Snapshot Technology: Improving Data Availability and Redundancy

 $\ensuremath{\textcircled{\sc c}}$ 2006 Infortrend Corporation. All rights reserved.

Table of Contents

3
3
6
7
7
7
7
9
9

Introduction

Today's RAID storage in the mid-tier to enterprise-scale environments is designed to provide 24x7, year-round data availability. Except for the planned downtime required to perform system maintenance and upgrades, interruptions to service is now considered unthinkable. However, data loss might still occur due to file system/data corruption, or simply the accidental disconnection of a cable. Rapid recovery of critical data, e.g., from a clustered server, is crucial for disaster control or even just to save time and effort for a routine backup. With the rapidly increasing flow of data inside and across every organization, an easier approach to data recovery is urgently needed in order to reduce the cost and complexity of the backup and restore operation.

Traditional backup operations, whether whole volume duplication or mirroring, usually take place when an application is shut down to avoid disrupting overall system performance. These methods are operated manually and consume a storage capacity equal to or larger than the original data source.

Furthermore, the drastically increasing size of data flow creates a demand for automatic and scheduled backup. A recurring backup/recovery schedule can be set to occur as frequently as every hour or every few minutes. Using conventional methods, a routine backup could require several times the storage capacity and management efforts, not to mention the costs of management tools.

Snapshot overcomes the deficiencies of conventional backup methods with its ease of operation, lower cost, and management efficiency. Snapshot technology is rapidly becoming an industry standard for protecting data at a fine granularity. As a leading provider of RAID storage arrays, Infortrend is designing snapshot functionality into its firmware so that it will consume a minimum share of system resources and disk capacity. This feature will be released on Infortrend systems in the fall 2006 timeframe.

Snapshot Overview

Snapshot is a "picture" of what the data looks like. This is usually referred to as a metadata copy in a reserved space, meaning the data itself is not copied to another location—only the pointers to where the data actually resides are copied. This is also called a PIT (point-in-time) copy of the source volume. Using a technology called copy-on-write (COW), if a snapshot is in place and then someone tries to over-write the data on the original volume, the snapshot software (either subsystem firmware or software running on the host) will first copy the original block of data to a new location (a pool of storage dedicated to copy operations) before it lets the write happen. It then maps pointers to the new location when an application references the original data, or back to the original location when the snapshot is referenced.



Figure 1: COW (Copy-on-Write) Snapshot Process

A snapshot is a view or logical copy of the source volume. It is available to servers on a storage area network (SAN) as a unique volume, but it always retains an association with the source volume. Using copy-on-write technology, snapshot creates logical copies of the source volume without consuming equivalent disk space. When a snapshot is initially created, no data is copied and the job is completed almost instantaneously.

If blocks are updated on the source volume, the old data blocks are moved to a location known as the snapshot reserve space. When a host sends a request to read the specific snapshot image, the system determines whether or not these requested blocks were changed on the source volume after the snapshot was created. If the blocks have been changed, the read request is satisfied from the source volume. If not, the read request is satisfied from the data stored in the snapshot.

1. Read current volume



2. Read the first snapshot volume



Figure 2: PIT (Point-in-Time) Image Mapping

If one or several snapshots exist, these snapshot images all depend on the source volume and the snapshot processing records. These snapshot copies protect stored data from accidental or malicious loss. With snapshot technology, data recovery takes place as fast as the backup process. The operation does not affect system performance or I/O access from applications running on the host. IT administrators or users can promptly retrieve lost data without considerable interruptions.



Read data of snapshot volume at T2: (With current data volume and snapshot data)



Retrieve data of snapshot volume at T1: (The original volume before snapshot)



Figure 3: Recovering Data from Several Snapshot Volumes

Snapshot Operation from Infortrend

Functional Description

Infortrend will implement snapshot as a group of firmware functionalities. Before using the functionality, the administrator should manually allocate a reserved space for storing snapshot images using any of the provided RAID configuration utilities. The snapshot reserved space functions as a storage pool for snapshot images and the associated configuration data. This flexible, easy-to-use RAID expansion features allows users to easily allocate available disk space for snapshot operation.

After allocating a reserved space, the administrator can now designate a disk volume to be snapped. Infortrend's snapshot implementation will allow all disk volumes to be snapped and these volumes use the same reserved space. This saves the trouble of designating an individual space for every disk volume.

© 2006 Infortrend Corporation Create an Advanced Backup Solution Using Snapshot Technology

Creating Snapshots

Snapshot functionality can be manually executed or set to operate automatically as scheduled tasks with configurable time intervals. Configuring snapshot as scheduled tasks facilitates data protection, makes it easy to manage, and protects online transaction data with fine granularities. Infortrend's snapshot policies will also include configurable options to prune and purge out-of-date or useless snapshots to minimize the disk space used for snapshot images.

Other Snapshot Operations

Infortrend's snapshot implementation will support operations for existing snapshots such as snapshot mapping (mounting), deletion, and reformatting. Snapshot images appear as separate "virtual" disk volumes. To validate the snapshot recovery function, the administrator needs to associate the virtual volumes with host ID/LUNs and avail the volumes in the operating system. If a user accidentally deletes a file, the mounted virtual volumes make it easier to recover the lost file. Infortrend's snapshot implementation will allow all existing snapshots to be associated with host ID/LUNs, mounted, and accessed simultaneously. In this way, performance impact is reduced and data recovery becomes flexible.

Users can also manually delete old snapshots or have them checked by preset algorithms to free the capacity used for snapshot operation. Like creating a snapshot, Infortrend will also offer configurable options so users can define the strategy for deleting or purging obsolete snapshots. Deletion can be configured to take place by time (obsolete time span), by capacity (as a percentage of actual size), or other specified requirements. The above conditions can also be prioritized. Freeing space by deleting old snapshots is completed instantly so the space is immediately available. Snapshot reformatting erases all data on a virtual volume but keeps the associated snapshot attributes intact.

System Rollback

To enable disaster recovery when the entire system files have been irretrievably corrupted or destroyed, an administrator can use the system rollback function for complete system recovery. Recovery time will vary depending on the amount of differential data between the dead volume and the snapshot volume. System rollback migrates all previously stored data blocks and removes unnecessary records and configurations. To minimize system downtime, the actual data migration takes place in the background.

Key Features

Snapshot benefits users by providing instant, disk-based, point-in-time images of any disk volume using proprietary technology. Snapshot's key features include:

Zero Impact Backup

As a giant step ahead of traditional tape backup methods, disk-based snapshot features rich functionalities to provide efficiency, flexibility, protection, and ease of management. Data access, backup, and recovery are faster and easier than tapes.

Snapshot is apt for online backup. Snapshot creates incremental and differential images with almost zero latency so there is no need to shut down system applications.

Maximum Storage Utilization at Minimum Cost

With snapshot technology, whole volume backup consumes only the resources required to calculate the incremental differential data against the previously saved volume images. This block incremental behavior minimizes the consumption of storage capacity to optimize the amount of usable capacity.

Instant Virtual Volume

When virtual volumes or snapshots are associated with host ID/LUNs, they appear as physical disk volumes to the host OS. Once configured, these live volumes are instantaneously available to multiple host systems.

Data Security

Virtual or snapshot volumes are read-only. Like a frozen disk, they remain completely static and incorruptible copies. This feature meets the requirements for absolute data integrity for purposes such as information archives or instant recovery. Users can directly access snapshot copies to retrieve or recover from accidental deletion, corruption, or malicious modification of data.

Scalability

A configurable number of snapshots can be created as point-in-time images or snapshot levels depending on the capacity of the reserved space. Infortrend's acclaimed RAID expansion capabilities will also apply to virtual or snapshot volumes. Virtual volumes can be dynamically expanded to create additional snapshots as needed.

Snapshots as Scheduled Tasks

Snapshot is made hands-free by ingenious management tools and comprehensive policies. Snapshots can be taken automatically according to the parameters previously set by a scheduling tool or manually by a system administrator. The configurable automation parameters include: the time interval between snapshots, priorities among various instances, and configurations for purging outdated snapshots.

Intelligent Rollback

Infortrend's snapshot is being designed for near-instantaneous operation and failsafe recovery. Whole volumes can be rolled back within seconds using any existing snapshot. Instantaneous operation is achieved through quick references to the target snapshot in a virtual volume and then storing the updates into yet another volume. During the rollback operation, the most up-to-date volume remains accessible.

Application and Standard Support

Infortrend's snapshot technology will be fully compatible with initiator-side methodologies, such as Microsoft Virtual Shadow Service (VSS), to allow snapshot to be fully application-aware. Although Infortrend's snapshot functionalities will operate without additional backup software, snapshot also works well with other vendors' backup applications, making it ideal for building integrated solutions with related products.

Infortrend Snapshot Benefits

- 1. Enhances data protection and retention of critical and confidential information.
- 2. Increases availability of storage resources and user data.
- 3. Eliminates cost to purchase backup software.
- 4. Instantly recovers massive volumes.
- 5. Delivers low total cost of ownership (TCO) and higher return on investment (ROI).

Conclusion

Snapshot provides an instant, economic, and convenient solution for failsafe protection and hands-free management. It looks and behaves like complete backups, and is mounted as real volumes. Read requests to snapshot volumes are fulfilled simultaneously without affecting the original volumes or the overall performance of the system. Massively corrupted volumes can be recovered within moments; and if necessary, free space can be easily released by deleting outdated snapshots. This ease and flexibility derive from the fact that these operations can be performed manually by a system administrator or automatically by using the embedded scheduling tools.

Snapshot will be an exciting addition to the features-rich functionalities of Infortrend RAID arrays. In addition to snapshot's basic concepts, such as PIT copy and COW, Infortrend is implementing enhanced snapshot scalability to manage virtual volumes as if they are RAID volumes. In keeping with our commitment to protect data, snapshot implementation will help build RAID storage subsystems that are reliable, efficient, and user-friendly.

###

www.infortrend.com