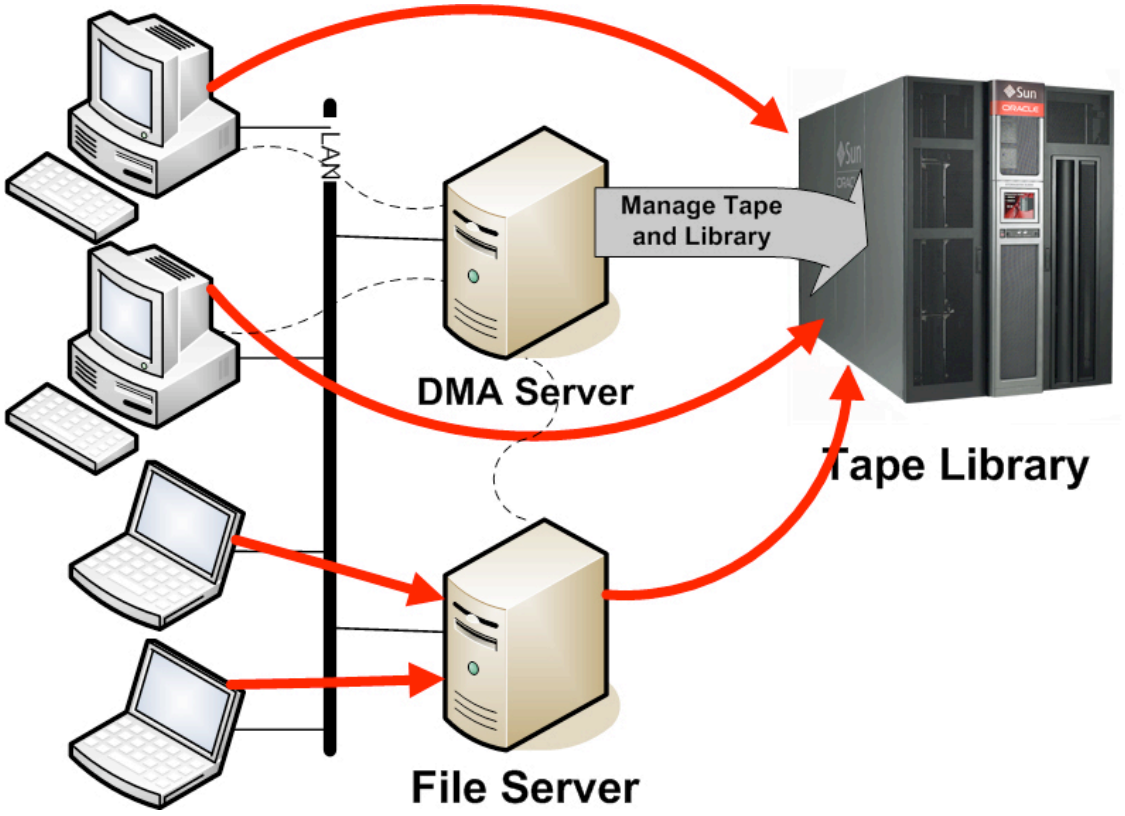


Figure 3: NDMP Backup Configuration

In almost all backup environments the typical sustained system throughput is 50-60 MB/s. Many environments experience sustained performance in the 20-40 MB/s range. As a result, the backup application is almost always the bottleneck in tape backup performance. The performance of the host interface and the tape drive head/media interface has a marginal impact, if any, on these systems.

## Tape Drive Performance in an Archive or Direct Attach Environment

Many solutions have been developed that optimize tape drive performance in an archive, tiered storage, or high performance server or mainframe environment. Figure 1 shows how these environments are typically setup. Oracle's Exadata Storage Server and Oracle's Sun Storage Archive Manager (SAM) are two examples of these solutions. Unlike backup applications, these solutions create metadata within the context of a file system or database that is tightly coupled to the application. When records are written to tape, the creation of metadata is very concise with little or no intrusive processing. If a backup application like Oracle Secure Backup is used, it has been optimized for the specific solution with direct attachment to a dedicated tape library. In addition, the data is usually compacted, encrypted or transformed in some fashion before it is stored to tape. In these applications, individual pieces of data like files or database elements are unique. As a result, tape compression in

both of these environments is less than a 2:1. A properly configured SAM environment writes data to tape at a sustained rate of 100-150 MB/s. In an Exadata environment, the sustained performance of the tape drive at the head/media interface limits the performance of the system. As long as the performance of the host interface exceeds the sustained performance of the tape drive head/media interface, the host speed has little or no impact on the performance of the system.

## StorageTek T10000C Tape Drive Performance

Oracle is committed to redefining tape storage, by introducing the StorageTek T10000C tape drive, with an unprecedented native throughput of 240 MB/s at the head/media interface. This drive performance combined with a 4Gb (400 MB/s) SCSI FCP interface out performs all other tape drive at the critical compression ratios of 2:1 or less and typical tape block sizes of 256K bytes or larger. The following charts compare StorageTek T10000C (240 MB/s throughput and 4Gb interface) performance to StorageTek T10000B (120 MB/s throughput and 4Gb interface) and LTO-5 (140 MB/s throughput and 8Gb interface) tape drives.[1]

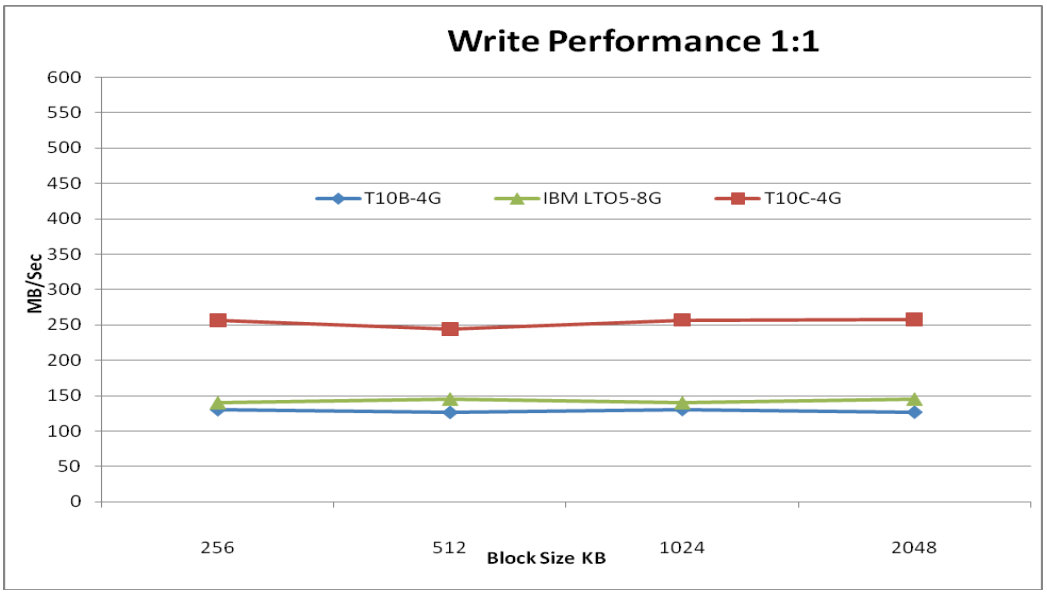


Figure 4: Write comparison at 1:1 compression ratio

[1] Performance charts generated using the following configuration: Server: Sunfire X4270, Processor: Intel family 6 model 25, 8 virtual processors, 2267 MHz, Operating System: Solaris 10, Generic_141445-09 i86pc, 4Gb HBA: Qlogic QLA2462, 8Gb HBA: Qlogic QLE2562, Application: Stkbench

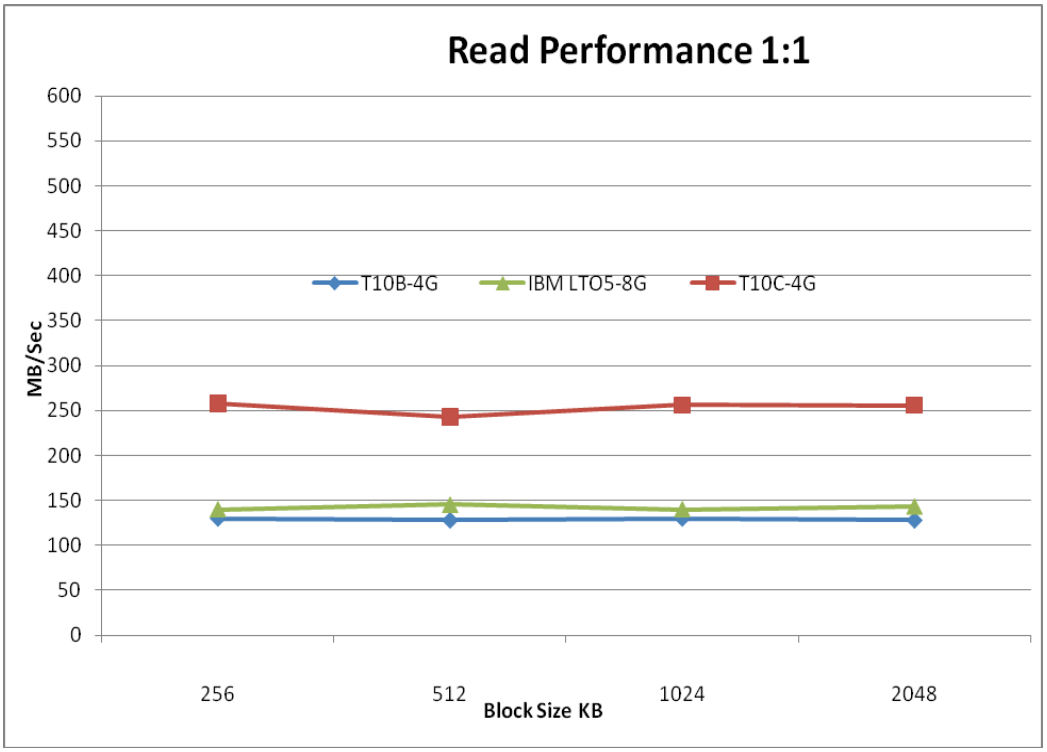


Figure 5: Read comparison at 1:1 compression ratio

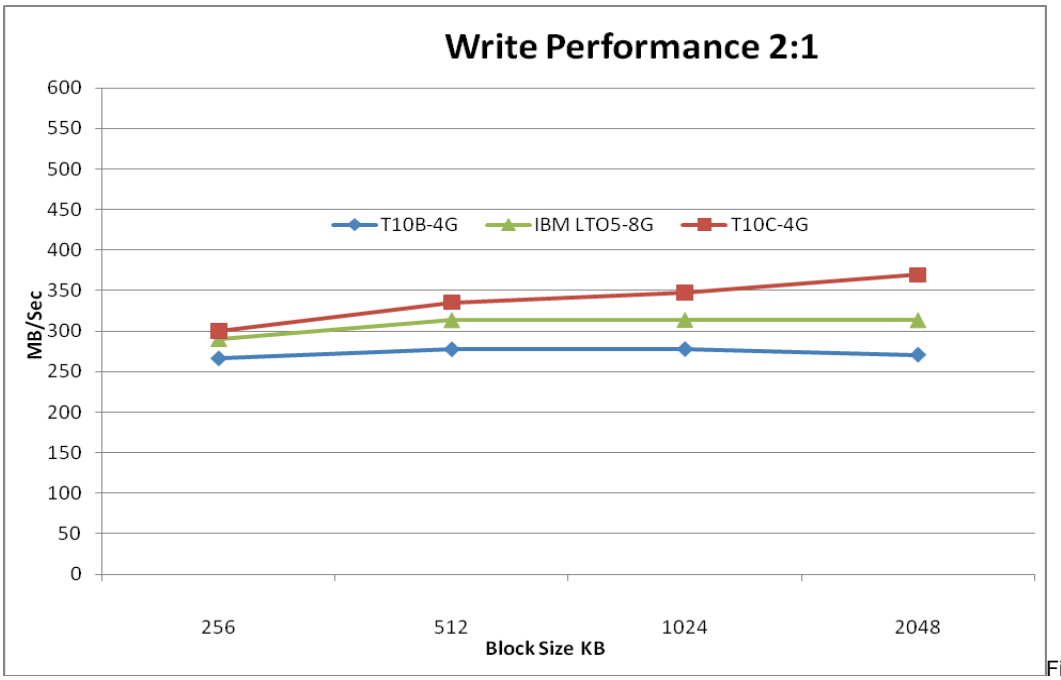


Figure 6: Write comparison at 2:1 compression ratio

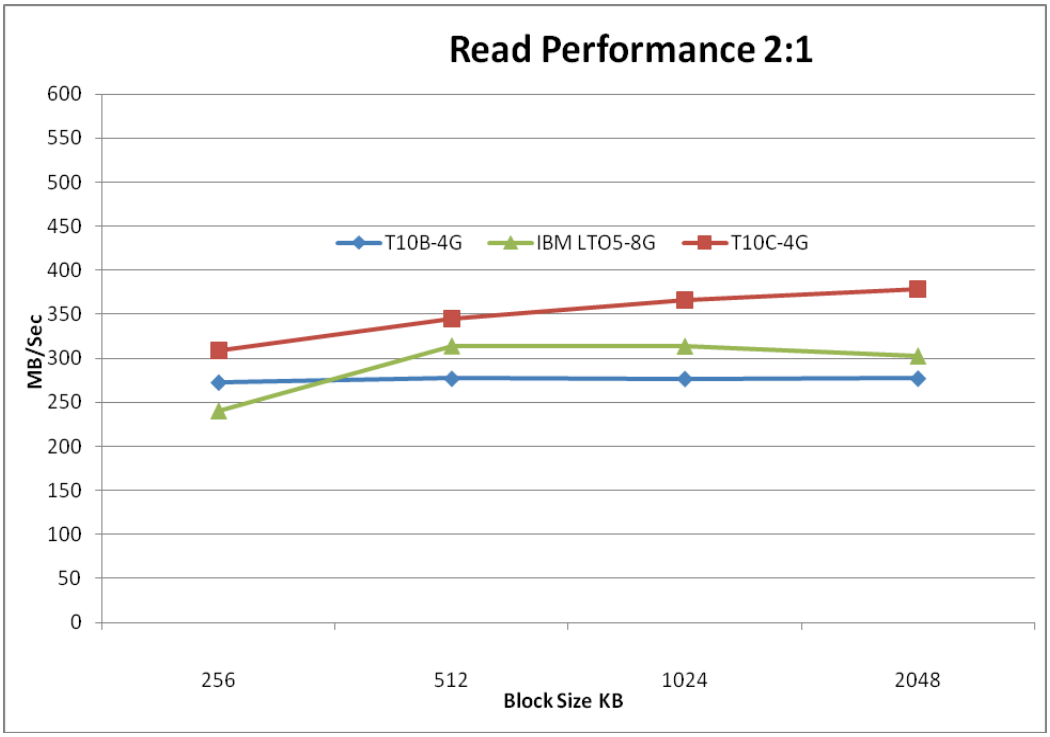


Figure 7: Read comparison at 2:1 compression ratio

We believe many applications are similar to SAM and Exadata and will experience a compression ratio of about 1.5:1 to 2:1. The StorageTek T10000C tape drive is optimized for this compression ratio.

## Conclusion

In most situations, either the storage applications' speed or the drive throughput speed are the limiting factors in tape drive performance. Typical storage application throughput of 50-60 MB/s falls far short of the 400 MB/s maximum speed of a 4Gb SCSI FCP interface. Even with compressed data, current tape drive technologies are also not able to write or read data any faster than a 4 Gb (400 MB/s) SCSI FCP interface can push the data.

Oracle is focused on delivering storage solutions rooted in true customer applications. By focusing on native drive performance, and doubling throughput on the newest StorageTek T10000 tape drive, Oracle is solving a real customer problem.

Evaluating Tape Drive Performance
Jan 2011
Author: Dwayne Edling

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**