

### Our address

24 The Bridge Business Centre  
Beresford Way  
Chesterfield  
S41 9FG

### Get in touch

Call us on: 01246 454222  
Email us: [web@coolspirit.co.uk](mailto:web@coolspirit.co.uk)  
Find us: [View location map](#)  
Web: [www.coolspirit.co.uk](http://www.coolspirit.co.uk)

### Office hours

mon - thurs 8:30am - 5:30pm  
fri 8:30am - 5pm  
sat - sun Closed

“

Boost your storage buying power...  
use ours!

”

Buy with confidence from  
Coolspirit your authorised  
Quantum Partner

**Quantum. ALLIANCE**  
*Platinum Partner*

**DLT-S4 Buffer Management**  
*Speedmatching*



## Introduction

In 2006, Quantum launched the market's highest capacity tape drive, the DLT-S4, with a record-breaking 800GB native capacity and 1.6TB compressed capacity. This high capacity drive is the first drive optimized for tiered storage.

This paper discusses 'speedmatching' and Quantum's buffer management approach used for optimizing performance of its new DLT-S4 drive. Quantum's DLT-S4 buffer management and speedmatching approach is compared to techniques used with the LTO-3 drives sold by HP and IBM. The minimal effects of repositioning on reliability are also discussed. The paper also publishes results received by an independent test vendor, Percept Technology Labs.

## Executive Summary

In the past, a tape drive was viewed as the speed bottleneck in data backup throughput. Several years ago, with the introduction of faster tape drives, IT managers were discovering that other system components, such as the operating system, ISV software, host adapters, servers and disk subsystems, were having trouble keeping up with the "streaming" performance capabilities now being delivered by these faster drives.

The performance variability among system components generally caused the tape drive to fall out of data streaming mode several times during a backup session, causing what is known as a write reposition, back hitching, or shoeshine. To help solve this performance inconsistency problem, tape drive vendors have devised various ways to maintain throughput under slow or varying host data transfer rate conditions.

Two such methods to stabilize performance are DLT's Digital Data Rate Agent (DDRA) and HP's Ultrium Data Rate Matching (DRM). DDRA is a pure digital solution, while DRM is a hybrid digital/analog/mechanical solution dependent on three drive subsystems for performance stabilization. IBM also uses a hybrid digital/analog/mechanical solution for performance stabilization.

Under a series of benchmark tests to compare the relative performance of various algorithms, it was found that DDRA enabled the DLT-S4 tape drive to deliver a performance advantage over the DRM-enabled HP LTO-3 Ultrium drive and the IBM LTO-3 Ultrium drive. Although the DLT-S4 native specification is 60MB/second compared to the HP LTO-3 and the IBM LTO-3 native specifications of 80MB/second, data received from the independent third party vendor shows the DLT-S4 capacity to be comparative to or better than the competitive LTO-3 drives in either a native or compressed environment. In fact, using 2:1 compressible data, the DLT-S4 performed up to 9% faster than the LTO-3 drives tested.

## Background

Today's midrange tape drives offer very high transfer rates. IT managers are finding that other system components, such as the operating system, ISV software, host adapters, servers and disk subsystems, are having trouble sustaining the high transfer rates required to get optimal tape backup speed "streaming" performance.

This is particularly true in environments where the tape drive is directly attached to the backup server and is responsible for backing up File, Print, and eMail application data. This type of data is generally comprised of mostly small (< 50K) and medium (< 100K) file sizes. In this configuration the speed at which the server transfers data to the tape drive can vary widely (see figure 1), affecting backup performance.

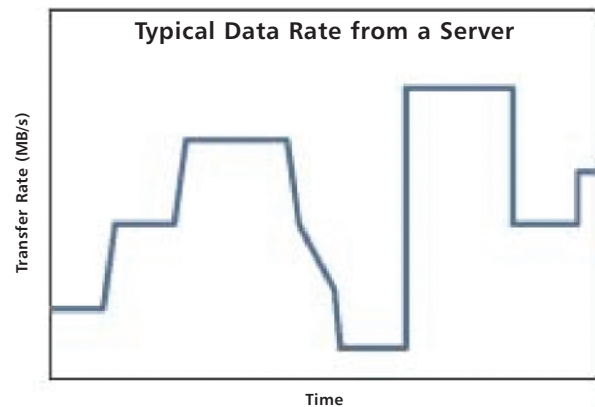


Figure 1. Typical Server Transfer Data Rate

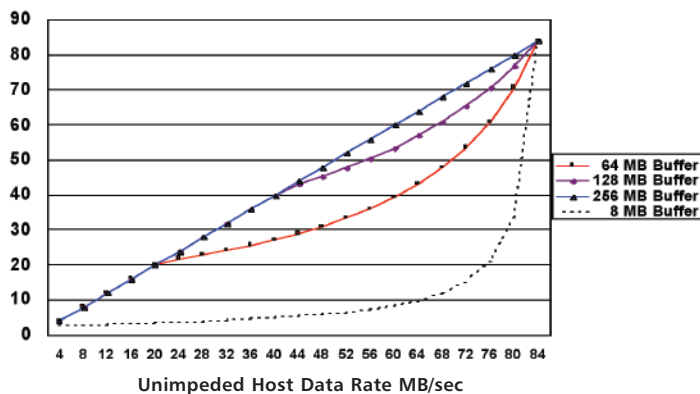
This performance variability can cause the tape drive to fall out of data streaming mode during a backup session and perform a tape reposition, or "shoeshine." Under normal operation, a tape drive's read/write heads are positioned precisely against the tape media which then "streams" past the heads. The tape drive takes data from the host and writes that data to the tape media. Because the tape media continually streams past the tape drive heads, the data transfer rate from the host must remain high enough to keep the tape streaming. If the transfer rate slows down too much the tape drive runs out of information to write and must stop in order to rewind and reposition itself for continued writing. Rewinding to the point at which it ran out of information takes a foot or two of actual tape. Once it reaches this point, it stops again and then restarts, accelerating so that it is up to appropriate speed by the time it reaches the last point it wrote data. The starting and stopping activities of this process looks much like the back and forth motion of a shoe shine, hence its name – shoeshining.

Tape shoeshining may degrade overall backup throughput. To solve the problem, tape drive vendors have devised various ways to maintain high backup throughput under slow host or varying host data transfer rate conditions. Quantum's DDRA speedmatching technique allows for enhanced performance of its drives during variable host data rates as described above.

### DLT-S4 Performance Speed Matching Solution

The DLT-S4 drive was architected with a pure digital buffering solution to achieve the highest performance for both speed and drive reliability. Cache sags caused by slow host data transfers and cache buffer mismanagement are the dominant reasons why tape drives fall out of streaming mode. Quantum uses both a large 256MB buffer (twice the size of the LTO-3 buffer) and an intelligent buffer management algorithm (DDRA) to minimize and eliminate these performance bottlenecks. In this case, studies have shown the 256MB buffer size ideal for optimized performance. Graph 1 below shows the enhanced linear transfer rate capabilities that are enabled by using a large buffer.

Buffer Size Comparison



Graph 1

### The DLT-S4 speedmatching buffer architecture objectives:

- Always provide buffered data access to the drive
- Maintain buffer level between full and empty
- Adjust device data rate to match compressed host data rate
- Always allow the host to transfer data

### Digital Data Rate Agent (DDRA) buffer management algorithm:

The DLT technology known as the Digital Data Rate Agent (DDRA) is used to dynamically manage the data flow rate from the host system to the DLTtape® media.

DDRA works by continually monitoring the data rate from the host system and intelligently buffering the data to the write/read channel to smooth out variability in the host transfer rate.

A technique known as 'watermarking' is one of the features used by the DDRA. Described below is an overview of the watermarking technique:

### Adaptive Write Watermarking:

For a fast host, the tape motion is started and writes data as soon as possible, this allows the host longer use of the cache since the host is rapidly filling the cache buffer. During this period, all data written to tape and verified is discarded as soon as possible, this allows the drive to fill the buffer with more data than the actual capacity of the buffer!

For a slow host, the tape motion is not started until the cache is almost full. This keeps the cache as full as possible before starting the write while the drive empties the cache rapidly, faster than the host is filling it. This provides a delay period and avoids costly data underruns minimizing the number of reposition cycles.

Some data may be retained in the cache to allow for quick repositioning near the end of data (EOD) to prepare for the next command. This is helpful when the host application requires a reposition to store multiple backup sets on the same tape.

### Adaptive Read Watermarking:

For a fast host, the tape motion is started and reads data as soon as possible; this allows the cache to become as full as possible while the host is rapidly emptying its buffer. All data transferred to the host is discarded as soon as possible, which allows the host to receive (as if cached) more data than the actual capacity of the buffer!

For a slow host, the tape motion is started when the cache is nearly empty. This allows the drive to read more data on each start without causing the host to underrun (run out of data to send). Eliminating host underruns improves the overall system performance.

Some data may be retained in the cache to allow for quick repositioning near the end of data (EOD) to prepare for the next command. This is helpful when the host application requires a reposition to store multiple backup sets on the same tape.

**Adaptive Watermarking Adjustments:**

This algorithm always favors the fast host. The default setting is to assume there will be a fast host. Adjustments in the fast host direction are very quick and decisive. Only when a slow host is observed will the drive deliberately adjust in the slow host direction.

**Reliability: Is Repositioning Really a Factor?**

In the past, for the tape drive industry there were valid concerns regarding the affects of repositioning on the drive and media reliability. Reliability is defined as the probability a product will function per specification without failure, under specified environmental conditions, for a specified period of time.

Incorporating design for reliability techniques, and today's newer component technologies, Quantum has virtually eliminated repositioning as a factor in reliability. This is demonstrated in both reliability testing and empirical field return data.

Some of these reliability concerns were based on older head and media technology that was more susceptible to wear than today's advanced technology. When using the newer MR (magnetoresistive) Head assembly and media technologies, this particular wear condition is not a consideration. In fact, repositioning can improve the performance of the drive. This improvement occurs due to the abrasive nature of the newer media acting as a cleaning mechanism on the MR Head assembly. This cleaning action can reduce tape stain buildup for stains that a cleaning brush is unable to remove. This additional cleaning action also improves the read/write signal strength. The improved signal strength actually enhances the performance of the drive by reducing repositions.

Another area of past concern was the tape transport system reliability. Again, using design for reliability techniques, in conjunction with newer technology in motors, rollers, and transport materials; Quantum has virtually eliminated repositioning as a factor for reliability.

Quantum verifies the drive and media meet all specifications for reliability with stringent testing, often using the worse-case conditions.

The only real consideration for today's drive and media technology when considering repositioning is the overall performance of the drive. Quantum has focused on improved performance while ensuring the drive and media reliability meets or exceeds specifications.

**Test Objective**

This White Paper describes the performance test results of tape super drives in slow host or varying host data transfer rate conditions. The objective of this test was to compare the performance of DLT-S4's Digital Data Rate Agent against the performance of the HP LTO-3 and the IBM LTO-3 tape drives adaptive tape speed. The test measured the write/read transfer rates with native (1:1) and then compression (2:1) modes, while limiting the host data transfer rate under various conditions. The specific host transfer rate conditions were:

- Random Host Transfer Rate
- Toggled Host Transfer Rate

Quantum contracted Percept Technologies, an independent Product Test and Development Firm, to perform the performance testing. Percept Technologies created the test methodologies and procedures, performed the tests and evaluated test results.

**TAPE DRIVE SPECIFICATIONS**

	DLT-S4	HP LTO-3	IBM LTO-3
Capacity (native)	800GB	400GB	400GB
Transfer Rate (native)	60MB/s	80MB/s	80MB/s
Cartridge Load Time (to BOT)	20 sec	19-22 sec	15 sec
Average File Access Time	70 sec	53 sec	54 sec
8TB Library Storage Density	10 cartridges	20 cartridges	20 cartridges
Interfaces	Ultra 320 SCSI 2 or 4 Gb FC	Ultra 320 SCSI 2 or 4 Gb FC	Ultra 160 SCSI 2 or 4 Gb FC
Track Density	2988 tracks/inch	1408 tracks/inch	1408 tracks/inch
Data compression	DLZ	LTO ALDC	LTO ALDC
Uncorrected Error Rate	1x10 <sup>17</sup> bits read	1x10 <sup>17</sup> bits read	1x10 <sup>17</sup> bits read
MTBF	250,000 hours @ 100% duty cycle	250,000 hours @ 100% duty cycle	250,000 hours @ 100% duty cycle

Table 1

**Test Method**

This test writes 5GB of data using selected compression ratio and block size combination. Each combination is maintained for five complete write/read passes. Following block sizes were tested:

- 16k
- 32k
- 64k
- 128k
- 256k

The following compression ratios (data-controlled) were used:

- 1:1
- 2:1

During each pass, test software will measure actual data transfer rate\*, while limiting maximum transfer rate\*\* to randomly selected rate from the following values:

- 1.5 MB/sec
- 5 MB/sec
- 30 MB/sec
- 80 MB/sec

The rate is changed at random every 5 seconds. All continuous transfer data rates are measured in 16-bit Wide Ultra SCSI-3, LVD mode.

\*Actual data transfer rate: Measured from the start of the command (Cmd) phase through the completion of the status phase (Bytes Read/Write)/(Read/Write time). All calculations are based on R/W commands that have completed. Read or Write time is calculated from the first completed Read/Write Cmd captured through the Cmd Complete message of the last Read/Write Cmd captured.

\*\*Maximum transfer rate - the program will limit the maximum amount of data transferred, in 1 second increments, to the limits specified by the case.

**Performance Testing Results**

**Performance (Transfer Rates)**


- 1:1 Compression Slow Host Backup (write) Results
- Quantum DLT-S4 performed up to 7% faster than its competitive LTO-3 products.
- 2:1 Compression Slow Host Backup (write) Results
- Quantum DLT-S4 performed up to 9% faster than the LTO-3 products tested.

Using a series of benchmark tests to compare to the performance of the DLT-S4 Digital Data Rate Agent against HP’s Data Rate Matching (DRM), Quantum Corporation’s DLT-S4 displayed a relative performance advantage over the HP Ultrium LTO-3 and the IBM Ultrium LTO-3 drives despite a lower speed specification design point.

DLT’s DDRA technology with the DLT-S4 drive proved to be at par or better than the hybrid digital/analog/mechanical speedmatching technology used by the LTO-3 drives tested when using slow or varying host data transfer rate conditions. Quantum’s DLT-S4 drive performed up to 4 % faster than LTO-3 over a wide range of block sizes (see below) in a native environment during a backup.


Note: The conclusions made are based on the test results obtained from the independent third party test facility, Percept Technologies Labs.

1:1 Compression Backup



Graph 2

2:1 Compression Backup (write)



Graph 3

### Summary

Incorporating a “design for reliability” approach with today’s newer component technologies, Quantum has virtually eliminated repositioning as a factor in reliability. This has been demonstrated with both reliability testing and empirical field return data.

The only real consideration for today’s drive and media technology when considering repositioning is the overall performance of the drive. Quantum has focused on improved performance while ensuring the drive and media reliability meets or exceeds specifications.

The digital/analog/mechanical technology solutions used by some LTO-3 products for data speed matching can at best only ‘match’ the performance of pure digital buffering solutions when used in the very best case host data transfers conditions. In less than optimal host conditions, the best performance is achieved using the large cache buffer and pure digital solution, as is the case for the Quantum DLT-S4 product.

### Notice

**This White Paper may contain proprietary information protected by copyright. Information in this White Paper is subject to change without notice and does not represent a commitment on the part of Quantum. Although using sources deemed to be reliable, Quantum assumes no liability for any inaccuracies that may be contained in this White Paper.**

**Quantum makes no commitment to update or keep current this information in this White Paper, and reserves the right to make changes to or discontinue this White Paper and/or products without notice.**

**No part of this document may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or information storage and retrieval systems, for any purpose other than the purchaser’s personal use, without the express written permission of Quantum.**

**This paper is not intended as legal advice on the regulations driving compliance. Quantum recommends that customers seek qualified expert advice to ensure certifiable compliance.**