



“ For over a decade Coolspirit have been supplying the UK’s top organisations with storage products and solutions so be assured we will meet your requirements head on.

It’s all about getting things right first time, quickly and simply! ”

Damon Robertson
Coolspirit Ltd

Our address

24 The Bridge Business Centre
Beresford Way
Chesterfield
S41 9FG

Get in touch

Call us on: 01246 454222
Email us: web@coolspirit.co.uk
Find us: [View location map](#)
Web: www.coolspirit.co.uk

Office hours

mon - thurs 8:30am - 5:30pm
fri 8:30am - 5pm
sat - sun Closed

“ Boost your storage buying power...
use ours! ”

Buy with confidence from
Coolspirit your authorised
Dot Hill Partner





SNAPSHOTS & DATA PROTECTION

THE R/EVOLUTION DIFFERENCE

white paper

Snapshots & Data Protection



SNAPSHOTS & DATA PROTECTION

THE R/EVOLUTION DIFFERENCE

TABLE OF CONTENTS

1 Executive Summary	3
1.1 Market Drivers	3
1.2 What is Snapshot?	4
1.3 What are the Key Benefits of Snapshots?	4
1.4 What are the IT Benefits of Snapshots?	4
2 The Snapshot Solution	5
2.1 Snapshots	5
2.2 Snapshots using Single Copy-On-Write	6
2.3 Use Scenarios	7
2.4 Implementing Snapshots	7
2.5 Master Volume Writes	7
2.6 Master Volume Reads	8
2.7 Snapshot Features and Functionality	8
2.8 Master Volumes	8
2.9 Snap Pools	8
2.10 Expansion Capabilities	9
2.11 Accessibility	9
2.12 Failover Support	9
2.13 Snapshot Rollback	10
2.14 Snapshot Reset	10
2.15 Policies and Thresholds	10
3 Summary	10

white paper

Snapshots & Data Protection

Snapshots & Data Protection

1 Executive Summary

It is very obvious that information and data protection is critical to business today. Hundreds of billions of dollars a year are spent on information technology hardware, software and services. And at its core, most of the spending surrounds an organization's desire to access, distribute and protect its informational resources. This fundamental objective serves to support, maintain, improve or grow business, government, education and other institutions around the world. In this paper we are going to discuss data snapshot technology - one of the key tools for accessing, distributing and protecting your data resources, the lifeblood at the heart of the information age.

Snapshots are vital tools in helping organizations increase their competitiveness. Snapshots improve IT operations for customers, suppliers and internal staff. This technology reduces both the time and

Software running in a disk array or storage network device accounted for revenue of \$1.7 billion in 2005, and array-based backup via snapshots technology is expected to exceed tape-based backup by the end of 2006

costs involved in the creation, development and release of new applications. It also supports business continuance and meets government regulations. Snapshot technology is fast, efficient and provides virtually no impact to ongoing business operations.

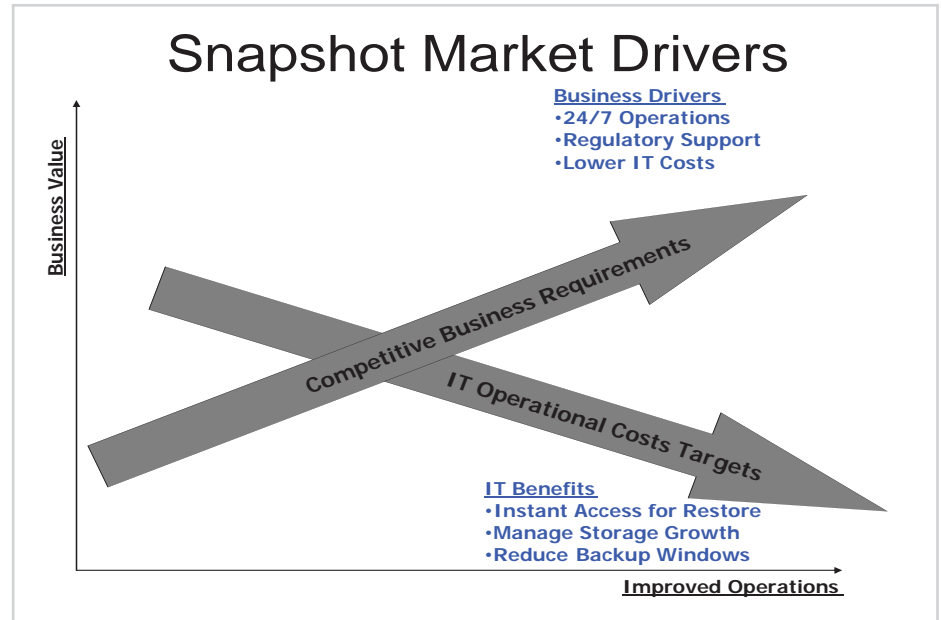


Figure 1. Snapshot Market Drivers

According to Gartner, software running in a disk array or storage network device accounted for revenue of \$1.7 billion in 2005, and array-based backup via snapshots technology is expected to exceed tape-based backup by the end of 2006. The backup and replication markets are continuing to merge into a single recovery-centric model. It is accelerating the growth of disk as the primary recovery media, since it offers significantly shorter recovery time with point-in-time snapshots. Increasingly, lines of distinction between "backup" and "replication" are blurring due the use of snapshots for data protection capabilities, including volume copies, and local or remote file or block-level replication. Ultimately, backup and replication are the means to address recovery requirements; and for disk-based data protection, snapshot is the enabling technology.

There are many market drivers and business benefits that will accelerate the use of array-based data protection. We will now cover these drivers and the technology benefits of a snapshot solution.

1.1 Market Drivers

Though the Internet bubble had burst, it left the legacy of a true 24/7 global business world. It created a fundamental change in how IT shops create and support the strategic goals of their organizations. One of the key changes was the need to keep customer and business services always available. To remain competitive or to lead markets, IT shops needed to find new tools to meet these real-time requirements, enabling them to rapidly bring new applications and services online, meet service level agreements, support business continuance, meet new business regulations and protect the data on which all of these business requirements depend.

Snapshots & Data Protection

1.2 What is Snapshot?

Snapshot technology is based on a simple concept. As the name implies, it is the capturing of an image or a copy of the data. This simple operation is very powerful because the snapshot copies can be used as a backup of the data or to distribute the data for other uses. Snapshot technology is the core tool for building a complete family of data protection services. The fundamental ability to rapidly make copies of data without impacting business operations is vital to managing the many uses and requirements placed on data in today's IT shops.

1.3 What are the Key Benefits of Snapshots?

Snapshot technology provides the ability to accomplish three key data management and protection capabilities. First, snapshots create instant point-in-time (PIT) file copies that can be used for backups to drastically reduce backup windows. Second, snapshots provide the ability to access and recover data instantly versus tape-based backup that must seek through tapes to locate and restore files. Lastly, snapshots create copies of data that can be used to support key IT functions like remote replication, business continuance, rapid application development and support for regulatory requirements.

1.3.1 Point-in-Time Backup - The Big Story on Snapshots

Perhaps the most significant function of snapshots is the ability to capture a moment in time of a disk volume. It allows the copy of data not only to function as a backup, but it provides the ability to move back to a specific point-in-time to recover systems to a moment before data corruption occurred, an application failed or other significant events. This is

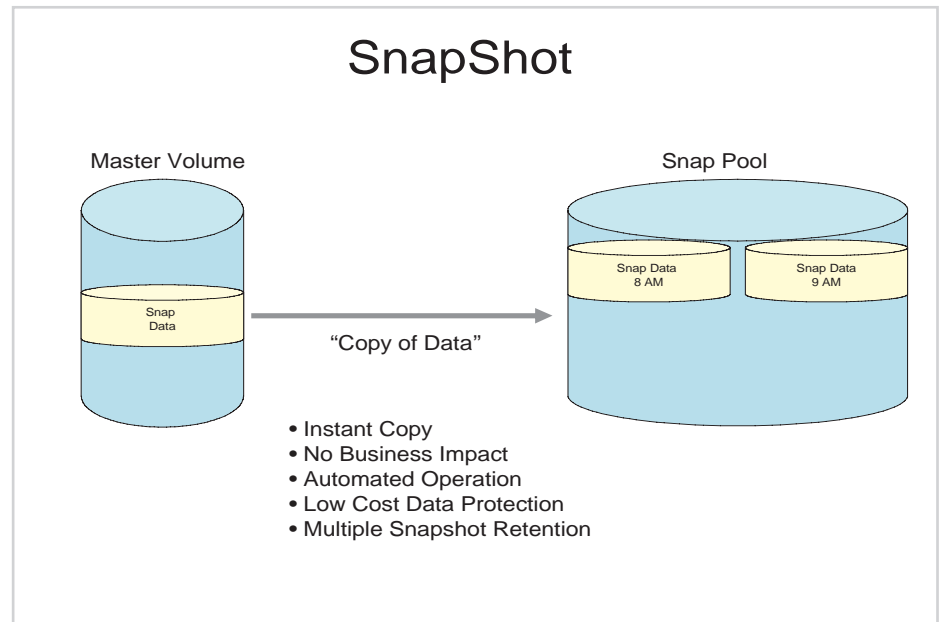


Figure 2. Snapshot Operation.

why the speed and lack of impact to operations of snapshots is so significant.

Snapshots can be run frequently and can be used to restore your business operations instantly by rolling back to the right point-in-time.

1.4 What are the IT Benefits of Snapshots?

Organizations worldwide demand a fundamental shift in how data is accessed, distributed and protected. The answer is to use more disk-based storage solutions in multiple configurations in order to protect, replicate and distribute data versus tape. The primary technology used to accomplish these goals is snapshots. A simple disk-to-disk copy supports five fundamental needs in the IT world:

- Drastically reduces backup windows by creating fast and frequent backup copies.
- Provides instant recovery for files, folders or complete volumes

- Supports data access and protection for regulatory requirements

- Provides local or remote real-time fail-over volumes for business continuance.
- Creates copies of production data to develop and test new applications

These snapshot capabilities help IT shops support their business' competitive goals, reduce costs, improve operations by reducing or eliminating the cost of tape backup, and meet the access and retention requirements of new regulations.

1.4.1 Drastically Reduced Backup Windows

A snapshot can be created in seconds and can be done without taking applications offline. The snapshot is made while the business continues to operate. This ensures that data is always protected and always recoverable; in other words, always available. This reduces the overhead of backing up data on production sys-

Snapshots & Data Protection

-tems, eliminates the risk of tape failure during backup and restore, and provides higher data availability that is verified immediately. If tape backup is required for archive or regulatory purposes, it can be done offline from the snapshot without impacting production systems.

1.4.2 Instant Recovery of Files, Folder and Volumes

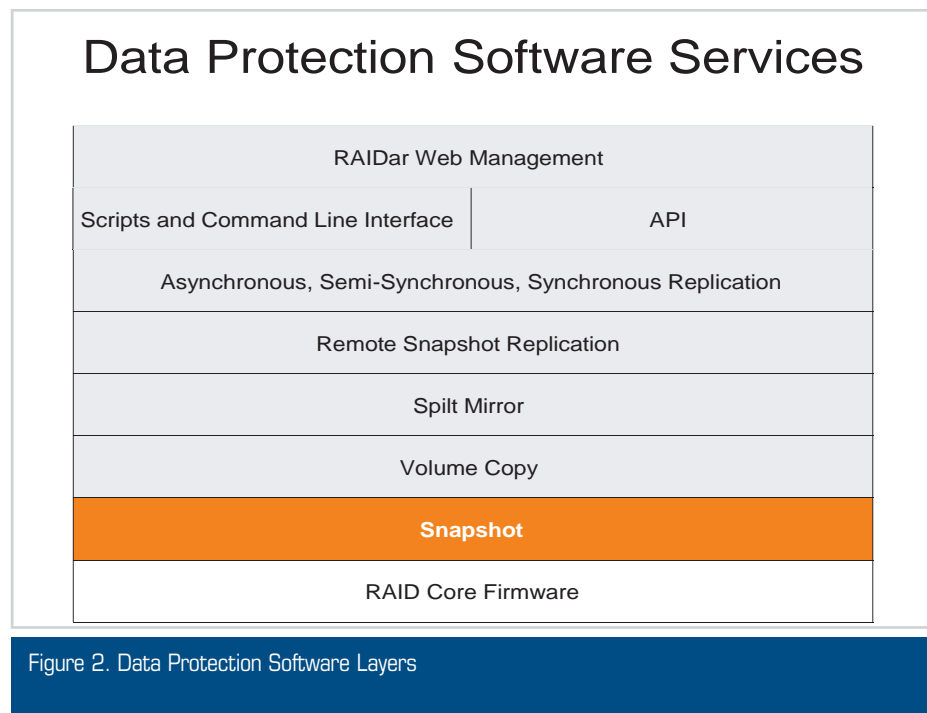
Since snapshots create backups on disk, they can be accessed instantly by users to restore from lost or corrupt data events, either back to the original application server or on an entirely different server. This reduces the cost and time to restore and access data, and helps meet the growing list of regulatory demands.

1.4.3 Regulatory Requirements

In virtually every country worldwide, a wave of new regulations demands the use of snapshots to provide point-in-time backup and recovery of information. All data can be considered material under regulatory law, including email, application files, financial data and medical records. Data must be tracked, protected, secured and accessible globally and on demand. In the U.S., regulations such as SEC 17a-4, Sarbanes-Oxley section 404 and HIPAA, to name a few, are driving new mission-critical storage strategies based on snapshot capabilities.

1.4.4 Business Continuance

Having multiple copies of snapshot data is a key way to achieve business continuance in order to protect the business from system failure, human error, natural disasters and a myriad of other occurrences. Snapshot copies are vital to providing key functions such as remote replication to hot sites, recovery to a specific point-in-time or failing over to a new vol



ume if required.

1.4.5 Rapid Application Development

Creating and releasing new versions of mission-critical applications needs to be executed flawlessly. Snapshots help IT shops achieve this in two ways. First, one of the key requirements in bringing a new application to production level is the ability to test it with "real" data. The ability to create a snapshot of the current production data, which can then be used to test applications, helps accelerate their development and deployment without impact to production servers and applications. Second, the point-in-time capabilities of snapshots allow them to go back in time if a new application release does not go as planned.

2 The Snapshot Solution

AssuredSnap™ is part of the R/Evolution Storage Architecture, and is at the core of any data protection strategy. Please refer to the white paper called

"Secure Data Protection with Remote Replication." for a high-level overview of remote replication service.

2.1 Snapshots

Snapshots are point-in-time (PIT) representations of what data existed on a master (snapped) volume at the time the snapshot was taken. Snapshots also capture the blocks that have changed since the last snapshot, similar to a traditional incremental backup providing a very efficient method of data protection. For read-only snapshots or snapshots that are not mounted, the data captured is constant, representing exactly what was on the volume at that point in time until deleted. Snapshots can be mounted read-write accessed as well.

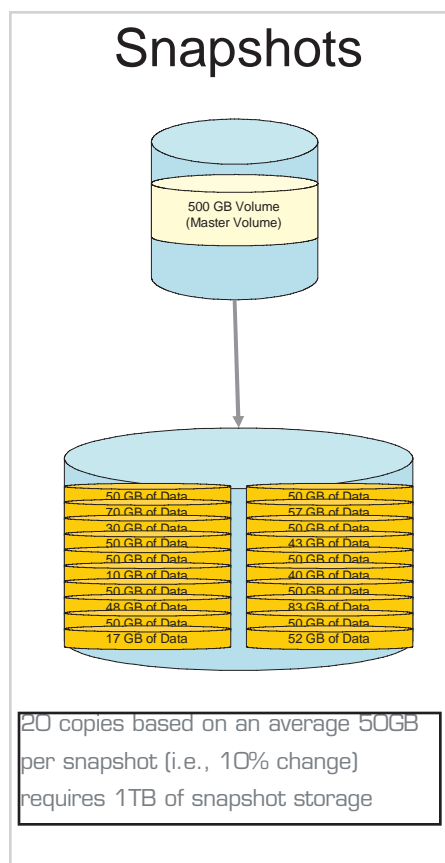
Snapshots are particularly vital for point-in-time recovery. They allow administrators to make frequent snapshots and support a wider range of point-in-time recovery options. As a result, snapshots require significantly smaller reserve space

Snapshots & Data Protection

and offer better ROI for data protection than other data protection methods such as volume copy and split-mirror.

formance of the snapshot process. Other snapshot methods can require two write operations.

been designated as 'snappable'. That is, the master volume has a snap pool associated with it and snapshots may be



2.2 Snapshots using Single Copy-On-Write

AssuredSnap uses a method called Single Copy-On-Write (SCW) to perform snapshots. This functionality is designed in such a way that for any write operation to the master volume, a single copy-on-write operation is required to preserve the snapshot data. If a block is to be overwritten on the master volume, and several snapshots depend on the existing data in that block, the data is copied from the master to the snap pool exactly one time. All snapshots dependant on that data are able to access the data from the same location. This has the effect of reducing the impact of snapshot on master volume writes. This also improves the write per-

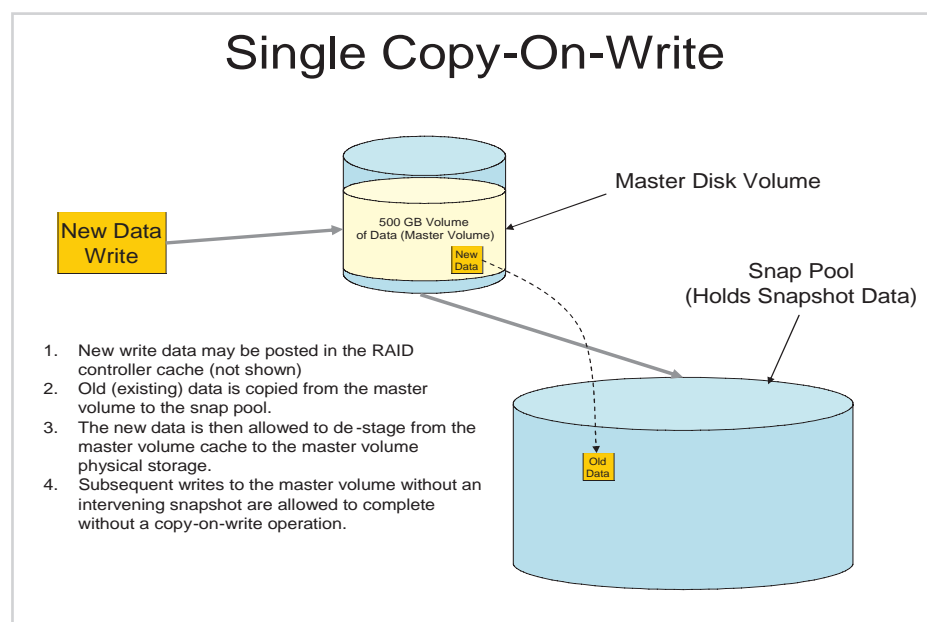


Figure 3 Snapshots using Single Copy-On-Write

The following highlights some of the key elements of the SCW method:

2.2.1 Snap Pool

A snap pool is a snapshot storage pool or snapshot reserve space. All snapshot data for a given master volume is written to the snap pool. There are two types of data that may be found on a snap pool:

- Preserve Data is data that results from a single copy-on-write (SCW) operation. Snapshot preserve data may be shared between multiple dependent snapshots.
- Write Data is data written to a specific snapshot. It is unique data - data that only applies to that specific snapshot. It is not shared with other snapshots.

2.2.2 Master Volume

A master volume is a volume that has

taken on this volume. Note that a snapshot may not be taken on a volume until it is configured as a master volume. This configuration can occur at creation, or an existing volume - referred to as a standard volume - may be converted to snappable status (i.e., a master volume).

At the time a volume is designated a 'master volume,' a snap pool must be associated with that master volume. This snap pool may be on the same disk or a different disk as the master volume. It may be smaller or larger than the master volume.

Snapshots & Data Protection

2.2.3 Snapshot Volume

A snapshot volume is a virtual volume that represents the data that existed on the master volume at the point-in-time that the snapshot was taken. If the snapshot mapped as writable, this data may contain modifications resulting from those writes.

Since the snapshot volume is virtual, all data represented as being on a snapshot volume actually exists elsewhere. It exists either on the master volume (data which has not been modified since the snapshot was taken) or on the snap pool (data which has been modified since the snapshot was taken.) This data could have been modified on the master volume or on the snapshot.

Note: Future diagrams in this document will not depict the snap pool. Instead it will be implicit: All data shown will be presented as a snapshot

2.3 Use Scenarios

Snapshots can be used to keep near-line backups of data at periodic intervals, sometimes at frequent intervals. Examples of use include the following:

- If a user inadvertently deletes a file in a tape backup environment, the file needs to be retrieved from tape, which is usually a long, tedious process. With snapshots, the file can be retrieved quickly and easily, by simply mounting and accessing the last snapshot that contains the data required.
- Snapshots are also used as the source volume for backup to tape operations - e.g. instead of backing up the master volume, backup the snapshot instead, drastically reducing backup windows.

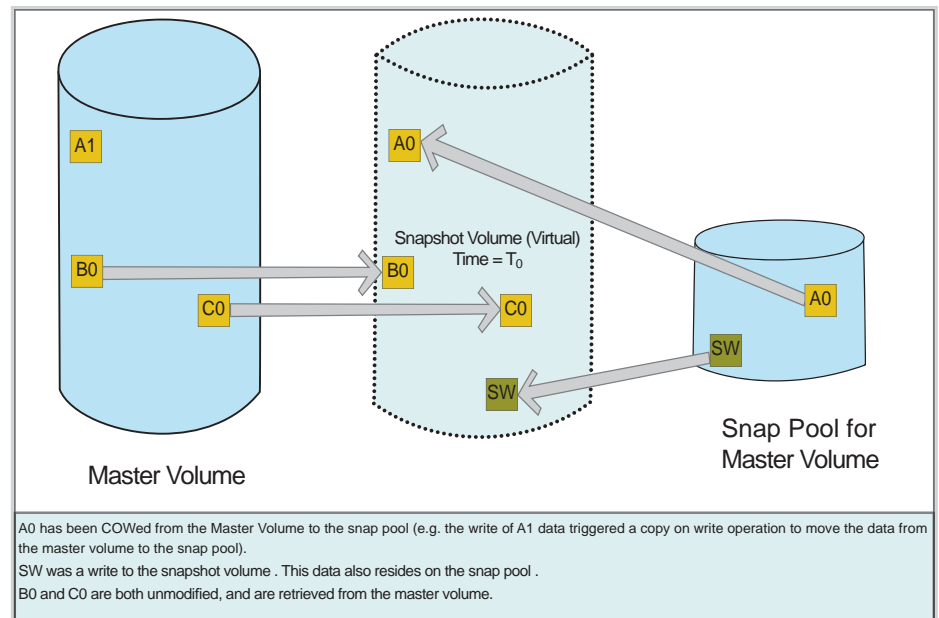


Figure 4 Master volume, snapshot volume and snap pool for master volume

- Cycle-end accounting: Here the snapshot represents exactly what data existed at the end of a cycle. No need to take certain transactions/operations offline in order to perform the cycle-end accounting.
- Writable snapshots can be used as test volumes for test organizations to run their tests on real (i.e., production) data, or for IT organizations to test out new patches, installations, etc.

2.4 Implementing Snapshots

Snapshots perform the data capture almost instantaneously and do this without copying any of the data from the volume snapped (the master volume) at the time the snapshot was taken. The snapshots are preserved by intercepting writes to the master volume and making a determination of whether or not the data needs to be preserved. The following sections will describe how snapshots affect various I/O operations.

2.5 Master Volume Writes

A master volume write is a write to a master (snapped) volume. Master volume writes are intercepted and the request is checked to determine if the data will overwrite data that needs to be preserved for a snapshot. If it does, the existing data must be copied to the snap pool before allowing the write request to complete. This is called the single copy-on-write (SCW) process. The SCW saves off (preserves) data from the master volume (the volume that was snapped) to a snap pool. It consists of a read from the master volume followed by a write to the snap pool. This data must be preserved prior to allowing a master volume write to overwrite data associated with a snapshot.

Snapshots & Data Protection

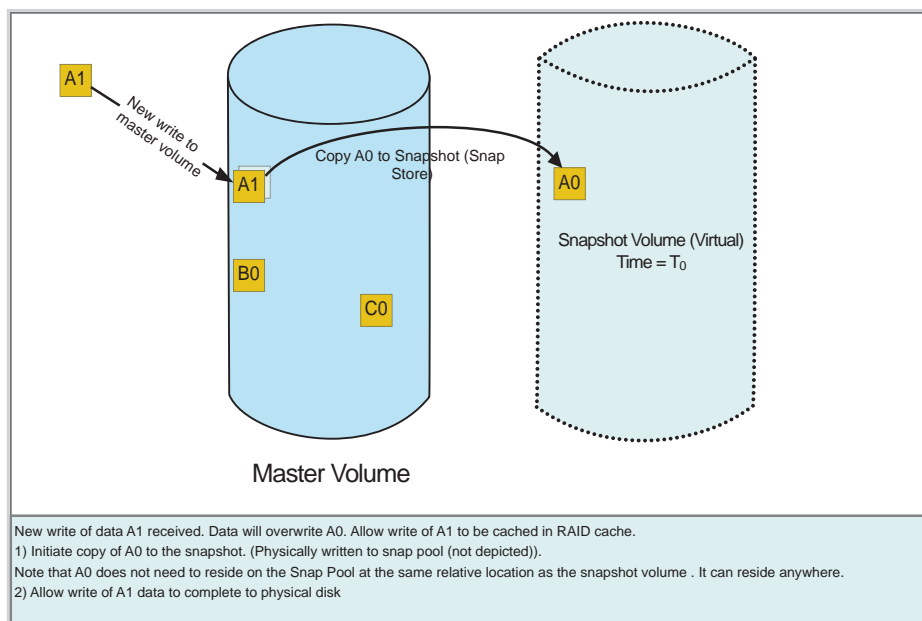


Figure 5 Master volume, writes

Note that for AssuredSnap, at most one SCW operation occurs for any piece of data on the disk. If another write to the same location on the disk occurs after the SCW operation occurs (assuming no intervening snapshots), the write will simply proceed. If multiple snapshots rely on the same data, it is copied to the latest snapshot only. (Figure 6 overleaf)

2.6 Master Volume Reads

Master volume reads are unaffected by snapshot functionality.

2.6.1 Snapshot Read

A snapshot read can be a multi-step (and sometimes slow) process. Since multiple snapshots to the same master volume are dependent on one another, a snapshot read may need to search for the data that it needs to read. This is performed by the following process:

- Check the metadata for the snapshot. If data is present, read the data and return it to the requestor.

- If the data is not present, repeat this process for each newer snapshot for the same master volume.

- If the data is not present in any of the snapshots, then read the data from the master volume and return to the requestor.

Since several snapshots may need to be searched to determine where the data is, the speed of this operation can be much slower than a read to a master volume (Figure 7 overleaf)

2.6.2 Snapshot Write

A snapshot write is a write to a snapshot volume. Because a snapshot volume is a virtual volume, the data associated with the snapshot write is written to the snap pool. This data is unique to the snapshot and is not shared by other snapshots.

Note: Because of the AssuredSnap design, a write to a snapshot volume does not require any SCW operations.

2.7 Snapshot Features and Functionality

AssuredSnap has been designed to run on a variety of platforms, each addressing a different market space.

Architectural maximums allow for it to be extendable and meet a variety of needs at different levels in the RAID solution space.

Snapshot functionality is tied to specific hardware features and limitations. For example, each snapshot, snap pool and master volume occupies a partition on the RAID subsystem. Therefore, the maximum number of snapshots that a system is capable of supporting is directly tied to the maximum number of partitions it supports. Other factors may also limit snapshot capabilities on a specific platform, including the amount of memory on the system.

2.8 Master Volumes

Any standard volume on the system may be configured as a master volume, so long as a snap pool can be mapped to that volume. In addition, a volume may be designated a master volume on creation of the volume.

2.9 Snap Pools

The architecture allows master volumes to be mapped to snap pools either in a 1-to-1 configuration or in a many-to-1 configuration. This allows the administrator maximum flexibility. For example, the administrator may allocate a large snap pool for use by master volumes that are typically infrequently accessed; alternatively, the administrator could have snaps taken and deleted on a frequent basis, such as for a volume that is only snapped for backup operations. Or the same administrator may want to take daily snapshots of another volume and allocate a dedicated snap pool for that master vol-

Snapshots & Data Protection

The snap pool and the master volume do not need to be on the same array. This means that the snap pool can use secondary storage, including drives of a different type, a different speed or configured in a different manner than the master volume. For example, the master volume may be configured as a RAID 5 logical drive of SAS disks, and the snap pool may be placed on a RAID 0 logical drives of SATA disks. Or alternatively, the master volume may be on high-speed fibre channel disks and the snap pool on slower, older fibre channel disks. Note: A snap pool is a hidden volume. It is never exposed to a host operating system.

2.10 Expansion Capabilities

Both master volumes and snap pools are manually expandable using the expand feature of the RAID controller. Support for automatic expansion of the snap pool when it nears capacity is planned. Snapshot volumes are not expandable at this time.

2.11 Accessibility

Snapshots may be hidden or may be mapped as either read only or read-write.

2.12 Failover Support

Snapshot functionality has been designed to support failover of the RAID controller. It is both failover aware and failover safe. That is, the snapshot services do not themselves initiate a failover, but if a failover occurs, the snapshot services are notified of the failover and are able to recover from the failover operation.

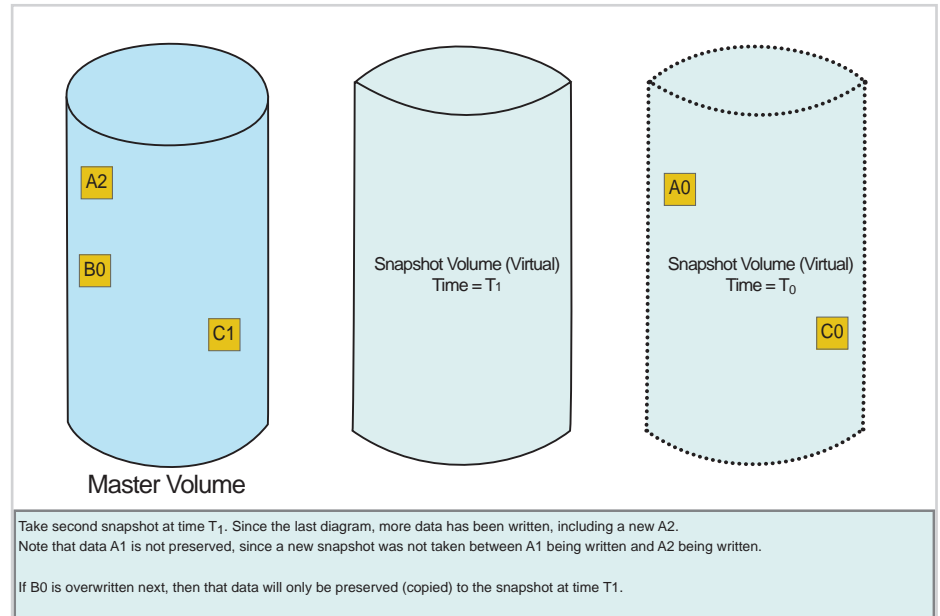


Figure 6. Second snapshot following master volume updates

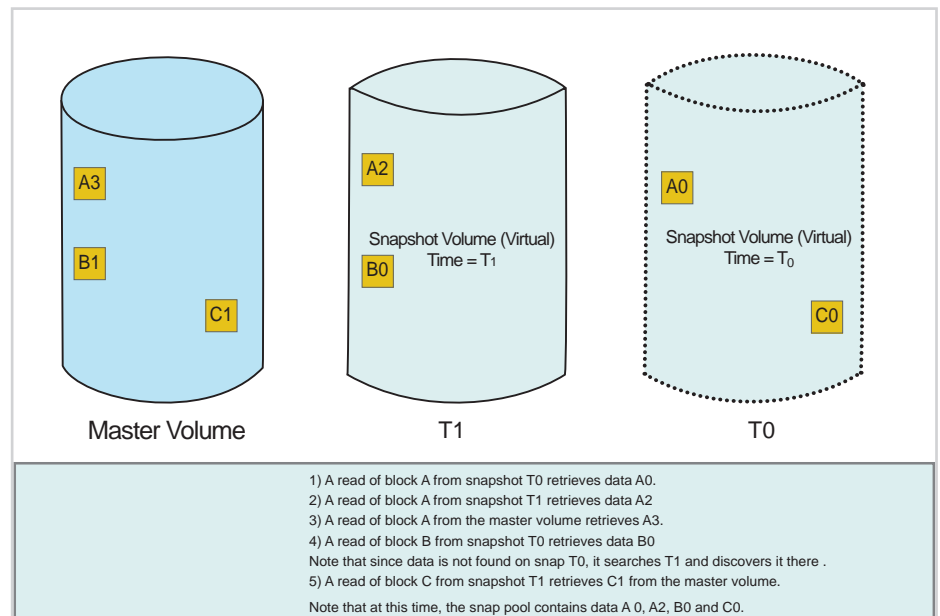


Figure 7. Snapshot reads

Snapshots & Data Protection

2.13 Snapshot Rollback

Snapshot functionality supports a rollback feature. This feature 'rolls' the master volume 'back' to the state represented by a specified snapshot. That is, it replaces all data on the master volume with the data as it existed at the time the snapshot was taken. The master volume must be taken offline briefly (dismounted) when the rollback is initiated, but can immediately be remounted and used. The rollback operation is performed as a background operation.

When the rollback is initiated, all snapshots (both older and newer) are kept. In fact, the data on the master volume as it existed immediately before the rollback was initiated may also be preserved by taking a snapshot on the master before initiating the rollback.

A distinctive feature of the rollback functionality is that it provides both a rollback and a roll-forward capability. That is, if a master volume is rolled back to an earlier point-in-time, since all snapshots (earlier and later) are still preserved, the master volume may be subsequently rolled back to an even earlier snapshot or rolled forward to a later snapshot.

While the rollback is in process, the master volume is available for both read and write-access. In addition, new snapshots may be taken on the master volume while this rollback is in process. There are a couple of restrictions on the rollback

process. Namely, the snapshot being rolled back to cannot be deleted until the snapshot rollback has completed, and a new rollback cannot be started until the first rollback has completed.

2.14 Snapshot Reset

Snapshot Reset is a means of keeping the same characteristics of a snapshot volume while throwing away the data it points to. That is, it is a logical delete-and-creation of a new snapshot with a transfer of all characteristics (volume name, attributes, LUN, etc) to the new snapshot. The reset function is very useful for backup operations, where the backup utility expects the volume to maintain the same characteristics between backup operations..

2.15 Policies and Thresholds

The snapshot function supports the setting of various policies and three thresholds which trigger these policies. The thresholds and policies correspond to the available capacity of the snap pool. The supported policies in the initial release include:

- Invalidate Snapshots (Default for Error Threshold)
- Fail I/O to the Master Volume.
- Delete Oldest Snapshot

In addition, there are plans to support an auto-expansion policy, which expands the size of the snap pool size by a specified amount in the second release of snap-

shot services. Note that manual expansion of the snap pool is supported in the initial release.

3 Summary

In this paper we discussed how AssuredSnap helps companies meet critical business needs including:

- Drastically reducing backup windows
- Providing instant restore for backup data
- Supporting regulatory requirements for data protection
- Enabling more robust business continuance
- Supporting rapid application development

AssuredSnap, part of the R/Evolution Storage Architecture, provides the core capabilities for protecting data and enabling a series of data management services. Snapshots are vital tools in helping organizations increase their competitiveness by improving IT operations for customers, suppliers and internal staff. This technology reduces the time and cost to create and release new applications, support business continuance and meet government regulations. Snapshot technology is fast, efficient and provides the best method for protecting your vital business data in real time.