Managing the information that drives the enterprise

STORAGE

Vol. 8 No. 10 February 2010

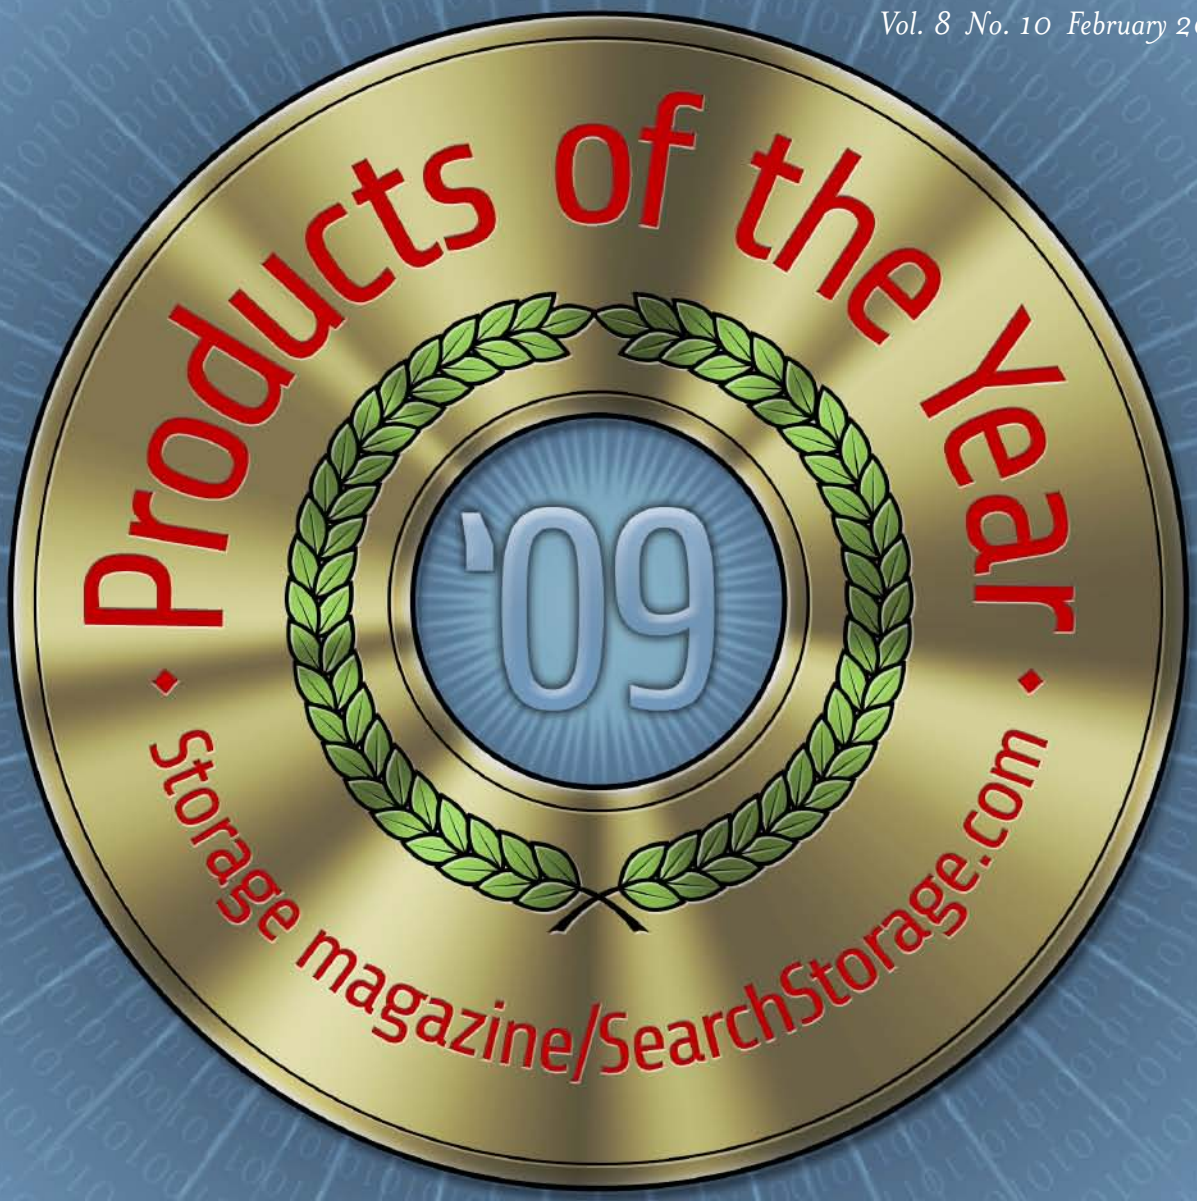

Our expert panel picks the top storage products of 2009 in five key categories. See what products take home the awards. P. 10

ALSO INSIDE

TechTarget
The Technology Media ROI Experts

HITACHI
Inspire the Next

pillar DATA SYSTEMS

syncsort
PERFORMANCE DRIVEN

tangent

Trend or Just Trendy?

2009 Storage Products of the Year

Get Control of NAS Systems

Storage Networks Get Virtual

10 Ways to Improve Data Protection in VMware Environments

SAS Drives Showing Up More and More

Vendor Resources

Cover illustration by ENRICO VARRASSO

Hitachi Data Systems

HITACHI
Inspire the Next

Storage has a new dimension.

Cloud enabled, multitier archiving for the next generation.

If you are considering cloud computing, it can be hard to see through the fog and clearly identify the cloud strategy that works for you.

Hitachi Data Systems and the Hitachi Content Platform delivers an integrated set of storage infrastructure to create a single pool of virtualized block, file and object storage, so you can adopt cloud services at your own pace.

For more information on Hitachi Content Platform and our cloud strategy, please visit:

www.hds.com/cloud

To arrange a discussion please call:

888-234-5601

www.hds.com/cloud

Hitachi Data Systems

Trend or just trendy?

What's in vogue in data storage today can become a vague memory tomorrow. Figuring out which new technologies have staying power tests the skills and instincts of enterprise data storage managers.

STORAGE blows hot and cold. Technologies that seem promising—even those that manage to rack up some measurable market share—can suddenly seem old hat.

Case in point: A lot of industry observers predicted that iSCSI storage systems would shake up the networked storage market and knock Fibre Channel (FC) off its throne, exiling FC arrays to a niche where only the well-heeled shop for expensive, high-performance storage gear. What we've actually seen with iSCSI is more of an evolution of data storage systems: a good idea that, instead of toppling the incumbent, has proven to be more of a gap filler and complement to the high-priced stuff. iSCSI adoption has been slow and steady; according to our surveys, approximately 40% of companies now have at least some iSCSI storage in place. But Fibre Channel still, undeniably, reigns supreme. In our most recent survey, where we asked storage managers about their plans for 2010, of those respondents who expect to buy storage systems, 62% said they'd go for FC while only 22% said they had iSCSI in their sights.

Fibre Channel didn't simply outrank iSCSI on that survey; there were four other storage types separating FC and iSCSI. If iSCSI is going to pose a threat to FC's hegemony, it's going to have to climb over multiprotocol arrays (40%), network-attached storage (NAS) (28%), direct-attached storage (DAS) (26%) and even solid-state storage (23%). So, is iSCSI going away? Nah, but maybe it's not the game changer that so many observers predicted.

And is solid state truly going to outrun iSCSI in 2010 as our survey suggests? I'd guess that in this case the 23% of respondents who said they'll deploy solid-state storage this year represent a lot of pilot projects, a handful of production implementations and—to steal a phrase from Alan Greenspan—just a bit of "irrational exuberance" about solid state in general.

Wide-area network (WAN) optimization is another example of the pace of change. It became the storage technology du jour a few years ago but, interestingly, it wasn't specifically intended for data storage environments. It was—and still is—primarily an application-enabling technology that just happens to work pretty well when used to back up remote offices by making it possible to send more data faster over

relatively thin pipes. But data deduplication now effectively accomplishes the same thing by paring down the data and limiting what needs to hit the network in the first place. As dedupe becomes more ubiquitous, with products that can fit into practically any part of the backup infrastructure, it seems that there will be less need for WAN optimization in backup environments.

And now deduplication is coming out from under the backup covers and showing up in primary data storage systems as well. Although there are only a few players right now, user interest is extremely high given the never-ending disk capacity battle, so 2010 should see more storage vendors getting into the game. Not long ago, another technology—thin provisioning—was seen as the most promising antidote to bloated disk capacities. Most array vendors, some grudgingly considering the impact it could have on their storage sales, now offer a thin provisioning option. Thin provisioning is a great idea based on the common-sense principle of only giving an application the storage it actually needs. However, the technology hasn't caught on like wildfire; its numbers are up on our storage priorities survey, with nearly twice as many respondents saying they've implemented it vs. a couple of years ago, but the already-deployed figure is still a fairly modest 19%. While 51% say they'll evaluate or implement thin provisioning this year, 60% say they have evaluation/deployment plans for primary storage dedupe. So one has to wonder if thin provisioning will take a backseat as the dedupe juggernaut rolls on.

Most array vendors, some grudgingly considering the impact it could have on their storage sales, now offer a thin provisioning option.

A lot of these hot-to-cold shifts in data storage are just the effects of the normal churn found in any technology industry. But even if they're just part of the natural selection process, your choices can have significant repercussions down the road. There's nothing wrong with iSCSI, solid-state storage, WAN optimization or thin provisioning—it's just that none of them is likely to solve all of your data storage problems or revolutionize your shop. Some technologies may seem like the answer to your prayers, but some will never live up to their hype. Finding the best balance of technologies for your environment is tough, but it's what makes managing storage so interesting. ⊙

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

** Click here for a sneak peek at what's coming up in the March 2010 issue.*

www.pillardata.com

Put an end to storage waste. Oversized and underutilized legacy storage systems are costing you way too much energy, floor space, and money.

Make the move to Pillar Axiom®. It's the world's most efficient storage machine, and the only true Application-Aware Storage™ system. Get the industry's highest utilization rate – up to 80%. Guaranteed. Dramatically cut power and cooling costs and floor space. Reduce TCO by as much as 50%. And increase user satisfaction immeasurably.

It's time to Stop Storage Waste and see how efficient you can be. **www.pillardata.com**

Download a complimentary white paper:
Bringing an End to Storage Waste.
www.pillardata.com/endwaste

pillar® DATA SYSTEMS

COMING IN MARCH

The New Backup

Backup applications have changed significantly over the last few years, incorporating capabilities that were previously available only in third-party applications. Features such as data deduplication, backup reporting, CDP and replication management, as well as interfaces designed specifically for disk-based backup and virtual server environments, have radically changed traditional backup applications.

Virtual Disaster Recovery

Virtualizing servers and storage can provide a level of flexibility that offers greater agility and responsiveness in case disaster recovery (DR) measures are needed. But virtualization can complicate matters. We look at the feasibility of an all-virtual environment for DR, including best practices, configuration tips, and suggestions on how to manage and maintain a virtual environment in your organization.

Quality Awards V: Enterprise Arrays

In this fifth edition of the *Storage* magazine/SearchStorage.com Quality Awards for enterprise-class arrays, users rank the leading systems for initial product quality, sales-force competence, product features, product reliability and technical support. See if past winners like EMC, Hitachi Data Systems and NetApp can repeat as leaders of the enterprise array pack.

And don't miss our monthly columns and commentary, or the results of our Snapshot reader survey.

***Storage* magazine**
117 Kendrick Street
Suite 800
Needham, MA 02494
editor@storagemagazine.com

Subscriptions:
www.SearchStorage.com

Ensure you're dealt a winning hand.

CIO Decisions Ezine provides CIOs and business strategists with expert and peer insight "from the trenches" on pressing issues decision-makers face every day, including:

- Project and Portfolio Management
- Virtualization
- Business Process Management
- Outsourcing
- And much more!

Subscribe Today!

CIO DECISIONS
ezine

2009 Storage Products of the Year

The best storage products of 2009 are highlighted by several repeat winners who keep improving their products, and by relative newcomers who are breaking ground with newer technologies.

By Rich Castagna, Todd Erickson, Chris Griffin, Beth Pariseau, Dave Raffo and Carol Sliwa

the results are in: 2009 produced a solid crop of new and upgraded enterprise data storage products. The Storage 2009 Products of the Year offers a mix of products that implement leading-edge technologies in innovative ways, and more traditional products that take tried-and-true technology to the next level.

The 15 awards presented here go to companies and products that stand out for their functionality, value, performance, innovation, ease of use and integration. In short, these products had the wherewithal to impress our panel of judges.

The technologies that these products advance cover a wide range—from those generating much of the current buzz, like data deduplication and solid-state storage, to bread-and-butter operations like data protection and storage-area network (SAN) storage.

As usual, we have a blend of older, more established vendors and those who can still be considered in startup mode. A number of our awards go to repeat winners—some of whom have won multiple awards—but a fair share also go to relative newcomers. That's a strong indication that successful enterprise data storage vendors continue to innovate even as the upstarts challenge them.

Congratulations to all of our *Storage* Products of the Year winners!

Backup and Disaster Recovery Hardware

gold: ExaGrid Systems Inc. EX10000E

ExaGrid's EX10000E sits at the nexus of two popular technology trends in 2009: scale-out network-attached storage (NAS) management for storage and data deduplication. Like ExaGrid's other products, the EX10000E offers global data deduplication for disk-based backup across multiple physical nodes in its clustered system, making it easily expandable while parallelizing performance.

Users have a choice of post-process or concurrent data deduplication, depending on how they prefer to balance the data deduplication process with backup windows. They also have the option of configuring the data to send it to tape or to replicate it offsite before dedupe is begun.

As data deduplication offerings have branched out from a backup focus to address primary data storage and other areas of the data center, and some data deduplication vendors have begun to broaden their target markets, ExaGrid Systems has remained focused on midmarket customers of backup-oriented data deduplication, working to add incremental upgrades to the product to support existing customers.

The company responded to another industry-wide theme in 2009: relentless data growth, regardless of company size, coupled with flat or declining budgets. With the release of the EX10000E, consisting of up to 10 10 TB server nodes for a total capacity of 100 TB, new customers can store more data in less space than with previous models. But customers with previously released smaller nodes won't have to rip and replace existing clusters; the new larger node can be mixed and matched with 1 TB to 5 TB nodes, with overall management via a single user interface.

Each node added to a cluster brings memory and processing power, boosting performance of data deduplication processing along with the additional capacity. In keeping with a theme of frugality in a down economy, the EX10000E consumes 50% less power than previous generations.

As one of our judges put it, "ExaGrid remains one of the few [vendors] thinking about how a system grows with the customer. Kudos on taking it to the next level."

Backup and Disaster Recovery Hardware

silver: Axxana Phoenix System RP

Many data storage administrators have begun to shift the focus of their disaster recovery (DR) plans from relying on the traditional method of sending tape backups offsite to newer methods that allow quicker recoveries. More DR plans now embrace technologies such as remote data replication, server and storage virtualization, and cloud storage services to advance recoverability of mission-critical applications.

Synchronous replication, which captures every change, and offsite disaster recovery have long been available to enterprises with the means and infrastructure to afford them, but in recent years we've seen a democratization of DR through products that help smaller enterprises get their hands on the technologies that can deliver short recovery time objectives (RTOs) and more aggressive recovery point objectives (RPOs).

Among those is Axxana, which offers what it calls "synchronous replication over asynchronous infrastructure." Its Phoenix System RP product, which first became generally available in the spring of 2009, combines EMC Corp.'s RecoverPoint asynchronous replication software for replication over distance with its own local synchronous recording mechanism. Should a disaster occur, the Phoenix System RP will send any uncommitted changes to the secondary site, bridging the gap between synchronous and asynchronous data protection.

To do this even if the entire primary site is wiped out, the Phoenix System RP wraps itself in a hardened "black box" that's resistant to flames, pressure, puncture, water pressure and electric shock, and has the ability to replicate those last bits of data using cellular signals if normal communications are down.

"I would hate to have to rely on cell infrastructure to move too much cached data in a DR event," one of our judges said. "Still, there is no magic solution to this problem and this is the most interesting idea as a solution."

Backup and Disaster Recovery Hardware

bronze: EMC Corp./Data Domain DD880 Enterprise Deduplication Storage System

A perennial winner in this category, Data Domain remains the most widely deployed data deduplication system on the market, a situation that's unlikely to change now that it's been brought under the wing of the industry's largest storage vendor, EMC Corp.

The nasty battle between EMC and rival NetApp Inc. this year with Data Domain as the prize, as well as the ultimate $2.1 billion acquisition price, are testament to Data Domain's popularity and the increasing importance of data deduplication as a table-stakes feature for storage backup vendors.

In the midst of the acquisition drama, Data Domain brought out its first quad-socket, quad-core processor system, and the increased CPU power upped aggregate throughput up to 5.4 TB per hour and single-stream throughput up to 1.2 TB per hour, with support for up to 71 TB of usable capacity. That's twice the throughput and capacity of its next largest box, the dual-socket, quad-core DD690. With the DD880, a DDX array with the maximum 16 controllers can perform at up to 86 TB per hour with 56 PB of usable capacity.

While plans for the DD880 were well underway before the acquisition took place, it's a sign of things to come from the new EMC subsidiary as it looks to bring dedupe to larger enterprises. The boost in throughput and capacity could alleviate contention issues enterprise users have experienced at times with smaller Data Domain boxes.

More performance and scalability are always needed in the data backup market, but with a starting list price of $400,000, the DD880 may not be affordable for some companies in the current economy, leading one of our judges to summarize the product this way: "Blistering performance remains. Blistering price remains." ⊙

Backup and Disaster Recovery Software and Services

gold: CommVault Systems Inc. Simpana 8

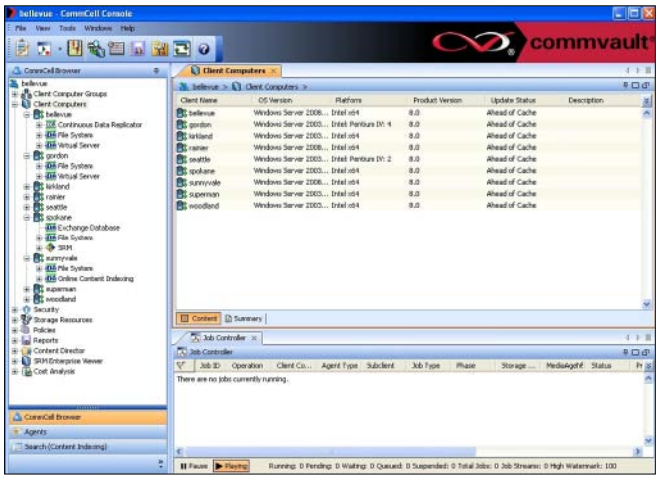

You can't accuse the folks over at CommVault of resting on their laurels. With this year's gold award for its Simpana 8 backup suite, CommVault is making its fourth trip to the winner's circle, having snared gold in 2002 (Galaxy 4.1), 2003 (QiNetix 4.2) and 2007 (Simpana 7.0 Software Suite)—not only a record for our Products of the Year, but probably one for product name changes, too.

CommVault's users are also impressed. The company has finished in the top spot three of the four times we've conducted our Quality Awards survey that reflects users' satisfaction related to quality, reliability and support.

CommVault continues to enhance its flagship product. Simpana 8 boasts a slew of new or beefed-up features. Version 8's advanced block-level data deduplication extends Simpana 7's capability by taking a global approach and building it into not only the backup process but the archiving module as well. In an industry first, Simpana 8 can dedupe data on tape to reduce the number of archive tapes a company needs to maintain. The product can also dedupe remote site backups and consolidate all deduplication results in a database.

Among Simpana 8's other enhancements are new data protection features for remote offices and mobile users. Continuous data protection can capture all writes to disk and then replicate those changes back to the data center. In addition, backups that are deduped at the remote site can be replicated to the central site. Remote data protection can extend to the edge. The Workstation Backup Agent can be installed remotely on PCs where it monitors writes and sends only changed blocks to a central server from which the data can be compressed and encrypted, and backed up or archived to a central site. CommVault says a single server can support hundreds of laptops and desktops, and end users can use a Web interface to do their own file recoveries.

CommVault's integrated toolkit approach to data protection has turned a lot of users into fans. One of our judges said Simpana 8 is "an incredibly functional product that integrates more into the complete suite than most vendors even offer as separate products."

Backup and Disaster Recovery Software and Services

silver: Continuity Software RecoverGuard v4.0

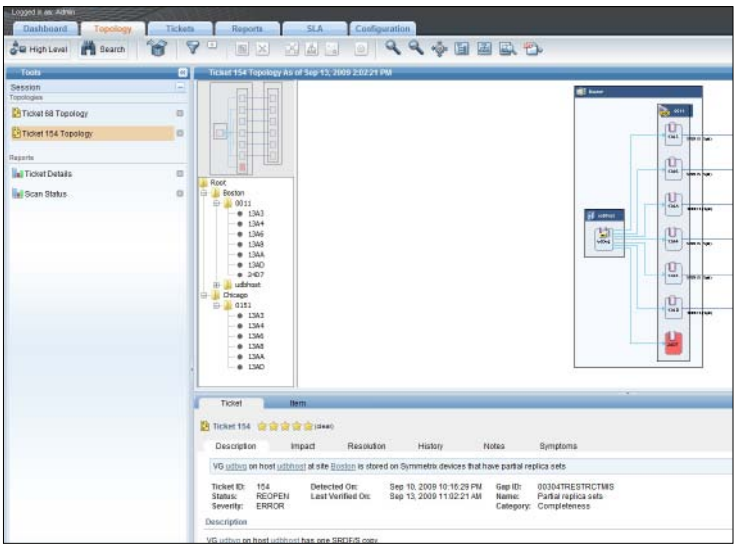

Getting a disaster recovery plan right requires testing, which can be an expensive, time-consuming and disruptive process. Some companies roll the dice and skimp on testing, but others have found that DR testing tools, like Continuity Software's RecoverGuard, can provide a measure of assurance to their DR plans.

Continuity was one of the pioneers in the still fairly narrow field of DR testing applications, and was also a silver award winner in 2007 for Version 2.0 of RecoverGuard.

RecoverGuard provides a sanity check for a DR plan as it detects potential gaps or conflicts in the processes that a recovery would depend on, such as incomplete replication configurations, or inadequate or incompatible remote site gear. To do this, RecoverGuard uses agentless technology to inventory an environment and then continually monitors the status to flag any changes that could knock a DR plan out of kilter. With that information in hand, the odds of crafting an effective disaster recovery plan vastly improve, or as one judge put it: "What better way to do DR is there?"

RecoverGuard 4.0 adds to its long list of capabilities with new support for high-availability (HA) clustered environments, enabling it to detect gaps in these often complex configurations. Version 4.0 supports HA clustering products from Hewlett-Packard (HP) Co., IBM Corp., Linux, Microsoft Corp., Oracle Corp. and Sun Microsystems Inc.

While RecoverGuard could always detect and report vulnerabilities, its new root-cause capability can now provide more specific remedial advice related to infrastructure changes that could jeopardize a recovery. Configuration "drift" (when systems slip out of synch) can also be detected for a variety of hardware and software conditions, and against desired service levels.

A RecoverGuard 4.0 license to protect 30 servers costs approximately $30,000. For approximately $1,000 a server that's a small price to pay to help ensure a DR plan is up to date and comprehensive. Continuity Software doesn't pitch RecoverGuard as a replacement for testing, but the app can certainly make testing easier and less costly.

Backup and Disaster Recovery Software and Services

bronze: AppAssure Software Inc. Replay 4

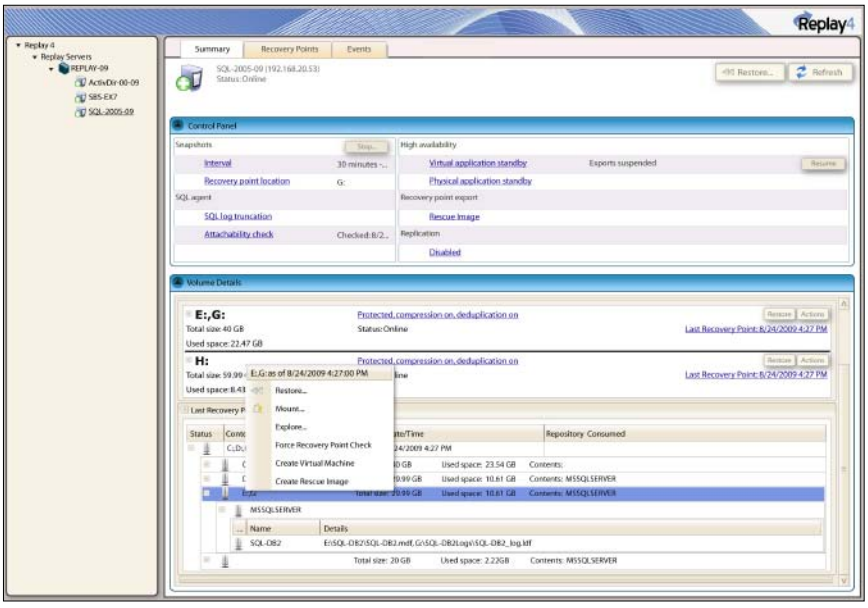

While only five years old, AppAssure is poised to make some waves in the small- and medium-sized business (SMB) backup market with its low-cost yet full-featured Replay 4 near-continuous data protection (CDP) suite. With an entry-level price of $1,079 per server, Replay 4 provides sophisticated features that are well within the reach of many smaller businesses. One judge called it "highly innovative" and said it "competes very effectively with [Symantec Corp.'s] Backup Exec, Acronis Inc.'s Backup & Recovery and other SMB products."

AppAssure touts the speed of Replay 4, which it says has been ramped up in this release, and the quick initial installation that requires only a few mouse clicks. Replay 4 is for Windows Server environments where it leverages Volume Shadow Copy Service (VSS) for its near-CDP capability and full-image capture feature. Replay's replication option allows sending compressed, deduplicated backup copies to a remote location or cloud storage service.

The company offers a variety of Replay versions, including those tailored specifically for three key Microsoft apps: Exchange Server, SQL Server and SharePoint. Version 4 adds support for Microsoft Hyper-V and VMware ESX virtualization platforms, in addition to Windows Server 2008 R2. Replay for Windows Desktops was also released with the Version 4 rollout. The desktop product can be managed centrally and will automatically copy changes and create bare-metal images for full system recoveries. Backups can be scheduled based on policies, and end users can recover their own backed up files.

Replay 4 is available on a long-term license or by subscription starting at $59 per month per server. A try-before-you-buy option lets users road test Replay 4 and other AppAssure products with free downloads of 30-day trial versions available on the company's site. ⊙

Disks and Disk Subsystems

gold: 3PAR Inc. InServ F400 Storage Server

3PAR's InServ enterprise data storage systems have always won high marks for their thin provisioning and management features, but the vendor's midrange offerings never gained much respect until the F-Class came out last April. The InServ F-Class Storage Servers have the same architecture and ASIC as the InServ T-Class enterprise system that launched in late 2008. That includes mesh-active quad-controller technology, thin provisioning ASIC and management software. The only difference is that the F-Class doesn't scale as high.

The F400 is the larger of the two F-Class arrays, with four controllers and 16 drives. 3PAR was a thin provisioning pioneer, and its latest ASIC with "Thin Built In" technology includes Thin Conversion and Thin Persistence features to help convert fat volumes to thin and keep volumes thin over time.

Our judges were impressed with the F400's features for a midrange system, giving it high scores for performance, ease of use and functionality despite finding it a bit pricey (starting at $86,000 for 2.3 TB of capacity).

"A good quality system that extends 3PAR's line somewhat downward ... good functionality with some innovation," one judge said.

Another added: "Expensive for mid-tier but a nice upgrade from their older [midrange] line. Still in many ways a leader in the thin storage market."

Disks and Disk Subsystems

silver: Dell EqualLogic PS6000 Series

EqualLogic had its foot firmly in the iSCSI door when iSCSI was still considered an ugly step-child of block storage, and long before Dell Inc. spent $1.4 billion to acquire EqualLogic and turned it into the iSCSI market leader. The PS6000 extends the Dell midrange array offerings by adding solid-state drive (SSD) support, more cache, faster processors, and enhanced management and data protection software.

The PS6000 represents a hardware redesign, and includes new eight-drive and 16-drive SSD modules, software integration with Microsoft Hyper-V, a new SAN Headquarters performance monitoring tool and a modified support structure. SAN Headquarters helps administrators track performance data over multiple groups of SANs instead of just within one local SAN group.

Judges were most impressed with the Ethernet-based system's ease of use and integration while noting its new features. "More innovative than previous generations," one judge said of the latest EqualLogic system. "Still a very fine, easy-to-use iSCSI SAN."

Disks and Disk Subsystems

bronze: Avere Systems Inc. FXT Series

Of all the solid-state drive products to hit the market in 2009, startup Avere Systems' FXT Series caught our judges' attention the most. Avere's appliances combine clustered NAS, file virtualization and automatic data placement with SSDs. "One of the best uses of solid state to date," one judge noted.

That means it's not just an SSD product, but rather it uses solid-state drives to solve customer pain points. FXT Series nodes are available in two models: the FXT 2300 and FXT 2500. Each FXT Series node contains 64 GB of read-only DRAM and 1 GB of battery-backed NVRAM. The FXT 2300 has a list price of $52,000, and includes 1.2 TB of 15,000 rpm SAS drives. The FXT 2500, at $72,000, comes with 3.5 TB of SAS disk. The nodes can scale out under a global name-space, and Avere said it has tested up to 25 nodes in a cluster.

The Avere OS also dynamically organizes data into storage tiers, automatically placing active data on FXT appliances.

"A very innovative application of file virtualization clustering, SSDs and automated tiering," another judge said. Pointing to its ability to attach to a third-party NFS NAS for data archival and backup storage, the judge added: "It has strong value with a high reuse of installed NAS systems." ⊙

Storage Management Tools

gold: Akorri Inc. BalancePoint 3.0

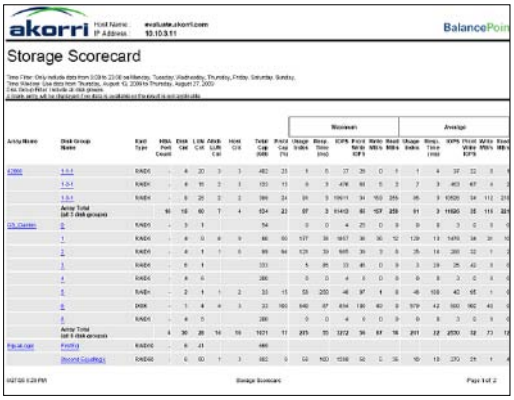

The general trend with storage management tools has been a shift to agentless monitoring and more end-to-end monitoring, from the applications and virtual servers to the storage arrays, to better troubleshoot performance problems.

Akorri Inc.'s BalancePoint has been ahead of the curve in agent-free monitoring, and the latest version of the product adds analysis for another piece of the infrastructure, SAN switches from Brocade Communications Systems Inc. and Cisco Systems Inc., as well as richer metrics for virtual machines and storage.

"Adding SAN switches is something they have needed to complete the picture," said one judge. Another said the tool "sets the bar for holistic capacity and performance-oriented infrastructure management."

Rich Corley, Akorri's founder, chief technology officer (CTO) and vice president of engineering, said BalancePoint, which also won gold in 2007, initially focused on managing the virtual storage infrastructure. But the company expanded the product's scope to provide greater end-to-end visibility into mapping and layout, as well as performance and availability, from the virtual machines to the arrays.

BalancePoint plugs into an Ethernet network and collects and stores information at 15-minute intervals using industry-standard protocols such as the Storage Management Initiative Specification (SMI-S), Windows Management Instrumentation (WMI) and Secure Shell (SSH), wherever possible, and going to the API level when necessary.

The tool delivers a visual representation that shows how infrastructure pieces interact with each other, analysis of how workload and utilization levels affect application performance, and actionable recommendations.

New troubleshooting and planning analysis metrics include Virtual CPU Efficiency (to identify and resolve CPU bottlenecks by comparing the guest OS with actual virtual machine usage) and a VM Performance Index (a physical-to-virtual migration metric to enable optimal sizing of each VM for its application).

BalancePoint 3.0 supports most major storage systems, with this version adding support for 3PAR's InServ Series. Version 3.0 also adds support for VMware Inc.'s vSphere 4.

Entry-level list pricing to manage 5 TB of storage and all attached applications and servers is $8,750. Pricing is based on terabytes of storage managed.

Storage Management Tools

silver: Seanodes Inc. Exanodes Virtual Machine Edition

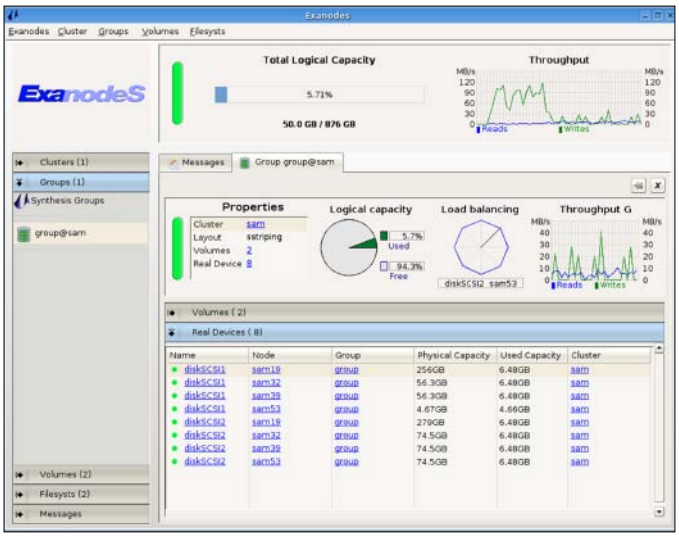

Virtual server technology remains wildly popular among IT organizations, but to take advantage of many of the advanced features, users need to have networked storage in place.

Seanodes Inc.'s Exanodes Virtual Machine Edition aims to remove the complexity and expense for cash-strapped companies with its innovative approach to creating shared storage, without requiring additional hardware or an external SAN or fabric.

Exanodes VM Edition software can turn the internal disks or direct-attached storage (DAS) of VMware Inc. ESX Servers into a virtual iSCSI SAN. That, in turn, makes it possible to take advantage of advanced VMware features, such as VMotion and Consolidated Backup, without having to invest in networked storage.

"Low cost/high reliability makes this a cost-effective solution," commented one judge. Another said the product "should be very attractive to the SMB market."

The Exanodes software creates a shared storage system using a server's typically underutilized hard disks. It uses the network interconnect as a data path and the nodes' memory and CPU to process storage access requests. The software integrates with server virtualization products to extend their capabilities to the internal disks of application servers, converting DAS into networked storage. The clustered parallel design and Redundant Array of Independent Nodes (RAIN) architecture help on the performance front. I/O processing is carried out simultaneously on a large number of nodes.

Seanodes notes that Exanodes provides file systems with volumes accessible in block mode, and the logical unit numbers (LUNs) generated by Exanodes have the same features as a LUN generated by a traditional storage array. The virtual array can be partitioned with partitions allocated to specific nodes or masked.

The product supports all types of storage devices accessible in block mode, such as SATA or SCSI disks, and software or hardware RAID. Users manage Exanodes Virtual Machine Edition through the command-line interface or a GUI-based console. Pricing for Exanodes VM Edition is $1,425 per ESX Server.

Storage Management Tools

bronze: Bycast Inc. StorageGRID 8

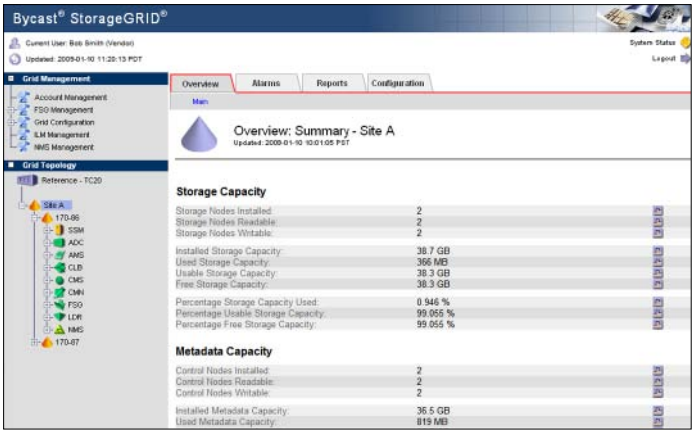

IT shops that are wary of the security or performance of cloud storage might take heart at some recent technology improvements. Bycast Inc.'s StorageGRID 8 software, which was traditionally strong in the medical archiving market, takes aim at the business and technical needs of private and public cloud storage providers. Key enhancements include discrete security partitions to enable service providers or IT departments to support multiple customers, or tenants, through a single system, as well as a Distributed File System Gateway (DFSG) to complement the product's distributed multitier object store.

Past StorageGRID technology supported NAS gateways for access to its storage systems, which make use of metadata and policies to organize data over geographically distributed locations. Version 8 uses a clustered file system to support *n*-way clustering of the gateways to boost performance.

The DFSG clustered file system front-end, coupled with StorageGRID's distributed archive, can help users scale their environments for a variety of applications. The product integrates a wide range of data storage assets, from high-performance disks to deep archive devices, in a single system that applications access through a standard file system interface.

Other new features include support for VMware Inc. ESX servers, and enhanced audit and reporting capabilities to enable storage service providers and IT departments to implement per-transaction, per-byte or per-asset chargeback models.

"The ability to segregate and audit data is critical," commented one judge. Another judge, noting the product's "excellent functionality," said StorageGRID is "a bit pricey for the market but probably worth it."

Licenses are sold on a perpetual basis or pay-as-you-go model. Perpetual license fees are tiered and based on managed capacity. Pay-as-you-go pricing has lower upfront fees; providers pay more as their capacity grows, but the price per terabyte decreases as system size increases. ⊙

Storage Networking Equipment

gold: Storspeed Inc. SP5000 Application-Aware Cache

Storspeed Inc.'s SP5000 application-aware cache is a gold award winner because of its ability to create a virtual performance layer and selectively apply administrator-created performance profiles to users and applications. The SP5000 started shipping in October 2009, and its $65,000 base price includes 80 GB of dynamic RAM (DRAM), eight bays for flash-based solid-state drives, a storage-usage metrics application and storage optimization tools.

Our panel of judges gave the SP5000 its highest marks when considering the product's innovation. Its ability to ensure that key users and applications will receive adequate resources allows administrators to confidently scale and tier data storage resources and use lower-cost capacity for archival data.

The SP5000 uses fast-packet inspection to determine how to populate a network cache and deliver predictable, high-performance storage to users and their applications regardless of the back-end storage infrastructure. It's standards based and invisible to applications, storage and network infrastructure because it installs between applications and network storage.

One judge commented that "NAS apps are often latency sensitive, so if [the SP5000] can really reduce round-trip time [RTT] without burning tons of CPU on the filer for key apps it will pay for itself." Another judge thought that it's a "great use of SSD."

Storage Networking Equipment

silver: QLogic Corp. 8100 Series Converged Network Adapters

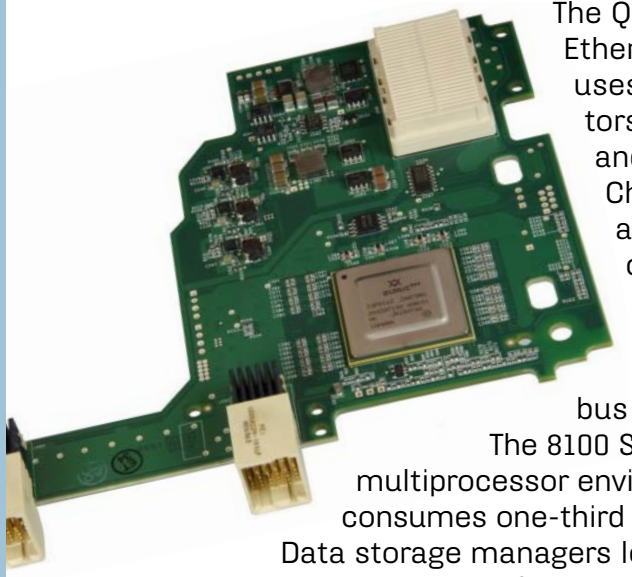

The QLogic Corp. 8100 Series Fibre Channel over Ethernet (FCoE) converged network adapter (CNA) uses a single chip set to allow storage administrators to extend server and storage consolidation and virtualization projects by converging Fibre Channel and Ethernet networks. The adapter's architecture integrates five previously disparate components into one ASIC, handles storage and data networking traffic at 10 Gbps, and appears to servers as a dual-port Ethernet network interface card (NIC) and a dual-port FC host bus adapter (HBA).

The 8100 Series is also specifically designed for multi-core, multiprocessor environments because it generates less heat and consumes one-third the power of existing CNA chip sets.

Data storage managers looking ahead to FCoE integration will appreciate the 8100 Series' ability to full hardware offload for protocol processing. One judge commented that QLogic's second-generation PCIe adapter is "one of the best CNAs on the market." The same judge added that the 8100 is "easier to integrate than most because of built-in driver compatibility." And another judge on our panel called it the "standard of HBAs" and "a great idea."

Storage Networking Equipment

bronze: VirtenSys Inc. VIO 4000 Series I/O Virtualization (IOV) Switch

The VirtenSys VIO 4000 Series made it into the storage networking top-three list by consolidating, virtualizing and dynamically sharing standard Ethernet NICs, Fibre Channel NICs or SAS controllers to relieve I/O bottlenecks. The VIO 4000 extends the internal PCI Express bus out of the server and moves existing off-the-shelf adapter cards into one top-of-rack unit or blade module, thereby replacing multiple networking and storage aggregation switches. In blade environments, the I/O virtualization modules incorporate directly into the blade chassis. The VIO 4000 also dramatically reduces the number of tasks, cabling, and power and cooling resources needed to manage and provision servers.

The VIO 4000 IOV switch provides up to 40 Gbps of bandwidth to each server, and appears as a physical I/O adapter to conventional or virtualized environments without requiring any changes to current software, firmware, physical connections or management tools. The VirtenSys IOV switch received the highest overall innovation marks from our judges in the storage networking equipment category.

Administrators can use a command-line interface or the Web-based GUI to dynamically configure and change I/O bandwidth based on workloads, service-level agreements or other criteria without any physical changes. With server and storage virtualization well entrenched, VirtenSys is one of a handful of companies now offering the benefits of virtualization for storage network infrastructures. ⊙

How we selected the Products of the Year

SearchStorage.com and *Storage* magazine invited enterprise data storage product companies to nominate their products for the 2009 Products of the Year awards. For previously available products, the upgrade must have incorporated major new features that were designed to meet new or evolving needs in the marketplace.

Products could be entered in five categories: backup and disaster recovery hardware, backup and disaster recovery software and services, disks and disk subsystems, storage management tools and storage networking equipment.

Products were judged by a panel of users, analysts, consultants, and *Storage* magazine and SearchStorage.com editors. Products were rated on six criteria: innovation, performance, ease of integration into environment, ease of use and manageability, functionality and value.

2009 STORAGE PRODUCTS OF THE YEAR FINALISTS

BACKUP AND DISASTER RECOVERY (DR) HARDWARE

- Axxana Phoenix System RP
- Barracuda Networks Inc. Barracuda Backup Service
- EMC Corp./Data Domain DD880 Enterprise Deduplication Storage System
- ExaGrid Systems Inc. EX10000E
- i365, A Seagate Company EVault Plug-n-Protect
- Spectra Logic Corp. BlueScale Version 10.6

BACKUP AND DISASTER RECOVERY SOFTWARE AND SERVICES

- Acronis Inc. Backup & Recovery 10
- AppAssure Software Inc. Replay 4
- CA ARCserve Backup r12.5
- CommVault Systems Inc. Simpana 8
- Continuity Software RecoverGuard v4.0
- Double-Take Software Inc. Double-Take Backup Version 5.2
- FalconStor Software Inc. FalconStor File-interface Deduplication System (FDS)
- IBM Corp. Tivoli Storage Manager (TSM) 6
- Vizioncore Inc. vRanger Pro 4 Data Protection Platform (DPP)
- VMware Inc. VMware Data Recovery (VDR) 1

DISKS AND DISK SUBSYSTEMS

- 3PAR InServ F400 Storage Server
- Avere Systems Inc. FXT Series
- BlueArc Corp. Mercury
- Dell EqualLogic PS6000 Series
- EMC Corp. Symmetrix V-MAX
- Fusion-io ioDrive Duo
- Pliant Technology Inc. Lightning LS 300 Enterprise Flash Drive (EFD)
- Sun Microsystems Inc. Sun Storage 7410 Unified Storage System
- Texas Memory Systems Inc. RamSan-20

STORAGE MANAGEMENT TOOLS

- Akorri Inc. BalancePoint 3.0
- Aptare Inc. StorageConsole 7.0
- Bycast Inc. StorageGRID 8
- DataCore Software Corp. SANsymphony 7
- ParaScale Inc. ParaScale Cloud Storage (PCS) Software, Version 1.4
- Seanodes Inc. Exanodes Virtual Machine Edition
- Tarmin Technologies Inc. GridBank 1.5
- Virtual Instruments Corp. VirtualWisdom Virtual Infrastructure Optimization Service
- Zetta Inc. Enterprise Cloud Storage

STORAGE NETWORKING EQUIPMENT

- Brocade Communications Systems Inc. DCX-4S Backbone Switch
- Dataram Corp. XcelaSAN
- NetApp Inc. Performance Acceleration Module (PAM) II
- QLogic Corp. 8100 Series Converged Network Adapters
- Storspeed Inc. SP5000 Application-Aware Cache
- VirtenSys Inc. VIO 4000 Series I/O Virtualization Switch

Your One Stop Shop for All Things Security

Nowhere else will you find such a highly targeted combination of resources specifically dedicated to the success of today's IT-security professional. Free.

IT security pro's turn to the TechTarget Security Media Group for the information they require to keep their corporate data, systems and assets secure. We're the only information resource that provides immediate access to breaking industry news, virus alerts, new hacker threats and attacks, security standard compliance, videos, webcasts, white papers, podcasts, a selection of highly focused security newsletters and more — **all at no cost.**

Feature stories and analysis designed to meet the ever-changing need for information on security technologies and best practices.

INFORMATION SECURITY®

www.SearchSecurity.com

Breaking news, technical tips, security schools and more for enterprise IT professionals.

SearchSecurity.com

www.SearchSecurity.com

Learning materials geared towards ensuring security in high-risk financial environments.

SearchFinancialSecurity.com

www.SearchFinancialSecurity.com

UK-focused case studies and technical advice on the hottest topics in the UK Security industry.

SearchSecurity.co.UK

www.SearchSecurity.co.UK

Information Security strategies for the Midmarket IT professional.

SearchMidmarketSecurity.com

www.SearchMidmarketSecurity.com

Technical guidance AND business advice specialized for VARs, IT resellers and systems integrators.

SearchSecurityChannel.com

www.SearchSecurityChannel.com

TechTarget®
The IT Media ROI Experts

Get control of NAS systems

With file storage sprouting up like weeds, data storage shops are grappling with managing multiple disparate NAS systems. But you can fight NAS sprawl with a number of technologies.

By Marc Staimer

AT THE END OF 2008, Framingham, Mass.-based research firm IDC reported that for the first time ever, more data was stored on network-attached storage (NAS) systems or filers than on storage-area network (SAN) storage. In addition, IDC's more recent forecasts predict an acceleration of this trend. It's not only the number of files growing, but their size as well.

All of this translates into more installed NAS systems. Adding more NAS systems is an understandable reaction to file growth as network-attached storage systems are typically self-contained and preconfigured for rapid installation, and are easy to implement, operate, manage and use. But most traditional NAS systems are also silos, so they contribute to NAS sprawl. The consequences of NAS sprawl can be summed up by the often-repeated adage, "I loved my first NAS filer, I really liked my second, but by my tenth I was pulling my hair out."

FIVE WAYS NAS SPRAWL CAUSES PROBLEMS

NAS sprawl generally creates five major IT challenges (these are the biggies; there are others as well). All of them are complicated by the limited number of tasks a data storage administrator can complete in a given time, and they're all pretty difficult.

1. **System management.** Even though NAS management is far simpler than SAN storage management, it still requires some care, feeding and time.
2. **Managing client and application access to data.** Each NAS system must be mounted on every server and workstation that requires access. Mounts are application disruptive so they require scheduling downtime for the server applications. With more NAS systems you have more mounts, and that adds up to more scheduled downtime.
3. **File location.** Policies for file placement must be set based on performance, accessibility, age, access frequency, storage cost, availability, data protection and so forth. Policy setting is the easy part, but actually moving the files to the appropriate NAS system is a time-consuming manual data migration process. And it's an ongoing one. When the migration is done, the originating application must be re-pointed at the correct NAS system; this isn't such a big deal with a couple of NAS systems, but it's compounded as NAS systems are added.
4. **NAS load balancing.** Load balancing is required to get better utilization or to meet applications' performance requirements. Because load balancing is also a manual process to set up and manage, it becomes a major time sink even if you have identically configured NAS boxes.
5. **Protecting, replicating and/or backing up files.** Different NAS systems have different methods for snapshots, continuous data protection (CDP), mirroring and replication. Some are well integrated with common backup vendors, such as Windows Volume Shadow Copy Service (VSS), VMware or Citrix Systems Inc.'s XenServer, but others aren't. So there are more tasks requiring more time, training and experience. Even with identical NAS systems, they still require separate touch points for each data protection setup, operation and management.

These challenges get more difficult, take more time and make it more likely that errors will occur as NAS sprawl grows.

TECHNOLOGIES THAT CAN HELP WITH NAS SPRAWL

The industry took note and recognized this sober situation. The result is the current availability of four technologies designed to solve some or all of these challenges, albeit in completely different ways. They include: operating system built-ins such as Microsoft's Distributed File System (DFS) for CIFS as well as Linux/Unix automounters for NFS; file virtualization systems; clustered NAS systems; and private cloud and grid storage. A brief analysis of each of these technologies illustrates what they do and don't do to meet the aforementioned challenges.

Operating system built-ins

Microsoft Distributed File System (DFS) is part of Microsoft's Windows 2003 and 2008 server operating systems; DFS was developed for the small- and medium-sized business (SMB) Windows-only (CIFS) market. DFS Namespaces enables multiple file servers' shared folders to be grouped into one or more logical namespaces. Users see the namespace as a single shared folder and are automatically connected to shared folders in the same available Active Directory domain services site. This sidesteps the need for LAN or WAN routing. DFS Replication can automatically synchronize folders between local file servers or remote NAS systems on a wide-area network.

Pros

- Easy integration with Windows environments
- Familiar to Windows administrators
- No additional licensing costs
- Low upfront total costs
- Solves user access, mount/unmount, load balancing, data migration as well as some data protection challenges of multiple CIFS NAS systems

Cons

- Requires a relatively high level of Windows expertise
- Loose file synchronization among different servers, especially when geographically dispersed; a user at a remote location may access a file before it's updated
- Only works with CIFS (not NFS)
- Limited scalability; not architected to scale to large numbers of file servers
- Doesn't provide file-level granularity
- Doesn't work with non-Windows-based NAS systems
- Poor storage utilization because of large numbers of duplicate files
- Can require additional hardware infrastructure to meet performance requirements
- Doesn't solve issue of managing multiple NAS or filer systems; doesn't address data migration and many of the data protection challenges

Linux/Unix automounters are intended for NFS users. Automounters mount and unmount directories from other systems on the network as they're needed. They get their mounting instructions from centralized maps, which can be flat files, NIS maps or sections of an LDAP directory. Automounters are far easier to use than managing multiple static NFS mounts. Automounter advantages are readily apparent when there's a service failure. If a remote file server becomes unavailable, an automounter will simply time out and unmount the directory without

alarming users. With static NFS, mounts will hang until the file server is back up and running again.

Pros

- Easy integration with Linux and Unix environments
- No additional licensing costs
- Low upfront total costs
- Eliminates server hangs when the mounted service fails
- Works in conjunction with most NFS NAS systems
- Solves many of the user and mount issues with multiple NFS NAS systems

Cons

- Requires significant Linux or Unix expertise
- Not easy to set up
- Only works with NFS
- Lacks built-in replication capabilities
- Doesn't work with Windows (CIFS)-based NAS systems
- No file-level granularity
- Doesn't solve many of the management problems with multiple NAS systems

File virtualization systems

File virtualization systems separate the physical location of a file from the representation of that file. File virtualization systems essentially eliminate the requirement for a user or application to know exactly where their files are stored as they see only a single global namespace (GNS.) Depending on how it's implemented, file virtualization allows transparent file access, load balancing, data storage tiering, file migration, and even snapshots and replication for multiple homogeneous or heterogeneous NAS systems.

File virtualization implementations can usually leverage Microsoft's DFS and/or Linux/Unix automounters by acting as a management layer. This allows them to automatically update the DFS Namespace to include NAS filers and file servers, while also providing common management for multiple dissimilar NAS systems. F5 Network Inc.'s ARX file virtualization appliance also provides available disk space monitoring, while others (Avere Systems Inc.'s FXT Series and EMC Corp.'s Celerra NS with FAST) provide storage tiering.

No additional software is required to leverage DFS Namespace and Linux/Unix automounters. If the file virtualization technology fails, the file maps for Windows and mounts for Linux/Unix remain intact, allowing users and applications access to their files. Not all the file virtualization systems work with DFS or automounters, and some that do don't necessarily require them.

There are two types of file virtualization products: shared path and split path.

Shared-path file virtualization systems share the control and data path, which means that all connections to the NAS and all data to/from the NAS flow through the virtualization system. Shared-path file virtualization systems are full proxies that touch every file and every packet in the path before it's written or read.

Pros

- Allows files to be migrated in real-time even when in use; the file virtualization system updates the global namespace with the new physical location of the file
- Intuitive
- Easy to operate
- Protects current investment
- Transparent retirement of older NAS or file systems
- Individual file-level granularity
- Heterogeneous NAS and/or file server support; eliminates NAS system lock-in
- Definable policies using file metadata such as file type, creation date or when last accessed

Cons

- Added latency to pass through file virtualization system can be a bottleneck affecting response times and IOPS
- Single point of failure; a dead-box failure cuts off all access to the NAS and/or file systems
- Scalability is limited by the throughput of the shared-path file virtualization system

Split-path file virtualization systems separate the control and data paths, so the NAS connections and all data to/from the NAS don't pass through the file virtualization system. Split-path file virtualization is typically deployed as an x86 appliance connected to the LAN switch. They manage the namespace to direct files to the appropriate NAS or file system without intercepting any packets.

Pros

- Nondisruptive implementation for applications/users
- Highly scalable
- File virtualization system failure won't cut off access to data
- Protects current investment in NAS and file systems
- Relatively easy file migration
- If it uses Microsoft DFS for the namespace, DFS will always have the most recent namespace configuration allowing users and applications to access their files

- Heterogeneous NAS support
- Easy to operate

Cons

- Usually requires agents on application servers and workstations for transparent file migration; agents must be managed and maintained
- Tends to be Windows (CIFS) focused with limited NFS support

Shared-path and split-path systems are typically mutually exclusive. But EMC's Rainfinity is primarily a split-path system except when moving files when it's configured as shared path. That eliminates the need for split-path agents for file migrations and the shared-path scalability, performance and single-point–of-failure issues.

Shared-path systems include Avere Systems' FXT Series and F5 Network's ARX series, and EMC's Rainfinity when performing data migration. Split-path options include AutoVirt Inc.'s AutoVirt 3.0 and EMC's Rainfinity.

File virtualization systems have continued to evolve, solving more network-attached storage sprawl issues. Avere Systems' FXT automates NAS storage tiering by hosting the most active files requiring the highest performance on its system of solid-state disk and 15K rpm SAS drives. Using policies, it automatically moves files to heterogeneous back-end NAS systems based on access frequency, performance, age, etc. EMC's Rainfinity provides similar functionality within its Celerra NS NAS systems. FAST (fully automated storage tiering) on Celerra NS uses the Rainfinity engine for transparent file movement (it currently doesn't support heterogeneous systems). F5 Network's ARX uniquely solves NAS sprawl data protection by managing snapshots and replication for distributed heterogeneous NAS systems.

Clustered NAS systems

Clustered NAS systems use a distributed file system running concurrently on multiple NAS nodes. Data and metadata can be striped across both the cluster and underpinning block (direct-attached storage [DAS] or SAN) storage subsystems. Clustering also provides access to all files from any of the clustered nodes regardless of the physical location of the file. The number and location of the nodes are transparent to the users and applications accessing them.

Although clustering appears similar to file virtualization, the key difference is that all system nodes must be from the same vendor and often configured similarly. Some exceptions to this include BlueArc Corp.'s Titan and Mercury series, and NetApp's Ontap GX.

Clustered NAS systems typically provide transparent replication and fault tolerance, so that if one or more nodes fail, the system continues functioning without any data loss. Clustered NAS systems are distin-

guished by their large file systems that can scale to hundreds of terabytes (or more) of addressable capacity.

Clustered NAS systems include BlueArc's Titan and Mercury series, EMC's Celerra NS-960 with Multi-Path File System (MPFS), Exanet Inc.'s ExaStore, Hewlett-Packard (HP) Co.'s Ibrix Fusion and StorageWorks Scalable NAS (previously known as PolyServe), Hitachi Data Systems' HNAS and 3200 series, IBM Corp.'s Scale-out File Services (SoFS), Isilon Systems Inc.'s IQ, NetApp's Ontap GX, Panasas Inc.'s ActiveStor and Scale Computing's SN Series.

Pros

- Linearly scale to many nodes and high capacities, with millions to billions of managed file objects; aggregate throughput and IOPS independent of one another
- Easy to grow
- Pay-as-you-go architecture
- Built-in fault tolerance
- Centralized management
- Easy data protection
- Simple file access

Cons

- Rip-and-replace solution; can't reuse current NAS systems
- No support for heterogeneous NAS systems
- No ability to migrate files from current NAS systems to the clustered
- Higher hardware and license costs, but may be offset by significantly lower management costs

Clustered NAS does a very good job of resolving most network-attached storage sprawl challenges. It eliminates or at least mitigates the multisystem management issue depending on the scale of the environment. User and application access is simplified with load balancing built in, and data protection and replication is also part of the architecture. Clustered NAS does fall a little short on storage tiering; it does make it easier, but doesn't automate the process (with the exception of EMC's Celerra NS-960 with FAST using Rainfinity).

Private cloud or grid storage systems

Private cloud or grid storage systems are somewhat similar to clustered NAS systems, but grid storage provides peer-to-peer clustering that enables it to provide single-image files over geographically dispersed, long-distance and cross-domain operations.

Geographic location "awareness" adds another dimension to NAS sprawl management by centralizing control, management and access for distributed environments. Based on access performance and/or data

ANTIDOTES TO NAS SPRAWL
How well key technologies deal with NAS proliferation

NAS sprawl challenge	Let sprawl continue	Linux/Unix automounters	Microsoft DFS	Shared-path file virtualization	Split-path file virtualization	Clustered NAS systems	Grid or private cloud storage
Multi-NAS management	**Extremely poor** and very time consuming	**Fair** if environment is primarily NFS and client/app access is from Linux or Unix	**Good** if using Windows file servers or Windows-based NAS	**Very good** plus simplifies migration to this environment	**Excellent** and has exceptional scalability	**Good** but not heterogeneous; rip and replace; excellent scalability	**Excellent** for local and geographically dispersed sites; exceptionally scalable
Client and application file access management	**Poor:** Frustrating, time consuming and error prone		**Very good** for CIFS, but not available for NFS	**Very good** for both CIFS and NFS depending on product	**Good** for both CIFS and NFS depending on product	**Very good** for both CIFS and NFS depending on product	**Excellent:** Policy based to meet user and app requirements regardless of location
Storage tiering	**Poor:** Manual and virtually non-existent	**Poor:** Manual and virtually non-existent	**Fair:** Still manual but a little easier	**Good to excellent** depending on product; much simplified	**Good** but requires agent on clients and servers or data movement is offline	**Good to excellent** depending on product, much simplified	**Good to excellent** depending on product, much simplified
NAS load balancing	**Poor:** Manual and very difficult	**Fair:** Manual only, no automation	**Fair:** Manual only, no automation	**Excellent**	**Excellent**	**Excellent**	**Excellent:** Goes beyond simple load balancing taking into account performance and locality
Data protection, replication, backup and recovery	**Poor:** Uncoordinated	**Poor:** Uncoordinated	**Good** with policy-based replication, VSS integration	**Good to excellent** depending on product; much simplified	**Good to excellent** depending on product; much simplified	**Excellent**	**Excellent:** Local and geographically dispersed; exceptionally scalable
Total cost of ownership	**Poor:** Escalating; rapidly out of control	**Pretty good** with no software costs	**Pretty good** with no software costs and significant duplicate hardware costs	**Very good**, although upfront costs may be high, overall costs much reduced	**Very good** with overall costs much reduced	**Very good**, although upfront costs may be high, overall costs much reduced	**Very good** with excellent payback timeframes

protection policies, files are replicated and moved to the geographic location that best meets the policy. Whether you have only a few remote or branch offices or hundreds, grid or private cloud storage can make a lot of sense.

There are currently two commercially available private cloud storage

systems: Bycast Inc.'s StorageGRID and EMC Atmos. The Bycast StorageGRID runs on x86 nodes that sit in front of standard DAS or SAN storage, so it can use already installed block storage. EMC Atmos also runs on x86 nodes but can only use its own JBOD storage. Bycast's product is a bit more mature with hundreds of installations and OEM deals with HP and IBM.

Pros

- Same pros as clustered NAS
- Same or lower cost than clustered NAS
- Management of geographically dispersed locations
- Distributed geographically aware access with centralized management, protection and replication of all files
- Geographically aware, policy-based file replication and movement
- DAS and SAN investment protection or use of very low-cost storage

Cons

- Limited number of vendors with mature technology
- No automated storage tiering at this time
- Startup costs can be more than other technologies (but long-term costs will likely be less)

SUMMARIZING NAS SPRAWL SOLUTIONS

File storage growth is bordering on the out of control, with many companies struggling to get a handle on their network-attached storage systems. This NAS sprawl creates serious management problems that can tax overworked IT staffs and jeopardize users' access to corporate data. But the four different technologies described above are available today and can resolve many of the issues and challenges created by NAS sprawl.

Take a pragmatic approach, and implement the least amount of new technology that best meets current and forecasted requirements. That will help minimize risk, lessen the strain on CapEx and OpEx budgets, and can make a world of difference with NAS management. ⊙

Marc Staimer is president of Dragon Slayer Consulting.

STORAGE NETWORKS get virtual

Servers are virtualized, storage is virtualized, but what about your storage network? New products now make it possible to pool and share networking resources.

By Dennis Martin

STORAGE NETWORK VIRTUALIZATION, also known as I/O virtualization (IOV) or I/O consolidation, comprises an emerging family of technologies that extend the concept of virtualization to the major types of input and output (I/O) handled by today's servers.

In recent years, data centers have been deploying server and storage virtualization technologies to more efficiently use underutilized computing assets and to create more flexible infrastructures. In decoupling the logical function from the physical hardware, virtualization allows hardware to be pooled and shared, thus improving utilization. Once in place, virtualization makes new server or storage deployments much quicker and easier, while making it easier to effect changes to the existing infrastructure. For example, it's much easier to deploy new virtual servers than physical servers. And when storage systems are virtualized, many of the data migration issues related to new array deployments can be avoided by adding the new capacity to the existing pool of storage resources.

Virtualization has long been applied to a number of different computing technologies. Although storage virtualization has its roots in the mainframe world, it's only now beginning to gain wider adoption. Server virtualization, on the other hand, has become the poster child of virtualization in the last few years. A handful of organizations are now applying similar virtualization techniques to the "connective tissue" that links servers to storage in enterprise environments.

But I/O virtualization isn't exactly a brand new idea either, with virtualization concepts already being used for some network I/O technologies today. For example, a virtual local-area network (VLAN) separates the logical and physical aspects of a network so that one physical network appears as and can be managed as several smaller logical networks. Network interface card (NIC) teaming combines two or more network adapters and makes them appear to function as a single adapter with increased bandwidth. In both cases, logic in the hardware and management software layers allows the decoupling of the logical functions from the physical hardware, making it possible to carve up the hardware and share it as separate units, or to combine it to present it as one larger unit.

PCI EXPRESS AND I/O VIRTUALIZATION

A server in an enterprise data center typically needs access to a LAN, a storage-area network (SAN) and local direct-attached storage (DAS). Some servers also

PCI EXPRESS BOOSTS I/O VIRTUALIZATION

The PCI-SIG, the special interest group responsible for PCI Express (PCIe) industry-standard I/O technology, announced the completion of the PCI-SIG I/O virtualization (IOV) suite of specifications in June 2008. These specifications enable virtualization solutions to tackle the most I/O-intensive workloads by removing performance bottlenecks in both software and hardware virtualization components. The IOV suite provides a set of technologies that can be used by providers of processors, chipsets and I/O fabrics, and has implications for hypervisors and operating systems. These technologies provide:

- Address Translation Services (ATS) so that I/O devices can take on various personalities
- Single-Root IOV (SR-IOV) for native I/O virtualization in existing PCI Express topologies within a single server
- Multi-Root IOV (MR-IOV) for native I/O virtualization in new PCI Express topologies where multiple servers share a PCIe fabric

This set of specifications promises to trigger new virtualization solutions that provide improved performance, lower power consumption and new terminology that will change the way we view I/O to and from a server.

need access to high-end graphics processing. A server's access to these resources usually comes by way of an internal system bus. In a newer multicore physical server with a high-speed PCI Express (PCIe) bus, all of these I/O "pipes" occasionally hit peak bandwidth, but rarely simultaneously or on a sustained basis (see "PCI Express boosts I/O virtualization," p. 38). With many virtualized servers running on a single physical server, these I/O pipes are busier, but aren't likely to be running at full bandwidth simultaneously or on a sustained basis.

What if, instead of installing separate network and storage adapters in every server, the PCIe bus adapters could be virtualized and shared across multiple servers? Consider the potential cost and power savings for NICs, host bus adapters (HBAs) and SAS/SATA disk controller cards that could be shared across a rack of servers. A rack full of servers could have only one cable for each server connecting it to a virtualized set of I/O adapters at the top of the rack. Then that top-of-rack unit could dynamically direct all LAN, SAN and DAS traffic to the appropriate location as needed, such as end-of-row switches for example, leaving the servers to focus on computing. This "rack-area network" (RAN) concept can allow an entire rack of servers to have some of the same benefits as blade servers, but without the limitations of a blade server chassis. The consolidation realized in this scenario would also mean that the size of the rack servers could be reduced to 1 rack unit (1U) or even one-half of a rack unit (1/2 U).

What if, instead of installing separate network and storage adapters in every server, the PCIe bus adapters could be virtualized and shared across multiple servers?

Consider the movement of a virtual machine (VM) from one physical server to another physical server. Typically, this requires a SAN, because SANs are separate from the physical server and can be accessed from any server, assuming all of the security, zoning and logical unit numbering (LUN) masking issues have been addressed. What if movement of virtual machines could be made to work with any storage, rather than requiring a SAN? I/O virtualization-capable adapters would run some of the hypervisor functions in hardware, offloading the host CPU and freeing up CPU resource that could be used to host additional virtual machines or applications.

I/O VIRTUALIZATION vs. OTHER NETWORKING TECHNOLOGIES

Ethernet Data Center Bridging (DCB) and Fibre Channel over Ethernet (FCoE) are a pair of technologies that are young, but slightly more mature than I/O virtualization in today's marketplace. Together, DCB and FCoE allow for hardware consolidation by combining lossless Ethernet

with Fibre Channel at the switch and at the host adapter. This DCB/FCoE technology combination provides some of the same type of consolidation that I/O virtualization provides, but it's actually complementary to IOV. Because the DCB/FCoE converged adapters run on the PCI Express bus, they can be used in an I/O virtualization environment and, therefore, could be shared across multiple servers. The host adapters that support DCB and FCoE currently support or will soon support IOV technologies such as Single-Root IOV (SR-IOV). An IOV environment can communicate with existing Ethernet, Fibre Channel and DCB/FCoE switches using existing adapters and, as far as the host servers are concerned, they're connected directly to those switch environments.

Because the DCB/FCoE converged adapters run on the PCI Express bus, they can be used in an I/O virtualization environment and, therefore, could be shared across multiple servers.

InfiniBand is another high-speed, low-latency network technology that's typically used in compute cluster environments for server-to-server communication. InfiniBand provides faster speeds than Ethernet today. The newer InfiniBand host adapters, known as host channel adapters (HCAs), run on the PCI Express bus and can support I/O virtu-

I/O VIRTUALIZATION PRODUCT SAMPLER

Company	Product	Description
Aprius Inc.	Aprius G80 I/O Gateway	Top-of-rack device that uses an extended PCI Express (PCIe) bus to virtualize NICs, HBAs, CNAs, SAS, etc.; sharing based on policies, QoS, bandwidth, etc.; first demos expected early 2010.
Mellanox Technologies Ltd.	ConnectX-2 adapter cards	Virtual Protocol Interconnect (VPI) adapters support multiple protocols over one device, including InfiniBand, Ethernet, Data Center Ethernet, FCoE and Fibre Channel over InfiniBand.
NextIO	N1400-PCM PCI Express High Speed Switch Module; N2800-I/O Consolidation Appliance	Extends PCIe bus to top-of-rack device for racked servers and blade systems; allows pooling and sharing of NICs, HBAs, HCAs, CNAs, SAS, etc.
VirtenSys Inc.	VIO-4001 and VIO-4008 virtualization switches	Top-of-rack unit that uses extended PCIe bus; consolidates and shares NICs, HBAs, CNAs, SAS, etc.
Xsigo Systems Inc.	VP780 I/O Director	Connects to servers using 20 Gb InfiniBand and creates virtual NIC and HBA interface cards.

alization. In addition, some vendors are developing I/O virtualization solutions built around InfiniBand technology, using InfiniBand as the high-speed carrier for the IOV infrastructure.

CURRENT I/O VIRTUALIZATION PRODUCTS

The general IOV approach that most current products take is to connect the local host servers into a top-of-rack unit that holds a variety of network, storage and graphics adapters that can act as a dynamic pool of I/O connectivity resources. The top-of-rack device acts as an I/O fabric for the servers in the rack, and can communicate with other servers in the rack or can connect to end-of-row switches for more distant resources. These IOV top-of-rack units may be less expensive than some of the newer high-speed top-of-rack switches.

Two specific implementation models for I/O virtualization are emerging: PCIe- and InfiniBand-based approaches.

At least three companies are working on products that extend the PCI Express bus into a separate box for the purpose of virtualizing I/O adapters.

One approach to IOV is to extend the PCI Express bus out of the server chassis and into a separate box or chassis populated with IOV-capable adapters that can be shared across multiple servers. The I/O virtualization box would be installed in a rack and would function somewhat similarly to a top-of-rack switch, except that instead of only supporting Ethernet or Fibre Channel, this IOV box would act as a type of fabric switch for all LAN, SAN, DAS and possibly graphics traffic. At least three companies are working on products that extend the PCI Express bus into a separate box for the purpose of virtualizing I/O adapters. One advantage to this approach is that servers today already support PCI Express. Some IOV vendors now have first-generation products available and some are publicly discussing products that will appear this year. Some of these products require support for SR-IOV or Multi-Root IOV (MR-IOV), but others don't have that requirement. These products are built around the PCI Express 2.0 specifications, and vendors already have PCI Express 3.0 plans in their product roadmaps.

Aprius Inc. is a small vendor that's building a PCI Express gateway device that will support almost any type of PCI Express adapter (including network cards, storage controllers and graphics coprocessors) that can then be shared across multiple servers. These adapters basically form an I/O resource pool that can be dynamically assigned to physical or virtual servers.

NextIO is a company that was involved with developing the PCI-SIG

I/O virtualization specifications and had some IOV products as early as 2005. NextIO is working in several areas, including the high-performance computing (HPC) market and is interested in virtualizing graphics coprocessing in addition to traditional networking and storage I/O traffic. They're partnering with several big name vendors for a variety of IOV applications.

VirtenSys Inc. extends the PCIe bus with its I/O virtualization switches that can virtualize the major types of server networking and storage connectivity, as well as interprocessor communication (IPC) for HPC compute cluster environments.

Another approach to I/O virtualization is to use an existing network interconnect technology such as InfiniBand or 40 Gb Ethernet as the transport for virtualizing I/O adapters. Two companies are building products to handle IOV in this fashion:

Mellanox Technologies Ltd., well-known for its InfiniBand products, provides its I/O consolidation solutions using either InfiniBand or 10 Gb Ethernet (10 GbE) as the transport for performing IOV. They're also building 40 Gb Ethernet adapters that are compliant with SR-IOV.

Xsigo Systems Inc. uses InfiniBand HCAs that connect to its I/O Director that provides the infrastructure for IOV-capable adapters. One reason for using InfiniBand is its high speed and very low latency. Inside the I/O Director are the same PCI Express network and storage adapters that would otherwise be installed in each host server. Xsigo's I/O Director has been available for approximately two

HAIRPIN TURNS

In an I/O virtualization (IOV)-capable environment, traffic can be sent out of one virtual adapter and into another virtual adapter without regard to the underlying physical hardware. This leads to the very interesting possibility that traffic could be entirely contained in a single physical adapter, which is known as the IOV "hairpin turn." The application for this might be a virtual machine (VM) communicating with another virtual machine through their respective virtual network interface cards (NICs), where the virtual machines reside on the same physical server and the virtual NICs reside on the same physical NIC. In this case, the physical NIC is functioning as a mini-switch. This analogy also works with the top-of-rack IOV units, where the physical adapter is external to the server, but functions in the exact same way.

A storage adapter could be made to operate the same way. Suppose a host server had a Fibre Channel, iSCSI or SAS adapter that was located in a top-of-rack IOV unit. A storage server could be located in the same rack and could, theoretically, use the same storage adapter in the IOV unit as its adapter to the outside world. The physical adapter in the IOV unit would have one virtual adapter configured as the initiator and another virtual adapter configured as the target. Interesting possibilities indeed!

years, and the company has established partnerships with a number of storage vendors, including Dell Inc. and EMC Corp.

Many network and storage adapter vendors are working on full support for I/O virtualization, especially for compliance with the SR-IOV and/or MR-IOV specifications. The vendor roster includes Emulex Corp., Intel Corp., LSI, Neterion Inc., QLogic Corp. and others. The big server vendors, including Dell, Hewlett-Packard (HP) Co. and IBM Corp., are beginning to demonstrate solutions that support I/O virtualization, either in their rack servers or blade servers, or both. Cisco Systems Inc. has also joined the movement with its Cisco UCS M81KR Virtual Interface Card. The big processor vendors, Advanced Micro Devices (AMD) Inc. and Intel, include virtualization technologies that help enable some of these IOV functions.

HOW AND WHEN TO IMPLEMENT I/O VIRTUALIZATION

Implementation of I/O virtualization technologies will most likely be a slow, deliberate process. That's because the work to make all the adapters function in this manner isn't complete yet, and because the top-of-rack IOV units are still in their early stages. For I/O virtualization to work properly, development work needs to be completed on the adapter hardware and firmware, drivers, operating systems and hypervisors. Several vendors will be announcing support for various forms of IOV in 2010, and it's anticipated that IOV will emerge as one of the top new technologies for the year. However, expect I/O virtualization to take a few years to become commonplace.

Look for 10 Gb Ethernet adapters to be the first to fully support IOV.

Look for 10 Gb Ethernet adapters to be the first to fully support IOV. Demonstrations of IOV-capable 10 GbE adapters were shown publicly in 2009 at a number of trade shows. After the Ethernet adapters, you can expect to see storage adapters such as Fibre Channel HBAs, FCoE CNAs and SAS/SATA non-RAID adapters to support I/O virtualization. The last category of storage adapters that will likely fully support IOV are the RAID controllers, due to the complexity of sharing RAID functions across servers. Separately, some graphics coprocessor adapters will support IOV, with some products possibly available in 2010.

Implementing IOV-capable adapters will require top-of-rack I/O virtualization units and either PCIe bus extender cards or InfiniBand HCAs for the host servers, depending on the implementation. The IOV-capable adapters are then placed in the top-of-rack IOV units and can be shared across servers. Drivers for these adapters will be needed, and few production-ready drivers for any operating system are currently available.

I/O virtualization should be implemented in stages, as with the adoption

of any other new technology. The IOV implemented stages should begin with pilot tests run on a small number of servers; the pilot implementation should run until the products operate in a stable manner and benefits can be shown. The Demartek lab will be testing various IOV solutions during 2010, and we'll be able to provide first-hand commentary and results.

A good candidate environment for I/O virtualization might be a virtual server environment that would benefit from sharing some higher-end 10 GbE NICs or similar high-speed adapters. One of the goals of IOV implementations may be to acquire the necessary I/O adapters based on the overall bandwidth needs of all the servers in a rack, rather than simply buying adapters based on raw server count. This will require adjustments to the planning process to account for applications and bandwidth usage, and may require more bandwidth measurements to be taken in the current environments.

MANAGEMENT ISSUES WITH I/O VIRTUALIZATION

Managing virtual pools of I/O resources will require some new thinking. The adjustment is similar to what was required to effectively manage storage systems when SANs and virtualized storage solutions were first deployed. You'll need to understand that the I/O adapters and paths will no longer be exclusively owned by a particular server, in the same way that storage on a SAN isn't owned by a specific server. Rather, these adapters and paths will be dynamically assigned to servers, and can be released or adjusted as needed. Each of the vendors providing top-of-rack IOV units will have their own management interface for the I/O virtualization unit itself, and some level of adapter management. In addition, each of the adapter manufacturers will provide some basic element manager, similar to what's provided today.

It remains to be seen how the operating systems and hypervisors will view these virtualized I/O adapters. Because ownership of the adapters will no longer be tied to a particular operating system or hypervisor, the management of these IOV resources will have to be aware that these resources can logically move around in the data center and that the I/O resources can have multiple personalities.

Dennis Martin has been working in the IT industry since 1980, and is the founder and president of Demartek, a computer industry analyst organization and testing lab.

10 ways to improve data protection in VMware environments

With backup vendors revising their offerings to take advantage of VMware vSphere 4's new features, an upgrade to vSphere in 2010 could protect your organization's data storage.

BECAUSE A VIRTUAL MACHINE (VM) is an encapsulation of an operating system, applications and data, protecting a production virtual disk is vital. IT organizations face several challenges when it comes to protecting their virtual machine disk image (VMDK) files. Finite shared physical resources on the host system and in-guest I/O-intensive backup processes have the potential to bring VM performance to a crawl. The increase in the amount of data in VMware environments can impact the backup window.

VMware Inc.'s release of vSphere 4 in the first half of 2009 introduced several data protection enhancements over VMware Infrastructure 3-era tools. And now that backup vendors have had several months to revise their offerings to take advantage of its new features, an upgrade to vSphere in 2010 could greatly improve your organization's data protection.

VMware Data Recovery (VDR) is an agentless, disk-based backup solution that employs snapshots and data deduplication.

Here are 10 backup-related reasons why you should consider a vSphere upgrade in 2010:

1) VMware vStorage APIs for Data Protection. VMware introduced VMware Consolidated Backup (VCB) command line utilities in ESX 3.0 to simplify data protection by offloading backup from ESX Servers to one or more centralized proxy servers. It allows a live system image to be captured without disrupting the VM-resident applications or overwhelming the host's CPU. With VCB, backups occur off-host so there's no "backup window" to contend with. This method also removes the need for an agent in every virtual machine and provides VM-level recovery. But it requires either a two-step backup or recovery to perform a file-level recovery. In vSphere, VMware completely rearchitected the facilitation of backup. Instead of acquiring VCB and implementing it with your backup application, your backup vendor now leverages vStorage APIs (virtual machine file system, or VMFS, drivers) to access data on VMFS volumes. This integration makes for a much better backup implementation.

2) VMware Data Recovery (VDR). For those who prefer to use a native backup tool, VDR is an agentless, disk-based (local or Fibre Channel, iSCSI or network-attached storage shared disk) backup solution that employs snapshots and data deduplication. Running in a virtual machine as a Linux virtual appliance, VDR backup and recovery tasks are launched through VMware vCenter Server. Because VDR takes advantage of the vStorage APIs for Data Protection, a VM-level backup is performed that enables fast backup and recovery. Granular (file-level) recovery is possible from the VM-level backup. Using block-level data dedupe, it keeps required disk space to a minimum. VDR is included with vSphere Enterprise Plus, Advanced and Essentials Plus editions.

3) Granular recovery. File-level recovery is now possible from a VM-level backup without a two-step process. This allows organizations to perform an image-level backup (with no impact on the CPU of the host system) and to later do either a bare-metal-like or individual file recovery.

4) Changed block tracking (CBT). In the past, copying VMDKs for backup purposes meant moving large files around your LAN, WAN and/or SAN. The effect? An increase in network traffic and backup time. VMware introduced CBT in vSphere 4 to markedly improve performance. CBT allows the VMkernel to track changed blocks of a virtual machine's virtual disk. Backup applications using vStorage APIs can immediately identify the blocks that have changed since the last backup and copy only those blocks—cutting down the time it takes to capture and transfer data, as well as reducing traffic on the network.

5) No physical proxy server required. Previously, a physical machine was required for the VCB backup proxy server. But organizations can now reduce their infrastructure commitment because the proxy system that mounts the LUNs can be a virtual machine.

6) vCenter Server plug-ins. Monitoring and management improvements are gained via vCenter Server integration. Backup vendors offering vCenter Server plug-ins can provide views of a virtual machine's backup status, including success/failure, method of backup and schedule. Some even automate the process of assigning backup policies to new VMs as they come online, eliminating protection gaps.

7) Direct backup from shared storage. When the backup data path includes a proxy system, it can create network bottlenecks and inefficiency. Now, a direct-to-target architecture is employed, allowing backup data to go from the ESX host to the storage target. Leveraging the vStorage APIs, virtual disk data is read directly from the SAN storage device.

8) iSCSI enhancements. Organizations using software iSCSI initiators to implement shared storage in virtual server environments could see performance improvements with vSphere. More disk I/O throughput and reductions in CPU overhead benefit I/O-intensive backup processes.

9) Thin provisioning. VMware made some improvements in vSphere for thin provisioning: the ability to overcommit disk space. Not only is it easier to create, maintain and monitor "thin" disks, but using them could help improve backup performance, as well as the capacity of space used on host's data stores. Here's why: When a VM is created, storage is allocated. Virtual machines typically don't use all of the storage space assigned, so the VMDK file could be unnecessarily larger. Backup applications are able to recognize "empty" file system space to avoid backing them up, but the process takes up valuable time. Thin provisioning results in smaller VMDK files and faster backups.

10) Data deduplication. Data deduplication is not a feature of vSphere unless you use VDR for backup. Therefore, it's likely that you will look to your backup application provider for this feature. Since taking advantage of vSphere features, such as CBT, vStorage APIs for Data Protection and vCenter Plug-ins, requires you to also upgrade your backup application, you'll likely gain dedupe features in the upgraded backup application. Given the tendency for duplicate data in backup processes, eliminating duplicates—at the "source" VM or the proxy server—will reduce the amount of data transferred and stored, and improve performance. ⊙

Lauren Whitehouse is an analyst focusing on backup and recovery software and replication solutions at Enterprise Strategy Group, Milford, Mass.

SAS drives showing up more and more

THE LATEST 6 Gbps serial-attached SCSI (SAS) disk drives are beginning to show up in data centers, but their 3 Gbps predecessors have already made significant inroads in enterprise data storage shops (see "SAS challenges Fibre Channel drives," *Storage* magazine, January 2010). Slightly more than 50% of respondents to our latest survey currently use SAS drives. Right now, SAS drives are most used for direct-attached server storage (63%), but more than half of our SAS users have the drives installed in arrays. When we last asked about SAS drives a year and a half ago, approximately 26% of those respondents said that SAS drives accounted for 40% or more of their installed disk capacity; this time, 40% of respondents said that SAS accounted for that much capacity. The two biggest obstacles to SAS growth seem to be coming from above and below: of those who aren't currently using SAS disks, 32% said they need the performance provided by Fibre Channel (FC) while 38% said SATA suits them just fine. The higher capacity drives for both FC and SATA that have recently become available are likely factors that have helped these two legacy protocols hold off the upstart SAS just a little bit longer.

—Rich Castagna

Do you use SAS drives?

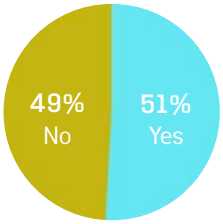

Where are you using SAS drives?
(Of those who said they currently use SAS drives)

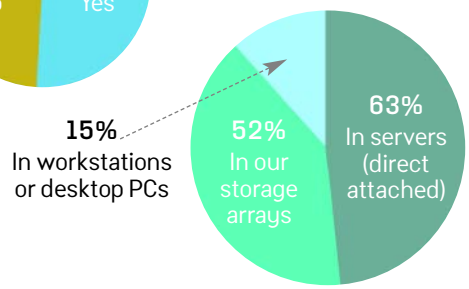

What are the reasons you don't use SAS drives?

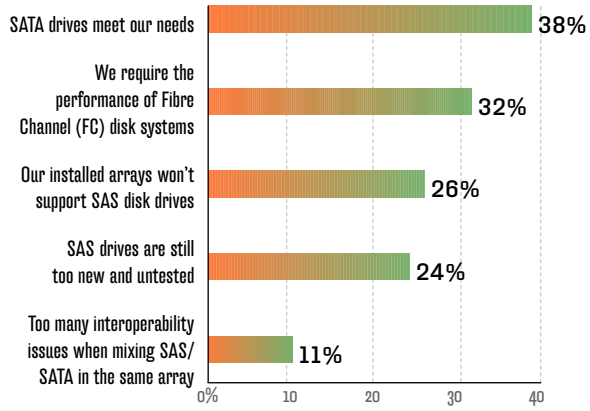

Why did you decide to use SAS drives?

56%	We're using a mixed SAS/SATA storage system
45%	SAS drives offer us adequate performance at a lower cost than FC drives
20%	The performance and/or reliability of SATA drives is inadequate for our needs
12%	We're using SAS drives to create a new storage tier

34%

of respondents say **SAS drives** account for 21% to 40% of their installed disk capacity

"SAS is the future—FC drives will disappear within the next 18 months."

—Survey respondent

Check out the following resources from our sponsors:

Hitachi Data Systems, page 4

The Benefits of Intelligent Tiered Storage and Hitachi Solutions

Economize your Storage the Hitachi Way

Pillar Data Systems, page 7

The Modern Virtualized Data Center

Complimentary White Paper: Bringing an End to Storage Waste

Syncsort

Virtualization, Deduplication, and the Data Protection Connection: Reinventing IT Infrastructure Consolidation with Efficient Data Protection

Four Critical Imperatives for Data Protection and Recovery: BEX Innovative Architecture Delivers Holistic Data Protection

Tangent

Email Archival Best Practices and Competitive Analysis

Data Cove Email Archival