



RELDATA SAN-level Snapshots

Getting a Business Return from Snapshots

A Technical Whitepaper
December 2006

Abstract

The RELDATA unified IP storage appliance incorporates SAN-level Virtualization technology that encompasses high performance, reliable, block-based snapshot functionality to help solve data protection issues and provide point-in-time volumes for business needs.

Businesses haven't leveraged the advantages of snapshots to their fullest because snapshots have traditionally been complex to implement and manage; each hardware vendor implements snapshots differently leading to inconsistent functionality and capabilities; they are difficult to integrate with business-critical applications or databases; and finally, they are labor intensive requiring lots of administrative personnel to manage, monitor and oversee snapshots.

RELDATA SAN-level Snapshots technology, part of its embedded SAN-level Virtualization software platform, delivers a simple way to create, present and manage snapshots of heterogeneous operating systems and applications within a single management console. By easing the administrative overhead normally associated with snapshots, RELDATA allows snapshots to become part of a comprehensive data protection and data management strategy.

RELDATA SAN-level Snapshots — Getting a Return from Snapshots

Content

The Challenges Facing IT with Snapshots	3
How RELDATA Snapshots Work	4
RELDATA Origin Volume	4
How Block Changes to a File are committed to Origin Volume	4
How Block Changes to a File are Committed to Copy on Write Area before Origin Volume is Updated.....	4
RELDATA Origin and Snapshot Volumes.....	5
How Older Data is Properly Read by Combining Origin and COW Volume Blocks versus Normal Disk Read.....	5
Multiple Snapshots with RELDATA SAN-level Snapshots.....	6
Business Use Case for Multiple Snapshots	6
RELDATA Multiple Snapshot Technology	6
RELDATA Writeable Snapshots.....	7
Business Needs for Writeable Snapshots.....	7
Use Cases for Writeable Snapshots	8
Rolling Back and Rolling Forward with Writeable Snapshots	8
Snapshot Renewal.....	8
Snapping a Snapshot.....	9
Summary	9

December 2006

Copyright ©2006 RELDATA Inc., all rights reserved. The copyright and all intellectual property rights in this document belong to RELDATA. It is strictly forbidden to copy, duplicate or otherwise use this document or any part thereof in any way shape or form without the prior written consent of RELDATA. RELDATA reserves the right to make changes to this document at any time, without notice, and assumes no responsibility for its use. This informational document describes features that may not be currently available. RELDATA and the RELDATA Logo are trademarks of RELDATA Inc.. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such.

U.S.A.:

RELDATA Inc.
www.reldata.com
Phone +1 (973) 644 2770
Fax +1 (973) 644 3385
Email sales@reldata.com

Europe:

RELDATA Europe GmbH
Waldburgstrasse 15
D-70563 Stuttgart, Germany
Phone +49 (711) 2272020
Fax +49 (711) 22720210
Email sales.emea@reldata.com

The Challenges Facing IT with Snapshots

IT organizations are facing increasing pressure from the business units to fully protect vital data and meet shrinking backup windows. In addition, regulatory and compliance needs combined with unrelenting growth of data stores demand that IT maintain accurate and time-sensitive copies of production data for protection, reporting, business intelligence, testing and application development uses. A properly implemented snapshot strategy improves Recovery Time Objectives (RTO) and Recovery Point Objectives (RPO). While all the major vendors promote their snapshot technologies, few IT organizations can fully exploit them due to their complexity and inconsistent capabilities. As a result, snapshot technology is underutilized by even sophisticated IT departments.

The reason is simple: snapshots can be file-based or block-based. File-based snapshots require significantly more integration with applications and databases to create a consistent snapshot image. In most cases, IT organizations are required to script the integration with applications or databases necessitating highly specific knowledge and heavy maintenance of the infrastructure to deal with inevitable changes of dynamic IT environments. RELDATA's block-based snapshots, an integral feature of its SAN-level Virtualization software platform, offers performance, consistency and simplicity advantages over competing solutions.

There are a number of reasons why businesses don't exploit snapshots:

Difficult to manage: snapshots require complex configuration parameters to be set and maintained in evolving IT environments. Many IT organizations lack the expertise to effectively management snapshots. RELDATA's web-based administrative interface makes it easy to create, track, destroy and oversee snapshots.

Inconsistent functionality: competitive vendors promote their snapshot capabilities for their FC SAN and NAS hardware, but associated snapshot capabilities are not consistent in terms of core functions, capabilities, even certifications. RELDATA's SAN-level Virtualization technology ensures users have the tools they need to meet their point-in-time copy requirements.

Heterogeneous support: the lack of cross platform support for Unix, Linux, and Windows including applications and databases restricts IT group's ability to use snapshots everywhere they might desire. RELDATA's SAN-level Virtualization technology ensures that users can deploy snapshot technology in heterogeneous environments, comprised of distributed operating system and applications, when and where business needs dictate.

RELDATA SAN-level Snapshots — Getting a Return from Snapshots

Labor intensive: lack of simple, easy administration of snapshots including set-up, scheduling and administration means IT must use complex, time consuming scripting process to manage snapshots. RELDATA eases the administrative burden by providing a web-based user interface to control and handle snapshots operations.

How RELDATA Snapshots Work

RELDATA Origin Volume

Let's start with the production or origin volume. It contains data that has been written to it, as files, documents or other information is saved. As new data is saved or written to the origin volume, it changes to reflect those changes.

How Block Changes to a File are committed to Origin Volume

When the new data is about to be written out to the origin storage volume, the content of the physical area about to be written over is saved to a separate, purposely allocated area, and a corresponding reference record is created. The RELDATA SAN-level Snapshots process is called copy-on-write (COW), and the save area is referred as the "COW volume". The following figure demonstrates how data from the origin volume is written to the COW volume before the origin volume is updated.

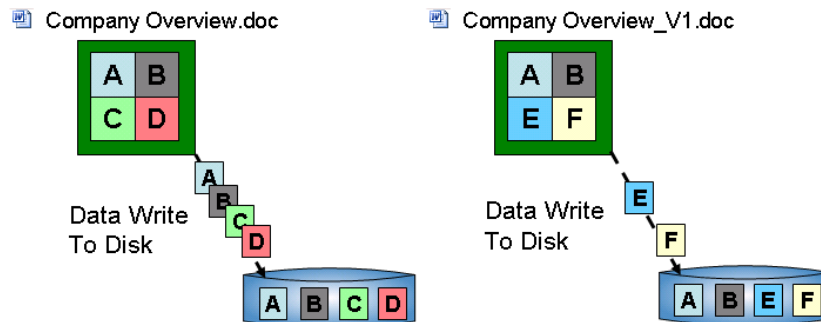


Figure 1: Origin Volume Changes As Data Writes Occur

How Block Changes to a File are Committed to Copy on Write Area before Origin Volume is Updated

Subsequent writes to the origin volume will not cause copy-on-write operations if the area being written has already been saved in the COW volume.

RELDATA SAN-level Snapshots — Getting a Return from Snapshots

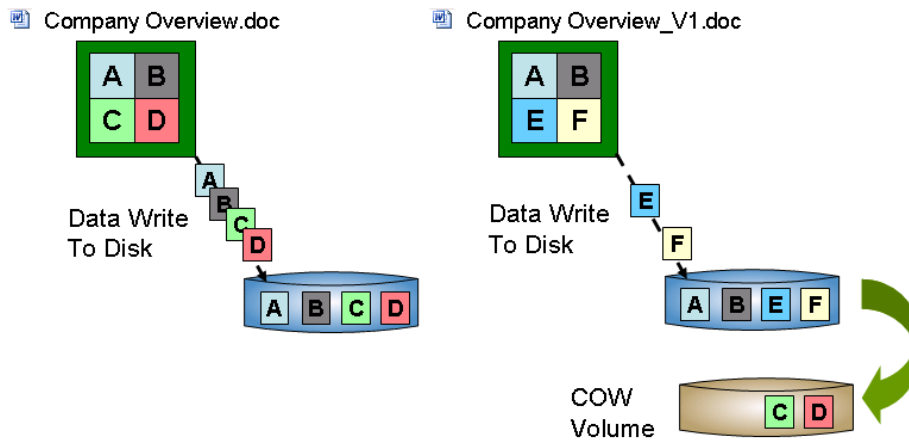


Figure 2: Copy on Write Volume Changes As Data Writes Occur

RELDATA Origin and Snapshot Volumes

For RELDATA unified IP storage systems, when the origin volume is read, the latest content from the real volume is returned – data as expected. When the snapshot volume is read, then the reference table is searched first. If the matching record exists, then the saved content from the COW volume is returned; if there is no record, then the content from the real volume is returned. Therefore, the origin volume will always return the latest data stored, and the snapshot volume will return the data frozen at the time the two volumes split.

How Older Data is Properly Read by Combining Origin and COW Volume Blocks versus Normal Disk Read

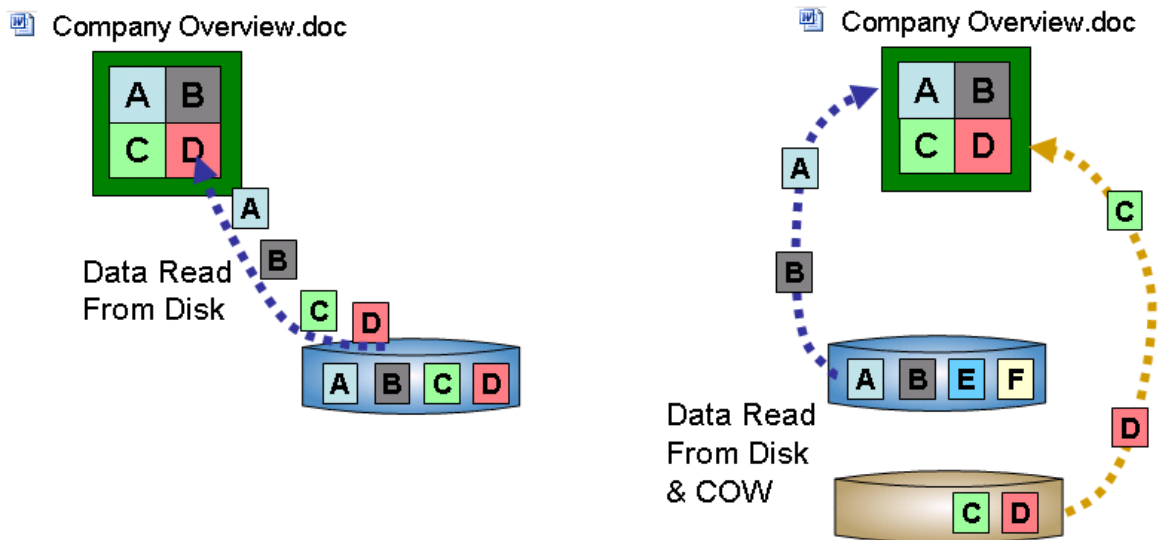


Figure 3: Data Read Combining Blocks from Origin and COW Volumes

Multiple Snapshots with RELDATA SAN-level Snapshots

In a modern IT environment the storage content changes quite rapidly. Users add, update, replace and remove various folders, files, documents and records. Data protection from an inadvertent change becomes an important concern.

A snapshot can capture and preserve the content of a storage volume at a certain checkpoint so that a “frozen” copy of the content would be available in parallel with the volume itself. The directories and files can be compared, the changes could be easily tracked, and a lost or corrupted file could be quickly restored.

Business Use Case for Multiple Snapshots

RELDATA snapshots have many business uses:

- Readily available backup of the volume as of the beginning or the end of the day,
- Creation of read only volume for testing, development or business intelligence purposes,
- Identify, track and reverse intra-day changes in case of inadvertent deletion of important data is discovered at a later time,
- Maintain several concurrent snapshots taken at different points in time, for example, each morning for the last five days, at the beginning of each week, at the beginning of each month, and so on for data protection purposes,
- Historical progression of the data reference points that makes it possible to turn back the clock and access the data as of yesterday, the day before, the week before, etc.

RELDATA Multiple Snapshot Technology

Multiple snapshots are facilitated by the same RELDATA SAN-level Virtualization copy-on-write technology. Each snapshot requires a separate reference table directing the snapshot reads to the corresponding areas of the COW volume. With each new snapshot, a new empty table is started, and is populated as the new write operations come across. The COW volume continues to fill up sequentially.

RELDATA SAN-level Snapshots — Getting a Return from Snapshots

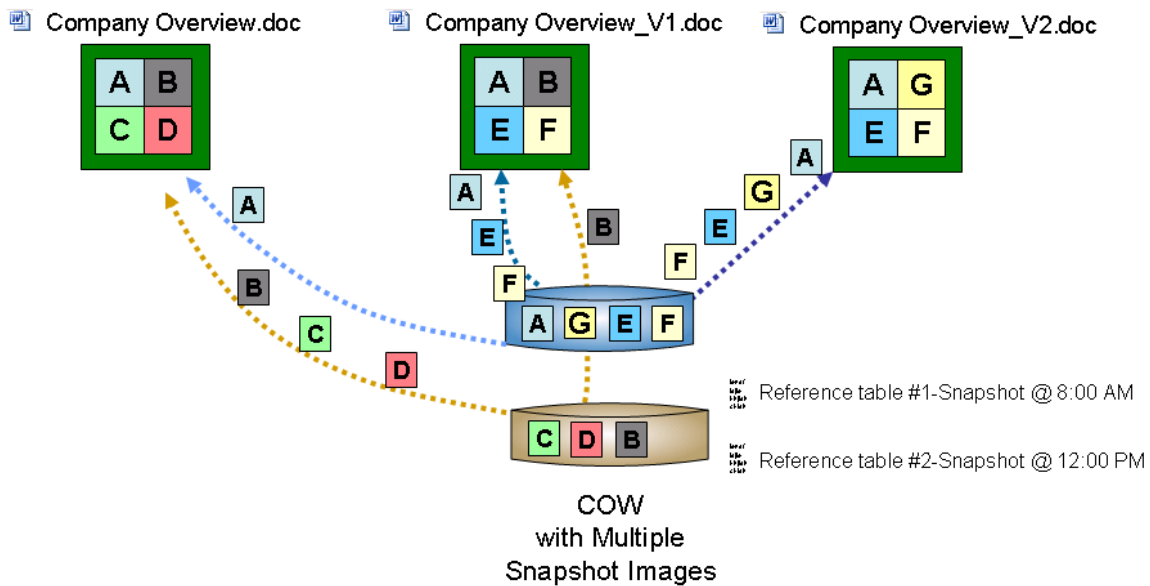


Figure 3: Multiple Reference Tables per COW Volume Retains Multiple Simultaneous Snapshots

RELDATA Writeable Snapshots

Business Needs for Writeable Snapshots

Many users wish they could alter the content of a snapshot. It's a common request, but a difficult thing to deliver. Typical business needs range from:

- Data warehousing
- Business intelligence
- Experimental software installation
- Testing
- Development of prototype applications

To date, storage array vendors have delivered writeable snapshots but at a considerable cost in terms of storage space. Whether the writeable snapshot volumes were called 'clones' or 'business continuance volumes', the result has been that users were required to pre-allocate two to three times the amount of space compared to the amount of production data.

RELDATA SAN-level Virtualization technology has pioneered writeable snapshots using space-efficient copy-on-write technology usually reserved for read-only snapshots. This provides business with a very cost-effective solution to a variety of important business needs.

RELDATA SAN-level Snapshots — Getting a Return from Snapshots

Use Cases for Writeable Snapshots

The use cases for writeable snapshots are only limited by the end user's imagination because it is simple and easy to leverage the power with RELDATA unified IP storage systems. Let's investigate some of the more common scenarios where RELDATA's writeable snapshot capability can save time and increase the manageability of data.

Here are just a few uses for RELDATA writeable snapshots:

- Mounting a snapshot image of an NTFS volume onto the same host as the origin volume,
- Rapidly undo failed software operating system or application upgrades,
- Protection during load and transform operations for a relational database,
- Migration of applications,
- Testing of new applications or development environments,
- Validation of disaster recovery scenarios.

Because of the flexible, high performance 'roll back' and 'roll forward' capabilities with RELDATA unified IP storage systems, users get many of the benefits of server virtualization without the investment in new hardware and software.

Rolling Back and Rolling Forward with Writeable Snapshots

Assume you are planning to apply a service pack (an upgrade) for an application. A smart thing to do would be to create a full backup that can be used as a last resort if the upgrade fails. Given the capacity of modern volumes, the backup can take many hours, often making this approach impractical, especially considering that the restore process will take at least as much time.

Instead, you can simply create a read-only snapshot of the volume attached to a RELDATA unified IP storage appliance. If things go wrong, copying the entire snapshot back to the source origin volume will get the data back. But there's a much faster method. Considering that the snapshot and the origin volumes differ only in the areas affected by changes to the origin (e.g. the service pack installation), it would be far more efficient to simply restore these areas from COW.

Since the service pack changed only a small percentage of the data on the volume, recovery will be significantly quicker because only a fraction of the total data needs to be moved. Once all everything is copied back, the origin volume will be "rolled back" to the snapshot, therefore undoing all changes made during the installation.

Snapshot Renewal

In a typical use case, there are five snapshots, one for each day of the week. On the subsequent week, new snapshots need to be created and the old ones

RELDATA SAN-level Snapshots — Getting a Return from Snapshots

deleted on their respective days. One can think about this operation as a “renewal” of the snapshots.

Regardless of the order of operations (remove first then create or create first then remove), the net result is the same – you have a renewed snapshot. Nevertheless there are a few important differences.

If the snapshot is deleted before the next one is taken then there would be a short period of time when there are no snapshots. It may or may not present a problem, and should be decided based on the data protection policy. The benefit of removing the snapshot first is that if it happens to be the last snapshot, then the entire COW volume is recycled during this operation.

Snapping a Snapshot

Scheduled renewals can interfere with any potential usage of the snapshot volumes. For example, if the volume is mounted, it can't be deleted, and the renewal will fail. At the same time, the renewal is desirable to capture the state of the origin volume at a specific point in time.

Another structure, called “snapshot clone”, was created to address this scenario. The snapshot clones are independent instances of a snapshot that can be created from any existing snapshot. The clones can be made writable and can be used for as long as necessary, even after the parent snapshot has been deleted or renewed.

In another scenario, the snapshot clones can be used to produce multiple copies of a snapshot (as opposed to origin volume). This would be helpful for purposes of loading multiple hosts (or server grid) with the same software without creating and maintaining full copies for each host.

Summary

RELDATA unified IP storage systems offer IT organizations powerful storage expansion and consolidation solutions to meet the challenge of data growth. Based on industry standard hardware, the RELDATA IP storage appliance delivers a cost-effective solution that encompasses a full range of snapshot capabilities not found in traditional ‘branded’ midrange or high end block storage systems.

The simplicity, flexibility, and manageability of RELDATA snapshot technology breaks new ground for IT to deliver fast, easy strategies for data protection, testing, development and deployment of operating systems, applications and databases.