

STORAGE SWITZERLAND

THE STORAGE HYPERVISOR



George Crump, Senior Analyst

Flexibility is a key objective for many data centers. They want to respond to the needs of the business without having to slow the business down. Flexibility like this is a primary deliverable of server virtualization and one of the reasons server virtualization projects continue to be a top priority for IT.

The flexibility of server virtualization comes from abstracting the application from the physical server hardware. If storage assets could be made as flexible as server virtualization has made server hardware, data centers could finally control storage growth and costs to make it an economical component of the data center.

The Storage Virtualization Past

Abstracting data from the physical storage systems has certainly been tried before, in the form of storage virtualization. It provided the ability to utilize any manufacturer's storage systems and then centralize those systems under a single storage services umbrella provided by a storage virtualization software application. Examples of this technology would be applications like DataCore and FalconStor.

The challenge that legacy storage virtualization applications face and one of the reasons for their very low adoption rate is that they re-create the same monolithic storage model that traditional storage has used. Instead of the storage software running on hardware inside of the storage system it's run on an external set of servers. While this gives it the ability to run on multiple vendors' hardware it doesn't really provide any new value.

While most of these software based storage virtualization solutions provided capabilities like thin provisioning, improved snapshots and flexible replication, most of the legacy storage systems were able to provide these features as well. Flexibility in mixing vendors, while interesting, was not in and of itself enough of a value for customers to justify taking a risk on a new way to deploy storage.

Additionally, the storage virtualization customer now had to provide at least two servers (for availability) on which to deploy the storage virtualization software. This solution took time to purchase and install the servers, to connect them to the data center infrastructure and then finally to load the software. A turnkey storage system already had that work done. As a result the software based virtualization path became too much of a "do it yourself project" for the typical data center.

The key limitation was the monolithic nature of these systems, where the servers running the storage virtualization software became the bottleneck. It was the same dual controller approach that the legacy vendors took, except that the controllers were now in an external package. Even though some of these companies have adopted some sort of scale-out model, there is overhead associated with that approach, and again this still means having to purchase, deploy and configure more servers.

What is a Storage Hypervisor?

A storage hypervisor is the concept of letting the server virtualization hypervisor run an increasingly larger portion of the storage services that are typically found in the array. The hypervisor already abstracts the physical CPU and network connections from the application, and to some extent, the physical storage as well. Why not let the hypervisor go the rest of the way and manage the storage completely?

There is already evidence of these capabilities today. Hypervisors can transparently move virtual machine disk images to different storage platforms while the VM is still running its application. For example, a VM that needs more performance can have its disk image migrated to an all-solid-state device to take advantage of this zero latency memory based storage. Or a less actively used VM can be migrated to higher capacity but lower cost storage, maybe even a system with power saving, spin down drives.

Running more of the storage services in the hypervisor also simplifies storage management for the administrators. Disk systems are now just bulk storage areas that the administrator can manipulate from within the virtual environment. Storage management is just another tab on the virtual management GUI, not a whole new interface to learn and keep track of. Compared to legacy storage virtualization, there are no new servers to deploy, the storage capabilities are in the hypervisor and it borrows resources from each host to accomplish its storage management duties.

This means that disk can be bought as a stand alone system to address the current need. The IT staff doesn't need to worry about how scalable it is. They just keep adding more storage systems and let the hypervisor be responsible for making them into a cohesive and easy to manage group.

Soon the hypervisor will be able to manage the storage environment in an automated fashion, similar to how it automatically load balances VMs today, from a processor perspective. This would allow the hypervisor to move storage performance demanding VMs to storage systems that have the most available IOPS and less demanding storage I/O workloads to more cost effective disk systems.

There is also a business reason for the hypervisor vendors to take on more and more of the storage services role. These vendors see storage as a key stumbling block in greater server virtualization adoption because the storage system becomes one of the most expensive and complex components of any virtualized implementation. Expect hypervisors to continue to add more and more storage capabilities with each successive version.

The biggest benefit that a storage hypervisor would be able to provide is an almost perfect scaling model. The addition of each host would mean another hypervisor to help out with storage services tasks. This scaling model may be more ideal than even using scale out storage, which needs to be expanded and managed independently from scaling the virtual host environment. The storage hypervisor would scale automatically with the host environment. Also, as stated earlier, this scaling comes without the need to purchase additional servers just to handle storage load.

The Shortcomings of a Storage Hypervisor

Hypervisors are not without their weaknesses, and today one of those is in providing advanced storage service features like snapshots, thin provisioning, cloning and replication. Both VMware and Microsoft Hyper-V see a significant decrease in performance when snapshots of virtual machines are active. Snapshots are the key foundation of other storage features like cloning and replication.

Thin provisioning is also a key feature of the modern storage system and is critically important in virtualized server environments to help curtail the cost of storage. Hypervisor managed thin provisioning volumes, like snapshots, perform significantly worse than thick or fully allocated volumes.

The result of poor snapshot performance and thin provisioning is that most administrators abandon having the storage hypervisor perform any but the most basic of these storage services tasks.

They are stuck using fully allocated volumes or moving to an enterprise storage system that has the features and the storage processing performance to handle their resource demands. In both workarounds the storage consumption of budget resources grows dramatically as does complexity.

Addressing the Shortcomings of the Storage Hypervisor

The solution, instead of abandoning the use of the hypervisor for storage services, may be to go ahead and leverage those services but to fill in the weak areas with third party software solutions. An excellent example of this is software products like [Virsto](#) that run as a filter driver or as a virtual machine. They can improve the performance of dynamically allocated volumes (thin provisioning) and snapshots while providing unlimited use of those features without performance loss.

In this scenario Virsto uses a structured log similar to what an enterprise database uses. Through the use of this log, Virsto sequentializes the very random, write-intensive I/O stream coming out of the virtual host, allowing whatever storage is used in it to operate at its streaming performance instead of its random performance. This generates a significant performance boost – from the point of view of the virtual machines on the host – from 35% to 1000%, depending on what storage is in the log and how it is configured. These writes are then asynchronously de-staged to a thinly provisioned “primary” storage area. Write acknowledgements flow back to the VMs when the write hits the log, not when they are written through to primary storage. This means that, from the point of view of the VMs, the write performance of the entire storage system is actually the performance sustained by the log, which is significantly faster than any device performing random writes.

An interesting side-effect of this architecture is that desirable storage operations such as thin provisioning and snapshot/ clones can be used without incurring the performance degradation normally associated with these operations. That’s because these operations all occur “behind” the log so that their performance – from the point of view the host – is the performance delivered by the log.

The behind the log side-effect allows thin provisioning and snapshots to be utilized in a nearly unlimited fashion which then means the resource utilization can now be risen to a new level. Thin provisioned volumes are important because they release captive disk areas which can comprise as much as 30% or more of a hard allocated volume. Snapshots can be used for multiple data protection copies to improve recovery granularity, and for clones to allow for the creation of a master VM which can then be cloned or sub-cloned for specific tasks. Given the likely similarities between VMs, cloning can lead to a 50% or greater increase in storage efficiency.

The net result is that thanks to the fact that the host is only dealing with the log file snapshots and thin provisioning can be used to their full extent. This leads to a potential 80% reduction in required storage needs while maintaining similar performance to hard allocated volumes and better data protection.

The combination of improved hypervisor capabilities for managing independent storage systems, plus the addition of software like Virsto’s, that extends the hypervisor’s limited capabilities in delivering services, give the administrator a powerful storage option, one in which they can select the most cost and performance correct storage system to solve today’s storage price / performance problem.

The Storage Hypervisor saves the organization from having to buy more performance or capacity than they need. When the time comes where a different performance option or more capacity is needed they simply add another storage system in pod-like fashion, then leverage the hypervisor to move virtual machines into that pod and eventually the hypervisor will manage those moves automatically.

Finally the Storage Hypervisor approach may change the economics of storage in virtualized environments. The days of relying on high-end arrays or legacy storage virtualization to deliver advanced storage management features with total costs typically in the range of \$20 - \$25 per GB, the storage hypervisor opens up an exciting future where lower-tier storage can deliver high performance features for \$6-8 per GB.

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>

Copyright © 2011 Storage Switzerland, Inc. - All rights reserved