

STORAGE SWITZERLAND REPORT

INTEGRATING SSD AND MAINTAINING DISASTER RECOVERY



George Crump, Senior Analyst

As stated in our recent article ["Visualizing SSD Readiness"](#) once the analysis has been performed and an organization has decided to implement Solid State Disk (SSD), thought has to be given to how to integrate SSD into the data center. How you integrate SSD into the data center will depend on what type of SSD system you buy.

While storage system capacity can be easily solved by increasing drive capacities, disk performance has not kept pace. Storage managers are now looking to companies like [Texas Memory Systems](#) to fix their performance-related issues. The historical, and still widely prevalent, method to integrate SSD into the data center is to simply manage it as a stand-alone storage platform. SSD is still highly unique in terms of performance and price and as a result, separate management is justified. Typically the data that is going to be placed on the system is highly visible to the organization as a whole.

An obstacle for companies implementing SSD is deciding how to replicate the data stored on the SSD. For integrated solutions where the SSD is just a hard drive replacement, data replication is straight-forward. Unfortunately, these architectures usually take away the performance and latency benefits that SSD offers. This article presents some options for providing disaster recovery services for stand-alone SSD solutions.

Don't Integrate - Replace

Stand-alone SSD adoption has been on the rise in recent years, thanks the big increases in SSD capacity and stunning price drops. It is now common for entire databases to be stored on SSD. In this case, SSD actually replaces the disk array rather than being integrated into that array. This is an excellent solution for applications that are at the core of an enterprises' profitability and growth.

One of the concerns from people who have bet their organizations' profits on SSD is what happens to the data and organization if the SSD fails. As discussed in a prior article, SSD reliability is higher than ever and in many cases has a higher life expectancy than mechanical drives. Further, these stand-alone SSD solutions can be architected to meet virtually any level of application availability requirements and have been doing so for years at high-profile stock exchanges and telecommunications sites.

As mentioned above, a real concern is how to move this data to a DR site, since it's outside of the protection of the existing array-based replication that many customers use. One option is to use one of the server-based continuous data protection (CDP) products to move that data, as it changes, to locally present mechanical media or asynchronously, to a remote site.

Log Shipping

Companies deploying SSD for application acceleration are typically using one of the enterprise relational database management systems, such as Oracle or Microsoft SQL Server. Enterprise relational database management systems offer log replication tools that enable local and remote DR. Replicating database logs is a time tested and inexpensive method to protect data stored on SSD.

Replication, when done inside the application, can be as much as 6X more efficient than if done by the storage array. It's also significantly less expensive, since no additional software needs to be purchased.

Preferred Reads

For organizations addicted to their array-based replication and whose performance problem is slow reads, a preferred read architecture offers a great way to accelerate read performance and maintain existing array-based replication strategies.

Preferred Reads are implemented in the file system, but some virtualization appliances have this ability as well. In a preferred read architecture, a volume is mirrored between two volumes on the mechanical disk array and a volume on the SSD. In a normal mirrored environment, reads come from both sides of the mirror in a round-robin fashion. For mechanical drives this is fine, but if the application can take advantage of SSD then the preferred-read file system or virtualization appliance can be set to have reads always come from one side of the mirror. In this case it would be the SSD. This preferred-read capability is available from companies like Symantec and in virtualization appliances like those from DataCore. Oracle's 11g database with Automatic Storage Management (ASM) also offers built-in preferred read capabilities.

Other operating systems like AIX and HP-UX do not allow you to specify which particular device reads can come from. Instead, they monitor the queue depth of the mirrored

volumes, unlike the traditional round-robin method. In this case, whichever volume has the lowest queue depth will be the one that fulfills the read request. As discussed in a prior Storage Switzerland article "[Visualizing SSD Readiness](#)" SSD technology virtually eliminates queue depth and, as a result, would always be servicing the read requests.

The advantage of using a mirror with preferred-read capabilities is that SSD can be added to the storage infrastructure with little to no affect on the storage operations, including disaster recovery policies. It simply accelerates the read performance. It also works effectively well for both structured databases and unstructured file systems.

The other advantage of a preferred-read architecture is that the configuration allows organizations to provide high availability for solid-state disk drives without having to purchase two solid-state units. As a result, preferred read technology not only improves integration of SSD into the broader storage environment, it also can drive down costs.

Virtualized Integration

The final option is virtualized integration. These are typically appliances or switches that lift the storage management function off of the array storage system and place it onto an appliance. The appliance then becomes the single point of storage management for any type of storage that lies underneath, including SSD. The unique performance characteristics this produces require that the appliance does not itself become the performance bottleneck.

The advantage of this type of integration is that all tiers of storage, from SSD to SATA, can now be managed from a single platform and interface. There are virtualized integration offerings for both NAS for file systems from BlueArc and NetApp as well as block I/O or SAN solutions from companies like DataCore. Texas Memory Systems recently announced the acquisition of the technology assets of Incipient which will allow them to offer this integration themselves, a first for an SSD vendor.

About Storage Switzerland

Storage Switzerland, is an analyst firm focused on the virtualization and storage marketplaces. For more information please visit our web site: <http://www.storage-switzerland.com>.