VERITAS™

VERITAS Storage Foundation for Windows FlashSnap Option

Snapshot Technology for Microsoft Windows Server 2000 and Windows Server 2003

August 13, 2004

VERITAS™

TABLE OF CONTENTS

Introduction

In today's network-centric business environment, profitability is increasingly dependent on productive, continuous access to enterprise applications and data. Snapshot technology helps businesses meet this availability demand by creating local, disk-based copies of data volumes as they existed at a particular point in time. These copies, known as snapshots, are stored separately from the production volumes and administrators use them for the fast restore of lost or corrupt files, maximizing end-user access to data. Snapshots are also available for transport to a different server, allowing resource-intensive applications to be run off the primary host. This helps to maintain the high-performance requirements of critical production applications.

Fast Data Recovery

Disk-based, locally stored snapshots support extremely fast data recovery because they eliminate the time and effort required to restore data from tape, or from a disk copy created through replication or disk-to-disk backup. For example, snapshots allow administrators to recover an entire Exchange Storage Group from the most current point-in-time copy, mount the database, and restore user service in less than 90 seconds. The snapshot functions like a full backup, but supports the high uptime requirements of today's enterprise service level agreements (SLA).

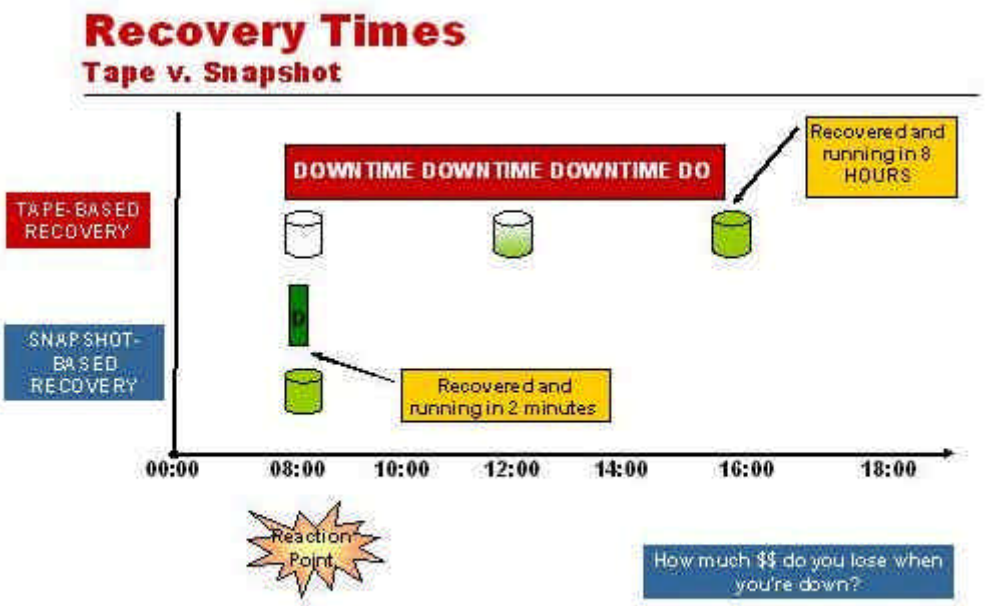

Figure 1: Restoring from a snapshot eliminates the time spent restoring from tape, speeding recoveries significantly.

Off-Host Processing

Some business-critical, resource-intensive applications such as backup, testing, and reporting do not require up-to-the-moment data currency. Moving these applications off the production server maximizes host system performance and ensures continued operations if the host system becomes unavailable. Running enterprise applications on a separate host using a snapshot of production data is known as off-host processing.

Backup

At a minimum, backup processing degrades the performance of other applications running on the server, and frequently requires a dedicated window of time when the system is unavailable to end-users. Snapshot technology allows the backup to run against a copy of the production data on a different server. This removes the performance overhead of backup processing from the production server and eliminates the need for a backup window.

Creating snapshots for off-host backup requires very little overhead and does not interfere with normal system operations. As a result, administrators can schedule snapshots frequently, minimizing potential data loss during a disaster. Off-host backups are a particularly good solution for organizations with stringent high availability requirements and large amounts of data.

Testing

Unit, system, regression, quality assurance (QA), and disaster recovery (DR) testing each require access to data that reflects the real-world of the production environment. But it is seldom possible to test applications on the production system — aside from the possibility of corrupting live data, the potential to negatively impact performance of critical business applications makes this scenario unlikely.

Off-host testing, using a snapshot of production data, allows developers and testers to run applications against real-world data without the integrity and performance constraints of working on a live production system. For example, a snapshot of production volumes in the evening allows QA engineers to run tests, locally or remotely, on data acquired the previous night.

Reporting and Analysis

Decision support applications often read through large numbers of database records to perform the analysis business users demand. This activity inevitably conflicts with the rapid-fire read and write requests of an online transaction processing (OLTP) system. Data Warehousing systems are designed to avoid this type of performance conflict by loading data into a separate database designed and tuned for analytic queries. However, the extract, transformation, and load (ETL) process that populates the warehouse must still read through the production database, potentially impacting the performance of other business applications.

Snapshots allow Data Warehouse ETL processing to run against a copy of the production database, instead of the live data. Taking a snapshot has minimal impact on the production system, so snapshots can generate fresh data for decision support applications as often as needed. This minimizes application performance degradation on the production system and maximizes the data currency for decision support applications.

Hardware versus Software Snapshots

Vendors of high-end storage devices include snapshot or third mirror split capabilities in the storage array hardware. However, array-based snapshot solutions must be created on or between hardware from a single vendor. This means that if network-wide snapshot capabilities are to be retained, IT purchasing options are limited to storage arrays from one vendor.

Array-based snapshot functionality also increases training costs as operators must be trained to use each vendor's unique solution. The inability to share snapshot data between arrays from different vendors also prevents storage administrators from leveraging the feature across other parts of the IT infrastructure. This limits off-host data sharing options because only servers attached to the array where the snapshot is stored can access the data.

In contrast, host-based snapshot solutions, running as software installed on the server platform, provide heterogeneous storage and server support, allowing snapshot data to be shared across the entire network. Copied volumes can span

storage arrays from multiple vendors, and both UNIX and Windows servers are supported. Host-based snapshot solutions offer greater flexibility, adapt easily to changing needs, and reduce administration costs.

Microsoft Windows Server 2003 VSS-enabled Snapshots

With the release of Windows Server 2003, Microsoft introduced a new storage management technology called Volume Shadow Copy Services (VSS). VSS offers a framework for applications and storage to perform coordinated data snapshots. Because the business application (VSS writer) storage (VSS provider) and backup application (VSS requestor) are tightly coordinated Microsoft is able to eliminate the need for a snapshot quiesce point while guaranteeing data integrity.

Microsoft's VSS technology allows backups to include all open files. Database files, held open exclusively by the database system, and files that are open due to operator or system activity can be successfully secured with a VSS backup. Any files that are open when the shadow copy is performed appear closed on the completed shadow copy volume. This ability to backup open files ensures the completeness of the backup — no files are skipped due to being open — and eliminates end-user downtime during the backup window.

Copy on Write

The Windows Server 2003 VSS feature includes a copy on write component. The VSS copy on write provider creates a non-persistent snapshot of an original volume when a backup is initiated.

Copy on write is different from traditional split mirror snapshot techniques because only changed data is maintained on the snapshot volume. As applications access and modify data the copy on write provider records pointers to changed data blocks and the changed data in a shadow copy volume. Backup processing is performed against both the shadow copy and the original volume. Unchanged data blocks are copied from the original location and changes are taken from the shadow copy volume. The backup then represents a consistent image of the data at the point in time of the snapshot. Applications and services continue running during the backup process, allowing backups to be performed at any time without interrupting users.

Copy on Write versus Traditional Snapshots

Copy on write snapshots take less time to create and demand less storage capacity than split-mirror snapshots, which require an amount of disk space equivalent to the original volumes contents to store each mirror. On average, a copy on write snapshot requires ten percent of the original volume's capacity — or a minimum of 105 MB — plus space for changed data. However, because copy on write snapshots represent changes to captured data, and not the full mirror of the original volume, they require an intact copy of the original volume to perform a successful restore. If the original volume becomes corrupt or unavailable, the snapshot cannot be used.

The process of copying data from the original volume during a copy on write snapshot adds overhead to the system being backed up. Split-mirror snapshots are independent, physical data copies and, as such, are immune to problems incurred by the original volume and do not affect source volume performance.

VERITAS Storage Foundation for Windows FlashSnap Option

VERITAS Storage Foundation for Windows™ provides mission-critical, enterprise Windows environments with sophisticated storage management features that reduce planned and unplanned outages, dramatically increasing end-user access to data. Storage Foundation maximizes administrator productivity, ensures high application data availability, optimizes storage I/O performance, and provides the heterogeneous support needed to protect current and future investments in storage technology.

VERITAS FlashSnap™ is a Storage Foundation option that creates independently addressable point-in-time copies of server volume mirrors, known as split-mirror snapshots. FlashSnap supports Windows Server 2003 Volume Shadow Copy Service (VSS): Microsoft's approved and supported snapshot backup technology.

Volumes of any size containing data of any type, including Storage Foundation logical database volumes, can be copied to a snapshot on any storage hardware. Heterogeneous operating system and storage platform support leverages current storage investments and provides a flexible environment for future expansion. Once created, snapshots can be easily moved to another server for backup or off-host processing, without affecting normal server functionality.

One of the most popular applications for snapshot technology is the very rapid data recovery that is possible when snapshots are stored on the same server as the original mirror. Data can be instantaneously recovered to the point of the last snapshot by restoring files to a point in time immediately before a failure.

Snapshots can be updated with changes to the primary storage volumes using VERITAS FastResync technology. FastResync streamlines mirror resynchronization by continually tracking changes made whenever the mirror copy is offline. The FastResync option exclusively reapplies updates to a mirror copy whenever the mirror is reattached to the original data volume. This lessens the time required to rejoin a split mirror with the mirror set, as well as reducing the server CPU cycles needed to complete the resynchronization. FastResync is ideal for data recovery, decision-support environments, data replication, batch processing, data warehousing or other applications that are typically performed offline to avoid impacting production servers.

Dynamic Disk Group Split and Join

The Dynamic Disk Group Split and Join (DGSJ) function is similar to a split-mirror snapshot, but allows entire groups of disks to be logically split from one group and joined to another. This feature is particularly useful for off-host processing where host applications demand the highest availability and perform heavy I/O.

To use the DGSJ feature, a group of disks is allocated to a dynamic disk group shared by multiple host servers. At any point in time, each individual disk in the group is under the control of only one server. However, when the data is needed by an alternate host, for example to perform backup processing, the mirror of a disk group is split and all volumes in the mirror are reassigned to the backup server. This process does not require any physical movement of data or disks. The ownership of disks and LUNs within the DGSJ is the only thing that changes. After off-host processing is complete, the split disks and LUNs are reassociated with the original disk group and resynchronized using FastResync.

DGSJ processing can also be used when reorganizing disk groups on one server. In this case, after the mirror is split, the reorganization processing takes places on the same server as the primary application. When the reorganization is complete, the mirror set is resynchronized.

How FlashSnap Works

The snapshot process begins by creating a volume mirror. This procedure generates an atomic copy, or mirror, which is similar to a full volume backup. Next, the mirror is detached from the original to create a new volume. This is the snapshot volume. The new volume is an exact copy of the original, as at the point in time of the snapshot. The snapshot copy is now available to be used for recovery, backup, or off-host processing.

Merging the snapshot back into the original volume is called resynchronization, or snapback. Snapback reattaches the snapshot volume to the original production volume. Storage Foundation uses FlashSnap's FastResync capability to perform snapback. FastResync ensures that only changes are copied between the original production volume and the snapshot.

At snapback, the administrator chooses whether the primary volume or the snapshot volume should become the single image. For example, if a logical error has occurred on the primary volume, then the snapshot volume represents a consistent point-in-time image that can be restored very quickly.

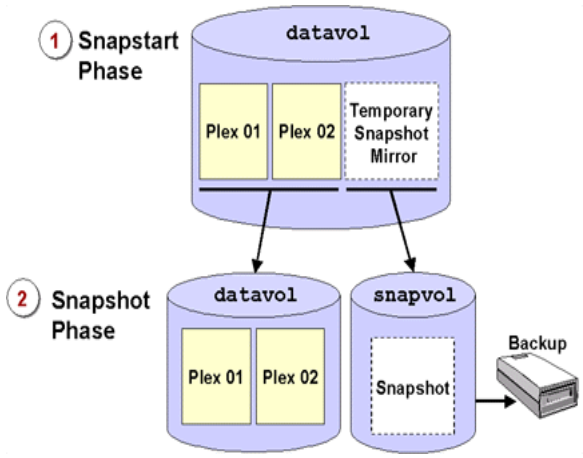

Figure 2: FlashSnap creates, detaches, and resynchronizes a snapshot mirror of an original volume.

FlashSnap VSS Provider

VERITAS Storage Foundation for Windows 4.1 enables Microsoft VSS snapshots by providing requester (VxSnap utility) and provider (VERITAS FlashSnap) components that are tightly integrated with other VERITAS solutions, such as Backup Exec, NetBackup, and Volume Replicator. The FlashSnap VSS provider creates split-mirror snapshots that are available to the Storage Foundation requester — VxSnap — or any other third party VSS requester.

Storage Foundation communicates with the VSS framework to perform the snapshot. When complete, the data is presented in its own volume and assigned a drive letter, or mount point if desired, just as if Storage Foundation were managing a non-VSS snapshot. By leveraging VSS, VERITAS Storage Foundation for Windows FlashSnap Option can create exact point-in-time copies of data to be used instantly for backup, quick recovery, or other off-host processing.

FlashSnap VSS Provider Advantages

The FlashSnap provider is software-based and hardware-independent, giving IT buyers the freedom to select storage and hardware solutions from any vendor. FlashSnap is also heterogeneous with respect to operating system, supporting Windows 2000 and Windows Server 2003, and allowing for mixed Windows environments. With Windows Server 2003, FlashSnap uses the VSS framework, taking advantage of Microsoft's guarantee of data consistency and integrity.

The split-mirror snapshots created by the FlashSnap VSS provider support data transport over a SAN, which facilitates off-host processing operations such as backup, data mining, data analysis, and testing. Backups take place at Fibre Channel speeds, because the data resides at the backup server and need not be pulled across the LAN. As a result, the process does not reduce the performance of the application server or the LAN during backup. Entire volumes can be recovered in less than 90 seconds, regardless of size. The split-mirror snapshot is also readable and writeable, which is required for off-host processing applications such as data mining and decision support.

The independence of split mirror snapshots from the original storage volumes also offers several advantages. Loss of the original volume does not affect the snapshot, and resynchronization from the snapshot volume can be performed even if the original volume is corrupted. Independent operation also means that split mirror snapshots offer faster performance than copy on write solutions, which require updates to reflect volume changes.

Summary

Fast recovery from data corruption has always been a critical issue for administrators, but heightened end-user intolerance of system downtime is raising the importance of the problem. In addition, the quest for greater application availability is squeezing the traditional backup window and forcing administrators to consider running more and more non-essential applications off-host. Taken together, these issues point to a fundamental need for technology capable of minimizing recovery time, eliminating the need for a backup window, and facilitating the efficient and effective processing of non-critical applications off the primary host. Snapshot technology provides the solution.

VERITAS FlashSnap augments the industry-leading storage management technology of VERITAS Storage Foundation for Windows to deliver a point-in-time snapshot solution for Microsoft Windows 2000 and Windows Server 2003 environments. Leveraging the data integrity guarantees of the Microsoft VSS framework, the VERITAS FlashSnap split-mirror process allows the rapid snapshot of storage volumes for use on or off the primary host server. FlashSnap allows rapid file recovery and dramatically reduces the downtime traditionally associated with sharing data between servers. Supporting heterogeneous storage resources, VERITAS FlashSnap allows administrators greater flexibility to maximize the use of available capacity, regardless of vendor.