

Streamlining Backup and Recovery Operations

Using Disk-based Protection

IT administrators are straining to protect massive amounts of data in the face of ever-shrinking backup windows. This article examines how VERITAS Backup Exec™ 9.1 *for Windows Servers* and VERITAS NetBackup™ 5.0 software can enable IT organizations to implement a disk-based data protection strategy that helps improve backup and recovery times while increasing system availability.

BY SCOTT KOSCIUK, MICHAEL PARKER, AND MARK THOMASON

To help meet stringent system availability requirements, administrators now can enhance traditional tape-only backup operations using fast, flexible disk-based backup and recovery techniques that do not encroach on business-critical applications. This approach can enable IT organizations to better maintain service level agreements (SLAs) and to help create more responsive, cost-effective data protection and disaster recovery strategies.

A high-performance data protection strategy integrates both disk and tape storage with an optimized backup application that can streamline backup and recovery operations. Although disk-based data protection is not likely to replace tape drives and tape robotics completely, it can enable more efficient backups and recoveries. For example, administrators can perform backups quickly from a primary disk to a backup disk and then copy the data from the backup disk to tape for long-term or off-site storage. In addition, the ability to restore data from snapshots that reside on the primary or backup disk can provide near-instantaneous data recovery.

Comparing disk-based to tape-based data protection

Although the throughput and capacity of tape devices has become competitive with that of disk drives over the past few years, tape is still a sequential-access medium and as such can be inflexible and cumbersome compared to disk media. Moreover, disk-based storage avoids mechanical delays that are inherent in tape libraries or devices, such as tape mounting, positioning, and availability. Given the fast random-access read performance of disk volumes—especially RAID volumes—administrators can achieve near-instantaneous disk backups and restores by leveraging *snapshots*, which provide a point-in-time image of a client's data on local or remote disk storage.

Using disk-resident snapshots, tape backups, or both, administrators can reduce network backup windows and free host CPU and I/O cycles to process business-critical applications. VERITAS® software enables administrators to back up data from the snapshot image, instead of directly from the client's primary data—thereby allowing client operations and user access to continue without

interruption during the backup. In addition, fast disk-based reads and writes allow administrators to schedule more frequent backups, which help to improve data protection by lowering the incidence of data loss.

When used as backup devices, disks can support simultaneous backup, recovery, and duplication operations. To back up multiple sources to a single tape drive, administrators traditionally use a multiplexing, or *interleaving*, approach. Multiplexing consolidates multiple streams into one stream while writing to a tape drive. This approach keeps the tape drive spinning, rather than starting and stopping to wait for additional data. Although multiplexing can greatly increase tape device efficiency, its main disadvantage can be slow recovery operations. Disk-based storage eliminates the need for multiplexing because disks are inherently random access devices.

Given the benefits of disk-based storage, tape may appear to have little future in data protection. However, tape still provides the best medium for long-term and off-site storage, which can make tape an important consideration for disaster recovery and business continuance planning. In addition, tape technology offers benefits that disk media has not yet achieved, such as:

- **Greater durability:** Tapes can survive large drops and tolerate rough handling better than disk drives, which need to be carefully packed for shipping.
- **Lower unit cost:** New large-capacity tapes are bigger and less expensive than today's disk drives, providing a lower cost per megabyte.

Understanding disk-based backup and recovery methods

Five methods of disk-based data protection are currently available: backup to disk, disk staging, inline copy, synthetic backup, and instant recovery (see Figure 1).

Backup to disk. The backup-to-disk approach writes the same data to a file on a disk volume as it would to a file on a tape volume. Therefore, when a backup-to-disk operation completes, a single file the size of the backup will exist on the target volume that contains all the files that were backed up. Products such as VERITAS Backup Exec™ 9.1 for Windows Servers and VERITAS NetBackup™ 5.0 software provide administrators with a graphical user interface (GUI) from which to configure backup-to-disk

Data protection application	Backup to disk	Disk staging	Inline copy	Synthetic backup	Instant recovery
VERITAS NetBackup 5.0	Yes	Yes	Yes	Yes	Yes
VERITAS Backup Exec 9.1 for Windows Servers	Yes	Yes	No	No	No

Figure 1. Support for disk-based data protection strategies

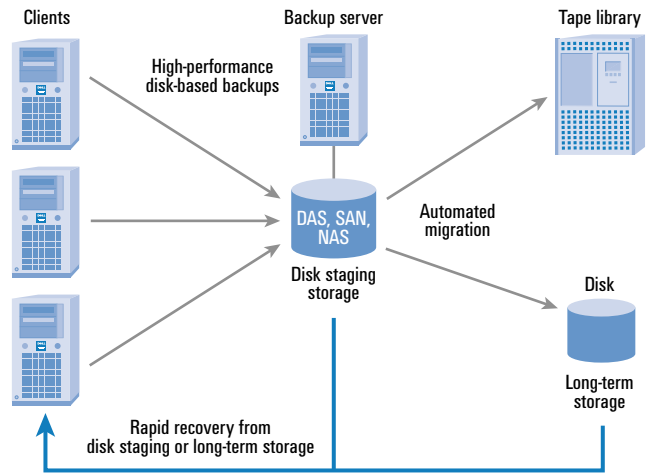


Figure 2. The disk staging process

policies and storage designations, and to perform full, archival backups as well as incremental or differential backups. Both products work with network attached storage (NAS), storage area networks (SANs), and direct attach storage (DAS).

Disk staging. The disk staging method writes backup data to a disk cache before sending the data to its final destination, disk or tape (see Figure 2). The purpose of disk staging is to use all available media to best advantage. For example, when staging a backup, administrators first copy the target data onto the disk cache and later move the backup image to tape according to the established disk staging schedule. Disk staging enables administrators to complete backups faster, shortening the backup window and thereby affecting business applications less than a direct backup-to-tape method.

The backup data remains in the disk staging storage unit until either the backup expires, based on the administrator-specified retention period, or another backup needs space in the disk staging location. When a backup application such as VERITAS NetBackup software detects a full disk staging location, it pauses the backup process, finds the oldest backup image that has been copied successfully to the final long-term storage destination, and deletes that data from the staging storage unit. For rapid restores, NetBackup can retrieve undeleted data from the disk staging storage unit before requesting the data from the secondary, long-term storage location.

Inline copy. The inline copy method writes backup data simultaneously to multiple destinations, such as disk and tape. Before the inline copy approach was introduced, administrators were required to duplicate backup application data as a secondary process, after they finished the initial backup from the client. By combining the backup and duplication operations, the inline copy approach can enable administrators to make an organization's electronic

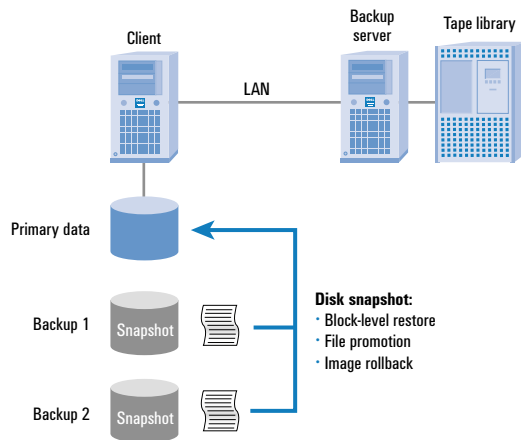


Figure 3. The instant recovery process

data-vaulting procedures more flexible and efficient, conserving time and IT resources.

Synthetic backup. When system availability requirements do not allow enough time to complete a full backup, administrators can create a synthetic backup that is identical to a current full backup. A synthetic backup assembles data from the system's previous full backup and subsequent incremental backups without involving network resources—the synthetic backup is completed on the media server. The ability to create synthetic backups can help administrators meet regulatory or SLA requirements for frequent full backups when performing actual archival backups over the network is too expensive or resource-intensive.

To create a synthetic backup, selected data is copied from the initial full backup and subsequent incremental backups to disk or tape as a single backup image. To construct a synthetic backup optimized for high-performance recovery operations, the VERITAS NetBackup software tracks not only which files were present on the volume for each backup, but also which data was moved or deleted

since the last full backup. A disk-based approach can dramatically increase the usability and efficiency of synthetic backups. For example, using disk storage for smaller, more frequent incremental backups can help administrators reduce, if not eliminate, the requirement to mount a large number of tapes.


Instant recovery. The instant recovery approach helps deliver the benefits of disk-based data protection without the need to perform backups or recoveries over the network (see Figure 3). By restoring

data from snapshots residing on a local disk, administrators quickly can resolve application corruption and end-user errors such as accidental deletions and overwrites. A scheduled backup creates a local snapshot on the client's disk as a background task, without interrupting the end user's access to data and without moving the backup data across the network to a backup server.

The instant recovery process enables administrators to perform three different types of high-performance recovery operations:

- **Block-level restore:** Moves only blocks that have changed since the client's primary file set was backed up.
- **File promotion:** Restores a file from one of the disk volume snapshots to the original volume. Data that has undergone several changes since the last backup can be recovered more quickly using file promotion than block-level restore.
- **Image rollback:** Instantaneously restores the entire volume to a previous state and time.

Improving backup and recovery performance

Products such as VERITAS Backup Exec 9.1 *for Windows Servers* and VERITAS NetBackup 5.0 software can help administrators implement a flexible, efficient data protection and disaster recovery strategy. Disk-based backup and recovery operations help administrators improve performance substantially over traditional sequential tape-based backups while offloading the host CPU to help increase system availability for business-critical applications. 

Scott Kosciuk (scott.kosciuk@veritas.com) is a product marketing manager for VERITAS NetBackup.

Michael Parker (michael.parker@veritas.com) is a product marketing manager for VERITAS Backup Exec.

Mark Thomason (mark.thomason@veritas.com) is a technical product manager for VERITAS Data Protection Solutions.

VERITAS Software Corporation (<http://www.veritas.com>) is a leading storage software company, providing data protection, application performance, storage management, high availability, and disaster recovery software.

FOR MORE INFORMATION

VERITAS Backup Exec *for Windows Servers*:

<http://www.veritas.com/products/category/ProductDetail.jhtml?productid=bews>

VERITAS NetBackup:

<http://www.veritas.com/products/category/ProductDetail.jhtml?productid=nbus>

VERITAS VISION Utility Computing Conference, May 3–7, 2004:

<http://www.veritas.com/vision>