Snapshot: The Core of Data Protection

ULTAMUS™ RAID

Overland STORAGE®

Simply Protected

The Core of Data Protection

Snapshot (an image or pointer-based copy of volumes) is the most fundamental tool in the data protection toolbox. It is the tool most frequently employed to backup data at specific points in time (PIT). Snapshots allow an IT manager to protect data with minimal impact to production environments, enabling a variety of IT development and application testing functions in real time.

Snapshots are increasingly being used by backup ISVs to offload backup processing overhead from application servers. Indeed, according to Gartner*, the primary use of snapshots is to optimize backup and protect data stored on the disk array without impacting production applications. This white paper reviews the ULTAMUS RAID snapshot feature. ULTAMUS RAID features a basic capability to take up to 4 snapshots at no additional cost.

Key Benefits of Snapshot

Snapshot technology enables six (6) key data management and data protection capabilities:

- **Eliminate backup windows with Point-In-Time (PIT) backups** Snapshots provide instantaneous point-in-time copies of a logical volume in an ULTAMUS RAID array. The snapshot preserves the logical volume's image as it appeared at the instant the snapshot was taken.
- **Improve operations with low impact backup** Snapshots enable backups to occur without taking production servers and applications offline. Once the snapshot is taken, it can be used for a variety of uses, e.g., staging other backups, application testing and other data management applications.
- **Fast recovery and rollback** Backup exists to enable restore. Since snapshots are slices of the current disk data, storage managers can recover instantly to specific points in time. This is an ideal process for dealing with new regularity requirements, recovering from human error, application development and other IT functions.
- **Regulatory requirement compliance** Across the globe, government regulatory requirements are changing and these changes lead to greater demand for data to be online and rapidly accessible. The data specified in a wide range of U.S. and international regulations must be tracked, protected and secured with accessibility requirements defined as well. Among the global regulations affecting data storage and protection are: SEC 17a-4, GDPdU and GoBS (Germany), Electronic Ledger Storage Law (Japan), HIPAA, FDA 21 CRF Part 11 11MEDIS-DC (Japan), ISO 18501/18509, BSI PD0008 (UK), NFZ 42-013 (France), Canadian Uniform Electronic Evidence Act, AIPA (Italy), Sarbanes-Oxley Act, PRO Government Standard (UK), European Union Regulations on Personal Data Processing and others.
- **Business Continuance** Many events are simply beyond our control. Ensuring that business continues without interruption is a leading priority for IT organizations worldwide. The requirements of business continuity have led to an increased use of disk-based, remote data replication as well as a groundswell in the use of disk instead of tape for backup and instantaneous recovery.
- **Accelerate new application development and deployment** One of the significant risks in developing, testing and deploying new IT solutions is the risk to production data. Snapshots are an ideal way to protect against mishaps. IT managers can now take a snapshot that can be mounted as a new volume to support these tasks. If something goes wrong, the snapshot can be discarded and no harm is done. Snapshots gained great popularity when Y2K testing was a top priority.

* Software Market Trends, Carolyn DiCenzo, Gartner PlanetStorage Summit 2005

What is an ULTAMUS RAID Snapshot?

ULTAMUS RAID snapshots feature an additional level of data protection while offering a means to improve production data utilization. Snapshot is designed for users whose data availability cannot be disrupted by routine management functions. Point-in-time images of logical drives are saved for near instantaneous roll-back of updates. Snapshot supports round-the-clock processing as it stages data for operations such as backup, data mining/analysis and work distribution.

Snapshot is especially useful in operations requiring a quick copy of data.

Snapshot images are fast and efficient, with minimal disruption to the user. A momentary suspension of processing allows application data to synchronize to a known state preparing the snapshot volume for use.

A snapshot logical drive is a virtual point-in-time image of a source logical drive. It is the logical equivalent of a complete physical copy, but is created much more quickly and requires less disk space. Snapshot logical drives appear and function as standard logical drives. They are host addressable and can be read or copied to create a real copy of a point-in-time.

Snapshot is an integral component of the ULTAMUS RAID firmware that runs in the RAID controller rather than on a host—maximizing performance while ensuring full availability to data as applications continue to process. Through its copy-on-write technology, snapshots preserve data in its original form in an "Overwrite Data Area". This functionality requires a minimal amount of storage capacity, typically between 10 percent and 20 percent of the source logical volume, enabling it to generate several snapshots within the space that would be required for a single mirror.

Fast recoveries to point-in-time versions of logical drives are accomplished via "snap-back" operations. A snap-back reverses all the updates made to all of the data stored on a source logical drive to the point-in-time the snapshot was established. Users can quickly back-out erroneous changes and recover critical data.

As shown in Figure 1, ULTAMUS RAID snapshots employ a Copy-On-Write (COW) process meaning that the snapshot copy and the primary data share the same data source, except when new data is created. The snapshot keeps the unchanged file and the primary data keeps the updated file. It is important to note that a snapshot is not a complete physical replication of the original disk when it is created, only a virtual copy, as all data written to the source volume will still reside on the source disk.

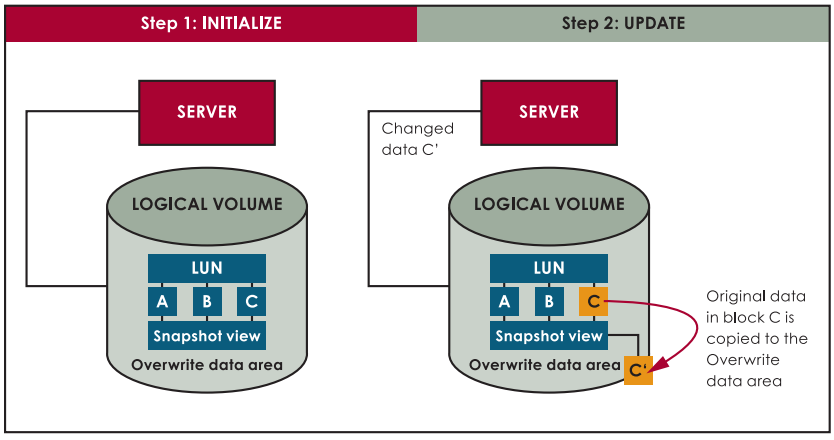

Figure 1. ULTAMUS RAID snapshot

ULTAMUS RAID snapshots take place above the RAID functionality, enabling snapshots to benefit from the redundancy and error recovery of the RAID system. Different RAID levels may be used for the source and snapshot logical volumes. This snapshot architecture also enables snapshots of primary data stored on SAS disk drives to be maintained on logical volumes constructed from lower-cost SATA disk drives.

When in operation, servers continue to read from the source volume and write to a special reserved area with the user unaware that the snapshot has occurred. Once the snapshot is released and the source volume is "unfrozen", data that was written to the source volume while the snapshot was active is re-integrated into the source volume. This activity is transparent to the user.

ULTAMUS RAID snapshots produce a result similar to a full volume copy, but snapshots are captured in far less time—and require far less disk space—than would be required for a full volume copy. Once a snapshot is taken, a server writing to the source logical volume transparently causes pre-defined segments of the source logical volume to be copied to the snapshot before allowing the write to continue, thus preserving the original data on the snapshot. Another server can then be allowed to "view" (or mount) the frozen data. A backup server can mount the snapshot volume and backup the frozen data to a REO™, NEO™, ARCvault™ or other backup device while the application server continues in production. Using snapshots creates a backup environment that dramatically minimizes the impact on day-to-day operations.

ULTAMUS RAID Snapshot Implementation

The storage manager configures an Overwrite Data Area (ODA) in a reserved area of the array. The ODA can be of any capacity, can be mapped to a drive array of any RAID level using any chunk size and stripe size. Each stripe in the ODA is 64KB in size. To minimize the performance impact, portions of the ODA are cached. The cached ODA segments are periodically committed to permanent media.

When a write to the source logical volume occurs, the snapshot metadata is checked to determine whether the copy segment to be updated has already been copied from the source to the snapshot. If the copy segment was previously copied, then the write proceeds to the source because the original data is preserved on the snapshot. Otherwise, the original data within the copy segment is copied from the source to the snapshot (Copy-On-Write snapshot).

The snapshot metadata (i.e., the mapping information) is stored in a hidden area of the ODA accessible only to the RAID controller firmware, so it enjoys the same level of protection as the source logical volume. The snapshot area is aggressively cached in the RAID controller's memory to minimize the impact on snapshot access performance.

Resource Considerations

The storage manager configures the Overwrite Data Area (ODA) in a reserved area of the array. The ODA can be of any capacity, can be mapped to a drive array of any RAID level using any chunk size and stripe size. Each stripe in the ODA is 64KB in size. To minimize the performance impact, portions of the ODA are cached. The cached ODA segments are periodically committed to permanent media.

A snapshot of a logical volume can be mounted for access through a dynamically created LUN called a "Virtual LUN". Each incoming read command will be scanned to determine from where the data is to be read.

Once the ODA area has been set up and initialized, the disk array is ready to receive snapshot preparation instructions. The sequence of events to prepare for a snapshot is as follows:

- Select the logical volume on which the snapshot is to be performed.
- Select the ODA to use for this and any future snapshots of the logical volume.

- Select if the snapshot data is to be exposed through a logical volume or kept internal only to be used by the controllers snap-back function.
- If the snapshot data is to be accessed through a logical volume, select the LUN number through which the resulting snapshot Virtual LUN will be made available.
- Select the protection stripe size to use for the copy-on-write. This defines how many sequential disk blocks that will be copied for each copy-on-write. The user can select a "default" option to let the controller select a recommended value or the user can select a value from a list of possible stripe sizes.

When a write to the source logical volume occurs, the snapshot metadata is checked to determine whether the copy segment to be updated has already been copied from the source to the snapshot. If the copy segment was previously copied, then the write proceeds to the source because the original data is preserved on the snapshot. Otherwise, the original data within the copy segment is copied from the source to the snapshot (Copy-On-Write snapshot).

The snapshot metadata (i.e., the mapping information) is stored in a hidden area of the ODA accessible only to the RAID controller firmware, so it enjoys the same level of protection as the source logical volume. The snapshot area is aggressively cached in the RAID controller's memory to minimize the impact on snapshot access performance.

Once a snapshot has been initiated for a Logical Volume, all incoming writes to that Logical Volume will be temporarily halted. During this pause, the firmware will check whether the data located at the address of the new incoming write has been saved in the ODA. If the data located at the address of the new incoming write has not been saved, then the controller will do so before allowing the paused write command to continue. If the data located at the address of the new incoming write has been previously saved, then the pause is cancelled and the write is allowed to continue.

Starting a Snapshot

Snapshots are started manually. Once a snapshot has been started, or triggered, the RAID controller will start monitoring all incoming commands. Commands received before the start of the snapshot will be processed normally and will not be affected by the snapshot. Commands received after the snapshot has been initiated will all go through the snapshot mechanism.

Any outstanding read or write command that was received by the RAID controller prior to the start of the snapshot - whether cached or not - will be executed as normal read/write commands. It is the snapshot initiation alone that determines if a command from a host computer will be affected by the snapshot process.

Snap-Back

Snap-back enables rolling back a snapshot onto its original logical volume. The sequence is as follows:

- Dismount the LUN that is to be snapped back.
- Select the snapshot to be snapped-back. This automatically starts the snap-back operation.
- You will be unable to access the Host LUN while it is being snapped back. Any SCSI commands issued to this LUN will be rejected with a check condition.
- When the snap-back completes regenerating the LUN it might be necessary to reboot the operating system.

After a snap-back operation completes, the controller will delete all snapshots that were taken after the snapshot is snapped-back. Snapshots that were taken at a time prior to the snapshot used for a snap-back will be left intact. The user can perform another snap-back operation to any of these older snapshots without risking data corruption.

Deleting Snapshots

Snapshots can be deleted manually at any time. Deleting one snapshot will not require the deletion of any other snapshots, including those that were taken at a time after the snapshot that was deleted.

Performance Considerations

Read from a Snapshotted LUN

Read commands will be unaffected by any snapshot operation. The RAID controller will continue with normal execution of read commands.

Write to a Snapshotted LUN

Write commands to a LUN that has not been snapshotted will not be affected and the RAID controller will continue with normal execution of these write commands. Write commands to LUNs that are being snapshotted are examined to determine if the current data will have to be copied over to the ODA before the new write data can be committed to disk. Before this can happen it will be necessary to pause the write command to temporarily prohibit further processing. This allows for further examination of the write command and possibly start a copy-on-write sequence.

Application Considerations

Since snapshots are virtual duplicates of existing logical volumes, care must be taken when presenting snapshots to hosts. Some file systems and operating systems write signatures on logical volumes and use these to uniquely identify those logical volumes. When a snapshot is taken of a logical volume with a signature, all of the data on the logical volume is duplicated including the signature. Because the signature is no longer a unique identifier the host can easily confuse the source and snapshot and unpredictable behavior may result.

Microsoft Window's Dynamic Disk Management (DDM) included in Windows Server 2003 is one such application. Because DDM places a footprint directly on the media, snapshots cannot be used with this utility.

It is very important for the application that initiates the snapshot to make absolutely sure that all data and caches are flushed to disk before the snapshot is triggered. This includes the applications own data caches, other applications running concurrently and their data caches and the operating system caches that might or might not be in use. If these caches are not flushed prior to triggering a snapshot, the resulting snapshot image will be out of sync with the application data and therefore corrupted. A snap-back from such a corrupted snapshot image will result in serious data corruption and loss of data.

The Windows NT and the NT filesystem (NTFS) perform "lazy writes" and "lazy commits" to improve performance. These "lazy" file system activities mean that when data is written, the "lazy" processes may cause the application to act as if the data has been committed to permanent media when it is actually still in cache. To avoid the effects of "lazy" file system processes, it is strongly recommended that file systems and I/O activity be quiesced before taking a snapshot of a logical volume containing Windows file system data. Similarly, it is also recommended that file systems and I/O activity be quiesced before taking a snapshot of a logical volume containing Linux or Unix file system data.

Snapshot Use Cases

Archival of Primary Data

The most common use for snapshots is to preserve the data of an actively used logical volume so that it can be backed up. In this case, a snapshot is taken of a source logical volume and the backup application backs up the resulting snapshot logical volume instead of the source logical volume. This allows writes to continue on the source logical volume while the snapshot preserves the data at the point in time of the snapshot.

Test Area

Another use for snapshots is as a "test area" where data can be modified and eventually deleted without altering the contents of the source logical disk. In this scenario, the user takes a snapshot of a database, uses it for testing (i.e., development or training) purposes, and then deletes it after completion of the test. This allows the source logical volume to remain unaffected by the test because all of the test writes are directed to the snapshot. While the test is being done, a production application can continue to modify the source logical volume.

Summary

ULTAMUS RAID supplies a limited snapshot capability that is supplied as part of the core RAID system, delivered to storage managers with less budget impact than competing products. ULTAMUS RAID snapshot technology is fast and efficient and raises the bar for data protection in the entry-level RAID space.

ULTAMUS RAID is not just a collection of new technologies; it is a well-balanced and flexible platform that can meet the performance, reliability and cost requirements of near-line applications. ULTAMUS RAID is a next generation RAID platform that incorporates the latest technologies based on the latest industry standards.

ULTAMUS RAID Features and Benefits

The architecture underlying ULTAMUS RAID couples comprehensive data protection with high performance and powerful, intuitive management. Overland's industry leadership delivers investment protection to storage buyers through leading-edge, "future proofed" technology. ULTAMUS RAID arms storage managers with the tools they need to implement storage strategies and build IT infrastructures that align with today's and tomorrow's business and data management needs.

Technology	Benefits	Overland Leadership
RAID 6 Hardware Accelerated	• Higher Availability	• Better for mission-critical application availability. • Overland's hardware accelerated RAID 6 allows storage to remain online in the event of two drives failing simultaneously.
SAS Drives	• Higher Performance • Lower Costs • Investment Protection • Higher Availability	• Ideal for email, web serving, database and other performance-intensive applications • provides the performance and reliability of Fibre Channel drives at 30% lower cost • SAS and SATA II drives can be combined in the same RAID chassis to support applications with vastly different capacity and performance requirements.
Enterprise SATA II Drives	• Higher Availability • Higher Performance • Better Price/TB • Investment Protection	• Better performance for D2D backup and media • Built for 7 x 24 operating environments • 1.4 million hour MTBF at 100% duty cycle—more than double the reliability of desktop SATA II drives • Directed Offline Scan monitors every write and performance diagnostic tests when drives are idle, enhancing data availability • Perpendicular Recording reduces the number of moving parts, for higher availability
SAS expansion & scalability	• Grows with your application requirements	• Effortlessly scales up to 60 drives in 12-drive/2U increments • No performance penalty
4Gb Fibre Channel	• Higher Performance • Investment Protection	• Better performance for all applications. • 4Gb FC improves IOP intensive applications with reduced latency and improves streaming applications with higher data rates
Snapshots	• Disaster Recovery • Data Protection	• Snapshot technology offers the best value in data protetion by creating low-impact copies of data that support point-in-time recovery
Pricing	• Low Acquisition Cost	• Excellent price/TB for RAID storage • Integrated RAID ASIC technology reduces controller costs and an optimized supply chain reduces platform costs.

Save

Worldwide Headquarters
4820 Overland Avenue
San Diego, CA 92123 USA
TEL 1-800-729-8725
TEL 1-858-571-5555
FAX 1-858-571-3664
EMAIL sales@overlandstorage.com

United Kingdom (EMEA Office)
Overland House
Ashville Way
Wokingham, Berkshire
RG41 2PL England
TEL +44 (0) 118-9898000
FAX +44 (0) 118-9891897
EMAIL europe@overlandstorage.com

France Office
Overland Storage
126 rue Gallieni
92643 Boulogne Cedex France
TEL +33 (0) 1 55 19 23 93
FAX +33 (0) 1 55 19 25 02
EMAIL europe@overlandstorage.com

Germany Office
Humboldtstr. 12
85609 Dornach Germany
TEL +49-89-94490-214
FAX +49-89-94490-414
EMAIL europe@overlandstorage.com

Asia Pacific Office
Level 44, Suntec Tower 3
8 Temasek Boulevard
Singapore, 038988
TEL +65.6866.3847
FAX +65.6866.3838
EMAIL asia@overlandstorage.com

www.overlandstorage.com

Overland STORAGE®
Simply Protected

 URAI-SNWP0207-04