



# Introduction to Disk and File Virtualization

By Paul Newton

**Disk and file virtualization, built upon storage area network (SAN) technology, can help you cheaply and effectively manage future storage capacity in the enterprise. This article introduces the technology and benefits of disk and file virtualization.**

## INTRODUCTION

Data storage efficiencies realized by z/OS® data facility storage managed subsystem (DFSMS/MVS®) are understood by the legacy mainframe management and storage managers. Since the introduction of legacy mainframe DFSMS/MVS®, the cost of storage per byte dropped significantly, and administration costs associated with the growth of legacy mainframe data was effectively contained as a result. The opportunities realized by the client/server era resulted in expanding growth of de-centralized enterprise data. Enterprise storage strategies must address cost associated with this business appetite for increased storage capacity with the wide variety of technology available to satisfy the appetite.

The objective of this article is to introduce disk and file virtualization technology embedded in the recent total cost of ownership (TCO) sensitive IBM® product announcements which exploit this concept. Those with a DFSMS/MVS® background can draw parallels and imagine future possibilities for this technology based upon history of DFSMS/MVS®. The accomplishments of DFSMS/MVS® for management, control, and TCO of legacy mainframe data storage growth can be accomplished for the remaining enterprise data storage by using disk and file virtualization.

## OVERVIEW

Disk and file virtualization is built upon storage area network (SAN) technology. SAN technology provides external expansion of data required by client/server business applications. While SAN permits ease of data growth, islands of host operating systems with associated data storage is the result. Each island requires its own administration. As technologic business enterprise opportunities advance, island sprawl must be addressed.

Disk and file virtualization addresses island sprawl.

IBM® SAN Volume Controller (2145) provides disk virtualization, and Storage Tank SAN File System (4146) provides file virtualization. Disk and file virtualization can be implemented independent of the other. However, combining the technologies offers the best positioning to eliminate island sprawl. IBM® SAN Integration Server (2146) is a prepacked SAN solution which includes SAN Volume Controller with SAN switches and storage devices.

The virtualization of disk storage involves pooling of SAN controller back-end disk storage into available storage extents. Extent sizes range from 16MB to 512MB, which is used to determine the maximum amount of data maintained by an individual SAN Volume Controller cluster, 64TB to 2PB. This allows for centralized control of storage pool expansion and ease of moving files to newly installed SAN controllers from older controllers scheduled to be de-installed.

The benefits of disk virtualization include transparent addition and migration of physical disks. Another benefit is the reallocation of spare capacity on underlying physical disks without impact on servers, irrespective of the server operating system or platform type. In addition, advance functions can be done at a single point in the SAN in a common way regardless of the underlying physical storage. FlashCopy®, peer-to-peer data copy and data migration can also be performed in a common way. This common platform will be used to provide other advanced functions over time like advanced security and quality of service capabilities.

The physical back-end disk storage enclosed in one or many SAN controllers is dynamically detected by SAN Volume Controller and labeled as unmanaged disks. Subsets of these individual unmanaged disks can be assigned to a managed disk group. The managed disk group represents a large pool of available storage extents. Virtual disks, vdisks, are the defined usable extents from the pool. Subsequently, these vdisks are mapped to one or many host systems previously detected and defined within the SAN fabric. An example of the mechanics to accomplish managed disk group assignment,

creation of vdisks and mapping vdisks to SAN fabric defined hosts can viewed at URL <http://www.storage.ibm.com/software/virtualization/svc/snapshot.html>. SAN Volume Controller currently supported hardware and software levels can be viewed at [www.ibm.com/storage/support/2145](http://www.ibm.com/storage/support/2145).

File virtualization makes the same data available to more than one host operating system including unlike operating systems within the SAN fabric. Additionally, creation of new files from any of these operating systems can be directed to designated back-end storage according to coded policies. The DFSMS/MVS® ACS routine used legacy mainframe z/OS® is a valid analogy. The SQL-like syntax is interpretive with a top-down execution structure and can be dynamically changed.

Both disk and file virtualization environments include a command line interface and a browser-based WebSphere application GUI for administration and control. The SAN Volume Controller and SAN File System consoles were both created from a common navigation and content presentation standard.

The virtualization technology resides within xSeries™ equipment, which is fibre channel attached to a SAN fabric switch and Ethernet-connected on a private TCP/IP LAN consisting of separate master console workstations. The disk and file virtualization software is a loosely coupled cluster of xSeries™ nodes running a specialized Linux kernel each with 4GB memory for processing and activity caching. Ancillary hardware is included with IBM® disk and file virtualization products to provide reliability, availability and serviceability.

Implementation of Storage Management Initiative Specifications, SMI-S, open standards from Storage Network Association Industry, [www.snia.org](http://www.snia.org), along with increased autonomic characteristics, is among the planned objectives for reaching desired maturity of a fully integrated approach to data and storage management. Open standards, automation and centralized control applied to disk and file virtualization will result in reduced administration costs by combining flexibility with operational economies of scale.

An enterprise with matured disk and file virtualization technologies would be in an advantageous position to exploit future technologies that will undoubtedly appear. For example, current research is underway for potential products in heat-assisted magnetic recording (HAMR), a significant advancement in total gigabits per square inch. Details of this research can be located at [www.insic.org/hamr.htm](http://www.insic.org/hamr.htm)

## SUMMARY

---

In summary, disk virtualization is one of many SAN controllers aggregated into a single pool of storage, divided into named managed disk groups, then available for mapping to one or more hosts. File virtualization is an aggregation of one or many file systems shared among one or many hosts. Total Cost of Ownership (TCO) sensitivity is a major reason to consider introducing disk and file virtualization into an environment that is expecting data capacity growth. The facilities available with disk and file virtualization will position the enterprise to exploit design point economies of scale as data capacity grows. Review SMI-S from SNIA, then consider a first step of introducing disk virtualization technology into your enterprise to cheaply and effectively manage future capacity. 

---

*NaSPA member Paul Newton is an IBM employee at the Developer Technical Support Center in Dallas, Texas.*

