White Paper

Rapid Recovery of Database Applications Using Low-Capacity Snapshots

Minimizing the Recovery Time Objective for Database Applications

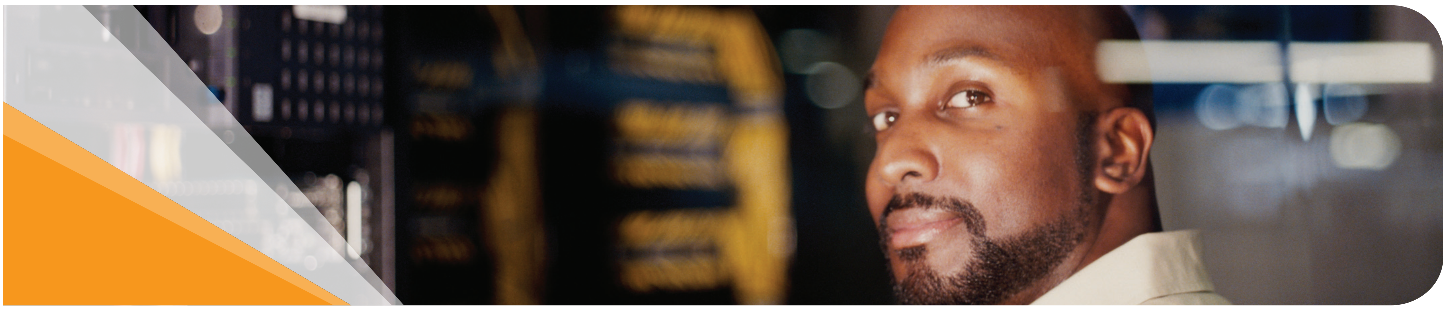

Executive Summary

For many years now, organisations have relied upon standard methods of backup and recovery using tape devices and backup agents or scripts. While this has provided a level of security and comfort to a business, the needs of the business have continued to change and standard methods of back-up cannot keep pace.

Technology has made the world a smaller place, people demand access to more and more information and are unwilling to wait for it. No-where is this seen more dramatically than in the explosive use of database technology to provide online or dynamic content.

The backup window for many organisations and applications has been reduced from hours to minutes and even down to seconds. How can an IT staff achieve the goals of the business without simultaneously pushing capital and operating expenditure into an ever increasing spiral upwards?

How does IT quantify the risks to an application and determine the steps required to reduce risk to an acceptable level for business continuity?

How much data can be lost from the recovery process, how many man hours are lost due to the unavailability of their critical database applications?

This white paper aims to cover the benefits of using snapshot technology, the differences among snapshot techniques for database recovery and then recommends how best to provide a rapid recovery of critical information.

Introduction

Many organisations today confuse business continuity with disaster recovery.

Business Continuity ensures that enough resilience has been built into the infrastructure to negate single points of failure to ensure "business as usual" availability. Disaster Recovery ensures that the Recovery Point Objective (RPO), the point where the application data can be successfully recovered and the Recovery Time Objective (RTO), the point where the application is live again are within the guidelines of the business.

The traditional methods of backup and restore make the RPO and RTO an expensive strategy for a business in terms of loss of business and people time. To address the changing needs of business, new ways of addressing the demand for "instant backup" of ever increasing data volumes need to be found.

In particular, this paper will discuss the use of LSI StoreAge SVM™ snap-shot-enhanced rapid application recovery solutions, and examine how using snapshot-enhanced technology can provide higher performance, high data integrity and faster recovery compared to other approaches.

Examining Snapshot Technologies

There are many forms of snapshot and distance replication technologies available, but the typical storage industry answer has been to provide snapshot capabilities within the storage array, or to provide volume managers with snapshot capabilities at the host level to reduce the backup window.

Copy on Write (COW) Snapshots

Host based snapshots rely upon the use of a host volume manager to manage the snapshots and provide some way of enabling the backup agent to stream the data to the backup server. This provides the ability for the backup agent to quiesce the database application and flushes the file system to disk and produce a snapshot of the volume. A view or image of this can then be presented to the backup agent and streamed across the LAN. This provides part of the solution in that the application is "frozen" for a few seconds while the data is flushed to disk and then thawed to resume normal operation.

Once a snapshot is in operation on a volume, there is a performance penalty incurred because before a block of data can be changed it must be read and written to the snapshot before the original write can complete, this is what the industry terms as a Copy on Write Snapshot. On a heavily loaded system this may have a severe impact to system performance and end user experience, especially if the backup agent is streaming the data at the same time. This also limits the practical number of snapshots allowed at any given time.

Many off-host based snapshots; occur within the RAID Array, thus relieving the Database Server from the burden of the Copy on Write overhead. The additional benefit of array based solutions is that the snapshot copy can be presented directly to the backup server, thus providing server-free or LAN free backup facilities. However, this will almost certainly cause an increased load upon the RAID Array controller(s). For each LUN where a snapshot exists the controller has the burden of the Copy on Write operation multiplied by the throughput of the hosts.

Storage Virtualisation has enabled the snapshot technologies to be brought into the SAN Fabric. The drawback remains with many of these solutions that using Copy on Write snapshots increases the load on these systems, thus requiring large CPU power and large cache subsystems to provide functionality even for basic volume management.

Redirect on Write (ROW) Snapshots

An alternative to Copy on Write snapshots is to provide highly scalable block level snapshots utilising Redirect on Write (ROW) techniques.

ROW Snapshots provide large scalability, allowing fully read writable independent volumes to be presented using low capacity rather than full capacity snapshots, these capabilities allow the business to use the snapshots for backup, application testing, patch testing, development and much more. A major benefit of this type of snapshot technology is to provide a larger number of snapshots per volume with rapid recovery of an application using snapshot rollback to restore to a previous point. This improved granularity serves to further reduce the RPO for applications using ROW snapshots.

Additional advanced capabilities such as using snapshots to restore data from view techniques can gain valuable time. For example Redirect on Write snapshots allow the data from a snapshot to be modified or repaired before being presented and mounted back to the original source, this ability reduces downtime (RTO). When database downtime is a premium this can help reduce the business impact of such an event to a minimum.

A requirement for the use of most snapshot technologies is the management of disks by grouping them together into "Consistency Groups", this aids good database design of separating database and log files onto separate disks. This function allows for the creation of consistent snapshots across all databases, log files and volumes. An application server can therefore be more efficiently managed and ease of snapshot automation can be provided via scripting. This is critical for large database installations in which multiple disk entities are presented to the database application. Failure to provide a solution that captures all of these components at the same point in time would leave the database application in an inconsistent and therefore unusable state.

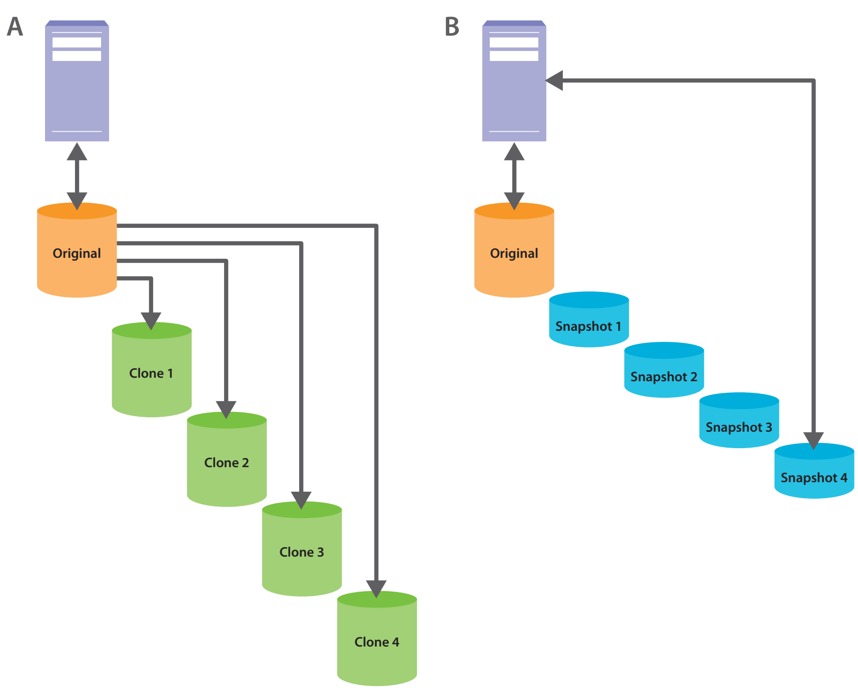

Full Capacity Snapshots (A)

Original volume size x number of snapshots + 1. Host writes to Original blocks, Copy on Write I/O overhead on Original Source Volume, limited number of snapshots supported.

Low Capacity Snapshots (B)

Host writes to new snapshot region only. More scalability on snapshots and features.

Database Application Challenges

Application Integration

One of the most essential factors for any database application is to guarantee that any solution provides a mechanism for ensuring data integrity.

Whichever solution a business decides to implement application integration is critical. If a business takes a snapshot of its mission critical database without having quiesced the application data, then the data is in a crash consistent state and no better than pulling the power chord from the server.

This leads to questions such as:

- Can the database recover?
- Does it depend upon what operation was being carried out?
- Can we roll forward any uncommitted transactions?
- How quickly can the database admin recover the application?
- How much data did we miss?
- What if we cannot recover the data?

Application integration permits a database to be transparently quiesced, allowing the buffers to flush to disk, thusly insures each snapshot will have integrity. The powerful SANAPI scripting application supports database applications such as Microsoft SQL Server, Microsoft Exchange and Oracle.

Database applications which only provide command line executables like the Sybase 12.5 "quiesce database" command to flush data to disk prior to creating a snapshot can also be utilised by integrating them into SANAPI scripts. The SAN CLI library provides interfaces for snapshot creation that provides near instantaneous snapshots of applications in a consistent state. Snapshots of volumes can then be presented to a backup server for back-up onto SATA disk, or VTL and then eventually traditional Tape Libraries. Immediate savings on database backup agents can be realised using this standard component of the package.

This is a powerful combination that saves time and resources in recovering applications to a known consistent state.

Application Recovery

Many businesses think about their backup strategy but do not think about their recovery strategy. Each application has its own merits and needs from a backup and recovery perspective. The business sponsors should agree on the recovery priority for the systems that they depend upon and a recovery strategy should be produced.

Application recovery can be achieved from traditional backup methods; this would take you back to the point at which the backup was instigated. With careful planning it may be possible to recover to a relatively recent state by replaying log files into a restored database.

Snapshot backups may allow you to present a view of the data at the time of the snapshot and copy the database files back across the damaged database; however this may take a large amount of time to complete the copy then roll forward the logs.

Some snapshot based solutions may only allow you to rollback to the point that the last snapshot was taken. The risk here is that if you rollback and the data is still corrupt you are forced back to restore from traditional backup methods which will greatly increase the recovery time objective.

Advanced solutions using StoreAge SVM™ low capacity snapshots allow for a full read / write copy of the data to be presented to the original or recovery database server to check the integrity of the data. If a snapshot is shown to be corrupt then a previous snapshot can be presented very quickly, without compromising the integrity of the original data. A test of any roll forward can then be implemented without putting at risk the original snapshot data, thus ensuring that if there is a problem with the roll forward it can be quickly reversed and any roll forward parameters tested again and again until a successful roll forward is achieved. If the roll forward is successful then a restore of the database over the original volume can be completed in a timescale not possible without snapshots.

Without these techniques, if the roll forward does not complete successfully, then the business would need to restore from tape, roll forward, test and then start over again using an earlier version of the back-up if it failed. This is an un-welcomed, expensive and time consuming event that would apply additional pressure onto the technical team.

If the business does not need to roll forward the database and is happy to accept that the Recovery Point Objective is that of the last snapshot time, then a simple rollback will be performed after testing via the StoreAge low-capacity snapshot application (StoreAge multiView™).

Many companies deploy cluster solutions to maintain the application availability. Database applications are installed onto the cluster nodes and an instance of the database will run on one of the nodes. In the event of a failure in the node running the database one of the surviving nodes will take ownership and start the database. This presents some complexities into the environment as the database application usually needs to be cluster aware and the resources that the database application depend upon must have the ability to move from one node to another. Furthermore the applications that backup these Database applications must be cluster aware and this generally makes the solution expensive. With the use of the snapshot technologies discussed here, the complexity of the backup solutions for clustered environments can be greatly simplified.

The solutions here address localised disasters, such as human error, local SAN failures and host failures, but what if we temporarily lose access to the environment due to adverse weather producing a brown out or terrorist action disabling local facilities.

Application Data Integrity with Distance Replication

Many businesses require that critical database applications must also be available at a secondary location and any corruption that occurs at the source will be presented to the secondary site. Advanced snapshot capabilities represent a method to validate data and do so without interrupting operations. This will serve to actually enhance recovery times from a corrupted database.

There are two issues that distance replication or mirroring must address when it comes to data integrity:

1. The vast majority of disasters are not a single, instantaneous event. Instead, disasters usually unfold over a period of minutes or even hours (intermittent power outages, communication link disruptions, disk drive failures). Intermittent failures are the most difficult to handle, since they can corrupt the integrity of data not just once but several times during the course of an unfolding disaster.
2. The total time needed to recover from a disruption. In a synchronous mirroring approach, all data – whether corrupt or not – is immediately replicated to the secondary storage device. In other words, a database or file system that is corrupted at one end will become corrupted at the other end as well. Recovering from this type of corruption typically takes hours or even days, and in some instances may be nearly impossible.

If the wrong recovery choice is made then the problem could be compounded or a total data loss could result.

If replication is used without snapshots then it is virtually impossible to go back to a known state, and traditional backup/restore methods must be employed.

If full capacity snapshots are used then a full data copy must occur and then the database consistent state may be so out of date that the business cannot recover.

Some snapshot technologies only allow for the snapshot to occur at either the primary or the secondary site, this then means that the application integrated snapshot remains local or the replicated snapshot is not application consistent

StoreAge SVM™ provides solutions to these problems. By using low capacity snapshots there are many points in which to check for consistent data sets. Once a consistent state is found the database can be brought online and if possible rolled forward. Application integrated snapshots are replicated to the DR site and the consistent data can be brought online either in the DR or Primary site very quickly without the need to recover from tapes. To further enhance the rapid recovery asynchronous mirror failback (StoreAge multiMirror™) facilitates true DR testing and enables replication of just the block changes from the DR site back to the Primary site without a complete block level synchronisation of all data. This saves money otherwise spent on communications link bandwidth to support data change rates.

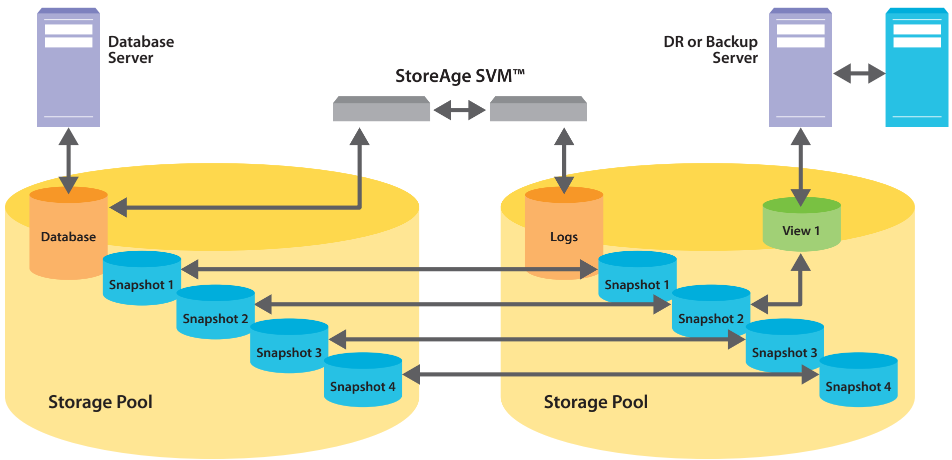

Database volume and snapshot volumes replicated from primary or DR site. DR site can be used for offsite backup or brought online and block differences are replicated to primary site.

For more information and sales office locations, please visit the LSI web sites at: lsi.com lsi.com/contacts

North American Headquarters
Milpitas, CA
T: +1.866.574.5741 (within U.S.)
T: +1.408.954.3108 (outside U.S.)

LSI Europe Ltd.
European Headquarters
United Kingdom
T: [+44] 1344.413200

LSI KK Headquarters
Tokyo, Japan
Tel: [+81] 3.5463.7165