

Virtualization: Tales from the trenches

http://storageMagazine.techtarget.com/magItem/0,291266,sid35_gci1208870,00.html

by: [Ray Lucchesi](#)

Issue: [Aug 2006](#)

While there's plenty of interest in storage virtualization, companies have delayed implementation for a variety of reasons. Here's how five companies took the virtualization plunge--with winning results.

Virtualization is the big buzz these days. Virtual servers are commonplace, and the "V" litany continues with virtual LANs, virtual SANs, virtual tape libraries and so on. Storage managers are clearly interested in virtualizing their storage systems, but while the technologies to do so have been around for a few years, companies are only now beginning to reshape their storage environments using virtualization.

We talked with storage professionals from five companies that have deployed, or are currently testing, storage virtualization to learn about their implementation experiences and how well the products are working in their production environments. The companies are using four different virtualization products: IBM Corp.'s SAN Volume Controller (SVC), StoreAge Ltd.'s Storage Virtualization Manager (SVM), Hitachi Data Systems (HDS) Corp.'s TagmaStore Universal Storage Platform (USP) 600 and EMC Corp.'s Invista (see "[Virtualizing with EMC Invista](#)"). Despite a few problems and some early hiccups, all of the products have fulfilled the promise of storage virtualization--much to the satisfaction of their users.

COMPANIES TEND TO RUN WITH ONLY A SINGLE VENDOR'S STORAGE. BEHIND THE VIRTUALIZATION ENGINE, ALTHOUGH THEY MIGHT MIX DIFFERENT CLASSES OF DISK

Defining terms

For the purposes of our review of these companies' activities, a virtual storage system is de-

finied as one that virtualizes block disk storage behind it. The virtualization products support heterogeneous storage devices and can pool disk storage from any supported attached storage.

Of course, there are several different ways to implement storage virtualization: on the host, in the SAN or at the controller. Block virtualization redirects an I/O request from a virtual location to a physical location, so data can be moved anywhere within the virtualized storage pool transparently to applications and users.

Other features common to virtual storage systems include:

- Snapshot or point-in-time copy across disparate or similar storage devices
- Remote replication across disparate or similar storage devices
- Data migration from one back-end disk group to another while applications continue to access the data
- Centralized management of all virtualized storage

Virtualization products may operate in-band or out-of-band. In-band products reside within the data path of the disk storage, where they intercept I/O requests and may cache responses to the I/O requests. Out-of-band devices typically don't have cache and, although they may redirect an I/O request, they don't alter the data path to storage.

Cloverleaf Communications Inc., DataCore Software Corp., FalconStor Software Inc. and IBM offer in-band products at the SAN level. In-band products that work at the controller are available from HDS, Network Appliance Inc. and Sun Microsystems Inc. (and its OEM partners). EMC and Fujitsu Computer Systems

Corp. offer out-of-band virtualization products that operate in the SAN. StoreAge and Symantec Corp. offer host-level, out-of-band products.

Why virtualize?

Many users virtualize their storage to gain better control and increase their capacity utilization levels. Whether faced with data center consolidation or just normal expansion, storage virtualization technologies promise to lower costs by managing more terabytes with fewer people.

Typically, companies tend to run with only a single vendor's storage behind the virtualization engine, although they might mix different classes of disk such as SATA and Fibre Channel (FC). Often, data migration features are used to migrate current storage to a virtualized environment while applications continue to access the data. Some companies are also using storage virtualization to create tiered storage because it makes it easier to move data among storage tiers without disrupting applications using the data.

Using IBM's SVC

Andrew Madsen, a senior storage administrator at Harley-Davidson Inc., says his company runs three IBM SVC clusters. One cluster virtualizes more than 50TB of IBM DS83xx storage with two I/O groups. There are two nodes per I/O group in the SVC providing hardware redundancy and data availability. All of the I/O groups within an SVC cluster are managed as a single entity sharing the I/O workload.

Harley-Davidson's second SVC cluster has 11TB of storage using a single I/O group, while the third is still being tested. The company implemented the SVCs as part of a data center consolidation effort.

Harley-Davidson installed its first SVC in November 2005. At the time, it was an EMC shop and initially used SVC to migrate from its EMC equipment to IBM storage while applications continued to access data. The company had to move 42TB of data, with LUNs ranging from 1GB to 2TB in size. Without virtualization, a significant operational outage would have been required.

Madsen is currently using flash copy (point-in-time copy) as part of his backup scheme, and data migration to move data from FC to SATA disk. He has plans to use IBM's Metro Mirror,

which provides synchronous remote replication over short distances, to provide for disaster recovery.

One misconception about in-band storage virtualization is that it can degrade performance. The reality is more complex. While in-band products add some latency to every I/O operation, they also add cache to compensate. Often, the additional overhead is so negligible that it's invisible to normal applications; the additional caching can also improve performance considerably, depending on the back-end storage. "Once it [SVC] was working correctly," says Harley-Davidson's Madsen, "it was equal to or better than direct-attached [storage]."

But there are some pitfalls to virtualized storage. Early on, Madsen had some configuration issues with SVC and EMC storage as his company's Clariion-to-SVC configuration wasn't done correctly and had to be modified. Also, the Symmetrix-to-SVC configuration was generating many I/O error retries, which caused hosts to take EMC volumes offline. At that time, SVC didn't see any problem with the volumes. Ultimately, it was determined that a fix to the host driver software resolved the problem.

One thing Madsen would like to improve is the way SVC handles a node outage. Each SVC node in an I/O group has an affinity to some of the LUNs assigned to the I/O group. If a node fails, its LUNs are moved to the other node in the I/O group; however, when this happens, they're put in "write-through mode" (no write caching). This hampers write performance and can cause volumes to be taken offline by hosts detecting bad performance, which results in application outages. Madsen would rather see the volumes moved to another I/O group (he can do this manually, but it's time consuming and is usually too late to avoid application outages).

The worst incident Madsen encountered was an outage that lasted more than eight hours for some SVC storage and 36 hours for the rest, which occurred while Harley-Davidson was still migrating data to SVC. Some of the SVC storage was in "image" or pass-through mode when the SVC went down, which meant it wasn't fully virtualized yet. Madsen was therefore able to bring that storage up outside the SVC within eight hours. The rest of the SVC storage had already been completely virtualized with no way

Virtualization with EMC Invista

While it's not yet in its production storage environment, a large legal technology provider in the Midwest is testing EMC Corp.'s Invista, a switch-based, out-of-band virtualization product. Steven McCaa, storage manager at the firm, says they've been testing the product since December 2005, with a test environment ranging from 1TB to 10TB of storage behind the Invista product. All told, the company currently has 2 petabytes (PB) of on-line storage, with 70% on Fibre Channel (FC) disk and the remainder on SATA.

The key motivations for the Invista pilot were business process improvement and data lifecycle management. With 1.4PB of FC disk, any substantive amount of data that can be moved to SATA can quickly pay for the Invista. The company installed pre-general availability Invista code and ran into some known bugs, but it's been a relatively straightforward install since then.

"You still have to be smart about what you are doing [with virtualization]," says McCaa, who only intends to virtualize the firm's Clariion-class storage. The company's Symmetrix storage is a highly tuned environment where the right LUNs are mapped to the right physical spindles, and McCaa doesn't want to upset that arrangement.

Virtualization products in general "don't give a good global map of virtualized storage, so that when something is broken you can't see end to end what is impacted," says McCaa. "If an ISL [inter-switch link] goes down and you see performance problems, you can tie the two together." Virtualization adds a layer of abstraction that needs to be understood to determine why performance went bad.

McCaa says the company hasn't seen a performance impact with Invista. He notes that there are four modes to data migration on Invista: trickle, very slow, slow and ASAP. "When you use ASAP," says McCaa, "you can tell right away" if performance degrades. In contrast, he indicated that migrating a 1TB volume at slow speed didn't impact performance, but did take a long time.

to back out, so Madsen had to wait until IBM fixed the SVC.

Despite a few bumps in the road to virtualization, Madsen likes SVC and thinks IBM has a real winner. "SVC will be here for a long time," he says, adding that it allows him to "better manage where storage sits" within his environment. In pre-SVC days, moving data between Clariion and Symmetrix storage required a long application outage, so it was never done. With SVC in place, Madsen can make the moves without any application outages.

FirstMerit Corp., Akron, OH, was one of IBM's first 100 SVC customers. Dave Samic,

senior network analyst at FirstMerit, says the bank installed SVC in the spring of 2004 and currently has 16TB operating through one SVC cluster with a single I/O group.

Samic had a couple of IBM TotalStorage FAST storage subsystems that were causing performance delays when the bank booted multiple virtual servers. In addition, FirstMerit had three disk LUNs with 146GB disks supporting a usable 280GB of space, but it could only use 100GB before the LUN began performing poorly.

IBM suggested implementing SVC as an alternative to buying faster storage. With SVC and its additional cache, FirstMerit can use much more disk space and boot 30 to 40 virtual servers without delay.

FirstMerit virtualizes its blade servers with VMware. Samic says the improved performance and manageability realized with SVC allows his firm to support 22 virtual servers per physical blade--considerably more than the 12 to 15 that VMware recommends. Samic also reports that FirstMerit is getting 60% to 75% blade utilization.

Both Samic and Harley-Davidson's Madsen say SVC's GUI is a bit slow. "[The GUI] only displays 10 disks per screen," says Samic, "and with 200 disks, that takes 20 screens."

Samic noted some internal resistance when he tried to include SATA disk under SVC. At the time, SVC licensing was based on the amount of storage under SVC control, and paying premium for cheaper SATA storage made little economic sense. However, Samic demonstrated that the bank could actually save money by taking better advantage of the SATA storage once it was under SVC. Since then, IBM has made some changes to SVC licensing, which would have made Samic's justification easier.

Virtualizing HP and HDS storage

After experiencing problems with an in-band virtualization product, Todd Wyman, Unix storage administrator at a Midwest manufacturer, says his company opted for StoreAge's SVM, an out-of-band, host-based product. The company has approximately 45TB behind SVM, with both HDS Thunder and Hewlett-Packard Co.'s Modular Smart Array (MSA) storage, a Brocade Communications Systems Inc. SAN infrastruc-

ture, and servers running HP-UX and Windows.

The StoreAge SVM was installed in January 2005 and Wyman says it took his firm approximately a year to migrate fully to the new environment, which was achieved without buying additional storage. Wyman's company carved up a bit of its current storage to virtualize it, migrated other storage to that and then virtualized the newly freed storage. In this fashion, they virtualized their whole complex of existing storage over time.

**IT'S NOT IF THERE WILL
BE ISSUES, BUT WHEN
AND WHEN THEY DO OCCUR
IT'S HOW WELL THEY RESPOND
THAT MATTERS**

Wyman says the company is using StoreAge's multiMirror (remote replication) in asynch mode for its disaster recovery site. It's also using point-in-time copy mirroring at the remote site and has added point-in-time copy to its other backup scripts.

Wyman encountered some issues during installation when SVM's host software agents wouldn't run on some older HP-UX systems to support SAN boot, but StoreAge remedied the problem. Some bleeding-edge Windows servers also needed to be certified with StoreAge SVM, which was a Windows certification issue not a StoreAge shortcoming. "It's not if there will be issues, but when," says Wyman. "And when they do occur, it's how well they respond that matters." With open systems, there's an almost infinite number of combinations of OS levels, host bus adapters, SAN hardware and storage--and, often, not all of it can be validated in advance.

Wyman has only praise for the level of service provided by StoreAge and the easier management SVM affords. "Single pane-of-glass management is a boon for anyone," says Wyman, who also likes the performance he's seen so far. He says performance has improved because he's now striping his virtualized LUNs across many more disk spindles.

Healthcare system taps TagmaStore

The University of Utah Health Sciences (UUHS) Center in Salt Lake City virtualizes its 135TB of

Virtualization Tips

Virtualization products can be a boon to your storage environment, but they come with their own level of added complexity. To make your transition to virtual storage easier, you should do the following:

CHECK OPERATING SYSTEM COMPATIBILITY. Many virtualization products require agents to be installed on your servers. Check to make sure your current operating system levels are supported by the agent software.

CHECK SAN AND STORAGE HARDWARE COMPATIBILITY. Some virtualization products support a wide range of storage and SAN hardware, but have only been validated in specific configurations. Verify that your SAN configuration and storage hardware are currently supported or can be brought into compliance with supported configurations.

CHECK YOUR VIRTUAL STORAGE CONFIGURATION. Virtualization products may have limitations or guidelines on how much storage is supported per virtualization node. Make sure you operate within the performance and capacity envelope the vendor recommends.

CHANGE YOUR WAY OF THINKING. You may be used to defining LUNs and assigning them to specific RAID groups, but with virtualization you won't have to anymore. Although these products will let you continue to do that, they work better if they decide which back-end RAID groups to assign to which virtualized LUNs

FC and SATA disk using an HDS TagmaStore USP600, which provides controller-based, in-band virtualization. Controller-based virtual storage products typically support both internal (within the virtual storage product cabinet) and externally attached storage. Jim Livingston, director of data resource center and technical architect at UUHS' Information Technology Services, and John Fagg, manager of storage management services, say their virtualized system supports three tiers of storage: high-end for critical apps, midtier for user and file space, and low-end for shadow image and disk-to-disk copy.

UUHS added a TagmaStore USP600 to a mix of four HDS Thunder models; some older IBM and Sun storage systems were also being phased out. They began with only 10TB behind the TagmaStore and in less than 20 months have grown its virtualized capacity to 135TB. Per-

formance and scalability were key reasons why Livingston and Fagg chose the TagmaStore product for their virtualization project. In addition, performance has improved considerably. "It varied from user to user, but [some people] saw a 300% improvement in some circumstances," says Livingston.

Tiered storage was also a part of their TagmaStore plans; before virtualizing, most of UUHS' data was on high-end storage. Eliminating "stranded" storage and simplifying storage management were also project goals. UUHS currently has a mix of 600 AIX, Sun Solaris Enterprise and Intel blade servers accessing the storage with no problems.

Livingston and Fagg say their virtualization project has been relatively problem free, as it was essentially a forklift upgrade that resulted in an all-HDS environment. UUHS also took the opportunity to reconfigure its SAN to add redundancy. Livingston says they've had plenty of

experience with other infrastructure upgrades, "but this has been the most successful [one] we have experienced," he notes.

Virtualization delivers

Regardless of the route to virtualization--in-band, out-of-band, at the host, in the SAN or in the controller--these companies have one thing in common: They're all happy with their virtualization implementations and don't want to return to "pre-virtualization" days. The problems they encountered while implementing these products were fairly minor and resolved in relatively short order. Although vendor lock-in is often cited as a deterrent to storage virtualization, these companies have, in some cases, managed to virtualize heterogeneous storage and are comfortable with the outcome. For these companies, at least, virtualized block storage has arrived and certainly proven its value.

About the author:

Ray Lucchesi is president of Silverton Consulting and since 2004 has helped startups to Fortune 500 companies improve storage product development and marketing. Ray has also helped end-users better select and use storage subsystems.

<mailto:info@silvertonconsulting.com>

<http://www.silvertonconsulting.com>

Storage magazine is a part of the TechTarget portfolio of enterprise IT-focused media. Copyright 2001-2006, TechTarget. All Rights Reserved.