Raidtec CORPORATION
WHITE PAPERS

A New Method of SAN Storage Virtualization

Table of Contents

1 - ABSTRACT
2 - THE NEED FOR STORAGE VIRTUALIZATION
3 - EXISTING STORAGE VIRTUALIZATION METHODS
4 - A NEW METHOD OF VIRTUALIZATION: Storage Virtualization Manager
5 - CONCLUSION

1 - Abstract

The fast growth in storage capacity and processing power in enterprise installations coupled with the need for high availability, requires a Storage Area Network (SAN) architecture to provide seamless addition of storage and performance elements without downtime. Virtualization provides consolidation and uniform management of storage resources for SANs with multiple servers and multiple storage sub-systems. The industry goal is for storage virtualization that is inexpensive, does not affect performance, and that is flexible enough to be applied to various types of systems and networks without requiring significant investments in time or resources. Currently, storage virtualization is implemented at three architectural levels: (a) at the storage sub-system, (b) at the host and (c) in the SAN fabric as a central management unit. Each of the methods provides specific advantages but is limited in its capabilities. This paper provides information on the available storage methods and on a newly developed virtualization method that provides the storage management benefits offered by existing techniques without the corresponding drawbacks.

2 - The Need For Storage Virtualization

The term "Storage Virtualization" is used to indicate the operation of dividing the available storage space into “Virtual Volumes" without regard to the physical layout or topology of the actual storage elements (e.g. Disk Drives, RAID sub-systems etc.). Typically, Virtual Volumes are presented to Operating Systems as an abstraction of physical disks and are used by these OS as if they were such disk drives.

Storage Area Networks impose the need for managing large amounts of storage, in a uniform way and from a central location. The fast growth in storage capacity and processing power in many enterprise installations, coupled with the need for high availability and 7x24 operation, requires the SAN architecture to enable seamless addition of storage without downtime. These goals are best achieved via virtualization of the storage. The main difference between a Virtual Volume and what is generally termed a Logical Unit (LUN), is that a Virtual Volume can be created, expanded, deleted, moved and selectively presented, independent of the storage sub-systems on which it resides. A Virtual Volume may consist of storage space in different storage sub-systems, with different characteristics (e.g. number and size of LUNs) and can be expanded with any available storage. The importance of this capability can be exemplified in handling peak activities of e-commerce web sites. Using a Volume Manager capable of adding/removing servers on the fly as well as instantly moving Volumes from server to server is the best solution to manage the large variance between normal and peak operating loads.

The independence of Virtual Volumes from the physical storage is essential to enable the migration and replacement of obsolete storage sub-systems without performing shutdown or causing data access interruptions. In addition, Virtualization of storage enables the use of software applications instead of human Administrators to perform Volume allocations. Data management is significantly enhanced through Storage Applications such as Snapshot, Remote Mirroring, Virtual Tapes, etc. that require the creation and expansion of Virtual Volumes for their use. In summary, a SAN cannot be useful without storage virtualization.

Storage Virtualization should provide the following functions without affecting inherent performance:

- ✓ **Universality** - the creation of a unified view of the virtual storage regardless of the physical elements or topology
- ✓ **Single Point of Management** - as dictated by the enterprise policy Server, Operating System and File System independence - to enable the construction of the most cost effective SAN
- ✓ **Storage sub-system independence** - to enable unbiased selection of best of breed system
- ✓ **Heterogeneous SAN support** - to enable smooth growth in changing market conditions
- ✓ **Flexible allocation of storage space to servers** - to effectively meet users' needs
- ✓ **Support the overall SAN criteria** - e.g. High Availability, Scalability, Security etc.

3 - Existing Storage Virtualization Methods

Virtualization at the Host Level

One method of virtualization is via storage management software that runs at the server level. The main advantage of this method is that it enables multiple storage subsystems to work in parallel with multiple servers. A key difficulty with this method is that it assumes a prior partitioning of the entire SAN resources (disks or LUNs) to the various servers. Virtualization is only performed on pre-assigned storage, losing a key advantage of SANs as well as the independence of Volumes from Servers. Typically, entire LUNs are assigned to specific servers, thus limiting the amount of servers that can use the same storage subsystem. Further, a Virtual Volume created over the storage space of, say, two LUNs will not be easily moved to another server, especially if there are other volumes created over the same LUNs. Virtualization at the host level typically also requires augmenting the management function with some parallel mechanism of zoning and LUN masking and relying on LAN connectivity for synchronization between servers, which may affect the reliability of the entire SAN.

Virtualization at the Storage Sub-system Level

This method was first implemented in mainframe environments in the 90's and is one of the most common storage virtualization solutions in use today. In this method, a uniform Storage Virtualization Manager is achieved by creating Virtual Volumes over the storage space of the specific storage subsystem. Pooling all SAN storage resources and managing Virtual Volumes across several storage subsystems requires that this method be augmented by other means which in general will only be practical in homogeneous SANs using a single type of RAID subsystem. Creating virtual volumes at the storage system level provides host independence, but with limited flexibility for growth. Raidtec has provided Storage Virtualization at the sub-system level for several years using its FC-Access volume management software, which is available with Raidtec's FC-Raidman and SANfinity storage management products. This continues to be a cost effective solution for entry-level SANs and even No Single Point of Failure clustering configurations where a single storage array can service all the hosts on the SAN.

Symmetric Virtualization

Virtualization in a separate hardware box that is placed between the servers and the storage solves many of the difficulties of the above two approaches but at a very high price. The advantage of this method is that the Volumes are completely independent of the both the servers and the storage subsystems on the SAN. The management software “sees" all the physical storage available and can create virtual volumes and allocate them as required. The key drawback of this architecture is that it achieves this independence at a very high price i.e. every I/O of every server is sent through this central unit causing significant performance degradation and a SAN bottleneck. Attempts to design high performance units result in expensive systems at the low end while more cost conscious solutions limit the upward scalability. This method may thus lead to poor performance, expensive solutions, or both. In addition, meeting some key SAN criteria such as High Availability may render very complex designs involving high performance clusters and similar (expensive) solutions.

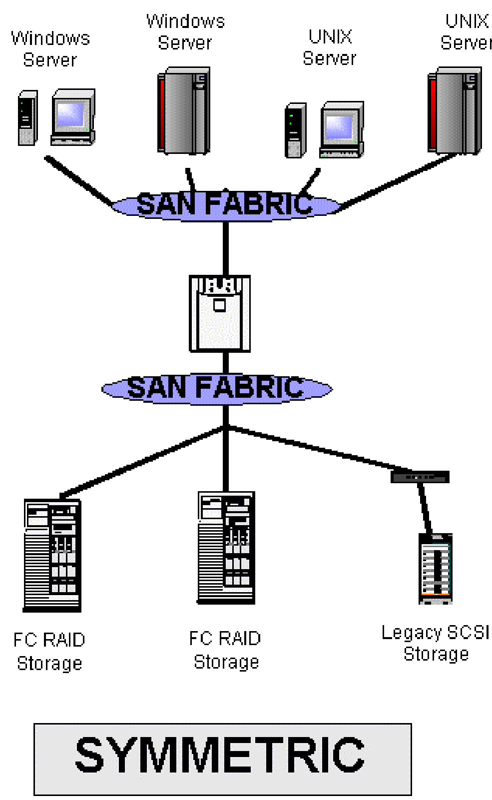

4 – Asymmetric Virtualization: Raidtec SAN Virtualization Appliance (SVA)

This method uses a combination of a Metadata Center and Volume Drivers for creating and managing virtual volumes while enabling direct data transfer between server and storage subsystems. This method ensures negligible impact on SAN throughput, even for very large SANs. Indeed, by allowing multiple storage subsystems to work in parallel with multiple servers, total SAN performance can be considerably enhanced.

The Raidtec SAN Virtualization Appliance (SVA) separates the handling of Metadata from the data path. It consists of a Metadata Center that "sees" the physical storage and allocates virtual volumes, and a Volume Driver at each of the servers on the SAN. The Volume Driver retrieves the volume configuration from the Metadata Center and presents virtual volumes to the operating system as if they were disk drives. When the operating system sends an I/O to the virtual volume, the Volume Driver intercepts the I/O, translates the volume address to the physical address, and sends the I/Os directly to the storage subsystems. The Volume Driver resides very close to the hardware level and can thus perform this data mapping with very low latency and negligible overhead on the host server.

The SVA architecture maintains the flexibility of virtualization in a central unit without incurring the degradation of performance or high cost of hardware associated with "in the data path" solutions. The Metadata Center can be a small and inexpensive unit, because it does not have to handle actual data transfer.

In the SVA based solution all storage is handled uniformly, independent of the type of storage subsystem or operating system. Virtual Volumes can be created, expanded, deleted, moved from server to server and assigned to multiple servers. This method also enhances the reliability of the entire system by avoiding parallel zoning and LUN masking management, because filtering is done at the Volume Driver according to mapping retrieved from the Metadata Center and not according to physical port locations or LUNs. Another important aspect of the SVA method is that not only human administrators (via the management GUI) but also Storage Applications can request volume allocations from the Metadata Center, giving the SAN new capabilities for improved data management. Such Applications as Snapshot and others can create, expand and delete their own volumes, communicating directly with the Metadata Center.

Key SAN criteria are easily met with the SVA concept. High Availability (HA) SANs are implemented through simple redundancy of the Metadata Center in an active/passive configuration. Very high degrees of scalability can be achieved through loosely connected storage domain architectures and similar designs, benefiting from the key SVA advantage i.e. that there are never pending I/Os stored in the management unit nor is there is a need for cache coherency or other complex mechanisms which limit the practical expansion of real life enterprise SANs.

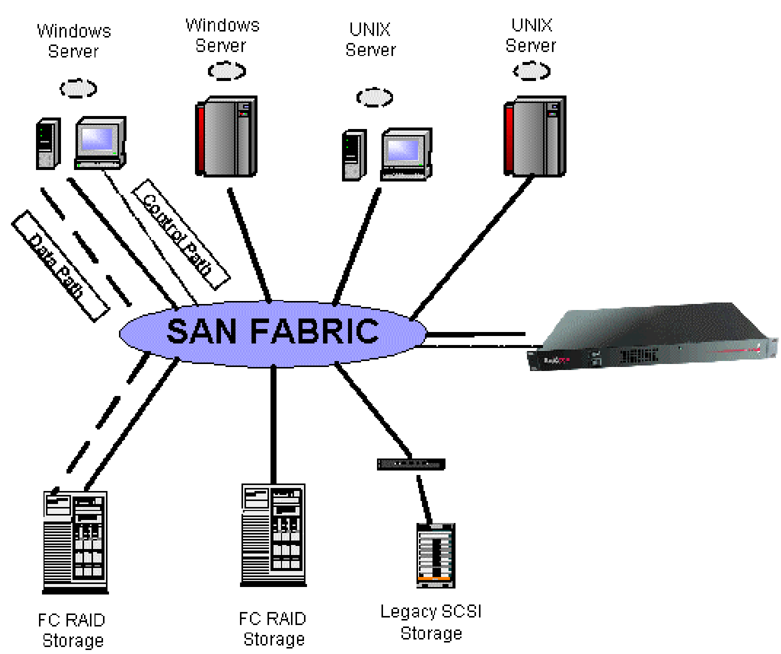

ASYMMETRIC

5 - Conclusion

Raidtec's SVA provides a central management point for storage virtualization while the data is directly transferred between the host computer and the storage devices. This architecture meets all key requirements from a good virtualization scheme and in particular the total independence of the virtual storage from both the servers and the storage subsystems, thus creating true heterogeneous SANs. This method provides significant scalability, performance, reliability and price advantages over alternative architectures.

Table 1 - Comparison of Various Virtualization Methods

Method ▸ / Criteria ▾	Server	Storage Subsystem	SAN Central Unit	SAN Appliance
Universal	✗	✗	✓	✓
Single Point of Management	✗	✓	✓	✓
Server, OS & FS independence	✗	✓	✓	✓
Storage subsystem independence	✓	✗	✓	✓
Flexibility	✗	✗	✓	✓
Performance	✓	✓	✗	✓
SAN Scalability	Poor	Poor	Poor	High
SAN HA Configurations	✓	✓	Complex	✓
Relative Management Cost	Medium	Medium	High	Low