

Technical Note
SNAZ™ OS Snapshot Technology

Raidtec SNAZ™ OS Snapshot Technology

©2003 Raidtec Corporation



TECHNICAL NOTE

SNAZ™ OS SNAPSHOT TECHNOLOGY

High Availability

Today business critical data is growing exponentially and companies are finding that an increasing proportion of their data is being stored digitally. High availability is now a business essential, given the high costs associated with network and system downtime. Data recovery plans must be in place to ensure the rapid recovery of data due to corruption, disaster or human error.

Raidtec SNAZ™ OS Snapshot Technology

SNAZ™ OS Snapshot Technology allows users to take instant, logical volume-based, point-in-time images of any logical volume on a SNAZ system.

Key features of SNAZ Snapshot Technology:

- **Instant virtual volumes** - the Snapshot Volume is an instant virtual copy of a live Logical Volume.
- **Point-in-time imaging** - the Snapshot Volume is a virtual image of the live volume as it appeared at the point in time when the snapshot was captured. It creates a nearly-instant map of pointers to the actual blocks of data found in the live volume.
- **Logical volume based** - the Snapshot Volume appears on the system as a separate Logical Volume and can be shared or exported to Windows or Unix clients. The blocks of data that remain unchanged following the completion of the Snapshot stay in the live volume while the original contents of data blocks that have changed are stored in a designated snapshot buffer.
- **Requires only a fraction of the capacity used by the live volume** – real disk capacity is consumed by the snapshot buffer only when data on the original live volume is changed. Typically only a small fraction of the original data is changed each day. The initial size of the snapshot buffer can be specified when the snapshot is created.
- **Snapshot buffer inflates on demand** - when a snapshot buffer is 70% full, it is automatically increased by 20% to prevent loss of snapshot data. System administrator is notified by email of these events.
- **Continuous client access to the live volume** - with the exception of a short pause as the live volume is frozen and captured in the snapshot image, network clients maintain read/write access.
- **Read-only access to the snapshot volume** - users have read-only access to individual snapshot volumes. Since these snapshot images are true point-in-time copies of the live volumes, users cannot manually modify individual snapshot images, protecting the snapshot from inadvertent user error.
- **Multiple concurrent snapshots** – Up to 25 concurrent snapshots per Logical Volume.

How Snapshot Works

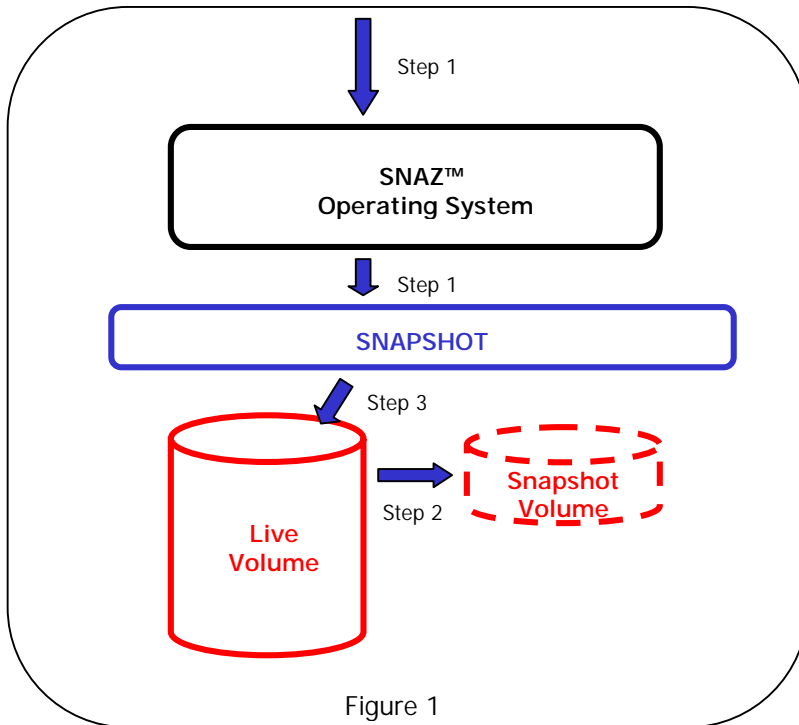


Figure 1

When Writing New Data

Step 1 - A new Write request is initiated by client.

Step 2 - The original data content in the Live Volume, which is about to be overwritten, is moved to the Snapshot Volume.

Step 3 - The Write request is then saved to the live volume.

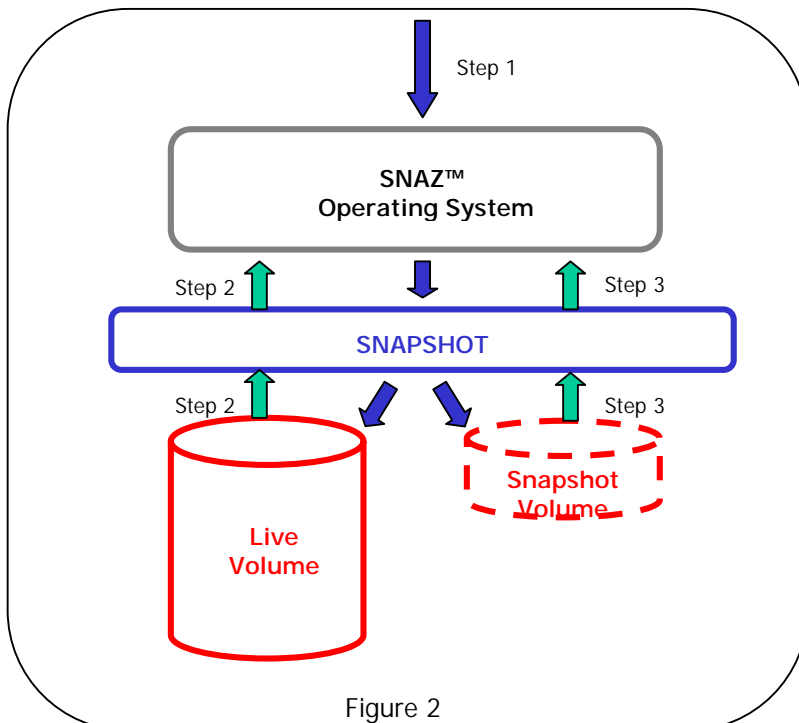


Figure 2

When Reading Data

Step 1 - Request to read data initiated by client from a snapshot image.

Step 2 - Unchanged data blocks are read from the live volume.

Step 3 - Where data blocks have been changed in the Live Volume (as in Figure 1 above), the original data contents are retrieved from the Snapshot volume.

Point-in-Time Image Creation

SNAZ™ OS Snapshot Technology uses copy-on-write execution to create point-in-time images of a Logical Volume. When a snapshot is initiated, the following steps occur:

1. The 'live' Logical Volume (LV) is momentarily frozen.
2. Almost instantaneously, a virtual copy of the LV is created. This is just a database of pointers to the actual blocks of data (which presently exist in the LV). Buffer space for the Snapshot Volume is also reserved in the same Volume Group (Storage Pool) in which the LV resides. This space will be used to copy blocks of frozen data before they are over-written on the LV. This buffer space can be specified when the snapshot is created. By default it is set to 15% of the LV size. When the buffer is more than 70% full, additional capacity equal to 15% of the LV size is reserved. If there is no more capacity available in the Volume Group, the Snapshot becomes invalid.
3. When write requests are made by the client, the snapshot technology intercepts these requests, reads the data in the blocks to be overwritten, and saves the original data in the Snapshot Volume (Figure 1).
4. Write requests initiated by the client are completed after a copy of the original data has been made. (Figure 1). Please note that for the sake of simplicity, Figure 1 shows only one Snapshot Volume, implying that only one Snapshot has been created. If multiple Snapshots are created and maintained at the same time, a separate Snapshot Volume will exist for each individual Snapshot image.
5. Steps 3 & 4 are triggered with further write requests to the LV, as a copy of the original data is made when new write requests are initiated. Please note that if the volume receives further requests to overwrite blocks of data which have already been overwritten and the original contents have already been captured in a Snapshot Volume, no further copy-on-writes will be initiated. Instead, only the very first overwrite of a set of data blocks will trigger a copy-on-write.
6. Upon receipt of client requests to read the contents of individual snapshots, SNAZ Snapshot Technology simply substitutes the original data blocks for the changed data blocks (Figure 2).

Powerful, Expandable NAS Solutions

Raidtec NAS Solutions with SNAZ™ OS Snapshot Technology can be employed in a number of different applications and environments and enables the user to carry out the following tasks:

- **Online backup** - each snapshot is a virtual copy of the live data volume at a given moment in time. SNAZ™ OS Snapshot Technology allows users to store up to 255 concurrent snapshot images as online logical volume-based back-ups. Additionally, users can share each of these snapshot images, and they can also be subsequently backed-up to tape for archive.
- **Simplified backup process** - typically, the network administrator performs regular backups of critical server data to tape for both backup and archive purposes. Usually, this can include any combination of daily backups and weekly/monthly complete backups. To reduce business impact, daily backup operations are usually performed with the servers taken offline after business hours, known as the 'backup window'. With the advent of 24x7 work environments in the global economy and the prohibitive cost of downtime, this 'backup window' no longer exists in many situations. The protection of mission-critical data in the 24x7 environment is paramount to economic survival. SNAZ™ OS Snapshot Technology eliminates the need for a 'backup window' by allowing a network administrator to perform daily backups and weekly/monthly complete backups to tape more speedily and efficiently. Snapshot images of the volume can be backed up throughout the course of the day as needed. Users will only experience a momentary pause in data access at the exact moment when the volume is frozen and the snapshot image is created.
- **Instant data recovery** – The network administrator can configure each snapshot image to be shared over the network, enabling individual users to access the snapshot images, from which they can copy any inadvertently modified, deleted and/or corrupted user data and files back into the live volume. For security, each user has read-only access privileges, and only to his/her own data in the snapshot.
- **Offline volume management** - the snapshot volume holds all completed and active file access data. Third-party reporting of storage activity by volume, file type, and application may be run on individual snapshot

images. The administrator is thus allowed to monitor and assess storage usage and plan for growth, while removing the access contention and performance overhead associated with running reports on a publicly accessed live volume. The facility to maintain a collection of snapshot images, corresponding to multiple points in time, makes SNAZ™ OS Snapshot Technology ideal for revision management and testing. Multiple copies of software code or documentation may be easily tracked, maintained, accessed, and audited for testing purposes.

Low Cost, High Availability

The benefits of SNAZ™ OS Snapshot Technology to administrators and users are significant and include:

- High availability of storage resources and user data. Multiple snapshots can be made available on-line as separate volumes to provide multiple consistent copies of data at various points in time.
- Better protection of mission-critical data. The use of snapshots as both logical volume-based virtual volumes and regular incremental backups of daily snapshots to tape provides multiple levels of protection for mission-critical company data.
- Immediate recovery of user data from disk. Snapshot functionality provides much speedier recovery of data than traditional data retrieval from tape. Users and administrators can retrieve specific data from logical volume-based snapshots almost instantly.
- Reduced total cost of ownership (TCO) and higher return on investment (ROI). The increased uptime and availability of mission-critical company assets and user access to data gives rise to higher productivity.

Ease of Use

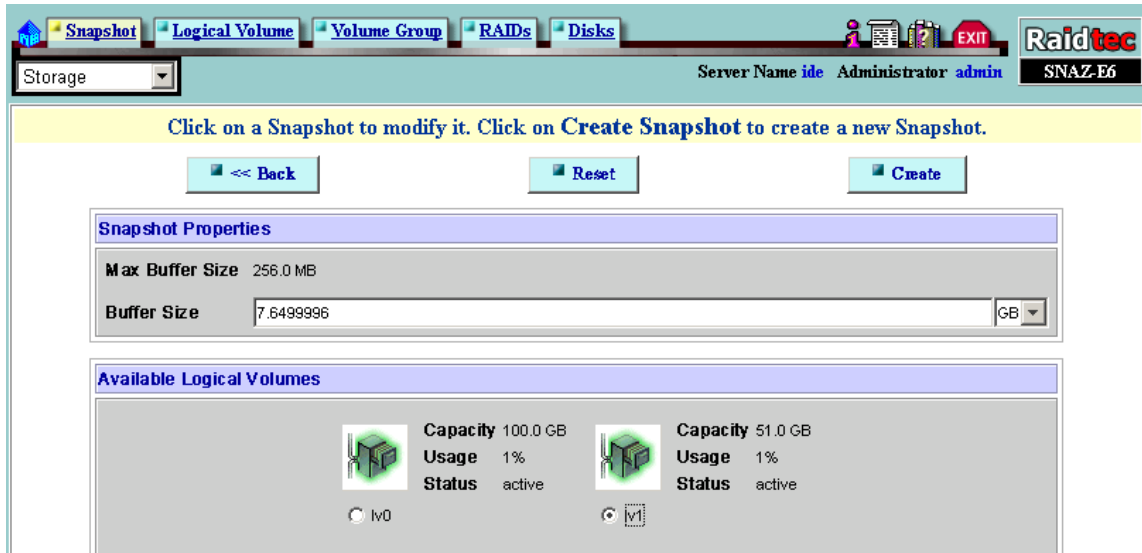
Snapshots can be created within a few clicks.

The network administrator is given the following options:

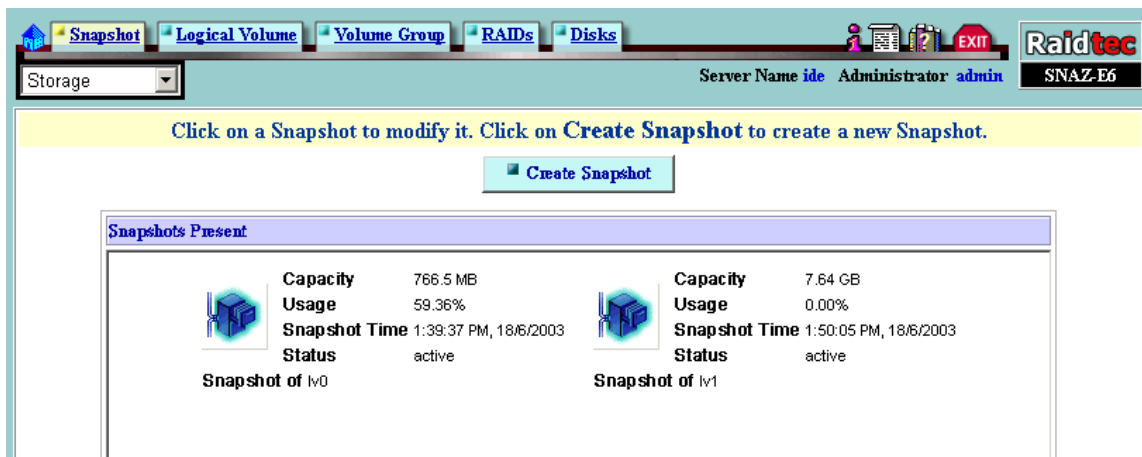
- Click on “Create Snapshot” button from the “Snapshot” screen in “Storage” tab.



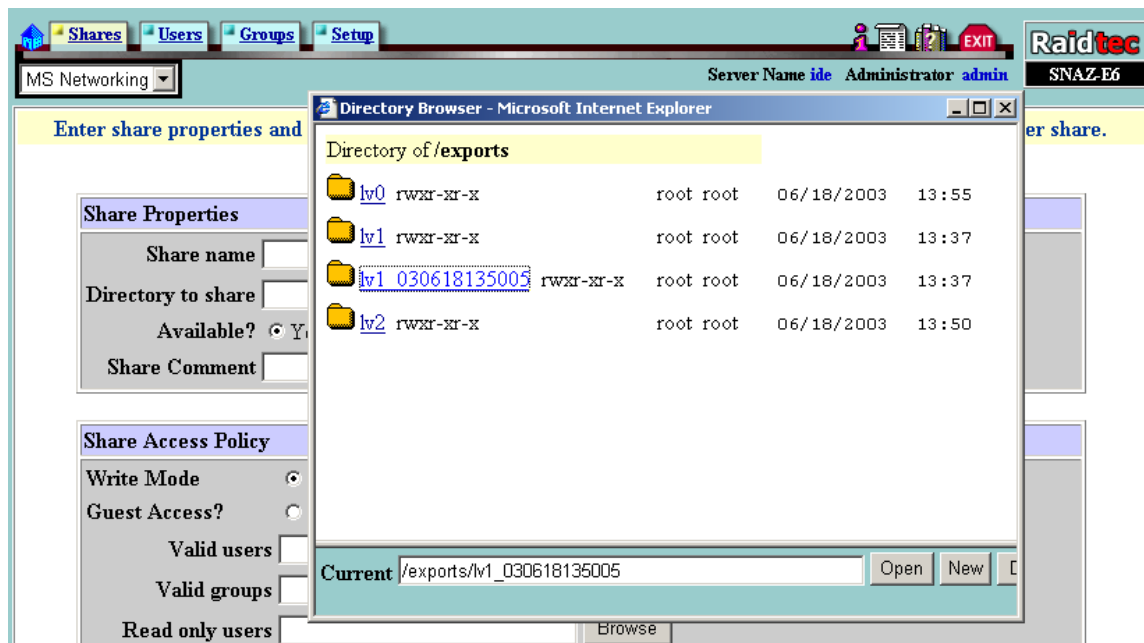
- Select a live volume to Snapshot. The administrator can specify any live logical volume in the SNAZ system.



- Enter the Snapshot buffer size. The administrator can specify any size up to the size of the live volume. The default buffer size is 15% of the live volume capacity.
- Click on “Create” button.



Exporting Snapshot as Windows Share or NFS Export



The administrator can export a Snapshot as a Share on a Windows network using the “Shares” screen in the MS Networking tab. A Snapshot name is created automatically using the live volume name plus the date and time the Snapshot is created. Once a Snapshot is shared on the Windows network file system, users can access the Snapshot in read-only mode. The Snapshot can also be used to backup data to any backup server on the network.

Simplified Daily Backup Process

Daily backup operations can be carried out by the network administrator any time, given that the snapshot process takes less than a second to complete. Once setup for daily repetition, the backup task will automatically backup individual snapshots: snap 1 on Sunday, snap 2 on Monday, etc. As a consequence, daily backups can be completed over 24 hours, thus making the ‘backup window’ a thing of the past.

Reduced TCO and High ROI

SNAZ™ OS Snapshot Technology provides easy-to-use features to facilitate the capture of point-in-time images of a live volume, thus allowing the network administrator to simplify the backup process and give all users the instant recoverability of lost or corrupted data. This, coupled with the hardware redundancy, enhanced system security, high performance, fault tolerance, and the ease of management, contributes to overall system and data reliability and availability, which leads to reduced Total Cost of Ownership (TCO) and high Return on Investment (ROI) for the user.

Technical Note

SNAZ™ OS Snapshot Technology



Raidtec Corporation (US)
1324 Union Hill Rd Bldg A,
Alpharetta, GA 30004 USA
T: 770-664-6066
F: 770-664-6166

Raidtec Corporation (Europe)
Castle Road, Little Island,
Cork, Ireland
T: +353 21 4353440
F: +353 21 4353799

For more information:
www.raidtec.com
info@raidtec.com

© 2003 Raidtec Corporation. All rights reserved.

No part of this publication may be reproduced in any material form (including photocopying, or storing it in any medium by electronic means whether or not transiently or incidentally to some other use of this publication) without the prior written permission of the copyright owner, or under the terms of a license issued by the copyright owner.

This document is for informational purposes only and does not set forth any warranty, expressed or implied, concerning any product, product feature, or service offered or to be offered by Raidtec. Raidtec reserves the right to make changes to this document at any time, without notice, and assumes no responsibility for its use.

Raidtec, *FlashLinux* and *SNAZ* are trademarks of Raidtec Corporation. All other marks mentioned in this material are the property of their respective owners.