

TECHNOLOGY BRIEF

September 1997

Compaq Computer Corporation

CONTENTS

Introduction.....	3
Test 1: Tape Drives in Relation to SCSI Cards — Windows NT	7
Test 2: Tape Drives in Relation to SCSI Cards — Novell IntranetWare 151 GB/hr	9
Test 3: 118 GB/hr Performance — Novell IntranetWare	11
Test 4: Local Backup of Windows NT to Various Numbers of DLT Drives	13
Test 5: Use of Concurrent Jobs to Improve Performance — Windows NT — 165 GB/hr	16
Test 6: Compressibility and 35/70 DLT to 15/30 DLT Performance	19
Test 7: RAIT-0 Versus RAIT-5 — Windows NT	20
Test 8: Network Performance — Simple Backup of a Remote Client — Windows NT	22
Test 9: Network Performance — Windows NT Server Pulling from an IntranetWare Remote Client	24
Test 10: Fibre Channel-Arbitrated Loop (FC/AL) Loop-Based Backup of 800 GB — Windows NT	26
Test 11: Feed Speed Relationships	29
Conclusion	29
Appendix 1: Compaq 35/70 GB DLT Drive Notes	29
Appendix 2: Compaq DLT Tape Array II Notes	29
Appendix 3: Standard High-Performance/High-Speed Backup Test Systems	29

Compaq DLT Tape Array II: High-Performance Backup of Enterprise Servers

EXECUTIVE SUMMARY

Since its introduction in February, 1996, the Compaq DLT Tape Array (which uses Compaq 15/30 GB DLT drives) has been an industry leader in backup performance capabilities. The DLT Tape Array II, which uses 35/70 GB DLT drives, sets new standards in high-speed backup performance.

Early in 1997, Compaq performed a series of tests to examine the performance of the Compaq 35/70 GB DLT drive and the DLT Tape Array II. The objective of these tests was to find the highest-performance configurations in a variety of backup scenarios on Compaq servers running Windows NT and Novell IntranetWare network operating systems.

The Compaq DLT Tape Array II provides high-capacity, high-speed tape storage for use in local and remote backup strategies. It uses the Compaq 35/70 GB DLT drive, which is engineered for superior performance. This brief addresses the technical aspects of high-performance backups that can reach and exceed 165 gigabytes per hour (GB/hr) with off-the-shelf Compaq hardware.




Figure 1: Maximum results in each backup scenario tested.

COMPAQ

Please direct comments regarding this communication to the ECG Technology Communications Group at this Internet address: TechCom@compaq.com

NOTICE

The information in this publication is subject to change without notice and is provided "AS IS" WITHOUT WARRANTY OF ANY KIND. THE ENTIRE RISK ARISING OUT OF THE USE OF THIS INFORMATION REMAINS WITH RECIPIENT. IN NO EVENT SHALL COMPAQ BE LIABLE FOR ANY DIRECT, CONSEQUENTIAL, INCIDENTAL, SPECIAL, PUNITIVE OR OTHER DAMAGES WHATSOEVER (INCLUDING WITHOUT LIMITATION, DAMAGES FOR LOSS OF BUSINESS PROFITS, BUSINESS INTERRUPTION OR LOSS OF BUSINESS INFORMATION), EVEN IF COMPAQ HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

The limited warranties for Compaq products are exclusively set forth in the documentation accompanying such products. Nothing herein should be construed as constituting a further or additional warranty.

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 and ProLiant are registered with the United States Patent and Trademark Office.

Netelligent is a trademark and/or service mark of Compaq Computer Corporation.

Microsoft, Windows, and Windows NT are trademarks and/or registered trademarks of Microsoft Corporation.

NetWare and Novell are registered trademarks and IntranetWare, is a trademark of Novell, Inc.

Pentium is a registered trademark of Intel Corporation.

DLT is a trademark applied for by Quantum.

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.

Compaq DLT Tape Array II: High-Performance Backup of Enterprise Servers

Second Edition (September 1997)
ECG075.0997

INTRODUCTION: THE NEED FOR SPEED

As companies migrate to standardized server solutions, there are significant changes in backup/recovery paradigms. The five-days-a-week, nine-hours-a-day operation has expanded to seven days a week and 24 hours a day. Backups must be performed concurrently (live) or during very small scheduled outage windows. In either case, the backup must be performed in a minimum amount of time: Speed is essential.

With the size of data sets and databases approaching and exceeding the terabyte mark, customers require improved performance for large backups. The chart below illustrates typical peak server backup speeds for the last five years. Today, Compaq passes 165 gigabytes per hour (GB/hr) with off-the-shelf technology.

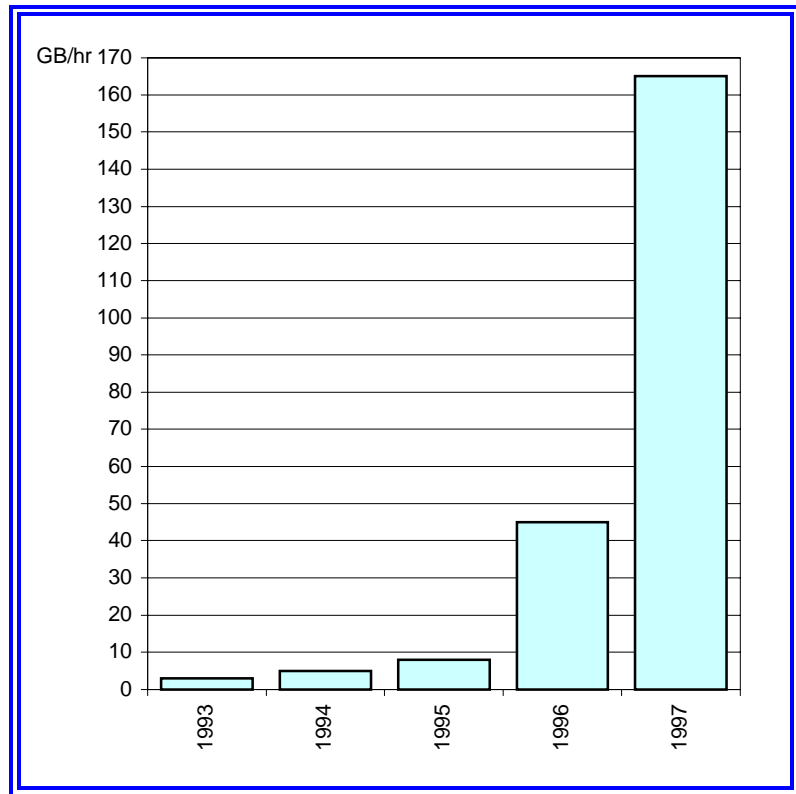


Figure 2: Backup speed improvement in the last five years. The 1997 figure represents results achieved by Compaq.

RESULTS: 165 GB/HR BACKUPS

Testing revealed three important factors in a high-speed backup solution:

- The base hardware capabilities of the system, centering on the PCI bus architecture.
- The speed at which the disk drive storage system can supply data to the tape system (referred to as *feed speed*).
- The tape technology selected to perform the backup.

In the tests, the system that provided optimum performance consisted of a Compaq ProLiant 6000 with four or eight drives in a DLT Tape Array II, which was fed data from a pair of Compaq SMART-2/P Array Controllers and associated SCSI disk drives. This configuration worked very well for Novell IntranetWare and Microsoft Windows NT network operating systems.

Test results show the DLT Tape Array II can exceed backup speeds of 150 GB/hr in multiple operating system environments.

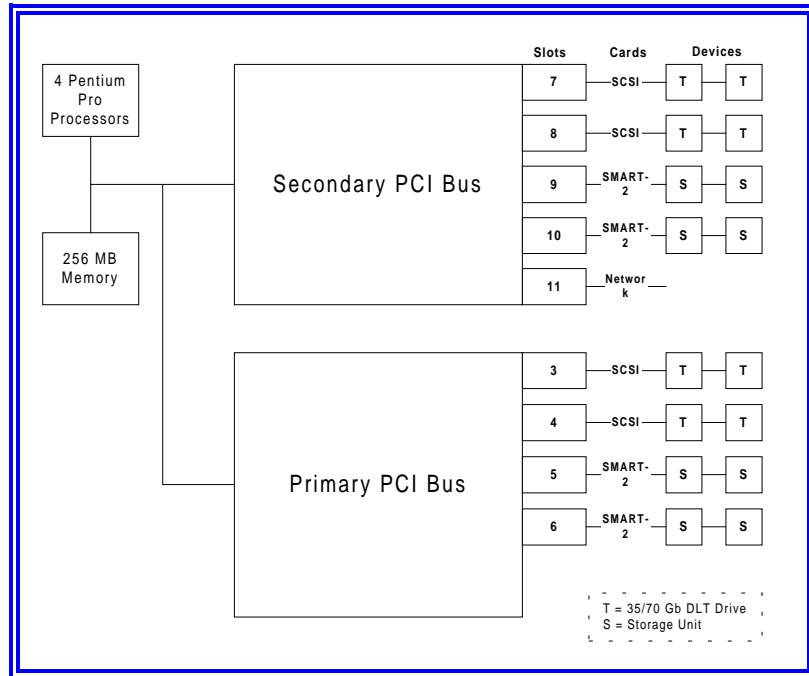


Figure 3: High-speed backup off-the-shelf-solution: ProLiant 6000 test system.

HIGH-SPEED BACKUP RESEARCH

Single and multi-processor systems were used to identify hardware factors that affect backup performance, including processors, memory, SCSI, and PCI channel bottlenecks. Some tests were run in a network environment; however, the focus was on local backups, because network bottlenecks significantly reduce feed speed. Future testing will address this problem and explore other ways to improve network-based backups.

Listed below are test machine configuration details, along with factors that were taken into consideration in developing the highest performance configurations.

Processor Quantity

The average backup uses from 6 to 25 percent of a single 200 MHz Pentium Pro (256K cache version) processor's power. The lower number is typical when using a single disk or tape drive; usage percentage grows substantially when striping data across a redundant array of independent tape drives (RAIT) in a RAIT-5 configuration. All testing was performed on Compaq ProLiant 5000 and ProLiant 6000 systems so that up to four 200 MHz Pentium Pro processors could be used. An assumption was made that most large systems would have at least two processors, and the majority of tests were performed with that number.

Memory

A system with 256 MB of RAM (regardless of the number of processors) was the standard for all tests, to eliminate the remote possibility of complications due to inadequate RAM.

PCI Bus

Although some testing was done with EISA bus-based systems, the higher throughput of PCI (up to three times faster) clinched the decision to use it for high-speed backup testing. The 32-bit x 33 MHz PCI bus peaked at 454 GB/hr in load tests with six Compaq SMART-2/P Array Controllers. To maximize bandwidth, tests were also performed to determine optimum card placement in the dual PCI bus system of the ProLiant 5000 and ProLiant 6000. The dual peer-to-peer PCI buses connect to a common 1,919-GB/hr, or 540 megabytes per second (MB/s) Gunning Transceiver Logic (GTL+) processor bus (the host bus), which permits a path to the processor that is twice the width of buses in many servers on the market today (Figure 4).




Figure 4: ProLiant 5000 peer-to-peer bus architecture. The ProLiant 5000 has seven PCI slots and three EISA slots across the two PCI buses (eight total; two are shared). The ProLiant 6000 has eleven slots: nine PCI and two EISA.

The dual peer-to-peer buses provide throughput that is superior to that of single PCI-buses typical of smaller servers. Many dual buses are stacked buses. They use a PCI-to-PCI-bridge, which feeds one PCI bus to the other and then goes through a host-to-PCI-bridge to get to the

host bus (Figure 5). Although this permits the use of additional cards by providing more slots, it does not increase performance. In contrast, the additional PCI card space in the ProLiant 5000 and ProLiant 6000 permits use of multiple SMART-2/P Array Controllers, multiple SCSI cards for tape drives, and multiple network interface cards. All of these devices run at peak PCI speeds, coexisting with minimal bottlenecks in performance. Although the PCI bus is rated at 133 MB/second, that is a burst data transfer speed. Actual measurements are typically from 60 MB/second to a maximum sustained peak of about 100 MB/second.

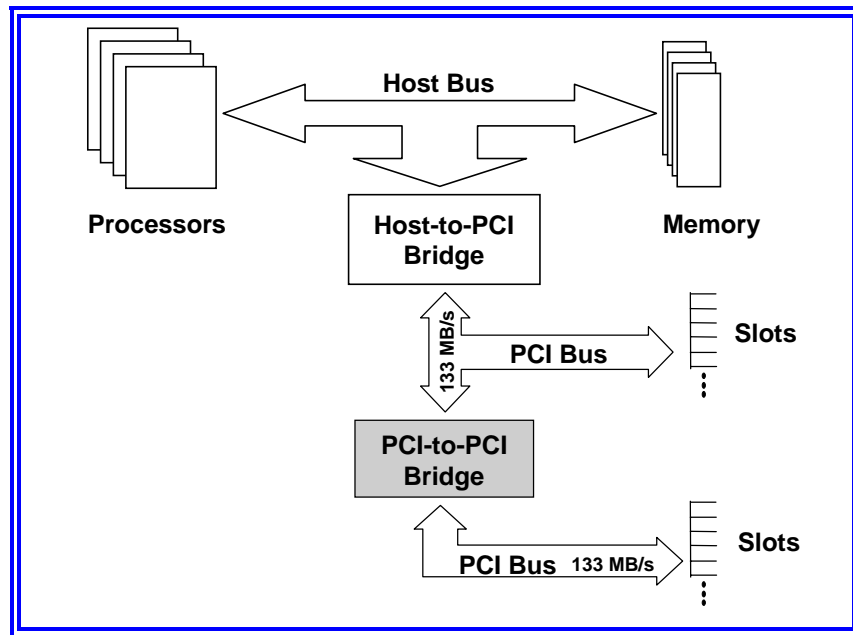


Figure 5: PCI-to-PCI bridge architecture.

Feed Speed

The disk drive subsystem is a critical component in high-speed backups. There is a direct relationship between the number of disk drives in the array and the speed at which the array can provide data to the processor (more drives allow greater speed). There is also a direct relationship between the number of RAID controllers and performance (more is better). For example, a server using a single SCSI card to access hard disks is limited to 10, 20 or 40 MB/s, depending on the type of SCSI cards and drives in use. On the other hand, an array controller can often deliver the maximum limit of performance that the PCI bus can deliver: 133 MB/s.

The Compaq SMART-2/P Array Controller is the best choice. It has two channels, and a minimum of four drives per channel should be used for high performance. If only six drives are available, the best performance will be achieved by assigning three drives to each channel. Adding a fourth drive to each channel improves performance significantly. Adding more than four drives per channel makes only a minor improvement in throughput, although it does add capacity.

Tests were performed with an assortment of narrow and wide drives ranging in size from 2 to 9 gigabytes (GB). Superior performance was seen on the Compaq Fast-Wide 2 GB, 4 GB and 9 GB drives. See Test 12 for more information.

Data Set

To simulate the types of file structures and data mixtures customers typically want to back up, mixes of data types included large files, large numbers of files, and large numbers of directory structures.

The size of a data set can affect backup speed. With a data sample of less than 1 GB, performance is sometimes affected by the tape drive size, as well as the tape drive cache and

disk cache. For example, a 32-MB file might appear to have been backed up in less than 1 second, simply because it went to the 8-MB cache in each of the four tape drives of the array. With data sets from 1 GB to 1 terabyte, performance was found to remain constant. The minimum test data set size for this project was 2 GB.

Compression

Hardware data compression is built into the 35/70 GB DLT drive. The default setting is to compress data. Compression has a direct impact on backup performance. Tests were performed with 2:1 and 4:1 compressible data sets, to simulate the most common levels of compressibility in actual working environments. Highly random sets of data may be less compressible, while large empty databases may be compressible to much higher levels. Determine how compressible the data sets are by testing them, usually by observing the speed of backup that the backup program reports. (The backup programs used for these tests [Cheyenne ARCserve and JetServe] report the compressibility of the data set being backed up. This number is reported to the backup application by the drive.) See Test 6 for more information.

Backup Types

Two types of backup techniques are generally available: *image* and *by file*.

In image backups, the backup program looks at the disk drive or file system as a single unit and takes a snapshot of it, backing it up in one session. Image is usually the fastest method of backup and restore for an entire file system; however, some image backups can only be restored as an entire unit, meaning the entire backup must be restored to obtain a single file.

By file backups are made on a file-by-file basis, retaining the information needed to retrieve a single file from the backup tape and restore it to the correct location on the system. The by file method usually takes much longer than the image method, because the full path/directory information is saved with each file. Nevertheless, it permits random access to the files for faster retrieval of individual files.

Both methods were included in Compaq's testing scenarios. It should be noted that some image backup programs (notably Cheyenne ARCserve for NT) have the ability do a by file restore from an image backup, but using that functionality usually extends the restore operation time significantly.

Test Sites

All tests except test 10 were performed on systems located in the Compaq Secondary Storage Systems Engineering lab. Test 10 was performed on the Max-Storage system in the Compaq Maximum Configuration Lab. All network loads were performed on lab networks.

**TEST 1 TAPE DRIVES IN RELATION TO SCSI CARDS —
WINDOWS NT**

Operating System	Windows NT 4.0
Backup Software	Cheyenne ARCserve 6.0 for Windows NT
Backup Method	Image and by file
Data Set	2 GB
Compression	4:1
Tape Format	RAIT-0
Hardware	ProLiant 5000, two Pentium Pro 200 processors, 256K cache, 256 MB RAM 80-GB array (two SMART-2/P Array Controllers with twenty 4-GB drives) 35/70 GB DLT drive(s) Wide-Ultra SCSI-3 cards (All Compaq off-the-shelf products)

Purpose

To observe the performance of various numbers of tape drives in relation to various numbers of SCSI cards under Windows NT 4.0.

Results

Of the configurations tested, two SCSI cards driving two 35/70 GB DLT drives offer the fastest speed. Although a typical configuration (two SCSI cards driving four 35/70 GB DLT drives) tested two percent slower, it doubles the capacity of the backup, and is therefore recommended to avoid tape changes during the backup of large data sets. Tape change delays might easily offset any performance advantage of using the smaller two-drive configuration.

There is no benefit to using a separate SCSI card for each drive when using the Wide-Ultra SCSI Controller (part number 272514-001). The bandwidth of this card can easily drive two 35/70 GB DLT drives. On a system with only Fast-Wide SCSI Controllers (part number 199533-001), splitting the four drives among four cards will help somewhat. It is also evident that image backups are considerably faster in a multi-drive configuration.

The ProLiant 5000 uses the slower Fast/Wide technology on its single embedded SCSI controller. It is recommended that this be used only for slower devices, not for multiple tape drives. The ProLiant 6000 has two embedded Wide-Ultra SCSI Controllers, of which the B-port is available for use with externally mounted tape devices. The A port can also be used; however, an optional cable is required.




Figure 6: Performance results from tests comparing various proportions of 53c875 SCSI cards to DLT drives (shown at the bases of the columns).

Number of SCSI Cards	Number of 35/70 GB DLT Drives	Backup GB/hr	Restore GB/hr
<i>Image Backup/Restore</i>			
1	1	34.38	21.89
1	2	51.91	31.09
2	2	56.54	56.54
2	4	51.91	31.15
4	4	49.73	30.21
<i>Backup/Restore by File</i>			
1	1	34.95	31.79
1	2	45.51	47.89
2	2	44.96	59.00
2	4	44.96	57.62
4	4	43.13	58.53

Table 1: Test data for tape drives versus SCSI card test — Windows NT.

**TEST 2 TAPE DRIVES IN RELATION TO SCSI CARDS —
NOVELL INTRANETWARE — 151 GB/HR BACKUP**

Operating System	Novell IntranetWare 4.11
Backup Software	Cheyenne JETserve 4.0 Plus
Backup Method	Image
Data Set	8 GB
Compression	4:1
Tape Format	RAIT-0
Hardware	ProLiant 6000, three Pentium Pro 200 processors, 512 K cache, 384 MB RAM 224-GB array (four SMART-2/P Array Controllers with fifty-six 4-GB drives) 35/70 GB DLT drive(s) Wide-Ultra SCSI-3 cards (All Compaq off-the-shelf products)

Purpose

To observe the performance of various numbers of tape drives in relation to various numbers of SCSI cards under Novell IntranetWare 4.11 on a very high-performance backup server configuration.

Results

In a typical DLT Tape Array II solution, two SCSI cards driving four 35/70 GB DLT drives offer a reasonable backup speed (121 GB/hr). A configuration of four SCSI cards and eight 35/70 GB DLT drives provides very high speed (151.58 GB/hr) and large capacity (more than 560 GB on a single set of mounted tapes). At 2:1 compression, storage capacity reaches 1.12 terabytes.

Restore speeds varied, but were generally about 60 GB/hr.




Figure 7: Performance of IntranetWare backups and restores by numbers of SCSI cards in relation to numbers of tape drives (shown at the bases of the columns).

Number of SCSI Cards	Number of DLT Drives	Backup GB/hr	Restore GB/hr
1	1	31.61	31.32
1	2	61.94	53.38
2	3	92.61	79.12
2	4	121.78	58.42
3	5	126.59	59.69
3	6	144.36	60.57
4	7	146.57	61.28
4	8	151.58	56.52

Table 2: Test data for numbers of tape drives in relation to numbers of SCSI cards under Novell IntranetWare.

The relationship between backup and restore speeds is inconsistent for a variety of reasons. Restore speed varies with the compression of the files. Compression and decompression times vary with file content and have different effects on CPU use.

Another influence on restore speed is the rate at which the target can receive and write data. The use of large, multi-megabyte caches on disk drive arrays occasionally allows restore speeds to exceed backup speeds.

**TEST 3 118 GB/HR PERFORMANCE TEST — NOVELL
INTRANETWARE**

Operating System	Novell IntranetWare 4.11 Build 2884
Backup Software	Cheyenne JETserve 4.0 Plus
Backup Method	Image
Data Set	2 GB
Compression	2:1, 4:1
Agents	None
Tape Format	RAIT-0
Hardware	ProLiant 5000, two Pentium Pro 200 processors, 256K cache, 256 MB RAM 80-GB array (two SMART-2/P Array Controllers with twenty 4-GB drives) 35/70 GB DLT drive(s) Wide-Ultra SCSI-3 cards (All Compaq off-the-shelf products)

Purpose

The initial purpose of this test was to find the fastest backup solution using Cheyenne JETserve and to examine performance variables. Long after this set of tests was performed, it was discovered that the addition of two SMART-2/P Array controllers and disk drives dramatically improved performance. Test 2 demonstrates this peak performance system.

Results

Of the configurations tested, a typical DLT Tape Array II solution with two SCSI cards driving four 35/70 GB DLT tape drives offers the fastest backup speed (118 GB/hr). The four card/eight drive configuration provides high speed (89 GB/hr) and very large capacity (more than 560 GB at 2:1 compression and more than 1 terabyte at 4:1) on a single set of mounted tapes. Numerous tests showed the two card/four drive setup to be the highest performer. The 2:1 versus 4:1 test results prompted other tests to measure the compressibility factor (see Test 6, page 19).

This set of tests illustrates the levels of performance that might be seen in this configuration. To achieve the highest speeds possible, a very fast source hard disk drive storage system is required. Other factors, such as CPU speed and bus speeds, also influence backup speed.




Figure 8: High-performance backup and restore under Novell IntranetWare with different proportions of SCSI cards to tape drives (shown at the bases of the columns).

*Hardware compression turned off.

**This test used 4:1 compression; all others, 2:1 compression.

Compression Ratio	SCSI Cards	Tape Drives	Backup GB/hr	Restore GB/hr
2:1	1	1	29.03	29.49
2:1	1	2	55.81	39.62
2:1	2	2	59.70	52.17
2:1	2	3	81.63	59.80
2:1	2	4	104.35	79.22
2:1 (but drive compression was turned off)	2	4	66.67	79.12
4:1	2	4	118.98	79.55
2:1	3	3	79.65	69.10
2:1	3	6	90.45	51.43
2:1	4	4	94.99	72.00
2:1	4	8	89.55	57.05

Table 3: Test data for performance test of backups under Novell IntranetWare.

TEST 4 LOCAL BACKUP OF WINDOWS NT TO VARIOUS NUMBERS OF DLT DRIVES

Operating System	Windows NT 4.0
Backup Software	Cheyenne ARCserve 6.0 for Windows NT
Backup Method	Image
Data Set	2 GB
Compression	2:1
Tape Format	RAIT-0
Hardware	ProLiant 5000, two Pentium Pro 200 processors, 256K cache, 256 MB RAM 40-GB array (one SMART-2/P Array Controller with ten 4-GB drives) 35/70 GB DLT drive(s) Wide-Ultra SCSI-3 cards (All Compaq off-the-shelf products)

Purpose

To determine the highest performance configuration for an image backup to DLT drive or RAIT-0 group of drives (RAIT-0 is not significant when using one drive). The test consisted of running a single job to back up a 2 GB file set to various numbers of DLT drives (from one to eight).

Results

Results show a limit of about 40 GB/hr while running a single job. Other tests with faster feed sources show an upper limit of 57 GB/hr for a single job writing to eight tape drives. This limit is based on the limit to the number of threads that a backup program can run under Windows NT. Additional tests will be run to show that if multiple jobs are run simultaneously, there is still available bandwidth from the disk drives, the CPU, the operating system, and the tape drives to permit a much higher aggregate speed.

In this configuration, the restore speed exceeded the backup speed in some cases, apparently because the SMART-2/P Array Controllers cache much of the data, making the restore data flow more consistently to the target.




Figure 9: Local backup and restore performance under Windows NT with various proportions of SCSI cards and 35/70 GB DLT drives (shown at bases of columns).

Cards	Drives	Backup GB/hr	Restore GB/hr
1	1	30.1	14.2
1	2	41.1	36.3
2	2	25.7	33.6
2	4	37.0	38.4
4	4	34.8	44.3
4	8	38.5	38.7

Table 4: Test data for local image backups to various numbers of DLT drives under Windows NT.

The following test results are included to demonstrate the performance increase achieved when using even larger systems with faster feed speeds. These tests differ only in that a ProLiant 6000 was used, with three Pentium Pro processors and two additional SMART-2/P Controllers with 56 drives.




Figure 10: Backup speed results on the ProLiant 6000 test under Windows NT, with various proportions of SCSI cards and 35/70 GB DLT drives (shown at bases of columns).

Jobs	Cards	Drives	2:1 Data Set Backup (GB/hr)	4:1 Data Set Backup (GB/hr)
1	1	1	29.88	32.17
1	1	2	52.70	55.17
1	2	2	58.92	67.38
1	2	4	76.54	86.50
1	4	4	78.93	84.50
1	3	6	84.96	86.23
1	4	6	84.51	85.02
1	4	8	83.56	82.48

Table 5: Test data for local image backups to various numbers of DLT drives under Windows NT using the ProLiant 6000.

TEST 5 USE OF CONCURRENT JOBS TO IMPROVE PERFORMANCE — WINDOWS NT — 165 GB/HR

Operating System	Windows NT 4.0
Backup Software	Cheyenne ARCserve 6.0 for Windows NT
Backup Method	Image — multiple jobs
Data Set	2 GB
Compression	2:1, 4:1
Tape Format	RAIT-0
Hardware	ProLiant 6000, three Pentium Pro 200 processors, 512 K cache, 384 MB RAM 224 GB array (four SMART-2/P Array Controllers with fifty-six 4-GB drives) 35/70 GB DLT drive(s) Wide-Ultra SCSI-3 cards (All Compaq off-the-shelf products)

Purpose

To determine whether throughput can be improved by performing concurrent backups on multiple drives.

Test Setup

To show what effect multiple tape drives would have, a large variety of situations were tested. In all cases, the Compaq Array Configuration Utility was used to configure the Compaq disk drive arrays. To join multiple disk array groups between adapter groupings, Windows NT Disk Administrator was used to group these as one striped drive.

Results

These tests demonstrated performance limits during image backups under Windows NT. The bottleneck has to do with the way an image is processed. To get around the bottleneck, concurrent backups of multiple drives were run. This technique works very well. It helps obtain higher overall backup speeds, but its use depends on the drive/data structure. It is of no advantage in an environment with a single large database, but it is of great value in many applications server or file server backups.

The following graph demonstrates one, two, four and eight jobs running under Windows NT 4.0, backing up in each case to a set of eight tape drives. The graph shows the high performance capability under Windows NT; however, the raw data below the chart shows that performance with a four-drive array is still very impressive.




Figure 11: Concurrent backup performance running Cheyenne ARCserve 6.0 (image option) under Windows NT.

Jobs	SCSI Cards	DLTs	RAIT Level	Smart Cards per Job	Drives per Job	2:1 Backup (GB/Hr)	4:1 Backup (GB/Hr)
1	4	8	0	4	56	83.56	82.48
2	1	2	No RAIT	2	28	54.02	57.71
2	2	2	No RAIT	2	28	60.05	68.79
2	2	4	0	2	28	97.03	107.11
2	4	4	0	2	28	118.36	136.52
2	4	8	0	2	28	135.64	137.15
4	2	4	No RAIT	1	14	107.71	112.73
4	4	4	No RAIT	1	14	121.88	135.04
4	4	8	0	1	14	154.27	147.14
8	4	8	No RAIT	1 channel	7	165.26	164.18

Table 6: Test data for concurrent backup tests under Windows NT.

- Jobs: Number of Cheyenne backup jobs started concurrently under Windows NT.
- SCSI Cards: Number of Wide-Ultra SCSI-3 cards used to connect to the 35/70 GB DLT drives.
- DLT: Number of tape drives used.
- RAIT Level: When a job used more than one drive, the drives were configured as RAIT 0.
- SMART Cards – Number of SMART-2/P Array Controllers used per job (to help determine loading factors).
- Drives per Job: Number of hard disk drives per job (to give perspective on feed source to each tape job running).
- Note: In all testing, care was taken to balance the loading of the two PCI buses.

TEST 6 COMPRESSIBILITY AND 35/70 GB DLT TO 15/30 GB DLT PERFORMANCE

Operating System	Windows NT 4.0
Backup Software	Cheyenne ARCserve 6.0 for Windows NT
Backup Method	Image
Data Set	2 GB
Compression	1:1, 2:1, 4:1, 8:1
Tape Format	RAIT-0
Hardware	ProLiant 5000, two Pentium Pro 200 processors, 256K cache, 256 MB RAM 96-GB array (two SMART-2/P Array Controllers with twenty-four 4-GB drives) 35/70 GB DLT drive(s) Wide-Ultra SCSI-3 cards (All Compaq off-the-shelf products)

Purpose

To compare the performance of the new 35/70 GB DLT drive with that of the 15/30 GB DLT drive. Though the new drive is faster and holds more, there are cases where the 15/30 GB DLT drive may still be a better match for the backup application.

Result

The 35/70 GB DLT drive, with its larger cache and higher throughput rating, is three times faster than the 15/30 GB DLT drive. In all compression categories, the 35/70 GB DLT drive is the top in its class in performance. However, in some backup scenarios (feeding the enterprise backup server with multiple 10MB Ethernet connections) the 15/30 GB DLT drive speed still exceeds the feed speed of the network link. In these cases, the 15/30 GB DLT drive is an excellent choice, and the use of the Compaq DLT Library is a very attractive and logical option.

The decision between the DLT Library and the DLT Tape Array II is based on backup requirements and the environment. The DLT Tape Array II is fast, but a single (4 tape drive) unit is only capable of storing 280 GB native (560 GB compressed) of tape online. Any backup exceeding those amounts risks significant delays while waiting for an operator to change tapes. The DLT Library has a 15-tape capacity. Eight single-drive DLT Library units (with 15/30 GB DLT drives) would deliver a capacity of 1,800 GB native (3,600 GB compressed).

The DLT Library can have two drives per unit. The speed and capacity of this system makes it an ideal solution for network backups. An example of a typical setup might be a central backup server with four DLT Library units (each with two tape drives). This will support eight simultaneous backup jobs from eight network-connected remote servers or systems. Using groups of tapes, this system (which has been tested by Compaq with Cheyenne software) can perform a weekly full backup of the remote systems, as well as a daily incremental backup. This is a totally automated backup procedure, with no human intervention required. Of course, sets of tapes can be rotated for archiving.

Data compression is an important factor in backup speeds, and in most cases, it is not something the customer can control. In investigations into real customer data set composition, it was found that most servers contain data with a compressibility between 2:1 and 4:1.

The backup speed for non-compressible data highlights the advantage of the data compression hardware built into the drive: The 2:1 data backed up nearly twice as fast as data backed up with the tape drive compression feature turned off. If data set compressibility is 2:1 or greater, customers will achieve significantly faster backup speeds using the 35/70 GB DLT drive, with its built-in compression. As the data set becomes more compressible, the advantage gained is proportionally less: Some of the tests documented in this brief were run with 2:1 data compression; others were run with 4:1 compression; others were run with 8:1 compression. The performance difference between these compression levels is only about 14 percent.




Figure 12: Data compression effects at various levels of compression (indicated at bases of columns).

Compression Ratio	15/30 Backup GB/hr	15/30 Restore GB/hr	35/70 Backup GB/hr	35/70 Restore GB/hr
1:1	3.90	3.93	15.35	15.39
2:1	7.51	6.97	29.50	26.46
4:1	9.68	9.35	33.32	23.07
8:1	9.99	9.80	36.13	28.64

Table 7: Test data for compressibility/performance tests.

TEST 7 RAIT-0 VERSUS RAIT-5 — WINDOWS NT

Operating System	Windows NT 4.0
Backup Software	Cheyenne ARCserve 6.0 for Windows NT
Backup Method	Image — single job
Data Set	2 GB
Compression	2:1, 4:1
Tape Format	RAIT-0 and RAIT-5
Hardware	ProLiant 5000, two Pentium Pro 200 processors, 256K cache, 256 MB RAM 80-GB array (two SMART-2/P Array Controllers with twenty 4-GB drives) 35/70 GB DLT drive(s) Wide-Ultra SCSI-3 cards (All Compaq off-the-shelf products)

Purpose

To determine the impact of RAIT-5 as it relates to processor performance and throughput, to discover how much CPU power is required to perform the RAIT-5 calculations, and to find out whether this overhead will seriously affect backup times.

Results

Tests show that the performance of RAIT-5 in a typical configuration (two SCSI cards connected to four drives) is only 1 to 7 percent slower (depending on data set compressibility) than the RAIT-0 backups. This suggests that RAIT-5 should be considered the optimum choice in a decision between speed and safety. RAIT-5 not only writes to a set of tapes as an array, but uses a parity stripe to provide checksum information, which permits file restoration even when there has been some data loss on a single tape in the array.

On the other hand, RAIT-0 connects the drives together logically and lets them be written to as a single drive. Because RAIT-0 does not perform the striping calculations, it is faster than RAIT-5. There is a significant increase in CPU use for RAIT-5 over RAIT-0 (7 percent versus 41 percent), which could be a problem with concurrent backups or backups made during heavy server activity. The impact of RAIT-5 on server resources can be reduced by adding processors to the system.




Figure 13: RAIT-5 and RAIT-0 performance with various proportions of cards to drives and two compression ratios (shown at bases of columns).

SCSI Cards	Tape Drives	Compression Ratio	RAIT-5 GB/hr	RAIT-0 GB/hr	RAIT-5 % Slower
2	4	2:1	56.16	59.88	7%
4	4	2:1	48.66	55.56	14%
2	4	4:1	60.84	61.5	1%
4	4	4:1	49.92	57.72	16%

Table 8: RAIT/compression ratio test results.
CPU use was 41 percent when using RAIT-5 and 7 percent when using RAIT-0.

TEST 8 NETWORK PERFORMANCE: SIMPLE BACKUP OF A REMOTE CLIENT — WINDOWS NT

Operating System	Windows NT 4.0 (server) pulling from Windows NT 4.0 (client)
Backup Software	Cheyenne ARCserve 6.0 for Windows NT
Backup Method	Single job
Data Set	2 GB
Compression	2:1
Tape Format	RAIT-0
Hardware	ProLiant 5000, two Pentium Pro 200 processors, 256K cache, 256 MB RAM 80-GB array (two SMART-2/P Array Controllers with twenty 4-GB drives) 35/70 GB DLT drive(s) Wide-Ultra SCSI-3 cards (All Compaq off-the-shelf products)
Network	10TX – Ethernet over CAT-5 cable through a simple hub 100TX – Ethernet over CAT-5 cable through a simple hub <i>All tests used TCP/IP protocol.</i>

Purpose

To document the average times for backing up a remote system over the network at Ethernet speeds of 10 and 100 megabits per second.

Results

The results clearly show that simple backups of remote systems are slow, due to the network connection. To provide greater overall bandwidth and minimize the impact on daily network traffic, a dedicated network card in the server is recommended for installations using this type of backup. Consideration should also be given to providing local tape drives to those large systems if better remote performance is desired. The 100TX theoretical maximum is approximately 40 GB/hr, but throughput can be significantly reduced by other factors, such as the protocol used (TCP/IP vs. NETBIOS, for example).

If network feed speed is less than the maximum rate at which the target drive can accept and write data on a continuous basis (or *stream*), the target drive drops out of write mode, and a significant amount of time is lost to shuffling the tape for the next sequential write.

The 35/70 GB DLT drive streams at more than 30 GB/hr. A network must run at that speed or greater to prevent the drive from running out of data and stopping. The 10 Mbit Ethernet moves data at less than 3.9 GB/hr, so a single job cannot exceed that, unless an agent is used to pre-compress the data. With pre-compression, the system might effectively move 4 or 5 GB/hr. During such a backup, the tape drive constantly starts and stops, wasting a large amount of time repositioning the tape for the next write.

Although a 100TX Ethernet permits up to 39 GB/hr feed speed, backups in these tests only achieved 19 GB/hr. This is not as fast as local backups because of the processing required to move data over the network. Still, at 19 GB/hr, the network has room for more data, and as seen in a test of four simultaneous backups over one network connection (Test 9), the backup rate climbs to 32.4 GB/hr.

As discussed in Test 6, an obvious match for many network backup requirements is the 15/30 GB DLT drive. This drive performs backups at about 10 GB/hr. A pair of these drives in the DLT Library can easily handle two concurrent backup over a single 100 TX Ethernet connection. As network load distribution is an important factor in backup speeds, it is not recommended that more than four systems be simultaneously backed up to two DLT Library units (with two drives in each). To back up more than four systems simultaneously, add additional network (100TX) connections as needed (Figure 14).

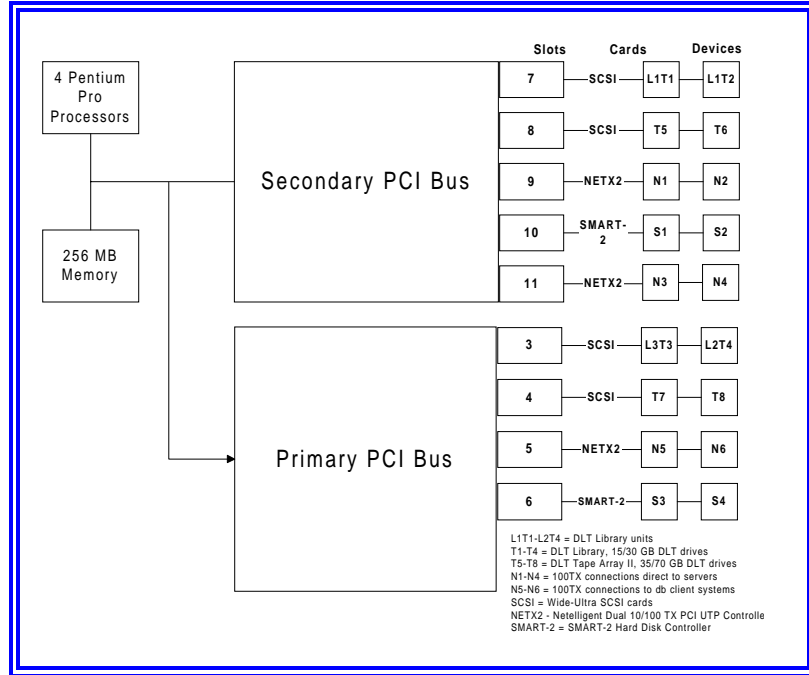


Figure 14: A typical backup configuration for a large network system using Compaq DLT Tape Array and DLT Library solutions.

Such a system can have more than 250 GB of space available for typical server applications. With the DLT Tape Array II, it can easily do high speed backups at night or concurrent backups at any time. The DLT Library units allow a server do double duty as a central backup server for any other servers over its four 100TX network links.. The final two 100TX network connections are used to ensure high performance access to the data on this server from client (user) machines on the network.

Network	Backup GB/hr	Restore GB/hr
10TX Ethernet	3.24	2.82
100TX Ethernet	19	6.9

Table 9: Test data for Windows NT simple backup of a remote client.

TEST 9 NETWORK PERFORMANCE — WINDOWS NT SERVER PULLING FROM AN INTRANETWARE REMOTE CLIENT

Operating System	Windows NT 4.0 (server) pulling IntranetWare 4.11 Server (as client)
Backup Software	Cheyenne ARCserve 6.0 for Windows NT
Backup Method	Single job
Data Set	2 GB
Compression	4:1
Agents	ARCserve 6 for NetWare high-performance push agent option
Tape Format	RAIT-0
Hardware	ProLiant 5000, two Pentium Pro 200 processors, 256K cache, 256 MB RAM 80-GB array (two SMART-2/P Array Controllers with twenty 4-GB drives) 35/70 GB DLT drive(s) Wide-Ultra SCSI-3 cards (All Compaq off-the-shelf products)
Network	10/100TX – Ethernet over CAT-5 cable through simple hub TCP/IP protocol

Purpose

To test the throughput of a multiple server backup over a single network connection.

Result

The throughput of more than 32 GB/hr for four jobs is close to the limit of the 100TX wire, though the NT client backup agent helps slightly. The Windows NT client does allow multiple backup streams; the network is still the limit, in most cases. For this reason, it is not surprising that the Windows NT numbers are very close to the Novell IntranetWare numbers in the network tests. Using multiple 100TX-controller cards on a switched hub to independent clients may allow network backups to approach local backup speeds; this would be scaleable to the limits of the tape drives.

IntranetWare client agent did not support multiple streams to the Windows NT backup server.

The reverse test, with IntranetWare 4.11 server pulling from a base of Windows NT Clients, was not considered for two reasons. If using Cheyenne JETserve, it is strictly an image backup, used as a local backup of the NetWare server's drives. Also, in comparing Cheyenne ARCserve for Windows NT to Cheyenne ARCserve for NetWare, it was noted that the Novell solution was not as robust a candidate. The NT based solution was initially chosen to best illustrate the objective. This is not to say that Cheyenne ARCserve for NetWare would not perform adequately in a Novell environment. Later tests show that all Novell solutions perform as well as the NT solutions.




Figure 15: Results of backups performed with Windows NT Server pulling from IntranetWare.

Test	Network Type	Agent Used	Number of DLT Drives	Backup Jobs on One Network Wire	Backup Rate (GB/hr)
A	10TX	No	1	1	3.6
B	10TX	Yes	1	1	3.9
C	100TX	No	1	1	19.0
D	100TX	Yes	1	1	21.1
E	100 TX	No	2	2 (2 clients)	25.5
F	100TX	No	4	4 (4 clients)	32.4

Table 10: Test data for Windows NT Server pulling from an IntranetWare remote client.

**TEST 10 FIBRE CHANNEL-ARBITRATED LOOP (FC-AL)
LOOP-BASED BACKUP OF 800 GB — WINDOWS NT**

Operating System	Windows NT 4.0
Backup Software	Cheyenne ARCserve 6.0 for Windows NT
Backup Method	By file — four jobs
Data Set	2 GB
Compression	4:1
Tape Format	RAIT-0 (not significant when using one drive)
Hardware	ProLiant 6000, four Pentium Pro 200 processors, 512K cache, 2 GB RAM 2.5-terabyte FC-AL drive array 35/70 GB DLT drive(s) Wide-Ultra SCSI-3 cards (All Compaq off-the-shelf products)

Purpose

To determine whether the DLT Tape Array II and its associated 35/70 GB DLT drives can successfully back up very large amounts of data (800 GB).

Results

This test consisted of running four simultaneous jobs, each backing up a separate logical drive under Windows NT 4.0. Each drive had 200 GB of data in 200,738 files; each file was a 1-MB 4:1 compressible data set. The files were stored in 6,693 directories, in a structure that was nested more than 100 levels deep.

The backup was successful. The backup went to two sets of four tapes (RAID-0), for a total of eight tapes. The tapes were verified and there was no data loss at any point.

This test was performed on the Max-Storage system in the Compaq Maximum Configuration Lab. The lab has multiple ProLiant 6000 systems with FC-AL storage systems. These FC-AL storage systems use 1-gigabit fiber optic links to connect large (2.5 to 3.5 terabytes) disk arrays to the servers (Figure 16).

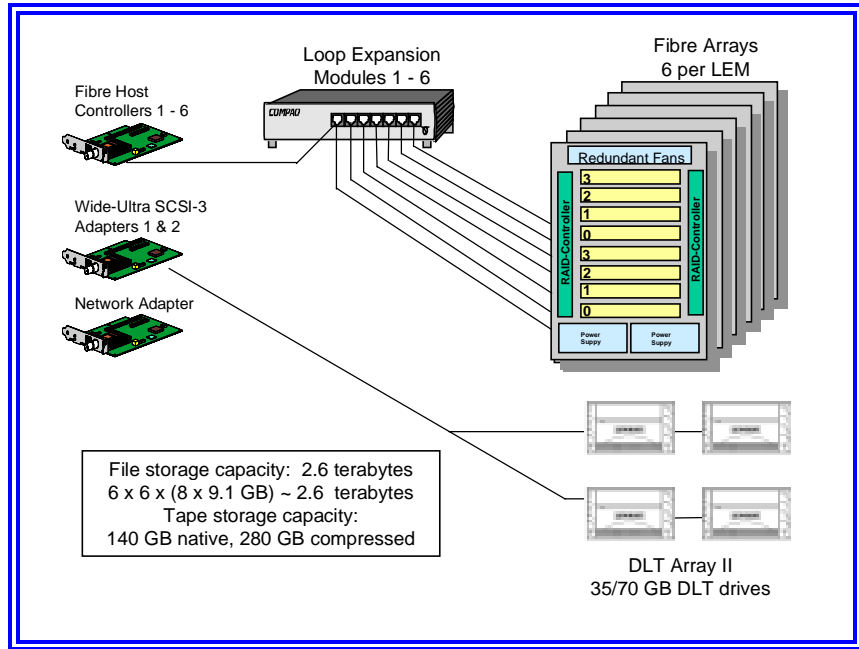


Figure 16: Configuration for 800-GB backup test.

Job 1	27.60 GB/hr
Job 2	26.44 GB/hr
Job 3	26.78 GB/hr
Job 4	26.64 GB/hr
Total backup rate of four concurrent jobs	107.46 GB/hr
Total subdirectories backed up	26,796
Total files backed up	801,024 (800 GB)
Duration (including tape changes):	8 hours

Table 11: Test data for very large data set backup.

TEST 11 FEED SPEED RELATIONSHIPS

Operating System	Windows NT 4.0
Backup Software	Cheyenne ARCserve 6.0 for Windows NT
Backup Method	By image, single and multiple jobs
Data Set	2 GB
Compression	4:1
Tape Format	RAIT-0 (not significant when using one drive)
Hardware	ProLiant 5000, two Pentium Pro 200 processors, 256K cache, 256 MB RAM; and ProLiant 6000, four Pentium Pro 200 processors, 256 K cache, 256 MB RAM Embedded or SMART-2/P SCSI (53C825) Controller (see details below) 35/70 GB DLT drive(s) Wide-Ultra SCSI-3 Cards (All Compaq off-the-shelf products)

Purpose

To demonstrate the relationship of feed speed to backup speed. Single and multiple job tests were performed to demonstrate the relationship of bandwidth to bottlenecks within the operating system

The tests were executed in four environments:

- The server used five 4-GB Fast-Wide SCSI drives, each formatted under Windows NT as a single logical drive. All drives were connected to the embedded SCSI (53C825) controller. Tests A and B used only one drive. Test C used four logical drives (each on a separate physical drive).
- The server data drives used the SMART-2/P Array Controller to operate twelve 4-GB Fast-Wide SCSI drives configured as a single RAID-0 drive. Windows NT was used to configure the array as a single drive letter (stripe set) for Tests A and B. Test C used four logical drives created by assigning three physical disk drives in RAID-0 to each logical disk using the array configuration utility and formatting them under Windows NT.
- The server data drives used two SMART-2/P Array controllers to operate twenty-four 4-GB Fast-Wide SCSI drives formatted as one RAID-0 drive (and one Windows NT-striped drive) for Tests A and B. Test C used four logical disks created by assigning one channel of each SMART-2/P Array Controller to each logical drive (six drives per channel) in a RAID-0. Each Smart-2/P Array Controller card was installed on a separate PCI bus within the server, for performance balancing.
- The server data drives used four SMART-2/P Array controllers to operate fifty-six 4-GB Fast-Wide SCSI drives formatted as one RAID-0 drive (and one Windows NT-striped drive) for Tests A and B. Test C used four logical disks, created by assigning one channel of each SMART-2/P Array Controller to each logical drive (seven drives per channel) in a RAID-0. Each SMART-2/P Array Controller card was installed on a separate PCI bus within the server, for performance balancing.

Test A

Perform baseline analysis using Cheyenne’s RAWREAD.EXE to obtain ratings of drive performance.

Test B

Perform a backup of a 2-GB data set per drive. This single job image backup to DLT Array II with four 35/70 GB DLT drives (driven by two Wide-Ultra SCSI-3/PCI cards) was written at RAIT-0.

Test C

Perform four backup jobs of four 2-GB data sets from four logical drives.




Figure 17: Results of tests A, B, and C in four environments.

Hard Disk Read Tests	A	B	C
1. ProLiant 5000, embedded SCSI hard diskdrives (53C825)	33 GB/hr	30 GB/hr	55 GB/hr
2. ProLiant 5000, one SMART-2/P Array Controller, 12 4-GB Fast-Wide SCSI drives	126 GB/hr	48 GB/hr	107 GB/hr
3. ProLiant 5000, two SMART-2/P Array Controllers, 24 4-GB Fast-Wide SCSI drives	251 GB/hr	65 GB/hr	115 GB/hr
4. ProLiant 6000, four SMART-2/P Array Controllers, fifty-six 4 GB Fast-Wide SCSI drives	454 GB/hr	84.5 GB/hr	135 GB/hr

Table 12: Relationship between disk drive feed speed and backup performance.

Results

The results clearly show this performance relationship: The faster the drive system, the faster the backup. Test B shows backup performance can be almost tripled (from 30 to 84.5 GB/hr) with hardware that provides data to the tape drives fast enough. Test C demonstrates that performance can be increased with multiple jobs, but only with adequate feed speed. When the read:write acceleration ratio was changed from 50:50 to 100:0 on the SMART-2/P Array Controller setup, there was no significant change. It is recommended the acceleration ratio setting remain at the default of 50:50.

.....
CONCLUSION

The DLT Tape Array II sets new standards in high-speed backup performance. The tests documented in this brief demonstrate that Compaq delivers very high speed backup solutions for an assortment of server environments, with the ability to back up data at speeds exceeding 150 GB/hr. Fast, reliable backups can now be made by any customer using similar configurations of high-performance Compaq products.

**APPENDIX 1: TIPS FOR CONFIGURING AND OPERATING THE
COMPAQ 35/70 GB DLT DRIVE**

Cable Quality

The cables provided with the various Compaq option kits are of very high quality. If they are ever damaged, replace them with identical Compaq parts. SCSI data transfers will operate poorly (or not at all) when signal conditions are degraded due to low-quality cabling.

Data and Equipment Protection

Always screw the external cable heads into the SCSI cards and DLT Tape Array II chassis to prevent damage to the connectors and possible loss of data.

Transfer Speed

If using an embedded SCSI controller to drive a single 35/70 GB DLT drive, ensure that it is a Compaq 32-bit Fast-Wide controller or install a 32-bit Fast-Wide SCSI card to improve performance. Use of a 32-bit Wide-Ultra SCSI-3 controller is *strongly recommended* if there will be two 35/70 GB DLT drives per card. To identify the card, look at the main logic chip in the center of the card and verify that it is a Symbios Logic 53c875 part. Cards with the NCR 53c825 part will not deliver maximum performance.

Disaster Recovery/Performance

For high performance, install four DLT drives (typically using the DLT Tape Array II chassis), connected by two cables to two 32-bit Fast-Wide SCSI PCI cards in the server. The drives can be configured as four separate drives running four separate jobs, or one as a single RAIT group running a single job.

Though performance may be better with a RAIT-0 configuration, RAIT-5 offers better protection from data loss by striping the tape with additional checksum/recovery information. For the highest margin of confidence, configure a set of mirrored drives (2x2 in one DLT Tape Array II or two DLT Tape Array II devices for a 4x4 set). Mirroring makes full duplicates of the tape and adds the advantage of storing tapes in two locations.

Cleaning

DLT drives have a "Use Cleaning Tape" indicator on the front panel. Clean the drive by loading a cleaning tape (PN 199704-001). It will clean the drive, then automatically unload. The cleaning tape is effective for 20 cleaning cycles. Do not exceed this. If a backup tape *causes* the cleaning light to come on (versus this occurring during a restore function) copy the tape or recover it and stop using it. Such a tape can contaminate the drive, which may then contaminate other tapes if the drive is not cleaned between loads. *Never* try to clean a DLT drive with anything other than a cleaning tape.

Operating Considerations

To avert tape leader failures, wait two seconds after the green indicator turns on before extracting a tape cartridge. Never move the drive while a tape is engaged. The tape moves past the head at 170 inches per second; moving the unit while it's operating can cause tape or drive damage. It is also important not to drop or physically shock the DLT tapes. Doing so can damage the tape, which in turn can damage the drive (as well as lose data). If the cartridge appears damaged, do not use it. It is best not to risk damaging the drive.

⋮ **DLT Media Considerations**

⋮ Although the drive can read half-inch DLT media labeled DLT Tape II, DLT Tape III, DLT
⋮ Tape III XT or DLT Tape IV, the DLT Tape IV delivers optimum performance. The DLT
⋮ Tape IV medium holds a minimum of 35 GB of uncompressed data. The 70 in 35/70 GB DLT
⋮ refers to the drive's ability to compress data, usually at a factor of at least 2:1. The drive can
⋮ compress some data at more than 8:1, but to be conservative, the industry uses 2:1 as an
⋮ average figure. When backing up servers containing an assortment of general end-user file
⋮ data, the DLT Tape Array II easily exceeds 100 GB per tape cartridge. The DLT Tape IV has
⋮ a minimum shelf life of 30 years if stored properly. Its rated usage is 500,000 passes (filling
⋮ the tape more than 9,000 times); however, it is suggested that tapes be rotated on a more
⋮ frequent schedule.

⋮ **Drive Duty Cycle**

⋮ Most tape drives have a 15 to 25 percent duty cycle, which means the manufacturer expects
⋮ them to actually be writing tape some small percentage of time during their installed lives. The
⋮ DLT drive is a 100 percent duty cycle rated device: While it won't be writing 100 percent of
⋮ the time, the drive is tougher and can be used more often, with greater trust in its reliability
⋮ than drives with a lower percentage duty cycle.

**APPENDIX 2: INSTALLATION NOTES FOR THE COMPAQ DLT
TAPE ARRAY II**

General

The Model 4 DLT Tape Array II comes with:

- Rack-mountable array cabinet with internal power supply.
- Four internally-mounted 35/70 GB DLT drives.
- Four DLT data cartridges.
- One DLT cleaning tape.
- Two internal SCSI cables.
- Two external 12-foot Fast-to-Fast SCSI-2 cables.
- Two external 12-foot Fast-to-Wide SCSI-2 cables.

It is ready to install as shipped. The internal cables are connected and the drives are pre-addressed.

Installation

The Model 4 (with four drives preinstalled) is factory configured and the addresses are preset. It is connected to the SCSI cards or ports. Other models require that the drives be addressed and mounted in the cabinet before mounting the cabinet in a rack system.

Reliability

Use the cable management arm. It adds life to the cables and reduces strain on the connectors.

Safety

Install the unit into a 19-inch rack by following the directions in the installation guide. Do not take any shortcuts. Use all the hardware provided to ensure a safe and reliable system.

APPENDIX 3: STANDARD HIGH-PERFORMANCE/HIGH-SPEED BACKUP TEST SYSTEMS

System	Compaq ProLiant 5000 and ProLiant 6000 with dual peer PCI bus system architecture
Processors	Two Intel Pentium Pro processors with 256K L2 cache
Memory	256-MB ECC DIMMs
Disk Array	Compaq SMART-2/P Array Controller(s)
Configuration	Two SMART-2/P Array Controllers feeding four channels, with a total of 24 drives
Disk Drives	Compaq 4.3-GB Pluggable 1" Wide-Ultra SCSI-3
SCSI Cards	Compaq Wide-Ultra SCSI Controller (PCI)
Network Cards	Compaq Netelligent 10/100MB/s TX PCI Ethernet Controller
Tape Drives	Compaq DLT Tape Array II using Compaq 35/70 GB DLT drives
Operating Systems	Windows NT 4.0 Build 1381 Novell IntranetWare Version 4.11 Build 2884
Software	Cheyenne ARCserve Windows NT Version 6.0 Build 358

This is typical of the system most tests were performed on. Performance can be improved with the addition of disk drives, disk controller cards, memory, and processors.