Managing the information that drives the enterprise

STORAGE

ESSENTIAL GUIDE TO

Storage Virtualization

Storage virtualization may not get as much attention as its server virtualization cousin, but it's a proven technology to lower equipment costs and solve your heterogeneous storage infrastructure issues.

INSIDE

TechTarget
The Technology Media ROI Experts

No more excuses

THE REVIEWS HAVE BEEN GOOD, but so far storage virtualization hasn't exactly been a box-office smash. It's more akin to that old TV commercial where the kids don't want to try some healthy-looking cereal so they get the little brother to try it. Then when he empties his bowl, they want some, too.

You probably have a friend at another company or met a storage manager at a conference who's told you all about their storage virtualization implementation trials and tribulations. Don't let those tales of woe deter you. Think of it as that kid with the bowl of cereal: They tried it, they probably like it now and you can learn a lot from their experiences.

It's not easy to decide where virtualization should reside, but there may never be a better time than now, when efficiency is the byword.

Then, of course, there's the classic conundrum of where to put virtualization: hosts, network or arrays. All three options have their appealing and not-so-appealing features, so it can be hard to weigh them on their technical merits alone. But the choices narrow when you assess your environment and requirements; in most cases, only one of the options will be anything close to viable when you consider your currently installed equipment, budget and operational environment.

But wait, there's more, you say. The dreaded vendor lock-in looms over storage virtualization as ominously as any of the other deterrents. I think that hazard is inflated, too. How many storage shops don't already enjoy a tight relationship with a select vendor or two? It happens because it works and benefits both sides. Storage is a long-term affair; you buy an array and it sits on your shop floor for three, four, five years or longer. It's the same with virtualization; three or four years from now, virtualization will be such an automatic feature that you won't even remember feeling locked in way back in 2009.

It's not easy to decide where virtualization should reside, and implementation's no snap either. It'll take some work and an investment in time and money. But there may never be a better time than now, when efficiency is the byword and you're trying to wring every last drop out of your installed storage capacity. ⊙

Rich Castagna (rcastagna@storagemagazine.com) is editorial director of the Storage Media Group.

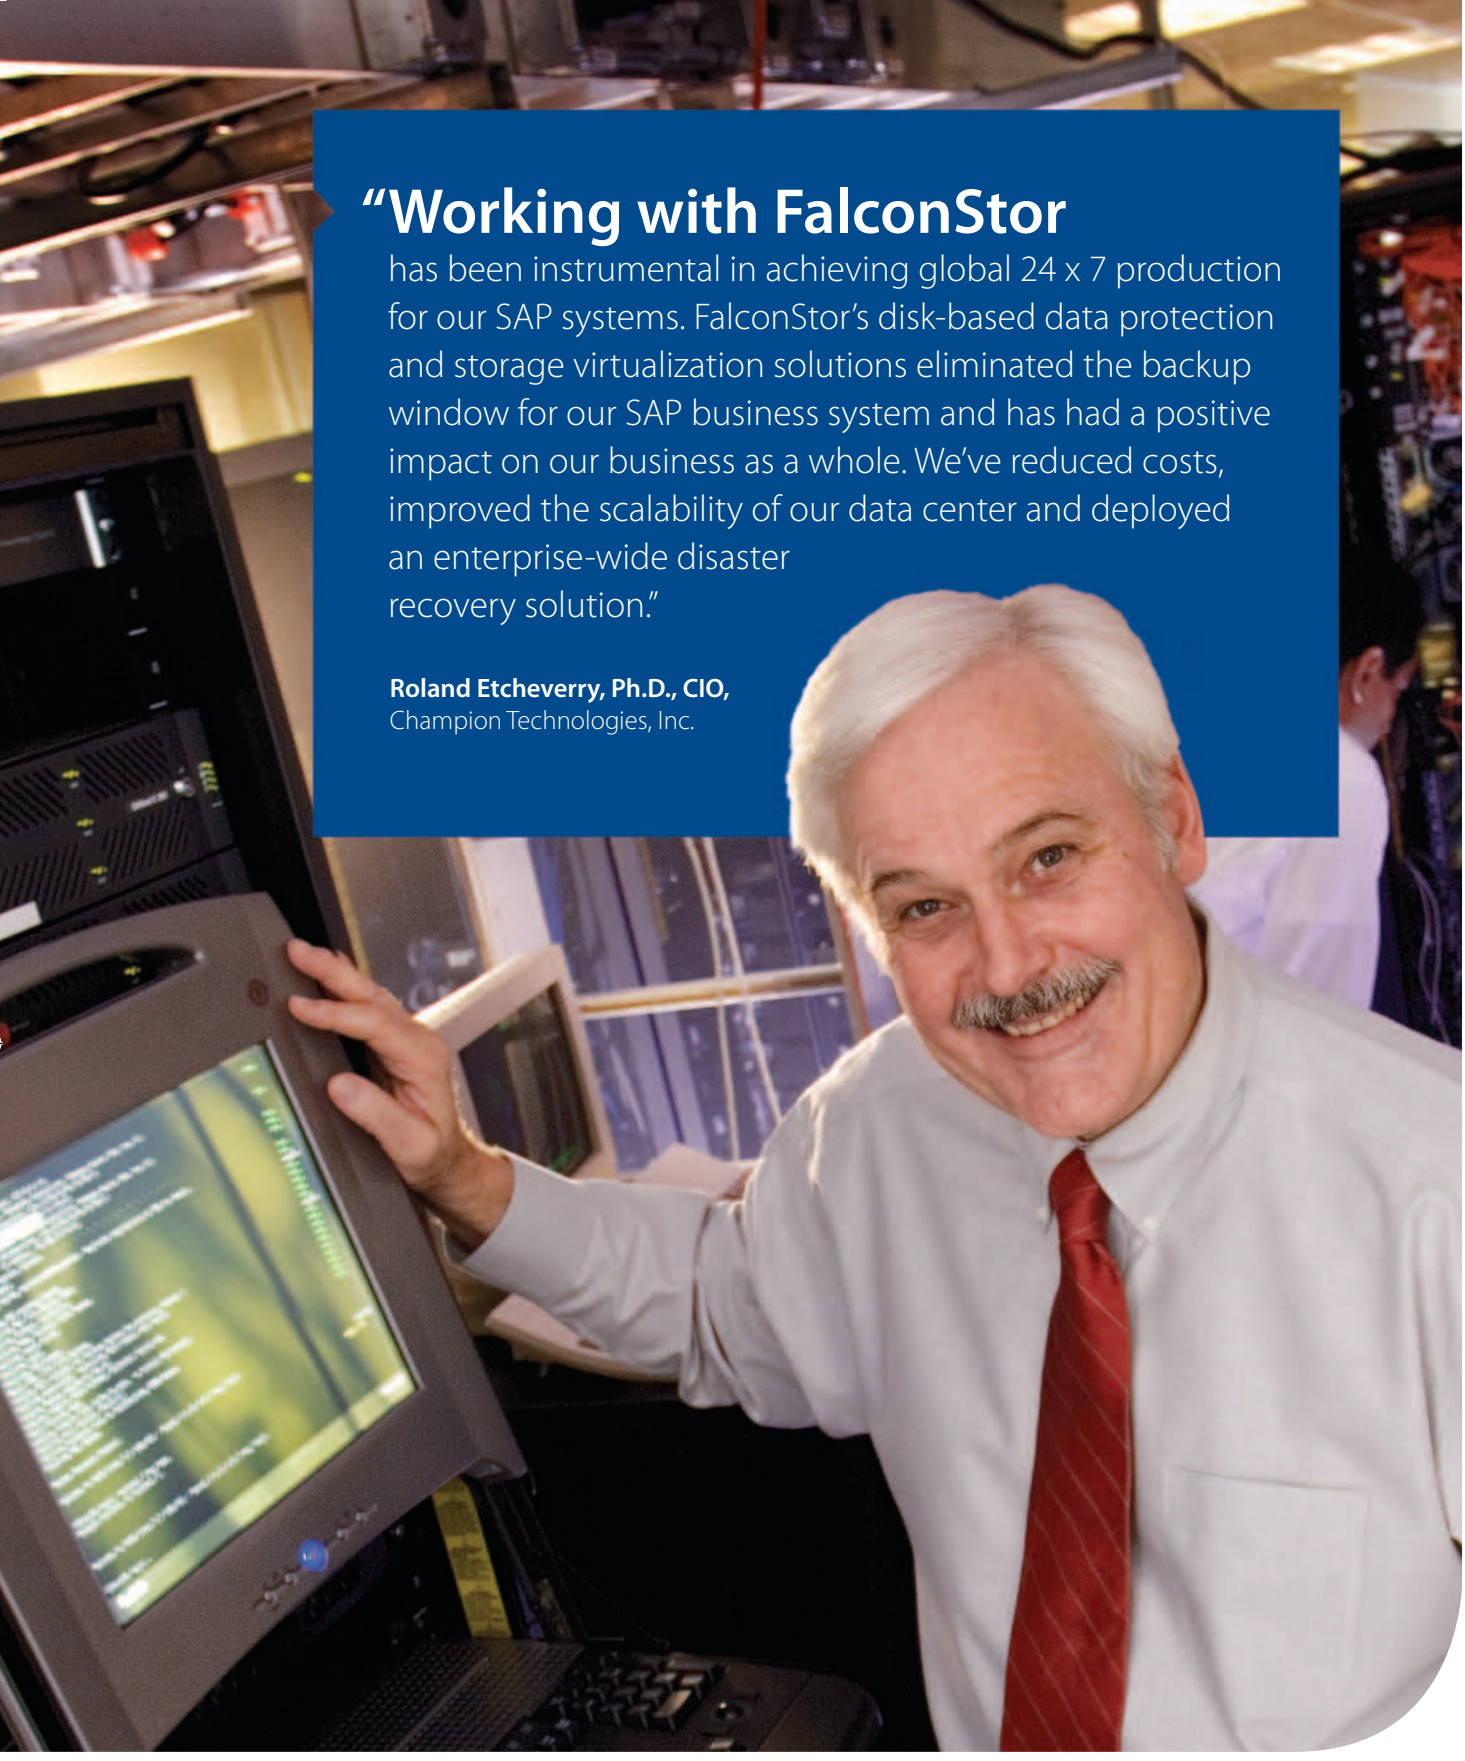

"Working with FalconStor

has been instrumental in achieving global 24 x 7 production for our SAP systems. FalconStor's disk-based data protection and storage virtualization solutions eliminated the backup window for our SAP business system and has had a positive impact on our business as a whole. We've reduced costs, improved the scalability of our data center and deployed an enterprise-wide disaster recovery solution."

Roland Etcheverry, Ph.D., CIO,
Champion Technologies, Inc.

To learn how to deploy enterprise-wide data protection and storage virtualization with thin provisioning:

Contact FalconStor at 866-NOW-FALC (866-669-3252) or visit www.falconstor.com

FalconStor Software

Update your storage system with server virtualization

Server virtualization and consolidation efforts might push storage virtualization efforts into the limelight.

By Steve Duplessie

A LOT OF STORAGE TYPES are put off by the overwhelming hoopla that's been generated by VMware Inc. and server virtualization. That's because storage intellectuals feel that it's unfair for server virtualization to get all of this attention considering application and implementation models have been seen in the storage market for years and years. But these people miss the point.

It's all about timing and the IT caste system. Storage consolidation, utilization and automation via virtualization started around 1988 with Veritas Volume Manager, and it's still going on. We went through RAID, logical unit numbers (LUNs), snapshots, continuous data protection (CDP), synthetic fulls and dedupe, and we're still not done. Comparatively, virtualization in the server layer had a big-bang arrival. Virtualization in the storage/data layer, meanwhile, has been a slow ride down a lazy river.

The reason storage never received the glory for all of its wonderful virtualization efforts is simple; in the grand scheme of things, the business still views servers as higher on the food chain and storage on the level of primordial ooze. It's an incorrect assessment, but a fact of life. This is because the server layer is the first infrastructure touchpoint associated with an application, which is what the business cares about.

Everyone says "our data is the lifeblood of the company," but it's all bull. Generally, business users view data as a part of their application interfaces; without it, data would be a bunch of 0s and 1s, which isn't all that useful to them. If business people truly believed that all of their shiny applications and gizmos would be entirely useless without data, they would make sure data was stored, protected, managed and delivered using at least a few processes and technologies from the current millennium. To some degree, it's our own fault. We've never been able to easily demonstrate our value in terms that business people understand—we're still geeks in storage. For some reason, it must sound smarter when a server guy talks about virtualization efficiency gains than when a storage guy does. Or maybe the storage guy is only talking to a filing cabinet in the basement? I don't know. What I do know is that everyone is paying attention to server virtualization.

I also know we should stop whining about how our industry was the pioneer of all the stuff that now has our business colleagues excited, and instead be leveraging the movement to our advantage. Consolidation is a good idea for servers and storage. Server virtualization allows us to finally get into the '80s, because any big disruptions to our infrastructure allow us (or force us) to take the time to correct legacy issues. Server virtualization will finally enable us to get rid of our direct-attached storage (DAS) and get our entire data infrastructure networked. It will allow us to re-evaluate our current, dumb ways of doing backup and do disaster recovery even if we're only a 20-person company. Server virtualization can have the glory, as long as it gets the core data architecture where it needs to be.

I don't care if the business guys think Microsoft invented "system" virtualization or that it was bestowed upon us by aliens three months ago. This is an opportunity for storage folks to evolve out of the sludge. Who knows, maybe next year we'll even get invited to the holiday party. ⊙

Steve Duplessie is founder and senior analyst at Enterprise Strategy Group. See his blog at http://esgblogs.typepad.com/steves_it_rants/.

nexsan TECHNOLOGIES

Save money

Save energy

Save all kinds of green

Alpine Bank cut its data storage footprint in half, reduced storage costs by 30% and sliced over 20% in power usage using Nexsan's energy saving AutoMAID™.

Visit www.nexsan.com/save to discover industry leading performance with next generation storage efficiency.

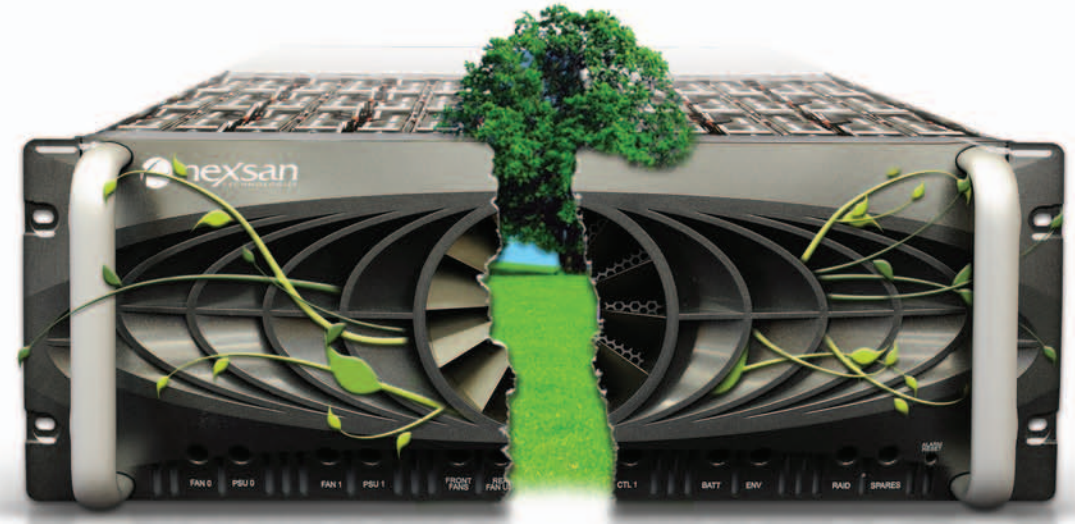

PROS AND CONS OF three architectures

Storage virtualization can reside in the fabric switch, in an appliance or in the array's controller. Each architecture has its pros and cons.

By Jacob Gsoedl

FABRIC-BASED VIRTUALIZATION PRODUCTS haven't been adopted as quickly as inline virtualization appliances like IBM Corp.'s System Storage SAN Volume Controller (SVC), but they're one of the most promising storage virtualization solutions. "Fabric-based virtualization is the best technical approach to storage virtualization, but it's not taking off," said Jim DeCaires, storage product marketing manager at Fujitsu Computer Systems Corp.

Switch-based virtualization brings many benefits to the storage-area network (SAN) fabric. Because the switch-based virtualization engine is out-of-band (out of the data path), there's no need for server agents, and it's the most scalable and highest performing of all virtualization architectures.

Storage virtualization sends data to physical arrays, from single or multiple vendors, as a single storage pool with one major benefit: Storage can be managed as if it were on a single array, using features such as provisioning to advanced functions like replication, snapshots and mirroring between pool arrays. To accomplish this, storage virtualization products map virtual volumes to physical devices; whenever a virtualized storage resource is accessed, the virtualization layer uses a mapping table to translate and redirect storage requests to designated physical storage units.

With three primary storage virtualization architectures—in-band appliances, storage controller based and fabric based—the location of where virtualization should occur is hotly debated. Each approach has its pros and cons.

IN-BAND APPLIANCES

These products, like switch-based virtualization products, perform virtualization within the network. They're located between arrays and servers, and all storage traffic passes through them. While fabric-based virtualization uses wire-speed switching to map and forward storage frames, in-band virtualization appliances require terminating incoming I/Os and initiating new I/Os to the actual storage target based on the mapping table information.

"The process of terminating, re-initiating and verifying I/Os adds significant latency to I/O processing," said Brian Garrett, technical director at Milford, Mass.-based Enterprise Strategy Group's (ESG) lab. To compensate for the overhead and performance penalty of spawning new I/Os, products like IBM's SVC depend on cache, which adds the complexity of ensuring data integrity and data consistency within the cache, a problem switch-based virtualization products don't have.

IBM SVC is the most prominent product in this category, and through scalable cluster configurations and plenty of cache, it has greatly reduced the performance and scalability concerns that have plagued in-band virtualization products in the past. A relatively low cost, simplicity and a rich feature set have greatly contributed to in-band virtualization being today's most widely deployed virtualization architecture.

"In-band virtualization products like IBM SVC or DataCore Software Corp.'s SANsymphony have the lowest entry cost; unlike fabric-based virtualization products, they don't require expensive intelligent switches," said Greg Schulz, founder and senior analyst at StorageIO Group, a technology analyst and consulting firm in Stillwater, Minn. Because products like IBM SVC work with any switch, in-band virtualization appliances have another advantage over fabric-based products like EMC Corp.'s Invista, which runs only on supported switch platforms.

STORAGE CONTROLLER-BASED VIRTUALIZATION

This architecture, championed by Hitachi Data Systems and used in its Universal Storage Platform (USP) V storage systems, performs virtualization within the storage controller of the array. A non-Hitachi Data Systems array can be virtualized by simply plugging it into a Fibre Channel (FC) port on the USP V. To third-party arrays, the USP V presents itself as a Windows server; once the third-party array is discovered by the USP V, it appears to other servers as a Hitachi

Data Systems array. Unlike switch-based virtualization, for organizations that have standardized on Hitachi Data Systems storage and already own USP arrays, the effort to enable virtualization is minuscule and relatively inexpensive. "About 50% of our USP V customers purchase a virtualization license, and with 9,200 USP V units sold in the past three and a half years, we have a significantly higher number of virtualization deployments than all fabric-based installations combined," claimed Claus Mikkelsen, chief scientist at Hitachi Data Systems.

For those who are using or who have standardized on array-based virtualization, vendor lock-in is high, even more so than for fabric-based virtualization. "You wouldn't buy a Hitachi USP V for virtualization if you're an EMC or NetApp shop; but USP V virtualization would be on top of your list if you had standardized on Hitachi storage," StorageIO Group's Schulz said.

Using arrays and virtualization software from a single vendor has the huge benefit of having a single point of support. In stark contrast, fabric-based virtualization products, namely those from EMC and Incipient Inc., require the orchestration of three different vendors (array, switch and virtualization software vendors), which clearly carries the risk of finger-pointing if problems arise.

VIRTUALIZATION APPROACHES COMPARED

Architecture	Pros	Cons	Products
Fabric-based, split-path architecture	High performance, high scalability	Expensive; high degree of vendor lock-in; more limited in advanced storage features	EMC Corp. Invista and RecoverPoint, Fujitsu Computer Systems Corp. Eternus VS900, LSI Corp. StoreAge Storage Virtualization Manager
Storage controller based	High performance, high scalability, rich set of advanced storage features	Array vendor lock-in; typically considered only by customers who have already standardized on the array vendor	Hitachi Data Systems Corp. Universal Storage Platform V and Universal Storage Platform VM
In-band appliances	Simplicity, cost-effective, rich set of advanced storage features	Higher latency than fabric-based products; more likely to hit a scalability limit than switch-based virtualization products	DataCore Software Corp. SANmelody and SANsymphony, FalconStor Software Inc. IPStor, IBM Corp. System Storage SAN Volume Controller

FABRIC-BASED VIRTUALIZATION

Fabric-based products map virtual storage to physical storage within the network, more specifically within an FC switch or director. Unlike network-based virtualization appliances like IBM's SVC, switch-based virtualization is typically implemented as a split-path architecture where the data and control paths are separate. With the virtualization logic running outside of the data path, I/Os pass directly through the switch without the speed bump of in-band solutions. The control-path software typically runs on switch CPUs, and it gets involved only if I/Os need to be redirected to instruct the switch where to route storage requests.

"In a split-path virtualization architecture, 90%-plus of the requests pass through the switch at wire speed; only if something special like migrating of data needs to be performed [does] the control-path controller get involved," explained StorageIO Group's Schulz. The separation of the data path and control path, combined with the low latency of switching for translating and forwarding virtualized storage requests, makes fabric-based virtualization the best performing and most scalable virtualization architecture today.

"In a split-path virtualization architecture, 90%-plus of the requests pass through the switch at wire speed; only if something special like migrating of data needs to be performed [does] the control-path controller get involved."

—Greg Schulz, senior analyst, StorageIO Group

On the downside, switch-based virtualization has the highest level of vendor lock-in of all virtualization approaches. Because the switch is used as the platform to run the virtualization software, it becomes very difficult for users to change switch vendors. Furthermore, as intelligent switches turn into multitasking platforms, concurrent storage services from the switch vendor and third parties make supporting these switches more challenging.

As long as there are no problems, it's a great concept. But if there are problems with the virtualization software or any of the third-party storage services, the concerted effort of all involved parties may be required. Besides the relatively high cost of intelligent switches, increased complexity and more challenging technical support are among the contributing factors for the cautious adoption of fabric-based virtualization. "In general, storage managers like to keep things simple and tend to go with more self-contained, easier-to-manage solutions like LSI [Corp.'s] StoreAge or even IBM SVC," said Nelson Nahum, who was chief technical officer at StoreAge Networking Technologies before it was acquired by LSI.

Without question, the low latency of fabric-based virtualization is a

big plus, but while it eliminates the use of cache, there's a downside: Virtualization solutions with cache, like IBM's SVC and Hitachi's USP V, use that cache to increase performance of the back-end storage. As a result, virtualization products with cache encourage the use of lower-cost, lower-performing storage tiers, with the cache boosting access performance. While the low latency of switch-based virtualization products is great for accessing fast arrays, its lack of cache actually turns into a disadvantage for accessing lower-performance arrays. "In switch-based virtualization, back-end disk performance shows unmasked," StorageIO Group's Schulz said.

A second and more profound implication of the stateless nature of switch-based virtualization is the more challenging support of virtualization apps that require information beyond the mapping information. Features such as remote replication and thin provisioning require memory to maintain certain state information. For instance, for a 2 TB thin-provisioned volume using 100 GB of physical storage, information about which 100 GB are actually used needs to be maintained. While products like IBM's SVC and Hitachi's USP V maintain this information in memory along with the cache, switch-based virtualization products don't have the luxury of cache memory, and their only option is maintaining this information on the SAN.

> "There's no complete solution for remote snapshots, remote mirroring and thin provisioning in switch-based virtualization products today because they're very difficult to implement without cache."
>
> —Jim DeCaires, storage product marketing manager, Fujitsu Computer Systems Corp.

"There's no complete solution for remote snapshots, remote mirroring and thin provisioning in switch-based virtualization products today because they're very difficult to implement without cache," Fujitsu's DeCaires said.

SWITCH-BASED PLATFORMS

Fabric-based virtualization products are offered by the following vendors.

- **EMC Invista:** Invista is the most prominent fabric-based virtualization product. While other vendors support only a single switch platform, Invista runs on Cisco Systems Inc.'s MDS multilayer SAN switch, as well as on Brocade Communications Systems Inc. switches and directors. On Cisco switches, Invista requires and runs on the Cisco MDS 9000 Storage Services Module (SSM), which provides 32 FC ports with embedded ASICs that perform the mapping and wire-speed switching of virtualized storage requests.

Invista also runs on Brocade's 7600 Application Platform, available as a switch blade and a standalone appliance, with 16 FC ports with embedded ASICs. For reads and writes, the ASICs in the Cisco and Brocade modules look up the virtualization mapping information from

the mapping table in memory and forward frames through the applicable FC port to the target at wire speed without the need of the CPUs on the intelligent switch module becoming involved. The control path of Invista consists of virtualization software running on the Cisco SSM or Brocade 7600, as well as the Data Path Controller (DPC) appliance. The virtualization software on the switch communicates with the DPC to receive information such as virtual disk configuration and directions for copy functions. Like all fabric-based virtualization products, Invista passes commands from the external DPC appliance to the intelligent fabric using the Fabric Application Interface Standard (FAIS) protocol.

- **EMC RecoverPoint:** RecoverPoint is another fabric-based virtualization product that complements Invista for those customers who need remote replication or continuous data protection (CDP). In 2006, EMC acquired privately held Kashya and subsequently released its technology as RecoverPoint. While Invista attempts to address a range of virtualization tasks, RecoverPoint's sole focus is on remote replication and remote site incremental snapshots via the underlying CDP engine. Invista's lack of remote replication prior to RecoverPoint is an example of the challenges fabric-based virtualization vendors face in adding features that require state information beyond the virtualization mapping table.

- **Fujitsu Eternus VS900:** Similar to Incipient's Incipient Network Storage Platform (iNSP), the Fujitsu Eternus VS900 doesn't depend on external control-path appliances. An external management server is used to upload and change configurations as well as for monitoring, but it isn't required to communicate with the virtualization software on the switch during normal operation. The Eternus VS900 continues to operate properly even if the management server is unavailable.

  The Eternus VS900 currently works only on Brocade switches. "It was developed as a collaborative effort between Brocade and Fujitsu, similar to what EMC has done with Cisco," Fujitsu's DeCaires said. Like Invista and Incipient iNSP, the Eternus VS900 currently lacks advanced storage features like remote replication and thin provisioning.

- **Incipient Network Storage Platform (iNSP):** Incipient iNSP is very similar to EMC's Invista with a few distinct differences. First, Incipient supports only Cisco MDS switches and directors. Data-path processing is identical to that of EMC, except Incipient calls it FastPath Processor. The most significant difference to Invista is that all virtualization software runs within the Cisco SSM module and isn't split to have dependent code running on external appliances. This eliminates dependencies outside of the switch, making it an overall less-complex solution.

- **LSI StoreAge Storage Virtualization Manager (SVM):** EMC, Fujitsu and Incipient virtualization products all run on intelligent FC switches. As a result, these products are expensive to deploy and bind the virtualization to a switch vendor (which creates vendor lock-in). They also increase the complexity of the SAN. LSI acknowledges the benefits of fabric-based virtualization, but realized early on that these disadvantages would hamper acceptance. To address this, LSI offers two virtualization products: SVM control-path software and the LSI 8400 data-path fabric hardware device, which comes from an OEM deal with QLogic Corp.

The LSI 8400 provides the data path and control path but, unlike Brocade and Cisco switches, it only provides the switching features for virtualization. This makes the LSI 8400 more cost-effective and complements existing FC switches rather than replaces them. "The 8400 is a virtualization appliance with switching capabilities, but it's not a switch with the huge benefit that we can connect to any switch," LSI's Nahum said.

The LSI 8400 provides the data path and control path but, unlike Brocade and Cisco switches, it only provides the switching features for virtualization.

From an implementation perspective, the LSI 8400 gets connected to an existing FC switch and the 16 switch ports become part of two zones: one contains initiator ports, while a second contains target ports. When a server accesses a virtualized volume, traffic is forwarded to the designated target port on the 8400 through a standard FC switch. The 8400 then performs the virtualization lookup and forwards frames to the appropriate storage device through one of its initiator ports. As the LSI 8400 connects through other FC switches, it adds two hops because a standard FC switch forwards traffic to an LSI target port and then receives frames from an LSI initiator port; however, the added latency is negligible.

VIRTUALIZATION APPLICATIONS

Storage virtualization is employed to solve specific business problems, such as simplified management, cost reduction or a need for nondisruptive data migration between arrays. "No one buys storage virtualization for virtualization per se, but to fill a very specific need; this is very different from server virtualization," said Robert Infantino, Incipient's senior vice president of marketing and alliances.

DATA MIGRATION

Data migration is the leading reason users deploy storage virtualization. As the virtualization layer controls the virtual-to-physical mapping, a virtualization product can forward storage requests to the correct physical device even while data is migrated and spread between a source and destination device. All virtualization products, regardless of their

underlying architecture, support data migration. Performance of data migration services will vary among the different virtualization architectures; Hitachi Data Systems claims to have a performance advantage because its virtualization resides within the storage controller.

PROVISIONING AND VOLUME MANAGEMENT

Provisioning of virtualized volumes is a core service in all virtualization products. Besides simplified storage management, centralized provisioning through the virtualization software enables higher storage utilization because storage is provisioned more granularly. Similarly, the volume management feature in virtualization products is instrumental in increasing storage utilization. By aggregating multiple physical disks to present them as a single large disk, and disaggregating large disks to present them as multiple smaller volumes, storage can be managed more effectively.

THIN PROVISIONING

Thin provisioning is currently not available in fabric-based virtualization products. The lack of cache and the stateless nature of fabric-based virtualization make it more difficult to implement. Hitachi was the first vendor to offer thin provisioning in its USP V product, providing the first 10 TB of thin-provisioned storage for free. IBM added thin provisioning in the recently released SVC 4.3 along with space-efficient snapshots and virtual disk mirroring. LSI released thin provisioning in October 2008, and Incipient has it on its roadmap. Fujitsu said it will add thin provisioning in 2009, but EMC has no intention at this point to offer thin provisioning in its virtualization product.

"We currently support thin provisioning in our arrays and haven't decided if and when we will support it in Invista," said Doc D'Errico, vice president of EMC's infrastructure software group.

SNAPSHOTS

Snapshots or clones are supported by all virtualization vendors except Fujitsu. However, space-efficient snapshots (snapshots that require disk space for changes between snapshots) are currently supported only by IBM's SVC, Hitachi's USP V and LSI's SVM. Similar to thin provisioning, space-efficient snapshots are more difficult to support in fabric-based virtualization products. While EMC's Invista supports only full-copy clones, users have the option to deploy RecoverPoint in addition to Invista to take advantage of space-efficient snapshots. LSI's virtualization proves that the application challenges of fabric-based virtualization can be overcome. Not only does LSI offer space-efficient snapshots, but its SVM product supports consistency groups with snapshots, as well as copy and mirroring, which enables entire applications to be snapped or copied at once and then recovered within minutes.

THE ABILITY TO SCALE

Performance and scalability are the two main benefits of fabric-based virtualization. Switch-based virtualization products can be scaled vertically if the switch vendor supports multiple intelligent fabric modules in a single switch, and horizontally by deploying additional intelligent switches. For example, storage architects can populate a single Cisco MDS switch or director with multiple intelligent line cards, or intelligent line cards can be inserted into multiple switches. Scaling is also achieved by deploying newer generations of intelligent line cards.

"Cisco is already shipping the MSM-18/4, which is Cisco's second-generation intelligent card," said Rajeev Bhardwaj, Cisco's director of product management for its Data Center Business Group.

There's currently no perfect virtualization product, and users need to carefully weigh the product that best fits their environment's requirements. For companies that have standardized on Hitachi Data Systems storage, Hitachi's array-based virtualization is likely at the top of their lists. Companies with a Brocade- or Cisco-based SAN should consider Invista or Incipient. They'll also need to weigh the performance and scalability benefits against some of these products' challenges, such as their relatively high cost, complexities and feature constraints. LSI deserves consideration as it's currently the only vendor offering a virtualization product that combines the simplicity and rich feature set of an in-band appliance like IBM's SVC with the benefits of fabric-based virtualization. ⊙

Jacob Gsoedl is a freelance writer and a corporate director for business systems. He can be reached at jgsoedl@yahoo.com.

Storage Virtualization Gets Real.

WOULD YOU LIKE TO GET MORE OUT OF YOUR EXISTING STORAGE ASSETS — AND FEEL SAFE DOING IT?

Hitachi Storage Virtualization transforms your multi-vendor storage assets into a common pool of shared resources. It lets you reclaim, utilize, and optimize 'lost' space.

Reinvigorate your infrastructure and get more out of what you already have with Hitachi Storage Virtualization:

www.hds.com/virtualization

Gartner rates Hitachi a Strong Positive.

"Hitachi already offers the most mature virtualization solution in the high-end enterprise disk array market."

Gartner MarketScope for High-End Enterprise Disk Arrays, 2008

Hitachi Storage Virtualization.
Reclaim, Utilize, Optimize.

Hitachi Data Systems

HITACHI
Inspire the Next

Buying guide & vendor comparison

Most consolidation initiatives now include storage virtualization. Consider these purchasing criteria and vendor offerings when you decide to take the virtual plunge. *By Stephen J. Bigelow*

TODAY, FEW CONSOLIDATION INITIATIVES are complete without virtualization—the use of software to create a layer of abstraction between hardware and applications. Virtualization allows organizations to handle more work with less equipment. With virtualization, a physical server can be segregated into logical servers to make better use of available CPU, memory and I/O resources.

Virtualization can also pool storage from across the data center, allowing disks from diverse storage platforms to appear as a single storage resource. This storage resource can then be allocated, provisioned, migrated, replicated and backed up without regard to where it's physically located. Not only does virtualization make storage easier to track and manage, it also makes better use of available storage, forestalling the expense of additional disk or storage platform purchases.

Still, virtualization has its tradeoffs.

- Organizations need to contend with another layer of software.
- The software needs to be compatible across the infrastructure and maintained as patches and updates are released.
- Virtualization must scale without impairing performance.
- Storage practices such as backups and replication will need to be modified to accommodate a virtualized environment.

Here is a list of the criteria for purchasing virtualization software:

WHERE WILL THE SOFTWARE BE RUNNING? Storage virtualization can be implemented with host-based, array-based or fabric-based products. Host-based virtualization runs software, such as Symantec Corp.'s Veritas Storage Foundation, installed on host servers. Dedicated appliances offer hardware acceleration for network-based virtualization software. Host-based products are the least costly and easiest to deploy, but they're also the least scalable.

Fabric-based virtualization runs software on intelligent switch devices such as Cisco Systems Inc.'s MDS 9500 series of multilayer directors. Fabric-based storage virtualization typically promises the greatest level of heterogeneity and scalability, but may require a new switch in the infrastructure.

Array-based virtualization integrates the technology in the storage array itself, such as Hitachi Data Systems' Workgroup Modular Storage and Adaptable Module Storage arrays. However, array-based virtualization typically uses software from the storage vendor, and generally isn't heterogeneous across different storage systems. (For more on this subject, see "Pros and cons of three architectures," p. 7.)

INTEROPERABILITY WITH YOUR CURRENT INFRASTRUCTURE. Interoperability is a critical consideration for virtualization technology. A virtualization product should accommodate all of your existing storage hardware, and meet the demands of future storage systems. Storage systems the virtualization product doesn't support will often remain in service, relegated to secondary storage tasks. Unfortunately, non-virtualized storage "islands" tend to fall into disuse, which wastes the valuable space that virtualization tries to organize. Vendor support matrices are a good place to start evaluating interoperability, but in-house testing can also be used.

WHAT SUPPORTS THE VIRTUALIZATION SOFTWARE? To support the virtualization software itself, you'll need host device drivers, path managers, agents and shims. IT staffers can get bogged down patching and updating a proliferation of storage virtualization servers when hardware is replaced or new versions become available. But not paying enough attention to maintenance can result in version disparity, leading to instability and performance problems. Evaluate any storage virtualization product from management and maintenance perspectives, and determine if the problems it solves outweigh the new issues it introduces.

HOW SCALABLE IS THE VIRTUALIZATION LAYER? A virtualization product can only manage x amount of storage, and storage performance may suffer as the amount of storage grows. You should understand the tradeoff between scale and performance, as 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. Consider scaling right from the start to help identify unacceptable products.

HOW WILL YOUR STORAGE PROCESSES CHANGE? The goal of storage virtualization is to consolidate a variety of storage resources into a single ubiquitous pool. This will invariably change the way that storage is organized, provisioned, migrated and protected. For example, a virtualization product may provide automatic provisioning, which may be a significant change for the IT organization. Storage administrators will also need to change backup or replication targets once virtualization is in place. This is another area where lab testing and vendor support can prevent problems before they start.

DEPLOY STORAGE VIRTUALIZATION IN PHASES. Deploying virtualization across the entire enterprise at once is risky. Perform a thorough up-front lab evaluation of any storage virtualization products. This should include a review of decommissioning drills. Once you decide what to purchase, you can start implementation on a small scale before systematically building out the environment. This conservative approach gives administrators ample time to get accustomed to virtualization management and prevents unforeseen problems from crippling an entire data center.

EXAMINE RESOURCE MANAGEMENT FEATURES. Storage virtualization products incorporate a growing range of resource monitoring and management features. For example, a storage virtualization tool can see every storage I/O, allowing the tool to track disk use to view performance metrics and monitor path configuration. Few virtualization products offer the range and sophistication of features found in full storage resource management (SRM) packages, but users can gain insights into resource management and automation without committing to yet another piece of sophisticated software.

CAN ONE UNDO VIRTUALIZATION? Performance issues, scalability limitations and interoperability problems are all reasons for decommissioning a storage virtualization product. An organization may also decide to discontinue one product in favor of a better one. Unfortunately, there's no simple way to undo virtualization once it's been deployed. Eliminating a problematic virtualization deployment is disruptive and time-consuming. Discuss any back-out options with the vendor before committing to a virtualization product.

Stephen J. Bigelow is a senior technology writer at TechTarget.

Managing the information that drives the enterprise

STORAGE

The Essential Guide to Data Deduplication

data deduplication

Data dedupe can reduce the amount of disk required for backups by removing redundant data, but there are a few things you need to know before implementing this technology.

Exclusive Extra for Storage Subscribers

The Essential Guide to Data Deduplication is a brand new expert resource brought to you by the Storage magazine team. This guide outlines the fundamental aspects you need to know about data dedupe and provides a number of tips and best practices from some of the leading experts in the field. Get a handle on data deduplication with 24/7 access to this new Essential Guide in an easy-to-access, easy-to-read, easy-to-print online format.

Get Your FREE Copy Today: **www.SearchDataBackup.com/Dedupe_Guide**

Storage Virtualization AND DR

Including storage virtualization as part of a disaster recovery plan can make the recovery process easier and lower equipment costs, but it also adds another layer of management and complexity.

By Alan Earls

VIRTUALIZATION IS EVERYWHERE THESE DAYS, it seems. But what is less clear, is how virtualization fits into disaster recovery (DR). Storage virtualization has many immediate and obvious benefits that make it helpful for DR purposes. But the devil is in the details. David Russell, vice president for storage strategies and technologies at Stamford, Conn.-based Gartner Inc., said storage virtualization offers clear positives, but adds a layer of management and complexity. “In a DR situation, you could argue that the more streamlined things are, the better,” he said. On the other hand, if you want to improve your cost structure so you can finally afford to implement DR solutions, storage virtualization can be the answer.

At its core, storage virtualization simply mobilizes storage capacity—making the data more portable and, ideally, faster to deploy. Those attributes are very relevant to DR scenarios. Nowadays, Russell noted, organizations aren’t as likely as in the past to have homogeneous environments with similar resources on both production and recovery sites. “With storage virtualization you can more easily overcome those differences in underlying architectures or even in vendor equipment,” he said.

Furthermore, one of the key side effects is that you can lower the cost of a DR solution because you can have the option of using higher-tier storage equipment at the primary site and lower-tier equipment at the backup site. “It might not be ideal for every workload, but for DR it can be sufficient,” he explained.

David G. Hill, principal at the Mesabi Group LLC, said virtualized storage isn’t necessarily an advantage or a disadvantage in terms of DR. Theoretically, he noted, having one logical pool of storage (which storage virtualization gives you) should make the recovery process easier because there’s only one management process. However, although all applications are affected in a disaster, individual apps may need to be restored over different periods of time (which could be hours, days or longer) and in order of priority, such as bringing up revenue-producing applications first.

“If all of the virtualization pool is in one physical array, you’re probably indifferent between having or not having virtualization from the perspective of going online again,” he said. However, if you have more than one array, you might want to bring up the arrays in a certain order based upon what applications run on each. “That may or may not be easier if you are running non-virtualized storage,” Hill said.

"If all of the virtualization pool is in one physical array, you're probably indifferent between having or not having virtualization from the perspective of going online again."

—David G. Hill,
principal, Mesabi Group LLC

Of course, virtualization has implications for the operation of your disaster recovery site. According to Hill, the advantage of storage virtualization is that you don’t necessarily need the same brand of storage arrays at the DR site and the local site. And while you do need the necessary physical capacity to hold the data, the backup site may not require as much physical capacity as the local, primary site.

Moreover, he noted, you may be able to use more cost-effective arrays at the DR site (such as higher-capacity disks). But, he said, “the tradeoff is that there could be application performance degradation such as longer response times to user online requests if a failover to the DR site is necessary, because less-expensive arrays don’t have the same level of performance.”

While there’s no absolute requirement that a DR site must also be set up for virtualized storage if the primary site is, Hill said, in practical terms, it makes far more sense to have that kind of match. In the event of a failure, not only does the DR site have to pick up the workload, it has to be the model for rebuilding the original site (in either the original or a different location). If storage virtualization is used at the original site, that should be the preferred mode.

Hill said there are a few other key points to bear in mind when

considering a virtualized DR investment. For instance, virtualization has implications when it comes to replicating to a DR site. He noted that if a virtual tape library (VTL) is used with data deduplication, the ability to replicate is enhanced because less bandwidth is used.

On the other hand, from a storage virtualization perspective, if all of the data in a virtualized storage pool would have had to be replicated to a DR site anyway, it shouldn't make much difference.

Finally, "the overall management process might be easier when recovering using both virtualized storage and virtualized servers," Hill said. ⊙

Alan Earls is a Boston-area freelance writer focused on business and technology, particularly data storage.

Check out the following resources from our sponsors:

FalconStor Software

WAN-optimized Replication: Revolutionize Your IT Efficiency

Simplifying Disaster Recovery (DR): Replication for Combined Physical/Virtual Data Centers

Webcast: Leveraging Virtualization for Greater ROI

HITACHI
Inspire the Next

Switch IT On with FREE Storage Virtualization Software. Limited Time Offer. Learn More.

White Paper: Solving the Top Six Enterprise Storage Issues with Virtualization

IDC Whitepaper: Storage Economics -- Assessing the Real Cost of Storage

nexsan TECHNOLOGIES

Nexsan iSeries™: Simple, Affordable, Powerful Enterprise Class iSCSI SAN for Your Business

SANs and Server Virtualization Podcast

Nexsan iSeries™ -- Enterprise Class iSCSI SAN