

TECHNOLOGY BRIEF

January 1997

Compaq Computer Corporation

CONTENTS

Introduction	3
Evolution of Tape-Backup Technology	3
DLT Technology	4
Compaq DLT Drive	6
Achieving Published Performance	7
Hardware Compression.....	8
Hardware Reliability	8
Error Checking and Correction	8
Firmware	9
Cleaning DLT Drives	9
DLT Media	9
Tape Durability	9
Archival Life	10
DLT Manageability	10
Drive Status Alert.....	10
Soft-Error Threshold.....	10
DLT Arrays	10
The Future of DLT	11

DLT Hardware Technology

EXECUTIVE SUMMARY

Previously the evolution of backup hardware technology has not kept pace with the backup demands of 24-hour data availability and of the new superservers, global networks, and enterprises. As a result, network administrators are spending more time monitoring backups, and the overall cost of backups is increasing. The emerging Digital Linear Tape (DLT) technology provides the needed backup solution for reducing the overall cost of backup while protecting critical data through regular backups of large servers, networks, and enterprises in a process that maintains the absolute integrity of the data.

This technology brief summarizes the evolution of tape backup and explains how DLT hardware technology provides the level of backup performance, capacity, reliability, durability, and manageability required for large systems, networks, and enterprises. It describes the current Compaq DLT product and DLT media, introduces Compaq DLT Array products, and predicts the future of DLT technology.



NOTICE

The information in this publication is subject to change without notice.

COMPAQ COMPUTER CORPORATION SHALL NOT BE LIABLE FOR TECHNICAL OR EDITORIAL ERRORS OR OMISSIONS CONTAINED HEREIN, NOR FOR INCIDENTAL OR CONSEQUENTIAL DAMAGES RESULTING FROM THE FURNISHING, PERFORMANCE, OR USE OF THIS MATERIAL. THIS INFORMATION IS PROVIDED "AS IS." COMPAQ MAKES NO REPRESENTATION OR WARRANTIES OF ANY KIND, WHETHER EXPRESS, IMPLIED, OR STATUTORY, AND EXPRESSLY DISCLAIMS THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR PARTICULAR PURPOSE, GOOD TITLE AND NONINFRINGEMENT.

This publication does not constitute an endorsement of the product or products that were tested. The configuration or configurations tested or described may or may not be the only available solution. This test is not a determination of product quality or correctness, nor does it ensure compliance with any federal, state or local requirements. Compaq does not warrant products other than its own strictly as stated in Compaq product warranties.

Compaq, Contura, Deskpro, Fastart, Compaq Insight Manager, LTE, PageMarq, Systempro, Systempro/LT, ProLiant, TwinTray, ROMPaq, LicensePaq, QVision, SLT, ProLinea, SmartStart, NetFlex, DirectPlus, QuickFind, RemotePaq, BackPaq, TechPaq, SpeedPaq, QuickBack, PaqFax, Presario, SilentCool, CompaqCare (design), Aero, SmartStation, MiniStation, and PaqRap, registered United States Patent and Trademark Office.

Netelligent, Smart Uplink, Extended Repeater Architecture, Scalable Clock Architecture, Armada, Cruiser, Concerto, QuickChoice, ProSignia, Systempro/XL, Net1, LTE Elite, Vocalyst, PageMate, SoftPaq, FirstPaq, SolutionPaq, EasyPoint, EZ Help, MaxLight, MultiLock, QuickBlank, QuickLock, UltraView, Innovate logo, Wonder Tools logo in black/white and color, and Compaq PC Card Solution logo are trademarks and/or service marks of Compaq Computer Corporation.

Other product names mentioned herein may be trademarks and/or registered trademarks of their respective companies.

©1997 Compaq Computer Corporation. All rights reserved. Printed in the U.S.A.

Microsoft, Windows, Windows NT, Windows NT Advanced Server, SQL Server for Windows NT are trademarks and/or registered trademarks of Microsoft Corporation.

DLT Hardware Technology

First Edition (January 1997)

395A/0197

INTRODUCTION

The designers of backup systems are confronted with a growing challenge. Network storage capacity has been doubling every year.¹ File sizes continue to increase as applications become more complex. Users want information at their fingertips and are reluctant to delete files. The amount of data to be backed up is growing rapidly.

On the other hand, as enterprises and networks become global, the demand for around-the-clock server availability continuously increases. The time when user demand is at a minimum and backups can be scheduled (the backup window) is shrinking. The time during which *complete* backups can be done is being reduced, in some cases, to near zero.

In addition, companies are consolidating servers to reduce network and administrative costs. When one superserver does the work of several smaller servers, the importance of keeping that superserver on line escalates. Backup technology has not kept up with the success of these new superservers, and backup has become more difficult. As a result, network administrators are spending more time monitoring backups. Because of this, the overall cost of backups and especially the associated administrative costs are increasing.

Mission-critical data routinely resides on networked data servers, high-end workstations, and PCs. Network and enterprise management must be certain that critical data is protected through regular backups in a process that maintains absolute integrity of the data. This need has never been more critical than it is today. This comes as no surprise to the developers of tape-backup hardware and software. To one extent or another, they have been striving to meet this need for years.

This technology brief summarizes the evolution of tape backup and shows how the emerging Digital Linear Tape (DLT) technology provides the needed solution for large networks and enterprises. The current DLT drive is introduced, along with details about its performance, reliability, and other important characteristics. The DLT media is described from the point of view of its durability and archival life. The benefits that Compaq Insight Manager adds to DLT manageability are covered, and DLT arrays are introduced. Finally, this document includes a brief look into the future of DLT.

EVOLUTION OF TAPE-BACKUP TECHNOLOGY

Tape drives have been the most cost-effective method for backing up large amounts of data from personal computers for two decades. In the late 1980s, the two most popular backup-tape technologies were Quarter-Inch-Cartridge (QIC) and helical scan. The QIC technology moves magnetic tape past a stationary read/write head just as in today's audio tape players. The original QIC Minicartridge used 0.25-inch tape (hence the name QIC). Later, *QIC-Wide* tapes increased the tape width to 0.315 inch (8 mm). *Travan*² was the name given to an extended QIC tape format. Travan drives read and write both Travan and QIC-Wide minicartridges. The popular QIC-80 minicartridge can store 125 Mbytes of uncompressed data; the equivalent Travan cartridge, 400 Mbytes of uncompressed data using the QIC-80 format.

In the mid-1980s, helical-scan tape technology was adapted from home videotape technology to produce a high-capacity backup-tape technology that met the needs of midrange and smaller systems.

Helical-scan drives are equipped with error-correction circuitry and higher-quality tape mechanisms; they use a rotating read/write head assembly (see Figure 1). The tape is pulled out of the cassette and wrapped around the rotating heads, and data is recorded in diagonal stripes. This approach results in high capacity and relatively slow tape movement (by today's standards).

¹ The *Meta Group* predicts that storage capacity needs will continue to increase at a compounded annual growth rate of 60 percent.

² *Travan* is a registered trademark of the 3M Corporation.

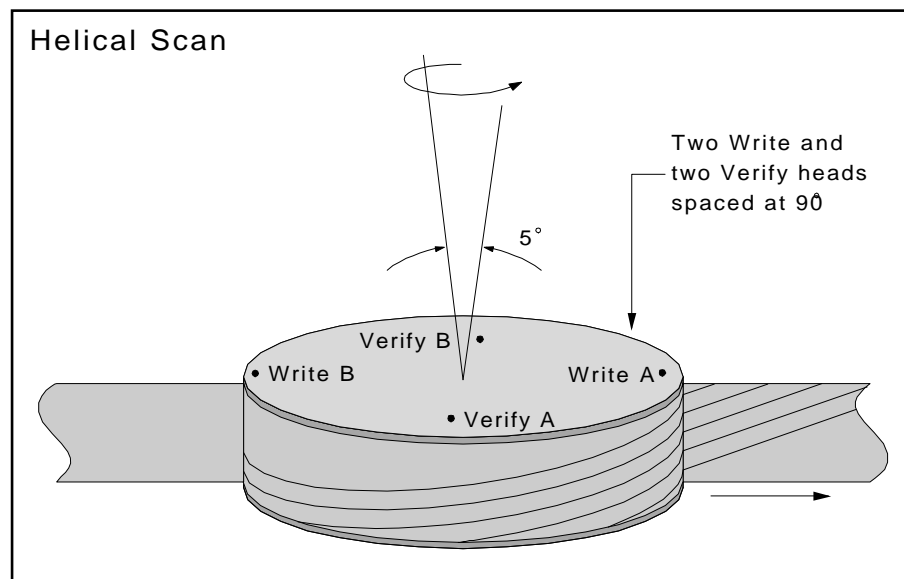


Figure 1. Helical-Scan Heads and Tape with Diagonal Tracks

Digital Audio Tape (DAT) and Exabyte 8-mm tape are two popular backup technologies that use helical scan. DAT was originally a CD-quality audio format that was adapted for computer use in 1988. DAT uses 4-mm tape cassettes with capacities from 2 to 12 Gbytes. Recently, DAT tape began replacing 8-mm tape in smaller systems with less demanding backup requirements. With the introduction of autoloaders, unattended backups became possible (provided the DAT heads were regularly cleaned). Unattended backups helped to reduce administrative costs. Today, DAT technology continues to be the appropriate choice for these smaller systems.

But networks continued to grow. Some enterprises became intercontinental. The storage capacity of servers and workstations grew rapidly. Individual superservers began taking the place of several high-capacity servers. Users placed around-the-clock demands on their systems, leaving less and less time for backup. Times had changed, and helical-scan technology could not meet the backup demands of larger networks and enterprises. Needing another solution, the industry turned to digital linear tape.

DLT TECHNOLOGY

The basic DLT technology had been around for more than a decade. It first appeared on the market as Digital Equipment Corporation's (DEC's) TK50 and TK70 drives. DLT appeared promising because of its inherent ability to deliver high performance, high capacity, and reliability. When coupled with the increasingly popular RAID technology, arrays of identical DLT drives offered a high degree of fault tolerance as well as high data accuracy. And when coupled with backup software able to do image backups (rather than file-by-file backups), DLT provided the backup solutions needed for today's large systems, networks, and enterprises.

DLT is faster than helical-scan technology (including DAT) because it records and reads multiple tracks simultaneously. As shown in Figure 2, DLT divides the tape into parallel, horizontal tracks and records data by moving the tape at high speeds³ past heads that remain stationary. As figure 2 illustrates, each set of heads consists of two write heads with a read/verify head between them;

³ Current DLT drives move the tape past the stationary heads at 100 to 150 inches per second during read/write operations and during faster search operations. As the technology advances, these tape speeds will increase.

and tape movement is bi-directional. During a write operation, the first write head of a set that encounters the tape writes new data to the track, overwriting existing data on that track. The middle head then reads the newly written data for verification purposes. (The DLT drive circuitry verifies the recorded data by comparing the data read from the tape with the data written to the tape.) The third head does nothing until the tape switches direction. When the tape reverses direction, the two DLT write heads switch roles. The same write-read-compare procedure is followed, as the tape moves in the opposite direction. Sets of DLT heads shift vertically to access other tracks.

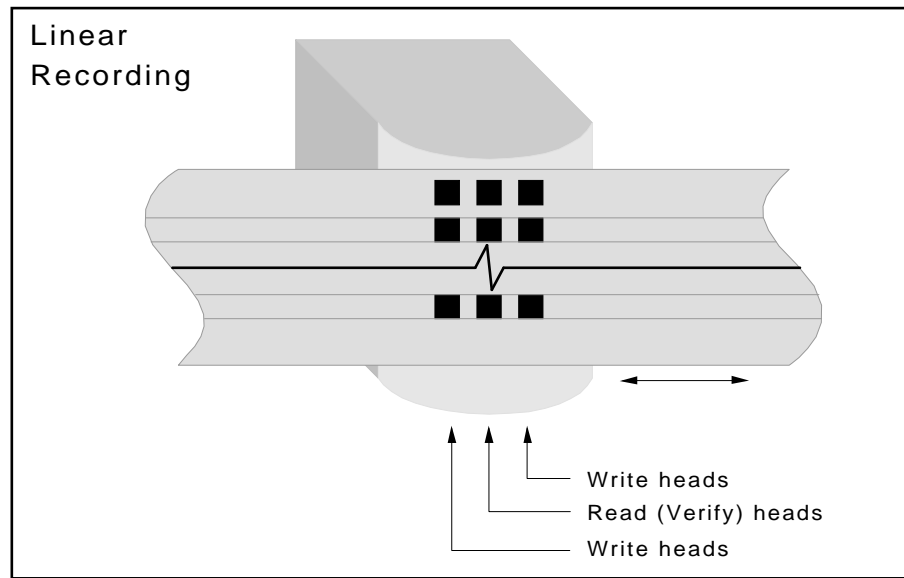


Figure 2. DLT Multiple Horizontal Tracks

In Figure 2, the heavier line with the engineering break in the center of the tape indicates that current DLT tape has more tracks than the number shown in the figure. As DLT technology continues to be developed, the number of tracks on the tape increases.

Because DLT tracks are horizontal, manufacturers can add more sets of read/write elements to the heads to increase data-transfer rates. Tape and head life are also increased because with DLT technology the tape is not pulled out of the cassette and wrapped around a rotating drum as it is with helical-scan technology.

DLT drives record tracks of data in a “serpentine” pattern (see Figure 3). If the drive is performing a read operation (reading multiple tracks) and reaches the end of the tape, it does not rewind the tape but continues the read operation by shifting the heads vertically to the next set of tracks and reversing the direction of the tape. File searches take place quickly using the direct track access (DTA) feature, even for image backups. For every backup, a DTA directory is built and recorded at the beginning of the tape. The directory includes the location of each file on the tape, including track numbers.

Figure 3 illustrates the layout of a typical DLT tape and shows the location of file ABC, which was backed up earlier. If the system requests the restoration of file ABC, the drive begins the search by looking up that file’s location in the DTA directory, reading the number of the track the file is on, and immediately shifting the heads to that track. It is necessary to search only through that one track to find file ABC. The average access time to find any backed-up file is currently

68 seconds. The maximum access time, from the Beginning Of Tape (BOT) mark to any file, is currently less than 90 seconds.

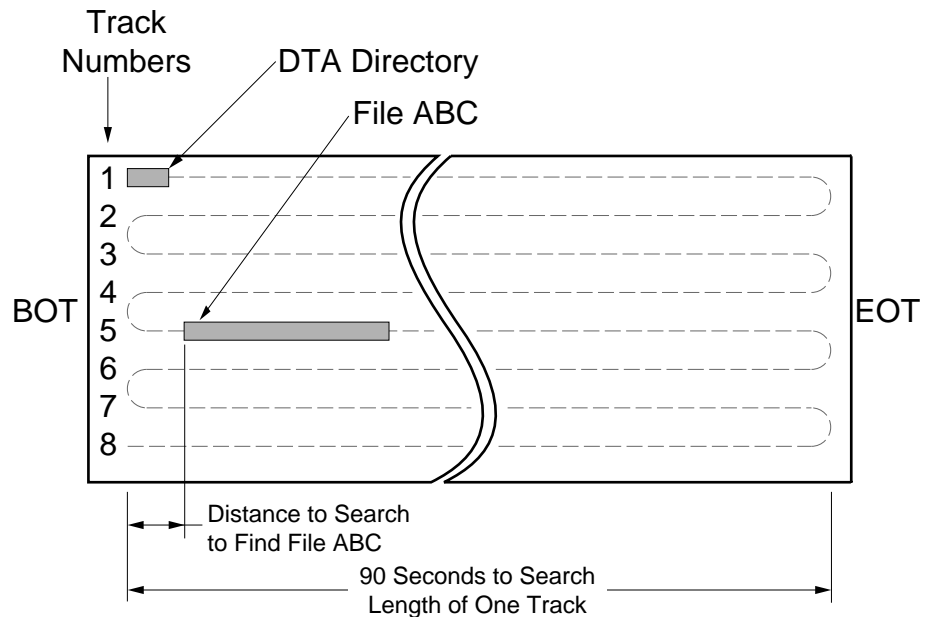


Figure 3. DLT Serpentine Track Layout

COMPAQ DLT DRIVE

Compaq currently offers the 15/30-Gbyte⁴ DLT drive in Compaq ProLiant and ProSignia servers. These drives offer a 50 percent increase in capacity over the earlier 10/20-Gbyte DLT drives. The 15/30-Gbyte drive specifications are listed in Table 1.

The *RLL* (Run Length Limited) encoding method specified for this drive is commonly used with magnetic tapes and disks. It is a modification of the *MFM* (Modified Frequency Translation) encoding method, which was used for the hard drives of the original PCs. When compared to *MFM*, *RLL* produces faster data-access times and increases the magnetic media's storage capacity.

The *run length* is the number of consecutive binary zeros before a binary one (1) is recorded. In *RLL 2,7* the sequences of binary zeros always comprise from 2 to 7 zeros. This requires fewer flux reversals (magnetic changes in the media) for a given amount of data, enabling more data to be placed on the tape or disk and bringing about a 50 percent increase in disk space over *MFM* encoding.

This remainder of this section describes the drive firmware, performance prerequisites, hardware compression, and error checking and correction of Compaq DLT drives. Also included is a description of head-cleaning requirements for DLT drives.

⁴ The first number (15) specifies the capacity of the DLT drive and cartridge to store *uncompressed* (native) data. The second number (30) is an *estimate* of the drive's capacity to store data that has been compressed 2:1. It has become a de facto industry standard to quote drive capacity using the native data rate (15) followed by an estimated data rate using 2:1 compression.

TABLE 1. SPECIFICATIONS FOR COMPAQ 15/30 DLT TAPE DRIVE

Recording Format	128-Track Serial Serpentine (variable block)
Number of physical tracks	128
Recording Density	62,500 bits per inch
Track Density	256 tracks per inch
Blocks per track	Variable
Blocks per frame	Variable
Bytes per block	Variable
Bytes per group	64,000
Data frame/group	16
Encoding Method	RLL 2,7

Achieving Published Performance

To achieve its published performance specifications, the Compaq DLT 15/30-Gbyte tape drive must have the recommended DLT tape cartridge installed and the tape must have been previously formatted. Table 2 lists the recommended cartridge and the two corresponding Compaq part numbers. Table 2 also includes capacity specifications for this cartridge and the *sustained transfer rate* that can be achieved by the tape drive with this cartridge. The *sustained transfer rate* is the rate at which the tape drive sends and receives data from the controller. This rate includes the time required for head switches and seeks and for system processing; it reflects the true performance of the tape drive.

For details on the durability and projected archival life of the recommended DLT tape, refer to the section “DLT Media” in this document.

TABLE 2. RECOMMENDED TAPE CARTRIDGE FOR COMPAQ DLT 15/30 DRIVE

Cartridge	DLT Tape IIIXT 242465-001 (single cartridge) 242466-001 (package of 7)	0.5-inch tape, 1800 feet long
Formatted Capacity	Native (uncompressed) With 2:1 Compression Average File-access Time	15 Gb (unattended operation) 30 Gb (unattended operation) 68 seconds
Sustained Transfer Rate	Native (uncompressed) With 2:1 Compression	1.25 Mb/s 2.5 Mb/s

Hardware Compression

The term *compression* refers to re-encoding digital data to take up less storage space on tape or disk. Digital data is compressed by finding repeating patterns of characters and recoding them. A string of 16 zeros, for example, would be replaced with code that indicates that 16 zeros follow a pointer. The more repetitive patterns that can be found in a file, the more that file can be compressed.

A compressed file that occupies 50 percent of the space it formerly occupied is said to be compressed 2:1. The actual space that compressed files will occupy can only be known for certain after the compression takes place. Generally, text can be compressed to about 2:1 and graphics files, to about 5:1. *Some files can hardly be compressed at all.* The amount of actual compression depends entirely on the type of file and the type of compression used.

The two most popular types of compression currently in use are Huffman coding and the Lempel-Ziv-Welch (LZW) compression algorithm. (This is the algorithm used by the Unix *compress* command.) Compaq DLT technology uses the DLZ (Digital-Lempel-Ziv) compression algorithm, which was developed especially for digital linear tape. DLZ is a high-efficiency variant of the LZW compression technique.

Hardware Reliability

The basic DLT design is inherently more reliable than those of DLT's counterparts from the older technologies. In DAT and other helical-scan drives, the tape is wound tightly around a drum and moves through a complex tape path, resulting in stress and abrasion to the tape and the heads, and in a need for frequent head cleaning. This serves to reduce the longevity of the tape, the heads, and the drive transport mechanism. The path that DLT tape takes through a DLT drive is less complex. As a result, the DLT tape operates at a lower constant tension, minimizing wear on the tape and read/write heads and significantly reducing the need for head cleaning. The head life of the Compaq 15/30 DLT drive is estimated by the manufacturer to be 30,000 hours at a 100 percent duty cycle (continuous operation).⁵ By comparison, typical helical-scan head life is 2,000 hours.

Error Checking And Correction

The current bit-error-rate specifications for DLT drives are:

- less than one unrecoverable hard error in 10^{17} bits of data, and
- less than one undetected soft error in 10^{30} bits of data.

This is achieved through a multilayered approach that interleaves the following during data recording:

- a 16-Kbyte Reed Solomon error-correction code (ECC) with every 64 Kbyte of data on tape,
- a 64-bit cyclical redundancy code (CRC) on each 4 Kbytes of data on tape,
- end-to-end 16-bit CRC on each record overlapped with parity from the SCSI bus, and
- internal parity checking on the cache buffer.

As an added reliability measure, Compaq DLT drives verify data by performing a read operation after each write command and by automatically re-recording data if a recording error is detected.

⁵ DLT reliability at high duty-cycle ratings generates a great deal of interest from Compaq customers. Compaq works closely with the DLT drive manufacturer in a continuing effort to improve quality and reliability at 100 percent duty cycle.

CRC refers to an error checking technique for ensuring the accuracy of transmitted digital data. The written data is divided into predetermined lengths which are divided by a fixed number. The resulting remainder is appended to and sent with the message. During the read operation, the remainder is recalculated. If it does not match the transmitted remainder, an error is detected.

Firmware

Compaq DLT-drive firmware contains diagnostic routines for troubleshooting, malfunction isolation, and recovery. These routines include a comprehensive off-line self-test, scratchpad memory, and error-reporting capabilities. The firmware recovery procedures include read retries, head repositioning, and tape tension adjustments. The firmware and the SCSI-2 interface parameters can be updated by tape.

Cleaning DLT Drives

DLT drives do not require frequent cleaning. When a DLT drive does need to be cleaned, the yellow **USE CLEANING TAPE** front-panel indicator lamp turns on. The system administrator then uses a special Compaq DLT cleaning tape.⁶ The cleaning process takes a little more than a minute. When the cleaning is finished, a beep sounds and a red indicator turns on to indicate that the cleaning tape can be removed. This is the only cleaning method approved for use with DLT drives.

DLT MEDIA

The DLT technology uses a high-grade metal-particle tape to achieve high densities. The current DLT tape is 0.5-inch MP2 (metal particle 2), and is 0.3-mil thick. The tape binder minimizes tape and head wear and resists the retention of airborne particles. The shock-resistant DLT cartridge measures 4.1 inch by 4.1 inch by 1 inch. DLT tape is very durable, and the cartridges have exceptional archival life.

Tape Durability

The life of the Compaq DLT 15/30 cartridge is rated by the drive manufacturer at an average of 500,000 passes. This amounts to an estimated average life of 10,000 cartridge uses. A *pass* is defined as a single movement of any given point on the tape past the head assembly. Usually, the number of passes is estimated to be the number of times the cartridge is written to or read from one end of the tape to the other. But this does not take into account the number of additional passes that occur because of the internal repositioning of the tape when an application does not “stream” the drive.

During streaming, the drive continuously writes information to the tape. The host sends the data quickly and steadily enough that the drive does not have to stop the tape to wait for more data. When it does become necessary to stop the tape, the drive then moves the tape backward to reposition the head adjacent to the last block written. This repositioning adds to the number of tape passes for that section of the tape. Later, when more data arrives, forward tape movement is resumed.

In laboratory tests conducted by the DLT drive manufacturer, DLT tape media have remained in good condition even after very high numbers of passes. DLT tape life can be shortened, not so much by the number of passes, but rather by use in an environment with excessive heat, humidity, or contamination levels. Currently, there is no suggested number of passes after which a DLT cartridge should be retired.

⁶ One Compaq DLT cleaning tape lasts for approximately 20 cleaning operations.

Archival Life

The archival shelf life of a Compaq DLT 15/30 tape cartridge is currently estimated by the DLT drive manufacturer at more than 20 years with less than 5 percent loss in demagnetization in a typical storage environment. This environment includes:

- a temperature between 68 and 82 degrees F (18 to 28 degrees C).
- non-condensing relative humidity between 40 and 60 percent
- cartridges stored in protective cases to keep out dust and contaminants.

DLT MANAGEABILITY

Compaq Insight Manager (supplied with Compaq ProSignia and ProLiant servers) makes it easy for administrators to manage their DLT resources because it automatically notifies administrators of key events. Two of the latest automatic-notification features available with Compaq Insight Manager are described below.

Automatic Notification: Drive Status Alert

Compaq Insight Manager issues an automatic alert if there is a tape-drive error, if the media fails, if a tape drive itself fails, or if the status of a tape drive changes in some other way. For example, if a tape drive is turned off while waiting to execute a backup, Compaq Insight Manager immediately alerts the administrator while there is still time to complete the backup.

Automatic Notification: Soft-Error Threshold Exceeded

When a tape drive tries to write to a bad block, the write operation fails. The drive then attempts to rewrite the block at a different tape location, repeating the write attempt up to 16 times. It stops only when it finds a good spot on the tape or after the 16th failed attempt. Typically, as any tape drive ages or the heads become contaminated, the number of these rewriting attempts (soft errors) increases. Compaq Insight Manager enables the administrator to set a soft-error-notification threshold. When this threshold is exceeded, the administrator is automatically alerted to the developing potential problem and can respond with preventive maintenance before a hard failure occurs. The soft-error-notification threshold can be set to any level. Typical settings are zero for mission-critical backups, half the normal level (usually 8) for notification when a tape is starting to fail, etc.

DLT ARRAYS

Compaq DLT Array products consist of two or more DLT tape drives driven by Cheyenne Software's JETserve 3.3 or ARCserve 6.0 software. These backup-software products enable fast image backups to be combined with RAID fault tolerance. (Currently, JETserve and ARCserve are the only software products on the market to offer these capabilities.)

Compaq DLT Arrays have set a new standard in backup performance. Compaq conducted extensive backup tests using two 4-drive DLT Arrays and four SMART-2 Array Controllers,⁷ one controller for every two hard drives for maximum performance. The drives were the Compaq 15/30-Gbyte model, and the software was JETserve configured for RAID 5 fault tolerance. The performance testing consistently demonstrated a backup throughput of 46 Gbytes per hour. At

⁷ Compaq DLT Arrays and the SMART-2 Array Controller are not compatible with each other. However, the highest backup throughput was achieved when the *hard disks* were configured with the Compaq SMART-2 Array Controller. This controller increases network data availability and enhances storage management capabilities. Its *Online Capacity Expansion* feature supports hot-pluggable drives and enables network administrators to expand storage without bringing the system down.

this rate and with this equipment, a network administrator can back up 210 Gbytes in less than 5 hours. If one of the tapes in a RAID 5 backup set is lost or damaged, all the backed-up data can be restored using the remaining tapes in the set.

Compaq DLT Arrays can attach to a Fast SCSI-2 Controller, a Fast/Wide SCSI-2 Controller, or a Wide Ultra SCSI Controller, but not to the Compaq SMART-2 Array Controller. In addition, Compaq's DLT Arrays can be used with other backup software packages. However, those currently on the market will treat a 4-drive array as four individual drives (rather than as one or more logical drives) and can only do file-by-file backups.

THE FUTURE OF DLT

For today's larger systems, networks, and enterprises, DLT is the Compaq technology of choice to perform reliable backups in relatively small backup windows. Even so, DLT is still a young technology. It offers room for development and growth in performance and capacity to support faster networks, systems, and expanded applications. It is conceivable that future DLT drives will be capable of achieving more than 200 Gbytes of storage capacity and transfer rates of 20 Mbytes or more per second.

Current DLT cartridges have three to six times the recording area of helical-scan cartridges, but still have room for considerable growth in capacity. The use of thinner tape will enable more tape to be packaged in a cartridge. Higher bit density will be achieved through the use of new read/write head technologies such as thin-film and MR (magnetoresistive) heads. This will increase capacity and transfer rates. In addition, higher data transfer rates will be achieved by increasing the number of tracks recorded simultaneously on DLT tape.

Nevertheless, future requirements for *large-capacity unattended* backups are expected to quickly exceed the capacity of single DLT cartridges. Having an operator standing by to change DLT cartridges is not an option. This leaves only two possibilities: (1) writing to multiple cartridges in multiple drives, or (2) finding a DLT library (autoloader) with reliability high enough to match the reliability of the DLT drives. A DLT library with this high level of reliability would enable unattended, high-capacity backups to take place over longer periods of time. (Multiple-drive DLT arrays would still be necessary for very high capacity backups in smaller backup windows.) Compaq is currently conducting extensive laboratory tests as part of the search for such a library.

DLT will play a key role in emerging technologies such as Hierarchical Storage Manager (HSM). Just as RAID combines several DLT physical drives into one logical unit, HSM combines multiple storage media (magnetic disk, optical disk, magnetic tape) into individual logical units. HSM then migrates data between media based on how often the data is accessed. DLT's fast performance, accuracy, and durability will make it a key part of- future HSM applications.