

SERIOUS ABOUT HIGH AVAILABILITY?



Eliminate Storage-related Downtime, Not Just Hardware Failures

Ever since hard disks were deemed critical to data processing, storage suppliers have devoted much effort to circumvent hardware failures. It started with basic disk mirroring and then evolved into the various RAID protection levels in attempts to reduce the cost of redundancy. As external disk subsystems became popular, vendors added redundancy to other components whose failure was considerably more catastrophic; fans, power supplies and disk controllers come to mind.


It is now commonplace to regard storage products as offering “high-availability” simply because they have internally redundant hardware. This narrow interpretation creates the expectation that you can always get to data on disks. This is not the case.

While internally redundant systems may continue to protect against data loss when a single disk fails or a fan goes out, they remain a constant source of planned and unplanned downtime.

Despite what the brochures may say, best practices dictate that even the world’s most sophisticated storage devices still have to be completely taken out of service before their firmware can be updated; even those claiming hot-swap features.

Best practices also call for a complete shutdown when doing major hardware reconfiguration or expansion. These precautions prevent unfortunate human errors from accidentally taking the whole array down, be it a technician tripping over a power cord or shorting the backplane doing routine maintenance. Of course, we have other accidents that render internally redundant arrays useless. Ever had a leaky bathroom fixture drip water on one of your storage devices from the floor above? These and other unexpected mishaps are far more frequent than those often cited...natural disasters. No matter how you look at it, housing storage under a common enclosure makes it a single point of failure and disruption.

In effect, better storage products have shifted the risk from hardware failures to data outages. Protecting against these outages is particularly important in heavily consolidated environments where numerous workloads depend on centralized storage devices. Take, for example, server and desktop virtualization. The costs and difficulties to schedule storage-related downtime are 10 to 30 times more than they were when all these tasks were spread across isolated servers with their own separate disks.




For these reasons, the requirements for highly available storage must encompass the prevention of downtime, not merely shielding against component failures.

Simple High Availability Solution - Double up and spread them apart.

A possible solution to these problems is to mirror disk I/O between separate arrays that are physically isolated from each other. Doing so ensures that users have continuous access to their data while technicians or nature force the other half of the mirrored drives offline. Some customers choose to split their mirrored arrays between different buildings on the same campus. Others with nearby branches will even push them out over a metropolitan fiber connection to the other end of town.

How can we entertain such mirrored HA configurations when budgets are tight?

A possible solution is found in storage virtualization. DataCore Software, as an example, offers a straightforward and affordable software solution for organizations of all sizes. It runs on standard x86 servers, or on a virtual machine in an existing server virtualized by one of the popular hypervisor products from VMware, Microsoft, Citrix, Parallels and Virtual Iron.




Solutions like these enable non-stop data access using commodity-priced storage devices from vendors of your choosing. Each side of the mirror can use different types of storage; they need not be from the same supplier. In fact, some of the solutions can reconfigure the equipment that you already have to eliminate storage downtime.

What are the performance implications when mirrors are stretched?

DataCore as an example employs advanced caching techniques that enable data to be written simultaneously to a perfectly mirrored pair at higher speeds than you would experience with the standalone array. When one side of the pair has to be taken down or suffers an unexpected outage, all the storage I/O is sustained from the second mirror without manual intervention. Nothing needs to be scripted or reconfigured.

With these solutions, preferred and alternate paths to the mirrored pairs can be configured ahead of time, just as they would for an internally redundant array, allowing the operating system or hypervisor multipath drivers to automatically fail over. Most importantly there is no suspense as to whether or not it will work. Essentially true HA is being verified every second of every day.



When an outage does occur on one of the units, whether scheduled or unplanned, the storage virtualization software keeps track of which disk blocks changed while the other half of the mirror was out. The virtualization software uses that information to resynchronize the pair when the hardware is brought back up. Following resynchronization, the client systems go back to their preferred paths.

The result is that you can confidently take half of your storage infrastructure out-of-service for routine maintenance any time of day without suffering downtime. And of course, you can sustain a major storage hardware failure without disturbing users!

Storage Virtualization Configuration

Architecturally, these virtual SANs are simple to configure. The solution provider will either supply hardware based appliances or in the case of software based solutions they'll select two standard x86 servers (physical or virtual machines) sized to your specific needs.

Be careful of the appliance based solutions. Often they are merely x86 servers with the software pre-installed and you may end up paying a premium for hardware that you already have or can purchase at a much lower cost. Additionally you possibly could introduce a foreign server into your environment.

As for appliance based solutions that use non x86 based system, i.e. proprietary hardware, it is fair to ask why they chose this approach. The x86 based solutions offer the same if not better performance, can be easily integrated with 3rd party products and are more rigorously tested in the mass market. In storage virtualization, performance is often a function of well written and efficient software that takes the best advantage of hardware resources like CPUs, memory and I/O ports.

The other challenge with non x86 based appliances is the lack of flexibility in configuration changes. For example if you are a smaller data center, there is no sense paying for an expensive platform that exceeds your performance demands. In fact, you may have surplus power within existing virtual servers to handle the storage virtualization load. In this case you are better off installing a software-based solution on those same machines and let the hypervisor manage the isolation.

At the other end of the spectrum, the performance needs of your data center may demand using the fastest available hardware. If so then it is likely that only the continuous performance improvement offered by standard x86 vendors will be able to keep pace with those demands, allowing you to swap out for faster servers in later years.

A software-based storage virtualization solution matches exactly what your performance needs demand and your budget allows.

Architectural Configuration

Putting storage virtualization software on standard x86 servers turn them into dedicated universal storage controllers to whom you attach current storage systems, either directly or as part of a SAN fabric. To create the ideal HA environment, you place a pair of storage virtualization servers some distance apart and connect them with high speed Fibre Channel or iSCSI/Ethernet connections. Each server will manage half of the storage pool using its collection of back-end disks. These may include internal drives and external arrays connected to the servers using any of the standard disk interfaces, including direct-attached IDE, SCSI, SATA, SAS as well as networked Fibre Channel and iSCSI devices. In essence, the storage servers are both protocol and storage agnostic allowing various types and tiers to be mixed throughout the environment as it makes sense for the business.

Less demanding applications may take advantage of less expensive SATA storage. But keep in mind that while the advanced caching and mirroring techniques will likely improve storage I/O performance across the board, some I/O workloads are bounded by how quickly one can write to disk, so balance the cost savings against the performance needs.

The application servers connect to the storage virtualization servers (and thus, the underlying virtual disks) using an iSCSI or Fibre Channel SAN. They never directly access the physical disks. As data is written to one storage server's cache, it is simultaneously mirrored to the companion storage server to ensure multi-cast stable storage.

Summary

The economic downturn is no time to compromise on data availability, especially when storage virtualization software has made it possible to eliminate costly, storage-related downtime without paying a premium.

If you are serious about high availability, look for storage virtualization software that provides the ability to stretch the two mirrored locations, potentially several kilometers apart yet still present mirrored virtual disks to applications as if they were single, high-performance, multi-ported drives.

About Storage Switzerland

Storage Switzerland, is an analyst firm focused on the virtualization and storage marketplaces. Storage Switzerland provides strategic consulting and analysis to storage users, suppliers and integrators.