Solid State Storage and Virtualization

Multiply the benefits of virtualization with solid state storage.

By: Khaled Amer, Chair of SSS TWG Performance and Benchmarking Committee, and Neal Ekker, Vice Chair SSSI Education Committee and VP Texas Memory Systems.

Introduction

Simply stated, solid state storage (SSS) is a data storage technology made from silicon chips instead of spinning metal platters or streaming tape. This technology, though available for decades, is finding new ways of providing cost effective solutions in a wide range of commercial and governmental information systems often referred to collectively as enterprise applications.

Because SSS has no mechanical moving parts, it handles random access patterns and sequential access patterns very well. This is in sharp contrast to mechanical spinning disks where the head actuators perform acrobatics in their struggle to keep up with random access patterns. Hence, the contrast in performance between SSS and traditional hard disk drives is more pronounced (by orders of magnitude) as the access patterns are more random.

Because of its outstanding performance characteristics, solid state storage can be used to leverage or multiply the capabilities and benefits of hardware and software in storage environments where it's deployed. For example, SSS can both help enable and leverage or multiply the desired results in server virtualization deployments. This is why it is important to begin fleshing out your SSS strategy as you move forward with your virtualization plans.

The technologies complement each other; SSS will both enable and increase the benefits you can realize from virtualization.

What is Virtualization?

In order to understand how solid state storage can multiply the benefits of server virtualization, it's important to clearly define what virtualization means. Virtualization is the creation of a virtual (rather than actual) version of something, such as an operating system, a server, a storage device, or network resources. Server virtualization is the masking of server resources (including the number and identity of individual physical servers, processors, and operating systems) from server users by creating virtual servers or virtual machines (VM). The intention is to spare the user from having to understand and manage complicated details of server resources while increasing resource sharing and utilization. Server virtualization allows a piece of hardware to run multiple operating system images at the same time. The usual goal of virtualization is to centralize administrative tasks while improving scalability and workloads. In the past several years, virtualization software has seen surprisingly fast adoption rates into both server side and storage side enterprise environments.

Solid State Storage and Virtualization

A major reason to virtualize your servers is to enable greater efficiencies and thus do more with fewer physical servers. As virtualization becomes main stream, higher performance applications are being moved onto VMs. Applications suffering from storage-related performance issues before moving onto VMs will still suffer from them after the move, perhaps even more so. With virtualization and server consolidation, a bigger number of data streams are being written to and read from

the consolidated storage. These data streams appear more and more like small random-access applications. Solid state storage does not have mechanical components that must move in sequence from one task to the next, so it normally performs much better that hard disks in random access environments. Because of this, applications with random access patterns benefit the most from SSS, thus virtualization and solid state storage make a natural and powerful fit.

An analogy would be the concept of having many stores scattered around vs. consolidating them in one place (like a mall) where the overhead is shared and clients can find all their needs there in a convenient consolidated place. This will tend to drive more traffic to the stores and provide a one-stop shopping place for clients. Sounds like a good idea. But imagine doing so while having very narrow aisles in that mall. It becomes a big problem for people to squeeze in and out of these aisles. They will queue up in lines waiting for a slot in one of the aisles. One solution is to dramatically increase the number of these narrow aisles to meet the demand (similar to deploying more hard drives in a storage system). While partially solving the problem, it's an expensive, inefficient solution. A better solution would be to develop a smaller number of wide aisles (similar to deploying SSS in a storage system) to allow all the traffic to go through smoothly without bottlenecks and enable the clients and stores to reap the benefits of the mall concept.

SSS and HDD Work Virtually Well Together

In order to implement virtualization solutions, you will need to centralize your data storage. Functionality such as virtual machine relocation, for example, requires sharable storage in order to support the movement of VMs from one physical host to another within the virtualized server resource pool or farm.

Centralizing your storage offers an excellent opportunity to add SSS to the mix. Even though the centralized storage can be all SSS, many storage experts now advocate a combination of solid state storage to address performance requirements and high density inexpensive hard disk drives to address capacity requirements. This provides the best of both worlds in one solution – high efficiency and low cost.

Notice how many solutions open up when a strategic mix of SSS and hard disk drives (HDD) storage is matched with a virtualized host environment. First, even though SSS can cost more than HDDs in terms of $/capacity, it costs substantially less in terms of $/performance, especially in random I/O operations and applications. As mentioned above, the storage access patterns in virtualized environments become more random as VM numbers rise and as high performance applications are virtualized. Moving the hot portions of a dataset (which can be as little as 2-5%) onto SSS can have a dramatic effect on performance in a more cost effective fashion than using HDDs alone.

Next, it's likely that the multiple VMs you set up will host applications with a range of I/O needs and response time (latency) sensitivity. The SSS/HDD storage mixture offers a similar range of performance capabilities, without resorting to inefficient and complex disk configurations that require adding many more unnecessary HDDs just to meet the performance requirements. If server consolidation in a cost-effective manner while meeting application performance requirements is a major goal, then this storage architecture enables it, most likely to a greater extent than virtualization alone would.

Finally, reliability is a major factor in mission critical enterprise environments, and believe it or not, solid state storage actually increases the reliability of your storage solution.

- SSS has no mechanical parts, which means no mechanical wear and tear
- Most reputable SSS products have robust ECC (error correcting codes) that go well beyond what is offered in typical HDDs
- When SSS fails, as all components eventually do, it fails "gracefully" with predictability and no catastrophic crashes.

The following are some examples of how SSS enables and multiplies the value of server virtualization.

Example 1: SSS & Virtualized OLTP Applications

Once you start down the path toward server virtualization, you'll want to eventually incorporate all your applications, but without the performance and bandwidth of solid state storage some common types of applications (that are I/O

The story is a simple one: as you move toward host-side virtualization to consolidate servers, increase efficiency, and save money, you can multiply the benefits by including solid state storage in your network storage plan.

intensive) will create challenges. For example, imagine when you decide to move your online transaction processing (OLTP) application onto a VM. OLTP applications are notorious performance hogs. Not only will they drag down your CPU pool with millions of small block transactions, but they will do the same with your newly centralized storage pool.

Because of the costs of provisioning so many extra disks devoted solely to performance, all-HDD storage solutions rarely have extra performance capabilities to share when query traffic suddenly spikes. SSS is the opposite; most often normal deployments don't come close to saturating its bandwidth or performance capabilities, especially for random I/Os.

Example 2: SSS & Batch Processing Applications

The same often holds true with that mission critical but enormous batch processing application resident for decades on the mainframe in the back room. There's so much to be gained by moving it to a virtualized processing environment. But when that batch job kicks in around midnight on top of other running applications, you'll be answering system alarms. That's unless you've included an appropriate SSS configuration on the storage side which will enable your CPUs and centralized storage to easily handle the performance requirements of adding the batch processing application to the load.

Conclusions

The story is a simple one: as you move toward host-side virtualization to consolidate servers, increase efficiency, and save money, you can multiply the benefits by including solid state storage in your network storage plan.

About the SSSI

The SNIA SSS Initiative has been created to foster the growth and success of the market for solid state storage in both commercial and consumer environments. The SSSI is responsible for coordinating and managing all activities within the SNIA related to the use of storage made from solid state devices, including those for education, technical development, business development, marketing, implementation and conformance testing. For additional information or to get involved visit the SNIA SSSI web site at www.snia.org/sssi.

About the Authors:

Khaled Amer currently serves as the Chairman of the SNIA SSS TWG Performance and Benchmarking Committee. Besides his participation within SNIA, he has been a recognized leader in many standards organizations such as the IEEE 802 (where he chaired several committees), JEDEC 64.8 standard committee for SSDs, the IETF (Internet Engineering Task Force) and others. Khaled also has extensive expertise in the analysis and optimization of architectures through queuing analysis & modeling and helped optimize many of the networking standards. During 9 years at IBM, Khaled was instrumental in the development and optimization of several LAN products. He also worked at Rockwell Semiconductor Systems as a Technology Planning Manager for the Division General Manager. He founded a consulting company which focused on standards and architecture optimization for about 10 years. His last job was the Director of Architecture and Standards at Fusion-io (a leader in enterprise SSDs). Khaled holds a M.S. in Electrical Engineering from the California Institute of Technology (Caltech).

Neal Ekker currently serves as the Vice Chair of the SNIA Solid State Storage (SSSI) Education Subcommittee. Recently he has co-authored white papers on various topics for SNIA SSSI. He works for Texas Memory Systems, an industry leader in enterprise solid state storage, as VP of Marketing. He has also held business development and operations management positions for software development and IT services companies with worldwide enterprises. Earlier, Neal worked in a wide range of fields, from uranium and gold mining operations to freelance journalism and technical writing, as well as helping create and administer a college campus on the Arizona Strip north of Grand Canyon. He holds BA and MA degrees from Stanford and Brown Universities.