

# Storage Virtualization-- Extracting Value from SAN-based Storage Systems

by Kent Winton,  
Director of Storage Software Products Division Channel Development  
IBM

The virtualization of storage is not a new concept or application. Many "self-contained storage systems," including EMC, IBM, HDS, have virtualization techniques embodied in their front-end processors. What is different in today's marketplace is the availability of "storage virtualization" that spans Open connectivity standard storage devices (individual branded storage systems) and provides SAN-wide functionality. Combining these two attributes allows for the provisioning of the SAN-attached storage assets to be managed, allocated, and used in an "economically optimized" fashion. Storage virtualization in a SAN environment can put in place the capability to optimize usage in a cost-effective sense from the storage assets deployed. What we will explore in this short article are some of the ways economic value is created by the virtualized SAN storage environment. The examples are derived from an IBM implementation called the SAN Volume Controller (SVC). Other implementations may vary in capability.

## SAN storage virtualization (the IBM SVC) provides the ability to:

1. Optimize the utilization of the SAN-based storage asset(s) across multiple systems and brands.
2. Provide an Open Connectivity standard based platform to implement SAN-wide functions such as "instant copy" and "remote copy"
3. Build a SAN management system that allows "total" control of the storage environment and minimizes down time (planned and unplanned)
4. Enhance the performance of previous generation storage systems and create longer asset useful lives
5. Simplify (and reduce the absolute COST in most cases) SAN storage administration

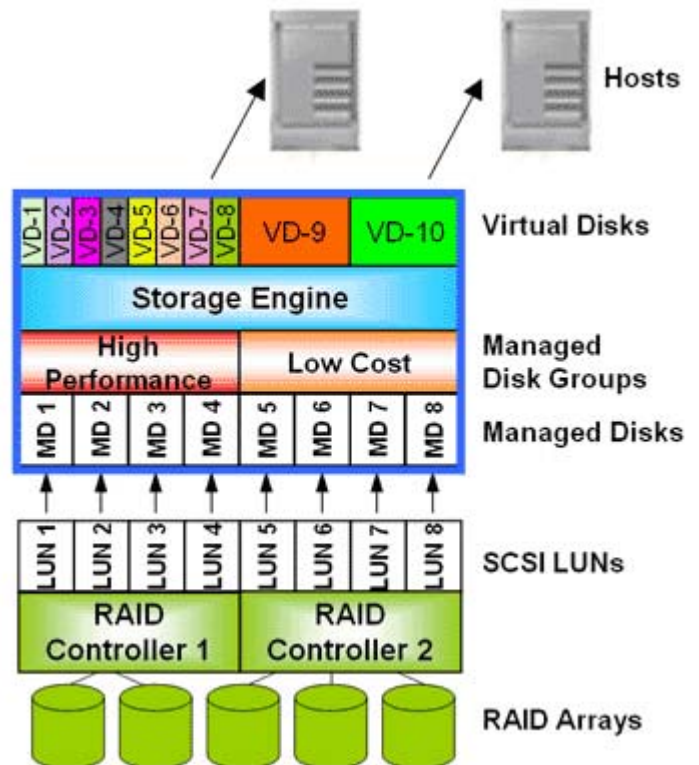
Exploring each of the above attributes of moving to a managed (through virtualization plus software functionality) SAN environment allows us to define how users can utilize and economically justify such a capability.

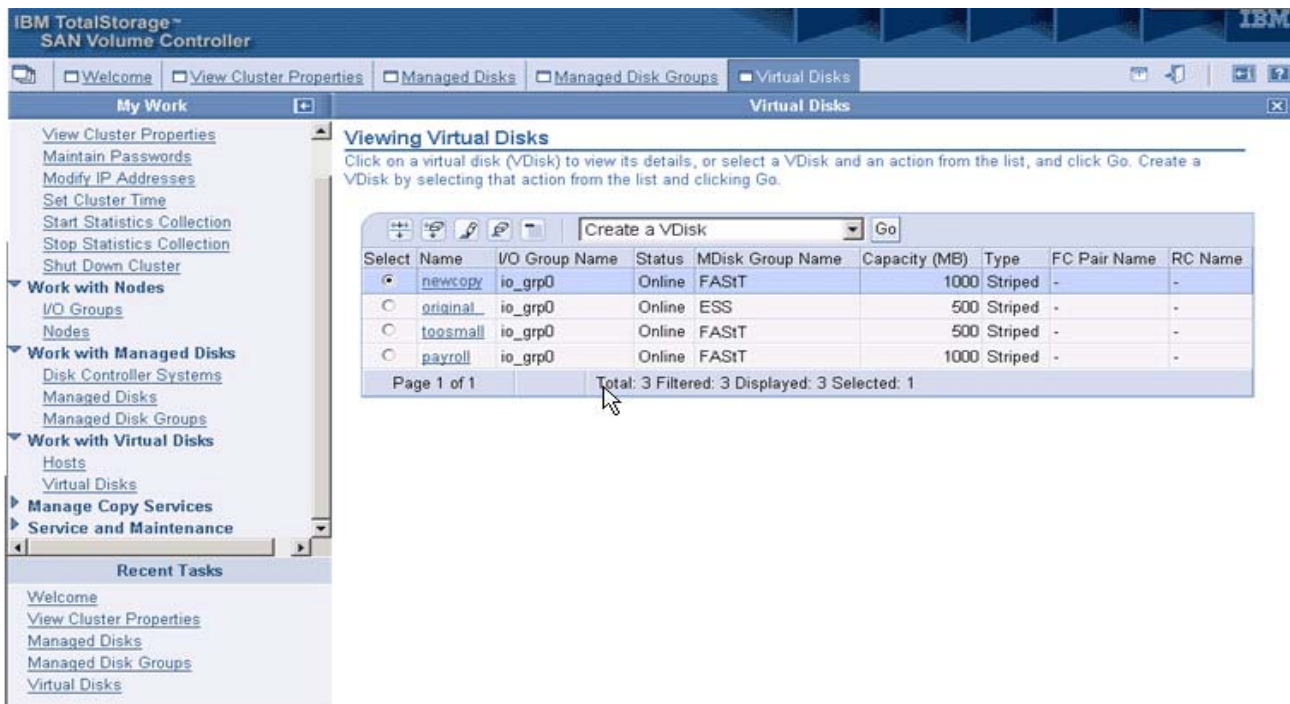
### Optimize Storage Utilization

When most SAN storage pools are assessed (using tools such as IBM Tivoli Storage Resource Manager), with regard to storage allocated to ACTUAL storage consumed (capacity), very interesting results are obtained. IN MOST CASES, LESS THAN 50% OF THE ACTUAL CAPACITY INSTALLED IS USED. In fact, many cases show capacity utilization is less than 30%. Recapturing more use of the storage asset INSTALLED is a major benefit of SAN Storage Management through virtualization. Storage is allocated, reallocated, moved ... under GUI administered control. No longer does an installation need to provide "excess storage in case of immediate need" but can analyze and program allocation based on actual usage. Many examples exist where much of the 50% plus of unused yet allocated storage, when managed, is recaptured and actually used. A very relevant case recaptured over 3TBs of such storage in a single 6TB capacity SAN; more storage capacity for users created without any new storage capacity asset deployed. Storage capacity though good storage management is the key.

### Build SAN-wide functions over different types of storage systems

A very important consideration in selecting storage systems today is their functional richness. Functions such as "instant copy", "remote copy", and "dynamic data migration" are no longer considered optional but are integral to the storage selection process. SAN management through virtualization techniques (specifically the IBM SVC) allows such functions to be implemented on a SAN-wide basis as opposed to being unique and special to each disk system. Early adopters of the SAN-wide functional implementations have found economic benefit in deploying their high performance/high resiliency (relatively higher cost) storage systems for mission-critical (e.g. transaction processing) applications and using lower cost storage devices as copy targets. And, the ability to implement "any to any" copy allows different brands of storage to be consolidated and data moved to the most appropriate devices. Finally, SAN management of the storage devices attached enables the migration of data from one storage device (older systems, devices to be retired, etc.) DYNAMICALLY without server downtime or inability to access data.





### Achieve total control of the SAN storage environment

The ability to integrate and manage different storage devices (even from different vendors) through a single set of commands via a "user friendly" GUI is truly a breakthrough in SAN storage deployments. One storage administrator looking at a "virtual pool" of storage and deciding what to do with which storage when can be of enormous economic and asset usage benefit. In most cases, SAN management will impact storage downtime (both planned and unplanned) in a very positive fashion. One only needs to consider movement of data dynamically from devices requiring upgrades to other devices without server interruption as a case in point. Or, consider the automatic SAN configuration of additional storage WITHOUT server or SAN outages. The example of scaling additional storage without reconfiguring servers or taking any devices DOWN can have tremendous IT operating benefits; in some cases, we have seen the storage pool grow from a few TBs to over 30 without (completely WITHOUT) any system outages.

### Enhance storage system performance

A sometimes hidden attribute of a very high performance SAN management engine that, among other functions, provides virtualization (the IBM SVC) is that caching in the SAN can have dramatic impact on existing storage systems. Systems without much cache, with lower performance attributes, MAY BE TURBO CHARGED by the SAN engine. So, not only can the older systems be used in less demanding storage applications (e.g. copy targets), but also their performance may be enhanced to the point that their service life is extended. SAN management through virtualization makes the best use of the assets employed (we have seen some 2X performance boosts in specific cache friendly applications).

### Dramatically extend storage administration capabilities

Imagine ONE GUI with ONE administrator looking at ALL the SAN-attached storage and managing it as a homogeneous resource. Preliminary indications are that with SAN management through virtualization, an individual storage administrator's productivity can be extended dramatically. We have seen several examples of SAN pools of over 100 TBs with several different storage devices that are now very effectively managed by ONE person. Implementing the copy functions SAN-wide reduces not only the cost of the copy services but extends the administrator productivity. More examples exist but one that we see over and over is the ability to deploy new storage "stuff" without new training, without new procedures, WITHOUT DOWNTIME. And, one person through a management GUI is pulling these triggers.

There are many more reasons SAN-deployed storage must have Management through virtualization. In our TSSC facility in Minneapolis, Datatrend can show you how to optimally plan, test (in our facility), implement, and support the new world of storage management in the SAN using virtualization methodologies (The IBM SVC is probably the industry leader; ask us about it).