

Managing the information that drives the enterprise

# STORAGE

Vol. 9 No. 1 March 2010

## Backup Apps Add Missing Pieces

Your backup application looks a lot different than it did a few years ago. See what new capabilities are being built in by backup vendors.

P. 13

### ALSO INSIDE

- 5 Gaps still remain in cloud storage
- 8 StorWars: Automated tiered storage
- 22 Virtualize your disaster recovery plan
- 31 Readers rank the best enterprise arrays
- 37 File storage that can really grow
- 40 Storage shops serious about solid state

**NEW**

# STORAGE

inside | march 2010



**HITACHI**  
Inspire the Next



**Quantum**

**Regional Solution Providers:**





## Cloud Storage's Missing Link

**5 EDITORIAL** A lot of storage pros seem to think cloud storage is ready for the enterprise. It might be, but there are still some gaps to fill. by RICH CASTAGNA

## No More Wasted Tiers

**8 STORWARS** We welcome back storage expert Tony Asaro to the pages of *Storage* magazine, where he'll share his unique view of the storage industry. This month, Tony explores intelligent tiered storage, which is becoming a necessity for many firms due to the massive amounts of data they're storing. by TONY ASARO



## Top New Features in Backup Apps

**13** Backup applications have evolved over the last few years and now incorporate features that were previously available only in third-party products. We'll look at some of the key backup technology advancements and describe how four of the leading backup vendors—CommVault, EMC, IBM and Symantec—have implemented these technologies. by W. CURTIS PRESTON

## Virtualize Disaster Recovery

**22** Virtualization can provide some clear advantages for disaster recovery. It can help you save money, time and effort, as well as make the often daunting task of designing and implementing a DR plan easier. But there are related challenges and costs. by ERIC SIEBERT



## Quality Awards V Enterprise Arrays: NetApp Alone at the Top

**31** In the last Quality Awards for enterprise arrays, NetApp and EMC finished in a dead-heat for first place. This time, NetApp ekes out a narrow victory over archival EMC. by RICH CASTAGNA

## Could 2010 be a Breakout Year for Scale-out NAS Architectures?

**37 HOT SPOTS** Scale-out NAS meets today's requirements for massively scalable and highly available systems, and it's generally a more efficient option than traditional scale-up architectures. But technology change introduces risk, and companies may not be ready for a switch. by TERRI MCCLURE

## Beginning of the End for Hard Drives?

**40 SNAPSHOT** Enterprise-ready solid-state storage hasn't been around for very long, but 34% of the respondents to our latest Snapshot survey have solid state up and running. Price is still an issue, but disk's days may be numbered. by RICH CASTAGNA

## Vendor Resources

**41** Useful links from our advertisers.

On page 26 of *Storage* magazine's February 2009 Products of the Year article, DataCore Software Corp.'s SANsymphony 7 was listed in the wrong product finalist category. It should be listed under Backup and Disaster Recovery Software and Services.



# EFFICIENT, AFFORDABLE & AUTOMATED DISASTER RECOVERY

MAKE DISASTER RECOVERY FOR THE VIRTUAL DATA CENTER A REALITY WITH DELL



The Dell™ EqualLogic™ PS Series SAN solution tightly integrates with your virtualized server deployments to automate data protection and disaster recovery. Its all-inclusive, intelligent feature set includes advanced hypervisor-aware snapshots and remote replication tools to efficiently create and maintain data copies online, at both local and remote sites, for rapid recovery of data not in days, not in hours, but in minutes.

The Efficient Enterprise runs on Dell.  
To learn more, visit [www.getmorevirtual.com](http://www.getmorevirtual.com)



# Cloud storage's missing link

*A lot of storage pros seem to think cloud storage is ready for the enterprise. Maybe so, but there are still some gaps to fill.*

**WHEN I FIRST** started thinking about writing this column I decided, no, write about something else, this whole cloud storage thing is already getting a little tired. Readers are probably getting clobbered with cloud-this and cloud-that from every direction as the hype machine revs into overdrive, so I should give them a break and can the cloud chatter for awhile.

But then a few surveys about cloud services popped up in my in-box. I'm a sucker for survey data, maybe because we do a considerable amount of survey-based research ourselves. So if you put a bar chart or line graph plotting user responses in an article, I'll probably read it. There's just something about quantitative data that turns me on.

What truly caught my eye were survey results that indicated that you, and your IT peers, are verging on going gaga over cloud services, including, of course, cloud storage. In a press release announcing the results of its survey, GlassHouse Technologies said, "60 percent of executives plan to implement cloud initiatives in the coming year." Another survey—this one commissioned by CommVault—finds that "52 percent of respondents are considering the use of cloud storage services now or in the future." The "in the future" leaves the door pretty wide open, but, still, that's a lot of folks gazing into the cloud.

Our own 2010 Storage Priorities survey produced a little more granular results. We found that approximately 4% of respondents were already using cloud storage for online or nearline data, with 9% planning to implement it this year and another 27% expecting to evaluate it in 2010. Add those numbers up, and 40% are using/about to implement/evaluating the technology. Not 60%, but still pretty impressive. Surprisingly, for data protection cloud services, the numbers were a bit lower, with figures of approximately 32% each for cloud backup, data archiving and disaster recovery (DR).

Whichever numbers you decide to believe, it looks like a lot of people are getting pretty serious about extending their storage operations into the cloud—despite cloud storage being almost universally dismissed (I was part of that universe) a few years ago as only a small-business play.

**Survey results indicate that you, and your IT peers, are verging on going gaga over cloud services, including, of course, cloud storage.**

Is it the year of scale-out NAS?

Quality Awards V: Enterprise arrays

Virtualize disaster recovery

Top new backup app features

Intelligent tiered storage

But attitudes have changed considerably, and you have to credit the cloud storage vendors with pleading a good case and addressing the key enterprise concerns. As the CommVault survey notes, despite the heightened overall interest, a lot of IT pros (75%) still feel some uncertainty about cloud storage, with most of that concern centered on data privacy and security. That, I think, is a relatively easy problem for cloud vendors to solve, as most of them offer encryption for data in flight and at rest.

But there are still some big question marks about cloud storage. One of them popped up during a meeting the other day with a new cloud storage vendor about to roll out its first product. It's a pretty neat product that looks like it'll do a good job of integrating internal operations with cloud storage, and not just for nearline storage, but for active file data. Like many other cloud vendors, this one isn't blowing big bucks on building its own data storage facilities but, rather, is content to lean on the biggies for this, including the usual suspects, like Amazon, Iron Mountain, Nirvanix, Rackspace, etc. It's a tidy package with consolidated billing and some pretty big names doing the hosting. But the issue that arose was kind of a Catch-22 related to service guarantees—essentially, there are none. The vendor is just selling software to make the integration work and make the cloud look like local disk. They're not providing the real disk, so there's no way in the world they can guarantee that your data will always be available, protected effectively and that the company giving it a home will be around a few years from now if a judge wants to see those old emails. That's a gap in the process that's just big enough for your data to fall through.

There's no question that cloud storage has covered a lot of ground over the past couple of years, going from a home-office/power-user service to something that warrants the attention of enterprise IT. But there are still some annoying little gaps that must be filled before we see the 40%, 50% or 60% of respondents who are interested in the technology turn into real cloud storage users.

**Attitudes have changed considerably, and you have to credit the cloud storage vendors with pleading a good case and addressing the key enterprise concerns.**

## WELCOME BACK, TONY

With this issue of *Storage* magazine, we're pleased to introduce a new column from an old friend and former contributor. Tony Asaro, CEO and founder of The INI Group, will share his considerable storage experience and insights each issue in his new StorWars column. Tony previously did a stint at the helm of our Storage Bin 2.0 column and has contributed articles to our websites. He was also a popular speaker at many of our Storage Decisions conferences. It's great to have you back, Tony. ☺

---

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

\* [Click here for a sneak peek at what's coming up in the April 2010 issue.](#)



# They're left following in our shadows.

Superior virtualization from Hitachi.

Hitachi Data Systems outruns the competition in many areas. Superior virtualization is just one of them.

Hitachi storage virtualization can transform your multivendor storage into one common pool of shared resources today. It also reclaims "lost" storage space and improves utilization of existing storage assets to significantly reduce capital expenditures or CAPEX.

To reduce your long-term storage ownership costs, Hitachi Data Systems solutions should be top of mind. Manage your data on one dynamic platform today. Hitachi storage virtualization reduces long term ownership costs and risks.

For more information on our superior virtualization, please visit:

[www.hds.com/go/virtualizetoday](http://www.hds.com/go/virtualizetoday)

To arrange a discussion please call:



**866-637-2763**





# No more wasted tiers

*With the massive amounts of data that companies are storing, intelligent tiered storage isn't a luxury—it's quickly becoming a necessity.*

**WAS SITTING** in a room of approximately 20 IT professionals from about 17 different companies and we were discussing the concept of intelligent tiered storage for SAN-based storage systems. Most of the people at this session worked for large companies with literally petabytes of capacity on the floor.

One of them was uncomfortable with the notion of the data storage system making its own decisions to move data to different tiers based on policies. He felt that some decisions should be made by humans. Two others in the group immediately disagreed. The “aha” moment came when one of them replied that with petabytes of data there was just no way they could make tiering decisions themselves. She said the storage system needed to be smart enough to move data based on metrics that would take people too long to analyze. She also pointed out that the problem will only get worse as their environment continues to grow.

Interestingly, this discussion was theoretical because at the moment none of these IT professionals have storage systems that can provide intelligent tiered storage at a granular enough level that is valuable to them. However, we've all been told by some of the leading storage vendors that this capability is coming in 2010.

Intelligent tiered storage is one of the most valuable capabilities on the horizon in the storage industry. Here's why:

- It's common for 60% to 80% of all data to become dormant 90 days after its creation. If you have 100 TB of data, then 60 TB to 80 TB of it is idly consuming lots of expensive primary storage each day.
- It's fairly universal that network storage systems have a capacity utilization of 50% or less.
- Storage capacity continues to grow for companies of all sizes and in all verticals. It was once unheard of for a small company to have terabytes of capacity and large companies to have petabytes, but it's now commonplace. This dynamic seems to be shaping up as a perpetual situation and, as such, tiering will become requisite.
- We know that tiering storage can save significant amounts of money. That's why we have different tiers to begin with. But tiering today is very rigid and requires innovation.
- Even though all of the above are true, very few companies have implemented intelligent tiered storage because it's not granular enough.



A number of storage systems allow you to set a policy to automatically move an entire volume, but that's pretty useless in most cases. If 10% of your volume is active and the other 90% is inactive, that volume is considered to be "active" and isn't moved to a lower tier. However, imagine if you could keep the active 10% on tier-1 storage and move the inactive 90% to a lower tier. The economics of granular tiering are essentially a "no-brainer."

Intelligent tiering is achieved by using technology that can "stretch" a volume across different tiers of storage. Tier 1 could be RAID 10 using solid-state or Fibre Channel drives. Tier 2 might be RAID 5 using SATA drives. The cost differential can be significant. For example, one of the IT professionals in our session actually converted all of their company's storage from RAID 1 to RAID 5 and saved more than \$100 million dollars! Not everyone has petabytes of storage like they do, but this example illustrates the impact tiering can have and the same relative savings for companies of all sizes.

This technology works today and is provided by emerging SAN storage vendor Compellent Technologies Inc. with a feature called Data Progression. 3PAR Inc., EMC Corp. and Hitachi Data Systems do intelligent tiering at a volume level, and I believe they'll all add support for more granular levels of tiering this year.

The IT professionals in this particular session were cautiously optimistic that data storage vendors are heading in the right direction. There was consensus that granular intelligent tiered storage was going to have a major economic impact in their environments. However, they all felt that it wasn't enough just to move data efficiently between tiers; it's also essential to have good reporting and to ensure that all service levels are met without a hitch.

Innovation is still alive and well in the storage world. Over the last two decades there have been milestone features such as mirroring, replication, snapshots, logical pools, thin provisioning and data deduplication. I believe intelligent tiered storage will be added to this powerful list with the potential to enable major economic improvements in SAN storage. ☺

---

Tony Asaro is a senior analyst and founder of Voices of IT ([www.VoicesofIT.com](http://www.VoicesofIT.com)).

Intelligent tiering is achieved by using technology that can "stretch" a volume across different tiers of storage.



# Your CEO, CFO and CIO Thank You.



Putting your data on Pillar Axiom® is the most efficient way to boost productivity, cut costs, and put more money to the bottom line.

Don't put up with the wasteful ways of legacy storage systems. They are costing you way too much in floor space, energy, and money.

Make the move to Pillar Axiom, the world's most efficient system, because it's truly Application-Aware. Get the industry's highest utilization rate – up to 80%. Guaranteed. Slash energy costs. Save 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](http://www.pillardata.com)



Download a complimentary white paper:  
*Bringing an End to Storage Waste.*  
[www.pillardata.com/endwaste](http://www.pillardata.com/endwaste)



Is it the year of scale-out NAS?

Quality Awards V: Enterprise arrays

Virtualize disaster recovery

Top new backup app features

Intelligent tiered storage

# STORAGE

## COMING IN APRIL

### Dedupe in Primary Storage

Data deduplication has taken its place in backup operations, with more and more shops adding this capability daily. But primary storage can also reap the benefits of data reduction, with the potential for even greater savings. Still, only a handful of companies offer primary storage data reduction technologies; see who they are, how they're doing it and what kind of results their implementations yield.

### Keeping Tabs on Storage Systems: Key Metrics to Manage

A few years ago, storage resource management (SRM) looked like the answer to the problem of monitoring and maintaining complex storage systems. But big, all-encompassing SRM tools were often too unwieldy for everyday storage administration. We describe the key storage system metrics you need to monitor and maintain, and the tools that can help you do it.

### Tapping the Power of Director-Class Switches

Director-class switches are a lot more than just really big switches with lots of ports. They provide great throughput, configuration flexibility and can host a number of advanced applications. This article surveys the director switches currently available, describes their capabilities, shows how they can improve network management and suggests where they can be used most effectively.

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

#### TechTarget Storage Media Group

### STORAGE

**Vice President of Editorial**  
Mark Schlack

**Editorial Director**  
Rich Castagna

**Senior Managing Editor**  
Kim Hefner

**Senior Editor**  
Ellen O'Brien

**Creative Director**  
Maureen Joyce

**Editorial Assistant**  
Francesca Sales

**Contributing Editors**  
Tony Asaro  
James Damoulakis  
Steve Duplessie  
Jacob Gsoedl



**Site Editor**  
Ellen O'Brien

**Senior News Director**  
Dave Raffo

**Senior News Writer**  
Beth Pariseau

**Features Writer**  
Carol Sliwa

**Senior Managing Editor**  
Kim Hefner

**Associate Site Editor**  
Chris Griffin



**Site Editor**  
Susan Troy



**Site Editor**  
Andrew Burton

**Associate Site Editor**  
Heather Darcy

**Assistant Site Editor**  
Megan Kellet

**Features Writer**  
Todd Erickson

**Executive Editor and Independent Backup Expert**  
W. Curtis Preston

### Storage Decisions

**TechTarget Conferences**

**Director of Editorial Events**  
Lindsay Mullen

**Editorial Events Associate**  
Nicole Tierney

**Storage magazine**  
275 Grove Street  
Newton, MA 02466  
[editor@storagemagazine.com](mailto:editor@storagemagazine.com)

Subscriptions:  
[www.SearchStorage.com](http://www.SearchStorage.com)

A man with white hair and a mustache, wearing a white dress shirt and a red striped tie, is smiling and looking towards the camera. He is standing in a server room, with his hand resting on a computer monitor. The monitor displays a green terminal window with white text. In the background, there are server racks and other computer equipment. A blue rectangular box is overlaid on the top left of the image, containing text.

## “Working with FalconStor

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

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

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

**Contact FalconStor at 866-NOW-FALC (866-669-3252)  
or visit [www.falconstor.com](http://www.falconstor.com)**

**FalconStor**  
Software

# top new features in backup apps

Backup applications have evolved over the last few years to incorporate features that were previously available only in third-party products.

*By W. Curtis Preston*



**n**EW BACKUP TECHNOLOGIES are ready for mass adoption, and they're not just for early adopters. Early adopters helped give these technologies a jumpstart because they were comfortable purchasing products from startups and didn't think twice about being the first company on the block to try something new. But pioneers are typically a small contingent with many more potential users choosing a "wait and see" approach. So, even as some of these newer products achieve technological acclaim, they may barely make a dent in the overall backup market.

But recent events have accelerated the adoption—and perceived maturity—of some backup technologies. Smaller vendors have been acquired by their bigger brethren, and enabling technologies have emerged that ease the implementation of these products.

We'll look at five key backup technology advancements:

- Data deduplication
- Data protection management
- Continuous data protection
- Synthetic backups
- Virtual server backup

We'll explain how these technologies have changed backup, and describe how four of the leading backup software vendors—CommVault Systems Inc., EMC Corp., IBM Corp. and Symantec Corp.—have implemented these new backup technologies. That's not to suggest that other backup vendors don't also offer these features, but these four are among the acknowledged backup software market leaders with the products readers ask about most often.

### **DATA DEDUPLICATION: DISK BACKUP GAME CHANGER**

It's hard to overemphasize the importance of data deduplication in today's backup systems. It's perhaps the biggest game changer since the introduction of network backup systems 15 years ago, and its popularity can be traced to a number of factors. First, data deduplication enables users to increase disk utilization in their backup system. Tape had always been significantly cheaper than disk as a target for backups, and while the cost of disk has decreased significantly in the last several years, so has the cost of tape. So disk was typically used just as a staging mechanism for tape, rather than for long-term backup or archive storage.

Deduplication changed that forever. The random-access capabilities of disk allow data deduplication systems to remove redundant segments of data and replace them with pointers without significantly affecting restore performance. (While there's some performance degradation, restores are still much faster than when using tape.)

Despite dedupe's indisputable benefits, a lot of users waited to see if the techniques employed in target dedupe devices would eventually make their way into backup software, making such special-purpose appliances unnecessary. While most experts don't believe that target deduplication appliances are no longer necessary, data deduplication has, indeed, made its way into mainstream backup software products.

EMC and Symantec were the first major backup software companies to integrate deduplication into their product lines, and both did it through acquisition. EMC acquired Avamar Technologies, and Symantec's PureDisk product line resulted from its acquisition of Datacenter Technologies. CommVault and IBM chose to "roll their own" deduplication products.

EMC and Symantec both offer source deduplication products. That is, you can install the Avamar or PureDisk agent on a computer and the client will communicate with the backup server to identify and eliminate redundant data before it's transferred across the network. Only new

bytes are sent with each backup, which makes source deduplication perfect for smaller remote offices and mobile data.

Both vendors offer their source deduplication products as standalone products, which means you don't have to purchase Symantec's NetBackup or EMC's NetWorker. So even if you weren't using Symantec or EMC backup apps, you could take advantage of their deduplication technology. But if you wanted the functionality of both the backup app and dedupe, you had to purchase and manage two products (i.e., NetBackup and PureDisk, or NetWorker and Avamar). Symantec is the first to change this with NetBackup 7, which has built-in source dedupe that doesn't require a separate PureDisk installation. While you can manage Avamar via NetWorker, and a single install of their client software supports both NetWorker and Avamar backups, Avamar still requires a separate server to back up to.

Target deduplication is also available from backup software vendors. Symantec was the first to do this by allowing NetBackup customers to send standard NetBackup backups to a media server where they would be deduplicated by PureDisk. (With NetBackup 7, this functionality is available without requiring a separate PureDisk installation.)

IBM entered the data deduplication space with the introduction of its post-process target deduplication feature in Tivoli Storage Manager (TSM) 6.1. TSM can natively deduplicate its backups stored on disk after they have completed. IBM's target deduplication offering is unique in that it's included in the base product; however, the deduplication ratios it achieves may be relatively modest compared to those of other products' options that you have to pay for.

CommVault's Simpana deduplication facility is difficult to categorize as target or source dedupe. Deduplication in backup software requires multiple steps: (1) slicing files to be backed up into segments or "chunks"; (2) creating a "hash" value (typically using SHA-1); (3) doing a hash table lookup to see if the value is unique; and (4) deciding whether or not to send the chunk to storage. Source deduplication products perform all four steps on the client; target deduplication appliances do all four at the target or backup server. With CommVault's approach, however, steps one and two are done at the client while steps three and four are done at the backup server (media agent in CommVault lingo). This is why it's difficult to classify the dedupe as source or target.

But if the real distinction between the two categories is whether or not

**While you can manage Avamar via NetWorker, and a single install of their client software supports both NetWorker and Avamar backups, Avamar still requires a separate server to back up to.**

Is it the year of scale-out NAS?

Quality Awards V: Enterprise arrays

Virtualize disaster recovery

Top new backup app features

Intelligent tiered storage

the original, native data ever leaves the client, then CommVault Simpana is best placed in the target deduplication category. Still, Simpana's unique practice of doing the first two steps on the client allows it to do something other target products can't do: client-side compression. Most target dedupe systems won't deduplicate your data well if you compress it at the client before sending it to the target because compression inhibits the deduplication system's ability to correctly chunk and fingerprint the data to identify duplicates. But because Simpana chunks and fingerprints the data at the client, it can compress it before sending it across the network with no negative effects. The compression doesn't save as much bandwidth as source deduplication, but it can be advantageous in some environments.

### **DATA PROTECTION MANAGEMENT: BEYOND SIMPLE BACKUP STATS**

Data protection management (DPM) was introduced several years ago by Bocada Inc., the first company to attempt to produce standardized reports on multiple backup products. A number of other startup firms soon entered the fray, including Aptare Inc., Tek-Tools Software Inc. (recently acquired by SolarWinds, Inc.), TSMworks Inc., Servergraph (now part of Rocket Software Inc.) and WysDM Software (now part of EMC). The big backup software vendors saw the potential of the DPM market: Symantec picked up a product called Advanced Reporter, which became Veritas Backup Reporter and then later Symantec's OpsCenter Analytics line; and EMC turned the WysDM product into its Data Protection Advisor.

All of these products offer far more than simply telling you which backups worked and which didn't, functionality that many believe should be included in any decent backup software. However, when it comes to things like trending, capacity planning, cross-product reporting and issues that go beyond traditional backups, standalone DPM products have carved out a unique niche.

Backup apps have begun to incorporate some of these capabilities. CommVault, in particular, has been vocal about how these reporting tools should be included in the base backup product. While it could be argued that the reporting included in Simpana is better in some areas than the reporting in other companies' base products, that's not to say Simpana users couldn't benefit from a DPM product. For TSM customers, IBM's response has typically been that everything you need to know is

Data protection management (DPM) was introduced several years ago by Bocada Inc., the first company to attempt to produce standardized reports on multiple backup products.



## Save money



## Save energy



## Save all kinds of green

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

Visit [www.nexsan.com/green](http://www.nexsan.com/green) to download our Green Storage Guide and discover how to save all kinds of green without giving up a thing.



in the TSM database so you just have to run a query. While that's true, it might be beyond the capability of many users. So while a few of the big backup vendors have incorporated some DPM features, users who need full data protection management functionality will likely turn to a third-party product.

### **CONTINUOUS DATA PROTECTION: STILL KICKING**

Only a few years ago there were a number of companies with continuous data protection (CDP) applications, but many of them are no longer around. Some simply went out of business, while others were acquired in fire-sale deals. Did CDP simply not work? Was it a bad idea? Or was it the *Star Trek* of backup products (a great idea before its time)?

Did CDP simply not work? Was it a bad idea? Or was it the *Star Trek* of backup products (a great idea before its time)?

CDP's rise and fall was probably a combination of all of the above. When CDP works as advertised, it's easily the best way to protect your most critical applications: zero downtime for back-ups, and recovery time objectives (RTOs) and recovery point objectives (RPOs) of zero. What's not to like? Unfortunately, storage managers tend to be most risk averse when it comes to their mission-critical applications, so few users opted to back up those mission-critical applications using a completely different method from a vendor that they'd never heard of before.

But attitudes toward CDP changed when major companies got into the game. Symantec bought Revivio and eventually released NetBackup RealTime. IBM came out with Tivoli Continuous Data Protection for Files and bought FilesX, which became TSM FastBack. EMC purchased Kashya and delivered RecoverPoint. CommVault built its own CDP functionality around its core Common Technology Engine. With these key players in the CDP game, users can now try it in their own environments without the fear that their CDP vendor may go out of business tomorrow.

### **SYNTHETIC BACKUPS: NO MORE FULLS**

A long time ago, TSM developers asked a simple question: Why are we backing up data that hasn't changed? This became one of the core elements of TSM design and what TSM would eventually refer to as "progressive incremental" and others would call "incrementals forever." Once a given version of a file has been backed up, it's never backed up again.

Other backup products have chosen to use the traditional full/incremental approach to backups, also referred to as the grandfather-father-son method. But the question persisted: Why are we backing up data that hasn't changed? Eventually, CommVault, EMC and Symantec all came to the

Is it the year of scale-out NAS?

Quality Awards V: Enterprise arrays

Virtualize disaster recovery

Top new backup app features

Intelligent tiered storage

same conclusion: instead of transferring data that's already been backed up across the network, just transfer it from one tape to another within the backup server. Because 90% of any given full backup is already on tape or disk somewhere, a "synthetic full" can be created by copying the data that's needed from the latest full to a new full backup. This provides the benefit of a full backup (fast restores via collocation of the necessary

## CAN BACKUPS BE TURNED INTO ARCHIVES?

IBM Corp.'s Tivoli Storage Manager (TSM) has a backup feature where backups are copied to what is officially called a "backup set." IBM occasionally also calls a backup set an "instant archive." This seems to go against the usual mantra that backups aren't archives, and simply holding onto backups longer doesn't magically turn them into archives. So are TSM backup sets truly archives?

To answer this question, let's take a look at a new feature in Symantec Corp.'s Backup Exec 2010. Backup Exec incorporates Symantec's market-leading Enterprise Vault engine, so users can create archives of their backups by copying them into this engine. But Backup Exec does more than just copy the data from one tape format to another; it actually creates an index of the content of the archived files or applications. This means that you can perform Google-like searches against these archives by searching for phrases that might appear in files or Exchange emails, and Backup Exec will extract that data for you.

CommVault Systems Inc.'s Simpana also has the ability to perform content searches against its backups. You can search for files or emails based on a particular word or phrase. Like Symantec, they have a more full-featured archive product as well, but you can perform archive-like searches against their backups.

Let's contrast this to what TSM is doing. A TSM backup set actually has fewer database entries than a regular TSM backup; its purpose is to "archive" older files that you no longer have room for in the TSM database. So instead of having more context than regular backups, a TSM "instant archive" actually has less. While it's now possible with some products to "turn a backup into an archive," calling a TSM backup set an "instant archive" does a disservice to the word archive.

But that's not to suggest that TSM backup sets have no value. They do allow for longer retention than what's possible in the TSM database, and they also allow for restores without having to install TSM.

data) without the downside of a full backup (unnecessary transfer of the data across the network).

All three products have implemented the concept of the synthetic full in a slightly different way (CommVault and Symantec call synthetic fulls “synthetic backups,” while EMC uses the term “save set consolidation”). However, all of them share one critical concept. Once a synthetic full is created, it’s essentially just like any other full: it will be used for restores and later incremental backups will be based on that full. The previous full is only necessary if you’re keeping it for longer retention.

Once a synthetic full is created, it's essentially just like any other full: it will be used for restores and later incremental backups will be based on that full.

TSM users may feel that TSM’s concept of a backup set is very similar to a synthetic full, but it’s actually quite different. Unlike synthetic backups, the contents of a TSM backup set aren’t tracked in the backup database. In fact, one of the main purposes of a TSM backup set is to create an “instant archive” of backups that you wish to keep for a longer period of time than your TSM database has room for (see [“Can backups be turned into archives?”](#) p. 19). Another purpose for the TSM backup set is to create a backup that can be used outside of TSM; a TSM backup set can be read without the aid of the TSM catalog. If TSM backup sets were kept in the TSM database and usable for standard restores, then they would be the same as a synthetic full.

### **VIRTUAL SERVER BACKUP: GETTING EASIER**

Server virtualization has been a boon for many data centers. Far too many applications required a “dedicated server,” when all they truly needed was to *think* they had a dedicated server. Their CPU and I/O requirements were easily met by sharing resources with the aid of a server virtualization product. But then there was backup to consider.

While most applications could be easily virtualized, backup would “not go gentle into that good night.” It wanted—needed—the full resources of both a beefy CPU and beefy storage capable of heavy throughput. It’s been said that backups are a great way to test your storage and network systems because they have to move everything from point A to point B every night. Despite potential I/O issues, most users back up their virtual machines (VMs) by simply pretending they aren’t virtual. They load the backup client into the virtual machine and back it up just like a standalone server. VMware Inc. introduced VMware Consolidated Backup (VCB) to help ease the pain of VM backup and to remove the I/O issues from the ESX server, but it also increased the complexity of VM backups. It required two-step backups and two-step

restores for image backups, as well as the use of a separate disk-staging area. Not surprisingly, few users implemented VCB. Users of other virtual server apps, like Microsoft Corp.'s Hyper-V, also tended to back up their VMs by pretending they were physical servers.

The backup outlook is a lot brighter for both products: VMware introduced vSphere and Microsoft rolled out Hyper-V's backup architecture (which doesn't have an "official" brand name). VMware vStorage APIs for Data Protection (VADP) replaced VCB, offering everything that VCB promised and introducing the concept of block-level incremental backups. Now users can perform an image backup without having to copy the data to a staging disk first, and they can perform an incremental backup by simply having the backup application ask the vStorage APIs what blocks have changed since the last backup. The APIs promise to make things much better for those attempting to back up VMware virtual servers. The first major backup product to fully support vStorage APIs was EMC's Avamar, followed shortly by Symantec's NetBackup. As of this writing, CommVault, EMC NetWorker and IBM TSM are all working on their integration with vStorage APIs.

Microsoft Hyper-V users simply need to make sure that their backup product knows that it's talking to a Hyper-V server. Although not quite as advanced in some ways as vStorage APIs, it does a very similar job, allowing you to back up Hyper-V virtual machines without performing guest-level backup inside the virtual machine.

Hyper-V does have one advantage over VMware because it offers full Microsoft Volume Shadow Copy Service (VSS) support and VMware doesn't. Hyper-V uses VSS to quiesce its applications and notify them that the backup was successful. This allows Hyper-V users to get an application-consistent backup of any application inside a Windows VM without having to load an agent on that virtual machine. In addition, the application will know that it has been backed up and can clear its transaction logs.

VMware can quiesce applications in Windows 2003, but the actual operation it performs (VSS\_COPY) doesn't notify the application that it has been backed up; therefore, you must manage the transaction logs yourself. In addition, it currently has no application support for Windows 2008. As of this writing, VMware is working on this limitation, but the company offered no comment on the timing of the roadmap. This limitation has created an opportunity for backup products to differentiate themselves and, so far, FalconStor Software, NetApp, PHD Virtual Technologies (esXpress), Symantec (BackupExec) and Veeam Software all offer workarounds to address this limitation of VMware. ☉

---

**W. Curtis Preston is the executive editor for SearchStorage.com and an independent backup expert.**



# virtualize DISASTER RECOVERY

*Virtualizing storage and servers can make disaster recovery easier, more flexible and less expensive.*

*By Eric Siebert*

# d

**DESIGNING AND IMPLEMENTING** a disaster recovery (DR) infrastructure is often complicated, expensive and challenging. Virtualization technologies—for both storage and servers—can help reduce the expense with unique approaches that differ from traditional DR methods and can provide increased flexibility and responsiveness. Server virtualization encapsulates an entire server into a single file, which makes transporting it to other locations much easier. Storage virtualization presents multiple storage devices as a single storage resource, which helps hide some of the back-end complexities of the storage devices and network. Either of these virtualization technologies will ease the implementation of a DR plan; used together, they can provide a very effective DR strategy.

For most companies, the type of DR environment they devise is typically determined by balancing the amount of money they have to spend on one-time and ongoing costs, with the required recovery time to ensure that any downtime is limited and doesn't significantly impact their business. Traditional DR scenarios usually called for maintaining a lot of physical servers at an offsite location and then using

tape backups/restores or storage replication to transfer data between sites. With virtualization, there are more options for DR and the hardware requirements for the recovery site are greatly reduced. Even if your production data center hasn't been virtualized, you can still leverage virtualization at your remote location and convert your physical servers into virtual machines (VMs).

### **A VARIETY OF VIRTUALIZATION APPROACHES**

We'll look at some of the methods involving both server and storage virtualization that can be used as a foundation for a DR strategy. Our focus is on products and processes related to VMware Inc.'s vSphere, but many are very similar for other hypervisors like Citrix Systems Inc.'s XenServer and Microsoft Corp.'s Hyper-V. Depending on the virtualization methods used, recovery times can vary from seconds to hours to days and, accordingly, the cost and infrastructure to implement these methods will also vary. The approach you choose may be determined by whether you want a cold, warm or hot recovery site. Cold sites have no network connectivity with the main site, and limited or no hardware. Warm sites have network connectivity, and server and storage hardware, but typically lack real-time synchronization. Hot sites are almost mirror copies of critical production systems and use real-time synchronization for minimal disruption of services. The cost and recovery times differ greatly from cold to hot, but all of these types of sites can benefit greatly from using virtualization technology in their design and implementation.

## **ABOUT VMWARE VCENTER SITE RECOVERY MANAGER**

VMware Inc. developed its vCenter Site Recovery Manager (SRM) product to help automate and simplify the recovery process to a disaster recovery (DR) site. SRM by itself isn't a complete solution for disaster recovery and relies on a supported third-party array replication application to handle the replication of virtual machine (VM) data to a DR site. To certify that storage arrays are supported and integrated with vCenter SRM, VMware works with many storage vendors, including 3PAR, Compellent Technologies, Dell Inc., EMC Inc., FalconStor Software, Hewlett-Packard Co., Hitachi Data Systems, IBM Corp., NetApp Inc., Sun Microsystems (now owned by Oracle Corp.) and Xiotech Corp.

vCenter SRM allows you to create recovery plans using vCenter Server, extend recovery plans with custom scripts, perform nondisruptive testing, automate execution of recovery plans with a single command and reconfigure VM networking at the DR site. vCenter SRM provides a nice front-end application that both integrates storage replication with virtualization and automates DR failover in VMware environments.

## VIRTUAL MACHINE REPLICATION

Virtual machine replication works at the server virtualization layer and relies on replication software that can copy all changes made to a virtual machine disk file (VMDK) to another host. It requires a warm or hot DR site with dedicated network connectivity linking the production and recovery sites. A snapshot is taken of the VM at the virtualization layer, which deflects writes to the virtual disk to a separate delta file. The virtual disk is then mounted by the replication software and any updates since the last replication cycle are copied to another identical virtual disk on a virtual host at the disaster recovery site. VMware vSphere's new vStorage APIs enhance this process because of the new Changed Block Tracking (CBT) feature. CBT provides much quicker incremental backups and replications because the VMkernel tracks which disk blocks have changed since the last replication. This allows shorter intervals between replication operations, resulting in nearly continuous data protection (CDP). A big advantage of this method is that any type of storage can be used on both the source and target virtual hosts. When it's necessary to cut over to the DR site due to an outage at the main site, you can power on the replicated VM at the DR site and begin using it; changed blocks are then tracked on the remote site VM so they can be replicated to the main site for failback. Applications that support this method include:

- **Double-Take Software Inc.'s Double-Take Availability** can replicate both physical servers and VMs to a virtual host at a DR site. Replication can occur either inside the guest OS or at the virtual host level.
- **PHD Virtual Technologies' esXpress** combines disk-to-disk backup with replication; it can do a simple full-VM copy to another site or incremental block-level updates.
- **Veeam Software's Veeam Backup & Replication** combines disk-to-disk (D2D) backup and replication in one product. It has built-in data deduplication and uses CBT to achieve near CDP; changed blocks are injected into the target VMDK during each replication cycle.
- **VizionCore Inc.'s vReplicator** is a dedicated replication product for virtual machines (vRanger Pro is their backup product). It supports CBT and Active Block Mapping (ABM) to detect white space in a VM so it can be ignored.

Virtual machine replication works at the server virtualization layer and relies on replication software that can copy all changes made to a virtual machine disk file (VMDK) to another host.

## STORAGE REPLICATION

Storage replication works at the storage subsystem layer and is mostly transparent to virtual hosts and VMs. This approach relies on storage



hardware or software that can do synchronous or asynchronous replication from one storage device to another. Because it happens at the storage layer, the virtualization layer is unaware of the process and all virtual machine data is copied to the disaster recovery site where it will be ready to be used by the virtual hosts if needed. Storage replication requires significant network bandwidth between the main site and recovery site because of the large amounts of data that must be transferred quickly. Many vendors employ technologies such as data deduplication and compression to reduce the amount of data sent over the network. Storage replication is commonly used to achieve near-CDP or CDP to allow for very fast recoveries. VMware's vCenter Site Recovery Manager (SRM) product was designed to work with this method; it relies on storage replication to copy data between the two sites and SRM handles the

## THE IMPORTANCE OF PROPER QUIESCING

**WHATEVER DATA PROTECTION** method you use—storage replication, backups or virtual machine (VM) replication—it's very important that you properly quiesce your VMs to ensure data integrity. Quiescing is the process of pausing the operating system and applications, and forcing all pending data in memory to be written to disk. There are two ways to quiesce. The first way is done at the operating system level and it tells the OS to write all pending data in memory to disk; however, because the OS isn't aware of what applications are doing, this could cause corrupt or incomplete application data. The second method is done at the application level where applications like Microsoft Exchange or SQL Server are notified so that they can complete any pending transactions before writing the pending data to disk. The latter is called "application consistent" quiescing, and it ensures that all application data is properly backed up without any loss of data.

Without any quiescing, a VM is considered to be in a "crash consistent" state, meaning the backup that's made is of a VM that's been powered off with any held-in memory data not accounted for. Microsoft Corp. VM's have a special Volume Shadow Copy Service (VSS) driver built into the OS to quiesce the operating system, but it often won't provide application quiescing. To achieve application consistent backups, you may have to install a special driver inside the guest OS. When choosing any backup or replication application make sure it includes the proper quiescing for the critical application data you're trying to protect.



# Welcome. Step Inside.

**Quantum now provides data protection for a virtualized data center by introducing deduplication and replication into a VMware environment.**

Quantum DXi™-Series disk-based data deduplication appliance combined with Quantum's latest data protection solution for VMware® environments is comprehensive and scalable, yet easy-to-use. This revolutionary solution adds unprecedented efficiency to the backup process by eliminating the need for additional physical servers while utilizing the virtual environment itself to backup more data in less time. Run your backups across your entire global virtual and physical environments, with minimal impact to performance. Set it and forget it.

To find out how you can reduce your storage needs and streamline your backup process, visit Quantum at [www.quantum.com/virtualization](http://www.quantum.com/virtualization)

cutover to the DR site by bringing up the virtual machines at the DR site in case of a disruption at the main site (see [“About VMware vCenter Site Recovery Manager,” p. 23](#)). Most storage arrays either have replication built-in or available as a software add-on; a sampling of products that support this method include the following:

- **EMC Corp.** has a wide variety of products that support replication, including its entry-level Celerra Replicator and MirrorView products, and higher-end RecoverPoint (journal-based) and Symmetrix Remote Data Facility (SRDF) products.
- **FalconStor Software** offers Network Storage Server (NSS), a storage virtualization product that supports replication, as well as Continuous Data Protector, a high-end CDP product.
- **Hewlett-Packard (HP) Co.** builds replication into its StorageWorks EVA and XP disk arrays, and offers add-on products such as Business Copy, Cluster Extension and Continuous Access software for both the EVA and XP product lines.
- **Hitachi Data Systems** has both a journal-based replication product called Universal Replicator and a high-end CDP product, TrueCopy Remote Replication.
- **NetApp Inc.** provides an affordable replication option with Metro-Cluster, and SnapMirror is the high-end flagship replication product.

### **DISK AND TAPE BACKUPS**

While tape backups are used less frequently today for disaster recovery, they're still useful for storing data offsite in a secure location. The most effective way to back up a VM is to back up the single large virtual disk file (image level) at the virtualization layer, rather than the traditional method of using an agent inside the guest operating system (file level). Image-level backups are very useful for disaster recovery as they provide a bare-metal restore capability for virtual machines. Instead of having to restore physical servers one by one, you can restore them all to a single virtual host. While using tape for DR is slower than other alternatives, it's still a low-cost way to restore multiple virtual machines. A disk-to-disk recovery is much faster than tape, and is very similar to VM replication as a virtual machine's virtual disk is mounted and then copied to another disk storage device. But, unlike replication, this approach is usually run on a scheduled basis and can be done incrementally or as a full backup. The disk target that's used can then be backed up to tape or copied to a DR site and used to quickly restore virtual

While tape backups are used less frequently today for disaster recovery, they're still useful for storing data offsite in a secure location.

machines as needed. Some apps that support this method include:

- **EMC's Avamar Virtual Edition for VMware** supports backups of both physical servers and virtual machines by operating at the guest OS or VM layer, and can also globally dedupe backup data. It can also do physical-to-virtual (P2V) and virtual-to-physical (V2P) recovery for maximum flexibility.
- **PHD Virtual Technologies' esXpress** also does both backup and replication, providing data protection and business continuity.
- **Symantec Corp.'s NetBackup** has very good virtualization integration and supports both disk-to-disk and disk-to-tape backups. It supports both physical and virtual servers, and can perform both image- and file-level virtual machine backups.
- **Veeam Backup & Replication** provides disk-to-disk backup and takes advantage of many of the new features in vSphere.
- **VizionCore vRanger Pro** is VizionCore's dedicated disk-to-disk backup product for virtual machines, and it supports many of the same features as their vReplicator app.
- **VMware Data Recovery** is included with some vSphere editions. While not as feature-rich as other products, it does provide dedupe as well as good integration with vSphere.

### **SIMPLE AND BUILT-IN METHODS**

There are some very low-cost and simple alternatives for virtual DR, as well as some built-in tools in vSphere. At the most basic level, you can use scripts to take a snapshot of a VM's disk to deflect writes to it and then copy the data using FTP/SCP to another disk target such as a CIFS or NFS share. The disk target could be as basic as a removable hard disk that can be transported off-site or a device at a DR site that you copy to over a network connection. Once the virtual disk files are at the DR site, you load them on a virtual host and you'll be up and running. VMware vCenter Converter is another tool that can be used to copy a physical server or a virtual machine to either a disk target or a virtual host; it's not very sophisticated, but it can be scripted and scheduled to make copies of servers. vSphere has some built-in high-availability (HA) and fault-tolerance technology, as well as VMware VMotion. Those features currently all require a local-area network (LAN) and aren't suitable for long-distance wide-area network (WAN) use. VMware has announced its intention to enhance the features to function over slower WAN networks.

### **WATCH OUT FOR VIRTUALIZATIONS GOTCHAS**

Using virtualization technology as part of your DR plan has some great benefits, but there are also related challenges and costs. It's often assumed that server virtualization will save lots of money on server hardware. Lower operational costs will save money in the long run, but you'll have some additional up-front costs in addition to new physical servers.

For example, using two or three physical servers with virtualization at your DR site in place of eight to 10 physical servers at your main site will obviously reduce hardware costs. But you'll have to consider the cost of virtualization software, management and data protection applications.

If you're already using virtualization at your main site, using it at your disaster recovery site is an easy decision. If not, expect a learning curve in understanding how to properly implement, configure and manage it. Also, virtual machines usually require management and backup applications designed specifically for virtualization that may not work with physical servers. So, you might need separate tools for virtual and physical environments, which increases costs and management complexity. Some apps, like Microsoft System Center, can manage both environments via a single interface; similarly, Symantec's Net-Backup can back up both environments.

There are some clear advantages to using server virtualization at a DR site. Disaster recovery rack space is often expensive and with fewer racks your ongoing costs will be lower. Fewer physical servers also mean fewer network port requirements and less gear to maintain. You can also replicate VMs running on hosts with shared storage at your main site to hosts with direct-attached storage (DAS) at your DR site, which can result in more savings. Server virtualization allows

## DR APPS FOR MICROSOFT HYPER-V AND CITRIX SYSTEMS XENSERVER

**MANY OF THE** same disaster recovery (DR) principles for VMware implementation also apply to Microsoft Corp.'s Hyper-V and Citrix Systems Inc.'s XenServer. There are also some applications designed specifically for Hyper-V and XenServer that can be used to implement a DR solution for those environments.

- **Citrix Essentials for Hyper-V** is similar to VMware Inc.'s vCenter Site Recovery Manager (SRM). It includes the StorageLink Site Recovery application that can automate DR processes and failover. It also provides integration with array-based storage replication, as well as integration with Microsoft Systems Center. A version for XenServer is expected to be released sometime in 2010.

- **Marathon Technologies Corp.'s everRun** suite of products provides a range of high- and continuous-availability protection for XenServer. The suite allows you to mix and match physical and virtual servers for maximum flexibility.

- **Neverfail Ltd.** offers a suite of continuous availability products that provide replication for Hyper-V and VMware. The products support both physical and virtual servers.

- **SteelEye Technology Inc.'s SteelEye Protection Suite** provides replication for Hyper-V and XenServer, as well as ESX hosts, at the VM level and is storage agnostic; it also works across a LAN or WAN.

physical hardware independence, so you can use any type of server hardware at your DR site without having to worry about operating system and application compatibility.

### **VIRTUAL DR OPTIONS**

There are many options that you can choose from when using virtualization. The route you decide upon will likely be dictated by the amount of bandwidth available between your main and DR sites. The benefits of using virtualization as part of your disaster recovery setup include:

- Fewer physical servers needed at a DR site reduces one-time and ongoing costs, and results in less idle hardware
- Lower-cost VM-level replication is storage independent and doesn't require expensive storage arrays
- Hardware independence allows for more hardware options without compatibility issues
- Encapsulation turns a VM into a single portable file for easier transport and deployment
- Snapshots provide an effective method for backup of virtual machines
- Automated failover and easier testing
- Easier server deployment; scripting can be used to help automate many configuration and operational tasks

Virtualization can provide some clear advantages for disaster recovery; help save money, time and effort; and make the often daunting task of designing and implementing a DR plan easier. ☺

---

**Eric Siebert is an IT industry veteran with more than 25 years of experience who now focuses on server administration and virtualization. He's the author of *VMware VI3 Implementation and Administration* (Prentice Hall, 2009).**

## QUALITY AWARDS V: ENTERPRISE ARRAYS

# NetApp alone at the top

*After tying with EMC in the last Quality Awards, NetApp prevails this time to nudge out EMC and IBM, which tied for second.*

By Rich Castagna



NetApp Inc. continues to shed its “NAS-only” image with another impressive outing in the *Storage* magazine/ SearchStorage.com [Quality Awards for enterprise arrays](#). After a tie with EMC Corp. for the top spot in the last Quality Awards survey for this product category, NetApp is alone in the winner’s circle this time around.

NetApp’s win was by a slim margin—just 0.06 points—and just barely ahead of EMC and IBM Corp., which tied for second with an overall rating of 6.45 compared with NetApp’s 6.51. At the top of the enterprise array heap this time, NetApp’s ratings have surged, especially considering its last-place finish the first time its product lines were included in this category just a few surveys back.

Just like the fourth edition of the Quality Awards for enterprise arrays, the scores for the seven qualifying vendors (of eight included in the survey) were high across the board, with five of seven finalists garnering overall



ratings of 6.00-plus. Rounding out the field after the top three were Hitachi Data Systems (6.28), Hewlett-Packard (HP) Co. (6.20), 3PAR (5.87) and with a 5.78, Sun Microsystems Inc. (now owned by Oracle Corp.). While those are all solid scores, the results are just a bit less impressive than our previous survey, when six of seven vendors finished in the 6.00-plus points realm.

For all vendors/products, the average total score was a very solid 6.22, which was just a bit lower than the 6.39 racked up in the last enterprise arrays survey. But the message is clear: Users are pretty darned satisfied with their enterprise-class storage systems. And that appears to be a distinction the array vendors have earned; the results reported for the first two surveys for this product category (in 2005 and 2006) showed users to be considerably less enthusiastic about these products as they doled out sub-6.00 scores across the board.

### **IT STARTS WITH THE SALES FORCE**

The working relationship that a vendor's sales team establishes with a data storage manager is, of course, pivotal to whether a sale is made or not. But it can be just as important after the contract is signed, as the ongoing vendor-user relationship is often shaped during the sales process. As one survey respondent, an EMC and IBM user, put it, "I feel storage vendors should also guide their customers on how they can best configure the storage as per applications needs."

IBM led the field for the sales-force competence rating category, with a 6.30 rating that put it in front of NetApp (6.25) and EMC (6.21). IBM's win in this category hinged on its top scores for three key statements. The 6.58 IBM earned for "The vendor's sales support team is knowledgeable" suggests well-versed sales engineers, while they showed off their deal-making chops with a score of 6.37 for "My sales rep is flexible" and a 6.14 for "My sales rep is easy to negotiate with."

## **ABOUT THE SURVEY**



The *Storage* magazine/SearchStorage.com Quality Awards are designed to identify and recognize products that have proven their quality and reliability in actual use. The results are derived from a survey of qualified readers who assess products in five main categories: sales-force competence, initial product quality, product features, product reliability and technical support. Our methodology incorporates statistically valid polling that eliminates market share as a factor. Indeed, our objective is to identify the most reliable products on the market regardless of vendor name, reputation or size. Products were rated on a scale of 1.00 to 8.00, where 8.00 is the best score.

A total of 252 respondents provided 392 system evaluations.



Although it fell short of IBM's mark for sales-force competence by a mere 0.05 points, NetApp picked up scores of 6.00-plus for all six rating statements in this category; it was rated highest in the group for the statement "My sales rep keeps my interests foremost" (6.13). The tight top-three finish in the category was rounded out by EMC, which showed particular strength on statements related to familiarity with its customers' businesses and industries.

### **FIRST IMPRESSIONS: INITIAL PRODUCT QUALITY**

After making a six- or seven-figure purchase of an enterprise array, getting it up and running and in production quickly and without mishaps is critical for storage managers eager to see a return on their sizable investments. The vendors apparently aren't disappointing their customers, as six of our seven finalists all scored above 6.00 in the initial product quality category; in fact, the average of all scores in the category was an impressive 6.21.

NetApp led the parade of strong scores with a 6.52, followed by IBM (6.46) and HP (6.40).

NetApp's best scores were for ease of use, getting its gear up and running, and the level of professional services required for its products. IBM's good showing in this category was largely due to its 6.68 rating for the statement "I am satisfied with the level of professional services this product requires," but its other scores for initial product quality were all consistently high, too. HP, while not in the top three overall, nonetheless had a very good showing in the initial product quality ratings with scores higher than 6.00 for all statements in the category. HP fared best for the statements "This product was installed without any defects" (6.58) and "This product was easy to get up and running" (6.53).

The highest score (6.75) for any single statement in this category went to EMC for "This product was installed without any defects."

## **PRODUCTS IN THE SURVEY**

The following products were included in the Quality Awards V Enterprise Arrays survey (the number of ratings responses are in parentheses).


- 3PAR InServ Storage Server S400/S800 (14)
- EMC Symmetrix 3000 Series/8000 Series, DMX/DMX-3/DMX-4, V-Max (92)
- Fujitsu Eternus8000 Series/Eternus DX8400/DX8700\*
- Hewlett-Packard StorageWorks XP Series (70)
- Hitachi Data Systems Lightning 99xx Series or USP/USP V/USP VM Series (46)
- IBM DS8000 Series or XIV Storage System or ESS 800 Series (64)
- NetApp FAS3000/FAS6000 Series or V6000 (81)
- Sun StorageTek 99xx series (20)

*\*Did not receive enough responses to be included in the final results.*




- 3PAR InServ Storage Server S400/S800
- EMC Symmetrix 3000 Series/8000 Series, DMX/DMX-3/DMX-4, V-Max
- Hewlett-Packard StorageWorks XP Series
- Hitachi Data Systems Lightning 99xx Series or USP/USP V/USP VM Series
- IBM DS8000 Series, XIV Storage System or ESS 800 Series
- NetApp FAS3000/FAS6000 Series or V6000
- Sun StorageTek 99xx series


### OVERALL RANKINGS




### SALES-FORCE COMPETENCE




### PRODUCT FEATURES




### INITIAL PRODUCT QUALITY




### PRODUCT RELIABILITY



### TECHNICAL SUPPORT



### WOULD YOU BUY THIS PRODUCT AGAIN?



Based on a 1.00-8.00 scoring scale

# ENTERPRISE ARRAYS

Is it the year of scale-out NAS?

Quality Awards V: Enterprise arrays

Virtualize disaster recovery

Top new backup app features

Intelligent tiered storage

## PRODUCT FEATURES

A less-than-pleasant sales experience or rocky initial setup may be forgotten or overlooked if the storage system performs as advertised or even exceeds expectations. While it can be argued that the feature sets of enterprise arrays are often roughly equivalent in vendor-to-vendor comparisons, there's still a great deal of differentiation in how well those features can integrate and work in a given environment. Judging from the results in this category, the vendors appear to be meeting expectations: six of the seven companies received scores of 6.20 or higher, with the seventh, 3PAR, pulling down a very respectable 5.82.

NetApp came out on top of the product features category with a 6.76 rating, highlighted by the highest single score for all statements in the survey, a 7.11 for "This product's snapshot features meet my needs." In fact, NetApp snared the top marks for all but one statement in this category—"This product's capacity scales to meet my needs"—where EMC (6.82) prevailed over NetApp (6.65). EMC's ratings for product features were all quite high, totaling 6.55 for a second-place finish; IBM also fared very well with a 6.49 that placed it third.

General satisfaction with enterprise array feature sets was reflected in the scores for the statement "Overall, this product's features meet my needs," where all product lines received ratings of 6.00 or higher—a score that's likely to warm the hearts of both users and vendors.

## THE LONG RUN: PRODUCT RELIABILITY

Time is the truest test of any product, especially workhorse systems like enterprise storage systems. Reliability is measured on the Quality Awards surveys by how well the systems meet service requirements, how easily and nondisruptively they can be upgraded, and how long they can chug along without major hiccups. "Never been powered off since we got them—which is as it should be!" declared one EMC user responding to the survey. That enthusiastic endorsement was joined by many others to propel EMC to a win in the reliability category with a 6.89, the highest rating we've ever seen

## WHO DOES THE HEAVY LIFTING?

Average installed capacity of the Quality Awards V Enterprise Arrays.

| Vendor  | TB  |
|---|-----|
| EMC Corp.   | 132 |
| IBM Corp.   | 120 |
| NetApp Inc.                                       | 114 |
| Hitachi Data Systems                              | 112 |
| 3PAR  | 111 |
| Sun Microsystems Inc. (now owned by Oracle Corp.) | 98  |
| Hewlett-Packard Co.                               | 75  |

for enterprise array reliability. EMC flirted with scores of 7.00 for four of the five statements in this rating category and nailed it with a 7.02 for the critical “This product experiences very little down time” statement.

This category again featured solid scores for all vendors, with only two barely racking up ratings below 6.00. NetApp’s 6.66 was good for the second spot, and Hitachi Data Systems posted consistent scores across all statements to come in third with a category rating of 6.53. Hitachi users were succinct in their praise; one simply noted, “99.999% reliability” while another said “They just plain work.”

### **WHEN HELP IS NEEDED: TECHNICAL SUPPORT**

Back when we ran our first Quality Awards for enterprise arrays—five surveys ago in 2005—the top score for the technical support rating category was 5.76. Not only wouldn’t that score be a winner in the current survey, it would trail six other products. IBM picked up its second category win with a 6.51 rating as it battled once again with EMC (6.40) and NetApp (6.35).

For IBM, it was the third time in a row that the company had the highest score for technical support. It’s an impressive string of wins and a strong testament to IBM’s support teams and infrastructure. EMC and IBM, along with Hitachi Data Systems, had ratings of 6.00-plus for all of the category’s eight statements. NetApp came close to that level with only one rating barely missing the 6.00 mark—a 5.99 for “The vendor provides adequate training”—and matched its third-place finish from the previous enterprise arrays survey. “NetApp problems have been resolved before we were even aware there was a problem,” said one respondent. “Replacement drives show up before we know there has been a failure.”

### **ENTERPRISE ARRAY ENCORE**

In each survey we also ask respondents, empowered now with the gift of hindsight, if they would make the same storage system purchase again. Often, these “buy again” responses seem to run counter to the category and overall rating results, but this time NetApp capped its overall win with the highest percentage of respondents we’ve seen to date in this product category (91.4%) saying they would opt for NetApp again. Hitachi Data Systems, which came in fourth place overall, apparently has loyal users as well, with 87% saying they would buy its product again.

Taking the last two Quality Awards for enterprise arrays together, it’s clear storage vendors have made great progress on all fronts and that their customers appreciate the results. And considering the commitment those customers have made to land that big iron on their data centers’ floors, it’s sure to account for some peace of mind all around. ☺

---

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



## Could 2010 be a breakout year for scale-out NAS architectures?

*Scale-out NAS meets today's requirements for massively scalable and highly available systems, is cost effective and generally more efficient than traditional scale-up architectures. But technology change introduces risk, and companies may not be ready for a switch.*

**T**HE INFORMATION WE store today is very different from the information we stored a mere decade ago. Every endpoint device has become a content creation and capture device that has enabled faster and more efficient business processes while also driving massive unstructured data growth. Nowhere has the impact been felt more than in the data center storage domain. And it seems no industry is safe. Across the board, file formats are richer and file sizes are growing exponentially.

Using traditional scale-up architectures to address this growth is unrealistic. IT organizations need more efficient storage technology, and they're frustrated by the complexity of current offerings. An alternative approach, scale-out NAS, is poised for a breakout year. It

not only meets today's requirements for massively scalable and highly available systems, but does so cost effectively. It's generally more efficient than traditional scale-up architectures and reduces complexity because it can scale to multipetabyte capacities within a single namespace. In other words, it enables more capacity with far fewer systems.

With independent scaling of storage capacity, processors and bandwidth, users can grow scale-out NAS systems as needed without buying racks and power supplies in advance of capacity requirements or buying extra spindles to stripe files across. In effect, scale-out NAS provides "just-in-time" scalability. And with most scale-out systems, many low-level storage management tasks are automated, such as expanding the file system when new physical capacity is added and load balancing performance across processors, significantly reducing management costs.

Until recently, scale-out NAS has been tucked away in a corner, used mostly in niche markets such as high-performance computing (HPC),

**IT organizations need more efficient storage technology, and they're frustrated by the complexity of current offerings.**

scientific computing, and media and entertainment environments. Scale-out architectures were originally designed and tuned to support the bandwidth-intensive applications found in these verticals and, like many technologies that made their mark in the past, they're now finding their way into mainstream IT shops.

In addition, major storage vendors are now putting skin in the game. In 2009, Hewlett-Packard (HP) Co. bought Ibrix and introduced a new scale-out line; IBM ramped up the volume on its General Parallel File System (GPFS)-based scale-out file services and scale-out NAS appliance; NetApp Inc. introduced its Ontap 8 operating system that combines scale-out and scale-up modes; and Hitachi Data Systems expanded its BlueArc-based Hitachi NAS scale-out portfolio with the addition of the BlueArc Mercury product. Even smaller players—like Bycast Inc., Isilon Systems Inc. and Panasas Inc.—that focus on scale-out NAS in the niche markets where it's become mainstream are seeing more interest and traction in commercial IT.

There's also evidence that the increased use of collaboration technology in today's enterprises is favorably impacting scale-out NAS. In ESG's recent 2010 data center spending survey, 28% of organizations citing new collaborative tools and business processes using Web 2.0 technologies (for example, blogs, wikis and social networking services) as a business initiative that will have the greatest impact on IT spending over the next 12 to 18 months will make significant investments in scale-out system technology for rapidly growing unstructured content. Among organizations that don't view collaboration as a key business initiative, only 14% will make similar investments.

Despite the rosy outlook for scale-out NAS in 2010, the shift to scale-out in commercial enterprises won't be immediate or wholesale; it will be a journey that will take some time. One reason: Change introduces risk—mostly risk of the unknown—so IT organizations will take a cautious approach. Plus, introducing new storage systems in the enterprise means training users on managing the system and laying out new data protection methodologies that work with the new technology. And tier-1 applications with demanding performance requirements will continue to need dedicated systems to support transactional performance, a good fit for the continued use of scale-up systems.

**Despite the rosy outlook for scale-out NAS in 2010, the shift to scale-out in commercial enterprises won't be immediate or wholesale; it will be a journey that will take some time.**

Is it the year of scale-out NAS?

Quality Awards V: Enterprise arrays

Virtualize disaster recovery

Top new backup app features

Intelligent tiered storage

Managing data growth is an ongoing challenge for IT. It's also the "low hanging fruit" with which CIOs can make an impact and reduce both costs and cycle times. Keeping up with data growth has become an ever-more-costly effort, as it's been historically limited by traditional inefficient and complex-to-manage scale-up architectures as capacity needs increase. These limited-scale architectures have created an environment in which any changes, even simply provisioning more capacity, can take six months or more thanks to a lengthy change management process. Deploying new applications in this type of environment is a long, drawn-out process that limits a business' ability to respond to changing market conditions. IT's ability to respond to business needs must occur in real-time, which in turn will drive IT to look at deploying newer scale-out technologies that can provide a platform for business agility, consolidation, ease of use and availability. ☉

---

Terri McClure is a storage analyst at Enterprise Strategy Group, Milford, Mass.

# Beginning of the end for hard drives?

**SOLID-STATE STORAGE** is getting a lot of buzz these days, with enterprise data storage managers envisioning data centers minus the hum of spinning disks. But plenty of technologies have been tripped up by their own hype, never making the transition from cool to real. Solid-state storage appears to be overcoming that stumbling block right now. In our latest Snapshot survey, 34% of respondents said they're using some form of solid-state storage, and some have had it in their shops for years. To keep things in perspective, it's hardly a solid-state tsunami that we're seeing, as the average installation adds up to 1.5 TB. That's equivalent to approximately 94 16 Gig iPods, but it shows that solid state has gone from zero to 60 mph at a pretty snappy pace. The chief beef about solid-state storage is price, as well as some lingering doubts about a technology that's still relatively new to enterprises. But the early results look good: 80% of solid-state users are either satisfied or extremely satisfied with their flashy new storage. *—Rich Castagna*

Is it the year of scale-out NAS?

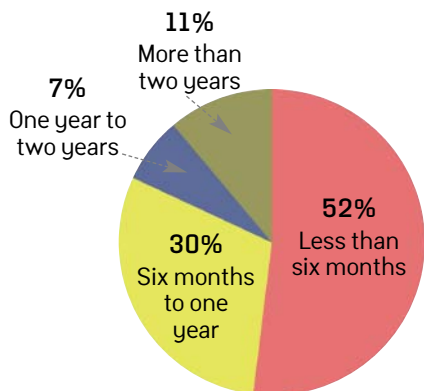
Quality Awards V: Enterprise arrays

Virtualize disaster recovery

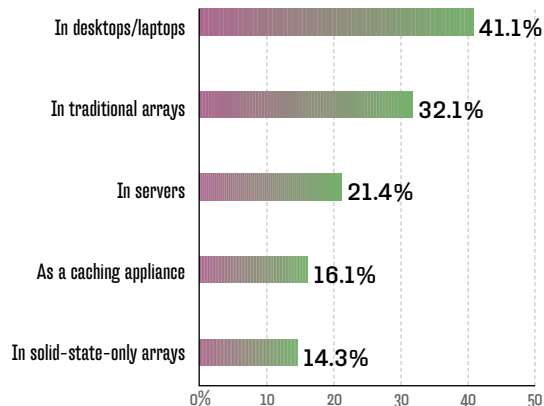
Top new backup app features

Intelligent tiered storage

How long have you been using solid-state storage in your shop?



How have you implemented solid-state storage?



What are the reasons why you don't use solid-state storage?

- 62%** Solid-state storage is too expensive
- 30%** Solid state is still too new and untested
- 28%** The performance of our hard disk systems is sufficient
- 23%** Solid-state storage capacities are too limited for our needs
- 21%** I don't know enough about solid-state storage at this time

**34**  
Percent of respondents already using solid-state storage

**“Hope the price comes down—we will use it in all the servers. The power consumption is great, and the speed is, too.”** *—Survey respondent*



## Check out the following resources from our sponsors:

### Dell Inc., page 4

[Building the Efficient Enterprise with Dell, Intel, and VMware vSphere 4](#)

[Comprehensive Data Protection & Disaster Recovery With Dell EqualLogic and VMware](#)

### FalconStor Software, page 12

[IDC Report: An ROI Analysis of FalconStor Storage Management Solutions](#)

[Book Chapter: SAN for Dummies -- Continuous Data Protection](#)

### Hitachi Data Systems, page 7

[Pile On the Savings: Free Storage Virtualization Software from Hitachi Data Systems](#)

[White Paper: The Economic Effects of Storage Virtualization](#)

### Nexsan Technologies, page 17

[The Guide to Green Storage -- Make an Ecological Difference](#)

[Making Cents of Tape vs. Disk](#)

### Pillar Data Systems, page 10

[Streamlining the Speed and Ease of VTL Deployment](#)

[Complimentary White Paper: Bringing an End to Storage Waste](#)

### Quantum Corp., page 26

[Optimize Your VMware Protection Strategy Leveraging Quantum Technology in a VMware Environment](#)

[Securing Data Stored on Tape With Encryption](#)