Mesabi Group

# Commentary

March 28. 2005

# Storage Virtualization 101: The "Cliff Notes"

*Block-level storage virtualization in the short term holds real promise to simplify management and facilitate non-disruptive data mobility. In the longer term, it will participate in a virtual infrastructure across servers, networks and storage. But the question of where the virtualization intelligence should reside has generated a lot of "religious" fervor. Buyers should avoid dogmatism and choose the alternative that best fits their business. This is a strategic decision; architecture and scalability are important as is heterogeneity and the ability to leverage existing functionality — such as array based remote replication.*

## Block-level Storage Virtualization — Time to Stop Talking and Start Doing

For many IT organizations, the question about block-level storage virtualization is no longer *if* they should do it, but rather *when* is the right time to select and implement virtualization, *where* the virtualization functionality should reside, *whether it will* augment existing solutions or require users to get rid of them, *whether or not it* will provide immediate and important benefits, and *how* to select from a number of options.

Within a few months, users will face a number of additional virtualization choices and a steady stream of marketing hype surrounding each. This paper will help users narrow the choices. It will examine the strengths and weaknesses of the major options. By the end of the paper, the strategic functional and architectural decisions should be far clearer.

## Home Sweet Virtualization Homes

Block-storage virtualization can occur in several locations (Figure 1): at #1 the host level, in the storage network (a.k.a. fabric) as #2 an appliance or using #3 the coming generation of switches, or #4 at the storage controller level for a controller that can manage multiple arrays.

## #1 Host-based Virtualization

The use of volume managers for host-based virtualization is a tried and true approach with which many IT organizations have high familiarity and positive experiences.

However, volume manager software has to reside on each host, and the only storage that can be virtualized is the storage that applications on that host can see. Effective storage-infrastructure-encompassing virtualization requires a


26 Country Lane
Westwood, Massachusetts 02090
Telephone: 781 326 0038
Fax: 781 326 0038
www.valleyviewventures.com

single view of a logical pool of storage that spans all the servers whose applications want to virtualize storage.

**Figure 1: Simplified View-of Where Storage Virtualization Intelligence Can Reside**

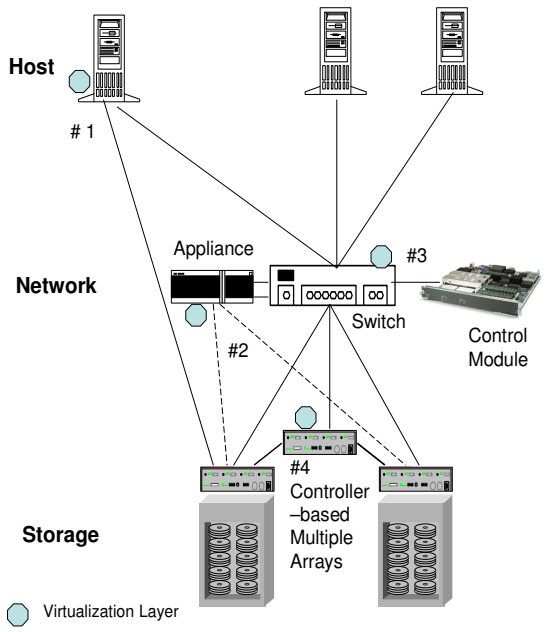


Source: Mesabi Group March 2005

The inability of host-based volume managers to do that effectively is the reason that IT organizations that are struggling with storage infrastructure complexity are seeking other alternatives.

## Network-based Storage Virtualization

SAN block-level storage virtualization is attracting a lot of attention. From the point where data traveling as I/Os leave their hosts until the point where they reach their physical storage destinations, they are in the network, and can be virtualized across the entire network. The two general network-based choices are:

- *Appliance-based* — I/Os flow through a switch either before or after the appliance, but the appliance controls the virtualization layer.
- *Switch-based* — the actual virtualization layer occurs in a switch, although control information to manage those I/Os can and in many respects should be processed in a storage control module "outside" the switch.

## # 2 Appliance-based Network Storage Virtualization

To the appliance approach goes the honor of first providing storage virtualization in the network. An appliance is a self-contained bundle of software and hardware dedicated to a single purpose — in this case virtualization. While this approach has proven to work well and has been well-tested, the impacts of the use of cache and the inherent limitations to scaling have to be considered.

*Caching Out.* The insertion of any new device in a network that has to examine data packets introduces a new node in the network through which all data packets must pass — creating "latency", or added time. In order to minimize the impact of any latency of going through the appliance, an appliance will likely introduce cache. The use of cache permits acknowledgements of writes back to host servers prior to the data actually having been written to physical storage. That solves the latency problem as a host can now issue a new write, but it introduces new issues.

Now the appliance, and likely several appliances, has the responsibility for managing cache. For protection purposes the cache has to be mirrored. Now data is in even more places. The appliance(s) now have to manage cache consistency


26 Country Lane
Westwood, Massachusetts 02090

Telephone: 781 326 0038
Fax: 781 326 0038
www.valleyviewventures.com

(cache and physical storage have to be synchronized) and cache coherency (the multiple caches have to be synchronized). Issues related to these cache issues are not trivial. Even though caching is a well-established technique, the burden and cost of adding additional layers of cache and the management thereof, especially among multiple appliances, is a potentially serious drawback to this approach.

*Scaling Out.* As the I/O rivers become larger to meet increased application demands, the appliance has to scale out in bandwidth to accommodate those demands, which means to add more processors to its cluster of general purpose processors. The issue with adding more processors is one of overhead and coordination that is typical of any clustering or distributed system environment. Understanding how this might impact application-related performance in your environment now or later should be a matter of consideration.

### # 3 Switch-based Network Virtualization

Each of the major SAN switch vendors has a strategy to enable applications such as virtualization to run on their switch hardware. The basis of this capability is specialized processing ASICs (Application Specific Integrated Circuits) at the switch ports and accompanying control modules where smart software can reside. The ASICs perform the translation of logical to physical addresses and the redirection of I/O that is the essence of storage virtualization. The control module runs the intelligent software that controls this virtualization process on the switch.

*The Choice is Yours.* Storage applications interface with control modules through an API (Application Programming Interface). Even though each vendor's control module and API is proprietary today, there is an emerging standard, FAIS (Fabric Application Interface Standard). IT organizations will have a choice of intelligent devices as long as storage vendors write their applications to the FAIS standard and the intelligent switches and devices support this standard.

*Cache Not Accepted.* Since the vast bulk of I/O traffic can flow directly through the very-high-speed ASIC (exceptions are handled by the control module), the latency problem can be made imperceptible to applications without the need to resort to caching.

*Tough Crowd.* The reality of actually writing the difficult software to create switch-based virtualization has been talked about for years. But, users should expect this to change significantly by midyear. The switch vendors and key storage vendors are moving to make it a reality. The show-me crowd will have the opportunity to see this approach in action and test drive as they consider virtualization options and timeframes.

By midyear, we expect to see scalable network based virtualization solutions that among other things enable non-disruptive data migration and mobility across the entire virtualized infrastructure. And, we expect the software to recognize and fully enable existing robust array based remote replication. In short, we expect these solutions to augment the existing infrastructure. If, as we expect, this comes to fruition mid year, you will have a well-designed solution that provides growth


26 Country Lane
Westwood, Massachusetts 02090

Telephone: 781 326 0038
Fax: 781 326 0038
www.valleyviewventures.com

investment protection and avoids retooling in the future.

### #4 Controller-based Virtualization across Disk Arrays

Controller-based virtualization is not necessarily limited to one array. If a storage controller can control the virtualization process across other "slave" arrays that can be from the same or different storage vendors, then it is a reasonable alternative.

*Rip and replace required?* This approach requires users to relinquish the robust capabilities and processes of their existing heterogeneous array-based solutions, including, but not limited to, replication. Getting rid of your hard fought investment may not be a realistic option.

*Economics.* Although the master controller does not necessarily have to have disk attached, having disks attached helps justify the cost. Yet commingling the storage virtualization decision with the storage purchase decision makes it difficult to ascertain the true cost of each.

*Potential downtime*. Controller-based virtualization has the potential to manage a very large pool of storage. This enterprise-class capability comes at a price. Imagine having to upgrade the controller or migrate to a more powerful controller (and that day will surely come). At worst all the slave storage arrays could suffer unacceptable downtime, and at best storage administrators would have to do considerable dancing to prevent any major inconvenience to the application user communities.

*Disaster protection*. Moreover, the controller represents a single point of failure from a disaster recovery perspective. If a SAN suffers from a catastrophe — attack or natural disaster — which could knock out the data center and controller that distant arrays depended upon. Disaster recovery planning teams should consider the implications very closely.

### Conclusion

Making a decision about block-level storage virtualization does not require knowledge of esoteric topics, rather it gets down to basics. Are you comfortable with the status quo? If not, this is a good time to consider virtualization. You will be faced with choices and a substantial dose of marketing hype around each one. We suggest that you look at this as a strategic decision and carefully consider the pros and cons outlined in this short paper. But whatever your choice, the one choice that will be unacceptable is not making a choice at all.

**David Hill**

Analyst Name: David Hill
Topic Area: Storage Management

Mesabi Group LLC
26 Country Lane
Westwood, MA 02090
www.valleyviewventures.com

*Mesabi Group LLC is an affiliate of Valley View Ventures that aims to provide thought leadership and sound advice to both vendors and users of information technology. .*

Phone: (781) 326-0038
email the author: davidhill@mesabigroup.com