


STORAGE

Vol. 9 No. 4 June 2010

Solid State vs. Disk

On a dollars-per-gigabyte basis, solid-state storage may still be out of reach, but that may not be the best way of weighing its cost-effectiveness.

P. 14




ALSO INSIDE

- 5 Backup is too hard
- 8 Good match: NFS and VMware
- 22 Protect SharePoint data
- 30 Users pick top backup apps
- 39 Symantec's OpenStorage API sets the pace
- 42 VMs strain storage systems
- 45 DR testing gets better


STORAGE

sponsors | june 2010



REGIONAL SOLUTION PROVIDERS





Backup is Too Hard

5 EDITORIAL Storage environments are getting more and more complex, which makes data protection even tougher. While backup software does a lot more these days, it's still frustrating the hell out of users. *by RICH CASTAGNA*

NFS a Good Match for VMware

8 STORWARS With server virtualization proliferating, most companies are hooking those virtualized environments into Fibre Channel SANs; but there may be a better alternative. *by TONY ASARO*

Making the Case for Solid-State Storage

14 The future of storage is solid state, but it also has a place in today's storage shops. With a variety of solid-state implementations available, it's time to take a serious look at how the technology could enhance your storage environment. *by DENNIS MARTIN*

Backing up SharePoint

22 Microsoft SharePoint is gaining in popularity as a corporate collaboration tool—it's great for office efficiency, but tough on backups. *by W. CURTIS PRESTON*

Quality Awards V: CommVault Continues to Dominate the Enterprise, While Acronis Leads the Midrange Field

30 CommVault returns to the winner's circle for the fourth time in our fifth Quality Awards for backup and recovery software; Acronis emerges as a new winner among midrange backup applications. *by RICH CASTAGNA*

Are More Symantec OST-like APIs Needed?

39 HOT SPOTS One of the by-products of Symantec's OpenStorage interface is a performance improvement in backup and recovery operations. *by LAUREN WHITEHOUSE*

Virtual Servers and Their Impact on Data Storage


NEW 42 READ/WRITE We introduce a new column written by analysts from the Taneja Group. This month, senior analyst and consultant Jeff Byrne opines on how almost every IT shop has virtualized at least some of its servers, shifting much of the burden of server consolidation onto data storage systems. *by JEFF BYRNE*

Companies are Better Prepared for Disasters

45 SNAPSHOT More than 50% of the respondents to our disaster recovery (DR) survey are very confident that their DR plans will save their company from disaster. But one-third don't even test their plans. *by RICH CASTAGNA*

Vendor Resources

46 Useful links from our advertisers.



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. | —
Learn more @ www.getmorevirtual.com



Backup is too hard

Backup software does a lot these days, including frustrating the hell out of users.

WAS GOING to refrain from writing about backup this month because we have the subject pretty well covered in this issue with W. Curtis Preston's article on backing up SharePoint and the results of the latest Quality Awards survey on backup applications. But while digging through that survey's data and compiling the results, I was struck by the comments made by some of the respondents. Approximately 65 of the 350-plus survey respondents took an extra minute or two to add a little narrative to their response, and of those 65, I counted fewer than 10 that were unqualified endorsements of a product. The rest were either flat-out pans of the app ("A royal pain in the __ to use" and "Never seems to work right") or tempered praise ("Good product; however, it requires more staff to manage well than I have on the team").

There's no question that backup keeps getting tougher for storage managers, with most companies adding tens of terabytes (or more) of data each year and changes to the data center environment—like virtualized servers—having an impact as well. Storage environments are growing more and more complex, and they can vary tremendously from one company to another. So, it's hardly reasonable to expect backup application vendors to write software that can address all conditions in all environments. But there are some things that are driving data storage managers and backup admins up a wall that vendors can address; and, in some cases, the solutions can be relatively simple.

But some backup application vendors still seem intent on trying to meet everyone's needs by building out huge, multifunctional backup applications that are designed to be all things to all users. It's a nice idea, but it rarely works out well. Users end up grappling with big, unwieldy apps as they try to pick their way through to find the features that could be useful in their shops. And bigger means more updates for users to deal with ("Too many patches and updates; very bloated software") and the likelihood for more programming miscues ("Upgrades are always a problem").

To be fair, vendors build in new functionality because their customers ask for new capabilities. But there should be a better process for upgrading apps,

Users end up grappling with big, unwieldy apps as they try to pick their way through to find the features that could be useful in their shops.

one that's not so disruptive. And I'm sure that some new features are shoved out the door before the code is fully baked—an age-old issue, for sure, but one that deserves some attention.

The problems of big apps and confusing operations can be debated, but the biggest bones of contention among our survey takers are (or should be) the cut-and-dried issues of licensing and support. Users are fed up with complex and often exorbitant licensing schemes, as well as inadequate support. How fed up are they? "They are damaging the industry [with] their licensing model and general mafia attitude," wrote one respondent. That might be a little extreme—I don't think backup software vendors are going around wearing shiny suits and breaking the legs of backup admins—but it certainly suggests that there's an awful lot of frustration out there. Another respondent noted that his backup app's "licensing model is very restrictive and expensive." Still another lamented that the "license model, license costs and annual maintenance are all too complex and costly for a business our size."

On the support side, storage pros struggle to learn these complex products and then feel cut adrift when they seek help. One respondent likes his backup application but said it was "a bit hard to learn"; another said his company needs "some training to use it better." Even after climbing that learning curve, help seems hard to find. "I've contacted support twice in the last five years and neither time were they able to help resolve my issues." And this one's my favorite: "Their support is good when you have just started out and don't know how to do things, but very poor when you have an actual problem."

C'mon, backup vendors, it's time to address the user issues that you, and *only* you, control: licensing and support. Fix those problems and you'll have a lot more satisfied customers. ☺

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 July/August 2010 issue.](#)

Users are fed up with complex and often exorbitant licensing schemes, as well as inadequate support.



Pillar Axiom with SSD

Racing Harness Not Included.



You don't have to put up with oversized, underutilized, energy-guzzling storage systems. End their wasteful ways that are slowing your service and busting your budget.

Make the move to Pillar Axiom[®], the only true Application-Aware Storage[™] system – now with SSD. Reap the combined benefits of the industry's highest performance and utilization rates, and the industry's best \$/IOPS. Guaranteed. You'll wait less and waste less, and have your applications screaming, instead of end users.

It's time to Stop Storage Waste and see how fast you can go with Pillar.

www.pillardata.com



Special Offer:

Download a complimentary white paper:
Accelerating and Managing Data
Using Solid State Drives.

www.pillardata.com/ssd





NFS a good match for VMware

With server virtualization proliferating, most companies are hooking those virtualized environments into Fibre Channel SANs; but is that the best choice?

W

WE DID RESEARCH over the last 18 months asking IT professionals about their major priorities and all of them emphatically stated that implementing VMware was one of their top initiatives to significantly drive down costs and improve operations. But these IT professionals also said the biggest challenge they face with their VMware environments is networked storage. With dedicated data storage managers buried under a mountain of projects and daily operations that consume so much of their time, taking on new tasks is nearly impossible.

Another important dynamic we've observed is that VMware administrators find themselves becoming quasi-storage administrators, which is somewhat akin to asking a brain surgeon to operate on someone's heart. This isn't due simply to the inherent complexity of most networked storage systems; it's an issue of specialization as well.

Ironically NFS, the protocol of geeks, is a means of simplifying VMware and networked storage. It turns out there are a number of shops using NFS with VMware and, as a result, they've eliminated much of the day-to-day complexity of managing network storage with VMware. Here are some of the advantages of using NFS with VMware rather than Fibre Channel (FC) or iSCSI SANs:

- It's very simple to add NFS datastores. With NFS there's no LUN management. This simplicity addresses configuration issues that might come back to bite you in a SAN environment. The storage provisioning process for SANs often requires a dozen or more time-consuming steps. Missing any one step may result in disaster and you might not know about it until it happens.
- You don't have to deal with all of the complexity of Fibre Channel, WWNs, zones, ISLs, etc.
- You can easily increase—or more importantly—decrease the size

With dedicated data storage managers buried under a mountain of projects and daily operations that consume so much of their time, taking on new tasks is nearly impossible.

of NFS datastores online for capacity reclamation.

- Because you're using NFS, you don't have to deal with VMFS or raw device mappings (RDMs). This means you can have volumes bigger than 2 TB. NetApp supports a 16 TB file system, and Isilon Systems and BlueArc can support hundreds of terabytes in a datastore (and, theoretically, even more).
- There's no single disk I/O queue with NFS, which means performance is dependent on network bandwidth and the data storage system itself. It also means NFS performance can even keep up with Fibre Channel in VMware environments.
 - You can back up whole virtual machines (VMs) or files within VMs.
 - Restoration of VMs is flexible, including individual VMs, multiple VMs or files within VMs.
 - The cost of Fibre Channel is higher, including the equipment and support. This is an argument for both NFS and iSCSI over FC.
 - One user pointed out that the tools to troubleshoot IP networks are much better than those for FC, another advantage for both NFS and iSCSI.
 - SAN expertise is more specialized than IP; it's harder to find and retain experts (and they're usually paid more).

However, there aren't many vendors offering storage systems that support NFS. There are literally dozens of SAN-based storage systems, but only a handful of NFS or NAS storage systems in the market. Customers want choices, and the fewer there are the less likely they will go down a particular path. And these are good reasons because competition fosters innovation, cost-effectiveness and better service.

NFS for VMware has the potential of changing the data storage landscape, but sadly, it probably won't. It's a major challenge to educate the market; there's an enormous amount of incumbent SAN storage systems; and there's no one taking up the mantle to fight this good fight. Even NetApp has lost its fervor for NFS and has instead taken a "we provide whatever protocol you want" attitude, responding to the market vs. driving it. That isn't to say NetApp isn't promoting the use of NFS with VMware, but it's not a core strategy and is only one of many things the company talks about. BlueArc and Isilon Systems are focusing on NFS for VMware, but they don't have the same resources or generate the same level of awareness as the big storage vendors. However, if NFS for VMware is going to take off, these two vendors are sowing some seeds and will be in a position to reap the rewards.

NFS is a more highly virtualized protocol than Fibre Channel and

Customers want choices, and the fewer there are the less likely they will go down a particular path.

even iSCSI. As a result, it works much more easily and efficiently with VMware. If NFS is to become the dominant VMware protocol, it's probably not the storage specialists that will make this happen but rather the VMware administrators who want highly virtualized networked storage without having to be a storage expert. The implications of this are significant as VMware continues to proliferate and those who manage these environments will have influence as they decide what they need to be successful. ☉

Tony Asaro is senior analyst and founder of [Voices of IT](#).

VMs strain
storage systemsUsers pick
top backup apps

Protect SharePoint data

Justifying
solid-state storageGood match:
NFS and VMware



Having trouble deciding whether to migrate to the cloud?

Iron Mountain Digital can help. We offer a rich set of on-premises, hybrid and cloud information management solutions for archiving, eDiscovery, data protection and compliance. We can help you select the right solutions to meet your specific needs so you can migrate to the cloud when and if you're ready.

Iron Mountain has a proven track record of intelligently managing enterprise information and pioneered cloud information management over 15 years ago. Our underground datacenters provide an unprecedented level of security, reliability and protection. Trusted by over 90% of the Fortune 1000, with 140,000 clients in more than 40 countries, we go to great lengths to protect our customers' information, their reputation and ours.

Email digital-info@ironmountain.com or visit www.ironmountain.com/cloudhelp to schedule an assessment of your information management needs.



STORAGE

COMING IN JULY/AUGUST

iSCSI Storage and vSphere

This step-by-step guide will walk you through setting up and configuring iSCSI storage for vSphere virtual machines. We also include best practices for using iSCSI storage, and compare the performance of using hardware initiators vs. software initiators, jumbo frames vs. no jumbo frames, and using paravirtualization vs. normal virtual SCSI adapters.

Continuous Data Protection Redux

Continuous data protection (CDP) products may provide the best data protection compared to other available technologies because they capture new data and changes to older data as they're committed and save them to a backup target. There's renewed interest in CDP, rekindled by virtual server environments and its integration with other backup apps, making it easier to implement and manage.

How to Really Use ITIL to Manage Your Storage Better

The Information Technology Infrastructure Library (ITIL) is a set of guidelines, processes and best practices for managing IT resources and services. For many storage managers, ITIL is a foreign concept, but in this article, Tom Woods, ITIL global service transition manager at Ford Motor Co., will describe the practical application of ITIL to storage management.

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
Greta Lindberg

Storage magazine
275 Grove Street
Newton, MA 02466
editor@storagemagazine.com

Subscriptions:
www.SearchStorage.com



Does Your Storage Need **Liposuction**?

Save **50%** with a 3PAR Storage Technology Refresh—**Guaranteed!***

There is no better time to make the move to 3PAR®. Simply purchase any 3PAR InServ® Storage Server with Thin Built In™ and use 3PAR Thin Conversion to non-disruptively “liposuction the fat” from fully allocated volumes on traditional storage—transforming them into new “thin” volumes on the InServ.

We guarantee that you’ll reduce your current capacity requirements by 50% or more, or we’ll make up the difference with free disk capacity and associated software and support.*

What could be simpler?

Get the 3PAR Get Thin Guarantee today by visiting www.3PAR.com/GetThin

*The above is intended to highlight certain aspects of our Get Thin Guarantee and does not contain the full terms, conditions, limitations, definitions, and other provisions (“Terms”) of the Get Thin Guarantee. The Terms shall be contained in a written Get Thin Offer which shall take precedence over the above. Qualification for our Get Thin Guarantee is subject to your acceptance of a Get Thin Offer containing the Terms and satisfaction of those Terms.



Get Agile. Get Efficient. Get 3PAR.

MAKING THE CASE FOR SOLID-STATE STORAGE

Interest in solid-state storage is high, and with a variety of solid-state implementations available and newer technologies emerging, it's time to take a serious look at how solid state could enhance your storage environment.

By Dennis Martin

DATA STORAGE PROFESSIONALS considering solid-state storage have myriad solid-state storage architectures to consider, including systems that use solid-state drives (SSDs) in various form factors, caching implementations and appliances. If that wasn't enough to ponder, those planning on implementing these systems need to decide whether to use a product that mixes solid-state storage and traditional disk drives, or to use SSD-only storage subsystems.

But perhaps more important than just choosing the hardware, enterprises need to decide what data to put on solid-state storage or consider using

some form of software automation to move the data onto solid-state storage to make the most efficient use of what is still a somewhat expensive resource. Deciding what data to place on solid-state storage and how to put it there makes choosing a solid-state storage option more complex, but your selections will have a long-term impact.

SOLID-STATE-ONLY SHOPS: NOT SO SOON

In a few decades, some form of solid-state storage may be the dominant and possibly only form of enterprise data storage. But given the present state of matters, that day is (at best) on the distant horizon. We might dream of replacing all of our electro-mechanical disk drives with solid-state storage if cost weren't a factor, but there's nowhere near enough semiconductor fabrication production capacity available today to satisfy the total storage capacity that's deployed in IT shops.

But there are some promising signs. Enterprise solid-state storage prices are dropping relative to enterprise hard disk drives (HDDs). Not that long ago, enterprise solid-state storage was as much as 40 times the price of an equivalent capacity of enterprise hard disk drive storage. The price comparison ratios are in the neighborhood of 25% to 50% of that today, depending on specific solid-state storage products.

As a result of this pricing and capacity disparity, data storage managers and administrators are finding that solid-state storage complements existing traditional forms of storage. They've deployed, or are planning to deploy, solid-state storage where high performance, low latency

FORM FACTORS AND INTERFACES

Solid-state storage comes in a variety of form factors, including nearly all the disk-drive form factors, as internal modules within a storage system or as a PCI Express bus card. The PCI Express bus form factor provides the potential for very high bandwidth storage access within a server or workstation.

Enterprise solid-state drives are available in 2.5-inch and 3.5-inch drive form factors that are compatible with today's servers and storage systems. The primary interfaces for these are SATA, SAS and Fibre Channel (FC). The SATA interface is available for many solid-state drives, especially for the consumer and desktop market. Fibre Channel has a long future as a SAN interface, but is approaching end-of-life as a disk drive interface. Disk drive suppliers and solid-state storage suppliers are moving away from Fibre Channel as a drive interface in favor of 6 Gbps SAS as an enterprise drive interface. We expect the Fibre Channel interface on 3.5-inch drives to stick around for a while to maintain spare parts on the relatively large number of 3.5-inch FC drives in enterprise disk subsystems. And we also anticipate that relatively few 2.5-inch enterprise drives will have a Fibre Channel interface.

or energy savings are needed.

There are two basic ways to implement solid-state storage technology:

- Use solid-state storage directly as a primary store
- Use solid-state storage as a cache in front of spinning disks

Each of these implementations has its advantages and disadvantages, and implementations vary among storage vendors. And some vendors offer one implementation now while planning to offer the other in the next six to 12 months.

USING SOLID STATE FOR PRIMARY STORAGE

For vendors that implement solid-state storage directly as a primary data store, many use the standard disk-drive form factor. This implementation method is simple to understand and is compatible with current subsystem designs and configuration processes. The one downside to this approach is that many of today's controllers and subsystems weren't designed for disk drives with an order of magnitude of faster performance at the drive level, so vendors typically don't support a large system completely full of solid-state disk drives. But this is changing as vendors design and build improved controllers that can handle many more solid-state drives. The good news is that significant performance gains can be achieved with a relatively small number of SSDs, often only one full or partial drive shelf. Some users are reporting five to eight times performance gains for some workloads with a relatively small amount of solid-state storage.

We're also seeing an increasing number of solid-state-only storage products available today and planned for release over the next several months. These systems are designed to use solid-state storage as the primary store, with capacities in the single- or double-digit terabytes today and larger capacities coming soon.

For users who have implemented solid-state storage as a primary store, the big question focuses on what data to put on the solid-state storage. There are some obvious candidates, such as database indexes, heavily accessed database tables or temporary scratch areas, log files or any other hot spot. However, this is often not a static solution. Some data that's hot today may not be hot tomorrow. So storage administrators, database administrators or other IT technicians may have to continually monitor data usage patterns and be prepared to make adjustments on a fairly regular basis. In some cases, this increased management burden may be too much work and operational expense to be worth the tradeoff for increased I/O performance.

For users who have implemented solid-state storage as a primary store, the big question focuses on what data to put on the solid-state storage.



Data Proliferation Creating Backup Problems?

Stop counting.

We've got everything you need to handle your data, no matter how fast it multiplies.

To get your free "Data Proliferation" poster, go to <http://offers.cptech.com/rabbits>



The answer is to provide an automated way for the storage system to identify the hot data and move it onto the solid-state storage automatically, then move it to slower storage when it no longer requires solid-state performance. Many vendors provide forms of tiering software that does exactly that. This software observes the I/O patterns for a time and then moves the data in a way that's transparent to the host applications. Many of these automated solutions allow the administrator to determine what activity level defines "hot" data, set the time period over which the observations are made, and then set a separate parameter that controls the frequency of data movement (anywhere from hourly to weekly). Some of this software has the ability to make recommendations about the data tiering based on the observations it has made, such as recommending a 10%/90% mix of solid state vs. spinning disk. Today, many of these automated data movers perform the data movement at the LUN level; sub-LUN-level data movement is expected from several vendors within the next six to 12 months.

The solid-state-only storage products eliminate the need to move data from faster to slower storage because all of the data is on fast storage. These systems appeal to customers who want to put an entire application and its data on solid-state storage. At today's price points, these solutions tend to be deployed for critical applications only. The decision (and budget) to acquire them tends to come from line-of-business owners or architects rather than from the IT department.

The solid-state-only storage solutions eliminate the need to move data from faster to slower storage because all of the data is on fast storage.

CACHING WITH SOLID STATE

The other basic implementation is to use solid-state storage as a cache in front of spinning disks. This method has the advantage of always accelerating the hot data in real-time, since only the hot data is likely to be in cache. And because the solid-state storage is acting as a cache, there's no need for an administrator to decide what data should be placed on it. The basic questions here are what size cache is appropriate and which workloads should be directed toward the cache to make the best use of the solid-state device.

Some solid-state caching solutions are built into existing storage systems, while others are delivered as external appliances. Adding flash memory as a cache inside a storage subsystem in effect provides a "level 2" cache not unlike the L2 cache found on many processors inside today's computers. This added cache capacity improves performance for most if not all operations. In addition, because flash memory is non-volatile, this cache provides some extra protection in the event

VMs strain storage systems

Users pick top backup apps

Protect SharePoint data

Justifying solid-state storage

Good match: NFS and VMware

of power loss. But issues such as cache coherency, and whether the cache is DRAM based or flash memory based, remain. Generally, a cache is tied to one processor or controller, and there are various cache management functions that can be applied to allow caches to work properly with multiple processors or controllers. In addition, storage systems that use caching can add special features to their internal OSes that are aware of the cache and can provide additional flexibility, such as the ability to assign different I/O priorities for I/O going to different volumes on the storage system.

The caching appliances add the benefits of cache without requiring changes to any existing servers or storage systems. These appliances fit easily into the storage network and can accelerate all I/O going through them, even sending data to different storage subsystems at the same time. Many of the appliances can be set to write-back, write-through or pass-through for any given volume they accelerate. Some of the caching appliances are constructed in such a way as to allow their memory modules to be hot-plugged, so maintenance or growth can occur without taking down the entire appliance.

The big question for a caching implementation is how much cache is enough. For many workloads and applications, a relatively small amount of cache (5% to 20%) relative to the total storage allocated to that application is enough to provide significant performance improvements. For other workloads, the cache needs to be large enough to hold the entire volume to achieve appreciable performance gains.

IT'S ALL ABOUT PERFORMANCE

Solid-state storage, however it's deployed, offers the promise of significant performance gains. We've seen results of seven to nine times overall performance gains in our lab testing for various real-world applications (email, database, etc.) when configured optimally for the application.

With performance gains of that magnitude possible, what's not to like? Certainly, pricing is a factor. However, consider some of the current methods that are used to increase performance for spinning disk drives, such as "short stroking" spinning disk drives. Short stroking spreads data over many disk drives by using only a portion of the capacity of each drive for data, so that as many "spindles" as possible can be applied to improve performance. To achieve desired performance goals, some users short stroke some of their enterprise disk drives using ratios of 7:1, 8:1 or 9:1, which means they're using only 1/7th, 1/8th or 1/9th of the available capacity on each drive. If the price of an enterprise SSD is 10 to 15

The caching appliances add the benefits of cache without requiring changes to any existing servers or storage systems.

times the price of the spinning drives being short stroked, it may make sense to move that application data to enterprise SSDs and get the required performance while using much less power and space.

Almost all data storage system vendors now offer configurations that use a combination of solid-state storage and enterprise SATA storage instead of arrays full of enterprise spinning disk drives. These new configurations typically offer higher performance, equivalent capacity, lower power consumption, smaller space requirements and lower total hardware costs. ☉

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



“In our London headquarters, recovery of key systems is now proven to an RTO of 15 minutes and full site recovery in four hours.”

Malcolm Todd, Head of Systems Delivery, Norton Rose LLP

“EMC deduplication storage provides superb reliability. Our backups and restores are completed with complete confidence.”

Service Delivery Manager, Ordnance Survey

“Unlike our older backup and disaster recovery processes, the EMC backup and recovery solution is fully integrated into our IT infrastructure. This provides us with much higher backup speeds, as well as quick data recovery and easy management.”

Benoit Perrussel, System Networks Manager for Honda France Manufacturing

Next-generation backup starts now.

No one instills more confidence in backup and recovery than EMC. Streamline operations, simplify management, and control costs now—with the leader in disk-based backup and recovery.

Next starts now. To learn more visit www.EMC.com/uk

BACKING UP SHAREPOINT

Microsoft SharePoint is gaining in popularity as a corporate collaboration tool; it's great for office efficiency, but tough on backups.

By W. Curtis Preston



Microsoft Office SharePoint Server is an interesting suite of applications. I've heard from a number of users who extol its collaboration methods, but from a backup perspective, SharePoint is somewhat analogous to VMware and other server virtualization technologies. SharePoint may be solving the world's problems, but how do you back this thing up? (Interestingly, the backup solutions for SharePoint and VMware are eerily similar; more on that later.)

The challenge with backing up SharePoint is that it's not just one application, but a suite of applications that work together. Each SharePoint portal consists of one or more Web servers, application servers, query servers and index servers, all of which store their data in multiple SQL Server databases (at minimum, one content database and one configuration database). In a very small environment these can all be placed on a single physical server, but they're typically configured across multiple servers to provide some scalability.

A LOOK INSIDE SHAREPOINT

SharePoint's configuration database obviously stores the configuration of SharePoint itself, including such things as:

- Internet Information Services settings, including IP addresses and Secure Sockets Layer (SSL) certificates
- Service accounts that are used to run various services such as search
- Search, connection, workflow, email, antivirus and logging settings
- Recycle bin settings, such as whether to have a multilevel recycle bin to protect against accidental deletion

Closely related to the configuration database is the administration database. Both of these databases are extremely important, which is why it's so surprising that most of Microsoft's built-in backup methods don't support restoring them. Oddly enough, they support backing them up, but restoring them isn't supported (see "[What's up with SharePoint's configuration database?](#)", p. 24).

In addition, some of SharePoint's customization is stored in files in the file system, not in a database at all. This means you have to back up both databases and file system data to fully back up SharePoint.

The content database is where all of SharePoint's collaborative content is stored. This includes Microsoft Office documents (e.g., Word, PowerPoint, Excel) and any communication related to those documents. One of the interesting things about how the SharePoint content database works is that as users share documents and store multiple versions of the same document in SharePoint, they significantly increase the amount of storage needed for their database.

Consider what you would do without SharePoint. You would put the file on a file share and turn Track Changes on. When you finished working on the file, you would send an email to your co-workers to take a look at it. They would review the document, make their changes and save them to that document. Track Changes keeps a record of all the edits and additions, and you didn't have to make a separate copy of the document. But SharePoint stores every version of the document, and it doesn't have deduplication enabled. This is an important point because if you thought deduping Exchange data received good data deduplication ratios, you're going to love the ones you get from SharePoint. (While we're focused here on backing up SharePoint, its versioning process also makes it a good candidate for primary storage deduplication.)

When planning your SharePoint backup and recovery system, you're

The content database is where all of SharePoint's collaborative content is stored.

obviously going to want to be familiar with the content databases, configuration databases and any other databases that are part of your SharePoint configuration. You also need to think about what you want to recover because the different backup and recovery options allow you to do things at different levels. In addition, some of the options allow you to recover at lower levels of granularity than others. A good place to start is with this [Microsoft TechNet article](#) that explains in detail the capabilities of the various backup and recovery options described here. The article focuses on using Microsoft tools like Data Protection Manager (DPM), but discusses other options as well.

NATIVE BACKUP AND RECOVERY OPTIONS

The following is a summary of the backup and recovery options that are available free of charge with any SharePoint installation.

SharePoint Central Administration. This is a GUI option available when running SharePoint Central Administration. While it can back up the entire site, it has three very big limitations: It doesn't have scheduling capabilities; it can't be used to restore the configuration or administration databases; and it can't back up site collections.

SharePoint stasdm.exe Command Line. The command line utility

WHAT'S UP WITH SHAREPOINT'S CONFIGURATION DATABASE?

None of the native backup and recovery tools for SharePoint support backing up and restoring the configuration and administration databases from a live system. The reason Microsoft Corp. gives for this is that these databases must be restored to the same point in time when the other databases were backed up or the results could be unpredictable. Therefore, while the native tools often allow you to back up the whole site, they can't (or shouldn't) be used to restore these databases because these tools have no facility for assuring that they're being restored to the same point in time. For more information on this oddity, please consult [this article](#).

The only supported way to solve this problem with native tools is to restore from a backup of a fully stopped farm. The procedure in [this article](#), which is about how to move all the databases from one server to another, can be used to back up and recover the entire site. (As long as you're not using single sign-on [SSO] databases, in which case you have to handle it separately using [this SSO procedure](#).)

And then, of course, there's the configuration and customization information that's not stored in the database at all. Those files need to be backed up and restored at the same time as the others. No wonder Microsoft says, "It's not supported." They even go so far as to tell you that you should document your configuration changes so that you can redo them, because there's a good chance that you won't be able to recover the configuration and administration databases.



The World's Fastest No Regrets SSDs

NEW


RamSan-630 SSD SAN, NAS

- + 4 to 10 TB SLC Flash & Direct Attach
- + Reads & Writes 500,000 IOPS
- + Throughput up to 8GB/s
- + 450 Watts
- + Fibre Channel



RamSan-20 SSD Workstations

- + 225 to 450 GB SLC Flash Servers
- + Reads/Writes 120,000/50,000 IOPS
- + Throughput up to 700MB/s
- + 15 Watts
- + PCIe



[>> Online Video](#) [>> Auto Price Quote](#)

Ask 7 smart questions to ensure your next solid state disk is a **flash of inspiration** and not a **regrettable flash in the pan!**

With so many new solid state disks (SSD) on the market, how do you choose one that will be fast and reliable today, and many years from now? You can start by asking your prospective vendors these 7 smart questions:

- Are the performance specifications based on burst or sustained data rates and how much of the capacity is actually available to me?
- Who has independently confirmed the performance claims?
- Is the SSD built with long life enterprise grade SLC Flash or MLC Flash that wears out quickly?
- Does the SSD have onboard controllers or does it need to tap the resources of my server?
- Does the SSD have onboard ultra capacitors for near instant recovery after a power outage or could recovery take hours?
- When did your firm ship its first production SSD and what generation is your current SSD?
- Can your firm provide a RAM-based SSD, PCIe and SAN-based SSDs when our needs change?

At Texas Memory Systems we'd love to answer those questions for you. Because when you compare the details you'll understand why the world's largest SSD installations, including major financial exchanges and 7 of the world's largest telecoms, rely on our award winning RamSan SSDs – the "World's Fastest Storage." For over **30 years** RamSan SSDs are the first choice for speed and reliability. They've always been the *World's Fastest No Regrets SSDs.*

TexasMemorySystems

White Papers and Self-Service Quotes | 1-713-266-3200 | TxMemSys.com/flash

(stasdm.exe) is very similar to the Central Administration option, but since it runs from the command line it can be used in concert with Windows Scheduled Tasks to provide scheduling of backups. It still can't be used to restore the administration or configuration databases. Unlike the Central Administration option, it can be used to back up site collections, but Microsoft warns against performing such backups because they say site collection backups can affect performance and should only be performed when the site collection is locked. Microsoft also notes that these types of backups can be particularly slow when dealing with site collections larger than 15 GB—a very modest size for a site collection. In addition, this utility doesn't seem to like backups that run longer than 17 hours, as it automatically restarts them after 17 hours. Given these issues, Microsoft recommends that to do site collection backups, you should just move that site to its own database and use database backup tools.

SQL Server Backup. Because SharePoint stores most of its information in SQL Server, you can use SQL Server backup tools to back up most of its information, including the configuration and administration databases. You can also use those backups to restore the databases, but it's not supported. Given the synchronization issue, it would seem that as long as you make sure to synchronize what you're restoring then it should work just fine. The key thing is to ensure that no configuration changes are made during your backup window. However, you'll still be missing any customization information stored in the file system if this is the method you choose to back up SharePoint.

Because SQL Server backup tools can be run from the command line, you can schedule this to run at convenient times using Scheduled Tasks. It does require that you manually reattach your databases to the appropriate Web application after a recovery.

What it can't do is back up the search database, and for an odd reason:

SAMPLER: BACKUP PRODUCTS FOR MICROSOFT SHAREPOINT

Check out the companies listed below for the latest information on these products' SharePoint backup capabilities.

Backup products with SharePoint agents

CA • ARCserve

CommVault Systems Inc. • Simpana

EMC Corp. • NetWorker

Hewlett-Packard (HP) Co. • HP Data Protector

IBM • Tivoli Storage Manager (TSM)

Microsoft Corp. • Data Protection Manager (DPM)

Symantec Corp. • Backup Exec

Symantec • NetBackup

Point solutions for SharePoint backup

AvePoint Inc. • DocAve Backup and Recovery

Idera • SharePoint Backup

Quest Software Inc. • Recovery Manager for SharePoint

the search indexes aren't stored in SQL Server. Because you can't synchronize the search database after a database-only backup, this backup approach isn't a viable option for that database.

Windows Server 2008 Backup. The native backup and recovery system for Windows Server 2008 can be used to back up all those things that aren't in the databases (such as the configuration and customization files), but it can't be used to back up the databases themselves.

It seems that the native tools have as many limitations as they have benefits, but it's possible to create a "workable" solution if all you have are the native tools—especially if you can do a regular shutdown of your farm. If you do a shutdown, you could do a SQL Server backup of all of the databases to a file system that's then backed up using the Windows Server 2008 backup system, along with the directories where customization and configuration information is stored.

THIRD-PARTY SHAREPOINT BACKUP OPTIONS

Obviously, to properly back up SharePoint, you need to back up all databases and some files in the file system, and you need to guarantee that these various backups are synchronized. A good recovery system would also allow you to restore the entire system, all configuration and customization data, as well as all content. In addition, it should be able to restore any of the above to various points in time, including the ability to recover individual pieces of content, such as a document.

The only way it seems that you're going to do all of that reliably is to invest in a commercial backup product—and it's likely that the backup application you're using now can handle the chore. Every major backup package has an agent for SharePoint.

The capabilities of each agent vary from one backup application to another, but they all have the same basic functionality. They're add-on agents to your backup software, much like a SQL Server or Exchange agent, that know how to talk to the SharePoint backup API. A well-written agent should only need to be told the name of the main SharePoint server, and it should be able to figure out everything from there. It should figure out the name(s) of any SharePoint farms associated with that server and back them up along with their configuration, administration and content databases, as well as back up any configuration data stored in the file system. All of this data is backed up directly to your backup system's preferred storage, be it disk, virtual tape library (VTL) or tape. Your backup application may actually be doing multiple types of backups under the covers (Microsoft SQL server, file system, etc.), but it should appear as one backup that works (or doesn't work) as a whole.

Every major backup package has an agent for SharePoint.

In addition to backup agents available for your favorite backup software package, there are products, such as AvePoint Inc.'s DocAve Backup and Recovery, Idera's SharePoint Backup and Quest Software Inc.'s Recovery Manager for SharePoint, that are "point solutions" designed just for SharePoint. These products are analogous to backup apps like PHD Virtual Technologies' esXpress, Veeam Software's Backup & Replication and Vizioncore Inc.'s vRanger Pro that are point backup solutions just for VMware. These are designed for firms that have a SharePoint installation, but are using a backup product that doesn't have a SharePoint agent; companies that don't like the capabilities of the agent; or organizations that can't afford the agent. These products tend to do everything you need a SharePoint backup product to do (they may even have more functionality than the agent offered by your backup app due to their specialized nature), but they don't integrate with your backup application. This typically means that their backups will be stored on disk; so if you want those backups to be put on your deduplication system or tape, you'll need to back them up with your other backup product.

A lot of commercial solutions use Microsoft's Volume Shadow Copy Service (VSS) to solve the synchronization problem. That is, they use the SharePoint VSS writer to quiesce SharePoint and the Windows VSS Writer to quiesce the system before backing up everything. That way everything that's backed up is synchronized to the same point. ☺

W. Curtis Preston is an executive editor in TechTarget's Storage Media Group and an independent backup expert. Curtis has worked extensively with data deduplication and other data-reduction systems.

 VMs strain
storage systems

 Users pick
top backup apps

Protect SharePoint data

 Justifying
solid-state storage

 Good match:
NFS and VMware



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

QUALITY AWARDS V:

CommVault continues to dominate the enterprise, while Acronis leads the midrange field



CommVault returned to the winner's circle in our Quality Awards for backup and recovery software for the fourth time as Acronis emerged as a new winner among midrange backup apps. By Rich Castagna



COMMVAULT SYSTEMS INC.'S Simpana suite of data protection apps is beginning to look like the 1927 New York Yankees of enterprise backup and recovery software, with yet another impressive win in our Quality Awards service and reliability survey. Users gave the company's flagship product the highest scores in four of the five rating categories on our survey and put CommVault back into the winner's circle for the fourth time in the five rounds of our survey. "CommVault is hands-down the best software I've ever encountered," noted one very pleased respondent.

Among midrange backup and recovery software products, the competition had been highlighted by a nearly as dominating performance by EMC Corp.'s Retrospect, which snagged top honors in three of our four previous surveys, including the last two. But Retrospect's run was snapped as the product failed to pick up enough responses in the current survey to be included among our finalists. The new king of midrange backup applications is Acronis Inc.'s Acronis Backup & Recovery, which, ironically, received a sufficient number of responses to qualify for the first time and, like CommVault, topped the field in four of five categories.

Four other enterprise products qualified as finalists; ranked in order behind CommVault's overall top score of 6.25 they were Hewlett-Packard (HP) Co.'s Data Protector (6.09), EMC NetWorker (5.63), IBM's Tivoli Storage Manager (5.58) and Symantec Corp.'s NetBackup (5.54).

In the midrange group, Acronis (5.64) was followed by Symantec Backup Exec (5.47) and CA ARCserve Backup (5.25).

USERS TOUGH ON SOFTWARE

This is the only software product group in our Quality Awards program, and over the many surveys we've conducted, software seems to be judged more critically than hardware products, with ratings that are typically (and sometimes significantly) lower. But if there's any solace for the software vendors, it's that the scores have been steadily, if modestly, climbing since our first survey.

In fact, HP's second-place score of 6.09 would have been good enough to win three of the four previous enterprise surveys.

The story's the same for midrange backup programs. Symantec's second-place overall score of 5.47 for Backup Exec would have bettered the winners of the first three Quality Awards surveys, including its own win in the second survey.

Other promising signs for backup software vendors are the improvements in the average scores for product features. Over the first three midrange surveys, those numbers languished in the

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 358 respondents provided 567 backup and recovery software evaluations.

VMs strain storage systems

Users pick top backup apps

Protect SharePoint data

Justifying solid-state storage

Good match: NFS and VMware

4.18 to 4.74 range, which, on our 1.00 to 8.00 rating scale, indicated the functionality of the rated apps was middling at best. The enterprise-class apps fared slightly better. Satisfaction with product feature sets, however, seems to be on the rise; for enterprise apps, the average feature scores on the last two surveys have been impressive, just barely exceeding 6.00 both times. Midrange products have also shown improvement, although less impressively so, with feature ratings finally climbing over the 5.00 mark.

BUYING BACKUP

Any IT purchase can be a gut-wrenching experience, not just because the buy might involve some big bucks, but because each purchase is a commitment. For backup software, the commitment goes beyond time, effort and other resources as you're also committing your company's data to the product. So the sales process is particularly critical for a data storage manager who's about to lock into a particular backup technology.

CommVault scored highest among enterprise backup applications in the sales-force competence judging category with a 6.20; HP, which proved to be a game competitor for CommVault, was second with a solid 6.10, well ahead of third-place finisher EMC NetWorker (5.85).

CommVault scored highest on six of the seven statements in this category, with its best ratings for "The vendor's sales support team is knowledgeable" (6.64) and "My sales rep is knowledgeable about my industry" (6.33). HP also did well on those statements (6.06 and 6.13, respectively), but topped CommVault—6.63 to 5.91—for when the sales process gets closest to signing on the dotted line ("The vendor's licensing formula offers good value"). All of the vendors fared well for having a knowledgeable sales support team, with EMC (6.37), IBM (5.77) and Symantec (5.65) getting their highest scores in this category for that statement.

The midrange products didn't fare quite as well in the sales-force competence category. Acronis picked up a couple of 6.00-plus ratings on its way to an category winning score of 5.65, but neither Symantec (5.36) nor CA (4.77) could muster one. Smaller businesses may be tougher customers for backup software vendors, but Acronis seems to be following the lead of its enterprise-class brethren, with its highest scores coming for those statements related to having a knowledgeable sales support team (6.11) and offering favorable licensing terms (6.11).

Any IT purchase can be a gut-wrenching experience, not just because the buy might involve some big bucks, but because each purchase is a commitment.

CLICK ON "START"

In four of the five evaluation categories for enterprise applications, HP nipped at CommVault's heels but ended up in second place. But in the initial product quality category, HP turned the tables to outdistance CommVault, 6.13 to 5.93, coming out on top for four of the seven statements in the process.

With software products, the key to long-term satisfaction is getting out of the gate fast with a quick and relatively pain-free installation. HP is apparently doing a good job of setting the tone early with its users by having a product that installs easily without customers requiring a lot of help from HP's professional services team. HP received its highest score in this category for the statement "This product was easy to install," and was the only vendor in the group to rise above a score of 5.00 for "This product did not require professional services."

Although it finished second to HP, CommVault was no slouch in the initial product quality category, with especially high ratings for "This product uses tape efficiently" (6.55) and "This product delivers good value for the money" (6.48). Its overall score was dragged down by a 4.52 for the statement about requiring professional services.

Arguably, midrange backup products should be easy to install and get up and running simply because the users of these products are likely to have fewer resources available to get through the initial stages of implementation. All three products showed their mettle in the initial product quality category, with Acronis (5.99) coming out on top, followed by Symantec (5.82) and CA (5.76); in addition, all three had 6.00-plus scores for the critical "easy to install" statement. Acronis scored very well for all seven statements in the category with one glaring exception—a 4.00 for "This product uses tape efficiently," which, ironically, was where Symantec racked up one of its highest marks on the way to a very respectable performance.

With software products, the key to long-term satisfaction is getting out of the gate fast with a quick and relatively pain-free installation.

FUNCTION FOLLOWS FORM

A no-sweat implementation may indeed set the tone for a user's overall experience, but a backup application still has to deliver the goods with the required features and functionality. With the highest score recorded for the product features category (6.70), CommVault proves that there's plenty of substance behind the style of its backup suite. Its victory in this category is a tour de force, as it scored the highest for every statement, highlighted by the only over-7.00 score in the survey, a 7.02 for "This product's file system backup features meet my needs." CommVault also

VMs strain storage systems

Users pick top backup apps

Protect SharePoint data

Justifying solid-state storage

Good match: NFS and VMware

scored impressively for database backups (6.88), for being a “complete solution” (6.81) and for its disk backup functionality (6.80).

Second place HP also did well for features, with its overall 6.14 score bolstered by strong showings for database backup (6.45) and file system backup (6.38). IBM’s Tivoli Storage Manager was hard on HP’s heels with a category average rating of 6.12, highlighted by a 6.52 for its disk-based backup features. Symantec’s NetBackup made a strong showing with an overall 5.89 built on three 6.00-plus statement scores.

Once again, survey respondents were a little tougher with their ratings for midrange backup apps, with the products earning scores lower than their enterprise counterparts. Symantec’s Backup Exec showed why it’s one of the most widely installed midtier backup applications by topping its competition on seven of the eight statements in the product features category for an average winning score of 5.68. Backup Exec did particularly well for file system backup features (6.05), backup to disk (5.99) and database backup (5.85); its lowest score—a 5.27 for archiving features—was just a hair shy of Acronis’ score of 5.31. Acronis’ Achilles’ heels in this category were database backup (5.17) and scalability (5.17).

IN FOR THE LONG RUN

The true test of a backup application comes with the daily grind of protecting a company’s data assets; first impressions and feature sets are all well and good, but long-term reliability and stability are paramount. CommVault was the only enterprise-class backup app with a 6.00-plus average for the product reliability rating category, picking up a 6.24 while leading on six of eight statements. Its strongest showings were for meeting service-level requirements (6.86) and operating system/platform support (6.84).

Second place HP (5.94 overall) also showed strength for those two reliability criteria (a 6.18 and 6.59, respectively) among its consistent marks in this category. For the two statements that CommVault didn’t have the leading score, EMC’s NetWorker led the pack with a 6.07 for “Requires very few unplanned patches/updates,” while IBM Tivoli Storage Manager garnered a top mark of 5.87 for “Very few bugs.”

Within the midrange product set, Acronis bested CA (5.38) and Symantec (5.34) for five of the eight reliability statements on its way to winning the category with a 5.68 average. This was CA’s sole second-place finish; its ARCserve Backup product came out on top for meeting service-level

The true test of a backup application comes with the daily grind of protecting a company's data assets; first impressions and feature sets are all well and good, but long-term reliability and stability are paramount.

VMs strain storage systems

Users pick top backup apps

Protect SharePoint data

Justifying solid-state storage

Good match: NFS and VMware

requirements (5.94) and for requiring few unplanned patches or updates (5.83). Backup Exec finished a close third with a winning score of 5.17 for “Vendor provides comprehensive upgrade guidance”—an area where all three vendors’ ratings were relatively low, indicating that some improvement in the upgrade process would be welcomed by users.

SOFTWARE SUPPORT

Backup software vendors have labored diligently to make their products easier to use with improved GUIs, wizards and other interface aids. But with an application as complex as backup, users will invariably need some help from time to time, so how well a vendor supports its product will figure significantly in a user’s perception of ease of use. After backup vendors as a group registered a disappointing overall average of 4.88 in the first Quality Awards survey for backup software, they have shown steady progress in the quality of support they provide, with the best numbers yet recorded this time around (a very healthy 5.89).

PRODUCTS IN THE SURVEY

The following backup and recovery software products were included in the survey. (The number of responses for finalist products is shown in parentheses.)

ENTERPRISE

- Asigra Inc. • Televaulting*
- BakBone Software Inc. • NetVault:Backup*
- CommVault Systems Inc. • Galaxy (or Simpana) (48)
- EMC Corp. • NetWorker (36)
- Hewlett-Packard (HP) Co. • Data Protector (43)
- IBM Corp. • Tivoli Storage Manager (81)
- Symantec Corp. • NetBackup (81)
- Symantec • NetBackup PureDisk*
- Syncsort Inc. • Backup Express*
- Zmanda Inc. • Amanda Enterprise*


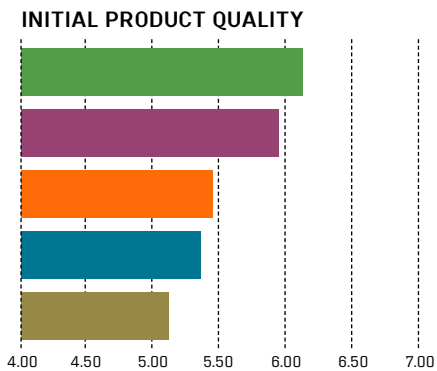


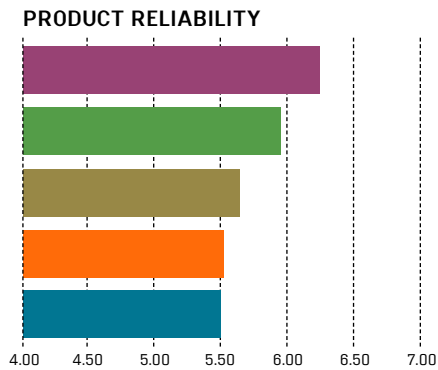

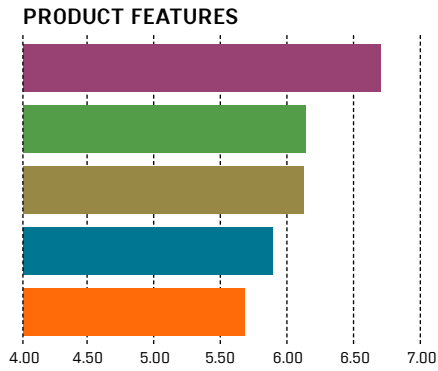
MIDRANGE

- Acronis Inc. • Backup & Recovery (33)
- Arkeia Software • Arkeia Network Backup*
- Atempo Inc. • Time Navigator*
- BarracudaWare • Yosemite Server Backup*
- BridgeHead Software Ltd. • HT Backup*
- CA • ARCserve Backup (43)
- Double-Take Software Inc. • Double-Take Backup*
- EMC • Avamar*
- EMC • Retrospect*
- FalconStor Software Inc. • Continuous Data Protector*
- Microsoft Corp. • Data Protection Manager (DPM)*
- PHD Virtual Technologies • esXpress*
- Symantec • Backup Exec (108)
- Veeam Software • Veeam Backup & Replication*
- Vizioncore Inc. • vRanger Pro*

**Did not receive enough survey responses to be included as a finalist.*



- CommVault Systems Inc. Galaxy (or Simpana)
- EMC Corp. NetWorker
- Hewlett-Packard (HP) Co. Data Protector
- IBM Corp. Tivoli Storage Manager
- Symantec Corp. NetBackup



Based on a 1.00-8.00 scoring scale

*% Yes

VMs strain storage systems

Users pick top backup apps

Protect SharePoint data


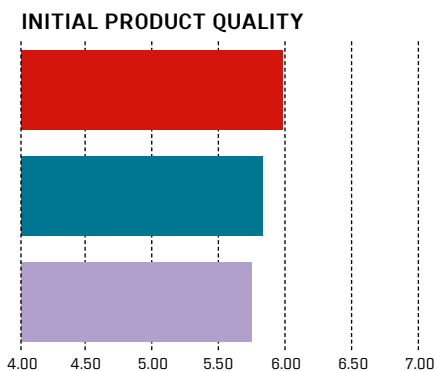





Justifying solid-state storage

Good match: NFS and VMware

ENTERPRISE BACKUP AND RECOVERY SOFTWARE



- Acronis Inc. Backup & Recovery
- CA ARCserve Backup
- Symantec Corp. Backup Exec



Based on a 1.00-8.00 scoring scale

*% Yes

VMs strain storage systems

Users pick top backup apps

Protect SharePoint data

Justifying solid-state storage

Good match: NFS and VMware

MIDRANGE BACKUP AND RECOVERY SOFTWARE

CommVault's 6.20 score for the technical support category was enough to nudge out HP's 6.12. IBM (5.82), EMC (5.73) and Symantec (5.58) turned up pretty good ratings as well. CommVault is apparently meeting users' support expectations, as its highest score in this category was a 6.86 for "Vendor supplies support as contractually specified." HP topped CommVault on the statements dealing with documentation (a 6.15 vs. CommVault's 6.00) and "Support issues rarely require escalation" (6.00 and 5.69, respectively). EMC was rewarded with a statement-high score of 6.27 for training its third-party partners well.

Again, we detect some grumbling among midrange users as they aren't nearly as satisfied with their vendors' support efforts. Although the overall support score average is higher than it's ever been, it's still a rather unimpressive 5.23.

Acronis topped the category with a modest 5.39 overall support score, but did manage to get the only 6.00-point rating in the category with an even 6.00 for providing support as contracted. All three vendors failed to reach the 5.00 plateau for providing adequate training, with Acronis earning a 4.80, followed by Symantec (4.74) and CA (4.59). And the three didn't fare much better for providing adequate documentation, with scores ranging from Acronis' 5.05 to Symantec's 5.08, with CA sandwiched in between at 5.07. The message to midrange backup software vendors seems clear: Provide more support and you're likely to see better all-around scores on future surveys.

BRAND LOYALTY

As on all Quality Awards surveys, we ask respondents if, given what they now know, they would make the same purchase again. Given its overall performance on this survey, it's not too surprising that CommVault had the highest percentage of users who said they'd do it all over again (83.7%). Only 62.9% of HP's users said they would buy Data Protector again, a surprisingly low number considering how well the product was rated overall.

For the midrange applications, 69.2% of Acronis' users said they're pleased enough to buy the product again, followed by Symantec Backup Exec (66.7%) and CA ARCserve Backup (57.1%).

Although not necessarily related to user satisfaction, it's interesting to see how heavily users are taxing these backup products. Among enterprise programs, IBM's Tivoli Storage Manager is the workhorse, backing up an average of 67.1 TB weekly. For the midrange apps, Symantec's venerable Backup Exec is the heaviest lifter with an average of 13.3 TB of data backed up each week. ☉

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



Are more Symantec OST-like APIs needed?

One of the by-products of the OST interface is better backup and recovery performance. Maybe that will get more backup vendors to develop their own APIs.

EMC CORP. RECENTLY announced its Data Domain Global Deduplication Array (GDA) that optimizes data deduplication in large-scale environments by aggregating the data storage capacity of two of its deduplication appliances to improve throughput performance and scale. In terms of delivery, one of the key enablers of GDA's ability to distribute deduplicate processing is Symantec Corp.'s OpenStorage (OST) API technology.

Symantec OST is an API for NetBackup (Versions 6.5 and higher) and Backup Exec 2010. Partners leverage the API to write a software plug-in module that's installed on the backup media server to communicate with the storage device, creating tighter integration between the backup software and target storage. In short, it's an interface that speeds up backup for NetBackup customers. The only problem with OST is that it highlights the fact that other backup vendors don't offer a similar capability.

SYMANTEC OST

Originally, the Symantec OST API was published to provide Symantec customers with a common interface to third-party disk targets. It allows backup data to be stored on disk with whatever protocol the target device uses, such as Fibre Channel (FC) or TCP/IP. Symantec backup software sees OST-enabled appliances as disk and enables features such as intelligent capacity management, media server load balancing, reporting and lifecycle policies.

It also delivers optimized data duplication—network-efficient replication and direct disk-to-tape (D2T) duplication that's monitored and cataloged by the backup software. Without Symantec OST, there are two scenarios: allow the storage device to transfer data without the backup catalog being aware of the copies, or transfer data from device to media server to device to keep the backup catalog aware of the copy. In the first scenario, the backup catalog is left out of the loop on the location of backup copies. This can create complexity and impede disaster recovery (DR) processes. The latter scenario increases LAN, WAN and SAN network traffic, and removes the benefits of deduplication in network transfer. Clearly, deduplication controlled by OST-enabled devices creates savings in both time and bandwidth requirements.

Because the catalog is aware of all copies, recovery of data from an OST-optimized duplicate copy is the same as recovery from another duplicate.

Through the backup application, the OST-optimized duplicate copy can be designated as the primary copy, and then a full or granular recovery can be initiated. The potential time savings when compared to recovery from a non-OST-optimized duplicate could be significant.

VENDOR OST ADOPTION

Many backup target device vendors have subscribed to the Symantec OST API, which isn't surprising given its benefits and Symantec's market share. Vendors with support for OST in conjunction with NetBackup and/or Backup Exec include EMC, ExaGrid, FalconStor Software, GreenBytes, IBM, NEC, Quantum (the only vendor so far to support OST direct-to-tape support with NetBackup) and Sepaton. It's also worthwhile to note that Symantec supports its own deduplication implementation in NetBackup and Backup Exec with OST.

One of the by-products of the OST interface is a performance improvement in backup and recovery operations, with some vendors claiming upwards of a 100% increase in performance. EMC's OST option for its Data Domain appliances was aptly renamed "Boost," a testament to its performance advantage. In creating its OST plug-in, EMC enhanced communications and optimized the packaging and transfer of data between backup media server and storage device, thereby improving performance.

EMC became more innovative with Data Domain GDA, taking advantage of the OST API to distribute a portion of the deduplication processing to the backup media server, which EMC claims lowers media server CPU utilization. And because deduplication occurs earlier in the backup data path, the implementation eliminates some redundant data at the media server, and reduces the network load between media server and storage.

In a similar move, NEC leveraged the OST API to optimize load balancing. While one of the inherent benefits of OST is to enable disk pooling for better overall backup system load balancing, NEC took things a step further. Hydrastor, NEC's storage platform offering data deduplication, has a scale-out grid architecture employing one or more logical storage units (data movers) and storage. Through OST integration, the backup application can now automatically distribute jobs to the logical storage units of the Hydrastor grid.

Many backup target device vendors have subscribed to the Symantec OST API, which isn't surprising given its benefits and Symantec's market share.

DISK AND DATA DEDUPLICATION IN DEMAND

Disk-based backup is becoming more pervasive in data protection strategies; ESG research finds that the number of organizations using only tape in backup

VMs strain storage systems

Users pick top backup apps

Protect SharePoint data

Justifying solid-state storage

Good match: NFS and VMware

operations dropped 27% between 2008 and 2010, with more organizations favoring a disk-to-disk (56% increase) or disk-to-disk-to-tape (42% increase) strategy. Data deduplication use grew more than 200% between 2008 and 2010.

When EMC launched Data Domain GDA, the company planned to move deduplication “upstream” via integration with EMC NetWorker. Unfortunately, the integration is likely to be hard-coded into NetWorker since NetWorker currently doesn’t have an OST-equivalent API—nor does any other backup vendor product.

It’s also unlikely that Symantec will make OST an open standard that other backup vendors could use (similar to how NDMP is utilized by backup vendors to back up filers). Looking ahead, it’s more likely we’ll see other backup vendors attempt OST-like APIs.

Of course, Symantec charges a fee to test and certify its OST partners’ solutions. So it could get expensive for a company like Quantum or Data Domain, for example, to certify its solutions with multiple backup products’ APIs. In turn, end users are charged a premium fee for OST enablement—from both the backup vendor and the target system vendor. In addition, an end user with multiple backup solutions (in this case Backup Exec and NetBackup) is likely to be required to pay license fees to Symantec for OST-enablement with each backup product. ☹

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



Virtual servers and their impact on data storage

Almost every IT shop has virtualized at least some of its servers, putting much of the burden of server consolidation on data storage systems.

MANAGING DATA STORAGE resources has proven to be a challenge in virtual server environments since server virtualization first took root in data centers approximately seven years ago. By consolidating multiple virtual machines (VMs) on a single physical system, server virtualization imposes additional pressure on storage and I/O resources, and complicates functions such as data protection and storage management. In a recent Taneja Group end-user research study, we learned how your storage management peers are coping with the impact of server virtualization, and how they're attempting to employ the technology to their advantage. A total of 360 IT managers from a mix of North American enterprise and midsized organizations gave us their perspectives via an online survey, supplemented by follow-up phone discussions with selected respondents.

MOVING BEYOND FIBRE CHANNEL

A majority of respondents use Fibre Channel SANs (FC SANs) as their primary networked storage platform, and it's still the leading virtual server storage protocol, used by 75% of respondents. But iSCSI (used by 36% of those surveyed) is the fastest growing storage choice for virtual servers. While most initially considered iSCSI because of its lower cost, they tended to be pleasantly surprised by the functional strengths and performance of current iSCSI offerings from suppliers such as Dell EqualLogic and Hewlett-Packard (HP) Co. (with its LeftHand Networks acquisition).

NFS is becoming more of a factor on the NAS side: While just 19% deploy it overall, NFS is used by three out of 10 enterprises (i.e., organizations with 10,000 or more employees) and four out of 10 companies with at least 1 PB of total storage capacity. Large firms are adopting NFS for two primary reasons: ease of management and its virtual server data protection capabilities. We believe enhanced functional support by vendors such as VMware and NetApp is also playing a role. In addition, the growing interest in NAS is driving significant innovation from scale-out NAS vendors such as Isilon Systems and HP with its StorageWorks X9000 Network Storage Systems.

TRADITIONAL BACKUP FOR VMs STILL PREFERRED

If you're using a physical system backup approach to protect your virtual server data, you're not alone. More than 75% of our survey respondents are

using one or more physical server data protection processes to back up their virtual servers. Most users said they opted for physical system backup because they wanted a proven and familiar backup approach; many also cited the cost and convenience of sticking with backup software from vendors they already work with. The most popular method is agent-based backup inside a guest OS (53%). File-level backup is the leading approach, which is consistent with the requirement (stated by 57%) that virtual server files be directly recoverable in a single step, rather than first having to restore a virtual machine image.

While current virtual server data protection practices suggest that many storage managers are averse to change, we found 30% of respondents are planning or strongly considering new backup approaches tailored to virtual servers, such as the use of hypervisor-based VM snapshots (currently used by 25%). Array-based snapshots and VCB-style proxy server backup (with nearly 30% and 27%, respectively) were also mentioned as data protection practices likely to be deployed in the next 12 months for at least some virtual server workloads.

Most users said they opted for physical system backup because they wanted a proven and familiar backup approach; many also cited the cost and convenience of sticking with backup software from vendors they already work with.

VIRTUAL SERVERS NOT YET KEY FOR DISASTER RECOVERY

Improving disaster recovery (DR) capabilities is often cited as one of the top reasons for virtualizing servers, but our survey revealed that virtual server-enabled DR is still in the early stages of adoption. Three out of four users have overall DR or business continuity (BC) plans in place; but despite virtualization vendors' best marketing efforts, only a minority of users have committed to DR strategies built around virtual server technology. For example, 40% of those surveyed are using capabilities like VMware VMotion to automate the movement of encapsulated workloads, and only a fraction of those are employing the technology to migrate workloads from a data center to DR site. Of course, technology constraints currently limit the effective range of such tools, and we expect their use for disaster recovery to increase significantly once those limitations are removed.

A large majority of end users also indicated that they're not yet able to rely on their virtual infrastructure to help them rapidly recover from a disaster, even though the frequency of outages in their virtual server environments is relatively high. Case in point: 33% have experienced at least one disk-related outage or failure in the past three years. Interestingly, most enterprise IT managers we interviewed still tend to think about disaster recovery from the

VMs strain storage systems

Users pick top backup apps

Protect SharePoint data

Justifying solid-state storage

Good match: NFS and VMware

storage side, focusing, for example, on remote replication capabilities in the array. It appears that server virtualization providers still have a lot of work to do to persuade and fully equip their customers to invest in virtual server-driven DR practices.

STORAGE MANAGEMENT REMAINS A CHALLENGE

Nearly 80% of respondents point to specific storage barriers that have slowed the progress of their virtual server deployments. The leading impediment is the increased complexity of data storage capacity and performance management (cited by 37%), followed by challenges in application-to-storage mapping and data protection (each noted by 28%).

To address the first issue, storage managers are increasingly using data reduction and related technologies to improve virtual server storage utilization and performance. For instance, in the next 12 months, respondents plan to deploy data deduplication (44% will deploy vs. 27% using it today); virtual server or array-based thin provisioning (46% will deploy vs. 33% today); and data compression for primary storage (9% will deploy vs. a little more than 5% using now).

VMs AND STORAGE: A WORK IN PROGRESS


Storage managers have made great strides in defining and enhancing their virtual server storage practices during the past few years, but a lot of work remains to be done. To gain the upper hand, we strongly encourage IT practitioners to view storage as a strategic and integral component of their virtual infrastructures, and to give storage equal weight with servers in planning and deployment. We'll continue to keep tabs on how users address these issues as the book on virtual server storage best practices continues to be written. ☺

Jeff Byrne is a senior analyst and consultant at Taneja Group. He can be reached at jeff.byrne@tanejagroup.com.

Companies are better prepared for disasters

WHATEVER YOU call it—an expensive insurance policy that you’ll probably never use or just a pain in the neck—disaster recovery (DR) planning and testing is a necessary evil for data storage managers. Two years ago when we first ran a DR survey, just slightly more than half of the respondents said their plans were tested regularly, but this time around that number jumped to 66%. Of course, “regularly” is relative; 59% on our current survey said they test at least twice year, but that’s down considerably from the 70% reported in 2008. So while more companies are actually testing their DR plans, they aren’t doing it as often. But perhaps testing less frequently has its own rewards: 47% of respondents said they met all of their recovery point objectives (RPOs) and recovery time objectives (RTOs) when they tested vs. only 31% a couple of years ago. That suggests some pretty good planning and testing, which is probably why 51% of those surveyed said they were very confident their plans would allow their firm to weather a disaster without significant business impact. So what about the 33% of those surveyed who don’t test their DR plans? People and money are the problem, with 27% claiming they lack the staff to do DR and another 15% citing lack of funds. *—Rich Castagna*


Do you perform regular tests of your company's disaster recovery (DR) plan?




52%

Have disaster recovery sites less than 100 miles from their main data center

Do you do full or partial DR testing?



Did you meet your recovery point objectives (RPOs) and recovery time objectives (RTOs) during your last test?



“DR site set up two years ago and apparently never tested . . . doh!”

—Survey respondent

VMs strain storage systems

Users pick top backup apps

Protect SharePoint data

Justifying solid-state storage

Good match: NFS and VMware

Check out the following resources from our sponsors:

3PAR, page 13

[3PAR -- Autonomic Storage Tiering](#)

[3PAR: Optimizing IO Service Levels](#)

Corporate Technologies Inc., page 17

[Storage Efficiency Strategies for Microsoft Application Environments: An ESG White Paper](#)

[Veritas NetBackup for Microsoft Exchange Server Solutions Guide: A Symantec White Paper](#)

Dell Inc., page 4

[Protecting a Paperless Office](#)

[Data Protection Examples Using PS Series Snapshots in a VMware Virtual Infrastructure](#)

EMC, page 21

[Optimizing Data Protection for Virtual Environments](#)

[Data Storage Priorities 2010: Data Backup, Capacity Growth and Disaster Recovery](#)

Iron Mountain Digital, page 11

[Seven Steps to a Successful Cloud Storage Strategy](#)

[Iron Mountain's Cloud-Based Server Data Protection for Servers](#)

Pillar Data Systems, page 7

[Accelerating and Managing Data Using Solid State Drives](#)

[A New Approach to Designing a Modern Storage Infrastructure](#)

Quantum Corp., page 29

[Quantum Goes Beyond Deduplication](#)

[Quantum DXi-Series and Symantec NetBackup OpenStorage](#)

Texas Memory Systems Inc., page 25

[The World's Fastest No Regrets SSDs -- RamSan by Texas Memory Systems](#)

[A Flash of Inspiration -- RamSan SSDs the "World's Fastest Storage" by Texas Memory Systems](#)