CommVault Systems®

Storage Virtualization: Seeing the Forest, Not the Trees

CommVault Systems
Corporate Headquarters
2 Crescent Place
Oceanport, New Jersey 07757-0900 USA
Telephone: 732-870-4000

For more information about CommVault Systems' enterprise data management solutions, please phone us or visit our web site at http://www.commvault.com

CommVault Systems White Paper

Storage Virtualization: Seeing the Forest, Not the Trees

Can't see the forest for the trees? In network terms, the trees are individual devices while the forest is a unified view of network elements. In its simplest form, storage virtualization technology reads data location information from storage devices, infrastructure services, and applications, and abstracts the resulting locations into logical data views for end-users. This allows end-users to safely ignore network complexities in favor of data-centric views. Storage virtualization also benefits IT administrators, who can more easily provision storage not strictly by device but according to logical virtual structures centered on applications. Enterprise virtualization architectures go even further, presenting enterprise-wide logical data views and easing administrative burdens.

And when administrators combine storage virtualization with policy-based data movement, the resulting virtual environments enable customizable and automated business continuance tasks that are keyed to the company's core business needs. In this scenario, virtualization software tools correlate device data and application data to create total virtual environments.

This scenario is not just for the large enterprises, but would benefit any company struggling to unify their server and storage environments. Even a mid-sized company might have five general-purpose servers, ten RAID devices, one tape library, and three NAS filers. A larger company might add servers galore, along with the contents of a data center and storage area networks. The end result is a horrendous combination of different locations, devices and volumes. Meanwhile, users aren't interested in remembering where the files are; they just want to find the information they need. And storage administrators are cracking under the load of managing myriad servers and storage devices.

Using storage virtualization, companies could create virtual volumes by many different schemas. It could be by department (Marketing is "Marketing" and Accounting is, well, "Accounting.") It could be by application (Exchange databases are on a volume called Mailbox), or by client (the company's biggest client is on a virtual volume called BreadNButter). In fact, just one of these virtual volumes may span three tape drives in the tape library, three disks out of the six-disk RAID array, and 3 gigabytes of storage on a network server.

And if the storage administrators can accomplish this with improved utilization and automated techniques, that would be a tremendous benefit as they struggle to corral their runaway storage resources. Some of their key storage management challenges include:

Over-Provisioning. Under-utilizing storage space is a scandal, since storage utilization rates average a pitiful 30% of disk space in NT environments and 40% disk space in UNIX. The result is storage over-provisioning as administrators try to stay ahead of the storage curve. However, the "throw on another box" mentality is an expensive proposition in terms of return on investment and human resources, as administrators try to manage scores of separate storage devices with no end in sight. Virtualization can help, since managed storage can achieve a much higher utilization rate.

Management Loads. Storage administrators often manually manage far more data than they can effectively handle. The result is a crisis mentality where it is impossible to manage storage environments. To relieve the tremendous burden, virtualization techniques should help to standardize and centralize storage management.

Common Virtualization Terms

In-band: Virtualization engines that sit between the hosts and their storage devices in the data path. Sometimes called "symmetric virtualization," they manage virtualization by receiving the hosts' I/O commands and redirecting the data to the assigned storage device. The engines range from switch-based appliances to devices such as routers.

Out-of-band: Dedicated virtualization devices that sit outside the data path. They maintain virtual volumes by separating metadata from data and processing it separately with location information to the correct storage device.

Host-based virtualization: Virtualization technique implemented in a host (server) instead of in a storage subsystem or appliance. Each host controls its pre-assigned storage.

Tape virtualization: Virtualizing tape drives or tape libraries.

Storage subsystem: An integrated storage system containing controllers and/or host bus adapters (HBA)s, storage devices such as disks, CD-ROMs, tape drives or libraries, and control software.

Storage virtualization: A series of techniques that can make physical storage locations appear as streamlined logical views. They can also shield administrators from complex operations and present unified management views. Virtualization can occur in layers.

Virtual volumes: Logical units of stored data, where actual physical storage locations are mapped to non-physical (logical) volumes. Virtual volumes are often created to support data-centric network views.

The Challenge

For all its benefits, comprehensive storage virtualization is hard to accomplish. Some storage virtualization techniques may be relatively new, but virtualization itself is not a new concept. Like many good mainframe technologies that have been painfully adapted for open systems, storage virtualization dates back to the mid-70s in both disk and tape. However, deploying virtualization techniques in the open systems world has been a huge challenge. And although most storage virtualization discussions center on the SAN, it is not limited to the storage area network.

But current development centers on storage area networks (SANs), our present-day notion of storage nirvana. The concept makes sense in that environment, which composites different storage subsystems and infrastructures in order to maintain centralized storage capacity. Unfortunately, this unified approach is part of the problem-vendors and customers cling to the idea that storage virtualization should be able to take all of these disk arrays and make them into one big, amorphous storage blob just waiting for slicing and dicing. Unfortunately, SANs are a good deal more complicated than that with their point-to-point switching architectures, endless storage device combinations and serious management issues. Combine these challenges with vendors slapping the storage virtualization label onto any old software, and the result is customers who won't even discuss the "V-word."

SAN virtualization techniques may be based on hosts, storage subsystems, in-band systems, out-of-band redirectors, or storage domain managers. Following is a breakdown of today's SAN virtualization technologies: host-based, storage system-based, out-of-band and in-band.

Host-Based Virtualization

Host-based storage virtualization in Fibre Channel SANs depends on a combination of zoning and LUN masking. When the hosts communicate with physical storage devices using SCSI protocol over Fibre Channel, the destination address is a combination of a target ID and a Logical Unit Number (LUN).

Some hosts grab any visible LUNs on Fibre Channel cables, so multiple hosts on the same LUN may compete for the same storage device. In order to work around these problems, administrators will commonly zone Fibre Channel switches in order to block other hosts from trying to access the same storage device. They also practice LUN masking, which presents a filtered view of available storage systems customized to each application server. Thus Server A may see two RAID subsystems on a particular LUN, but not the JBOD and tape library the administrator has assigned to another LUN, which only Server B can see.

Zoning and LUN masking is tedious and time-consuming work for the administrator and denies servers' respective storage to each other. A server-based storage virtualization application will only work on the server's pre-assigned storage, losing the SAN's ability to integrate multiple storage devices. It also limits volume independence from servers and makes it impossible to virtualize all storage devices within the storage area network. Administrators must also install virtualization drivers on each host, which makes it difficult to practice change control throughout different versions and platforms.

Storage Subsystem-Based Virtualization

Multi-host storage subsystems effectively practice virtualization by presenting multiple internal devices as one large volume. Large-scale disk arrays have long used virtualization technology to segregate a monolithic array into individual partitions. Administrators can then parcel out the partitions to different host systems and create virtual disk volumes from numerous physical drives. In the tape world, virtualization allowed the tape library to act as a shared storage pool where administrators could aggregate several tape devices to appear as a single resource to host systems. Tape systems have since added disk drives to act as the backup cache. The virtual volumes are not tied to the actual capacity of the internal devices, though the subsystem's monolithic enclosures limit the size and makeup of the resulting storage pool.

Out-of-Band Redirectors

Sometimes called "asymmetric," this approach splits the data path and control path, with a separate server (out-of-band) processing the resulting metadata (control data) in order to assign data to virtual volumes. The advantage of this approach is that the volumes remain independent of the servers and the specific storage subsystems in a SAN. The management software discovers all available storage, creates virtual volumes and allocates data. This all sounds quite wonderful, except that out-of-band virtualization, by forcing metadata to travel through the virtualization unit, severely degrades SAN performance. This necessitates buying more expensive equipment to compensate for the performance hits, and defeats the very reason that many people want to virtualize storage-to lower their equipment costs with improved utilization rates. Because of these performance hits, redirectors are often used primarily for offloading disk I/O from the LAN to the faster SAN, and only secondarily for virtualization.

In-Band Virtualization

Sometimes called "symmetric," a dedicated virtualization engine sits directly in the data path (in-band) between the servers and the storage devices. The engine may take several forms: it may be a simple appliance that requires additional switches and storage devices, a self-contained appliance with embedded switches, a disk-based platform similar to multi-host arrays, or even a switch or router that comes with virtualization software and administers the virtual volumes. The engine administrates the virtual volume and assigns the data, so storage administrators to not need to install drivers on separate hosts. Since SAN performance may suffer from processing all data through the in-band device, many in-band virtualization products provide sophisticated caching techniques and alternate pathing support.

Storage Domain-Based Virtualization

A storage domain server (SDS) is an application server platform that handles virtualization and storage allocation duties in a storage area network. Sometimes called network-based virtualization, the virtualization function is contained in a storage control layer running on the platform's native operating system. The network approach allows the SDS to leverage the operating system's networking, volume management, device interoperability and security features. Some products allow administrators to link multiple SDSs in order to distribute storage pool load and management while maintaining centralized administration.

Looking Ahead

Storage virtualization is breaking out of the SAN ghetto into the enterprise. In these virtualized network environments, enterprise data management companies, like Commvault Systems, can link physical, logical and application layers to present logical, virtualized network views. Users will not have to remember data locations across arcane physical views, but will be able retrieve their data logically and intuitively. Storage administrators will be able to provision application storage space much more effectively and easily, significantly improving utilization rates across the enterprise. Both IT administrators and end-users will spend less time finding and managing data, and more time using it. And in combination with policy-based data movement, enterprise virtualization architectures can integrate and automate vital business continuance tasks such as backup/recovery, replication, snapshot and data migration.

Maybe we can start saying the V-word again.