Managing the information that drives the enterprise

STORAGE

Storage Virtualization Workbook

Virtualizing storage offers configuration flexibility, better disk utilization and more. Here's how to get started.

Inside

TechTarget
The IT Media ROI Experts

Storage virtualization doesn't have to be hard

By Rich Castagna

THE CAPACITY OF YOUR STORAGE systems is efficiently allocated and used, data migration occurs seamlessly and when you do have to add more storage it simply becomes part of a pool of available capacity. Right? If you answered "Yes" to all three of these statements, you're in the midst of a good daydream, flat-out kidding yourself or one of a small but growing group of storage managers who have discovered the advantages of storage virtualization.

The idea of treating all installed storage—no matter whose label is attached to the box—as one big system has been around for a while. The benefits of virtualization are apparent, so I won't dwell on them here. But for many storage managers, there's been kind of an approach-avoidance conflict surrounding virtualization. The benefits are tangible, but isn't it hard to do?

There's no single, well-defined route to virtualization—there are at least three of them—and that adds to the confusion. But, in most cases, one's existing storage environment will all but dictate which form of storage virtualization will work best.

Storage virtualization has also earned a reputation as being hard to implement and manage. No marketing hype from me—storage virtualization isn't exactly plug 'n play. But its deployment difficulties are often overblown, approaching urban legend status. It's not exactly "Let's virtualize all of this storage and go grab some lunch," but it's certainly

something any savvy storage manager can handle.

Getting tied into a single vendor is another common concern about virtualization. But, again, it's likely your storage environment will be the key to overcoming this issue. Many shops have already gone the one-vendor route, while others divvy up their vendor preferences by storage tier. If you've made that sort of commitment, it's not much of a stretch to stick with that vendor to make better use of the storage it supplies. A back-out strategy is smart planning, but if you start small within controlled environments, you can lessen the risk considerably and give yourself some breathing room to iron out the wrinkles.

To dispel some of the apprehensions surrounding storage virtualization, we've put together this Storage Virtualization Workbook. It'll help you get a handle on the products and processes involved in a virtualization project. Drop us a line and let us know how your virtualization efforts fared. ⊙

Rich Castagna (rcastagna@storagemagazine.com) is Editorial Director of the Storage Media Group.

Purchasing tips

Key points to consider when buying a storage virtualization product.

By Stephen J. Bigelow

THE STORAGE DEMANDS of users and applications are spiraling out of control. Keeping pace with these demands usually requires adding disks and storage subsystems. But increasing storage resources entails a management penalty: Storage administrators must remember which applications are tied to what storage, allocate sufficient storage to address future growth and manually track performance. Over time, the inefficiency of this manual process leads to wasted storage space.

Storage virtualization alleviates these problems by implementing a layer of abstraction between applications and physical storage, allowing storage to be combined and treated as a ubiquitous resource, regardless of location. While storage can be virtualized in several ways, the ultimate goal is to improve storage utilization (reducing or forestalling capital expenditures) and storage management. Aside from the obvious issues of pricing and support, here are the key points involved in purchasing storage virtualization products.

1 **What does storage virtualization bring to the enterprise?** Storage virtualization offers many potential benefits, but which one(s) does your organization need? Improved storage utilization, capital cost savings and ease of management are three common ones, but there are

others that appeal to storage administrators. For example, virtualization can ease disaster recovery (DR) or business continuance planning by allowing the use of heterogeneous hardware between sites. It can handle data migration among storage platforms, tiers and sites. Virtualization also eases storage capacity expansion, often automating the manual tasks needed to allocate storage to applications.

2 **How much complexity will be added?** Many virtualization products require additional hardware or software in the infrastructure. Beyond new servers or switches, determine what host device drivers, path managers and agents you'll need to support the prospective virtualization product. Maintenance is usually first to feel the pinch here; IT staff can easily get bogged down in patching and updating a bunch of storage virtualization servers when hardware is replaced or new versions become available. Not paying enough attention to maintenance can result in version disparity, leading to instability and performance problems. Evaluate any storage virtualization product from a management and maintenance perspective, and determine if the problems it solves are outweighed by the new issues it introduces.

3 **Where does storage virtualization best fit in your environment?** In terms of implementation, storage virtualization can be host-based, array-based or fabric-based. Host-based virtualization relies on software (installed on host servers) that monitors data traffic and storage. Symantec Corp.'s Veritas Storage Foundation is a host-based virtualization product. Dedicated appliances, such as IBM Corp.'s SAN Volume Controller (SVC), are similar. Array-based virtualization integrates the technology in the storage array itself, such as a TagmaStore array from Hitachi Data Systems. Growing attention is being paid to fabric-based virtualization, such as EMC Corp.'s Invista and Incipient Inc.'s Network Storage Platform, which are software products that run on intelligent switch devices. Each approach offers advantages and disadvantages in terms of performance, scalability and cost.

4 **Scalability.** A virtualization product can only manage a finite amount of storage. You should understand the tradeoff between scale and performance. This is particularly treacherous because many virtualization initiatives begin as test or pilot deployments before being deployed throughout the enterprise. Scaling issues may not appear until later in the deployment cycle. Evaluating scaling up front can help weed out unacceptable products and deployment surprises.

5 **Interoperability.** The promise of cross-vendor storage utilization has been a compelling feature of virtualization technology, but true heterogeneous support is still lacking in virtualization environments. Array-based virtualization typically locks in the array vendor. Host- and fabric-based virtualization products also impose a certain amount of vendor lock-in with the software or appliance that embeds the software. You should determine the compatibility of any virtualization products with your current environment, as well as likely upgrades or updates.

6 **Test and start small.** Most analysts recommend a thorough lab evaluation of any storage virtualization product prior to any purchase commitment. This should include a review of decommissioning drills. Once a purchase decision is made, you should start implementation on a small scale before building out the virtualization systematically. This conservative approach gives administrators ample time to get accustomed to virtualization management, and prevents unforeseen problems from creating wide-scale problems or even crippling an entire data center.

7 **Know how to undo the implementation.** Storage virtualization isn't perfect. Performance issues, scalability limitations and interoperability problems are just a few reasons to decommission a virtualization product. An organization may also decide to discontinue one product in favor of a more appropriate one. Unfortunately, backing out of virtualization is disruptive to applications and confusing to administrators. Before committing to a virtualization product, discuss any back-out options with the vendor. ⊙

Stephen J. Bigelow is a features writer at TechTarget.

COUNT ON YOUR DATA TO BE THERE

BECAUSE YOUR iSCSI SAN IS SUPPOSE TO BE RELIABLE

Continuous access to your data and applications should be a given. The Dell EqualLogic™ PS Series is designed to deliver enterprise-class, total system reliability.

YOUR CHOICE IS SIMPLE

With a fully-redundant system architecture, automated load balancing, and routine self-monitoring and self-management, the PS Series of intelligent, all-inclusive iSCSI SAN solutions offers online serviceability to help provide uninterrupted access to your data. And, should disaster strike, you can count on the PS Series auto-snapshot and auto-replication capabilities to enable rapid restore. Find out what reliable storage from Dell can mean for your business.

EQUALLOGIC PS SERIES FROM DELL — SIMPLY RELIABLE

Let us show you what simplifying networked storage can mean for your business.

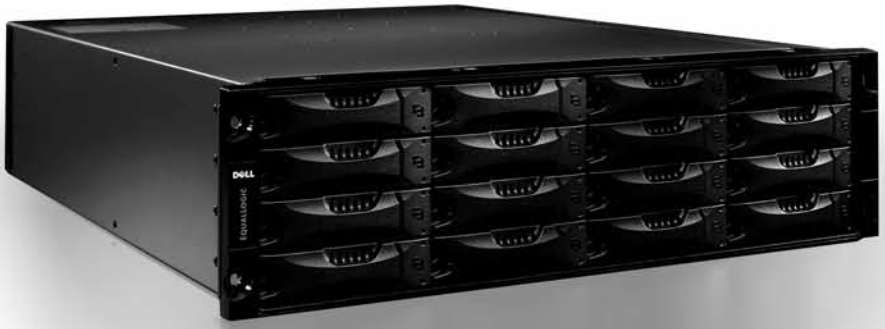

SIMPLIFY YOUR STORAGE, VISIT DELL.COM/DellEqualLogic

DELL EQUALLOGIC

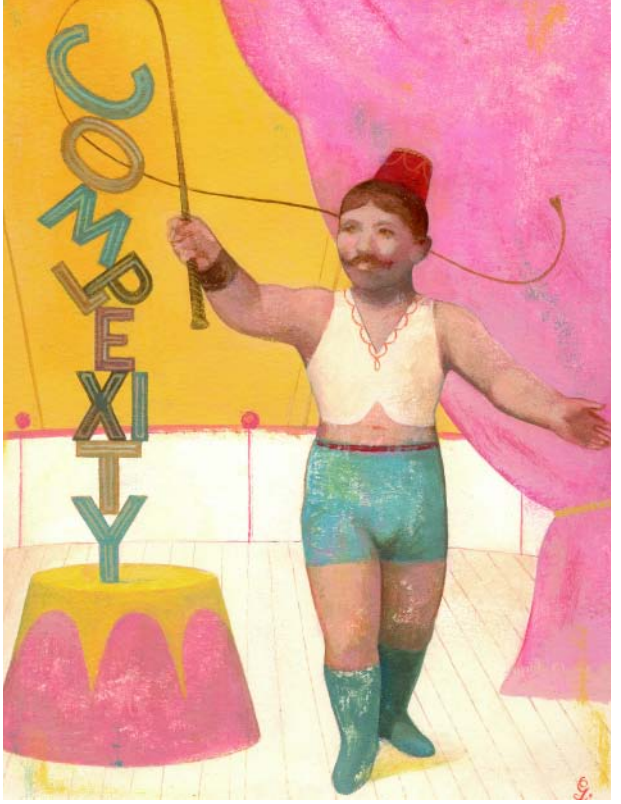

Three ways to virtualize

There are several ways to embed virtualization into your storage infrastructure. Here are the pros and cons of the various approaches.

By Jerome M. Wendt

VIRTUALIZATION IS SHOWING UP in all layers of the storage infrastructure. Host, network and storage system-based block virtualization can each simplify volume management, ease data migrations and decrease storage provisioning times, but all of them increase the complexity of managing a storage environment. As the storage infrastructure becomes more abstracted, organizations may no longer be able to confidently pinpoint what data resides on what storage device or how to isolate the root cause of problems.

Where to embed virtualization and when to use it depends on the size of the infrastructure, the type of apps running in it, and the levels of control and visibility required by administrators.

Today, storage managers need to balance the benefits of virtualization against the complexity it brings. We'll look at seven of the largest storage virtualization providers in North America. Each vendor offers

Illustration by Gérard DuBois

storage virtualization at one or more layers in the storage infrastructure, as well as complementary storage resource management (SRM) software that provides visibility, reporting and management across the multiple layers of the virtual storage infrastructure (see "Virtualization guidelines," below).

There's no right answer regarding the best place to implement virtualization, or at what layer and how much to virtualize. Basically, users

Virtualization guidelines

- **Match the intelligence of your Fibre Channel (FC) SAN infrastructure to your levels of server and storage virtualization.** Using VMware's VMotion, virtual servers may move to different server hardware, but as servers move so do their virtual initiators and targets. FC SAN switches need to recognize these changes so server multipathing and provisioning continue to function across the FC SAN when moves occur; otherwise administrators will need to manually make zoning and LUN masking changes.

- **Partitioning should be available and configurable at all layers.** Partitioning assigns each layer of storage virtualization its own administrators, memory, network bandwidth and processing power. Partitioning helps prevent secondary applications from consuming physical resources needed by mission-critical production applications. Administrators can also make changes for test applications in one partition without the changes impacting production applications in another partition.

- **Virtualize performance-intensive application servers cautiously.** Sometimes less virtualization is better when virtualizing performance-intensive applications. Too much virtualization can make it difficult for administrators to quickly diagnose the underlying resources used by the application. Before virtualizing, administrators should verify that SRM software can integrate with the specific storage virtualization software.

- **Virtualize servers and storage with low performance and utilization first.** Servers with low performance and storage utilization are those most likely to provide the highest return on your virtualization investment and to require the least amount of management. Consider virtualizing these servers using iSCSI SANs, as they can take advantage of some of the virtualization benefits inherently provided by TCP/IP.

have a choice of three approaches to address their combined storage virtualization and SRM software needs:

- Single vendor
- Host-based
- Software mix

SINGLE-VENDOR APPROACH

The single-vendor approach means that if you buy all of your storage gear from one vendor, virtualization will work just fine. EMC Corp. and IBM Corp. actively promote this road of dependency when customers buy EMC Symmetrix and IBM System Storage DS8000 storage systems and then install EMC PowerPath and IBM Subsystem Device Driver (SDD) path management software on their servers.

This path management software complements host-based virtualization by identifying the specific characteristics of the LUNs presented by EMC's and IBM's storage systems, such as which storage controller is the primary controller in a dual-active configuration, and sending I/O traffic to that controller. The path management software then works with native operating system volume managers to combine storage system LUNs that are presented down two or more Fibre Channel (FC) paths so they appear to the host volume manager as the same LUN. It also provides load balancing and path failover down the different FC paths to the LUN.

There's no right answer regarding the best place to implement virtualization.

But once a specific vendor's path management software is installed, it becomes more difficult to add another vendor's storage systems or network-based virtualization products. IBM's SDD software only works with IBM storage systems or IBM's network-based SAN Volume Controller (SVC) storage virtualization product. PowerPath supports EMC's Symmetrix and Clariion storage systems, as well as high-end storage from Hewlett-Packard (HP) Co., Hitachi Data Systems and IBM.

The limited ability of path management software to work with other storage systems literally pushes users to adopt EMC's Invista or IBM's SVC to virtualize their environments. Though Invista and SVC somewhat undermine EMC's and IBM's push for companies to buy only their storage systems, organizations may end up selecting their network-based storage virtualization products based on the path management software they already have in place.

The impetus behind EMC's and IBM's network-based storage virtualization and SRM software initiatives isn't entirely about virtualizing and

managing other vendors' storage systems. Part of the push to virtualize their own storage systems is to make it easier for users to migrate and place data on other vendors' storage systems owned by the user and to then use EMC's or IBM's SRM software for the ongoing end-to-end monitoring and management of the user's virtualized app, data and storage infrastructure.

According to Jim Rymarczyk, IBM fellow and chief virtualization technologist, most of the storage infrastructure costs approximately 10 years ago were associated with the acquisition of storage devices. Now as much as 70% of those costs are associated with management. While virtualization and SRM software can alleviate these management costs, trying to do end-to-end management in a heterogeneous storage network is fraught with problems.

The most acute problem virtualization creates in heterogeneous storage networks is the inability to map exactly what application data resides on which specific disk. As each layer of virtualization further abstracts the data from the underlying disk to simplify management, information such as storage system RAID levels or controller configurations (needed to optimize performance and troubleshoot problems) is lost. Accessing this information and mapping it back to the app almost always requires access to the storage systems' APIs.

With a heterogeneous virtualized SAN, it's impossible to map exactly what application's data resides on which disk.

These problems have led EMC and IBM to develop their respective ControlCenter and TotalStorage Productivity Center (TPC) SRM software suites to integrate most tightly with their own hardware and software products. While the EMC and IBM SRM software supports SMI-S and APIs from other storage systems, Rymarczyk admits that "customers will lose some of their freedom to pick other vendors' storage" if they use IBM's SRM software.

EMC and IBM took significantly different approaches to the architectures of their network-based storage virtualization products. EMC's Invista is a split-path architecture that places the storage virtualization code on a management or control path workstation that resides outside of the FC SAN. Virtualization settings are configured on the control path workstation, which then uploads the code into a cacheless FC switch called the data path controller, such as Brocade Communications Systems Inc.'s AP7420, or a director blade like the Cisco Systems Inc. MDS 9000 Family Storage Services Module (SSM).

Doc D'Errico, VP and general manager of EMC's infrastructure soft-

ware group, says this "stateless" approach preserves the intelligence on storage systems because storage systems perform other tasks such as replication and data optimization. If all of the intelligence was removed from the storage system and placed on an appliance, users would lose some of the inherent benefits provided by storage systems. Keeping the intelligence on the storage systems also prevents users from making a long-term commitment to network-based virtualization. "Users can transition in and out of network-based virtualization more easily using Invista," claims EMC's D'Errico.

The new 2.0 release of Invista addresses some of the deficiencies of the first release. In the first release, the lack of physical redundancy in its Control Path workstations was considered a potential liability. Invista 2.0 creates a Control Path Cluster (CPC) that's physically separated by FC distances; so if CPC workstation 1 fails, CPC workstation 2 can take over. Invista 2.0 also takes advantage of PowerPath's load-balancing features so it can dynamically load balance between the data path controller and back-end storage systems.

Both Invista and SVC require new devices to be inserted into the data path. To install the devices, a storage admin needs to halt application processing, physically rearrange the cables of the FC SAN, and then change FC SAN zoning and storage system LUN masking settings to introduce the network-based virtualization software into the data path.

The absence of cache on Invista presents other longer term challenges. It will minimally delay, if not preclude, Invista from supporting features like asynchronous replication or thin provisioning because these technologies typically rely on cache to work. EMC plans to add thin provisioning to Invista in late 2008, but recommends users adopt its RecoverPoint product if they need to do asynchronous replication between different storage systems.

IBM's SVC more closely resembles a storage system controller because it uses cache in its architecture. The SVC storage virtualization code resides on Linux servers that are deployed in clustered configurations and mirror-write I/Os over FC ports between the cache in the clustered pair. It supports four clustered pairs of servers in a logical configuration, with each clustered pair operating independently of the others. IBM recently added thin provisioning to the SVC.

Chris Saul, IBM's SVC marketing manager, recommends users first go through a capacity-planning exercise before implementing SVC. IBM is aware of congestion problems that can arise if users insert SVC into an existing FC SAN fabric and fail to isolate congestion-causing devices like tape from the SVCs. Sometimes the SVC is inserted in a core-to-edge design that puts excessive strain on inter-switch links (ISLs) which, says Saul, "increases the chances of ISL congestion."

HOST-BASED APPROACH

Storage managers often can't standardize on one or two vendors' storage; instead, they'll have to virtualize and manage the amalgamation of storage their company has accumulated over the years. In these types of cases, their best option is to select a storage virtualization product and SRM software that virtualizes and manages storage resources on the host.

For example, Symantec Corp.'s Veritas Storage Foundation provides a common way to virtualize a heterogeneous environment of storage system- or network-based storage virtualization products at the host. By using Veritas Storage Foundation, companies can virtualize and manage their storage in the same way across all of their platforms without needing to learn operating system-specific volume managers. Users can then manage the applications and storage devices using Symantec's Veritas CommandCentral Storage SRM software while gaining additional management benefits on hosts running Veritas Storage Foundation.

In 2006, Symantec significantly upgraded its Veritas Storage Foundation 5.0 Dynamic Multi-Pathing (DMP) path management feature. Prior to Version 5.0, DMP provided only a round-robin path management algorithm; now admins can choose from seven different algorithms to do path management. Its new default is to look at the minimum queue length on each FC path, identify the least-busy path and then send I/O traffic down that path.

DMP 5.0 also improves its error detection and volume discovery by circumventing the OS when managing specific FC paths and communicating directly with the FC host bus adapters (HBAs) using their APIs. Working with the HBAs, DMP can identify specific SCSI timeouts or commands issued by storage systems. These are normally received by the FC HBA, but not passed on to the volume managers or OS; however, DMP can spot specific FC path problems or storage system trespass errors, and use alternative paths to access LUNs on back-end storage systems.

DMP can also detect how different storage virtualization products present LUNs to the host. LUNs may be presented by storage controllers in active-active (A/A) or active-passive (A/P) states, which affects how Veritas Storage Foundation's Volume Manager treats them. A/A LUNs are simpler to manage because if the LUN is unavailable on one path, Volume Manager can simply try accessing the LUN on an alternative path.

Conversely, A/P LUNs are assigned to and managed by a specific storage system controller; if that controller becomes unavailable, however, it's not as simple as switching to another path as trespass errors on the storage system can occur. By monitoring FC commands

issued by storage system-based or network-based virtualization and received by the FC HBA, DMP can ascertain which alternate path to use to access a LUN without causing a storage system trespass error.

The growth of server virtualization also drove Symantec to port Veritas Storage Foundation's Veritas Volume Manager (VxVM) directly to the hypervisor level starting with the resource manager on Sun Microsystems Inc. Solaris Logical Domains (LDoms) (see "Improving virtual server backups," below). By default, the native Sun Solaris LDom

Improving virtual server backups

Virtual servers must share physical computing resources such as memory, network bandwidth and processing power. This may create issues because idling or excess resources on physical servers are no longer available for backup. Here are some tips for backing up virtualized servers.

- **Disk-based backup.** Disk-based backup should almost always accompany the introduction of server virtualization. More host, network and storage system-based storage virtualization technologies support one or more technologies like asynchronous replication, continuous data protection (CDP) and snapshots, which remove the server and network performance overhead associated with backup and place it on the back-end storage system.

- **Virtualize the I/O.** Sometimes performance issues surface when backing up a bunch of virtualized servers because there's insufficient bandwidth. To prevent performance issues, consider virtualizing the physical server's I/O using network protocols like InfiniBand or 10Gb Ethernet, which minimize the number of server network connections needed while increasing I/O throughput on the port.

- **VMware Consolidated Backup.** The new VMware Infrastructure 3 (VI3) includes VMware Consolidated Backup (VCB) to perform snapshots of virtual servers. VCB orchestrates the creation of a crash-consistent snapshot by queuing write I/Os across all of the virtual servers on the physical VMware server. Once writes are suspended, a snapshot is created that can be transferred to an external backup client. A key disadvantage of VCB orchestrated backups is that the resulting Virtual Machine Disk Format (VMDK) file encapsulates all of the virtual servers on a physical server so VMDK file recoveries are of the entire VMware server, not individual virtual servers or virtual server files.

resource manager virtualizes the storage volumes and FC HBAs presented to it, which it then re-presents as virtual volumes and HBAs to its virtual hosts.

Symantec's Veritas Storage Foundation Manager, used in conjunction with VxVM, lets administrators manage up to 3,000 hosts from a single Web console, including hosts like the Sun LDom. Also, when using Veritas Storage Foundation with Veritas CommandCentral, storage administrators can map what virtual resources are assigned to what physical resources.

At this point Symantec still has no official timeline as to when integration with VMware's ESX Server will be complete; Sean Derrington, Symantec's director of storage management, says Symantec is still waiting for access to VMware ESX Server's APIs before it can port its VxVM to the ESX hypervisor. Once Symantec has the ESX APIs, "we will be there," says Derrington.

Some companies may find they need a mix of storage virtualization and SRM software to virtualize and manage data at multiple points within the storage infrastructure.

SOFTWARE MIX APPROACH

Some companies may find they need a mix of storage virtualization and SRM software to virtualize and manage data at multiple points within the storage infrastructure. Hitachi Data Systems, HP, NetApp Inc. and Sun use storage system controllers to virtualize their own and other vendors' storage systems. While NetApp uses its own V-Series platforms to virtualize other systems, Hitachi, HP and Sun all use Hitachi's Universal Storage Platform (USP) V.

HP has an engineering and OEM relationship with Hitachi in which the two firms jointly develop the XP and USP V families. Inputs, suggestions, fixes and contributions received from HP are incorporated into a single version of Hitachi firmware that's then released into all versions of the XP/USP V products. The only differences between the software on the XP and USP V are ID strings embedded into the firmware. "HP uses these IDs to create disaster recovery solutions that work exclusively with the XP," says James Wilson, HP's XP product manager.

Hu Yoshida, CTO at Hitachi Data Systems, says his firm has specifically chosen to stay out of the network fabric because it's trying to satisfy specific customer high-availability and performance requirements, and deliver specific functions that network-based appliances can't. Unlike network-based appliances, which lose 50% of their performance should an appliance fail and have a fixed amount or no

cache, the USP V can create a cache that spans multiple controllers and processors.

"A fully configured USP V can lose up to four processors and still access the common cache without data loss," says Yoshida.

Hitachi Data Systems, HP, NetApp and Sun offer one set of SRM products to manage their own storage systems and another set of SRM tools to manage heterogeneous SAN environments. HP breaks its SRM software into three categories: unified storage and server management, element management and performance management. If you only need to manage storage products from HP, then HP's element management and performance management classifications will likely meet your requirements. Conversely, if a company intends to manage a heterogeneous storage environment, then they'll need to introduce HP's Storage Essentials product to provide this level of server and storage management.

Vendors are moving closer toward tying their storage virtualization and SRM software together. But because of the time and effort required to implement virtualization and complementary SRM software, you should deploy virtualization products gradually while keeping expectations at a modest level. ⊙

Jerome M. Wendt (jerome.wendt@att.net) is lead analyst and president of DCIG Inc.

File virtualization

File virtualization addresses the critical limitations of many NAS filers. While the field is limited, there are some significant players offering file virtualization products. Here's what they offer and how the products differ.

By Beth Pariseau

FILE VIRTUALIZATION PRODUCTS have emerged because of a common pain point with NAS systems: file-system size limitations. Many traditional NAS devices, such as NetApp Inc. filers, as well as individual servers that act as NAS storage, have upper limits of between 2TB and 16TB. As file data grows, these limitations require lots of painstaking storage management to keep capacity balanced on multiple NAS systems.

File virtualization vendors have taken two paths to address that issue, according to Brian Babineau, an analyst at Enterprise Storage Group (ESG), Milford, MA. In the past year, some of the more high-profile file virtualization products have added a layer of abstraction that allows storage administrators to manage multiple, existing heterogeneous NAS devices from one console. Other vendors have chosen to provide a virtual file system known as a global namespace, which pools data from individual NFS or CIFS devices together. Some of these products also provide NAS connectivity to block-based SAN systems.

Many traditional NAS devices have upper limits of between 2TB and 16TB.

The file virtualization marketplace has seen some consolidation, with some of the startups that began making a name for themselves

over the last two years being snatched up by larger vendors. Several of these vendors, in turn, have repackaged those products to make this type of file virtualization a feature of a networking product or to better target offerings around data migration, which so far has been the most popular use case.

The few players left in this market are listed below, along with a thumbnail description of their products' functionality and history.

VENDOR: **Attune Systems Inc.**
PRODUCT: **Maestro File Manager**
One of the last independent file virtualization vendors left, Attune's claim to fame is a hardware appliance that can move files in-band or out-of-band. The device is based on an embedded Windows Server 2003 OS and also offers heterogeneous snapshots.

VENDOR: **Brocade Communications Systems Inc.**
PRODUCTS: **StorageX, File Management Engine (FME)**
Brocade branched out from SANs into file virtualization with the acquisition of NuView in March 2006. The first product to spring from that merger was StorageX, an out-of-band Windows file virtualization software product. Brocade reports 600 users of StorageX, as well as reseller deals with IBM Corp., NetApp and Hitachi Data Systems. StorageX uses an embedded Windows Server 2003 OS.

Brocade's newer FME product sits in the network and automates migrations of locked and open files according to policy. While the StorageX software can manage files down to the folder level, FME can manage individual files because it sits in the data path.

VENDOR: **EMC Corp.**
PRODUCTS: **Rainfinity Global File Virtualization, Rainfinity File Management Appliance**
EMC was the first major vendor to make an acquisition in this space when it bought startup Rainfinity in 2005. EMC's Rainfinity products are attached to an intelligent network switch and can move in-band and out-of-band, depending on their function. The File Management Appliance, a pared-down repackaging of Global File Virtualization, was announced last year to refocus the product around data migration. Rainfinity uses a proprietary Linux-based OS.

VENDOR: **F5 Networks Inc.**
PRODUCT: **F5 Acopia ARX series**
Networking specialist F5 Networks bought Acopia Networks last year for $210 million, with the eventual goal of adding file-system management to its existing networking products. So far, F5 has added virtual

heterogeneous snapshots to Acopia's system, meaning the system can instantaneously create a snapshot across multiple NAS devices from different vendors. Those snapshots are still limited to EMC Celerra and NetApp FAS systems. Acopia uses a proprietary operating system called the FreedomFabric Network OS.

VENDOR: **Njini Inc.**
PRODUCT: **Njini Information Asset Management Suite**
Njini's software receives all file operation requests—such as open, close and save—by NAS protocols like NFS and CIFS, and maps these protocol operations to a virtual file system, including data classification information, to manage each file according to policy.

In addition, Cisco Systems Inc. acquired file virtualization software startup NeoPath Networks in 2007, but immediately took the NeoPath File Director products off the market. They have yet to be re-released, and Cisco has been vague about its plans for the Linux-based in-band systems.

GLOBAL NAMESPACE

Players in the global namespace marketplace include makers of clustered NAS systems, such as Isilon Systems Inc. and ONStor Inc.; SAN file-system vendors, including Ibrix Inc. and Quantum Corp. (with the StorNext file system it acquired with ADIC); and high-performance NAS vendors, including Exanet Inc. and Silicon Graphics. Microsoft Corp. also has its own global namespace software for Windows systems, formerly known as Distributed File System (DFS) and now known as DFS Namespaces.

ESG's Babineau predicts that the two product categories in file virtualization will converge as file storage continues to grow. "These two tracks in the market both have longevity just because of the amount of data that's going to be stored as files keep growing and growing," he says. "We're seeing more and more money fly that way. There could be collisions along the way." ⊙

Beth Pariseau is the senior news writer at SearchStorage.com.

Is your business prepared to cope with overwhelming amounts of data?

Meet the challenges of rapid data growth with Storage Virtualization solutions from Hitachi Data Systems.

Storage Virtualization is a practical and effective way to manage the increasing complexity of data and growing demands for high availability storage. With the right solutions from Hitachi Data Systems in place, Storage Virtualization can help you respond to exponential data growth in a manner that cuts operational costs, improves efficiency and increases your return on investment.

Download our Virtualization Briefing Paper today and learn how to:

- Leverage existing storage assets
- Increase utilization rates
- Simplify storage management

Visit www.hds.com/virtualization for more information.

HITACHI
Inspire the Next

Tech tips

Here are seven questions to consider before deciding on a virtualization product for your storage environment.

By Arun Taneja

SERVER VIRTUALIZATION PROMISES to bring huge benefits to storage fabrics. But making storage and virtual machines work well together is still a work in progress. The following seven questions will help you decide what type of virtualization products best fit your storage environment.

1 **Who offers storage virtualization? What are the differences?** Numerous companies offer storage virtualization products, which can be broadly placed into three buckets—in-band (typically appliances), out-of-band and Split Path Architecture (SPA)—depending on their architecture and where they sit in the network fabric.

In-band products include DataCore Software Corp.'s SANsymphony, FalconStor Software Inc.'s IPStor and IBM Corp.'s SAN Volume Controller (SVC).

LSI Corp.'s StoreAge Storage Virtualization Manager (SVM) is an out-of-band SAN appliance (it resides outside the primary data path) that provides virtual volume management in heterogeneous environments. The SVM product runs on an Intel server-based appliance.

SPA products, such as EMC Corp.'s Invista and Incipient Inc.'s Network Storage Platform, are similar to those with out-of-band architectures. However, both these products are designed to work in the network with an intelligent switch from Brocade Communications Systems Inc. or Cisco Systems Inc.

Hitachi Data Systems' Universal Storage Platform, which can virtual-

ize the storage of any storage array connected to it, represents a small fourth category of storage virtualization.

2 **Are the products ready?** The short answer is that the products are absolutely ready. Remember that the product examples mentioned in the first question are all designed to work across dissimilar hardware (although they'll certainly work across homogeneous environments). Some vendors, such as 3PAR Inc., Compellent Technologies Inc., Pillar Data Systems Inc. and Xiotech Corp., to name just a few, have done a terrific job of developing products for homogeneous platforms.

Storage virtualization has also been applied to files and blocks (e.g., iSCSI). On the file side, there are the V-Series products from NetApp, as well as offerings from Isilon Systems Inc. and PolyServe Inc. (now owned by Hewlett-Packard Co.). iSCSI (block) virtualization products come from Dell Inc.'s EqualLogic boxes, and companies such as Intransa Inc. and LeftHand Networks Inc.

3 **Does storage virtualization impair storage performance?** You're adding a layer of abstraction between physical storage and the SAN. The fundamental principles of computer science suggest that the added work will take CPU cycles away and draw upon other computing resources that would otherwise be devoted to storage and applications. This undoubtedly adds latency and will result in a performance impact.

It all depends on what type of virtualization product is deployed in your fabric. Overall, it's generally safe to assume that there'll be some performance impact, but you should question the vendor about performance and evaluate the net effect on your environment. Switch-based virtualization mitigates much of the performance impact found in other deployments.

4 **Should storage virtualization be done in the network or the array?** The only example of controller-based storage virtualization is from Hitachi Data Systems. There are pros and cons to this. Virtualization is being done inside the storage system (e.g., USP TagmaStore), committing the user to a homogeneous Hitachi environment; all storage traffic, whether from a Hitachi or a non-Hitachi box, has to flow through the virtualized storage array controller.

If you have a highly heterogeneous storage environment, a network-based solution will probably be more suitable. Network-based storage virtualization can be accomplished through an appliance based on an Intel server, an intelligent switch or a purpose-built appliance. All three techniques have a place in the network, but your choice of approach depends on your current storage vendor and the level of network performance you require.

5 **Once a layer of storage virtualization is added to the storage infrastructure, how difficult is it to remove?** Storage virtualization is a very strategic decision for a company to make. Reversing virtualization isn't easy. Once you go into a virtualized environment, the relationship between what's presented as storage to the server and where that physical storage is coming from is completely abstracted—an IT storage administrator won't know where storage is coming from. You'll probably live with a vendor's storage virtualization product for years to come, so it's absolutely crucial to pick your vendor carefully.

You can back out of storage virtualization, but it's a painful project. You'll basically have to back up your data somewhere else, remove the virtualization layer and then bring the data back onto known physical storage devices. It's not a course to be undertaken lightly.

6 **Is there storage that shouldn't be placed under the storage virtualization umbrella?** Suppose you have an isolated application that works fine, and requires few upgrades and minimal maintenance. In such a case, there's no need to impose storage virtualization. But aside from these relatively rare "fringe" applications, there's no reason to avoid virtualization. However, be cautious of any storage virtualization product that demands 100% control of the storage environment. Virtualization should be applied in a careful, measured fashion.

7 **How does the number of virtual machines affect storage?** The idea behind server virtualization is to create multiple logical "machines" on a single physical server. If I create 20 virtual machines on a physical server, and the server only has one or two host bus adapters (HBAs) and all of the storage is connected through those HBAs, there'll be a serious I/O bottleneck at the storage controller. The storage industry has been addressing such bottlenecks using, for example, N_Port ID Virtualization. This approach essentially divides the HBA into logical pieces that can be connected to each virtual machine, eliminating the storage bottleneck.

Server virtualization vendors, like VMware Inc., are becoming more sensitive to the storage needs of their virtual machines and incorporating storage-centric features, such as virtual consolidated backup that's designed to make backups work across virtual machines without licensing and installing a unique copy of the backup software on each virtual machine. A lot of work is being done to help storage and virtual servers work better together, but this work is still in its infancy. ⊙

Arun Taneja is founder of Taneja Group, Hopkinton, MA.

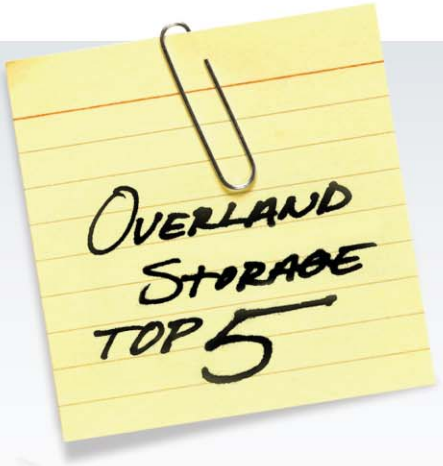

Top 5 Reasons you need a solid Business Continuity Plan

5. Downtime is more expensive than Paris Hilton's bar tab.
4. Sleepover parties are no fun when you're rebuilding a server.
3. "My dog ate the backup tapes" won't fly as an excuse anymore.
2. Your CEO won't learn your name for the first time during an outage.
1. You might just make it home for dinner tonight. And every night.

The success of your company rides on your ability to keep vital business operations running in the event of an infrastructure failure. Extended downtime can be fatal, and the longer you wait, the more damage it can do to your business processes and more importantly, your bottom line.

Overland Storage solutions keep key operations going strong. Our rock-solid tiered data protection works at every level of defense—all within your current infrastructure.

Let Overland Storage help you create a winning Business Continuity plan that makes sense for your business and gives you peace of mind.

Get your FREE white paper today!
www.overlandstorage.com/bc

Overland STORAGE®

Check out the following resources from our sponsors:

attune SYSTEMS

6 Steps for Simplifying File Management with Global Namespace Solutions
Case Study: Optimizing IT Resources through File Virtualization
Increasing the Value from Your Dedupe Solutions: Attune Systems Maestro File Manager and File Virtualization Improve ROI of Dedupe Solutions in Windows-Compatible Environments

DELL

New Deployment Methods for Desktops
Storage Virtualization: The What, the Why and the How
Register Today for a Live Dell EqualLogic PS Series iSCSI SAN Solution Demo

HITACHI
Inspire the Next

Solving the Top Six Enterprise Storage Issues: Learn How to Meet the Challenges of Rapid Data Growth with Storage Virtualization
Tiered Storage and Virtualization in the Real World: Calculating the ROI and Cost Savings of a Move to Tiered Storage by Fidelity National Information Services
Using Storage Virtualization to Meet the Challenges of Rapid Data Growth: Leverage Existing Assets, Reduce Costs, and Simplify Storage Management

NetApp

Building an Effective Virtualization Infrastructure from Server to Storage
Combining Deduplication and VMware Disaster Recovery: Cascading Savings Improves Cost Effectiveness
Managing Virtual Environments: What You Should Know about Backing up, Managing and Tracking Virtual Machines

Overland STORAGE

Business Continuity Planning Resource Page
REO SERIES is the World's Best-selling Disk-based Backup and Recovery Appliance
Calculate your Cost of Downtime Now!