

A First Look: Windows 2000 Tape Backup on Dell

By Charles Curtis and James Nall

Microsoft Windows 2000 incorporates tape backup software that will please many system administrators. The Windows 2000 Backup utility is simple, yet powerful, and works well to protect the individual server with directly connected tape drives, autoloaders, and libraries. This article describes key features of the Windows 2000 Backup utility and discusses some key considerations affecting tape backup performance with Dell servers.

The Dell Enterprise Storage Secondary Storage Development group has performed limited testing with Windows 2000 on an assortment of Dell servers, tape drives, and libraries. The group's mission is to perform certification and qualification of fully released tape backup products to ensure that they work well with Dell systems. It usually tests *only* production-level software.

Because of the significance of Windows 2000, however, the group began testing the system software to better understand its new features. This article describes the results of limited testing of the backup software embedded within Windows 2000. The tests used final code for Windows 2000.

In coming months, Dell *Power Solutions* will publish additional articles focusing on selected aspects of tape backup on Dell systems.

The Windows 2000 Backup Utility

The Windows 2000 Backup utility is a graphical backup utility that protects data on hard

disk drives by transferring a copy of that data to magnetic tape, recordable CD-ROM drives, or removable disks. The utility was designed to offer an effective single-server safety net. All testing was done using directly connected tape drives. This software does not currently support the Dell storage area network (SAN) connections. Figure 1 shows the Welcome screen for the Windows Backup utility.

Because of the significant new features in the Windows 2000 operating system, most current backup software will not work correctly unless it has been certified for Windows 2000. The differences between Windows 2000 and older versions of the operating system may not be visible immediately, but backup software companies must deal with them. Backup application software must be specifically designed to account for the improved architecture of Windows 2000.

The Windows 2000 Backup utility is best suited to running a single backup to a single drive.

Features Set

The Windows 2000 Backup utility does not equal the level of sophistication offered by the

major players in the backup industry, although it can certainly create and restore a complete backup of any single server. It includes fewer features than the new Windows 2000 versions of VERITAS® Backup Exec® and Computer Associates ARCserveIT.

The Windows 2000 Backup utility will be effective for many environments. The VERITAS and Computer Associates products should be considered for the following situations:

- Higher performance
- Multiple jobs
- Redundant Arrays of Independent Tape (RAIT)
- Redundant Arrays of Independent Libraries (RAIL)
- Open files issues coverage for non-Microsoft applications
- Database agents for SQL Server, Oracle, or SAP
- Back up other systems over the network
- Use of backup agents

Microsoft has included with the Backup utility a simple but effective Task Scheduler service, shown in Figure 2. This flexible arrangement makes it possible to schedule backups over the weekends or during periods when system activity is at its lowest.

When running a backup, it is important to note that a dialog box will ask the user to “Set Account Information.” To determine the backup schedule, enter the name and password of an account with the correct account security permissions for directing the backup job. The selected account must have the Backup Files and Directories privilege. If not set properly, the backup job will not have the rights needed to back up the files; thus, the job becomes ineffective.

The Windows 2000 Backup utility is a graphical backup utility that protects data on hard disk drives by transferring a copy of that data to magnetic tape, recordable CD-ROM drives, or removable disks.

Windows NTFS 2000 Volumes

The structure of the new Windows 2000 NT File System (NTFS) volume has many new features, such as disk quota information, mounted drive information, remote storage information, additional permissions, and encrypting file system (EFS) settings. The Backup utility handles these additional functions automatically.

Microsoft cautions users to be aware of these features and to realize that backup data should be restored to a compatible file system. Attempting to restore a Windows 2000 NTFS volume to a Windows NT Server 4.0 NTFS or a file allocation table (FAT) volume would result in loss of attribute information and possibly the loss of the data itself.

System State Data

The components collectively called System State Data are central to the Windows 2000 operating system architecture and must be handled carefully by any backup program that is used. These components include the following items:

- The system startup files (ntldr and ntdelect.com) and the Registry
- The SYSVOL folder (a shared, hidden structure that stores scripts and Group Policy Objects for the current domain and the enterprise), which exists on all Windows 2000 domain controllers
- The databases at the core of the operating system, which are:
 - The COM+ Component Services Class Registration database (Common Object Model is a method of linking large and sophisticated application program pieces together)



Figure 1. The Welcome Tab

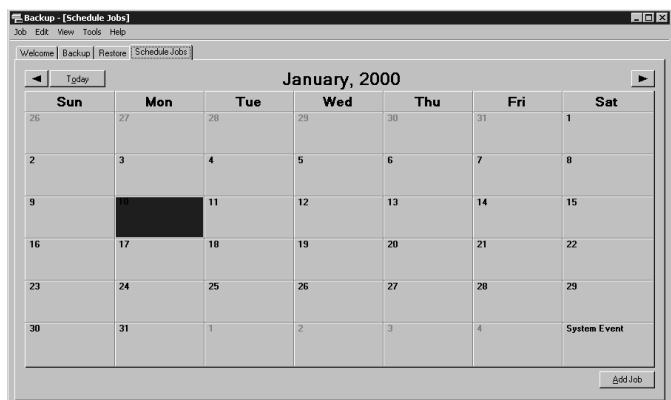


Figure 2. The Schedule Tab

- The Certificate Services database, which stores the authentication and secure exchange data needed to communicate securely on the Internet
- The Active Directory services (which stores network resource information, such as databases, groups, security policies, servers, and other objects)

The Windows 2000 Backup utility handles these components as part of its design. This is a formidable task, in that it must use special system calls (instructions) that permit access to this data while the system is running. All System State Data is backed up automatically if the System State Data box is checked.

Restoring Data

The restore function permits users to easily restore folders, files, or the contents of an entire drive to the original location or to an alternative location, including restoration to a single folder. Users can also choose whether or not to replace existing files, which is a safety feature that helps to prevent the loss of existing data.

Another feature is the ability to restore Active Directory services, which can be useful during replacement of a failed domain controller. Active Directory is a structure that Windows 2000 uses to track and locate any object on the network. One notable feature is that after Active Directory services are restored, Windows 2000 will automatically update the directory services and the File Replication Service with data from their replication partners. It also will perform a consistency check and re-index the Active Directory database. This kind of intelligence has been missing from previous products, and it is just one example of how well Microsoft understands what administrators really need to recover quickly after an outage.

Removable Storage Manager

The Removable Storage Manager is a media management tool that works with the Backup utility to access and track tape and CD media. It is found under Programs -> Administrative Tools -> Computer Management.

This application is complex, as evidenced by the 30+ pages of help text we had to read before using it. Consequently, this article covers only its major features.

The Removable Storage Manager is a media management tool that works with the Backup utility to access and track tape and CD media.

The Removable Storage Manager permits users to:

- Enable or disable a library
- Enable or disable a tape drive
- Change media in the library
- Inventory the library
- Open a library door
- Set door and port time-outs
- Mount, dismount, prepare, and eject media
- Manage operator requests and the work queue
- Manage security for media and permissions
- Clean a stand-alone tape drive or a library drive
- Manage media pools, identification, and media states

Notes:

- Users must create media pools before using the Removable Storage Manager. The process is described in the Checklist: Configure Removable Storage provided in the Help file for Removable Storage.
- The help text says, “Keep your media counts under 1000.” Our testing suggests that users should follow this advice. This should not present any capacity problems, considering that 1,000 DLT tapes (native capacity) would equate to 35 terabytes (35,000 gigabytes).

Remote Storage Manager

The Remote Storage Manager application is used to migrate files from managed volumes to increase effective free space. This application uses the Removable Media Manager and can control its own assigned tape drives, libraries, or both. This article describes only the highlights of this complex feature. It seems to be a powerful tool, and it should be among the first new features in Windows 2000 that storage managers test and investigate.

In our preliminary tests, the ease of use and actual performance of the Remote Storage Manager were impressive. When we requested that it return a drive so that we could test additional backups, it reported the amount of data on that particular tape and offered us an opportunity to migrate the data back onto the hard drive before it unmounted the tape. The process of migrating the data and unmounting the tape worked flawlessly in our limited tests.

The databases at the core of the operating system are the Component Services Class Registration database, the Certificate Services database, and the Active Directory services.

Windows 2000 Backup Utility in Operation

The Dell Enterprise Storage Secondary Storage Development group installed Windows 2000 Server on several Dell servers and encountered only one problem: Windows 2000 reported in one case that it could not find a device driver for the QFDDMM 2000 tape drive (should have been Quantum® 7000). We investigated further and discovered that the server had a bad SCSI cable with two bent and shorted pins in one cable head. The problem disappeared when we replaced the cable, deleted the device, and allowed Windows 2000 to discover it a second time.

To run Backup utility, simply select:

Start -> Programs -> Accessories -> System Tools -> Backup

Browse for and select the files, drives, or folders to back up, and then click on the **Start Backup** button. An even easier option is to use the Backup and Restore Wizard supplied with Windows 2000. Microsoft has also placed the Emergency Repair Disk option in the same folder for easy access to utilities with similar functionalities.

We determined some configuration options that we recommend for the various tabs under the Tools menu, as shown in the sidebar on page 43.

The type of backup to be performed or scheduled to be performed will depend on your desired recovery goals. If you have time and tape space for a full backup each day, then the normal backup run daily is the best option. If not, some combination of normal, differential, and incremental backups is appropriate.

Differential versus Incremental

Suppose your backup strategy is to make a normal backup on Sunday and an incremental backup each night of the week for the next six nights. A data drive crashes on Thursday morning and recovery is necessary. To do this, replace the drive and run a restore of that drive data from the Sunday tape with the full backup. Then restore the changes from the Monday, Tuesday, and Wednesday tapes to make it current. This is necessary because incremental backup tapes contain only the changes from one day to the next.

A differential strategy would make the normal backup on Sunday, then perform differential backups each night. Using the previous scenario, restore the Sunday tape, then the Wednesday tape, because each differential is a backup of everything that has changed since the last normal (full) backup.

The determining factor in choosing the appropriate technique is based on how much time, tape, and changing data is in the environment, and how fast the restore must be handled. If all three are in abundance, a full backup each day is best. Some rotation of media off-site for safety and duplication is mandatory.

Test Results

We compared the results of hundreds of hours of testing over the past year on Dell systems, mostly with Windows NT Server 4.0, with the results of preliminary testing of backups with Windows 2000. Some test results are described below.

Test 1: Single Backup Job to a Single Drive

The first test was to back up 1 TB to a Dell 130T tape library—a single backup job to a single drive test. The objective was to verify that the Windows 2000 Backup utility could handle a large number of files and a very large amount of data.

The challenge in creating this test was that our system had only four physical hard drives totaling 36 GB of storage and the system was not on a network. The solution was to take advantage of one of the new features of Windows 2000: a mount-point function similar to the function provided on UNIX-like operating systems.

This new feature allows Windows users to exceed the “letters of the alphabet” limit. Using Disk Management (found under Computer Management), select the **G:** drive, and then select the **Change Drive Letter** and **Paths** option. Select **Add** and choose an NTFS folder in which to mount the drive. Once an NTFS folder has been selected, the G: drive letter can be removed from the volume. With the letter removed, Windows 2000 can still access a hundred

volumes in a folder under one drive letter. In other words, if we do not have terabytes of storage space available, we can fake it.

To complete the test, we created a large number of folders and placed the same 18 GB drive over and over again into each folder as a mount point (without a drive letter). We created a dummy set of repeated links back to that single drive from each of those folders, and we were able to complete the test successfully. This approach saved days of creating the data, setting up the real storage structures, and so on. We would never have discovered this possibility if we had not checked the box under the General tab labeled “Back up mounted drives.”

The restore function permits users to easily restore folders, files, or the contents of an entire drive to the original location or to an alternative location, including restoration to a single folder.

RECOMMENDED CONFIGURATION

We tested a variety of configuration options for the various tabs under the Tools menu. Figures 3, 4, and 5 show several tabs in the Tools menu. Our preliminary recommendations include:

General Tab

- Compute selection information (desirable to save time)
- Use the catalog on the media to speed restores (desirable)
- Verify data after the backup completes (mandatory)
- Back up mounted drives (varies with the environment)
- Show alert if Removable Storage (RS) not running (desirable)
- Show alert if compatible import available (desirable)
- Show Alert message when new media appears in RS (desirable)
- Always put new media into the Backup media pool (questionable)

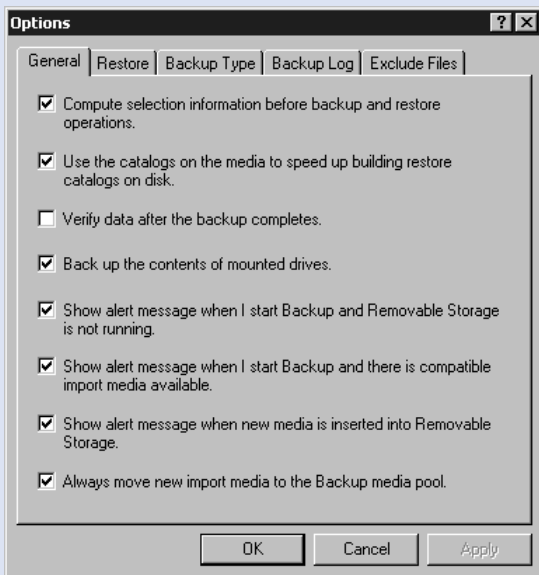


Figure 3. The Options Tab

Restore Tab (when restoring a file that is already on the computer)

- Do not replace the file on my computer (recommended)
- Replace file on disk only if the disk file is older (questionable)
- Always replace the file on my computer (dangerous)

Backup Type Tab

- Normal (Full)—backs up all files and clears (resets) the archive bit to mark the files as having been backed up (recommended)
- Copy—backs up all files and does clear the archive bit. A copy backup is slightly faster than a normal backup because it does not need to touch each file attribute to clear the archive bit; sometimes used as a performance technique

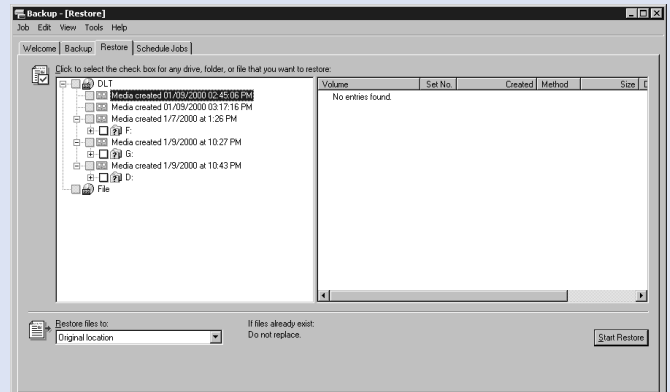


Figure 4. The Restore Tab

- Differential—backs up all files that were changed or created since the last normal backup (backs up files that have the archive bit on), and does not clear the archive bit
- Incremental—backs up files that have been accessed, changed, or created since the last normal or incremental backup; clears the archive bit after a file is backed up
- Daily—backs up only files that have changed today, but does not clear the archive bit

Backup Log Tab

- Detailed—logs *all* information, including the names of all the files and folders (log size could be excessive)
- Summary—logs only key operations (recommended)
- None—logs no information (dangerous)

Exclude Files Tab

- Identify files to be excluded for all users (used to exclude such items as folders containing application software that is duplicated on CD-ROM and not needed)
- Identify files to be excluded for the Administrator user account only (as needed)

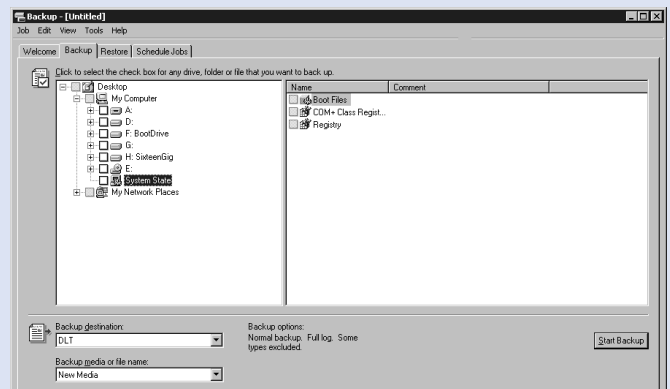


Figure 5. The Backup Tab

TEST NO. 1: WINDOWS 2000 BACKUP UTILITY

Test Parameters

- 1 TB (1000 MB) of data
- 255,000 files (assorted CDs copied onto a single 18 GB drive)
- 20 DLT type IV tapes (capacity achieved was about 50 GB on each)

Test Environment

- Dell 4100 PowerEdge Server
- Two 200 MHz Pentium Pro chips and 192 MB RAM
- Memory utilization was typically 111 MB
- CPU utilization before backup (idle) less than 2 percent

Results

- Time required: 54 hours, 25 minutes
- Average backup rate: 5.6 MB/sec (338 MB/min, 19.8 GB/hr)
- CPU utilization during backup ranged from 26 percent to 33 percent

Figure 6. Test No. 1, Windows 2000 Backup Utility

TEST NO. 2: WINDOWS 2000 BACKUP UTILITY

Test Parameters

- Back up a single 4 GB file that measures as 1:1 (not compressible) onto a DLT type IV tape on a DLT7000 tape drive
- Back up a single 4 GB file that measures as 2:1 compressible onto a DLT type IV tape on a DLT7000 tape drive

Test Environment

- Windows 2000 Backup utility
- Dell 4100 PowerEdge Server
- Two 200 MHz Pentium Pro chips with 192 MB RAM
- F/W SCSI used to access a single IBM 18 GB 10K drive
- Memory utilization typically 111 MB
- CPU utilization during backup approximately 30 percent

Results

Test	Backup Rate		
	MB per Second	MB per Minute	GB per Hour
1 Job, 1:1 Data	4.6	278	16.3
1 Job, 2:1 Data	6.1	367	21.5

Figure 7. Test No. 2, Windows 2000 Backup Utility

This tab permits you to exclude mounted drives. This function would normally be left unchecked to avoid the typical server backup from including the data residing on other remotely (LAN) connected systems in its normal backup. This is a simple, yet smart option in this backup utility.

Figure 6 summarizes the test results and testing environment. The utility did not achieve exceptional speed, but the backup rate was reasonably good for a single job using a single drive.

Test 2: A Single Job to a Single Drive

Figure 7 shows the typical results obtained in our tests—a single job to a single drive—with the Windows 2000 Backup utility. This test was to determine the results of compressible versus noncompressible data.

Test 3: Single and Multiple Jobs and Drives

Although the results given in Figure 7 indicate acceptable performance, additional testing under Windows 2000 with one of the major commercial backup products revealed significantly better performance (see Figure 8). Furthermore, the Windows 2000 Backup utility is best suited to running a single backup to a single drive. The beta version of the major vendor's backup software not only performed much better in general, but could also handle up to six jobs in parallel.

The ideal performance for any individual 2:1 data backup would have been a little over 30 GB per hour; 28.2 GB per hour is close to this ideal. Although increasing the number of jobs somewhat increases the overall backup rate, our use of a single SCSI channel within this particular server as the primary storage source limited the performance. An ideal rate for four jobs would be about 120 GB per hour with a faster primary storage source, such as a Dell PowerVault Fibre Channel-based storage system.

After Active Directory services are restored, Windows 2000 will automatically update the directory services and the File Replication Service with data from their replication partners.

Windows 2000 Backup Does the Job

Our first look at the embedded backup software within Windows 2000 was very pleasing and should be adequate to roll Windows 2000 quickly into production. Although this application is not the most elegant or versatile backup

TEST NO. 3: BETA VERSION OF A MAJOR BACKUP PRODUCT

Test Parameters

Perform one to four backup sessions using the following datasets:

- A single 4 GB file that measures as 1:1 (not compressible)
- A single 4 GB file that measures as 2:1 compressibility

Test Environment

- Beta version of a major vendor's backup software
- Dell 4100 PowerEdge Server
- Two 200 MHz Pentium Pro chips with 192 MB RAM
- F/W SCSI used to access four Dell DLT7000 drives on a single SCSI channel
- Drives 1–3, Seagate® 7200 RPM, 9 GB
- Drive 4, IBM 18 GB 10K
- Memory utilization of 110 MB to 160 MB (increased based on number of jobs)
- CPU utilization: 1 job at 33 percent; 4 jobs at 82 percent

Results¹

Test	Backup Rate (Local SCSI-Connected Tape)		
	MB per Second	MB per Minute	GB per Hour
1 Job, 1:1 Data	4.6	277	16.1
1 Job, 2:1 Data	8.8	481	28.2
2 Jobs, 2:1 Data (simultaneously)			
Job 1	7.6	455	26.7
Job 2	6.9	413	24.2
Both Jobs	14.5	868	50.9
4 Jobs, 2:1 Data (simultaneously)			
Job 1	6.4	385	22.6
Job 2	6.4	384	22.5
Job 3	6.5	388	22.7
Job 4	7.2	431	25.3
Totals	26.0	1588	93.0

¹The datasets used in these tests were created specifically for testing Dell 130T libraries that use the Quantum DLT7000 drives. The primary tests for maximum performance use one of two possible 4 GB files. The DLT7000 drive reports exactly 1:1 and 2:1 compression on these files. This was done to simulate the two capacity numbers given for DLT tape, 35 GB and 70 GB of capacity, based on the 1:1 and 2:1 compression ratios.

Figure 8. Test No. 3, Beta Version of a Major Backup Product

application, it performs its intended functions reliably when limited to a single job and a single drive at a time.

We suggest that although this product will certainly be able to handle a large number of server installations with directly connected tape drives or libraries, there will still be many installations that would greatly benefit from the utilization of a more sophisticated backup application.

The need for multiple jobs or performance backups using RAIT and RAIL as well as the need for specific database applications support will all require a more sophisticated backup application. Look for the new Windows 2000 version of VERITAS Backup Exec and Computer Associates ARCserveIT as the leaders in this arena.

A Final Hint

Test your restores. Backups mean nothing—only successfully tested restores matter. ♦

Test your restores.

**Backups mean
nothing—only
successfully tested
restores matter.**

Charles Curtis (charles_curtis@dell.com) is a software engineer and senior consultant in the Secondary Storage Development area of the Dell Enterprise Systems group in Round Rock, Texas. He has worked on computers and in the storage industry for over 25 years. He has been involved with storage solutions for PC, midrange, and mainframe systems. Charles has published numerous white papers about backup performance tuning and disaster recovery.

James Nall (james_nall@dell.com) is the product marketing manager for tape in the Storage Marketing unit of the Dell Enterprise Systems group in Round Rock, Texas. In his 30 years of working with computer hardware and software, Jim has held positions as a systems engineer, field salesperson, and marketing development manager for IBM, Microsoft, Memorex Telex®, Compaq®, and Dell.