


Cost-Effective Disaster Recovery with Snapshot-Enhanced, Any-to-Any Data Mirroring

Maximizing The Data Integrity, and Performance of Asynchronous Mirroring in Fibre Channel SAN Environments



For many years now, mirroring of production data for Disaster Recovery (DR) purposes has been available for both mainframe and open systems computing environments. Unfortunately, due to the cost and complexity of these types of solutions, they have mostly been deployed in larger organizations with larger IT budgets.

However, the demand for cost-effective disaster recovery solutions has never been higher, as organizations of all sizes realize the value of their stored data and the high costs associated with any type of downtime. In fact, many organizations today have established a DR strategy or requirement, but have not actually implemented a solution due to budgetary or other restrictions.

Fortunately, a new generation of affordable data mirroring solutions has emerged that brings sophisticated DR capabilities to virtually any size organization. Some of the key features of these solutions include:

- Compatibility with a wide range of existing storage devices
- Ability to mirror data between different devices from different vendors
- Use of snapshot-enhanced mirroring to ensure data integrity and rapid recovery after a disaster or other disruption

In particular, this paper will discuss the use of StoreAge's multiMirror™ snapshot-enhanced asynchronous mirroring solution, and examine how using snapshot-enhanced asynchronous mirroring technology can provide higher performance, high data integrity and faster recovery compared to synchronous approaches.

With synchronous mirroring, each time an application attempts to write data to disk, the transaction is sent to both the local and remote storage devices in parallel.

With asynchronous mirroring, each write transaction is acknowledged as soon as the local storage device completes the request, even if the remote system has not yet received and/or processed the request.

LSI offers both synchronous and asynchronous mirroring solutions. Because of this, we are able to impartially analyze each customer's requirements before recommending an appropriate solution.

One of the most critical factors in selecting a mirroring solution is the ability to ensure the integrity of the data being replicated between sites. Obviously, it makes little sense to mirror data unless you are confident that the data will be usable when needed.

Synchronous vs. Asynchronous Mirroring – A Definition

In a synchronous mirroring environment, each time an application attempts to write data to disk, the transaction is sent to both the local and remote storage devices in parallel. It is not until both devices have committed the write to disk that the system acknowledges that the transaction is complete. In other words, the application that initiated the write must wait until it receives the acknowledgement before it can continue on to the next task.

In an asynchronous environment, each write transaction is acknowledged as soon as the local storage device completes the request, even if the remote system has not yet received and/or processed the request.

Performance Issues

From a performance standpoint, a synchronous approach will always incur some level of performance degradation – even when the two storage devices are nearby – simply because both systems must complete each transaction before the application can continue.

On the other hand, since an asynchronous approach acknowledges the write request without waiting for confirmation from the remote storage device, the performance of the system is virtually identical to that of a non-mirrored system.

Cost Issues

From a cost standpoint, a synchronous approach usually requires higher bandwidth and more equipment in order to maintain acceptable performance. This is true for several reasons:

- *Bi-directional traffic:* Since each write transaction must be transmitted to the remote system and an acknowledgement received back, the transmission infrastructure must have sufficient bandwidth and performance to avoid becoming a bottleneck in this process.
- *Latency during peak periods:* A worst-case scenario should be factored into the design of the transmission network, since spikes in data activity could degrade overall performance, or cause application time-outs due to extended latencies.
- *Scalability:* SANs are designed to support multiple host servers, but as the number of hosts in a SAN increases, the synchronous mirroring infrastructure may not easily or economically scale to accommodate the increased data traffic.

As a result, a synchronous solution usually requires some level of over-provisioning of both the bandwidth and the available switch ports, in order to ensure sufficient performance during peak periods.

On the other hand, an asynchronous solution usually requires minimal bandwidth, as bi-directional traffic is significantly lower and communication latencies do not affect application performance. In addition, asynchronous solutions are designed to flexibly adapt to spikes in activity by buffering transactions in a queue until sufficient bandwidth becomes available to complete each transaction.

Finally, keep in mind that LSI offers both synchronous and asynchronous mirroring solutions. Because of this, we are able to impartially analyze each customer's requirements before recommending an appropriate solution.

Data Integrity During Mirroring

One of the most critical factors in selecting a mirroring solution is the ability to ensure the integrity of the data being replicated between sites. Obviously, it makes little sense to mirror data unless you are confident that the data will be usable when needed.

There are two issues that 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). And, 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 filesystem 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.

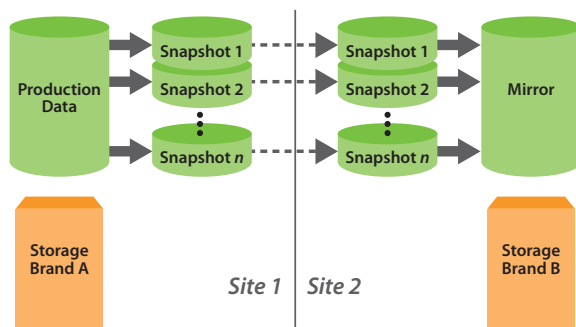
Snapshot-Enhanced Asynchronous Mirroring

To address the two data integrity issues discussed above, LSI has introduced StoreAge multiMirror snapshot-enhanced mirroring. This technology combines platform-independent, any-to-any, asynchronous mirroring with low-capacity, instant point-in-time snapshots to ensure data integrity between sites while enabling rapid recovery after a disaster.

There are several important differences in the snapshot capabilities utilized by StoreAge multiMirror. Compared to approaches used in other systems, LSI snapshots offer the following advantages:

1. Rather than create full-size copies of volumes, LSI creates instant volume snapshots that begin at zero capacity.
2. Disk space for snapshots does not need to be pre-allocated or reserved, enabling more efficient use of existing capacity.
3. Consistency groups allow snapshots to be created of logical groupings of volumes, such as the data and log files in a database.
4. A scheduling feature allows the user to specify how frequently snapshots are created (i.e., every few minutes).
5. Application-aware data consistency capabilities allow applications such as databases to be quiesced prior to creating snapshots, ensuring the data integrity of each snapshot's contents.

StoreAge multiMirror's snapshot-enhanced mirroring occurs as follows:



- The initial zero-capacity snapshot of production data is created (Snapshot 1 above).
- Snapshot 1 begins accumulating a copy of any production data that changes.
- On a user-defined schedule, Snapshot 1 is "frozen" and the next snapshot is automatically created (Snapshot 2 above).

LSI has introduced StoreAge multiMirror™ with snapshot-enhanced mirroring. This technology combines platform-independent, any-to-any, asynchronous mirroring with point-in-time snapshots to ensure data integrity between sites while enabling rapid recovery after a disaster.

Rather than create full-size copies of volumes, LSI creates instant volume snapshots that begin at zero capacity.

Disk space for snapshots does not need to be pre-allocated or reserved, enabling more efficient use of existing capacity.

Snapshot-enhanced mirroring ensures that each site has an identical copy of data as of a specific point-in-time.

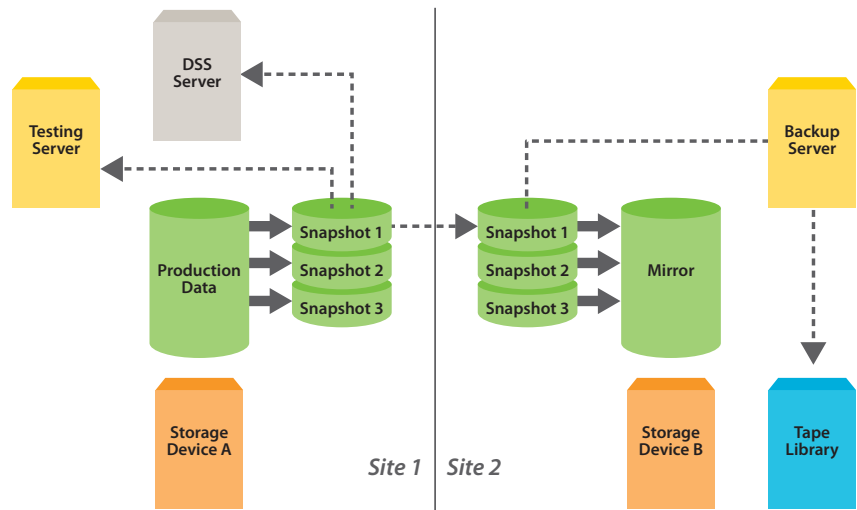
- The contents of Snapshot 1 are mirrored from Site 1 to Site 2, and Snapshot 1 is then retained at both Sites for a user-defined length of time.
- Each site is now assured of having an identical copy of data as of a specific point-in-time.
- The above process is repeated for each subsequent snapshot.

Using Snapshots of Mirrored Volumes for Zero-Downtime Backup, Non-Disruptive Application Testing, Decision Support and Other Critical Tasks

One of the biggest benefits of using snapshot-enhanced mirroring is the ability to utilize the same snapshots for other purposes as well. Since each snapshot is a separate read/write volume and is instantly available for use, these alternate uses of snapshots can include:

- *Zero-downtime backup* – backups may be done “in the background” using snapshot copies rather than production data. Backups can be started anytime, and finish anytime, without impacting normal operations or applications.
- *Application testing* – snapshots can be used in application testing without disrupting production data or applications.
- *Decision support* – snapshots can be used in refreshing data warehouses and other decision support systems, once again without disrupting production data or applications.

One of the biggest benefits of using snapshot-enhanced mirroring is the ability to utilize the same snapshots for other purposes such as zero-downtime backup and non-disruptive application testing.



- Instantly available, read/write snapshots are used to mirror data between sites.
- Any snapshot at any location may also be used for non-mirroring activities, such as zero-downtime backup, non-disruptive application testing, decision support system (DSS) updates and more.
- An unlimited number of snapshots may be retained for future use, and may also be deleted at any time when no longer needed.

Optimizing Performance Over Limited Bandwidth Connections

One of the major costs of any mirroring solution is the ongoing monthly fee for maintaining a communication link between data centers. Generally speaking, the lower the bandwidth of a connection, the lower the monthly cost. Therefore, mirroring solutions that can utilize a low bandwidth connection while still maintaining acceptable performance can significantly lower the total ownership costs.

However, a key issue that may limit the use of low bandwidth connections is something commonly known as “I/O ordering.” This issue occurs when data is transmitted over any type of IP connection, since IP does not guarantee in-order delivery of each I/O.

For example, here is a common “I/O ordering” situation that mirroring solutions must contend with:

1. Two I/O's that change the same block of data occur within a few seconds of each other.
2. The I/O's are mirrored to another location using IP.
3. The remote location receives the I/O's out of order (i.e., the second I/O is received before the first I/O).
4. Unless the remote site re-orders the I/O's, they will be applied “out of sequence” and the remote database will no longer be consistent with the source database.


In this example, the remote mirror is now “out of synch” with the source data, and if the primary site experiences a failure or other disruption, the remote database would be corrupted and not suitable for use.

To ensure that I/O's are correctly applied to remote data sets, and to also minimize the transmission bandwidth required between sites, StoreAge multiMirror™ incorporates an innovative “Last Block Changed” algorithm that examines the transactions contained within each snapshot prior to transmission. If the same block of data within a snapshot has changed multiple times, the Last Block Changed algorithm only transmits the last known change for each block of data. This Last Block Changed technique not only reduces the amount of data that must be mirrored, but also ensures that both sites will have consistent copies of data between sites, even if the data arrives out of sequence.

In addition to eliminating the “I/O ordering” problems encountered in data mirroring, Last Block Changed technology also prevents the following problems typically associated with asynchronous mirroring:

- Since duplicate changes to the same data are discarded, the size of each snapshot is minimized. This reduces the need for large buffers, or the possibility of having a buffer overflow disrupt the mirroring process.
- Latencies resulting from long transmission distances do not affect data integrity.
- During any transmission outages, snapshots continue to accumulate data changes until transmission links are restored, and then synchronize changes with the remote site(s).

An illustration of the Last Block Changed technique is shown below.



- Snapshot 1 accumulates changes to data over a period of time.
- During that time, Block 1 may change several times, while Block 2 and Block 3 may only change once.
- The snapshot is “frozen” and Last Block Changed technology only transmits the last change to each data block. This significantly reduces the amount of data that must be mirrored, minimizing the bandwidth needed between sites and eliminating the problems associated with out of sequence I/O delivery.
- Once the snapshot's contents are received and processed at the remote site, both sites now have copies of data that are consistent with each other at the point in time that the snapshot was originally “frozen.”

Mirroring solutions that can utilize a low bandwidth connection while still maintaining acceptable performance can significantly lower the total ownership costs.

StoreAge multiMirror's Last Block Changed technique not only reduces the amount of data that must be mirrored, but also ensures that both sites will have consistent copies of data, even if the data arrives out of sequence.

StoreAge multiMirror™ also enables application-aware snapshots to be created for each volume or group of volumes associated with an application.

StoreAge multiMirror's vendor-independent design allows it to perform "any-to-any mirroring," where data can be mirrored from any device, to any device, at any location.

StoreAge multiMirror's network-based asynchronous mirroring approach avoids burdening servers and storage devices with data replication tasks. This not only improves overall performance, but also centralizes all mirroring functions for easier management.

Application-Aware Data Consistency

In addition to scheduling the frequency of snapshots, StoreAge multiMirror also enables application-aware snapshots to be created for each volume or group of volumes associated with an application. For example, when creating a snapshot for a database, both the data and log volumes will be included, and the application will be temporarily quiesced to make sure any "in flight" transactions are completed prior to the snapshot creation. Once the snapshot is complete (just a few seconds), the application is returned to normal operation.

Application-aware snapshots can be managed using StoreAge multiMirror CLI's and/or third-party management applications.

Affordable Any-to-Any Mirroring

Because StoreAge multiMirror is a network-based mirroring solution, it is a device-independent layer that resides within the Fibre Channel switched fabric that connects host servers to storage devices. This independence allows StoreAge multiMirror to perform "any-to-any mirroring," where data can be mirrored from any device, to any device, at any location.

Any-to-any mirroring can be significantly less expensive than proprietary mirroring solutions that only mirror data between identical devices from the same vendor. With any-to-any mirroring, you have the freedom to select the most appropriate storage devices for each location without worrying about vendor-imposed restrictions on data movement and replication.

Offloading Mirroring From Servers and Storage

Traditional mirroring solutions often come with hidden performance penalties. For example:

- Server-based mirroring uses the host's CPU to manage data replication, which negatively impacts the performance of any application running on the server. In addition, each server's mirroring process must be managed individually, which increases the burden on storage administrators.
- Storage-based mirroring either burdens the storage controllers with the replication processes, or requires the use of dedicated mirroring controllers, which in turn limits the number of controllers that can be used for day-to-day processing tasks. In addition, data can typically only be mirrored to an identical storage device from the same vendor, limiting the ability to use less expensive devices at remote locations.

On the other hand, the StoreAge multiMirror asynchronous mirroring solution avoids these problems by using network-based appliances to handle data replication processes. Since these appliances work independently of the servers and storage devices in use, several key benefits are realized:

- Data mirroring occurs without involving or impacting the performance of the servers or storage devices.
- Data can be mirrored between storage devices from any vendor (any-to-any mirroring).
- All mirroring processes are centrally managed.

Other Benefits of Snapshot-Enhanced Asynchronous Mirroring

In addition to the advantages discussed above, here are some other reasons why snapshot-enhanced mirroring presents a highly reliable yet cost-effective method of disaster recovery:

- In the event of a temporary communication link disruption, a snapshot is added to the mirroring queue. Until the link is restored, additional snapshots will continue to be created and added to the queue. Once the link is restored, all snapshots in the queue are transmitted to the secondary sites.
- In the event of planned or unplanned downtime, the secondary sites can be quickly brought online using the last received snapshot. This ensures that the secondary site commences operations using a complete and up-to-date copy of the primary site's production data.
- In the event of data corruption, administrators can quickly "rollback" the secondary sites to the last known good point in time. Recovering a system using online snapshots can reduce recovery time to minutes, instead of hours or days.

Summary

LSI StoreAge multiMirror™ snapshot-enabled asynchronous mirroring combines the cost and scalability advantages of asynchronous mirroring with the data integrity and online recovery benefits of point-in-time snapshots.

Compared to other mirroring solutions, the key benefits of implementing StoreAge multiMirror with snapshot-enhanced mirroring include:

- Lower total cost of ownership
- Higher performance using less bandwidth
- Any-to-any mirroring between different storage devices from different vendors
- Higher levels of data integrity in the event of data corruption or "rolling disasters"
- Rapid online recovery to the last known good point in time

StoreAge multiMirror and StoreAge multiView™ are key products in the comprehensive StoreAge™ product suite, a unique set of SAN storage management and data protection solutions.



This document is provided by SGI,
an authorized reseller of LSI Corporation

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

LSI, the LSI logo design, StoreAge, StoreAge multiMirror and StoreAge multiView are trademarks or registered trademarks of LSI Corporation. All other brand and product names may be trademarks of their respective companies.

LSI Corporation reserves the right to make changes to any products and services herein at any time without notice. LSI does not assume any responsibility or liability arising out of the application or use of any product or service described herein, except as expressly agreed to in writing by LSI; nor does the purchase, lease, or use of a product or service from LSI convey a license under any patent rights, copyrights, trademark rights, or any other of the intellectual property rights of LSI or of third parties.

Copyright ©2007 by LSI Corporation. All rights reserved. 0707 10093

