



Compaq Backup and Anti-virus Solutions for Domino Server

TechNote

First Edition (January 1998)
Part Number ECG052/0198
Compaq Computer Corporation

Notice

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

COMPAQ COMPUTER CORPORATION SHALL NOT BE LIABLE FOR TECHNICAL OR EDITORIAL ERRORS OR OMISSIONS CONTAINED HEREIN, NOR FOR INCIDENTAL OR CONSEQUENTIAL DAMAGES RESULTING FROM THE FURNISHING, PERFORMANCE, OR USE OF THIS MATERIAL. THIS INFORMATION IS PROVIDED "AS IS" AND COMPAQ COMPUTER CORPORATION DISCLAIMS ANY WARRANTIES, EXPRESS, IMPLIED OR STATUTORY AND EXPRESSLY DISCLAIMS THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR PARTICULAR PURPOSE, GOOD TITLE AND AGAINST INFRINGEMENT.

This publication contains information protected by copyright. No part of this publication may be photocopied or reproduced in any form without prior written consent from Compaq Computer Corporation.

© 1998 Compaq Computer Corporation.
All rights reserved. Printed in the U.S.A.

The software described in this guide is furnished under a license agreement or nondisclosure agreement. The software may be used or copied only in accordance with the terms of the agreement.

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

Compaq, ProLiant, NetFlex, and ProSignia are registered with the United States Patent and Trademark Office.

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

Microsoft, MS-DOS, Windows, Windows 95, and Windows NT are registered trademarks of Microsoft Corporation.

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

Compaq Backup and Anti-virus Solutions for Domino Server

First Edition (January 1998)
Part Number ECG052/0198

Contents

Chapter 1

About this TechNote

Objective	1-2
Additional Resources	1-2

Chapter 2

Backup Basics

Overview	2-1
Why are Backups for Domino Data Essential?	2-1
Backup Terminology	2-2
Off-line and On-line	2-2
Complete and Partial Backups	2-2
Image and By File Backup Techniques	2-4
Backup Strategy	2-4
What Data Should be Backed Up?	2-5
How Much Data will be Backed Up?	2-5
What is the Backup Window?	2-6
How Frequently Should Backups Take Place?	2-6
How Should Backups Tapes be Managed?	2-7
Restore Strategy	2-7
Tips for Implementing a Domino Server Backup and Restore Strategy	2-8

Chapter 3

Backup Solution Software

ARCserve 6.5 for Windows NT Features	3-1
ARCserve 6.5 Restore Options	3-7

Chapter 4

Backup Solution Hardware

Compaq Enterprise Backup Solutions Provided by DLT Drive Technology	4-2
DLT Tape Format	4-3
DLT Tape Cartridges	4-3
DLT Interface	4-4
DLT Tape Library	4-5
DLT Device Drivers	4-6

Backup Solution Hardware *continued*

Compaq Backup Solutions for Small and Medium Size Business	4-6
Performance and Capacity	4-7
Hardware and Media Costs	4-8
Hardware Reliability	4-9
Compaq Small and Medium Backup Solution Tape Drive Alternatives	4-10
Compaq Small and Medium Business Backup Solution Summary	4-12
Compaq Backup Hardware Summary	4-12

Chapter 5

Compaq DLT Backup & Restore Performance Results

Local Backup Testing	5-2
Tape Drive and SCSI Controller Configuration Combinations	5-2
RAID Level	5-5
Data Compression	5-8
Processor Usage	5-12
Feed Speed	5-14
Data Set	5-14
Block Size	5-15
Other Performance Factors	5-15
Remote Backup Performance Testing	5-16
Concurrent Backup & Restore of Multiple Servers Across a Single 10Base-T Segment	5-17
Concurrent Backup of Multiple Servers over a Single 100Base-TX Segment	5-18
Remote Backup Conclusions	5-21
Backup Conclusions	5-22
Small and Medium Business	5-22
Enterprise Customers	5-22

Chapter 6

Virus Protection Security

Necessity of On-line Virus Scanning Utilities	6-1
Common Viruses	6-2
Traditional Anti-virus Utilities	6-2
CA Anti-virus Software Solution	6-3
Advanced Virus Protection	6-3
InocuLAN 4	6-4
AntiVirus Agent V2.0 for Lotus Notes	6-8
Protection Suite 2.0 for Lotus Notes	6-11
Summary of Anti-virus Protection	6-12

Index

TABLE OF FIGURES

FIGURE 3-1. ARCSERVE MANAGER'S CLASSIC QUICK ACCESS MENU	3-3
FIGURE 4-1. CAPACITY OF COMPAQ SMALL AND MEDIUM BUSINESS TAPE DRIVERS	4-8
FIGURE 4-2. 2/4-GB DAT BACKUP AND RESTORE PERFORMANCE	4-11
FIGURE 5-1. TAPE DRIVE AND SCSI CONTROLLER CONFIGURATION COMBINATIONS	5-4
FIGURE 5-2. 35/70-GB DLT BACKUP AND RESTORE PERFORMANCE COMPARISON USING RAID 0, 1, & 5 AND PROCESSOR UTILIZATION RATE	5-6
FIGURE 5-3. 10/20-GB DLT BACKUP AND RESTORE PERFORMANCE COMPARISON USING RAID 0, 1, & 5	5-7
FIGURE 5-4. BACKUP PERFORMANCE IMPACT OF DATA COMPRESSION RATIO	5-10
FIGURE 5-5. DATABASE PROPERTIES INFORMATION TAB PROVIDES % USED VALUE	5-11
FIGURE 5-6. COMPAQ 35/70 DLT PROCESSOR UTILIZATION FOR RAID 5 BACKUP TEST	5-12
FIGURE 5-7. COMPAQ 35/70 DLT PROCESSOR UTILIZATION FOR RAID 0 BACKUP TEST	5-13
FIGURE 5-8. CONCURRENT BACKUP & RESTORE OF MULTIPLE SERVERS ACROSS A SINGLE 10BASE-T SEGMENT	5-18
FIGURE 5-9. CONCURRENT BACKUP & RESTORE OF MULTIPLE SERVERS ACROSS A SINGLE 100BASE-TX SEGMENT	5-20

TABLE OF TABLES

TABLE 4-1 TAPE CARTRIDGE SUPPORT FOR THE COMPAQ DLT DRIVES	4-3
TABLE 4-2 TRANSFER RATES OF COMPAQ SMALL/MEDIUM BUSINESS TAPE BACKUP ALTERNATIVES	4-7
TABLE 4-3 COMPAQ TAPE DRIVES, NATIVE CAPACITY, DRIVE TRANSFER RATES, AND TIME REQUIRED FOR 10-GB BACKUP	4-13
TABLE 5-1 TEST SYSTEM CONFIGURATION	4-3
TABLE 5-2 TEST SYSTEM CONFIGURATION	4-5
TABLE 5-3 TEST SYSTEM CONFIGURATION	4-9
TABLE 5-4 TEST SYSTEM CONFIGURATION	4-17
TABLE 5-5 TEST SYSTEM CONFIGURATION	4-19
TABLE 6-1 INOCULAN FEATURES AND BENEFITS	4-6
TABLE 6-2 ANTIVIRUS AGENT V2.0 FOR LOTUS NOTES FEATURE & BENEFITS SUMMARY	4-10

Chapter 1

About this TechNote

This TechNote provides information that the customer, value-added reseller, system integrator, and/or network administrator can use to fully understand, size, and implement backup and restore solutions for the Compaq server family running Lotus Domino/Notes. The information contained in this paper enables the reader to implement a sound backup and anti-virus protection strategy within an organization. The following chapters are included as part of this backup and anti-virus protection solution:

- Backup Basics
- Backup Solution Software
- Backup Solution Hardware
- Compaq DLT Backup & Restore Performance Test Results
- Virus Protection Security

Compaq conducted all performance tests included in this paper. The information and recommendations presented in this paper are based on technical knowledge of Compaq performance engineers and the analysis of performance data relating to the testing performed in Compaq labs. Performance testing was conducted running Lotus Domino Server 4.6 and the Computer Associates (CA) and ARCserve 6.5 for Windows NT. The Backup Agent for Lotus Notes Version 2.0 and other options were included in the backup testing performed. InocuLAN and the AntiVirus Agent for Notes are examined as an anti-virus security solution.

This TechNote will help customers, value-added resellers, system integrators, and network administrators fully understand Compaq tape backup drive offerings. The solutions range from the DAT drive as a small and medium customer backup solution to the 35/70 DLT as an enterprise customer backup solution. This paper is a supplement to the *Compaq Hardware Reference* documentation, *the ARCserve*, *Backup Agent for Notes*, *InocuLAN*, and *AntiVirus Agent for Notes* documentation, and the *Lotus Domino Server 4.5* documentation.

Objective

The objective of this TechNote is to provide technical information regarding backup and anti-virus solutions for Domino Server 4.6 running on Compaq servers. The paper provides performance, capacity, and functionality information when using ARCserve 6.5 for Windows NT as well as discussing CA Backup Agent for Lotus Notes Version 2.0. CA Innoculan 4 for Windows NT and the AntiVirus Agent for Lotus Notes option was also tested. The data provided illustrates the performance of different hardware configurations and backup scenarios. Customers can use this data to determine which backup hardware configuration will best suit their business needs considering price and performance information.

Additional Resources

Consult the following resources for additional information on obtaining the best possible performance and throughput with Lotus Domino Server for Windows NT:

- Lotus Domino 4.6 documentation

The Lotus Domino 4.6 documentation set provides a comprehensive set of documents covering installation, reference, and an administrator's guide with detailed information on Lotus Domino 4.5x.

- World Wide Web on the Internet

<http://www.compaq.com/support/techpubs/> for Compaq technical publications

<http://www.compaq.com/products/servers/platforms.html> for Compaq system information on Compaq server offerings

<http://www.compaq.com/internet.html> for Compaq Internet and Intranet solutions

<http://www.cai.com/cheyenne> for Computer Associates backup and virus protection product information

<http://www.lotus.com> for Domino product information

- Compaq Hardware Reference documentation

The Compaq Hardware Reference documentation provides information similar to that available on the website for all Compaq server and option offerings

NOTE: This list of available materials is not intended to be all-inclusive, but access to these materials will be of benefit to the reader.

Chapter 2

Backup Basics

Overview

Data backup and restore capabilities are two of the most important aspects of business-critical application servers such as the Lotus Domino Server running on Compaq servers. This document will present and analyze various backup and recovery solutions available for a Windows NT and Domino Server, from both the hardware and software perspective, and provide choices based on performance, cost, capacity and functionality. Based on the information presented, the customer will be able to make an informed decision when selecting a Compaq tape drive for protection of data.

Why are Backups for Domino Data Essential?

The backup process copies important information (in many companies, this is vitally important information) onto magnetic tape or other disks. This enables the restoration of anything, should the need arise, from one file to the entire system. Backups have helped companies recover from data losses caused by power surges and outages, static electricity, lightning strikes, simple accidents (such as a spilled cup of coffee), sabotage, equipment malfunctions, viruses, and so on. Data recovery tools and services exist, but they are limited and can be expensive. While users might be able to recreate some lost data, retrieving all of the lost information is unlikely. Complex application and network configurations, customized setups, even passwords and IDs will be difficult and expensive, perhaps even impossible, to recreate.

The sudden loss of a mission-critical server that stores and maintains corporate records and data (one of a company's most valuable assets) can be financially disastrous. In most companies, just the down-time before recovery can be much too costly. A well-designed backup system safeguards crucial information, providing the most efficient and cost-effective insurance against a potentially disastrous loss of data, time, and money.

2-2 Backup Basics

Organizations depend highly on the messaging and workflow automation provided by the Compaq ProSignia 200 with LISP. Many organizations store gigabytes of e-mail messages, document databases and mission-critical applications on a server or even 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 as well as anti-virus protection is vital to a successful implementation. A well thought-out backup strategy can reduce lost productivity due to hardware or software failure. A proven restore strategy is also important to ensure successful restores when required. Backup management will be thoroughly discussed in Chapters 2 through 5, while anti-virus protection is addressed in Chapter 6.

Backup Terminology

Backups are classified by the status of the network server or servers (off-line or on-line) when the backup takes place, and by the amount of information that is backed up (complete or partial backups), and by the backup technique that is used (image or by file).

Off-line and On-line

For an off-line backup, the system administrator's first step is to take the server off-line, making it unavailable to users for the duration of the backup operation. The typical off-line backup takes place when user activity is at its lowest.

An on-line backup takes place with the server on-line and available to users. Depending on the network architecture, users may see network performance degradation while an on-line backup is taking place. In addition, there can be a danger to data integrity caused by file contention. Selecting the appropriate backup software can minimize this data integrity danger.

Complete and Partial Backups

Complete and partial are terms used to describe the amount of information that is copied during a backup. A complete backup is a full backup of the entire server or PC client hard drive. For a server, this includes all volumes, directories, and files. For a PC client, this includes all drives, directories, and files. A partial backup can be one of the types listed below. All backups, whether complete or partial, can be done on-line or off-line.

There are many partial backup alternatives that differ in which files are included in the backup. Partial backup alternatives include:

- **Differential** - All files that were changed since the last complete backup are copied. Differential backups are useful when it is important to have the latest version of each file. If the same tapes are used for consecutive differential backups, the newer versions of backed up files are often allowed to overwrite older versions of the same file on the tape. Typically, backup programs do not reset the file's archive bit after a differential backup; the archive bit remains turned on until the next complete backup.
- **Incremental** - All files that were changed since the last backup are copied, regardless of what kind of backup took place. This type of backup is used when each revision of a file must be maintained. If the same tapes are used for consecutive incremental backups, the newer versions of backed up files are not allowed to overwrite earlier versions. Rather, the newer files are usually appended to the backup medium. Typically, backup programs reset the archive bit following each incremental backup.
- **User-defined** - A user-defined set of files are copied during the backup. Often this is a special backup requested by a group of employees on a mission-critical project.

When choosing exactly which files will be included in the backup, there are several options. The information that is backed up can be:

- **All Applications** - This type of backup saves all files in the area defined by the user, including settings, customizations, passwords, etc. Application backups are particularly useful after a major change or upgrade in software.
- **Applications and Data** - This type of backup creates a standalone copy of the user's information base. Application and data backups allow easy restoration of the user organization's records. These backups can also be used to migrate information to another server.
- **Data Only** - This type of backup includes only data, which may be segregated by project or department, or which may include all information created within a certain timeframe, or both.

2-4 Backup Basics

Clearly, the amount of backed up information varies with the type of backup selected. This, in turn, directly affects the overall strategy in terms of capacity and transfer rate.

Image and By File Backup Techniques

Two types of backup techniques are generally available: image and by file. Both methods were included in Compaq's testing scenarios.

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.

Note that some image backup programs (i.e., ARCserve for Windows NT) have the ability to do a by file restore from an image backup, but using that functionality usually extends the restore operation time significantly.

Backup Strategy

Developing a successful company-wide backup strategy requires an understanding of the system's network architecture and the demands placed on the system by its users. A manager or an administrator can take a step toward developing a sound backup strategy by answering the following questions:

- What Data Should be Backed Up?
 - How Much Data Will be Backed Up?
 - What is the Backup Window?
 - How Frequently Should Backups Take Place?
 - How Should Backup Tapes be Managed?
-

Backing up your Lotus Domino data, as well as other important data on your system, should become part of a daily routine. When determining the precise backup strategy, your answers to these questions will help you build a sound backup strategy.

What Data Should be Backed Up?

All data that is important to the operation of a business should be backed up. Ultimately, all company data and programs should be backed up, so that the entire system can be restored in the event of a catastrophic disaster. In addition, certain groups of users may have special backup needs. For example, a group of key users responsible for developing mission-critical applications may require special backups because the data is so important to the proper operation of business. This extremely critical information is often scheduled for a special backup that takes place in addition to the regularly scheduled full-system backup. To develop an effective backup strategy, all corporate and user information needs must be identified and a backup plan determined. Domino Server databases need to be classified and a backup plan implemented according to the criticality of the data.

How Much Data will be Backed Up?

To answer this question, the following factors need to be taken into consideration.

- Determine the size of the complete system backup
- Determine the additional capacity required to handle redundant backups resulting from special, user-defined backups required for extremely critical data
- Project future data growth

Knowing the total amount of information to be backed up helps to determine the appropriate backup technology from two perspectives:

2-6 Backup Basics

1. **Storage Capacity** - The total amount of information to be backed up indicates the capacity required of the drive and the media. If planned backups will be unattended, then the selected backup device must have enough capacity to hold the full amount of information to be backed up. If the amount to be backed up requires more than one tape, then a backup tape drive with an autoloader may be appropriate. (The auto-loading feature uses a robotic mechanism to change tapes, thereby reducing administrative costs.)
2. **Performance** - The selected drive's typical backup rate together with the appropriate backup software (and several other factors) must make the backup system capable of accomplishing the entire backup during the time that the network is available and the server can be taken off-line (referred to as the backup window). Otherwise, off-line backups will not be possible.

Projecting the future data requirements is important. The amount of data stored on the typical server doubles every year for some companies. As imaging and multimedia applications become more common, the trend for increasing future storage demands will likely continue. A successful backup strategy will accommodate this growth. Proper planning can help ensure that the current backup solution can grow to meet next year's requirements as well.

What is the Backup Window?

Administrators typically perform backups when user demands on the server are at the lowest. Ideally, this time period, the backup window, is when user access can be restricted or the server shut down. As more and more companies move to 24-hours-per-day, 7-days-per-week operation, backup windows are shrinking. For many companies with worldwide operations accessing their servers, no clear backup window exists. The system administrator must determine how to get the backup done without impacting the productivity of users or seriously degrading network performance.

How Frequently Should Backups Take Place?

Backups must be performed regularly. The actual frequency of backups will be determined by considerations such as:

- the acceptable amount of work that could be lost, if any, in the event of a catastrophic failure
-

- the allowable down-time for recovery from this failure
- the volume of update transactions that normally take place

An effective backup strategy should also incorporate redundancy.

How Should Backups Tapes be Managed?

The Grandfather-Father-Son (GFS) tape-rotation scheme is the most commonly used and requires a weekly backup capacity of at least double the server's storage capacity. This rotation scheme uses three levels of backup to provide redundancy and security. Among other things, this scheme allows for different levels of data retention. The system administrator can choose which generation of tapes to store temporarily and which to archive.

The name GFS refers to these three levels of backup: the monthly "grandfathers," weekly "fathers," and the daily "sons." Typically, the system administrator performs a full backup every Monday (father) and does incremental backups on Tuesdays, Wednesdays, and Thursdays (sons). The administrator performs another full backup at the end of the week (father) and yet another at the end of the month (grandfather). The media containing the weekly and monthly backups are usually stored in a location away from the site of the server. To help reduce media costs, many companies reuse older weekly backup tapes.

Of course, the implementation of this GFS backup strategy varies from company to company. Many enterprise customers will develop a backup plan based on a 24-hours-per-day, 7-days-per week production schedule. However, a small district office may have a scaled-down plan based on Monday through Friday office hours. The enterprise will incorporate special, well-defined backup into the plan, while the small district office uses its scaled-down plan. This may involve performing the one daily backup Monday through Thursday, and a weekly backup on Friday.

Restore Strategy

A sound restore strategy is as important as a good backup strategy. The restore strategy should include the period testing of tapes to ensure that a successful restore is possible. There are often multiple methods of restoring files, depending on the backup and restore software selection. The restore methods typically include an option that controls whether the restored file replaces an existing file with the same name in the backup source directory, or whether the restore file may be placed in a specified directory different than the source.

Careful consideration should be made before implementing one of the available restore methods of your software. In a Lotus Domino environment, there are databases that are critical to the operation of the Intranet applications. In addition there may also be other critical databases designed by third-party developers.

Perhaps the most critical database of all from an administration standpoint is the Public Name and Address Book, also known as NAMES.NSF. Some software packages support restoring an important file such as the Public Name and Address Book to a pre-determined directory other than the original directory. Of course, this software would also support restoring the file to its original location. The restore method selected would depend entirely on the situation at hand.

For example, suppose the Notes administrator inadvertently deletes a view from the Public Name and Address Book that gets replicated throughout the entire domain before the damage is discovered. If this view is not preventing users from authenticating with the server and accessing databases or mail files, then the best option would be to restore the Public Name and Address Book to an alternative directory. This will save time on the restore process. All that is required then is to copy the restored NAMES.NSF to its original data directory at a later time, replicate it across the domain, and restore the file to its previous state.

Tips for Implementing a Domino Server Backup and Restore Strategy

When finalizing your backup and restore strategy, the following tips for implementing a sound backup and restore strategy may prove useful.

- After your backup system is in place, test it by restoring a file, the contents of a directory, and the contents of a drive.
- Remember that Lotus Domino requires that the following files be open when it runs: LOG.NSF, NAMES.NSF, MAIL.BOX, and the server ID file. Your backup solution must be able to address these open file issues.
- Store a recent set of backup files off-site in the event of fire, flood, or theft.
- Copy the server ID file to disk and store it in a secure area.
- Make a replica of the Public Name and Address Book on a workstation and keep it up-to-date by replicating the local replica with the server replica. Then if the Public Name and Address Book becomes corrupted, you can quickly restore it by creating a new replica from the local workstation replica. Even if you do this, continue to back up the Public Name and Address Book to tape.
- Make a replica of each production database, if possible. When you create a replica for this purpose, copy the entire database (access control list, views, forms, and documents) to the new replica. If problems occur in a database, you may be able to restore the database from a replica. Remember that you should never rely on replication as your sole method of database backup.

If your system uses a shared mail database, back up the shared mail database(s) along with user mail files.

Chapter 3

Backup Solution Software

One software product available today, ARCserve from Computer Associates (CA) Software, backs up Domino data files and other data whenever possible and enters an error in the log to notify the administrator that the file's protection is suspect. The administrator can then take the appropriate action, which may be to restore an earlier, non-suspect version of the file. Although often referred to as backup software, the same application supports restore operations as well.

Other products that protect open files during backup include ARCserve Database Agents, which are available from Compaq. The Backup Agent communicates with Lotus Domino and ARCserve to back up the open Lotus Domino databases on a local or remote server.

ARCserve 6.5 for Windows NT Features

ARCserve 6.5 Enterprise is a powerful storage management product that offers a rich set of features, especially when combined with its option software. This chapter also provides a detailed description of the ARCserve Manager, the Backup Agent for Lotus Notes Version 2.0, and ARCserve restore options. The following list describes some capabilities of the ARCserve products.

- ARCserve Manager – The user-friendly centralized menu-driven mechanism for controlling most operations.
- ARCserve Job Engine - Automation is made possible by the ARCserve Job Engine, which supports scheduling backup/restore jobs based on customized repeat methods (time interval, days of the week, etc.). In addition, a Tape Rotation scheme can be configured for the backup jobs.
- Database Agent - The Backup Agent for Lotus is available as an option. Otherwise, ARCserve must perform off-line backups of the database device files. The Agent functions as a Windows NT service, and allows backup at the database level.
- Information Repository - The ARCserve Database Engine maintains complete historical information on such things as jobs that have been completed, a record of files/directories/drives/machines that have been backed up, and the media that was used.

3-2 Backup Solution Software

- Tape Spanning - ARCserve 6.5 allows the creation of 'groups' of media, should multiple tape drives be available. A single backup job can continue uninterrupted across all the tape drives in that group. ARCserve supports up to eight tape drives per server.
- Parallel Streaming - If tape drives are separated into different groups, then multiple backup jobs can be run (one to each group) simultaneously.
- Fault Tolerance/Striping - This option provides a Tape RAID (or 'RAIT', with the 'T' implying *tape* rather than *drive*) system so that multiple tape drives can be placed in a group and written to as one 'tape array' by a single backup job. Fault tolerance can be provided to the array through RAID 1 (mirroring) or RAID 5 (striping with parity), so that the backup operation can continue if a tape drive fails, or so that a backup set can be recreated if a tape is lost. The array can also utilize RAID 0 striping, for a purely high-performance backup without fault tolerance.
- Centralized, Remote Management - The Enterprise version allows the management of multiple servers from a single machine, using a 'tree' view display. All storage management tasks for these servers, including submission of backup and restore jobs, can be performed from a single location.

ARCserve Manager

The ARCserve Manager is the CA program that is used for controlling most operations. After the startup of the ARCserve Manager, the Quick Access window displays as shown in Figure 3-1. This allows you to select the operation you want to start. You can choose from the following operations:

- Job Status – Global ARCserve servers' jobs and activities management
 - Backup – Custom and automated backup of networks
 - Restore – Complete data recovery of volumes, directories, and files
 - Data Migration – Manages data migration operations
 - Media Pool Management – Media pools and media configuration
 - Device Management – Real-time format, erase, and status display of backup devices
-

- Database – Manage and view the Database (Jobs, Media, etc.)
- Reports – View and print logs and reports

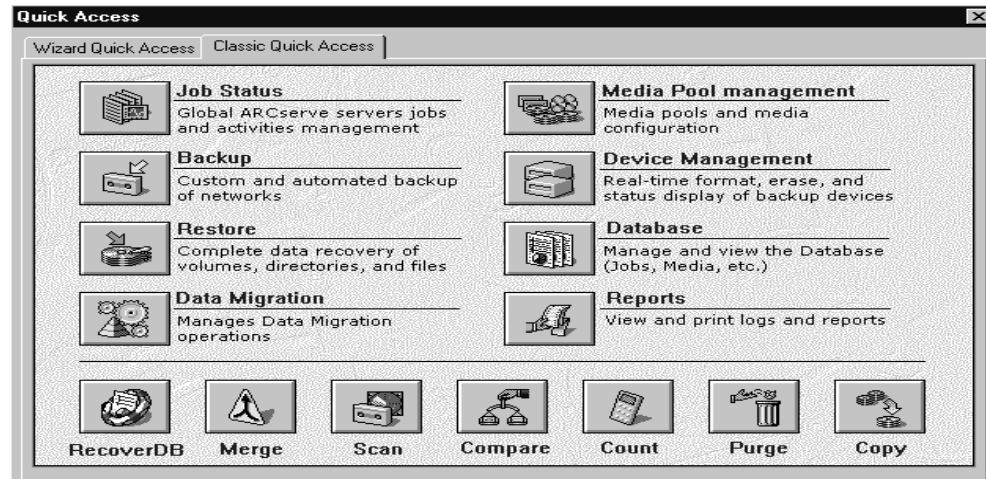


Figure 3-1. ARCserve Manager's Classic Quick Access Menu

Of the many ARCserve operations that can be started from this ARCserve Manager's Classic Quick Access menu, the tools that support ARCserve Tape Label Handling capabilities, such as Device Management and Media Pool Management, are very flexible and efficient. The Device Management and Media Pool Management options are described in the following section, as well as two tools that are not started from the Quick Access menu – ARCserve Engine Administrator and the ARCserve RAID Engine.

Device Management

If you have more than one tape device connected to your Windows NT machine, ARCserve lets you separate them into two or more "groups". Establishing tape groups is a key component of ARCserve's flexibility and efficiency. The Device Manager is very easy to use.

The tape groups enable you to perform the following options:

- **Parallel streaming** – Supports several tape operations occurring simultaneously, one at each tape group configured for your system.
- **Automated tape spanning** – Supports automatic spanning to another tape in another device within the same tape group if, during a backup session, ARCserve detects that the tape has become full.

3-4 Backup Solution Software

By default, ARCserve is installed with each tape device assigned to its own group. ARCserve finds identical tape drives during the installation, placing them automatically in the same group. Later, you can use the Device Manager to regroup your devices. Similar to the limitation of no more than seven tape devices connected to each SCSI bus in your machine, ARCserve supports up to seven tape groups per SCSI board.

There is only one rule when assigning tape devices to the same group; the devices in the group must be identical. That means they must be the same make and model.

The Device Manager is also used for formatting tapes. By formatting a tape you can specify a tape name, which represents a kind of tape label when combined together with the tape ID. The tape ID is four digits chosen by ARCserve.

Media Pool Management

A media pool is a collection of media managed as a unit. Tape rotation or Grandfather-Father-Son (GFS) methods of backup are based on media pools. Each media pool is assigned a name, and the media is organized according to serial numbers. The serial numbers assigned are permanent. If you are using a device with a bar code reader, the bar code labels will be used for the serial number of the media. The media pools are organized by the range of serial numbers of the media they contain and are divided into two sets, the Save Set and the Scratch Set.

- **Save Set** - The set of media containing important data that cannot be overwritten.
- **Scratch Set** - The collection of media from the Save Set that has met criteria such as minimum number of media to save and specific retention period. Therefore, this media can be re-used and overwritten, and is placed in the Scratch Set.

Each time a media is written to, it moves from the Scratch Set to the Save Set. That media will move back to the Scratch set once the specified criteria have been met.

Media pools apply to every media, regardless of the selected backup type and method. All rotation backup jobs will create their own media pool based on the name entered in the Media Pool Name field of the Schedule dialog box for simple rotation, or the prefix entered in the Media Pool Name Prefix field with GFS enabled. The GFS backup uses three media pools: Daily, Weekly and Monthly.

These three pools are also based on the information entered in the Media Pool Name Prefix field. For example, if you enter ACCT as the Media Pool Name Prefix for a GFS backup job, the Daily media pool name would be ACCT_DLY. For a simple (single media pool) rotation, you specify the complete name for the media pool.

ARCserve Engine Administrator

The ARCserve Engine Administrator can be found in the ARCserve Windows NT group. With the Engine Administrator, different ARCserve engines (Job, Tape, and Database Engine) can be started and stopped. In the Engine Administrator applet, you can also see the status and some more data of these engines. Normally you do not need this utility because you can start and stop the engines with the ARCserve Manager and the Windows NT Service Manager.

ARCserve RAID Engine

The ARCserve RAID Engine creates one or more logical tape drives from a group of physical tape drives by applying one of the RAID algorithms that are well known from hard disk arrays. The same kind of tape devices can be grouped into a logical tape drive. It is not possible to mix different drives such as DLT-tapes with DAT-tapes. However, a group of DLT-tapes and a group of DAT-tapes can be defined. The first implementation of RAID allowed for up to eight physical tape drives to belong to one logical tape drive. RAID V2 supports up to 32 tape drive devices in one logical tape drive. Keep in mind that no more than two DLT drives should ever be connected to one SCSI controller for performance reasons. In the test environment, Compaq engineers used up to four 32-Bit Fast-Wide SCSI-2 Controllers for the eight DLT drives test scenarios included in the backup performance chapter.

3-6 Backup Solution Software

ARCserve provides the capability to combine multiple physical tape drives into one physical RAID array. Data is distributed equally over all tapes of the RAID system. Since this increases the probability of failure (if one tape of the RAID system is damaged, the entire backup is lost), additional mechanisms for data safety are included in the RAID technology. There are three RAID levels supporting data integrity and varying degrees of fault tolerance support. The three RAID levels are:

- **RAID 0** – Only data is distributed over all tapes of the RAID system. It does not provide any features for error correction but provides a higher data throughput than single devices. 100% of the tape net capacity can be used.
- **RAID 1** – The mirroring technique is used to reduce the probability of data loss in case of a damaged tape. Data is written on two tapes in parallel. For a RAID 1 tape array, only two tape devices are needed. The disadvantage of RAID 1 is that only 50% of the net tape capacity can be used for the backup. The other 50% is used for fault tolerance mirroring support.
- **RAID 5** - Checksums (parity information) are generated for the data to be backed up, and the checksums, as well as data, are distributed over all tapes of the RAID system. If one tape is damaged, data can be restored from the remaining tapes. This fault tolerance method requires less tapes per usable capacity than RAID 1. A RAID 5 array must include at least three tape drives. RAID 5 needs one additional tape, i.e. net capacity plus one tapearray to support the parity and data being distributed across all tapes in an array.

Backup Agent for Lotus Notes Version 2.0

The Backup Agent communicates with Lotus Notes and ARCserve to back up open Lotus Notes databases on a local or remote server. The Backup Agent for Lotus Notes Version 2.0, an option that uses open files technology, specifically provides support for backing up open Lotus Notes databases.

A list of features of the Backup Agent for Lotus Notes Version 2.0 follows:

- 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 Domino Server off-line to backup the database. The Backup Agent Version 2.0 works in conjunction with ARCserve Backup Agent for Open File Version 3.1 for Windows NT.
- Browse and backup of encrypted databases use the same procedures as standard database files. The data is backed up and restored as encrypted files.
- The restore to alternate directory feature permits an administrator to specify a destination for restored files. This capability is useful when testing restore procedures, moving files between servers, or changing system configurations.
- Backup and restore of shared mail in the shared mailbox is supported instead of multiple individual copies.

A list of known limitations of the Backup Agent for Lotus Notes Version 2.0 can be obtained through Computer Associates.

ARCserve 6.5 Restore Options

ARCserve provides you with five methods for selecting the data you want to restore (the source) which are:

- Restore by Tree
- Restore by Tape
- Restore by Query
- Restore by Tape Media
- Restore by Image

The first three methods use information logged in ARCserve’s database, while the fourth method does not. In addition, if you have purchased separately the Image Backup Option, Restore by Image will also be available. You also have several methods for selecting a destination when restoring data.

Restore by Tree

The Restore by Tree view allows you restore a specific directory or drive. This view displays a machine tree of files and directories that were backed up with ARCserve. It is also possible to restore individual files from an Image Backup.

Restore by Tape

Restore by Tape allows you to select the tape, then the session, and finally, the files and directories you want to restore. This source view is useful if you know the tape that holds the data you want to restore, but you are not sure of the session location.

Restore by Query

The Restore by Query source is used to restore files when you already know the name of the file or directory you want to restore, but you do not know the machine it was backed up from or the tape(s) it was backed up to. This view also uses the ARCserve database. As long as the file and directory information are logged in the database, you can effectively use this view to select files to restore.

Restore by Tape Media

The Restore by Tape Media allows you to restore complete sessions on a tape to a tape device. All files in this session are restored to the destination, unless filters are added to the job. This source is useful in case information from your ARCserve database is unavailable, or in case you try to recover your network after a disaster.

Restore by Image

Image mode restore ("Restore by Image" mode in ARCserve Restore Manager) overwrites the target volume data.

Chapter 4

Backup Solution Hardware

Compaq offers a line of tape drives that include the Compaq DLT (Digital Linear Technology) family of tape drives, SLR (Single-channel Linear Recording) drives, DAT (Digital Audio Tape) drives, and supports other older tape technologies, such as the QIC (Quarter Inch Cartridge). The four tape drive technologies are described below:

- **DLT** (Digital Linear Tape) drives use simultaneous, multichannel/multihead read/write technology to achieve capacities up to 35-GB without compression. A DLT drive is the appropriate high-end backup solution for systems with 35- to 70-GB of data.
- **SLR** (Single-channel Linear Recording) The Compaq SLR Tape Drive features backward compatibility to older tapes and growth to future products ranging from 125-MB to 32-GB, making the Compaq 4/8 GB SLR Tape Drive a good growth choice.
- **DAT** (Digital Audio Tape) drives are usually the appropriate choice for servers with 2- to 8-GB capacity. Standard DDS1 DAT tape drives (2/8-GB DAT) can store approximately 2-GB without compression. DDS-2 DAT (4/16-GB Turbo-DAT) tapes can store 4-GB without compression. An Autoloader for use with DAT tapes is available from Compaq. Autoloaders reduce administrative costs by using a robotic mechanism to load and unload tapes.
- **QIC** (Quarter Inch Cartridge) drives meet the half-height form factor of desktop computers. QIC tapes are virtually industry-standard for standalone machines. However, with capacities limited to 1.2-GB, they are not generally suitable for backing up servers with storage capacities equal to or greater than 2-GB of storage capacity.

4-2 Backup Solution Hardware

In Chapter 5 of this paper, the performance of the Compaq DLT family of tape drives, which includes the 10/20-GB DLT, 15/30-GB DLT, and 35/70-GB DLT tape drives, is evaluated. The DLT family of tape drives is positioned to provide a backup solution to the larger corporations and enterprise environments. The performance of the other tape technologies such as SLR (Single-channel Linear Recording) and DAT (Digital Audio Tape) drives are examined in this chapter since the configuration options are more limited than the DLT tape drive options. The SLR and DAT drives are positioned to provide a backup solution for small and medium size businesses. DLT performance will be discussed separately in Chapter 5, "Compaq DLT Backup and Restore Performance Test Results".

Compaq Enterprise Backup Solutions Provided by DLT Drive Technology

The Compaq Digital Linear Tape (DLT) drives were designed to meet the growing backup needs of large computing environments, providing greater performance, capacity, and reliability from a PC server tape solution than ever before. The DLT ensures data integrity through multiple checks including: error-correction-code (ECC) on every 64-KB of tape data, cyclical-redundancy-check (CRC) on each 4-KB of tape data, end-to-end CRC on each record overlapped with parity from the SCSI bus, and internal parity checking on the 2-MB memory buffer. DLT drives record and read data using multiple, parallel tracks in a serpentine (alternating end-to-end) fashion. The DLT tape media supports a 128-track format, and DLT cartridges are specified to last 500,000 passes over the read/write head, resulting in the use of the media equaling up to 7,812 backups using the entire tape capacity. The head life of the drive itself is estimated at 30,000 hours under continuous operation.

DLT Tape Format

All Compaq DLT drives provide built-in hardware data compression using the DLZ (Digital Lempel-Ziv) algorithm. The 10/20-GB DLT drive supports 3 density formats, the most efficient of which yields 10-GB of storage space uncompressed and 20-GB with 2:1 compression. The 15/30-GB DLT drive supports 4 density formats, the most efficient of which yields 15-GB of storage space uncompressed and 30-GB with 2:1 compression. The 15/30 supports a superset of the 10/20 formats, so that it can read any tape recorded by a 10/20; but the 10/20 cannot read a tape that was recorded on the 15/30 using the highest (15-GB) format. The 35/70-GB DLT drive supports a superset of the other drives' formats, and offers a native format providing 35-GB of capacity, or 70-GB with 2:1 compression.

The 10/20-GB and 15/30-GB DLT drives do not differ in performance, and are both capable of 4.5-GB/hr (1.25-MB/sec) throughput with data compression disabled (up to 9.0-GB/hr with 2:1 compression). The 35/70-GB DLT drive however, offers not only enhanced capacity but also enhanced performance, being capable of up to 18.0-GB/hr (5.0-MB/sec) throughput with data compression disabled (up to 36.0-GB/hr with 2:1 compression). These high-performance, high-capacity and high-duty-cycle backups make the DLT ideal for storage of large, mission critical databases like those found in many Lotus Notes Server environments.

DLT Tape Cartridges

The table below summarizes the tape cartridge support of the Compaq DLT Tape Drives, which were used in the performance evaluation discussed in Chapter 5 of this document.

Table 4-1
Tape Cartridge Support for the Compaq DLT Drives

Tape Drive	Tape Cartridge(s) Supported
Compaq 10/20-GB DLT	DLT Tape III (10.0-GB)
Compaq 15/30-GB DLT	DLT Tape III (10.0-GB) DLT Tape IIIXT (15.0-GB)

continued

Tape Cartridge Support for the Compaq DLT Drives *continued*

Compaq 35/70-GB DLT	DLT Tape III (10.0-GB) DLT Tape IIIXT (15.0-GB) DLT Tape IV (35.0-GB)
---------------------	---

Referring to the table, the correct selection of tape media is critical to being able to support a specific density of data on tape. The DLT Tape IV is required to write the maximum density of data to tape.

DLT Interface

The Compaq 10/20 and 15/30 DLT drives are available with either an internal or external standard SCSI (8-bit) interface. These drives can, however, be connected to a Fast-Wide SCSI-2 controller with the appropriate cable, although half the bandwidth of the Wide SCSI-2 bus will remain unused. EISA systems such as the retired ProLiant 4500, using a Fast SCSI-2 or a Fast-Wide SCSI-2/E controller, will yield the same performance. For best performance on all Pentium Pro based PCI systems, Compaq recommends that the Compaq Wide-Ultra SCSI controller be used with these drives. The reasons for the recommendation of Wide-Ultra SCSI controller pertains to PCI bus utilization issues which can prove significant in terms of system throughput, especially when using multiple such controllers, not to the enhanced bandwidth of the Wide-Ultra SCSI bus. However the Fast-Wide SCSI-2 controller embedded on the ProLiant system board, can be used in conjunction with additional Wide-Ultra controllers without degrading performance.

The Compaq 35/70 DLT drive is available with either an internal or external Wide SCSI-2 (16-bit) interface. Like the other DLT drives, the 35/70 should be used with the Compaq Wide-Ultra SCSI controllers on Pentium Pro based ProLiant systems if peak performance is desired.

Multiple drives can exist on the same SCSI bus by daisy-chaining them with the appropriate cables. Compaq also provides DLT Tape Array cabinets, which support up to four internal DLT tape drives. The cabinets have two SCSI channels so that the drives can be attached to two separate SCSI controllers (two drives each). The drives can optionally be attached individually to four controllers using additional cables. For best performance, Compaq recommends not attaching more than two DLT drives per controller.

NOTE: The DLT does not interface to the Compaq SMART SCSI-2 Array controllers, which are intended only for disk drive connections

DLT Tape Library

The highly reliable, low-cost rack-mountable DLT Tape Library solutions provide customers with significant enhancements in backup capacity and automation. For unattended backup and restore functionality Compaq offers the DLT 15-Cartridge Library. The DLT 15-Cartridge Library supports up to fifteen 15/30-GB or 35/70-GB cartridges (3 magazines with 5 cartridges in each) for a storage capacity of up to 450-GB with the 15/30 drive or 1050-GB with the 35/70 drive.

The addition of an optional, second DLT Library Ready drive for either Library allows independent, simultaneous backup of separate databases and also provides On-Line Repair functionality which allows the second DLT drive to take over the backup function, in the event the first drive fails. The failed drive can then be replaced without taking the machine down.

The DLT 15-Cartridge Library has a rack-mountable chassis (8U height) that easily integrates into a 19" industry-standard rack. A conversion Kit is also available allowing the rack tape library to be converted to a table-top configuration.

The Library features a robotic arm for automatic tape changing, as well as an integrated bar code reader to uniquely identify each cartridge and create a complete catalog of all archived data - in conjunction with supporting software. Software that supports the DLT Library includes backup products such as ARCserve 6.5 for Windows NT and its tape changer option.

The requirement for faster backup becomes a higher priority as the amount of data in the enterprise continues to grow while the backup window diminishes. Thus the robotic arm meets the requirement for unattended backup operations, which can now be performed during off-peak hours. To ensure that a 7-days-per-week, 24-hours per-day operation's "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.

4-6 Backup Solution Hardware

Used in conjunction with ARCserve for Windows NT, the ARCserve RAID Option delivers high performance and provides tape fault tolerance. The RAID Option supports an array of like tape drives that behave as a single, high performance drive. This option delivers high performance by enabling ARCserve to stripe data across multiple drives, drives within a tape library, or across tape libraries. Thus 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 level, RAID Level 5, as well as RAID Levels 0 (striping), and 1 (mirroring).

DLT Device Drivers

When operating in a Microsoft Windows NT 3.51 or 4.0 environment, the Compaq 10/20-GB and 15/30-GB DLT drives use the Digital TZ86, TZ87, DLT2000 driver (DLTTAPE.SYS) that ships with the standard Windows NT product. This driver is not available on the Compaq Software Support Diskette for Windows NT (NT SSD). The Compaq 35/70-GB DLT drive uses a separate driver (DLT7000.SYS) that is available from Quantum Corporation. Some third-party backup applications (such as ARCserve for Windows NT) provide their own driver for the DLT family of tape drives. Thus when using one of these applications, the Windows NT driver must be removed or disabled. If the DLT is interfaced to a 32-bit Compaq SCSI-2 controller such as the integrated controllers found on ProLiant and ProSignia servers, the Compaq 32-Bit SCSI-2 Controllers' driver (CPQ32FS2.SYS) is installed from the Compaq Windows NT SSD. Lastly, the Quantum driver, *qntmdlt.sys*, available of the Quantum website, works with the Compaq 10/20, 15/30, and the 35/70 DLT drives.

Compaq Backup Solutions for Small and Medium Size Business

Compaq offers a line of tape drives that include the Compaq DLT family of tape drives, DAT drives, SLR drives, and supports other older tape technologies such, as the QIC.

The decision for selecting the appropriate hardware should be carefully thought out. After determining the company's backup needs, the system administrator determines the specifications for an appropriate, cost-effective backup solution that will best meet those needs in terms of:

- Performance and Capacity
- Hardware and Media Cost
- Hardware Reliability
- Compaq Small and Medium Business Tape Drive Alternatives

Performance and Capacity

Performance is measured by dividing the amount of information (in gigabytes) that must be backed up by the length of the backup window (in hours). This simple calculation yields the required performance as an overall transfer rate expressed in gigabytes per hour (GB/hr).

The following table provides a listing of the transfer rates of the Compaq small/medium business backup alternatives, providing performance information that can be used in making the backup solution decision.

Table 4-2
Transfer Rates of Compaq Small/Medium Business Tape Backup Alternatives

Tape Drive	2/4-GB DAT	4/8-GB SLR
Native Maximum	650-MB/HR	1.36-GB/hr
Native Typical	300-MB/HR	1.0-GB/hr
2:1 Compressed Ratio Maximum	1.4-GB/hr	2.7-GB/hr
2:1 Compressed Ratio Typical	1.0-GB/hr	2.0-GB/hr

The following chart provides capacity information for the Compaq 2/4-GB DAT and 4/8-GB SLR drives that should be considered as an appropriate backup solution for the small and medium business.




Figure 4-1. Capacity of Compaq Small and Medium Business Tape Drivers

By comparing the results of the required-performance calculation with the typical backup performance rates listed in Table 4-1, the administrator can determine if the required backup performance is achievable. Capacity information is presented in the chart above. If one of the drives will provide the level of performance and meets capacity requirements, then local, off-line backup using that device is an appropriate choice.

Hardware and Media Costs

The needed drive performance and drive capacity must be evaluated in view of current drive and media costs and your company's budget. A realistic budget for the purchase of the appropriate drives and media is important. If too little is budgeted for drives, the company will probably incur increased labor costs.

An all-too-common example of a poor backup solution decision involves a system administrator who must work overtime every evening changing tapes on a drive without an autoloader. Another example of a poor solution decision involves a company that ends up with a drive that is simply too slow to complete the backup operation during the backup window. This could result in degradation of server performance outside the backup window. The unit will be in use more than its intended design, and will prematurely wear out a tape drive.

On the opposite end of the price/performance spectrum, a drive with ten times the required capacity may have the advantage of low-cost per gigabyte, but the initial purchase price would be difficult to justify. The Compaq DAT and SLR tape drives mentioned previously fall within the typical backup solution price/performance requirements for small and medium size customers.

Hardware Reliability

The reliability of any backup device is directly related to its duty cycle (the number of hours per day that the device is in use). The hardware device which will back up the small and medium business server should be selected so that the reliability or duty cycle rating satisfies the expected and, more importantly, actual level of backup activity. For example, if a tape drive designed for 1-GB backups is being used to back up a 10-GB server, the impact on that drive will probably include the following:

- A need for accelerated preventive maintenance, especially head cleaning
- Premature aging
- Reliability problems

The best method for building hardware reliability into a backup strategy is to ensure that the backup hardware is matched to the server(s). If the company needs special, partial backups in addition to the routine backups, it might be appropriate to select the next larger drive size.

If a 2/4-GB DAT drive appears to be the appropriate choice for a given system, the projected duty cycle of the drive becomes an important consideration. Consider the example of an organization that plans to do an unattended, 4-GB complete backup every night, using one 2/4-GB DAT Tape Drive. The 4-GB of data are well within the capacity of the drive. However, at a backup rate around 1-GB/hour, that tape drive will be in operation for 4-8 hours every night. With this amount of usage, the tape drive's read/write heads should be cleaned every other night. Assuming that a person can be found to do this, these backups could hardly be called "unattended." If the drive heads are not cleaned when they need cleaning, soft errors will increase (as will overall backup time) to the point at which something will fail. For this particular customer, even though a 2/4-GB DAT drive would seem to be the appropriate choice, a SLR 4/8-GB, or perhaps even a DLT drive, would be a better choice.

Compaq Small and Medium Backup Solution Tape Drive Alternatives

As previously mentioned, Compaq offers a line of tape drives that include the Compaq DLT family of tape drives, DAT drives, SLR drives, and provides support for other older tape technologies, such as the QIC.

The most appropriate drives for the small and medium business backup solution are the DAT and SLR drives. These drives are most likely to provide the performance, capacity, and reliability needed at a tape drive and media cost that is appropriate for the small or medium sized business.

- **DAT** drives are usually the appropriate choice for servers with 2- to 4-GB capacity. Standard DDS1 DAT tape drives (2/4-GB DAT) can store approximately 2-GB without compression. An Autoloader for use with DAT tapes is available from Compaq. Autoloaders reduce administrative costs by using a robotic mechanism to load and unload tapes.
- **Compaq SLR Drives** feature backward compatibility to older tapes and growth to future products ranging from 125-MB to 32-GB, making the Compaq 4/8-GB SLR Tape Drive a best growth choice.
- **QIC** drives meet the half-height form factor of desktop computers. QIC tapes are virtually industry-standard for standalone machines. However, with capacities limited to 1.2-GB, they are not generally suitable for backing up servers with 2-GB or more storage capacity.

Compaq Small and Medium Business Tape Drive Backup and Restore Performance

Performance testing involved on-line backup and restore using ARCserve 6.5 for Windows NT on a Compaq ProSignia 200 with 128-MB of memory. The system included Domino R4.5 and the Lotus Notes Backup Agent 2.0. Also included were CA InocuLAN 4 for Windows NT and CA AntiVirus Agent Version 2.0 for Lotus Notes. The system was running Windows NT Server 4.0. The testing involved local on-line backup and restore operations.




Figure 4-2. 2/4-GB DAT Backup and Restore Performance

The 2/4-GB DAT Tape Drive's backup performance is 1.32-GB/hr. The same tape drive's restore performance is 1.36-GB/hr. The backup performance of the 4/8-GB SLR Tape Drive exceeds the performance of the 2/4-GB DAT Drive. The 4/8-GB SLR represents a low-cost, reliable, higher capacity alternative to the 2/4-GB DAT. The 4/8 SLR supports 4-GB of native data capacity and 8-GB of 2:1 compressed data.

Compaq Small and Medium Business Backup Solution Summary

Companies continue to entrust their mission-critical data to their computer systems, networks, and enterprises. As long as the possibility of a catastrophic failure exists, a well thought-out backup strategy is crucial. Developing the optimum backup strategy for a particular corporate situation requires an understanding of the performance, capacity, life, and costs of the various backup solutions available today and in the near future. A sound restore strategy is also important.

In general, the system administrator must weigh performance, capacity, projected duty cycles, and cost factors against the volume of backup information when choosing an appropriate and cost-effective backup technology.

4-12 Backup Solution Hardware

The Compaq 2/4-GB DAT drive and Compaq 4/8-GB SLR Tape Drive are the appropriate tape drives for a small company. These tape drives will meet the capacity and performance requirements of most small business systems. The prices of these tape drives are less than \$1,000, placing both drives within the budgetary limitations of most small businesses. These tape drives can back up more than 1-GB/hr using the ARCserve Backup technology.

Compaq Backup Hardware Summary

The success of any backup and restore strategy implementation depends largely on the selection of the appropriate tape drive component for the backup solution. The following tables will help you to focus on two important factors to consider when making the hardware decisions:

- Performance – the data transfer rate of the drive
- Capacity – the total storage of the drive on a single tape

The following table provides a list of the Compaq tape drives, the drive's native (uncompressed) capacity, and the drive transfer rate without compression, and the amount of time required to perform a 10-GB backup.

**Table 4-3
Compaq Tape Drives, Native Capacity, Drive Transfer Rates, and Time Required for 10-GB Backup**

Compaq Tape Drives	2/4-GB DAT	4/16-GB DAT	10/20-GB DLT	15/30-GB DLT	35/70-GB DLT
Drive's native (uncompressed) capacity	<2-GB	<4-GB	<10-GB	<15-GB	<35-GB
Drive's (compressed) capacity	<4-GB	<8-GB	<20-GB	<30-GB	<70-GB
Drive transfer rate without data compression (GB/hr)	1.3	1.44	4.5	4.5	24
Time required for a 10GB backup (HR)	8	7	2.5	2.5	0.5

Properly matching the tape drive unit to a given system requires using capacity and performance information together to make an appropriate tape drive selection. The transfer rate as well as capacity must be considered when selecting the appropriate tape drive. A drive with sufficient capacity, but slow performance could mean that the daily backups take eight hours to complete. The drive would thus likely be worked much harder than its design intended. This results in a tape drive that needs to be cleaned much more frequently than when used for its designed workload. This could also result in a shorter product life than product specifications indicate due to extended use of the read/write heads and other components.

Additional performance information pertaining specifically to the DLT tape family is located in Chapter 5, "Compaq DLT Backup and Restore Performance Test Results".

Chapter 5

Compaq DLT Backup & Restore Performance Results

The enterprise customer has a different set of criterion on which backup solution decisions are based. While the small/medium business usually has a limited number of servers, so that the backup operations take place locally at the Domino server, the enterprise customer's environment is typically much more complex. A high-end backup solution alternative may be selected and used to back up data from multiple servers using the network to perform the remote backup. Thus the factors that are considered when making an enterprise backup decision are quite different than the factors weighed in the backup solution decision for a small or medium sized business.

This chapter discusses several factors that affect performance of backup and restore operations that can be controlled through the configuration of the Compaq backup solution. The performance numbers presented in the following sections are to be used in relative comparison between different configurations, and are included to give the reader an idea of which method and/or configuration achieves better throughput rates. The throughput rates can vary with the nature of data being backed up and with other variables. This is particularly true when using data compression. The nature of data being backed up determines the effectiveness of data compression algorithms, and, in turn, will affect the backup throughputs.

There are a lot of factors that affect backup performance including processors, memory, SCSI type, PCI channel bottlenecks, disk subsystem, network bottlenecks, the number of DLT tape drives, RAID support, data compression ratios, block size, and other minor factors such as the amount of time required to rewind the tape. This chapter contains local and remote backup performance information for the DLT tape drive.

Local Backup Testing

Local backup testing examined the impact on performance of the following factors:

- Tape Drive and SCSI Controller Configuration
- RAID Level
- Data Compression Ratio
- Processor

The chapter also includes a discussion of the potential impact due to the environment feed speed, the data being backed up and/or restored, block size, as well as many minor factors that impact performance. Remote backup performance testing examines the impact on performance of the speed of the network, comparing a network operating at 10-Mb/s to one operating at 100-Mb/s.

Tape Drive and SCSI Controller Configuration Combinations

The purpose of this test is to observe the average throughput performance of 1 to 4 Compaq 35/70-GB DLT Tape Drives used with 1 or 2 Wide-Ultra SCSI controllers running under ARCserve 6.5.

Table 5-1
Test System Configuration

Operating System:	Windows NT 4.0
Backup Software:	ARCserve 6.5 for Windows NT; ARCserve Backup Agent for Lotus Notes Version 2.0
Data Set:	2-GB Notes Database
Tape Drive:	Compaq DLT-35/70GB Tape Drive
Tape Format:	RAID 0
Hardware:	ProLiant 6000, one Pentium Pro 200 processor with 512K L2 cache, 128-MB RAM, 8-GB array, Wide-Ultra SCSI-3 controllers

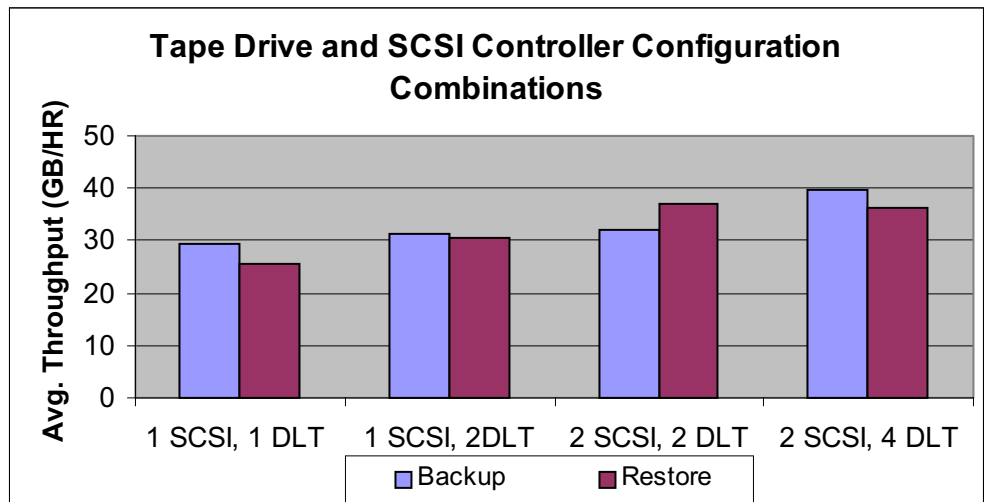


Figure 5-1. Tape Drive and SCSI Controller Configuration Combinations

In the previous chart, two SCSI controllers driving four 35/70-GB DLT drives offer the fastest backup average throughput. Although it is slightly slower than the two SCSI controllers driving the two 35/70-GB DLT drive's restore performance, the backup capacity is doubled. Thus two SCSI controllers driving four 35/70-GB DLT drives is the recommended configuration to avoid tape changing during the backup of large data sets. Four 35/70-GB DLT drives are recommended, especially for enterprise customers who have backup plans handling more than 100-GB of data. Tape change delays would 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. The bandwidth of this card can easily drive two 35/70-GB DLT drives. On a system with only Fast-wide SCSI Controllers, driving four 35/70-GB DLT drives with four cards provided the optimal backup performance.

RAID Level

The purpose of this series of tests is to determine the performance differences of various Tape RAID levels running ARCserve 6.5 with the ARCserve Backup Agent for Lotus Notes Version 2.0. ARCserve 6.5 provides a Tape RAID (or 'RAIT') system so that multiple tape drives can be placed in a group and be written to as one 'tape array' by a single backup job. Fault tolerance can be provided to the array through RAID 1 (mirroring) or RAID 5 (striping with parity), so that backup operations can continue if a tape drive fails, or so that a backup set can be recreated if a tape is lost. The array can also utilize RAID 0 (striping with no parity) for a purely high-performance backup without fault tolerance.

The tests were run to determine the impact of various RAID levels as related to processor performance and throughput rates, to discover how much processing power is required to perform the various RAID calculations, and to determine how the introduction of fault tolerance affected backup and restore times.

Table 5-2
Test System Configuration

Operating System:	Windows NT 4.0
Backup Software:	ARCserve 6.5 for Windows NT; ARCserve Backup Agent for Lotus Notes Version 2.0
Data Set:	2-GB Notes Database
Tape Drive:	4 Compaq DLT 35/70-GB Tape Drives or 3 Compaq DLT 10/20-GB Tape Drives
Tape Format:	RAID 0, 1 and 5
Hardware:	ProLiant 6000, one Pentium Pro 200 processor with 512K cache, 128-MB RAM, 8-GB array , Wide-Ultra SCSI-3 controllers

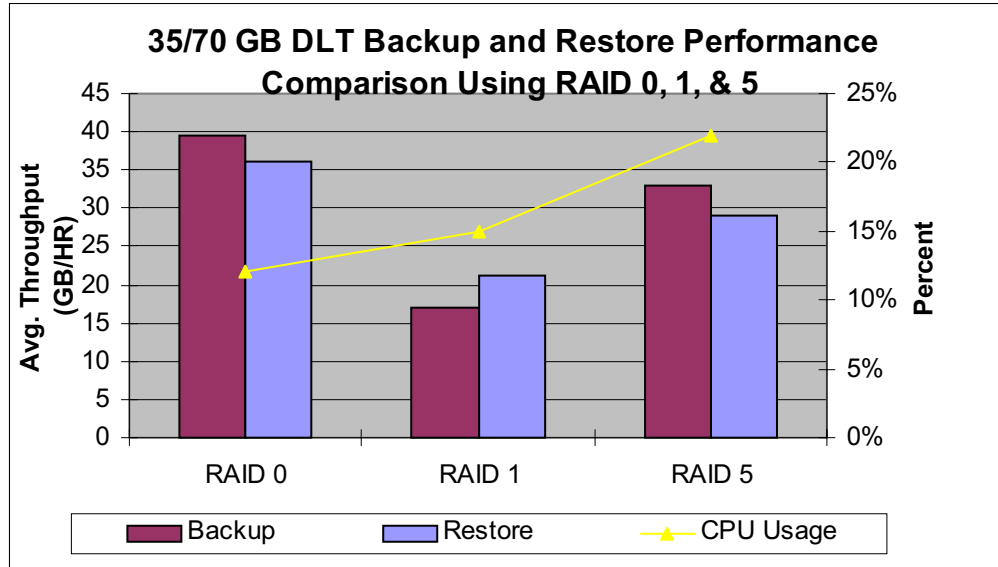


Figure 5-2. 35/70-GB DLT Backup and Restore Performance Comparison Using RAID 0, 1, & 5 and Processor Utilization Rate

The previous chart illustrates the backup and restore performance of the Compaq 35/70-GB DLT Tape Drive configured using RAID 0, RAID 1, and RAID 5. Average throughput and CPU utilization can be compared for each of the three backup and restore tests run. The CPU Utilization increased from 12% when configured with RAID 0, to 15% for the RAID 1 configuration. RAID 1 provided fault tolerance at a very low cost in terms of processor utilization. However, note that the difference in RAID levels had a significant impact on the average throughput of both the backup and restore operations. The RAID 0 backup operated at a rate of 40-GB/hr while the RAID 1 backup average throughput was only 17-GB/hr. The restore operation resulted in similar results. The RAID 0 restore operated with a average throughput rate of 36-GB/hr while the RAID 1 restore rate was 21-GB/hr.

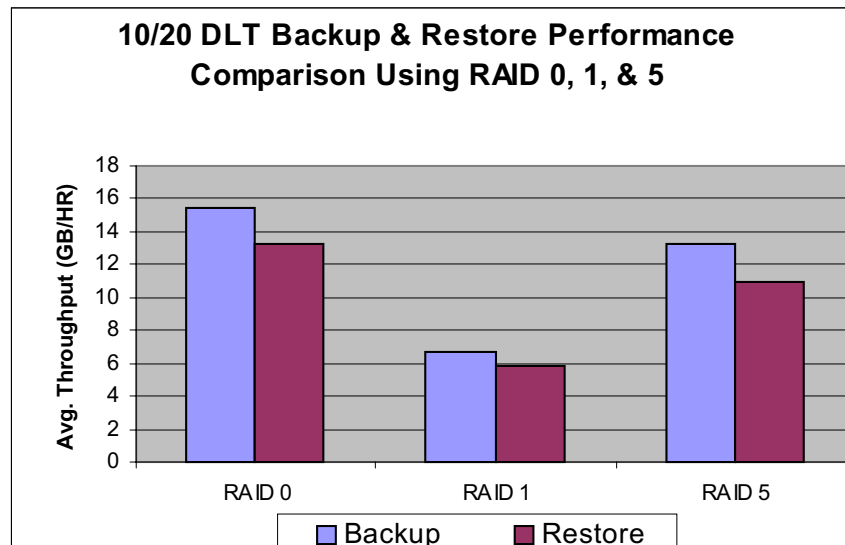


Figure 5-3. 10/20-GB DLT Backup and Restore Performance Comparison Using RAID 0, 1, & 5

The previous chart illustrates the backup and restore performance of the Compaq 10/20-GB DLT Tape Drive configured using RAID 0, RAID 1, and RAID 5. Average throughput can be compared for each of the three backup and restore test runs. Note that the difference in RAID levels had a significant impact on the average throughput of both the backup and restore operations. The RAID 0 backup operated at a rate of 15.4-GB/hr while the RAID 1 backup average throughput was only 6.7-GB/hr. The RAID 5 backup operated at an average throughput of 13.3-GB/hr. The restore operation resulted in similar results. The RAID 0 restore operated with a average throughput rate of 13.3-GB/hr while the RAID 1 restore rate was 5.8-GB/hr and the RAID 5 restore rate was 10.9-GB/hr. For both backup and restore, RAID 5 provided fault tolerance support while maintaining a throughput rate very close to the rate maintained by the RAID 0 configuration, which provided no fault tolerance support.

In summary, both 35/70-GB DLT and 10/20-GB DLT tests show that the performance of RAID 5 in a configuration using two SCSI cards connected to four drives is roughly 20 percent slower (depending on data set compressibility) than the RAID 0 backups. Thus, RAID 5 should be considered as the optimal choice in a decision between speed and safety. RAID 5 not only writes to a set of tapes as an array, but also uses a parity stripe to provide checksum information. This checksum information permits file restoration even when there has been data loss on a single tape in the array.

5-8 Compaq DLT Backup & Restore Performance Results

RAID 0, on the other hand, logically associates the drives and permits them to be written to as a single drive. Because RAID 0 does not perform the parity calculations, better throughput is achieved than when configured to support RAID 5. Thus additional processing power is required to support RAID 5 than RAID 0. This could create a problem if performing concurrent backups during heavy server activity. Adding processors to the system may help reduce the impact of RAID 5 on server resources.

Data Compression

Data compression is typically available via hardware (tape drive or tape drive controller) or software (off-line backup application). Speed, effectiveness (compression ratio), and levels of support vary from one solution to another and from one data set to another. Data compression will reduce the time and storage required at the expense of adding some processing overhead.

Hardware data compression is available with the DLT tape drives when using 10.0, 15.0, or 35.0 density formats. In this mode, data is always compressed when written to the tape cartridge. However, the drive will read both compressed and uncompressed tapes. You can choose not to use data compression by manually de-selecting the compression option on the drive, or by using tape software that supports data compression control.

Hardware data compression is typically more effective than software data compression, when both are operating on the same data set. This is an advantage for sites that want to conserve backup storage, and are not pressed to perform backups within the smallest time window possible.

Software data compression can sometimes be faster than hardware data compression, since it can utilize a more powerful and faster system processor (depending upon the processor type, of course), and because it actually sends less data to the tape drive controller. Performance of software data compression is an advantage for sites that have a need to back up as fast as possible, with less concern over the backup storage consumed.

The major disadvantage of software data compression is the consumption of system processing power. File-based backup software, with software data compression turned on, will typically utilize the processor heavily. To achieve optimal backup throughput in this case, make sure you have plenty of processor power (a high-speed Pentium or Pentium Pro processor, or dual processors if the compression algorithm is multi-threaded). Use the Performance Monitor to determine your CPU utilization level. If your processor becomes saturated, then your performance can suffer greatly. Hardware data compression, on the other hand, does not utilize the system processor(s) at all, and only the system I/O imposes some overhead on the processors.

Table 5-3
Test System Configuration

Operating System:	Windows NT 4.0
Backup Software:	ARCserve 6.5 for Windows NT; ARCserve Backup Agent for Lotus Notes Version 2.0
Data Set:	2-GB Notes Database
Tape Drive:	Compaq DLT 35/70-GB Tape Drive
Tape Format:	Single Stream, Regular
Hardware:	ProLiant 6000, one Pentium Pro 200 processor with 512K L2 cache, 128-MB RAM, 8-GB array, Wide-Ultra SCSI-3 controller




Figure 5-4. Backup Performance Impact of Data Compression Ratio

The backup speed for non-compressible data highlights the advantage of the data compression hardware built into the drive. The 2.32:1 data backed up roughly 40% faster than the same 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 capabilities. As the data compression ratio approaches 1:1, the advantage gained becomes proportionally less.

The nature of the data determines the effectiveness of data compression, and will directly impact backup performance. Highly compressible data, which we will characterize below, will yield substantially higher backup throughput rates than data that is not compressible.

Predicting throughput rates and backup duration is extremely difficult without knowledge of the compressibility of your data. Lotus provides a way to determine the compressibility of Lotus Notes Database. From the Notes Client, use the Menu command **File\Database\Properties, Select the Information Tab**, to enable you to view the properties dialog box illustrated below. Click the **% used** button to request Notes to calculate the percentage that actual data represents as part of the overall file size.

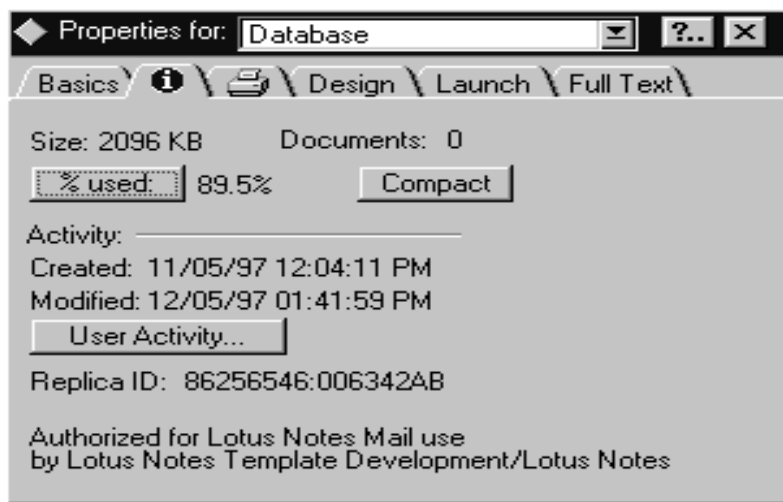


Figure 5-5. Database Properties Information Tab Provides % Used Value

Logically, a smaller number implies a higher compressible ratio. For example, in the previous chart, the % used of the database whose compressible ratio is 2:32 is 75%. A low % used result means that the data occupancy of the storage space used for total database storage space is not optimal, and could be compressed using a compression algorithm. A high % used result, on the other hand, means that the data occupancy of the total storage space is very optimally used, and therefore the data could not be compressed much more. For example, in the previous chart, the % used of the Lotus Notes databases whose compressible ratio was 1.11 is 99.9%.

In summary, the hardware compression capabilities of the 35/70-GB DLT tape drive provide a significant improvement in backup rates for data that is highly compressible. For data that is already highly compressed, turning on the hardware compression feature will not provide an advantage from a throughput perspective.

Processor Usage

The average backup utilizes from 12 to 20% of a single 200 MHz Pentium Pro 512K L2 cache processor. The lower end of the spectrum is typical of using a single disk or tape drive. This usage percentage grows substantially when striping data across a redundant array of independent tape drives (RAID) in a RAID-5 configuration. All testing was performed on 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. Therefore, the majority of tests were performed with two processors.




Figure 5-6. Compaq 35/70 DLT Processor Utilization for RAID 5 Backup Test

The average processor utilization is 20% for the Compaq 35/70 DLT using RAID 5.

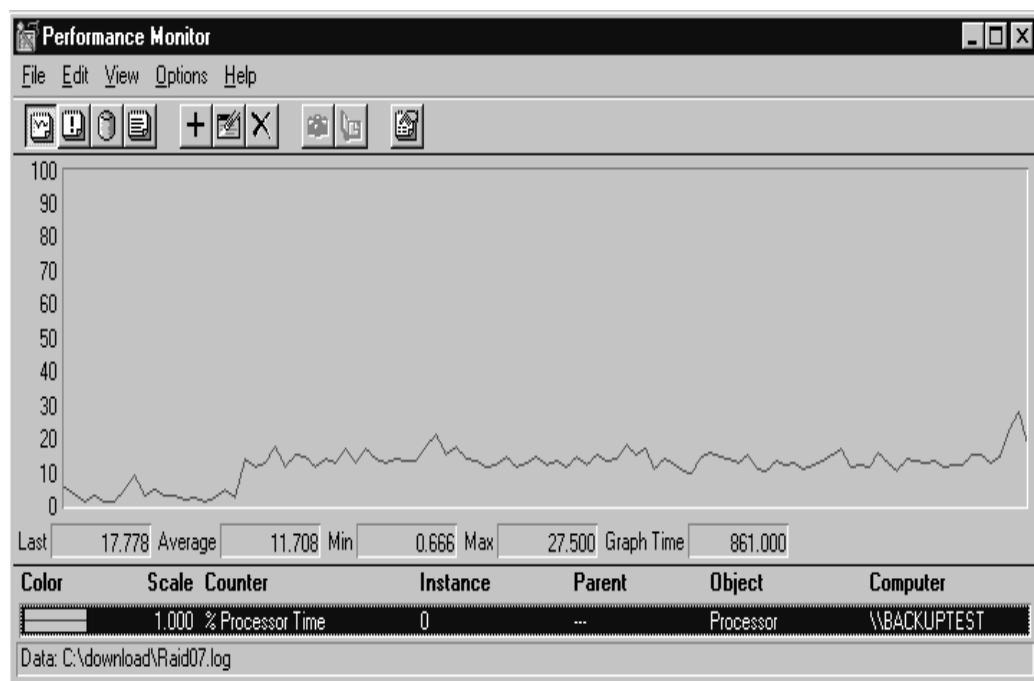


Figure 5-7. Compaq 35/70 DLT Processor Utilization for RAID 0 Backup Test

The average processor utilization is 11.7% for the Compaq 35/70 DLT using RAID 0 during this backup test.

Based on the previous processor performance monitor chart, the average processor power usage is from 12% with RAID 0 to 20% with RAID 5. Processing power is a factor that can affect the backup performance, although even one Pentium 133 or Pentium 166 processor is sufficient to provide acceptable backup performance. If a company is using remote backup, the backup server performance will typically encounter a bottleneck with network resources and does not typically experience any processor performance problems. If you are using the local backup, an optimally scheduled backup task that avoids the busy times for the Notes server will help alleviate the burden encountered by the system processor or the Notes server during the backup activity.

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 than 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.

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 one 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.

Block Size

When data is written to tape media, it is formatted for a certain blocking factor. The number of bytes per block is usually determined by the application when sending the data to the tape drive. When working with DLT Tape drives, this block size is an important consideration, as it can have serious performance implications. Ideally, the DLT drive can benefit from as large block size as possible, while the minimum block size recommended is 8-KB. However, many applications do not use the larger block sizes due to in-memory buffering constraints. Some applications unfamiliar with DLT technology even use block sizes smaller than 8-KB because most other kinds of tape drives perform well with block sizes that can be much smaller (512 bytes or 1-KB).

A well-designed application will first query the tape drive with the appropriate API call to determine the kind of tape drive before attempting to set the block size. Actual SCSI data transfers are done to/from the DLT drive in 64-KB sizes across the SCSI bus, but are further formatted for an 8-KB blocking factor by the DLT before writing them to the tape media. The drive is set for fixed block mode, in which it sends or receives a fixed amount of 'blocks' in each data transfer.

Limitations will still be encountered, however, even with an 8-KB block size. These limitations are most apparent when using compressible data sets. The overhead involved in formatting the data for an 8-KB block size establishes a ceiling on how fast the DLT drive can write the data onto the tape. Thus the additional advantage gained by having the drive compress the data is not realized here. However, when an increase is made to a larger block size, the drive is able to write the data much faster so that the advantages of hardware compression become apparent.

Other Performance Factors

A number of additional factors that can all influence performance of backup and restore operations exist and should be mentioned. These factors include the following:

- Rewinding tapes
- Changing tapes
- Searching for the end of the backup set, when applicable (appending new data)

5-16 Compaq DLT Backup & Restore Performance Results

- Writing additional data onto tape, such as registry and/or catalog information
- Cataloging of tapes
- Backing up across the network when the network utilization is high (available network bandwidth therefore low)
- User activity on the database
- Diagnostics and replacement of the failed hardware
- Creating and/or formatting of lost partitions, if necessary

Re-creation of the database devices and the database itself, if necessary

Remote Backup Performance Testing

Remote backups involve backing up a database which resides on one server, by transferring data across a network interconnect to storage devices which reside on another server. This method allows for the ‘pooling’ of tape or disk storage devices on a central storage server, to which data is sent from other servers on the network. The need to have local backup devices on each server in the environment is thus eliminated. When multiple tape or disk drives exist on the storage server, the remote backup strategy can also be employed to back up multiple database servers simultaneously.

Database backup operations are highly throughput-intensive operations, requiring fast data transfer rates. With the powerful systems and applications available to us, very fast backups can be achieved to high-speed storage devices connected locally, as we have already shown. In most remote backup scenarios however, the network devices end up becoming the bottleneck. During the remote backup and restore performance testing, the goal was to determine the impact of a network bottleneck when backing up and restoring multiple servers across a single 10BaseT segment and a single 100Base TX segment. The objective was to alleviate or minimize the network bottleneck by using strategies involving high speed 100-MB/s Ethernet (in addition to slower, 10-MB/s Ethernet), multiple segments for concurrent backups, and multiple Network Interface Cards (NICs) per segment.

Concurrent Backup & Restore of Multiple Servers Across a Single 10Base-T Segment

Table 5-4
Test System Configuration

Operating System:	Windows NT 4.0
Backup Software:	ARCserve 6.5 for Windows NT; ARCserve Backup Agent for Lotus Notes Version 2.0
Data Set:	2-GB Notes Database
Tape Drive:	Compaq DLT 35/70-GB Tape Drive
Tape Format:	RAID 0
Network:	1.Compaq NetFlex-3 and Netelligent NIC (Network Interface Cards). 2.Compaq Netelligent 10Base-T Repeater. 3.TCP/IP Protocol
Storage Server:	ProLiant 6000, one Pentium Pro 200 processor with 512K L2 cache, 128-MB RAM, 8-GB array, Wide-Ultra SCSI-3 controller

The performance was compared for backing up and restoring a single remote server located on the same 10Base-T network segment to providing simultaneous backup and restore services to two remote servers located on the same 10Base-T network segment.

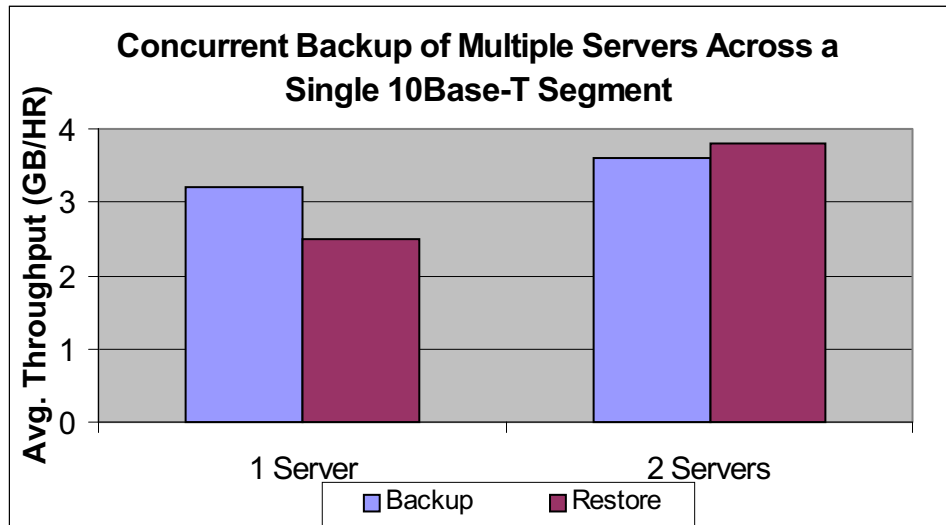


Figure 5-8. Concurrent Backup & Restore of Multiple Servers Across a Single 10Base-T Segment

The average throughput for the single-server remote backup was 2.5-GB/hr while the two-server remote backup throughput averaged 3.6-GB/hr. The average throughput for the single-server remote restore was 2.5-GB/hr while the two-server remote restore throughput averaged 3.8-GB/hr. The restore throughput performance increased more for the two-server restore operation than the two-server backup operation.

Although not displayed, the CPU usage at both servers for all transfers was very low, averaging around 5%. Since plenty of system resources are left at both servers for processing other jobs, the network is clearly the bottleneck in this backup and restore scenario. Given the network saturation, such backup jobs should be scheduled during hours of low network use by other applications to avoid interference, or across a dedicated network link.

Concurrent Backup of Multiple Servers over a Single 100Base-TX Segment

Table 5-5
Test System Configuration

Operating System:	Windows NT 4.0
Backup Software:	ARCserve 6.5 for Windows NT; ARCserve Backup Agent for Lotus Notes Version 2.0
Data set:	2-GB Notes Database
Tape Drive:	Compaq DLT 35/70-GB Tape Drive
Tape Format:	RAID 0
Network:	<ol style="list-style-type: none"> 1. Compaq NetFlex-3 and Netelligent NIC (network interface cards). 2. Compaq Netelligent 100Base-TX Repeater. 3. TCP/IP Protocol
Storage Server:	ProLiant 6000, one Pentium Pro 200 processor with 512K L2 cache, 128-MB RAM, 8-GB array, Wide-Ultra SCSI-3 controller

The performance was compared for backing up and restoring a single remote server located on the same 100Base-TX network segment to providing simultaneous backup and restore services to two and then three remote servers located on the same 100Base-TX network segment.

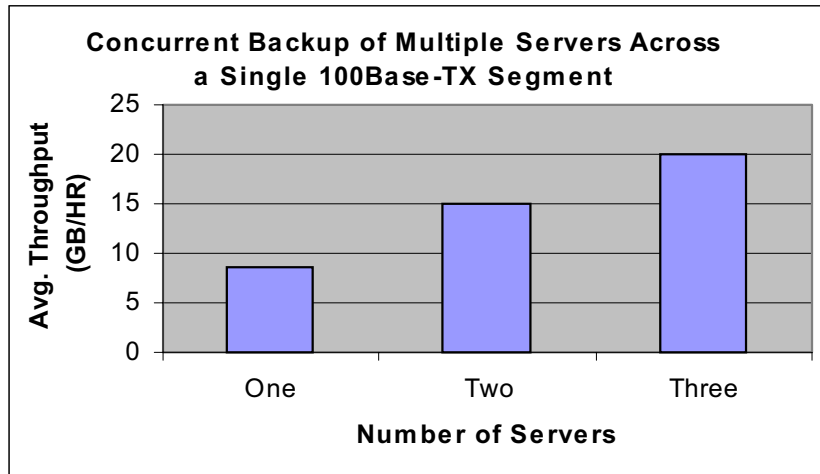


Figure 5-9. Concurrent Backup & Restore of Multiple Servers Across a Single 100Base-TX Segment

The benefits of this concurrent strategy are obvious, as greater throughput is derived from the tape server and more of the network bandwidth is utilized, for each additional database backed up from a remote system. Eventually, the overall throughput reaches 20-GB/hr with three simultaneous channels of data being transferred into the server and written out to tape, and 60% of the network bandwidth in use.

The performance increase is fairly linear as the number of servers is increased from one to two and finally to three. The single-server scenario supported 8.7-GB/hr, the two-server scenario supported 15-GB/hr, and the three-server scenario supported 20-GB/hr.

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 the network 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 will drop out of write mode. A significant amount of time is then 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 rate 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.

Remote Backup Conclusions

Compaq Engineers recommend that customers select a 100Base-TX Repeater for the remote backup if the DLT 35/70-GB Tape drive is being used, since only the 100Base-TX repeater will feed the data at a rate fast enough to match the DLT Tape D drive capabilities. Compaq engineers strongly recommend using the Compaq Netelligent 100Base-TX Repeater for remote backup if there is a requirement to back up multiple concurrent servers as fast as possible across the network.

The 10Base-T Repeater was observed providing support for up to 3.2-GB/hr backup average throughput performance, so the DLT 10/20 or even DAT 12/24 tape drives are fast enough for a 10 Base-T Repeater network implementation. At the same time, the DLT 10/20 and 12/24 will utilize most of the available network bandwidth if the 10 Base-T Repeater is used. Thus The DLT 10/20-GB Tape Drive or DAT 12/24-GB Tape Drive will be a smart choice for a similar implementation using 10Base-T Repeaters that can bear the slower backup rate. Another strong recommendation is to set the backup schedule during the hours of low network use by other applications to avoid interference, or configure the backup to operate across a dedicated network link.

Backup Conclusions

Companies continue to entrust their mission-critical data to their computer systems, networks, and enterprises. As long as the possibility of a catastrophic failure exists, a well thought-out backup strategy is crucial. Developing the optimum backup strategy for a particular corporate situation requires an understanding of the performance, capacity, life, and costs of the various backup solutions available today and in the near future. A sound restore strategy is also important.

5-22 Compaq DLT Backup & Restore Performance Results

In general, the system administrator must weigh performance, capacity, projected duty cycles, and cost factors against the volume of backup information when choosing an appropriate and cost-effective backup technology.

Small and Medium Business

The Compaq 2/4-GB DAT drive and Compaq 4/8-GB SLR Tape Drive are the appropriate tape drives for a small to medium size company. These tape drives will meet the capacity and performance requirements of most small business systems. The prices of these tape drives are less than \$1,000, placing both drives within the budgetary limitations of most small businesses. These tape drives can back up more than 1-GB/hr using the ARCserve Backup technology. As the names imply, the capacity of the 2/4-GB DAT drive is 2-GB native (uncompressed) and 4-GB using a 2:1 compression ratio. The 4/8-GB SLR similarly has a capacity of 4-GB native (uncompressed) and 8-GB using a 2:1 compression ratio.

Enterprise Customers

The DLT drives are Compaq's highest performing tape backup solution, and when configured as an array with the use of ARCserve RAID Option this performance is scaled even further. By adding enough of the drives, we can know that we can attain backup throughput that is limited only by the ability of the disk subsystem to deliver the data. Of course, disk subsystem data-transfer rates will continue to increase as more advanced SCSI and Fiber Channel technologies emerge, and these combined with high-end DLT drives, such as the 35/70-GB DLT, will allow us to attain unprecedented backup throughput even beyond what is demonstrated in this paper.

The DLT drives also provide us with the type of capacity sufficient for any size company. Two DLT Tape Arrays (eight 15/30-GB DLT drives) can be combined into a single array that provides at least 120-GB of storage space (no compression) on a single set of tapes. With just one DLT Tape Array II (four 35/70-GB DLT drives), at least 140-GB of space (no compression) is available. With the DLT 15-Cartridge Library, storage capacity is extended to at least 450-GB (these numbers are increased if compression is enabled).

The DLT hardware and media are rated for superior reliability, even under continuous operation. However, it is important to consider that, as multiple tape drives are combined into a stripe set, the Mean Time Between Failure (MTBF) of the tape array as a whole is compounded. This is due to the dependency upon a greater number of devices (e.g: an array of four tape drives is four times more likely to fail than a single tape drive). Thus, fault tolerance for the tape drives or tape sets may become a consideration. The Data Warehousing or IT department seeking a fault tolerant backup solution should certainly look at the ARCserve RAID Options, which are available directly from Compaq.

Using disk drives instead of tape drives for local backup operations may offer certain benefits in performance that can be useful for those who have a limited time interval in which to complete the backup of a large database. The performance of disk array dumps can be matched with a sufficient number of striped 35/70 DLT drives. However, dumps to disks are generally not practical due to capacity limitations. In addition, disks do not offer readily removable media that can be inexpensively archived. Rather, the use of disk drives for local backup should be used in conjunction with a tape solution, where the database can be dumped rapidly to disk and then the dump device files backed up to tape at leisure using a file-based backup utility. The advantage here is that interruptions to user transactions on the database, caused by a rapid on line dump, would be relatively brief.

Chapter 6

Virus Protection Security

As applications on computer systems become major contributors to the daily productivity of employees, maintaining the safety and usability of the data on those systems becomes increasingly critical. An earlier chapter discussed the importance of backing up data to permit restoration of the data if it somehow became lost or corrupted. This chapter focuses on the detection and protection against viruses on your system.

Viruses can occur on the client's workstations or on the server. Both servers and workstations should have some type of anti-virus protection software. Many of the major firewall vendors support anti-virus scanners; however, it is extremely important to implement some type of anti-virus protection on the server if the company has not implemented a firewall solution (because the corporate network is not yet connected to the Internet). The company should provide a recommendation regarding the implementation of virus scanning utilities for the server as well as for employee desktops as part of the overall company security policy.

Necessity of On-line Virus Scanning Utilities

Many virus-scanning products are available for the workstation or desktop. The selection of a standard utility for your company systems is highly recommended. A virus that has infected one workstation can be spread throughout the entire company as fast as a raging fire across a field of dry grass.

Organizations demand robust and full featured backup and virus protection tools that are easy to use and manage. In the past, administrators were required to shut down or bring the messaging servers off-line in order to back up the entire system with full data integrity. Although this was a possible solution, it obviously greatly impacted (temporarily eliminated) user access. In today's environment, customers demand robust and full-featured virus protection tools that are easy to use and manage, *without* impacting user activities.

Common Viruses

Viruses have threatened the computer user for years. The type of virus that is most common has transitioned from the traditional boot-sector viruses to macro and polymorphic viruses. Boot-sector viruses place their code in the sector whose code the machine will automatically execute when booting, so that when the machine boots, the virus loads and runs. After the boot sector virus is finished loading, it loads the original boot code, which has been previously moved to another location.

Macro viruses have become the biggest threat to today's computer user, overcoming the traditional boot-sector viruses. Macro viruses are able to attach to and infect common documents, such as Microsoft Word files and Excel spreadsheets. This ability makes the spread of a macro virus a very simple process that can affect many unsuspecting recipients. The macro virus can be spread by simply sending or posting an e-mail with an infected document to one or many unsuspecting users. The damage of these viruses can grow into a serious problem very quickly. This damage accelerates as each individual user is affected by the macro virus, and administrators experience lost data and /or wasted time.

Because macro viruses are written in a macro language, a given macro virus will run in any environment which can understand and interpret those macro instructions. Macro viruses are always application-specific. A macro virus consists of a nothing more than a set of macros. These macros must be stored wherever the application stores macros.

Polymorphic viruses are growing in popularity. A polymorphic virus has the ability to mutate by changing code segments to look different from one infection to another. This type of virus is a challenge for anti-virus detection methods.

Traditional Anti-virus Utilities

Traditional anti-virus utilities are unable to protect messaging systems from virus infection. The utilities are unaware that viruses may be lurking inside server-based databases as document attachments. Viruses are able to avoid virus detection and can spread rampantly throughout an organization; relying on the messaging system as a virtual "safe harbor" and worldwide distribution mechanism. The only users safe from infection from viruses in messaging systems are those who detach all document attachments to their PCs running the latest anti-virus software.

CA Anti-virus Software Solution

Many client/server anti-virus solutions can protect users from viruses when documents are detached from messages. However, not many anti-virus solutions can completely scan and cure server-based messaging systems such as Lotus Notes.

Additionally, the integration of CA InocuLAN and ARCserve provides a solution for virus-free backups. CA's AntiVirus Agent Version 2.0 for Lotus Notes integrates with CA'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.

Advanced Virus Protection

Viruses today are more sophisticated and capable of spreading undetected via file servers, e-mail document attachments and Internet downloads. The National Computer Security Association (NCSA) has reported that over 98% of corporations have had a virus incident within the last year, resulting in considerable amounts of wasted time and destroyed data.

CA protects servers and clients in both real-time and with full scheduled virus scans. Whether a potential virus is being copied to a server, downloaded off the Internet, accessed on a floppy diskette, backed up to tape, or detached from an e-mail message; InocuLAN and CA AntiVirus are products whose purpose is detecting and curing viruses as you work.

6-4 Virus Protection Security

InocuLAN has a virus scanning engine which is certified by the NCSA to detect 100% of viruses in the wild. The scanning engine includes an extensive and constantly expanding library of virus signatures, or fingerprints, for detecting and curing known computer viruses. InocuLAN's polymorphic analyzer studies suspected viruses for suspicious activity and an integrity checker monitors critical system files. The latest InocuLAN technology, the macro virus analyzer, addresses the threat of macro viruses which attach to and infect standard documents.

Network-oriented anti-virus technologies also abound, including InocuLAN's virus wall, virus quarantine and virus source tracking features which protect servers from virus infection, isolate infected clients and identify the sources of virus infection. All the InocuLAN virus-fighting technology is updated every month, free of charge, via signature update files.

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

In addition, the integration of CA InocuLAN and Arcserve provides a solution for virus-free backups. A detailed description of each of the CA products relating to security and anti-virus are described in the following section.

InocuLAN 4

CA InocuLAN 4 for Windows NT is a second-generation anti-virus solution for a Windows NT network. Available options, which will be discussed shortly, provide protection for Lotus Notes messaging systems.

InocuLAN 4 provides solid enterprise management and virus protection that installs on a single Windows NT server. CA provides hands-free signature update download and distribution capabilities. InocuLAN is integrated with separate CA AntiVirus clients and unique AntiVirus Agents to protect both workstations and mission-critical messaging systems, such as Lotus Notes, from viruses.

InocuLAN addresses the needs of enterprise network managers with a host of unique features, including automated installation and hands-free updating of InocuLAN servers and networked clients. Both InocuLAN for Windows NT and NetWare servers can be easily grouped and managed from one console with unique domain management features that greatly simplify viewing, configuring, and scanning other InocuLAN servers.

InocuLAN's extensive network traffic management, server utilization and alerting controls enable customization of the solution for your network. InocuLAN's tight integration with Windows NT and ARCserve storage management software help to further ensure maximum compatibility and ease of management.

With InocuLAN, organizations can effectively manage anti-virus efforts across the enterprise, saving time and money, while ensuring the highest level of virus protection.

InocuLAN Feature Highlights

This section describes some of the highlighted features of InocuLAN. The table which follows provides complete product feature and benefit information.

- **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 cannot 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:** Diskettes are the most common source of virus infections, and InocuLAN fully protects the enterprise from floppy-based viruses. As soon as a diskette is accessed, such as looking at the disk contents in your 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 occurs 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:** InocuLAN can protect your Lotus Notes mail systems with its messaging options. Even attached ZIP files are scanned. The options are discussed in a section that follows.

InocuLAN Feature & Benefit Summary

Table 6-1
InocuLAN Features and Benefits

Feature	Benefit
NCSA Certification	Meets National Computer Security Association guidelines for virus detection, including 100% of viruses in the wild.
State-of-the-Art Detection	Detects virtually all known (and many unknown) virus types, including: boot sector, master boot sector, memory resident, file

	multipartite, macro, stealth and polymorphic.
Macro Virus Analyzer	Detects and cures the rapidly spreading Microsoft Word macro viruses.
Domain Management	Greatly eases network administration through Windows-based domain grouping of servers and workstations.
Hands-Free Updating	Automatically dials, downloads and updates all InocuLAN servers and workstations with the latest virus signature update files; even automatically uploads and reloads itself on the server.

continued

6-8 Virus Protection Security

InocuLAN Features and Benefits *continued*

Monthly Updates	Free updates posted monthly on CA's BBS, WWW site and CompuServe forum for the latest protection.
Network Enforcement	Workstations that are not employing real-time virus protection can be remotely logged off the network.
NDS Aware	Supports Novell's NetWare Directory Services for enterprise network management.
Backup Integration	Combined with ARCserve (CA's storage management software), InocuLAN provides the most efficient solution for fast, virus-free backups.
Utilization Gauge	Controls virus scanning to prevent degradation of network performance.
Real-Time Scanning	Scans each file in the background as it is accessed to provide a secure barrier against viruses.
Critical Disk Protection	Guards essential client system files from viruses and enables quick disaster recovery from diskette or server.
Compressed Files	Scans compressed file archives and Internet downloads in ZIP, ARJ and Microsoft Compressed formats.
Command Line Support	Enables automating anti-virus efforts using batch files, scripts or macros.
Detection Options	Variety of ways to handle viruses, including: cure, rename, move, report, delete or purge.

continued

InocuLAN Features and Benefits *continued*

Extensive Alerting	Automatic virus alerts can be sent to selected users via e-mail, pager, fax, SNMP, network broadcast or printed trouble ticket.
Scanning Logs	Complete virus scanning history (for a workstation, domain, network or enterprise) including files checked, viruses found and responses taken.
VxD Technology	Windows 3.x and Windows 95 clients both utilize Virtual Device Drivers (rather than TSRs) for superior performance with less memory requirements.
Windows Interface	Easy to install, learn and use with a common interface across PC platforms and direct scanning capability from the Windows 95 Explorer.
Customer Support	On-screen, context-sensitive help, virus encyclopedia and technical support from CA's worldwide offices.
Cross-Platform Solution	InocuLAN is available for NetWare and Windows NT servers and Windows 3.x, Windows 95, DOS and Macintosh workstations.

AntiVirus Agent V2.0 for Lotus Notes

The CA AntiVirus Agent for Lotus Notes works in conjunction with CA's AntiVirus solution, InocuLAN. The Agent scans for viruses in document attachments and shared databases on Notes servers. Infected files are automatically cured and users are notified directly through e-mail. This solution ensures that the messaging server is virus free. To keep the server virus free, automatic scans can be scheduled to run periodically to check new messages.

AntiVirus Agent V2.0 for Lotus Notes Feature Highlights

This section describes some of the highlighted features of the CA AntiVirus Agent for Lotus Notes. The table, which follows, provides complete product feature and benefit information.

- **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 can be configured. 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 4.x release of Lotus Notes/Domino Server.
-

**AntiVirus Agent V2.0 for Lotus Notes
Feature & Benefit Summary**

**Table 6-2
AntiVirus Agent V2.0 for Lotus Notes Feature & Benefits Summary**

Feature	Benefit
Real-Time Scanning	Automatically scans and cures messages and file attachments for viruses when messages enter the server.
Automatic Updates	Automatically updated to scan for the latest viruses.
Encrypted Database	Able to scan files attached to encrypted Notes databases.
Selective Scanning of Notes database and mailboxes	Options to scan entire e-mail databases or individual post boxes.
Incremental Scanning	Quickly scans only new and modified messages.
Integrated Scheduled Scanning	Scans can be scheduled to run at user-specified intervals.
Integrated AntiVirus Interface	One common interface for configuring and managing anti-virus efforts. The AV Agent is tightly integrated with CA's other Enterprise AntiVirus software for complete protection and system management and automation.
Cross Platform Support	Support for Lotus Notes running on multiple platforms.

6-12 Virus Protection Security

The AntiVirus Agent for Lotus Notes tested on Compaq servers resulted in a very insignificant impact on system resources when comparing a system running the CA AntiVirus Agent for Lotus Notes to the same system not running the Agent. The recommendation is clearly to implement a virus protection program for your servers and client workstations. CA AntiVirus Agent for Lotus Notes provides a full range of features along with sound protection from all known viruses.

AntiVirus Agent for Notes Limitations

There are some known limitations of the CA AntiVirus Agent for Lotus Notes. The following information was taken from CA reference material:

- 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 database; CA has so far identified two cases where they occurred in: 1) Any action other than Report Only will generate this type of error message. 2) Notification via Attachment will generate this type of error message.
- The Agent will only support scanning of explicit server-level encrypted database. .NSF databases encrypted locally 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.

Protection Suite 2.0 for Lotus Notes

The CA Protection Suite for Lotus Notes offers complete 24-hours-per-day, 7-days-per-week, protection for groupware messaging servers. CA's unique AntiVirus and Backup Agents work in conjunction with InocuLAN and ARCserve to provide full backup/restore capabilities and virus protection for on-line Lotus Notes servers.

The Protection Suite includes:

- ARCserve 6 for Windows NT Enterprise Edition
- InocuLAN 4 for Windows NT Single Server
- AntiVirus Agent for Lotus Notes 2.0
- Backup Agent for Lotus Notes 2.0
- InocuLAN Client

Summary of Anti-virus Protection

The threat of viruses requires that systems be protected as much as possible by backing up data as well as by providing virus protection and prevention software solutions.

The CA InocuLAN 4 for Windows NT provides the foundation component needed to provide enterprise systems with virus protection. The CA AntiVirus Agent works in conjunction with InocuLAN to detect and cure viruses attached to Lotus Notes e-mail and in document databases. The AntiVirus Agent is thus another component of CA's enterprise virus protection solution. Performance testing that took place in Compaq labs indicated that the system resource requirements for supporting InocuLAN and the Anti-Virus Agent were minimal.

The CA Protection Suite 2.0 for Lotus Notes provides a combined data protection suite of products. Enterprises that use a CA solution for backup and virus protection could utilize the CA Protection Suite 2.0 for Lotus Notes.

The Protection Suite includes:

- ARCserve 6 for Windows NT Enterprise Edition
- InocuLAN 4 for Windows NT Single Server
- AntiVirus Agent for Lotus Notes 2.0
- Backup Agent for Lotus Notes 2.0
- InocuLAN Client

Index

A

- About this TechNote 1-1
- Additional resources 1-2
- Advanced virus protection 6-3
- Anti-virus
 - protection summary 6-12
 - software solution 6-3
 - utilities 6-2
- AntiVirus Agent V2.0 for Lotus Notes
 - 6-8
 - feature and benefit summary 6-10
 - features 6-9
 - real-time scanning and cure 6-9
 - scanning of encrypted databases 6-9
 - scanning of shared mail 6-9
 - support for Lotus Notes 4.5 6-9
 - limitations 6-11
 - purpose of 6-3
- ARCserve
 - Database Agents 3-1
 - default 3-4
 - integration 6-3
 - RAID engine 3-5
 - RAID option 4-6
 - RAID system 3-6
- ARCserve 6.5
 - for Windows NT 4-5, 4-11
 - features 3-1
 - restore options 3-7
- ARCserve capabilities
 - centralized, remote management 3-2
 - Database Agent 3-1
 - fault tolerance 3-2
 - information repository 3-1
 - Job Engine 3-1
 - Manger 3-1
 - parallel streaming 3-2
 - tape spanning 3-2
- ARCserve Engine Administrator 3-5
- ARCserve engines
 - database engine 3-5
 - job 3-5
 - tape 3-5
- ARCserve Job Engine 3-1
- ARCserve Manager 3-1, 3-2
 - operation options 3-2
- ARCserve operations
 - ARCserve RAID Engine 3-5
 - backup 3-2
 - Backup Agent for Lotus Notes 3-6
 - data migration 3-2
 - database 3-2
 - device management 3-2, 3-3
 - Engine Administrator 3-5
 - job status 3-2
 - media pool management 3-2, 3-4
 - reports 3-3
 - restore 3-2
- ARCserve Windows NT
 - Engine Administrator 3-5
- Autoloader 4-10
- Automated tape spanning 3-3

B

- Backup 3-2
 - basics 2-1
 - capacity 4-7
 - conclusions 5-22
 - high capacity 4-3
 - of all applications 2-3
 - of applications and data 2-3
 - of Public Name and Address Book 2-9
 - off-line 4-8
 - performance 2-6, 4-7
 - storage capacity 2-5
- Backup Agent
 - for Lotus 3-1
- Backup Agent for Lotus Notes V 2.0
 - features 3-6
- Backup basics
 - by file 2-4
 - choosing files for 2-3
 - complete 2-2
 - differential 2-3
 - Domino data backup 2-1
 - Domino implementation tips 2-8
 - frequency of backup 2-6
 - image 2-4
 - incremental 2-3
 - of data only 2-3
 - of shared mail database 2-9
 - off-line 2-2
 - on-line 2-2
 - overview 2-1
 - partial 2-2
 - restore strategy 2-7
 - strategy 2-4
 - terminology 2-2
 - user-defined 2-3
 - which data 2-5

- Backup conclusions
 - enterprise 5-22
 - small/medium business 5-22
- Backup files
 - off-site 2-8
- Backup performance
 - factors affecting 5-1
- Backup solution hardware 4-1
 - DLT drive technology 4-2
 - for enterprise 4-2
 - for small/meium business 4-6
- Backup solution software
 - ARCserve 6.5 for Window NT 3-1
- Backup strategy
 - points to consider 2-4
- Backup tapes
 - management of 2-7
- Backup testing
 - local 5-2
 - products used in 1-1
- Backup window 2-6
- Block size 5-15
- By file backup 2-4

C

- CA
 - AntiVirus Agent 4-11
 - AntiVirus Agent V2.0 for Lotus Notes 6-3
 - anti-virus software 6-3
 - ARCserve 6-3
 - InocuLAN *See* InocuLAN
 - InocuLAN 4 for Windows NT 4-11
 - Causes of data loss 2-1
-

- Centralized management 3-2
 - Common viruses 6-2
 - Compaq
 - 32-Bit SCSI-2 Controller 4-6
 - 35/70 DLT drive 4-4
 - DAT drives 4-1
 - DLT 15-Cartridge Library 4-5
 - DLT drives 4-1
 - hardware reference
 - documentation 1-3
 - Internet and Intranet solutions 1-2
 - ProSignia 200 4-11
 - ProSignia 200 with LISP 2-2
 - SLR drives 4-1
 - Software Support Diskette for Windows NT 4-6
 - solution hardware summary 4-12
 - system information on server offerings 1-2
 - technical publications 1-2
 - Wide-Ultra SCSI controller 4-4
 - Complete backup 2-2
 - Computer Associates
 - ARCserve 3-1
 - backup and virus protection product information 1-2
 - Computer Security Association *See* NCSA
 - Conclusions
 - about backup 5-22
 - about remote backup 5-21
 - Concurrent backup
 - multiple servers 5-19
 - Concurrent backup & restore
 - multiple servers 5-17
 - CRC
 - end-to-end 4-2
 - Critical databases
 - NAMES.NSF 2-8
 - Public Name and Address Book 2-8
 - Cyclical-redundancy-check *See* CRC
- ## D
- Daily media pool 3-5
 - DAT 4-1, 4-10
 - 2/4-GB 4-8
 - 2/4-GB DAT restore performance 4-11
 - Data
 - backup 2-5
 - compression 5-8
 - loss, causes of 2-1
 - migration 3-2
 - set 5-14
 - Database 3-2
 - Database Agent 3-1
 - Device management 3-2, 3-3
 - Differential backup 2-3
 - Digital Audio Tape *See* DAT
 - Digital Lempel-Ziv *See* DLZ
 - Digital Linear Technology *See* DLT
 - Digital TZ86, TZ87, DLT2000 driver 4-6
 - DLT 4-1
 - 10/20-GB 4-3
 - 15/30-GB 4-3
 - 35/70-GB 4-3, 5-2
 - device drivers 4-6
 - interface 4-4, 4-6
 - multiple checks 4-2
 - Tape Array cabinets 4-4
 - tape cartridges 4-3
 - tape format 4-3
 - Tape Library 4-5

Index -4

DLT backup and restore
 backup conclusions 5-22
 local testing 5-2
 performance results 5-1
 remote performance testing 5-16

DLT drive
 10/20-GB 4-6
 15/30 storage capacity 4-5
 15/30-GB 4-6
 35/70 4-4
 35/70 storage capacity 4-5
 35/70-GB 4-6

DLT7000.SYS 4-6

DLZ 4-3

Documentation
 for Lotus Domino 4.6 1-2

Domino
 data backups 2-1
 product information 1-2

Domino R4.5 4-11

Driver
 Quantum 4-6

E

ECC 4-2

Encrypted databases
 scanning of 6-9

End-to-end CRC 4-2

Engine Administrator 3-5

Error-correction-code *See* ECC

External interface 4-4

F

Fault tolerance 3-2
 through RAID 0 3-2
 through RAID 1 3-2
 through RAID 5 3-2

Feed speed 5-14

Floppy-drive protection 6-5

Frequency of backups 2-6

G

GFS backup 3-4

 Daily media pool 3-5

 Monthly media pool 3-5

 Weekly media pool 3-5

GFS tape-rotation 2-7

Groupware messaging anti-virus
 options 6-6

H

Hardware

 costs 4-8

 reliability 4-9

I

Image backup 2-4

Implementation

 Domino server backup 2-8

 restore strategy 2-8

Incremental backup 2-3

Information repository 3-1

InocuLAN

 feature and benefit summary 6-6

 features 6-5

 floppy-drive protection 6-5

 groupware messaging anti-
 virus options 6-6

 Internet-enabled 6-6

 network-drive protection 6-
 5

 real-time scanning mode 6-
 5

 virus quarantine 6-5

virus wall 6-5
purpose of 6-3
InocuLAN 4 6-4

Index -6

- Interface
 - external 4-4
 - internal 4-4
 - Internal interface 4-4
 - Internal parity check 4-2
 - Internet resources 1-2
 - Internet-enabled 6-6
- J**
- Job status 3-2
- L**
- Local backup testing 5-2
 - block size 5-15
 - data compression 5-8
 - data set 5-14
 - feed speed 5-14
 - performance factors 5-15
 - processor usage 5-12
 - RAID level 5-5
 - tape drive & SCSI Controller configurations combination 5-2
 - Lotus Domino 4.6 documentation 1-2
 - Lotus Notes 4.5, support 6-9
 - Lotus Notes Backup Agent 2.0 4-11
- M**
- Management
 - centralized, remote 3-2
 - of backup tapes 2-7
 - of media pool 3-4
 - Media
 - costs 4-8
 - Media pool management 3-2, 3-4
 - save set 3-4
- Media pool management *Continued*
 - scratch set 3-4
 - Mirroring 3-2, 3-6, 4-6, 5-5
 - Monthly media pool 3-5
 - Multiple servers
 - concurrent backup 5-19
 - concurrent backup and restore 5-17
- N**
- NAMES.NSF 2-8
 - NCSA
 - virus studies 6-3
 - Necessity of on-line virus scanning utilities 6-1
 - Network-drive protection 6-5
 - Notes Server
 - scanning encrypted databases on 6-9
 - NT SSD 4-6
- O**
- Objective of this TechNote 1-2
 - Off-line backup 2-2, 4-8
 - On-line
 - backup 2-2
 - virus scanning utilities 6-1
- P**
- Parallel streaming 3-2, 3-3
 - Parity
 - information 3-6
 - striping 4-6
 - Partial backup 2-2
 - alternatives to 2-3
 - Performance factors 5-15
 - Performance of backup data 2-6
-

Performance testing 4-11
 products used in 1-1
Processor usage 5-12
Production database replication 2-9
Protection Suite 2.0 for Lotus Notes
 6-11
Public Name and Address Book 2-8
 backup of 2-9

Restore by Tape 3-8

Q

QIC 4-1, 4-10
Quantum driver 4-6
Quarter Inch Cartridge *See* QIC

R

RAID 0 3-6
RAID 1 3-6
RAID 5 3-6
RAID Engine 3-5
RAID level 5-5
Real-time
 scanning and cure 6-9
 scanning mode 6-5
Reliability of hardware 4-9
Remote backup
 conclusions 5-21
 performance testing 5-16
Remote management 3-2
Reports 3-3
Resources
 additional 1-2
 for Lotus Domino 4.6 1-2
 World Wide Web 1-2
Restore 3-2
Restore options 3-7
 Restore by Image 3-8
 Restore by Query 3-8

Index -8

Restore options *Continued*
 Restore by Tape Media 3-8
 Restore by Tree 3-8
Restore strategy 2-7

S

Save set 3-4
Scanning
 of encrypted databases 6-9
 of shared mail 6-9
 real-time 6-9
Scratch set 3-4
SCSI Controller
 configuration combinations 5-2
Shared mail database backup 2-9
Single-channel Linear Recording *See*
 SLR
SLR 4-1, 4-10
 4/8-GB 4-8
Software solution
 anit-virus 6-3
Storage capacity 4-5
 of backup device 2-5
Strategy for backup 2-4
Striping 3-2, 4-6
 with no parity 5-5
 with parity 3-2, 5-5
Supplemental reading 1-2
Support
 for Lotus Notes 4.5 6-9

T

Tape drive
 10/20-GB DLT 4-3
 15/30-GB DLT 4-3
 2/4-GB DAT backup
 performance 4-11

Tape drive *Continued*
 2/4-GB DAT restore performance
 4-11
 35/70-DLT 5-2
 35/70-GB DLT 4-3
 4/8-GB SLR restore performance
 4-11
 alternatives 4-10
 backup and restore performance
 4-11
 configuration combinations 5-2
 DAT 4-1
 DLT 4-1
 QIC 4-1
 SLR 4-1
 storage capacity 4-5
Tape ID 3-4
Tape spanning 3-2
Tape-rotation 2-7
TechNote
 information on 1-1
 objective of 1-2
Terminology for backups 2-2

U

User-defined backup 2-3
Utilities
 anti-virus 6-2

V

Virus protection
 advanced 6-3
 AntiVirus Agent V2.0 for Lotus
 Notes 6-8
 anti-virus utilities 6-2
 CA anti-virus software 6-3
 InocuLAN 4 6-4

necessity of scanning utilities 6-1

Virus protection *Continued*
Protection Suite 2.0 for Lotus
Notes 6-11
security 6-1
summary of anti-virus protection
6-12
Virus quarantine 6-5
Virus scanning utilities
necessity of 6-1
Virus wall 6-5
Viruses
common 6-2
polymorphic 6-2

W

Weekly media pool 3-5
Wide-Ultra SCSI controller 4-4
World Wide Web on the Internet 1-2