# WHITEPAPER

May 1997

Prepared By
Intranet Solutions
Engineering

Compaq Computer
Corporation

## CONTENTS



COMPAQ

591A/0697ECG

# Backup and Anti-virus Solution for Lotus Notes

*Cheyenne is a major partner of Compaq for backup solutions. Cheyenne's backup software, ARCserve has proven its reliability with Compaq hardware in performance tests conducted at Compaq on platforms such as Novell NetWare, MS Windows NT and with applications such as SAP R/3. Recently, Cheyenne has released the Protection Suite for Lotus Notes for Windows NT. This software solution is intended to assist enterprises running Lotus Notes 4.x and Domino servers with two of the biggest concerns facing administrators: backing up open files and preventing viruses in email attachments.*

*Compaq has previously tested the functionality of ARCserve 6 for NT using Compaq DLT tape drives and DLT arrays. Compaq DLT Tape Array Models 0 and 4 are designed to allow backups of 10 megabytes of data uncompressed — 20 compressed on Model 0 and 15 megabytes of data uncompressed and 30 compressed using Model 4.*

*The focus of this paper is to determine the functionality of all the components in the Cheyenne Protection Suite for Lotus Notes for Windows NT. Compaq is also providing performance data useful in understanding backups and restores of Lotus Notes servers on a specific testbed using Compaq DLT Tape drives and arrays in a given scenario.*

Compaq has delivered White Papers on ARCserve 6.0 for NT and NetWare, Backup of NT SAP R/3 Systems with ARCserve and DLT Technology. You can use the search engine at www.compaq.com to find these technical documents. The focus of this White Paper is the functionality and performance considerations of the Cheyenne AntiVirus Agent for Lotus Notes and the Cheyenne Backup for Lotus Notes.

## 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 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.

# INTRODUCTION

Organizations depend highly on the messaging and workflow automation provided by groupware. Many organizations store gigabytes of email messages, document databases and mission critical applications developed across multiple servers. Data loss can be catastrophic in most environments, resulting in the loss of days or weeks of productivity. For these reasons, backup management and anti-virus protection is vital to a successful groupware implementation. A well thought-out backup management strategy can reduce lost productivity due to hardware or software failure. A proven anti-virus software solution acquisition can also deter painstaking, tormenting hours of "cleaning" up your data network volumes and client stations.

## Understanding Messaging/Groupware Systems

Electronic messaging systems such as Lotus Notes are a common way for today's corporations to communicate. Quite often, the messaging system becomes an essential method for sharing information and documents. Unfortunately, these systems can provide a "safe harbor" for viruses to reside and rapidly spread through an organization—posing risks to both data and productivity. The data residing in these types of messaging servers is most often in the form of e-mails, threaded discussions and document attachments. Document attachments can range from trivial text messages to detailed sales proposals to complex financial spreadsheets. The ability to centrally post files or documents for collaborative editing and distribution among various workgroups is one of the major benefits of messaging systems. Voice mail, video and other types of data are increasingly routed through the messaging systems straight to user PCs.

In general, the messaging server is the central storage point for all data thus creating large, server-based databases. Although data is often mirrored on desktop and laptop PCs, the server-based storage of the data enables the essential data sharing capabilities. The databases often contain gigabytes of data in compressed and proprietary formats.

## Protection of Messaging Data

As organizations continue to revolve business processes around groupware messaging systems, this data becomes extremely time-sensitive and business-critical. Administrators are responsible for establishing "safety nets" to prevent against the loss or corruption of this data. However, the decentralized nature of messaging systems makes administrating, managing and specifically, the protecting of this data a logistical nightmare. Equally challenging is ensuring that the entire organization has access to messaging data -- 24 hours a day, 7 days a week.

In the past, administrators would shut down or bring the messaging servers off-line to effectively backup the system with full data integrity. Although this was an effective method, it greatly impacted user access. It was also impossible to detect, much less cure viruses in documents that were attached to messages or stored in groupware databases. In today's global environment, organizations demand robust and full feature backup and virus protection tools that are easy to use and manage, without impacting user activities. Cheyenne has answered this challenge by providing comprehensive backup and Anti-virus solutions for messaging and groupware systems. These solutions are designed to address the unique requirements of the leading systems - notably Notes.

Many client/server anti-virus solutions (Cheyenne InocuLAN included) can protect users from viruses when documents are detached from messages; however, only the Cheyenne anti-virus messaging option can completely scan and cure server-based messaging systems such as Lotus Notes.

In addition, the integration of Cheyenne InocuLAN and ARCserve provides a solution for virus-free backups.

Anti-virus software that offers real-time protection checks each file as it is written to the backup media. With thousands of files to backup, the back-and-forth process of the anti-virus and backup programs significantly impacts server performance and dramatically lengthens the time to complete the backup.

With most advanced backup solutions, every file is actually opened several times during a backup. ARCserve 6.0, for example, opens each file once during the estimation process, once during the backup process and potentially once during the verification process. An anti-virus product (other than InocuLAN) will scan the file each time it is opened, which significantly slows the backup. Most anti-virus products will try to scan each file as the backup occurs -- forcing the server to "thrash" back and forth between two CPU-intensive operations.

Approaches used to avoid these problems are to temporarily turn off real-time virus protection, to schedule separate times for backups and virus scans or to check only incoming (and not outgoing) server files for viruses. Unfortunately, all three approaches defeat the benefits of using anti-virus software in the first place and significantly increase the chance of including a virus in your backup.

While image backup products can help speed the backup process, they do not replace the need for virus protection. A separate virus scan is still necessary before running an image backup to ensure that the backup is virus-free. When using automatic scheduling, such an approach requires constantly estimating the time to complete the virus scan.

Cheyenne's ARCserve and InocuLAN provide an integrated solution for virus-free backups.

When ARCserve begins a backup, it checks to see if InocuLAN is loaded. If so, ARCserve passes a list of the files to be scanned to InocuLAN. InocuLAN then quickly scans the specified files, passing back the results of the scan to ARCserve. Detected viruses are handled and alerts are sent as specified. ARCserve continues the backup without any further intervention from InocuLAN. The two products will continue in this manner, until the backup is complete and virus free.

This method is significantly faster than the file-by-file approach that occurs with other backup software in conjunction with any other anti-virus software.

## Additional Procedures

Performing regular tape backups is only a small part of backup management. Network administrators should develop a backup management strategy that includes the following:

- Completing an evaluation of the network and data to be backed up, and developing a risk assessment based on the nature of the data and the amount of risk incurred by data loss.
- Establishing a tape backup plan, including the frequency of backups, the amount of data to be backed up, and a backup schedule for each server.
- Organizing network directory structures and defining a data management plan.
- Establishing a data restoration procedure.

Network administrators must also:

- Follow through with periodic downloads of updated Anti-virus signature files from the software vendor.
- Instruct end users and suggest implementing company policy regarding the handling of transferable media such as 3.5" diskettes.
- Develop "disinfecting" procedures for detected software viruses.

Cheyenne ARCserve Protection Suite provides the software components to aid a groupware administrator in implementing the noted steps above. Compaq provides the DLT hardware technology to complement ARCserve's software backup component. This paper focuses on Cheyenne Protection Suite for Lotus Notes for NT.

## ARCSERVE 6.0 FOR WINDOWS NT

ARCserve for Windows NT 6.0 is a data management and backup program that allows you to back up data from your Windows NT workstation or other machines attached to your network.

ARCserve's advantages include:

- Scheduled backups: schedule repeating backup jobs that are custom-made for your environment.
- Tape rotation: plan your tape rotation schemes using a friendly interface.
- Tape spanning: large jobs automatically span multiple tapes until the job is complete.
- Parallel streaming: runs jobs concurrently when using multiple tape devices.
- Reports: view job history and activity on the machines, directories and tapes.
- Cheyenne anti-virus integration: automatically scan files during a backup operation.
- User-defined scripts: configure jobs once and re-use them when needed.
- Intelligent restore: the ARCserve database system allows quick location and restoration of archived data.
- Tape support: back up your files to a wide variety of SCSI 4mm, 8mm, DLT and QIC tape drives.
- Single Server: reduces network traffic because the data can stay locally.
- Remote Management of ARCserve Server: the Enterprise version allows management of multiple servers from a single machine; gives the ability to back up, copy and restore a Windows NT machine and Novell NetWare server in the network.
- Event-driven notification system: the Cheyenne Alert notification system informs of key events in your operations.
- Enterprise client agent support: back up remote machines using the Cheyenne client agents (such as Client Agent for NetWare), which employ the latest PUSH technology.
- Application agents: allows backup of complex databases and groupware such as MS-Exchange and Lotus Notes.
- Windows NT 3.51/4.0 support: backup compressed files, directories and drives on Windows NT File System (NTFS) formatted drives. Also supports Service Security Impersonation.
- Auto Changer (optional): use an autochanger with the backup rotation.

- NetWare compatibility: back up, copy and restore any Windows NT machine and Novell NetWare server in the network.
- Easy-to-use interface as seen in Figure 1.

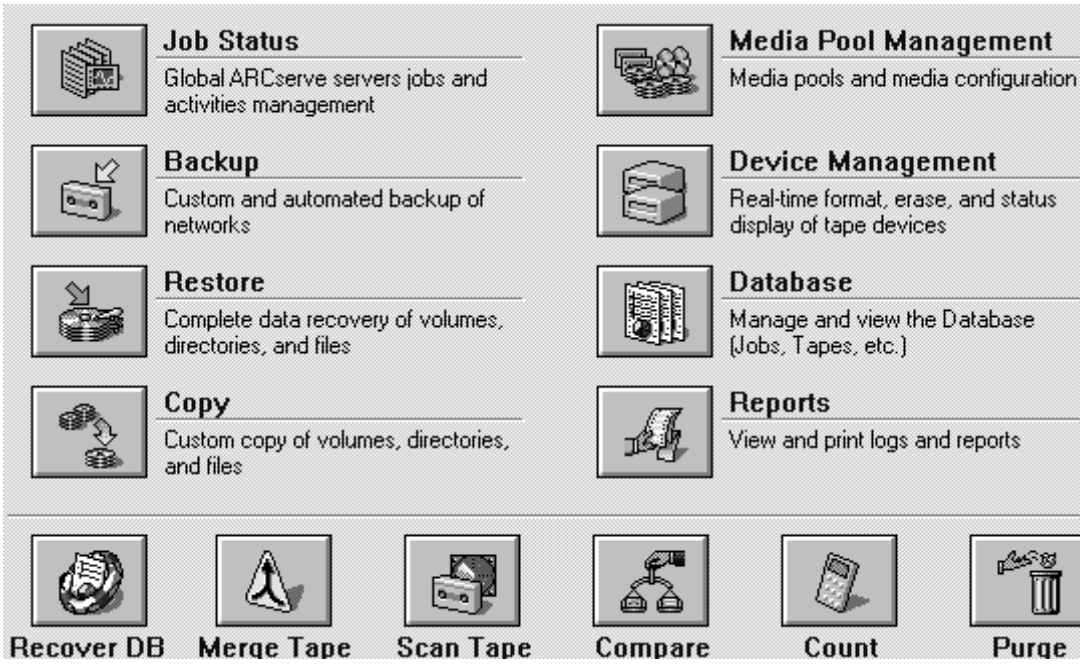


*Figure 1*

## BACKUP AGENT FOR LOTUS NOTES V2.0

The Backup Agent communicates with Lotus Notes and ARCserve to back up the open Lotus Notes databases on a local or remote server.

The following are the core benefits of using the Backup Agent for Lotus Notes included in the Protection Suite for Lotus Notes:

- The Lotus Notes server does not need to be stopped to perform the backup. This is ideal for businesses that operate 365 days-per-year on a 24-hour basis.
- The Backup Agent 2.0 uses "open file technology", which allows users to backup an on-line database while users are accessing and manipulating data. This streamlines the backup process by eliminating the need to replicate the database and to take the Lotus Notes server off-line to back up the database. The Backup Agent v2.0 works in conjunction with ARCserve Backup Agent for Open File v3.1 for Windows NT.
- Access to a wide array of tape storage devices.
- Perform full backup jobs on Lotus Notes databases.
- Flexible scheduling capabilities through ARCserve. Submit a job for a specified date and select a repeat method, an interval for the job to be run again, or select a rotation scheme, a pre-set weekly backup strategy consisting of overlapping patterns of full, incremental and differential back up jobs.
- The backup is complete, usable and consistent.

- Safely back up the Lotus Notes database from any remote server running Lotus Notes and the Backup Agent.

New features in v2.0 of the Cheyenne Backup Agent for Lotus Notes include:

- High Performance and Reduced Disk Requirements — Use of Cheyenne advanced technology to dramatically improve backup throughput and reduce the disk space required for Notes server backup.
- Backup of Encrypted Databases — Browse and back up encrypted databases using the same procedures as standard database files (Note: security is maintained — the data is backed up and restored as encrypted files).
- Restore to alternate directory — Permits administrator to specify destination for restored files. This is useful when testing restore procedures, moving files between servers or changing system configurations.
- Backup and restore of shared mail — Backup of the shared mailbox instead of multiple individual copies.
- Support for Lotus Notes 4.5 — Supports the new Lotus Notes Domino server.

### Known Limitations

(Documented in Backup Agent v2.0 for Lotus Notes Server Release Notes):

- The Backup Agent 2.0 does not support Lotus Notes v3.x or Domino *Partitioned* servers.
- Restoring a database with a database pointer will restore the database file only and **not** the pointer (text) file. If this file was not explicitly selected for backup (from the ARCserve files system backup browser), then it will not be restored; the pointer file must be recreated.
- Restoring shared mail databases is a complex task. If you are unsure about the requirements of backing up these databases, please consult your Notes Administrator guide or call Cheyenne Technical support **before** you back up your shared mail databases.
- Browsing Lotus Notes databases in the ARCserve Manager may cause the program to hang if shared mail is enabled. To avoid this, change the value of the Backup Agent registry entry "SecondaryServer" from 0 to 1.

## INOCULAN 4

Cheyenne InocuLAN 4 for Windows NT is a powerful, second-generation anti-virus solution for a Windows NT network. Available options provide protection for Lotus Notes and Microsoft Exchange messaging systems.

InocuLAN uses a number of real-time components to protect all avenues of entry into the Windows NT enterprise, including:

- Real-time Scanning Mode: All files going to and from the server are scanned for viruses, including compressed files. Viruses will not spread through your network with InocuLAN in real-time operation.
- Virus Wall: A little-known but very dangerous security leak that many anti-virus products can not stop is the infection of a server by a workstation. InocuLAN stops any infected file from being copied to a server and replaces the clean version of the file, thereby keeping enterprise security intact.

- Virus Quarantine: Users who try to copy infected files to a server are automatically suspended from the machine, isolating the infection before it can spread. A message is sent listing the name of the user who tried to move an infected file.
- Floppy-drive protection: Floppy diskettes are the most common source of virus infections, and InocuLAN fully protects the enterprise from floppy-based viruses. As soon as a floppy diskette is accessed, such as looking at the disk contents in My Computer, InocuLAN scans the boot sector, preventing the spread of dangerous boot viruses. When a file is opened or copied from the floppy, InocuLAN scans it before it moves to the hard drive.
- Network drive protection: Another little-understood but common way of spreading viruses happens when files are copied from one mapped drive to another. Even though no file passes through the hard drive of the local machine, InocuLAN will still scan all files moving between mapped drives.
- Internet-enabled: The newest source of virus infections is the Internet. As users gain nearly limitless access to computers worldwide, the probability of downloading infected files grows exponentially. With InocuLAN running, all file downloads are automatically scanned for viruses before they can infect a machine. This includes support for compressed files. InocuLAN works with browsers from Netscape and Microsoft.
- Groupware Messaging Anti-virus options: More than ever, companies are communicating electronically. As more data is being exchanged, more viruses are spreading by hiding in mail attachments and database files. InocuLAN can protect your Lotus Notes or Microsoft Exchange mail systems with its messaging options. Even attached ZIP files are scanned.
- Support for Windows NT 4.0
- Includes shell extension integration, providing right-click scanning from any volume, folder or file.
- Multi-platform support: InocuLAN NT version for Intel, Digital Alpha, NEC MIPS and Motorola Power PC.
- Microsoft BackOffice support: InocuLAN NT carries the "Designed for BackOffice" logo.
- NCSA Certification ensures protection against 100% of the computer virsuses, as certified by the National Computer Security Association (NCSA).
- New Multiple Source Browser: A new Explorer-like browser makes viewing and selecting servers, directories and files faster and easier. Multiple sources can be selected for scanning.
- NetWare domain management allows administration of InocuLAN NetWare servers through the Windows NT console.
- Real-time Copy Cure option: Makes a copy of the infected file before curing it.
- Automatic Software Download, Distribution and Update: Hands-free downloading and distribution of the latest signature files and search engines using modem or FTP downloads. Supports multi-language, multi-platform networks.
- Point-to-Point Management: InocuLAN Servers can be managed by entering the machine name. This means InocuLAN can communicate with all InocuLAN servers, even across segmented LANs where broadcasts are filtered out.
- Compressed Files: Scans compressed files and Internet downloads in .ZIP, .ARJ and Microsoft compressed formats.
- Scheduled Scanning allows administrators to scan networked servers at predetermined times.

- Domain Support allows configuration of servers into InocuLAN domains. Multiple servers can be configured at one time.
- Updated Alert System immediately notifies selected users of a virus threat through network broadcast, print queue/trouble ticket, Microsoft Mail, Microsoft Exchange, SNMP and pager.
- Remote System Event Log Support: Uses Alert 4.0 to forward Alarm information to remote server's system event logs.
- Flexible Reporting includes scanning results, virus incidents, configuration changes and status reports. Reports are completely automated and centralized across InocuLAN domains.

## Lotus Notes Anti-virus Agent

The recently released Cheyenne AntiVirus Agent v2.0 for Lotus Notes integrates with Cheyenne's InocuLAN to scan and detect viruses in documents attached to e-mail messages and Lotus Notes databases. Infected Lotus Notes attachments can be automatically cured. The users are notified through the host messaging system or through InocuLAN's Alert system.

*NOTE:*
*You need InocuLAN build 241 or higher for AntiVirus Agent functionality.*

In order to take advantage of this new technology, Cheyenne has updated the InocuLAN system. The CD-ROM included in the Protection Suite package contains updates for InocuLAN v1.01 for Windows NT and InocuLAN v4.0 for NetWare.

### New Features in Cheyenne AntiVirus Agent v2.0 for Lotus Notes

In addition to the existing full and incremental scanning, scheduled scanning and extensive cure and alerting options, the AV Agent for Lotus Notes now offers:

- Real-time Scanning and Cure: e-mail attachments are automatically scanned at the point of entry into the messaging system. Upon virus detection the file can be cured in real-time, or other actions such as deletion or copy. The sender, recipient and/or administrator will be notified so that corrective actions can be taken to prevent future transmissions of infected mail.
- Scanning of shared mail: Transparently scans messages sent to multiple recipients in a single operation (rather than scanning when each user accesses the message).
- Scanning of encrypted databases on the Notes Server: Detects and cures viruses in encrypted databases, including mail databases on the Notes server.
- Support for Lotus Notes 4.5: Supports the latest Lotus Notes Domino server.

InocuLAN will scan e-mail attachments in three ways: Real Time scanning, Scheduled scanning and Local or Immediate scanning.

The Real Time scanner can be configured to notify the Notes Administrator, the Message Sender or the Mailbox Owner. The Agent can also be configured to include a text file attachment in the infected e-mail detailing the action that was taken on the virus. To protect the Lotus Notes Server in Real Time, the Agent "hooks" into the mail router and scans any e-mail attachments that pass through it.

The Local scanner can be configured to scan the entire e-mail system or individual mailboxes. In addition to the notification capabilities that are available in the Real Time scanner, the Local scanner can be configured to scan all messages, messages after a certain date/time, and on an incremental method.

Scheduled scanning is done through InocuLAN's Domain Manager. In addition to having the same notification options available to you as in the Local scanner, it allows setup of a scheduled

job to scan the entire Notes Server. The Domain Manager can even be used from a machine not running the Notes Agent to setup and configure Lotus Notes scanning. These machines can also be configured into an InocuLAN Domain to utilize its centralized management capabilities.

### Known Issues

(Documented in Cheyenne AntiVirus Agent v2.0 for Lotus Notes Release Notes)

- Local scanner will only scan the first 1000 e-mails in any individual mailbox.
- Shared e-mail databases cause MDA (134) error messages! These are due to a limitation of shared mail databases; Cheyenne has so far identified two cases that they occur in:
  - Any action other than Report Only will generate this type of error message.
  - Notification via Attachment will generate this type of error message.
- The agent will only support scanning of explicit server-level encrypted databases. .NSF databases encrypted on the local will be skipped by the agent and no virus scan will be available.
- Mail messages sent by encrypt option will not be scanned by the agent.

## DLT TECHNOLOGY

There are many configurable combinations of Compaq hardward compatible with Lotus Notes and Cheyenne ARCserve Protection Suite. The setup a customer has depends on his/her business needs and Compaq assures there is a server for every size company.

The tests of the Cheyenne Protection Suite in this paper were performed on Compaq ProLiant 4500 rack mountable servers with one Compaq DLT 4 Drive Tape Array 10/20.

The basic DLT technology has 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. DLT divides the tape into parallel, horizontal tracks and records data by moving the tape at high speeds past heads that remain stationary while reading a track or writing to it. (DLT heads do shift vertically to access other tracks.) During a write operation, the first (vertical) set of write heads encountered by the tape writes new data to multiple tracks on the tape, overwriting existing data on these tracks. The middle set of heads 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 vertical set of write heads does nothing until the tape switches direction. When the tape reverses direction, the two sets of DLT write heads switch roles. The same write-read-compare procedure is followed, but for the tape moving in the opposite direction.

As DLT technology continues to be developed, the number of tracks on the tape increases.

Because DLT tracks are horizontal, manufacturers can add more read/write elements to the heads to increase data-transfer rates. Tape and head life are also increased because the tape is not

pulled out of the cassette and wrapped around a rotating drum as it is with helical-scan technology.

DLT drives read and write multiple parallel tracks in a "serpentine" pattern. (See Figure 2.) 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 2 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 then only necessary to search 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.

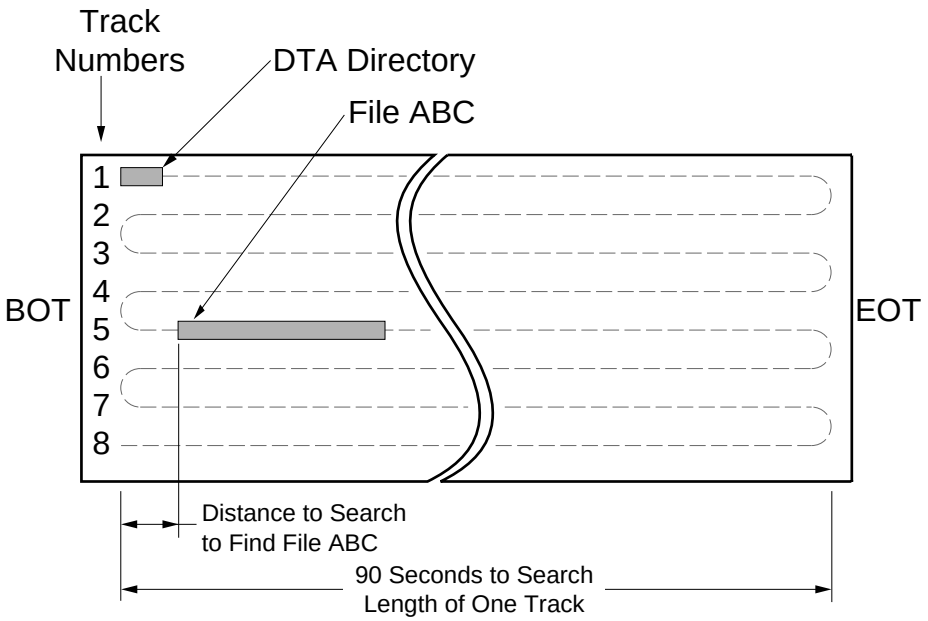


*Figure 2: DLT Serpentine Track Layout*

## Compaq's Current DLT Drive

Compaq currently offers the 15/30-GB[1] DLT drive in ProLiant and ProSignia servers however, you can also purchase the Compaq 10/20-GB drive. These drives offer a 50% increase in capacity over the earlier 10/20-GB DLT drives. The 15/30-GB drive specifications are listed in Table 1. This section describes the drive firmware, performance prerequisites, hardware compression, and error checking and correction. Also included is a description of DLT-drive head-cleaning requirements.

[1] 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 performance 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 |

The *RLL* (Run Length Limited) encoding method 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 0s before a binary 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.

## Achieving Published Performance

To achieve its published performance specifications, the Compaq DLT 15/30-GB 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 entitled DLT MEDIA in this brief.

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

| | | |
|---|---|---|
| **Cartridge** | *CompacTape IIIXT*[2]<br>• 242465-001 (single cartridge)<br>• 242466-001 (package of 7) | 0.5-inch tape, 1800 feet long. |
| **Formatted Capacity** | Native (uncompressed) | 15 GB (unattended operation) |
| | With 2:1 Compression | 30 GB (unattended operation) |
| | Average File-access Time | 68 seconds |
| **Sustained Transfer Rate** | Native (uncompressed) | 1.25 Mb/second |
| | With 2:1 Compression | 2.5 Mb/second |

## Hardware Compression

*NOTE:*
*ARCserve itself does not perform any software compression. Any compression done during backups is done by the tape drive, at the hardware level. Compression must be physically enabled on the tape drive. ARCserve can be configured to indicate whether or not hardware compression is enabled and utilized.*

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 re-coding 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% 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; graphics files, 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.

## Checking Compression via ARCserve

In ARCserve's Device Management screen, highlight the tape unit. If the compression icon is available for selection, compression is enabled on the drive. You can also look at the detail tab that appears in the right panel when the drive is highlighted to determine the status of compression. A field labeled Compression will indicate that compression is either On, Off, or N/A.

## Enabling Compression

Make sure there is a blank, unformatted tape in the drive. Click the Compression icon. Compression on the tape will be toggled on or off, depending on the current status of the tape drive's compression. Once it is flagged as "on", you can do a manual format or just let the backup execute at its scheduled time and it will format the tape to use compression.

[2] *CompacTape* is a registered trademark of the Digital Equipment Corporation.

## Hardware Reliability

The basic DLT design is inherently more reliable than those of DLT's counterparts from the older technologies. 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% duty cycle (continuous operation).[3] By comparison, the estimated head-life of 8-mm helical-scan drives 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 multi-layered approach that interleaves the following during data recording:

- A 16-Kbyte Reed Solomon error-correction code (ECC) with every 64 Kbytes 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.

*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 periodic cleaning. When a DLT drive does needs 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.[4] 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.

---

[3] 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% duty cycle.

[4] One Compaq DLT cleaning tape lasts for approximately 20 cleaning operations.

### 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, DLT tape media have remained in good condition even after very high numbers of passes. [5] 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.

### 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% 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%.
- 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 by automatically notifying administrators of key events. The following are two of the latest automatic-notification features available with Compaq Insight Manager:

[5] In these lab tests, 15/30 DLT cartridges have exceeded 500,000 head passes with only a minimal increase in the number of soft errors detected.

- 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 MODEL 0 AND MODEL 4

Compaq DLT Array products take advantage of the fact that DLT drives read and write to multiple parallel tracks simultaneously. Compaq DLT Arrays comprise 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. Cheyenne has recently released ARCserve RAID Option v1.0 for NT. You can find additional detailed information further in this paper.

Compaq DLT Arrays have set a new standard in backup performance. Model 0 is the 10/20 GB capacity drive array and Model 4 is the 15/30 GB capacity drive array. Compaq conducted extensive backup tests using two 4-drive DLT Arrays and four SMART-2 Array Controllers; one controller for every two drives for maximum performance. The drives were the Compaq 15/30-GB model, and the software was JETserve configured for RAID 5 fault tolerance. The performance testing consistently demonstrated a backup throughput of 46 GB per hour. At this rate and with this equipment, a network administrator can back up 210 GB 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 recommends that customers use one controller for every two DLT drives. This will provide maximum performance. See chart below for more details:

| # of DLT Tape drives: | 1 SCSI controller: | 2 SCSI controllers: | 3 SCSI controllers: | 4 SCSI controllers: |
|---|---|---|---|---|
| 3 | 14.0 GB/hr. | 15.0 GB/hr. | 15.0 GB/hr. | |
| 4 | 11.0 GB/hr. | 21.0 GB/hr. | 22.0 GB/hr. | 22.0 GB/hr. |
| 6 | | 24.0 GB/hr. | 33.0 GB/hr. | 34.0 GB/hr. |
| 8 | | 23.0 GB/hr. | 33.0 GB/hr. | 46.0 GB/hr. |

## Compaq DAT Drives

Compaq recently relaunched the Compaq DAT Family to reset expectation on capacity and to introduce new features. These drives have been renamed to use a naming convention consistent with other Compaq Tape drives. Compaq has also added new capabilities that provide greater durability and reliability for servers with up to 8GB in size. The 2/4-GB DAT and 4/8-GB DAT drives incorporate a new enhanced internal tape head cleaning mechanism which extends the time between head cleanings, reduces errors in reading and writing tape, and extends the life of the tape drive itself. Even "dirty" tapes can now read more reliably. The 4/8-GB DAT drive offers an additional feature of a 25% faster data transfer rate.

TABLE 3: RELAUNCH OF THE COMPAQ DAT FAMILY

| 2/4-GB DAT Drive: | 4/8-GB DAT Drive: |
|---|---|
| Formerly the 2/8-GB DAT Drive | Formerly the 4/16-GB DAT Drive |
| Reduced cleaning requirements | Reduced cleaning requirements |
| Protects servers up to 4 GB | 25% faster - 1.8 to 3.6 GB/hr. |
| 2 GB Native | Protects servers up to 8 GB |
| 4GB with 2-1 compression | 4 GB Native |
| 650 NB to 1.4 GB/hr | 8 GB with 2-1 compression |
| Internal only | Internal only |

# BACKUP PLAN

The most successful backup plan provides for some sort of daily backup. Choosing less frequent backups, then the frequency of backups should be proportional to the amount of work performed on the hard drive and the criticality of the data. This is a matter of experience and judgment. Use the following criteria to decide how often to back up the hard drive subsystem:

- Value and volatility of data
- Risk of losing data
- Reliability of backup hardware and software

## Tape Backup Plan

Establish a tape backup and tape storage plan to ensure timely, routine backups (incremental or differential and full backups). Keep tapes readily available and clearly labeled to ensure that when backed up files must be retrieved, the tapes can be obtained and the files retrieved efficiently.

The following formula can be used to establish tape backup performance expectations for a single-stream backup server using the 4-/16-Gigabyte TurboDAT Drive with data compression enabled.

$$\frac{\text{Data (in MB) on Local Server}}{29} + \frac{\text{Data (in MB) on Remote Server}}{14} = \text{Backup time in minutes}$$

In the formula, 29 represents the average amount of data per minute that can be backed up locally to tape. The value of 14 represents the average amount of data per minute that can be backed up to tape over the network.

For example, the amount of time needed to back up a local server with 3 gigabytes of disk space and a remote server with 1 gigabyte of disk space is calculated in the following manner:

$$\frac{3{,}000}{29} + \frac{1{,}000}{14} = 175 \text{ minutes}$$

In this example, the server requires 2 hours and 55 minutes to perform a backup.

If the tape backup storage capability is limited, use the following guidelines:

- If local server backup storage capacity is limited, increase the backup capacity by adding more tape backup devices or the Compaq TurboDAT Autoloader.
- Plan multilevel backups. Some volumes require a full daily backup, while some need only differential or incremental backup.
- Enforce storage management plans.

## Backup Manager

The most logical candidates for tape backup responsibility are users with ADMINISTRATOR-level access, such as system administrators and system operators. Select both a primary backup manager and a secondary backup manager to provide support services in the event of a server failure.

Ensure that the backup managers are provided adequate training in managing the network and backup as well as performing backup procedures. Also ensure that the backup manager remains familiar with the data restoration process.

## Scheduled Backups

Schedule backups for periods when network traffic is low and server activity is minimal. Performing remote backups adds traffic to the network, potentially affecting network performance.

## Tape Rotation Scheme

A full daily backup of all files in the system may seem like a guarantee of data integrity, but it can also be time-consuming and is often impractical for high-volume hard disk activity. One alternative to daily full backups is a weekly full backup with differential or incremental daily backups.

## Tape Cartridge Management

Properly label each cartridge and tape with its name for easy retrieval and for enforcement of the rotation scheme. When storing tape cartidges on site, place them in a secured location, such as a locked room.

### Disaster Recovery

A reliable tape backup plan allows retrieval of unintentionally damaged or deleted files. To provide this kind of service to network users, incorporate an off-site rotation and retrieval process and store a copy of the most recent weekly or monthly full backup on site.

*IMPORTANT!*
*Periodically restore tape data to a server hard drive to check backup hardware components. This also allows the administrator to become familiar with the restoration procedure. When data must be restored, time is usually critical.*

### Tape Backup Hardware Redundancy

Data is usually restored as soon as possible after data loss. With a regular backup routine, the data is secured on the tape. However, if the backup system has also failed, can the company afford to perform without the crucial data until a new tape backup system arrives?

A well-designed backup system should cover all aspects, including backup hardware redundancy. Providing hardware redundancy is expensive, but the time and data that are lost can potentially cost far more. A single network server may require two backup systems to provide better data security against data disaster. For multiple servers, each network server should include at least one backup subsystem, so that another one can be used if one subsystem fails.

### Security

Include backup security measures whenever possible. Data security can be compromised during the backup process, since tapes are small and easily moved. Assess the potential risk of stolen data, and minimize or eliminate its possibility. Use physical security measures to protect data security, such as keeping tapes in a secure area not accessible to unauthorized personnel.

## CHEYENNE ARCSERVE RAID OPTION

As the amount of data in the enterprise continues to grow, the window of opportunity to perform backup operations diminishes. The requirement for faster backup becomes a higher priority. Subsequently, it is necessary for unattended backup operations to be performed during off-peak hours. To ensure that the "lights out" backup is successful, it becomes more important to have a fault tolerant system in place. With the growing value of corporate data, it is critical to prevent data loss due to media failures. ARCserve for Windows NT and its RAID Option meet the high performance and fault tolerant needs of enterprises with these requirements.

Cheyenne has made available the RAID Option v1.0 for Windows NT. This is not included as part of the Cheyenne Protection Suite, however it is definitely recommended.

Major features include full RAID tolerance on a tape drive array. For example, Level 5 RAID protects the continuity of backup or restore jobs in the event that a tape drive fails during operation. In addition, if one of the tapes in the set is damaged after the job is complete, the remaining tapes in the RAID set will preserve the entire backup job.

The RAID Option supports RAID 0, 1, and 5. Because of its ability to send concurrent data streams across all drives in the RAID tape array, RAID optimizes throughput to deliver extremely high performance.

Used in conjunction with ARCserve for Windows NT, the ARCserve RAID Option delivers extremely high performance and provides tape fault tolerance. The RAID Option integrates multiple independent tape drives into an array, providing unprecedented throughput due to the concurrent streaming of data across all drives. Using parity striping, the ARCserve RAID Option protects the continuity of backup or restore jobs in the event that a tape drive fails during the operation. In addition, should one of the tapes in the set be damaged after the job is complete, the

remaining tapes in the RAID set will preserve the entire backup job. The RAID Option supports the most commonly used RAID, RAID 5, as well as RAID 0 (striping), and RAID 1 (mirroring).

### RAID 5 for High Performance and Fault Tolerance

Three or more tape drives can be configured as a RAID 5 array. The ARCserve RAID Option will stripe the data across the drives and will produce rotating parity redundancy on the tapes in the array. If a drive should fail during the operation, backup or restore will continue with the remaining drives in the array. If, during storage, one of the tapes in a RAID 5 set should be damaged or lost, then all data stored in the set will be recoverable from the other tapes in the set.

### RAID 1 for Fault Tolerance

For RAID 1 two tape drives can be configured as identical mirrors of each other. The RAID Option will deliver identical data to both drives, simultaneously. If a drive should fail during the operation, the data will continue to be delivered to the other drive. One set of tapes can be maintained on-site while the mirrored set can be sent to off-site storage for precautionary measures.

RAID 1 also provides an ideal method of sharing or transporting data via tape without time-consuming tape copy utilities.

### RAID 0 for High Performance

When one or more drives is defined in a RAID 0 array, RAID delivers concurrent data streams to the drives without generating parity. This striping creates the highest possible throughput, but does not provide fault tolerance. When one tape drive is defined, RAID can send multiple data streams, simultaneously, to that drive. RAID interleaves the streams at the block level and writes them to the tape. RAID 0 is an ideal solution when, optimum throughput, rather than fault tolerance, is the objective; several sources must be backed up to a single device with the highest possible performance; or data is being streamed to an extremely high-performance tape device.

All of ARCserve's features including centralized administration, flexible scheduling, advanced restore choices, device management, and MIS reports, are available when ARCserve is used with the RAID Option. Using the RAID Option can further enhance the performance of backup and restore in conjunction with the Image Option. This enables image backup with direct streaming of data from disks to a tape array, avoiding the file system bottleneck.

The following table lists the major features of ARCserve RAID and describes their benefits.

## TABLE 4: CHEYENNE RAID OPTION FEATURES

| Feature: | Benefits: |
|---|---|
| High Performance | RAID ability to send simultaneous data streams to the RAID array optimized throughput to achieve very high performance. |
| Multiple RAID Support | ARCserve's RAID Option will support multiple RAID arrays on one server. They may be a combination of Level 0, 1 or 5. |
| RAID Set Compatibility | RAID will recognize and read tapes made in any RAID level or in any size RAID array. |
| Flexible RAID Configuration | Drives can be easily configured to any RAID level. Once assigned to an array, the drive is fully dedicated to that array. |

*Continued*

## TABLE 4: CHEYENNE RAID OPTION FEATURES (CONTINUED)

| Feature: | Benefits: |
|---|---|
| RAID Device Management | RAID allows the devices to be managed as a single unit, but reports comprehensive device management information on each individual drive. RAID also delivers a range of device management features, including Erase, Format, Eject and Change Compression. |
| RAID Size | RAID supports up to 8 drives per array, with a minimum of 8 drives per server. |
| Multiple Host Adapter Support | Devices in RAID array can be attached to separate SCSI host adapters for even higher performance and greater fault tolerance. |

## Known Issues

(Documented in ARCserve RAID Option for Windows NT v1.0 Release Notes)

For performance reasons, the RAID Option changes the default tape block size of DLT drives from 16 KB to 64 KB. When two DLT drives are attached to the same SCSI host adapter, there is a significant difference in performance between using 16-KB block size and 64-KB block size.

If you wish to change the default tape block size back to 16 KB, you can do so by modifying settings in the Windows NT registry.

The settings are listed in the HKEY_LOCAL_MACHINE window under the following key:

**SOFTWARE\Cheyenne\ARCserve\CurrentVersion\TapeEngine**

1. Run REGEDT32.EXE from the \%winroot%\system32 subdirectory
2. From the HKEY_LOCAL_MACHINE Hive go to SOFTWARE \ Cheyenne subkey.
3. Select ARCserve \ CurrentVersion \ Tape Engine
4. Select the Device number that the tape drive is on
5. From the Menu Bar select Edit | Add Value
6. In the Value field type: DefaultBlockFactor
7. In the Data Type field choose: REG_DWORD, click OK.
8. Input one of the following values in the String field: (0 - 5 )
   0 = 512 (default)
   1 = 1024 (1K)
   2 = 2048 (2K)
   3 = 4096 (4K)
   4 = 8192 (8K)
   5 = 16384 (16K)

**DefaultBlockFactor —** Set DefaultBlockFactor:REG_DWORD:5 in the DLT device key to change the default block size to 16-KB.

- Sessions after a repaired session must be individually merged.
- Before performing overwrite-any-tape backup (either for first tape or span tape), make sure the tapes in the drives are erased or formatted with the current RAID device, or formatted with a RAID device of the same level and size.

## CHEYENNE PROTECTION SUITE TESTBED

The tests of the Cheyenne Protection Suite were performed on Compaq ProLiant 4500 rack mountable servers with one Compaq DLT 4 Drive Tape Array 10/20. Refer to Table 5 for complete list of the hardware and software used in testing.

### TABLE 5: COMPAQ LAB SETUP

| System Component: | Compaq Lab Setup: | Tech Notes: |
|---|---|---|
| Operating system | MS Windows NT Server 4.0 SP2 | |
| Lotus Notes | Lotus Notes 4.5 | Cheyenne ARCserve Backup Agent v2.0 does not support Lotus Domino partitioned servers. The major benefit lost is that of being able to backup open files on Domino partitioned server other than the Domino partitioned server whose server.id is in the data directory specified by the registry value of the subkey shown in Image 2. |
| Cheyenne ARCserve | ARCserve 6.0 SP2 for NT | ARCserve needs the latest Cheyenne service pack installed before installing the Backup Agent v2.0 for Lotus Notes. At time of test, service pack 2 was the latest available. There are some installation notes (sp3notes.wri) which can be found in the patches and upgrades area at **www.cheyenne.com**.<br><br>The Install Notes contain resolutions to possible problems that you may encounter during the install. |
| ARCserve host server | Compaq ProLiant4500R | The Compaq ProLiant 4500R listed here served as the host server for both the Cheyenne Protection Suite and the Lotus Notes application.<br><br>The server has 128 MB RAM and 2 P5 processors |
| Domino host server | Compaq ProLiant 4500R | The Compaq ProLiant 4500R listed here served as the host server for both the Cheyenne Protection Suite and the Lotus Notes application. |
| DLT Tape Array | Compaq DLT Tape Array Model 0 | **Note:** The host SCSI adapter used to connect the tape array should be found in the Cheyenne in the ARCserve 6.x for Windows NT Certified Device List and in the Microsoft Windows NT Hardware Compatibility List.<br><br>**Note:** Windows NT does not support EISA Fast and Wide SCSI Controller — Compaq Assembly No. 003529-001 Revision D. However, Windows NT does support the Fast and Wide Embedded SCSI controller. |
| DLT Tape Drives | 4 Compaq DLT 10/20 in DLT Tape Array | Tape drives were in DLT Tape Array 10/20-GB Model 0<br><br>DLT Tape Array Part No: 199860-001.<br><br>DLT Tape Cartridge Part No: 199702-001 |
| DLT Tape Array Host Adapter | Compaq Fast SCSI —2 Controller and Compaq Fast and Wide SCSI — 2 Embedded Controller | Fast SCSI — 2 Assembly No. 002682 |
| Data Drives | ProLiant Storage Unit | Two 4.3 GB Fast and Wide SCSIs configured w/RAID0 via Compaq SMART Array Controller. |

*Continued*

TABLE 7: COMPAQ LAB SETUP (CONTINUED)

| System Component: | Compaq Lab Setup: | Tech Notes: |
|---|---|---|
| Data Drive Host Adapter | Compaq SMART Array Controller | |
| Backup Agent for Lotus Notes Server | V2.0 | Released in late 3/97, v2.0 contains new features previously listed in section titled Cheyenne Backup Agent for Lotus Notes v2.0<br>Works in conjunction with Cheyenne Open Files Agent.<br>**Note**: You do not have to purchase the Cheyenne Open File Agent separately. |
| AntiVirus Agent for Lotus Notes Server | V2.0 | Works in conjunction w/Cheyenne InocuLAN 4 for Windows NT.<br>The AntiVirus Agent can be set to scan in Real Time, Immediate or Scheduled via InocuLAN. |
| InocuLAN 4 for NT | | In order for the AntiVirus Agent to work, you must have InocuLAN 4 Build 241 or higher. To find out what build level of InocuLAN 4 you have installed, open InocuLAN and click on the Help — About. You will see a splash screen appear with the build number in the title bar. |

## Test Methodology

Using the hardware and software components listed in Table 5, Cheyenne ARCserve was used to test the functionality and performance of the Cheyenne Protection Suite for Lotus Notes for Windows NT components.

All components were installed and the backups were performed unattended with the results noted from the Cheyenne ARCserve Backup logs for each job.

### Single Stream Backup

| Single Stream | |
|---|---|
| Avg. Throughput- MB/min | Source |
| 145 | Local Domino Server |
| 149 | Remote Domino Server |

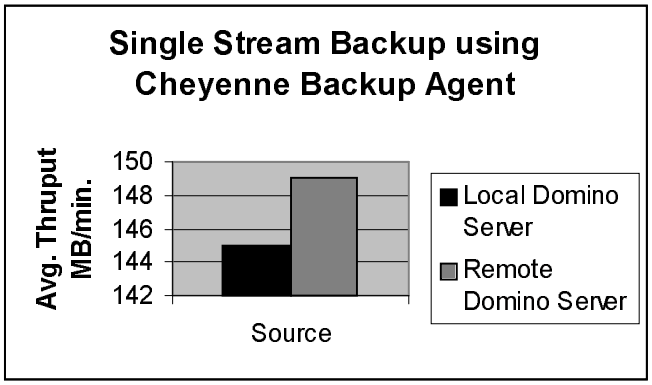

### RAID 0 Backup

| RAID 0 | |
|---|---|
| Avg. Throughput MB/min | Source |
| 338 | Local Domino Server |
| 340 | Remote Domino Server |

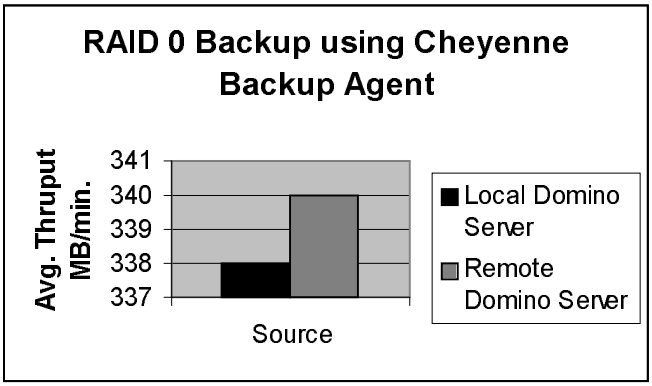


### RAID 1 Backup

| RAID 1 | |
|---|---|
| Avg. Throughput- MB/min | Source |
| 101 | Local Domino Server |
| 103 | Remote Domino Server |

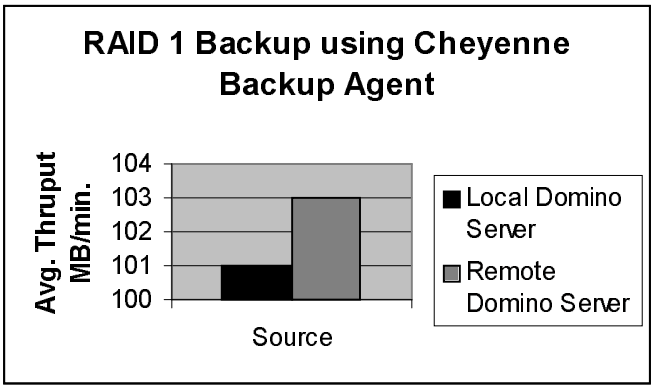

**RAID 5 Backup**

| RAID 5 | |
|---|---|
| Avg. Throughput- MB/min | Source |
| 251 | Local Domino Server |
| 254 | Remote Domino Server |

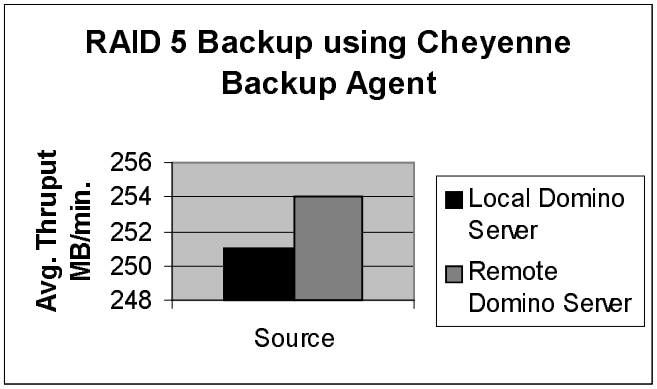


In each of the backups performed, the remote backup was slightly faster due to the configuration of the source remote machine. The source remote Compaq ProLiant 4500 had a tape array with no fault tolerance versus the mirroring in the source local Compaq ProLiant 4500. This was determined to be the cause when monitoring performance of the disk drives during the backups.

## Backup Agent Configuration

The Cheyenne Backup Agent for Lotus Notes v2.0 operates as a Windows NT service and can be configured to start automatically. This allows the Backup Agent to run without requiring a user to be logged in.

The user can customize the product by modifying settings in the Windows NT Registry. There are settings that can be set by using the Windows NT REGEDT32 utility. Backup Agent settings are listed in the HKEY_LOCAL_MACHINE window under the following key:

SOFTWARE\Cheyenne\DSAgent\CurrentVersion\agent\dbanotes

Several optional settings are provided for the Backup Agent under the dbanotes key as shown below:

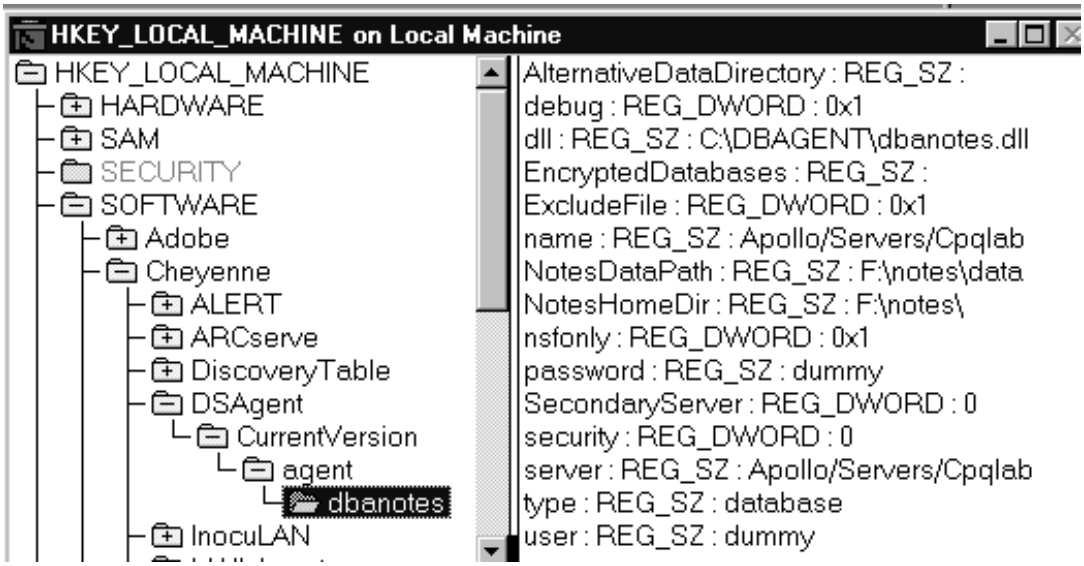


*Image 2: Registry tree of Lotus Notes Database Agent*

**AlternativeDataDirectory —** Specifies the desired location to restore your Lotus Notes databases if you do **not** want to restore your databases to the original directory

**Debug —** Set the debug value to *debug:REG_DWORD:1* (TRUE) to enable debug for the Backup Agent. Enable this when requested by Cheyenne's Technical Support. The default value is 0 (FALSE).

*NOTE:*
*Specify the path of the Notes data directory using the value NotesDataPath as shown in Image 1. If the Notes data directory is NOT in the NT path as well, the Backup Agent will not function. This error displays: Backup Agent RPC Service and the Backup Agent RPC service will shut down.*

**dll —** Specifies the location of the Backup Agent files.

**EncryptedDatabases —** Specifies the FULL path statements and filenames for any locally encrypted databases on the Notes server. If there is more than one, they should be separated by commas.

**nsfonly —** Specifies if only databases with the *.NSF extension will be backed up

**SecondaryServer** — If shared mail is enabled on the Notes server and the ARCserve Manager hangs when browsing the Notes server, this option should be set to 1.

## BACKING UP LOTUS NOTES DATABASE FILES

Performing a backup of the Lotus Notes database files consists of using the ARCserve Backup Manager. As with any type of backups via the ARCserve Backup Manager, select to perform an immediate backup or a scheduled backup of your Lotus Notes servers.

*NOTE:*
*For shops with more than one Notes server, the Cheyenne Backup Agent for Lotus Notes must be installed on each Notes server. It is not necessary to install Cheyenne ARCserve for NT on each Notes server. Cheyenne ARCserve will successfully backup a remote Lotus Notes server that has the Backup Agent loaded.*

By default, all databases will be displayed when selecting the Lotus Notes Server via the ARCserve browser in the Backup window. It is possible to specify that only Notes databases be displayed in the browser as noted in section 'Backup Agent Configuration'.

In addition, the name of the Lotus Notes object from the default 'Lotus Notes Server' can be changed to the actual name of the Notes server where the Backup Agent has been loaded. The

name change can be made in the Windows NT Registry for the value name in the subkey as shown in Figure 3.

The ARCserve browser will appear as shown in Figure 3.

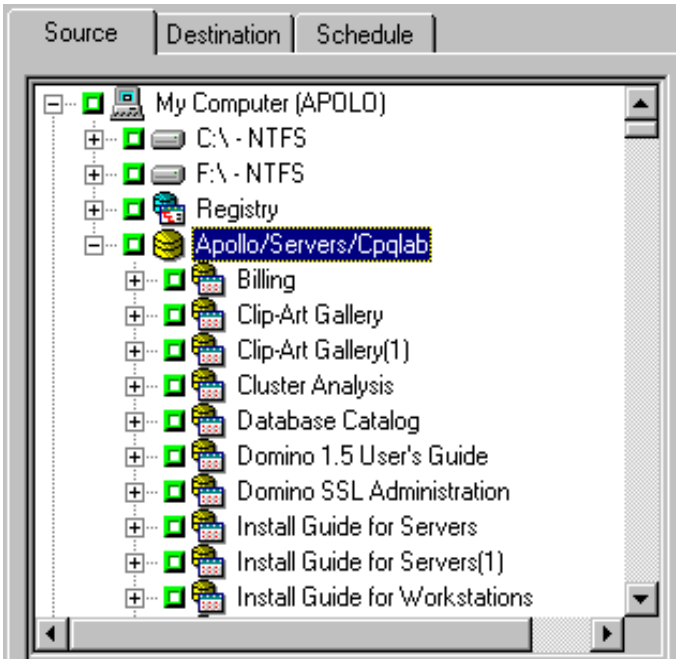


*Figure 3: ARCserve browser listing Notes server*

*NOTE:*
*Lotus Notes must be running on the server to be able to backup any databases.*

## Performance Measurements

To manage performance effectively, evaluate backup performance regularly to optimize the efficiency of your Lotus Notes server or servers. Although it is more common for backups to be performed overnight and therefore host server performance is not critical, overnight is the middle of the day for an office on the other side of the world. Therefore, for Lotus Notes mission critical databases that are prone to being in use hence open at any time during the 24 hours of a day, an optimized ARCserve host server will ensure those databases are backed up as quick as the hardware permits.

It is beyond the scope of this paper to discuss all the possible tuning parameters that can be manipulated to optimize your backup performance. However, as in any host server optimization attempt, there are key areas to observe:

- System processor
- Disk channel
- NIC channel
- Memory Utilization

Windows NT Performance Monitor is an excellent tool for observing the above system components.

As previously noted, Compaq has conducted performance tests with different disk subsystem configurations. In one test conducted, Compaq used the Completion port I/O test utility provided with the Compaq Resource Kit for Windows NT.

Basically, 2 configurations were tested. In one case, Compaq just operated an increasing number of disk drives on 1 SCSI channel of a SMART-2 controller. In the other case, Compaq distributed an increasing number of disk drives equally over the 2 SCSI channels of the SMART-2 controller. Since RAID 5 is a very cost effective fault tolerance mechanism, Compaq started out with RAID5 configurations. However as write accesses to a RAID5 volume naturally carry the burden of the additional read and write access for the parity generation, Compaq also looked at a few RAID1 configurations.

The following diagram depicts more details of the 1 and 2 SCSI channel configuration with the example of a 6 disk configuration.

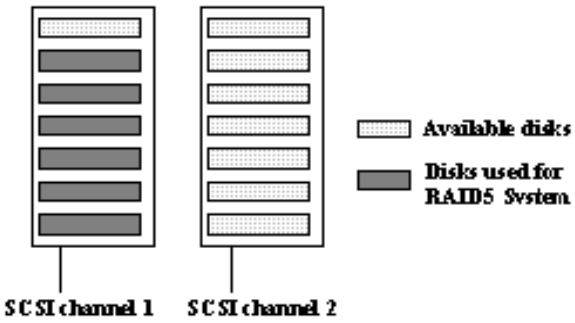


- 6 X 4.3 GB disks
- SMART-2 controller
- 6 disks connected to one SCSI channel

In the tests with two SCSI channels, Compaq distributed the disks equally to the two disk cabinets.

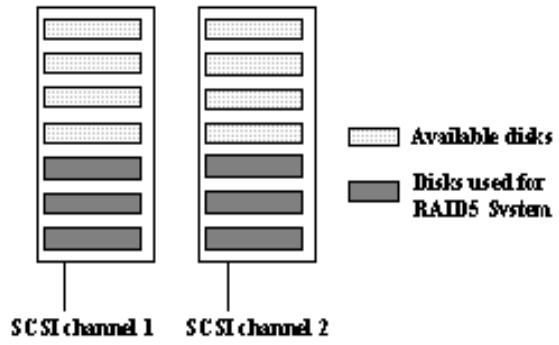


- 6 x 4.3 GB disks
- SMART-2 controller
- 3 disks at SCSI channel 1
- 3 disks at SCSI channel 2

Compaq did tests from three disks, which means the minimum for a RAID5 configuration was up to ten disks.

## TABLE 6: DISK SUBSYSTEM PERFORMANCE ON RAID 5 CONFIGURATIONS

| RAID 5 disk configuration | Read performance in GB/h | Write performance in GB/h |
|---|---|---|
| 3x4.3 GB disks on 1 SCSI channel | 45 | 12 |
| 3x4.3 GB disks on 2 SCSI channels | 51 | 12 |
| 4x4.3 GB disks on 1 SCSI channel | 46 | 15 |
| 4x4.3 GB disks on 2 SCSI channels | 73 | 15 |
| 5x4.3 GB disks on 1 SCSI channel | 48 | 16 |
| 5x4.3 GB disks on 2 SCSI channels | 72 | 17 |
| 6x4.3 GB disks on 1 SCSI channel | 49 | 17 |
| 6x4.3 GB disks on 2 SCSI channels | 83 | 17 |
| 7x4.3 GB disks on 2 SCSI channels | 81 | 18 |
| 8x4.3 GB disks on 2 SCSI channels | 91 | 19 |
| 9x4.3 GB disks on 2 SCSI channels | 86 | 19 |
| 10x4.3 GB disks on 2 SCSI channels | 94 | 19 |

To get the maximum performance, Compaq also did some tests with RAID0 and RAID1 configurations. It was concluded that in order to obtain the best writes performance, you have to set the Stripe Factor of the Compaq SMART Array controller to 32. For RAID5 configurations, the default setting is already 32. For RAID0 and RAID1 configurations the default value is 256. To change the Stripe Factor of the Compaq SMART-2 Array controller, the Compaq utility CONFIG.EXE was used. Table 7 shows the maximum numbers optimized for the write performance.

## TABLE 7: DISK SUBSYSTEM PERFORMANCE OF RAID0 AND RAID1 CONFIGURATIONS

| RAID0, RAID1 disk configuration Stripe Factor 32: | Read performance in GB/h: | Write performance in GB/h: |
|---|---|---|
| 12x4.3 GB RAID0 | 90 | 32 |
| 14x4.3 GB RAID0 | 99 | 33 |
| 12X4.3 GB RAID1 | 63 | 26 |
| 14X4.3 GB RAID1 | 74 | 26 |
| 7X4.3 GB RAID0 on 1 SCSI channel | 50 | 27 |

As you can see in Table 7, the best write performance that is important for the restore process is 33 GB/h with 14 disks in RAID0 configuration — no fault tolerance.

The fastest fault tolerance configuration that can be implemented is a RAID1 configuration with 14 disks. With the RAID1 configuration, Compaq got a read rate of 73 GB/h and a write rate of 26 GB/h. The last two numbers show that the overhead in the write process of a RAID1 configuration is not very high (26 GB/h vs. 27 GB/h).

In utilizing a SMART-2 Array controller with your hard disks, note that there is also a feature to change the cache settings of the Compaq SMART-2 Array controller. For all these tests, the SMART-2 Array controller cache was set to 50% read and 50% write. Changing the SMART-2 cache configuration did not improve the disk performance in the particular test environment as Table 8 shows.

TABLE 8: COMPARISON OF DIFFERENT CACHE CONFIGURATIONS

| Cache configuration: (10x4 GB disks on 2 SCSIs) | Read performance in GB/h: | Write performance in GB/h: |
|---|---|---|
| 50% read, 50% write | 93 | 19 |
| 0% read, 100% write | 93 | 19 |
| 100% read, 0% write | 94 | 13 |

## RESTORING A DATABASE

In a Notes environment, there are many databases which are critical to the operation of a Notes domain and in addition, anywhere where Notes is deployed, there will be critical workflow databases whether designed by in-house developers or third party developers.

Perhaps the most critical database of all, at least from an administration standpoint, is the Public Name and Address Book, also known as NAMES.NSF.

ARCserve provides Restore by Tree, Tape and Query. The selected items are restore file by file.

Using the Cheyenne Backup Agent for Lotus Notes, you can restore an important file such as the Public Name and Address Book to a pre-determined directory other than the original directory. You can select the directory to which you want your Notes database restored using REGEDT32 under the subkey *dbanotes* using the value *AlternativeDataDirectory* shown in Image 2. Of course, you can also choose to restore the file to its original location. The method you choose would depend entirely on the situation at hand.

If, for example, one of your Notes administrators inadvertently deletes a View from the Public Name and Address Book which gets replicated throughout your entire domain before the damage is discovered, but this View is not preventing your users from authenticating with the Notes server and accessing databases or mail files, then you would probably opt to restore the Public Name and Address Book to an alternative directory. This way you will save time on the restore process. Then all you would have to do is copy the restored NAMES.NSF to its original data directory at a later time, replicate it across your domain and you're back in business.

To restore a Lotus Notes database, do the following:

1. Select Restore from the Quick Access box.
   The Restore Manager appears. A part of it is shown in Figure 4.

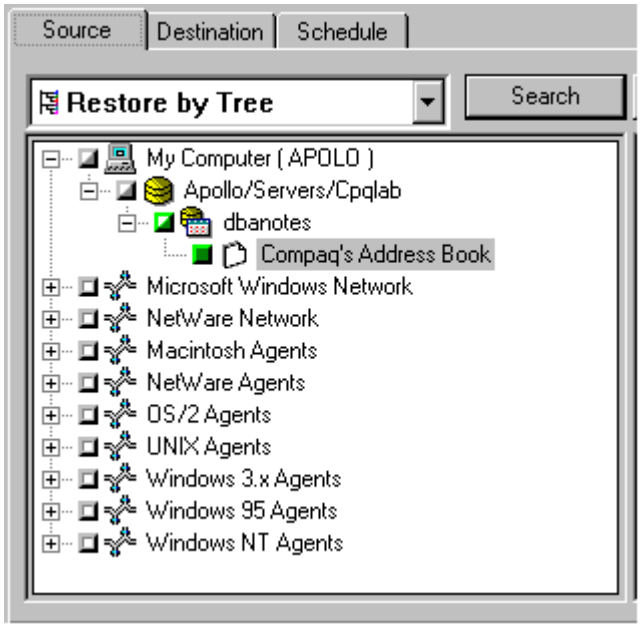


*Figure 4: Screenshot of the Restore Manager*

*NOTE:*
*This view allows you to view and select the individual databases you backed up on tape media.*

2. Select the Restore by Tree option.
3. From the Source screen, select the icon for the backed up database you want to restore, then select the individual database files in the expanded tree.
4. If you have previously backed up this database, click the Version History dialog box on the Restore Manager.

5. On the Destination tab, select the Restore files to their original location check box. You may optionally choose to deselect this check box as aforementioned and as shown in Figure 5.

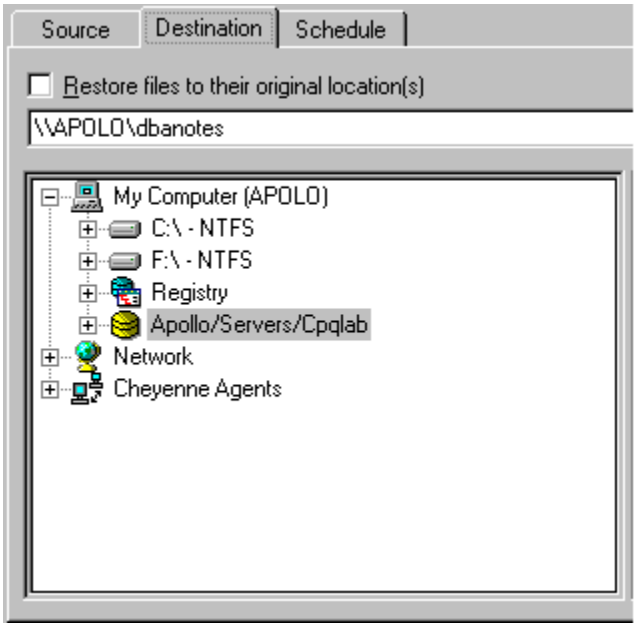


*Figure 5: Selecting Destination for Restore*

The following results were obtained using the configuration testbed shown in Table 5. Results will vary depending on the configuration you choose, whether it be RAID0, RAID1 or RAID 5 on your DLT tape array. In addition, you will have the option of choosing the number of host SCSI adapters to use with your DLT tape array.

**RAID 5 Restore**

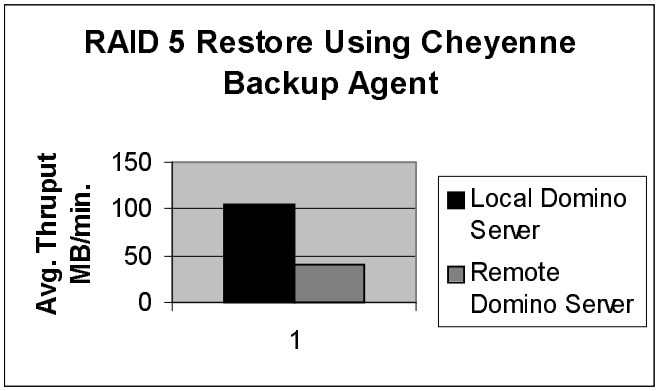

**RAID 1 Restore**

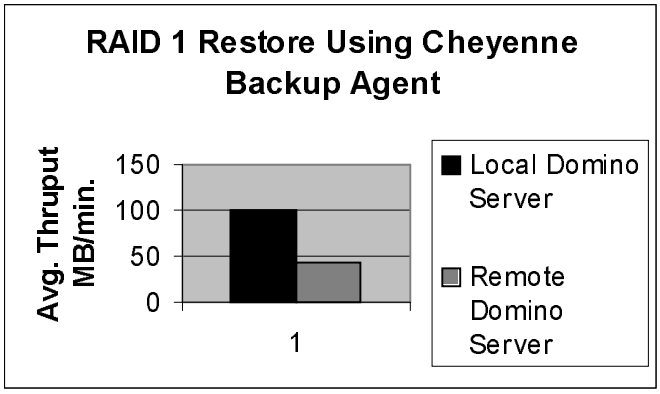


**RAID 0 Restore**

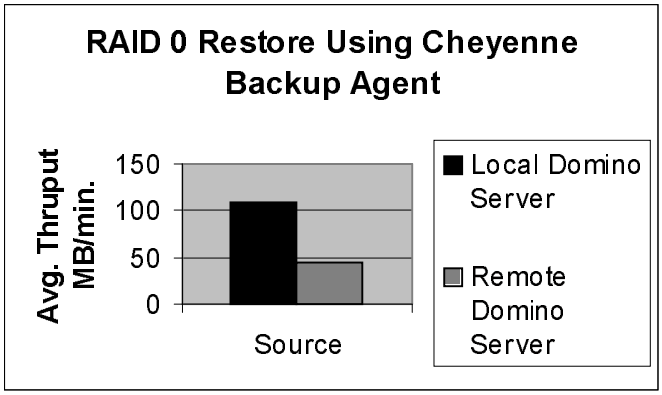

## SUMMARY

Cheyenne has delivered the Protection Suite for Lotus Notes for Windows NT. The goal of this package is to resolve two major concerns Notes administrators face: backing up open files and protecting their messaging system from viruses.

The backing up of open files was successful. Backups of all the Lotus Notes data directory .NSF files (Notes databases) were performed repeatedly with different tape array configurations.

- Compaq DLT Tape Array configured in RAID 0 using ARCserve RAID Option delivered the highest throughput logged at 340 MB/min. using the testbed documented in this paper.
- Compaq DLT Tape Array configured in RAID 1 via ARCserve RAID Option delivered the highest throughput logged at 109 MB/min. using the testbed configuration in this paper.

The restore rate is limited by the write performance of the disk subsystem, which is derived from your hard drive configuration. Disk channel, CPU channel, memory, network utilization and DLT tape array configurations affect the backup rate.

Cheyenne's RAID Option for ARCserve Windows NT Edition is a must if you are using a DLT tape array.

The hardware configuration will depend on the size of the business. Using the performance guidelines in this paper, Compaq customers should get a feel for what type of configuration they require.