

Managing the information that drives the enterprise

STORAGE

Vol. 10 No. 3 May 2011

Automated Storage Tiering

Smart storage systems put data in its place to make better use of high-performance or high-capacity drives.

ALSO INSIDE

Every cloud has a green virtual lining

Needed: A new way to protect files

Shining a spotlight on unified storage

Use of cloud is evolving

Capacity control and keeping legal top email archiving list



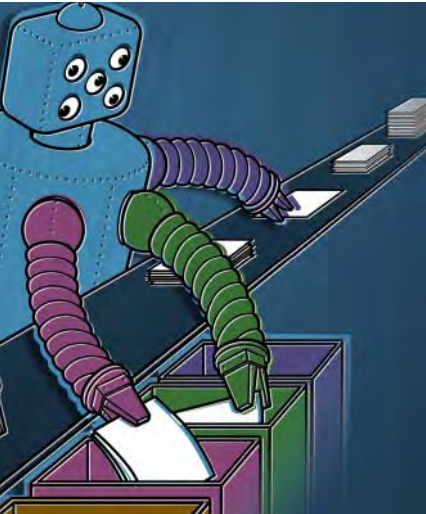
STORAGE

inside | May 2011



STORAGE

inside | may 2011



Every cloud has a green virtual lining

- 5 **EDITORIAL** Buzzwords are taking over the data storage industry, so it's probably asking way too much of storage vendors to just tell us what their products can—and can't—do. by RICH CASTAGNA

Reinventing file storage protection and recovery

- 9 **STORWARS** Traditional backup just won't cut it anymore as file data grows and grows; we need some new thinking and an updated approach to replication. by TONY ASARO

Automated storage tiering: Higher performance AND lower cost?

- 12 Automated storage tiering is an effective way to make efficient use of installed data storage resources, and to take advantage of the high performance of solid-state storage. by PHIL GOODWIN



Blueprint for cloud-based disaster recovery

- 21 Cloud storage and computing services offer a number of alternatives for cloud-based DR, but your options will depend on the recovery time and recovery point objectives you need. by JACOB GSOEDL

Storage budget recovery on a roll

- 29 Storage budgets continue to recover from their recessionary pounding. But while storage managers might have more money to spend, they'll need yet more capacity to meet new demands. by RICH CASTAGNA

Virtualization, cloud shine spotlight on unified storage

- 38 **HOT SPOTS** Unified storage adoption is starting to ramp up as data storage pros see the need for simplifying storage so it can be powered, cooled and managed in one pool. by TERRI MCCLURE



Where is the cloud storage market headed?

- 42 **READ/WRITE** Break down the cloud storage services market and you'll find players both big and small jockeying for position in key segments. by JEFF BYRNE

Capacity and the law drive email archiving

- 45 **SNAPSHOT** In our most recent Snapshot survey, we asked respondents why they were archiving: 28% say they're doing it for legal issues, while 26% use it for capacity management. by RICH CASTAGNA

From our sponsors

- 47 Useful links from our sponsors.



**SERVER ROOMS
THAT REQUIRE
GPS NAVIGATION.**

SOLVED.

We get that virtualization can drive a better ROI. Highly certified by Microsoft, VMware, HP and others, we can evaluate, design and implement the right solution for you.

We'll get you out of this mess at CDW.com/virtualization





Every cloud has a green virtual lining

Buzzwords are taking over the data storage industry, so it's probably asking way too much of storage vendors to just tell us what their products can—and can't—do.

I'M ABOUT TO have another Peter Finch moment—specifically, when he played the slightly demented newscaster Howard Beale in the movie *Network* and exhorted the masses to proclaim, “I’m as mad as hell, and I’m not going to take this anymore!”

OK, maybe I’m not quite *that* ticked off and, yes, I’ve used this Peter Finch reference once before in a column. I think I may have regressed to the “mad as hell” theme because the thing that set me off the first time—vendors carelessly tossing around marketing mumbo jumbo—hasn’t gone away. In fact, it somehow managed to rev itself up into an even higher gear, achieving new heights of nonsense and non sequitur. But . . .

“Hey, vendors, we’re not dummies!”

Now that I’ve gotten my “mad as hell” rant out of the way, let’s get down to specifics. I don’t know if storage vendors actually think we’re dummies or if they just kind of treat us that way because they don’t know any better. And it’s not that they’re cheating people or selling bad stuff or anything like that—what they’re doing, they’re doing with words. And some of those vendors seem to live in an alternate universe where things become true just by saying them. Mostly, though, they just delude themselves and don’t fool the ones they’d really like to convince, like storage managers.

Let’s face it, storage is a tough business. Making the things on which you just store stuff sound exciting or novel isn’t easy, especially with the competition always nipping at your heels with catchier catchphrases.

But it does take a certain amount of talent to take a word or phrase and by sheer misuse (and repeated use) render it meaningless. Most data storage vendors are doing a great job with “cloud”; they’ve managed to simultaneously

It takes a certain amount of talent to take a word or phrase and by sheer misuse (and repeated use) render it meaningless.

A barrage of buzzwords

New way to protect files

Automated storage tiering

Cloud-based DR

Storage budgets on the rise

Spotlight on unified storage

Use of cloud is evolving

Email archiving

Sponsor resources

render it meaningless while making it impossible to define. Cloud is headed straight to the Buzzword Buzzkill Hall of Fame to take its place next to “compliance” and “green.”

“Virtual,” the poster child for storagespeak in 2010, has been nudged out of the spotlight by the first significant catchphrase of 2011 that appears to have some legs: “big data.” This is an interesting one because in the short time it’s been bandied about, its meaning has already morphed into something that’s essentially the complete opposite of what the term originally referred to. That’s record-breaking obfuscation in my book, and it should probably earn members of the tech marketing intelligentsia nominations to the Cunning Marketers Hall of Fame.

I don’t know where “big data” came from, but at first it was used fairly innocuously (and accurately) to describe really huge files like video or research data that put a strain on storage gear. But then EMC went and bought Isilon, and Joe Tucci, EMC’s top guy, said “big data” was the key to the acquisition, and that was enough to trigger an avalanche of “big data” me-too-ism. When it comes to buying stuff, EMC doesn’t screw up very often (ever?), and if Joe says “big data,” everyone listens.

EMC's Joe Tucci said "big data" was the key to the Isilon acquisition, and that was enough to trigger an avalanche of "big data" me-too-ism.

And soon—what a surprise!—*everyone* had storage systems that were perfect for big data. Except now “big data” also means lots and lots of not-necessarily-really-huge files. So everyone does big data even if they don’t really do big data, and big data includes both large and small files. Got that?

We’re not out of the “big data” woods by a long shot. For some reason, a lot of analysts and consultants are on the big-data bandwagon now, and everyone’s talking about it *as if it actually meant something*.

But don’t worry; as soon as “big data” wears out, the buzz will shift again. I’m betting on “_aaS,” as in, something-or-other as a service. We already have SaaS, which can mean storage as a service or software as a service; IaaS, which is infrastructure as a service; and PaaS, or platform as a service. And just the other day I saw DPaaS: data protection as a service. The _aaS engine is just heating up so you can expect more and more of this stuff, especially as people start to get tired of hearing the word “cloud.”

I guess I’m just old-fashioned and expect storage vendors to push the truly unique aspects of their products, like being faster or bigger than the others, or doing as much for less money. Instead, storage marketers obscure any real

A barrage of buzzwords

New way to protect files

Automated storage tiering

Cloud-based DR

Storage budgets on the rise

Spotlight on unified storage

Use of cloud is evolving

Email archiving

Sponsor resources

outstanding qualities of their products by paying lip service to the same buzzwords that everyone else is working over. These vendors seem to think they can be all things to all people: “Sure, we do big data!” I wonder if they’ve even asked storage managers what “big data” means to them. And I wonder how storage managers would define “big data.”

Can we start talking about real features and capabilities again? When everyone uses the same term to describe things that are basically (and very obviously) unlike, it can only be confusing. Let’s try to get this stuff out of our systems once and for all and get back to reality.

So tell me again, what exactly does your Green Virtual Cloud Big Data as a Service Compliance Edition system do? ☺

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 June 2011 issue.](#)

[A barrage of buzzwords](#)

[New way to protect files](#)

[Automated storage tiering](#)

[Cloud-based DR](#)

[Storage budgets on the rise](#)

[Spotlight on unified storage](#)

[Use of cloud is evolving](#)

[Email archiving](#)

[Sponsor resources](#)

SHAREPOINT
GOVERNANCE

QUARTERLY
FISCAL REPORTS

1998
CUSTOMER
CONTRACTS

VITAL RECORD
SECURITY

EMAIL
TO BE ARCHIVED

SCANNED
INVOICES

Your Information is at Risk. Protect What Matters Most.

As the amount of information your organization has to manage and protect continues to grow, the challenge of managing the potential risk increases exponentially. How can you ensure your organization's information is not at risk? Partner with the company thousands have trusted to store, protect and manage their information regardless of format – Iron Mountain. With unmatched experience, putting us at your side makes information easier to manage. We can do more, together.

Safeguard your Information. Visit us at ironmountain.com.



CATEGORIZE ▶ ARCHIVE ▶ IMAGE ▶ DISCOVER ▶ DESTROY



Reinventing file storage protection and recovery

As file data growth surges, traditional backup just won't cut it anymore; we need some new thinking and an updated approach to replication.

ACCORDING TO IDC, the amount of new file storage growth between 2009 and 2014 is expected to be about 160.35 exabytes. That's approximately 300% more than the growth of every other data type combined, including database and email, over the same period of time. That kind of file growth has a number of negative ramifications, not the least of which is protecting it all. Traditional data backup approaches are no longer practical because of the sheer mass of file storage.

In many cases, IT professionals don't create file systems larger than 2 TB because they don't want backup data sets to be too big. This means that if you have one petabyte of NAS storage, you'll have at least 500 file systems you have to back up. There are companies with thousands of file systems out there; over time, that kind of situation will become more and more commonplace.

Although the market tends to hype and value large file systems, they're difficult to protect. If you have a file system that's 100 TB, then backing up the entire file system becomes extremely impractical. This is also true of object-based storage systems that have a flat name space. Vendors that provide these products often recommend you replicate to disk vs. backing up. However, that doesn't provide an easy way to recover data. The challenge is that most of the storage-based replication solutions are block based, so you really don't have any efficient methods to recover data at the file level. And even if some of those systems provided file-level replication, they have no recovery app for users to find and retrieve the files they're looking for. As file storage increases, it makes the needles in our ever-growing haystack harder and harder to find.

Block-based replication has never been, and never will be, an adequate replacement for backup for a number of reasons. Storage-based solutions are vendor specific and therefore don't provide a universal method for data protection. Additionally, these solutions are typically confined to single storage systems; they're stovepiped. If you have 100 NAS systems, it will be a nightmare to manage remote mirroring for all of them. This approach is also costly because it's usually a paid-for option, it increases maintenance charges and replicates data onto the same vendor's storage, which isn't necessarily a low-cost solution. Perhaps most importantly, recovering specific files is a difficult if not impossible

A barrage of buzzwords

New way to protect files

Automated storage tiering

Cloud-based DR

Storage budgets on the rise

Spotlight on unified storage

Use of cloud is evolving

Email archiving

Sponsor resources

task. Remote mirroring is not well suited for granular recoveries; it's better suited for recovering entire systems.

A better and smarter approach is an intelligent file-level replication solution with the following capabilities:

- The ability to replicate data to and from any file system
- The ability to replicate entire systems, individual file systems, directories and sub-directories, and at the file level
- It's essential it provide search and recovery capabilities so users can find what they're looking for efficiently
- The ability to scan the file systems for any changed or new files, and to replicate only those to the system
- Must be able to scale to petabyte environments, including discovery, replication and search with high performance

A software company called Digital Reef Inc. is doing all of the above. However, it's also important to find a lower cost and easy to manage storage tier to replicate this data. There are a number of scale-out file storage systems that fit this requirement, including HP Ibrix and IBM SONAS. EMC Isilon isn't really a lower cost solution, but there are configurations where it would certainly be more attractive price-wise than tier 1 NAS. Dell Exanet should be available as an option for this tier as well. There also seems to be an uptick in interest in the Symantec file system, which sounds good on paper. There are a number of open-source file systems, including Gluster and Hadoop, and we can't forget ZFS—it's not scale-out, but you can throw Gluster in front of it to provide that capability. However, whenever you're using an open-source file system there's typically some handholding that's required by the user.

The return on investment would be significant. In some cases, you could even stop backing up your file systems altogether. Consider the impact on your infrastructure and resources of eliminating file backups. You can also reduce your reliance on storage-based mirroring and minimize the cost and management of these solutions. Reserve remote mirroring technology for mission-critical files and leverage file replication to a lower cost, extensible storage tier for everything else.

The world has changed and yet we're still using the same tools to manage our file data. That's neither practical nor sustainable. Not unless you have an unlimited budget, endless floor space and a deep pool of skilled people who don't mind doing mundane work while putting out fires all the time. ☹

Tony Asaro is senior analyst and founder of Voices of IT (www.VoicesofIT.com).

The world has changed and yet we're still using the same tools to manage our file data.

A barrage of buzzwords

New way to protect files

Automated storage tiering

Cloud-based DR

Storage budgets on the rise

Spotlight on unified storage

Use of cloud is evolving

Email archiving

Sponsor resources

EMC²



BIG DATA

IT'S ON  ISILON


SCALABILITY
PERFORMANCE
CAPACITY
SIMPLICITY

The Buzz on Big Data is real. What is it? How will it impact your business? And can you keep up? EMC Isilon's new generation of scale-out NAS hardware and software is the answer to every Big Data storage question.

www.isilon.com

BIG DATA: IT'S ON





Automated storage tiering is an effective way to make efficient use of installed data storage resources, and to take advantage of the high performance of solid-state storage.

Automated STORAGE TIERING: HIGHER PERFORMANCE AND LOWER COST?

BY PHIL GOODWIN

REMEMBER THOSE light beer commercials back in the 1980s with competing contingents shouting “Tastes great!” and “Less filling!” at each other? The idea was that a beer could have fewer calories without sacrificing taste. Perhaps advocates of automated storage tiering (AST) are taking a similar approach: its two goals—lower cost and higher performance—seem to be just as diametrically opposed. Historically, if you wanted higher I/O performance (data throughput) you bought high-end Fibre Channel (FC) arrays and disk devices. If budget was a bigger issue, you gravitated toward IP storage and SATA drives.

[A barrage of buzzwords](#)

[New way to protect files](#)

[Automated storage tiering](#)

[Cloud-based DR](#)

[Storage budgets on the rise](#)

[Spotlight on unified storage](#)

[Use of cloud is evolving](#)

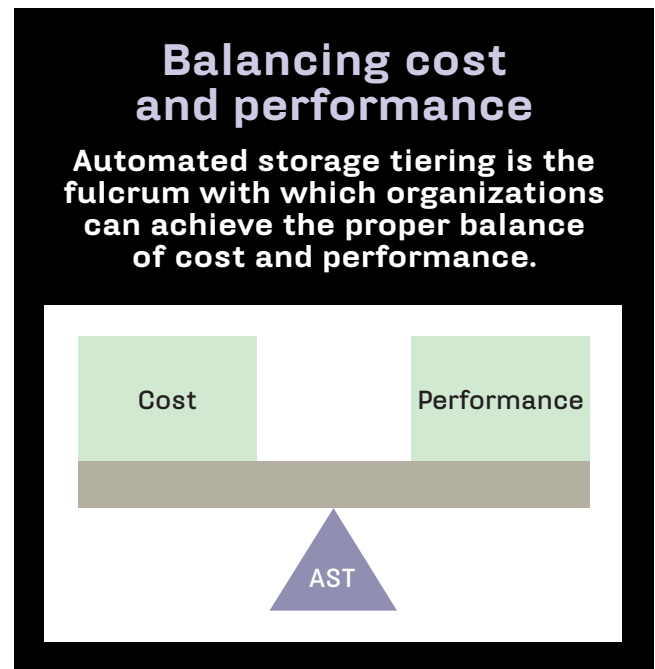
[Email archiving](#)

[Sponsor resources](#)

In practice, most companies use both types of storage in an effort to match application throughput requirements with budget constraints. That effectively represents tiered storage, and how that tiering is managed boils down to whether the staff chooses de facto manual tiering or implements an automated system. Given the increasing complexity of data storage environments, data growth and the typically poor utilization of storage, it's hard to imagine how manual tiering management is tenable for the long term.

A DELICATE BALANCE: COST AND PERFORMANCE

When storage vendors speak of their AST solutions, they all tout higher performance and lower cost. Given the dichotomy between lower cost and higher performance, one wonders whether they've somehow discovered a way to repeal the laws of physics. Fortunately for Newtonian science, the answer is no. In fact, AST can't deliver both lower cost and higher performance simultaneously. What it can do is deliver the performance needed by the application at the lowest possible cost. Thus, it's more a balancing act between the two objectives (see "Balancing cost and performance," above).



STORAGE TIERING REVIEW

Most IT professionals generally understand storage tiering, but it's worth a brief review of the concept. Tiers are defined predominantly by the performance characteristics of the underlying media. Solid-state drives (SSDs) and flash memory are referred to as tier 0; high-speed FC drives such as 15K rpm disks are tier 1; 10K rpm FC and SAS disks are tier 2; and less than 10K rpm SATA disks are tier 3. These aren't absolute rules, but they're typical tier differentiators.

Tiers are implemented in two different ways. The first is intra-array, in which a single array is populated with two or more media types. The second is inter-array, in which arrays with different media types are associated to facilitate

A barrage of buzzwords

New way to protect files

Automated storage tiering

Cloud-based DR

Storage budgets on the rise

Spotlight on unified storage

Use of cloud is evolving

Email archiving

Sponsor resources

data movement. It's also possible to have both simultaneously in the same configuration.

AUTOMATING THE TIERING PROCESS

Neither storage tiering nor AST are new technologies. In fact, Hewlett-Packard (HP) Co. claims to have implemented automated storage tiering in 1996. Nevertheless, the adoption of AST has been relatively slow. That's because the earliest implementations required a significant effort to classify data and develop the policies that governed data movement between tiers. Most often, data was moved based on age, which is rarely the best arbiter of value.

Current AST implementations use sophisticated algorithms that calculate the usage of data chunks ranging in size from a 4 KB block up to a 1 GB block, depending on vendor and settings. This calculation is done based on access demand relative to other chunks, as there's no definition of "high demand."

Data can be elevated to a higher tier during high demand periods and demoted when demand lessens. The quality of the algorithm determines the value of the product and the size of the block determines workload suitability. Smaller block sizes are generally better for random I/O, while larger sizes are better for sequential I/O.

Both established vendors and emerging vendors offer AST capabilities. Some of the newer vendors, such as Dell Compellent, have made automated storage tiering a cornerstone of their product architecture. With the company's Storage Center product line and its Fluid Data Architecture, there's only one array architecture and AST is an integrated part of it. Fluid Data Architecture data movement block size is a relatively granular 2 MB.

Similarly, for Avere Systems Inc., AST isn't an optional feature in its FXT appliances. However, it adds the ability to use any network-attached storage (NAS) or JBOD array as tier 3 storage. Thus, Avere offers both inter- and intra-array tiering. In addition, Avere uses its own file system, which gives it an additional measure of control over data movement in its algorithm. FXT is a "set-and-forget" model that doesn't allow user modification of movement policies, although tiers can be scaled separately to match workload changes.

Current AST implementations use sophisticated algorithms that calculate the usage of data chunks ranging in size from a 4 KB block up to a 1 GB block, depending on vendor and settings.

A barrage of buzzwords

New way to protect files

Automated storage tiering

Cloud-based DR

Storage budgets on the rise

Spotlight on unified storage

Use of cloud is evolving

Email archiving

Sponsor resources

For Greg Folsom, CIO at Arnold Worldwide, simplicity is the key issue. According to Folsom, Dell Compellent systems are “drop-dead easy” to install and manage. Arnold Worldwide, a Boston-based ad agency, uses a three-tier strategy with two different storage policies. “These things are so easy that even I can be talked through managing them when our storage manager is away from the office,” he joked.

Chris Elam, Arnold Worldwide’s senior systems engineer, began using Dell Compellent’s default automated tiered storage policies but tweaked them over time. Dell Compellent’s Enterprise Manager utility helped Elam identify usage patterns. “Enterprise Manager helped us to see exactly how data is accessed in the system. With this information, we created a tier 1-2 policy for some apps and a tier 2-3 policy for other applications. We’ve been using the system for more than four years and we haven’t had to change the policies in a long time,” Elam said. New volumes are simply assigned to one of the policies at creation time.

SOLID-STATE STORAGE COMPLEMENTS TIERING

Xiotech Corp. offers another example of a “set-and-forget” AST implementation. Xiotech’s Hybrid ISE product combines SSD and hard disk drives in a sealed

AUTOMATED TIERING: Buying considerations

Shopping for automated tiering for your data storage environment? Keep these key points in mind:

- Understand your application’s data usage characteristics
- Examine management tools to keep the system tuned over time
- Determine the integration of the proposed automated storage (AST) tiering capability with existing tools and vendors
- Decide if you want a “set-and-forget” or customizable AST product
- AST is a true price-to-performance play measurable in the monetary savings of devices

A barrage of buzzwords

New way to protect files

Automated storage tiering

Cloud-based DR

Storage budgets on the rise

Spotlight on unified storage

Use of cloud is evolving

Email archiving

Sponsor resources

Expert content that allows IT pros to get their job done right.


IT pros turn to **SearchDataCenter.com** for the information they require to effectively manage the complex environment of their data center. We're the only information resource that provides immediate access to breaking industry news, case studies, tutorials, technical tips, videos, a selection of highly focused data center newsletters and more—all at no cost. Discover key trends in shaping today's data center and benefit from data center-specific expertise focused on:

- Management strategies and mainframe migration policies
- Virtualization best practices for successful cloud computing
- Unix systems administration and management
- Maximizing ROI through new technologies
- And so much more

To find out more and activate your **FREE** membership today, visit:
www.SearchDataCenter.com/Free



GET ANSWERS FROM INDUSTRY EXPERTS AND YOUR PEERS



At IT Knowledge Exchange, we provide IT professionals with the necessary resources to assist in their day-to-day work. The community includes everything related to Enterprise IT, including data center design and infrastructure, cloud computing, systems monitoring, server management, and more.

- **Collaborate with your peers.** Use the discussion feature to chat about a task you've been working on and just aren't sure of the best approach.
- **Ask questions.** For all those tasks, large or small, that you can't seem to figure out, ask fellow members how they've resolved the same issue.
- **Answer questions.** Pay it forward by providing your expertise on an issue that you know the answer to and earn badges for your expertise.
- **Read expert blogs.** Our bloggers are writing about the topics that you deal with on a daily basis, including virtualization, data center cooling, and server performance.

To get answers, share knowledge, and collaborate with peers, visit IT Knowledge Exchange at **www.ITKnowledgeExchange.com**



14.4 TB 3U container. Of the 14.4 TB, 1 TB is SSD and the rest comprises 900 GB 10K rpm SAS drives (tier 2). Controller-level software, called Continuous Adaptive Data Placement, automatically manages data placement from the moment of deployment. Although the company provides a graphical ISE Analyzer utility to highlight I/O activity, in practice a user can't adjust any of the parameters or configuration. The company says it designed Hybrid ISE to never need tuning.

Among the vendors offering more configurable architectures, NetApp Inc. stresses the ability to scale performance and capacity separately. The firm's Flash Cache (PAM II) product is analogous to tier 0 SSD in other product lines. Though it can support multiple tiers, NetApp said in many cases the tiers can be simplified to two: Flash Cache and either tier 2 or 3. That's because they've found data tends to be either "hot" or "cold" and rarely in between. Buffer cache is used to buffer write activity to avoid performance degradation. Data block movement size is the most granular at just 4 KB. Although this architecture may require more flash disk than other systems (10% to 20% of total capacity), the elimination of relatively expensive tier 1 hard disks and spreading cold data across more SATA drives can result in the same performance at a lower total cost. Moreover, NetApp combines AST with deduplication and compression on the spinning disk for even greater space efficiency. Because data is managed through the WAFL file system and Data Ontap, it doesn't need to be "rehydrated" when being elevated from a lower tier to tier 0 as the data becomes hot. The same automated storage tiering capabilities apply across all NetApp product lines.

CERN, the European Organization for Nuclear Research in Geneva, uses NetApp's Flash Cache on Oracle RAC databases. "Prior to using Flash Cache, we had to size everything based on IOPS regardless of storage utilization," said Eric Grancher of the CERN IT department. "Now, we can optimize both IOPS and capacity. We have moved from expensive Fibre Channel drives to less-expensive SATA drives. This has resulted in a substantial savings for the organization." Grancher has found the NetApp system to be very adaptive to workloads resulting in simple management. His experience has determined that overall performance is better when the flash memory is in the storage rather than in the servers. "It makes more sense to have the stable NetApp systems cache the data rather than the database servers, which are restarted

Among the vendors offering more configurable architectures, NetApp Inc. stresses the ability to scale performance and capacity separately.

A barrage of buzzwords

New way to protect files

Automated storage tiering

Cloud-based DR

Storage budgets on the rise

Spotlight on unified storage

Use of cloud is evolving

Email archiving

Sponsor resources

more frequently for patching or updates. A data cache on the storage server is already ‘warmed up’ and so eliminates the inevitable periods of poor performance we would suffer with cold server-based caches after each restart,” he said.

EMC Fully Automated Storage Tiering (FAST) is another example of a more configurable system. FAST has an install wizard that allows you to implement default configurations for simple deployment, which EMC says the majority of users find sufficient in most cases for “set and forget.” Other users tap into FAST Tier Advisor, a utility that collects usage statistics over time. Those statistics can be used to apply optimized policies for specific applications. Users can also set the size of the data movement block from 768 KB to 1 GB, depending on whether the reads tend to be random or sequential.

EMC recommends that users start with approximately 3% of capacity in tier 0, 20% in tier 1 and 77% in tier 3. Tier Advisor will track usage and, over time, tier 1 should be minimized as little more than a buffer between the higher and lower tiers. In any event, Tier Advisor lets users optimize any of the tiers based on actual usage patterns.

INTER-ARRAY TIERING

Hitachi Data Systems’ (HDS) AST supports the same tool set across all product lines for inter-array tiering. It begins with virtualization to abstract and partition workloads. In fact, HDS recommends application and workload classification rather than data classification. “Organizations should avoid starting out too complex in their tiering strategy,” said Sean Moser, vice president of software at HDS. “Don’t use too many tiers and over-optimize individual applications.” Although HDS supports three tiers, as a practical matter the middle tier becomes a “shock absorber” between higher and lower tiers.

HDS offers a Data Center Management suite that includes configuration management, tuning management and tiered storage management. It provides alerts and a dashboard that gives details by volume, storage pool, service-level agreement (SLA) and peak periods. Using these tools, users can fine-tune the system over time. HDS can also incorporate other vendors’ arrays into the storage pool whereby older systems can be repurposed and used as a data archive. HDS can use spin-down drives for the archive tier to reduce power

"Organizations should avoid starting out too complex in their tiering strategy. Don't use too many tiers and over-optimize individual applications."

—SEAN MOSER,
vice president of software, HDS

A barrage of buzzwords

New way to protect files

Automated storage tiering

Cloud-based DR

Storage budgets on the rise

Spotlight on unified storage

Use of cloud is evolving

Email archiving

Sponsor resources

and cooling requirements.

HP is more traditional in its approach to automated storage tiering. Perhaps because some of its arrays come via a partnership and acquisitions, the AST capabilities vary between product lines. Its high-end P9500 systems, OEM units from HDS, behave very similarly to HDS's AST implementation, and you can use the P9500 to virtualize other arrays.

HP's 3PAR product line is a relative newcomer to AST, having rolled out those capabilities approximately a year ago. 3PAR supports three tiers, but it's largely up to users how to configure them. HP recommends monitoring the applications for usage patterns and then determining what tiers at what sizes to implement. Its Adaptive Optimization tool is available to help with the monitoring and sizing of tiers.

HP's x9000 scalable NAS uses its own AST as well. In this case, all policies are user generated. HP says automated storage tiering evolves from user policies to automation over time.

IBM's Easy Tier product is supported on its Storwize V7000, DS8700, DS8800 and SAN Volume Controller products. Currently, Easy Tier supports two tiers, one of which must be solid-state drives. Once every 24 hours, the product analyzes performance metrics and generates a plan to relocate data appropriately. Data relocation occurs in 1 GB extents, which are migrated no more often than every five minutes to avoid performance interruption. Easy Tier is a function of the array and is a no-cost option.

AUTOMATED TIERING MARKET STILL DEVELOPING

The good news about automated storage tiering is that the market is robust with many options. The bad news is that the options make comparing implementations rather bewildering. Jerome Wendt, lead analyst and president at DCIG in Omaha, Neb., has some practical advice for evaluating the appropriate solution. "First, users should match the performance needs of the application to the architecture of the product," he said. "This includes understanding the size of the data block being moved, how often it's being moved and how it's moved between tiers." Wendt further advises that file systems are fairly safe candidates for AST, but that Microsoft Exchange and databases should be approached more cautiously. ☉

Phil Goodwin is a storage consultant and freelance writer.

A barrage of buzzwords

New way to protect files

Automated storage tiering

Cloud-based DR

Storage budgets on the rise

Spotlight on unified storage

Use of cloud is evolving

Email archiving

Sponsor resources

More bang for your backup.

Quantum's DXi-Series Appliances with deduplication provide higher performance at lower cost than the leading competitor.



Quantum has helped some of the largest organizations in the world integrate deduplication into their backup process. The benefits they report are immediate and significant—faster backup and restore, 90%+ reduction in disk needs, automated DR using remote replication, reduced administration time—all while lowering overall costs and improving the bottom line.

Our award-winning DXi®-Series appliances deliver a smart, time-saving approach to disk backup. They are acknowledged technical leaders. In fact, our DXi6500 was just nominated as a “Best Backup Hardware” finalist in Storage Magazine’s Best Product of the Year Awards—it’s both faster and up to 45% less expensive than the leading competitor.

**Get more bang for your backup today.
Faster performance. Easier deployment. Lower cost.**

Contact us to learn more at (866) 809-5230 or visit www.quantum.com/dxi

Preserving The World's Most Important Data. Yours.™

©2011 Quantum Corporation. All rights reserved.

Quantum.




Blueprint for cloud- based disaster recovery



Cloud storage and computing services offer a number of alternatives for cloud-based DR depending on the recovery time and recovery point objectives a company requires.

BY JACOB GSOEDL



CLOUD COMPUTING, along with mobile and tablet devices, accounts for much of the high-tech buzz these days. But when it comes to hype, the cloud seems to absorb more than its fair share, which has had the unintended consequence of sometimes overshadowing its real utility.

Although the concept—and some of the products and services—of cloud-based disaster recovery (DR) is still nascent, some companies, especially smaller organizations, are discovering and starting to leverage cloud services for DR. It can be an attractive alternative for companies that may be strapped for IT resources because the usage-based cost of cloud services is well suited for DR where the secondary infrastructure is parked and

[A barrage of buzzwords](#)

[New way to protect files](#)

[Automated storage tiering](#)

[Cloud-based DR](#)

[Storage budgets on the rise](#)

[Spotlight on unified storage](#)

[Use of cloud is evolving](#)

[Email archiving](#)

[Sponsor resources](#)

idling most of the time. Having DR sites in the cloud reduces the need for data center space, IT infrastructure and IT resources, which leads to significant cost reductions, enabling smaller companies to deploy disaster recovery options that were previously only found in larger enterprises. “Cloud-based DR moves the discussion from data center space and hardware to one about cloud capacity planning,” said Lauren Whitehouse, senior analyst at Enterprise Strategy Group (ESG) in Milford, Mass.

But cloud-based disaster recovery isn’t a perfect solution, and its shortcomings and challenges need to be clearly understood before a firm ventures into it. Security usually tops the list of concerns:

- Is data securely transferred and stored in the cloud?
- How are users authenticated?
- Are passwords the only option or does the cloud provider offer some type of two-factor authentication?
- Does the cloud provider meet regulatory requirements?

And because clouds are accessed via the Internet, bandwidth requirements also need to be clearly understood. There’s a risk of only planning for bandwidth requirements to move data into the cloud without sufficient analysis of how to make the data accessible when a disaster strikes:

- Do you have the bandwidth and network capacity to redirect all users to the cloud?
- If you plan to restore from the cloud to on-premises infrastructure, how long will that restore take?

“If you use cloud-based backups as part of your DR, you need to design your backup sets for recovery,” said Chander Kant, CEO and founder at Zmanda Inc., a provider of cloud backup services and an open-source backup app. Reliability of the cloud provider, its availability and its ability to serve your users while a disaster is in progress are other key considerations. The choice of a cloud service provider or managed service provider (MSP) that can deliver service within the agreed terms is essential, and while making a wrong choice may not land you in IT hell, it can easily put you in the doghouse or even get you fired.

DEVisING A DISASTER RECOVERY BLUEPRINT

Just as with traditional DR, there isn’t a single blueprint for cloud-based disaster recovery. Every company is unique in the applications it runs, and the relevance of the applications to its business and the industry it’s in. Therefore, a cloud disaster recovery plan (aka cloud DR blueprint) is very specific and

A barrage of buzzwords

New way to protect files

Automated storage tiering

Cloud-based DR

Storage budgets on the rise

Spotlight on unified storage

Use of cloud is evolving

Email archiving

Sponsor resources

distinctive for each organization.

Triage is the overarching principle used to derive traditional as well as cloud-based DR plans. The process of devising a DR plan starts with identifying and prioritizing applications, services and data, and determining for each one the amount of downtime that's acceptable before there's a significant business impact. Priority and required recovery time objectives (RTOs) will then determine the disaster recovery approach.

Identifying critical resources and recovery methods is the most relevant aspect during this process, since you need to ensure that all critical apps and data are included in your blueprint. By the same token, to control costs and to ensure speedy and focused recovery when the plan needs to be executed, you want to make sure to leave out irrelevant applications and data. The more focused a DR plan is, the more likely you'll be able to test it periodically and execute it within the defined objectives.

With applications identified and prioritized, and RTOs defined, you can then determine the best and most cost-effective methods of achieving the RTOs, which needs to be done by application and service. In the rarest of cases, you'll have a single DR method for all your applications and data; more likely you'll end up with several methods that protect clusters of applications and data with similar RTOs. "A combination of cost and recovery objectives drive different levels of disaster recovery," said Seth Goodling, virtualization practice manager at Acronis Inc.

Identifying critical resources and recovery methods is the most relevant aspect during this process, since you need to ensure that all critical apps and data are included in your blueprint.

CLOUD-BASED DISASTER RECOVERY OPTIONS

Managed applications and managed DR. An increasingly popular option is to put both primary production and disaster recovery instances into the cloud and have both handled by an MSP. By doing this you're reaping all the benefits of cloud computing, from usage-based cost to eliminating on-premises infrastructure. Instead of doing it yourself, you're deferring DR to the cloud or managed service provider. The choice of service provider and the process of negotiating appropriate service-level agreements (SLAs) are of utmost importance. By handing over control to the service provider, you need to be

A barrage of buzzwords

New way to protect files

Automated storage tiering

Cloud-based DR

Storage budgets on the rise

Spotlight on unified storage

Use of cloud is evolving


Email archiving

Sponsor resources

Up to 85% of computing capacity sits idle in distributed environments. A smarter planet needs smarter infrastructure.

Let's build a smarter planet.

ibm.com/dynamic



absolutely certain it's able to deliver uninterrupted service within the defined SLAs for both primary and DR instances. "The relevance of service-level agreements with a cloud provider cannot be overstated; with SLAs you're negotiating access to your applications," said Greg Schulz, founder and senior analyst at Stillwater, Minn.-based StorageIO Group.

A pure cloud play is becoming increasingly popular for email and some other business applications, such as customer relationship management (CRM), where Salesforce.com has been a pioneer and is now leading the cloud-based CRM market.

Back up to and restore from the cloud. Applications and data remain on-premises in this approach, with data being backed up into the cloud and restored onto on-premises hardware when a disaster occurs. In other words, the backup in the cloud becomes a substitute for tape-based off-site backups.

<h2 style="text-align: center;">Cloud-based DR approaches side-by-side</h2>			
	Managed primary and DR instances	Cloud-based backup and restore	Replication in the cloud
Instances	<ul style="list-style-type: none"> • Salesforce.com CRM • Email in the cloud 	<ul style="list-style-type: none"> • On-premises into the cloud • Cloud to cloud 	<ul style="list-style-type: none"> • On-premises into the cloud • Cloud to cloud
Merits	<ul style="list-style-type: none"> • Fully managed DR • 100% usage based • Least complex 	<ul style="list-style-type: none"> • Only requires cloud storage; cloud virtual machines are optional • Usually less complex than replication 	<ul style="list-style-type: none"> • Best recovery time objectives (RTOs) and recovery point objectives (RPOs) • More likely to support application-consistent recovery
Caution	Service-level agreements define access to production and DR instances	Less favorable RTOs and RPOs than replication	Higher degree of complexity
Implemented via . . .	N/A	Backup applications and appliances	<ul style="list-style-type: none"> • Replication software • Cloud gateways • Cloud storage software such as EMC Atmos and Hitachi HCP

A barrage of buzzwords

New way to protect files

Automated storage tiering

Cloud-based DR

Storage budgets on the rise

Spotlight on unified storage

Use of cloud is evolving

Email archiving

Sponsor resources

When contemplating cloud-based backup and restore, it's crucial to clearly understand both the backup and the more problematic restore aspects. Backing up into the cloud is relatively straightforward, and backup application vendors have been extending their backup suites with options to directly back up to popular cloud service providers such as AT&T, Amazon, Microsoft Corp., Nirvanix Inc. and Rackspace. "Our cloud connector moves data deduped, compressed and encrypted into the cloud, and allows setting retention times of data in the cloud," said David Ngo, director of engineering alliances at CommVault Systems Inc., who aptly summarized features you should look for in products that move data into the cloud. Likewise, cloud gateways such as the Cirtas Bluejet Cloud Storage Controller, F5 ARX Cloud Extender, Nasuni Filer, Riverbed Whitewater and TwinStrata CloudArray, can be used to move data into the cloud. They straddle on-premises and cloud storage, and keep both on-premises data and data in the cloud in sync.

The challenging aspect of using cloud-based backups for disaster recovery is the recovery. With bandwidth limited and possibly terabytes of data to be recovered, getting data restored back on-premises within defined RTOs can be challenging. Some cloud backup service providers offer an option to restore data to disks, which are then sent to the customer for local on-premises recovery. Another option is a large on-premises cache of recent backups that can be used for local restores.

"I firmly believe that backups need to be local and from there sent into the cloud; in other words, the backup in the cloud becomes your secondary off-site backup," said Jim Avazpour, president at OS33 Inc.'s infrastructure division. On the other hand, depending on the data to be restored, features like compression and, more importantly, data dedupe can make restores from data in the cloud to on-premises infrastructure

a viable option. A case in point is Michigan-based Rockford Construction Co., which uses a StorSimple appliance for cloud-based protection of its Exchange and SharePoint infrastructures. "In case of a disaster, we'll pull VMs [virtual machines] from the cloud; with StorSimple's deduplication we pretty much have to only pull down one full VM copy and the differences for others," said Shaun Partridge, vice president (VP) of IT at Rockford Construction.

Back up to and restore to the cloud. In this approach, data isn't restored back to on-premises infrastructure; instead it's restored to virtual machines

With bandwidth limited and possibly terabytes of data to be recovered, getting data restored back on-premises within defined RTOs can be challenging.

A barrage of buzzwords

New way to protect files

Automated storage tiering

Cloud-based DR

Storage budgets on the rise

Spotlight on unified storage

Use of cloud is evolving

Email archiving

Sponsor resources

in the cloud. This requires both cloud storage and cloud compute resources, such as Amazon's Elastic Compute Cloud (EC2). The restore can be done when a disaster is declared or on a continuous basis (pre-staged). Pre-staging DR VMs and keeping them relatively up-to-date through scheduled restores is crucial in cases where aggressive RTOs need to be met. Some cloud service providers facilitate bringing up cloud virtual machines as part of their DR offering. "Several cloud service providers use our products for secure deduped replication and to bring servers up virtually in the cloud," said Chris Poelker, VP of enterprise solutions at FalconStor Software.

Replication to virtual machines in the cloud. For applications that require aggressive recovery time and recovery point objectives (RPOs), as well as application awareness, replication is the data movement option of choice. Replication to cloud virtual machines can be used to protect both cloud and on-premises production instances. In other words, replication is suitable for both cloud-VM-to-cloud-VM and on-premises-to-cloud-VM data protection. Replication products are based on continuous data protection (CDP), such as CommVault Continuous Data Replicator, snapshots or object-based cloud storage such as EMC Atmos or the Hitachi Content Platform (HCP). "Cloud service provider Peak Web Hosting enables on-premises HCP instances to replicate to a Peak Web HCP instance instead of another on-premises HCP instance," said Robert Primmer, senior technologist and senior director content services, Hitachi Data Systems.

Replication to cloud virtual machines can be used to protect both cloud and on-premises production instances.

NEW OPTIONS, OLD FUNDAMENTALS

The cloud greatly extends disaster recovery options, yields significant cost savings, and enables DR methods in small- and medium-sized businesses (SMBs) that were previously only possible in larger organizations. It does not, however, change the DR fundamentals of having to devise a solid disaster recovery plan, testing it periodically, and having users trained and prepared appropriately. ☉

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

A barrage of buzzwords

New way to protect files

Automated storage tiering

Cloud-based DR

Storage budgets on the rise

Spotlight on unified storage

Use of cloud is evolving

Email archiving

Sponsor resources



Compellent

Manage data differently with Fluid Data storage.

Always put your data in the right place at the right time for the right cost.

See a live webinar
featuring Fluid Data
technology at
compellent.com/ttdemo.



Don't have time for a webinar?

Scan and see a snapshot of customer success
with Dell Compellent storage.

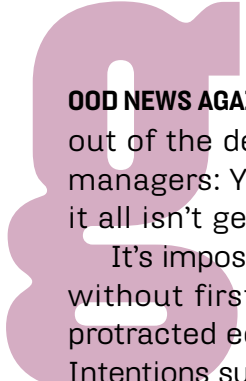
For scan app visit get.beetag.com



Storage budget

recovery on a roll

Data storage budgets continue to recover from their recessionary pounding. But while storage managers might have more money to spend, they'll need yet more capacity to meet new demands. BY RICH CASTAGNA



GOOD NEWS AGAIN for data storage managers: Your budgets are continuing to climb out of the depths of the 2008-2009 recession. Not such good news for storage managers: You'll have more data than ever to deal with, and finding a place for it all isn't getting any easier.

It's impossible these days to have a discussion about storage technology spending without first acknowledging just how deeply affected most shops were by the protracted economic downturn. But the results of our exclusive storage Purchasing Intentions survey show that this spring storage budgets—on a year-over-year basis—are up for the fourth consecutive survey (covering a two-year period).

The storage managers who participated in our survey expect an average 1.8% [increase in their storage budgets](#), led by larger companies (more than \$1 billion in

[A barrage of buzzwords](#)

[New way to protect files](#)

[Automated storage tiering](#)

[Cloud-based DR](#)

[Storage budgets on the rise](#)

[Spotlight on unified storage](#)

[Use of cloud is evolving](#)

[Email archiving](#)

[Sponsor resources](#)

revenue), where budgets are expected to increase by 3%. Even small businesses, which have struggled to get their budgets out of negative territory, will see a modest yet encouraging gain of 1.2%. To be sure, the budget change numbers are well shy of the hikes we saw routinely a few years back, but the upward trend is heartening.

On an actual dollar basis, the average data storage budget recorded on the survey is \$3 million, which is approximately the same as reported last year. As expected, that average is tilted toward the high end by larger companies, which averaged budgets of \$8.3 million.

NO RELIEF IN SIGHT FOR DATA GROWTH

As usual, storage managers will have to find ways to wring every cent from their budgets, both to accommodate new capacity demands and to implement newer technologies that will help ease the annual capacity crunch through greater efficiencies.

Storage managers already have their hands full, with the average shop now managing 263 TB of disk capacity. That's a pretty big number and it's rising fast; it's the highest average we've seen in the two years we've asked respondents about installed capacity, and it's 5% higher than last spring. With an average of 77 TB of installed capacity, small businesses boast storage environments that would have rivaled most enterprises 10 or so years ago.

And if managing more than a quarter of a petabyte of storage weren't enough, those surveyed said they'll **add an average of 43 TB** in 2011. After a dip in 2009 and relatively lower numbers on a couple of subsequent surveys, it looks like the capacity machine is cranking up again. Enterprise-class outfits are planning to add an average of 94 TB, mid-sized companies won't be far behind with plans to add 43 TB, and even small companies are looking at an additional 22 TB of disk capacity.

ABOUT THE STORAGE PURCHASING SURVEY

The *Storage* magazine/Search Storage.com Purchasing Intentions survey is fielded twice a year; this is the ninth year the survey has been conducted. *Storage* magazine subscribers and SearchStorage.com members are invited to participate in the survey, which gathers info related to storage managers' purchasing plans for a variety of data storage product categories. This edition had 833 qualified respondents across a broad spectrum of industries, with the average company size measured as having revenue of \$1.4 billion.

A barrage of buzzwords

New way to protect files

Automated storage tiering

Cloud-based DR

Storage budgets on the rise

Spotlight on unified storage

Use of cloud is evolving

Email archiving

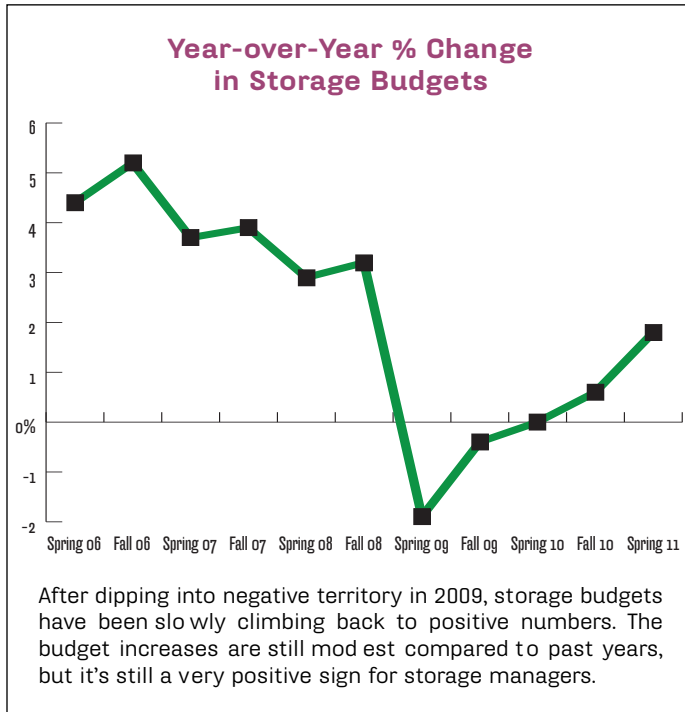
Sponsor resources

DISK SYSTEMS: A MIX OF NEW AND OLD TECHS

Most of the installed (and anticipated) disk capacity is still residing on network-attached storage (NAS) and Fibre Channel (FC) arrays. By capacity, 61% of respondents reported that they're using NAS while 59% said they were FC storage users—numbers that are in line with the results of both of last year's surveys. But taken in a larger context, there's a slow and steady shift taking place on the block storage side, with iSCSI progressively horning in on Fibre Channel's turf. In the spring of 2008, FC accounted for 70% of installed capacity vs. iSCSI at 27%; this time, iSCSI has narrowed the gap considerably with 43% (vs. 59% for Fibre Channel).

Forty-one percent of respondents have either installed a new iSCSI system or plan to this year, a figure comparable to the results on the past few surveys, suggesting the technology has attained a level of market maturity. This trend is bolstered by the high confidence levels those users demonstrate: 47% said they'll run some of their mission-critical apps on iSCSI storage. That's the highest number we've ever seen.

The sustained interest in iSCSI storage is also reflected in how storage managers plan to divvy up their systems-buying budget dollars. Nearly half (49%) will go to midrange systems, the highest percentage we've seen, as buying plans for high-end systems continue to recede. But, as we've seen for the past four-plus years, buying new systems isn't a priority at many companies as 35% of all



market maturity. This trend is bolstered by the high confidence levels those users demonstrate: 47% said they'll run some of their mission-critical apps on iSCSI storage. That's the highest number we've ever seen.

TOP 5 Purchased/Plan to purchase a disk system from	
EMC	36%
Hewlett-Packard	26%
NetApp	26%
Dell	25%
IBM	24%

- A barrage of buzzwords
- New way to protect files
- Automated storage tiering
- Cloud-based DR
- Storage budgets on the rise
- Spotlight on unified storage
- Use of cloud is evolving
- Email archiving
- Sponsor resources

TOP 5

Purchased/Plan to purchase storage switches from

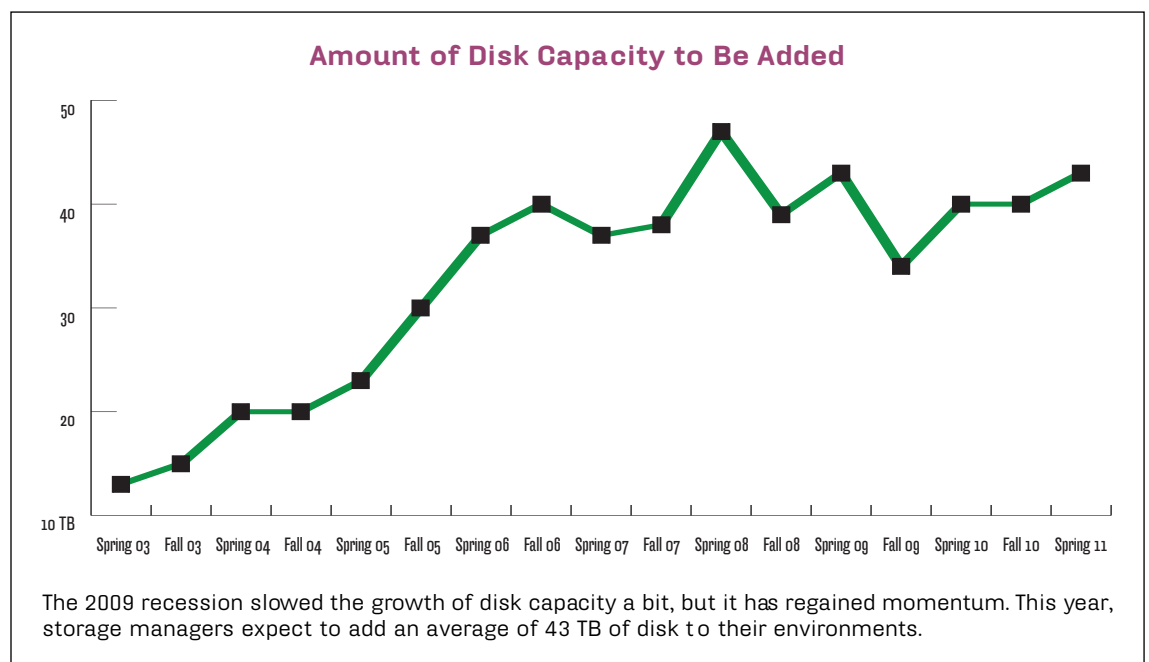
Cisco	54%
Brocade	39%
QLogic	12%
Hewlett-Packard	4%
Sanrad	1%

money going toward disk storage products is earmarked to buy additional drives for already installed systems. We expected that trend to shift back to new systems as firms refreshed their storage technologies, but it looks like the recession may have pushed those refreshes out two or three years.

File data is the fastest growing data type, but despite the increasing burden it places

on most data storage operations, we haven't seen much of a shift from traditional file storage methods. In 2011, money for file storage will be spread over several technologies, with the top three—direct-attached storage (DAS) (22%), NAS systems (22%) and NAS gateways fronting storage-area networks (SANs) (18%)—overshadowing newer techs like file virtualization and NAS clustering.

Regardless of what type of storage a shop might be looking to acquire, the price of the product will be a major factor in the purchase decision. When we asked respondents what the most important factor was in their choice of a primary disk system vendor, features and functions (as usual) came out on top



A barrage of buzzwords

New way to protect files

Automated storage tiering

Cloud-based DR

Storage budgets on the rise

Spotlight on unified storage

Use of cloud is evolving

Email archiving

Sponsor resources

A barrage of buzzwords

New way to protect files

Automated storage tiering

Cloud-based DR

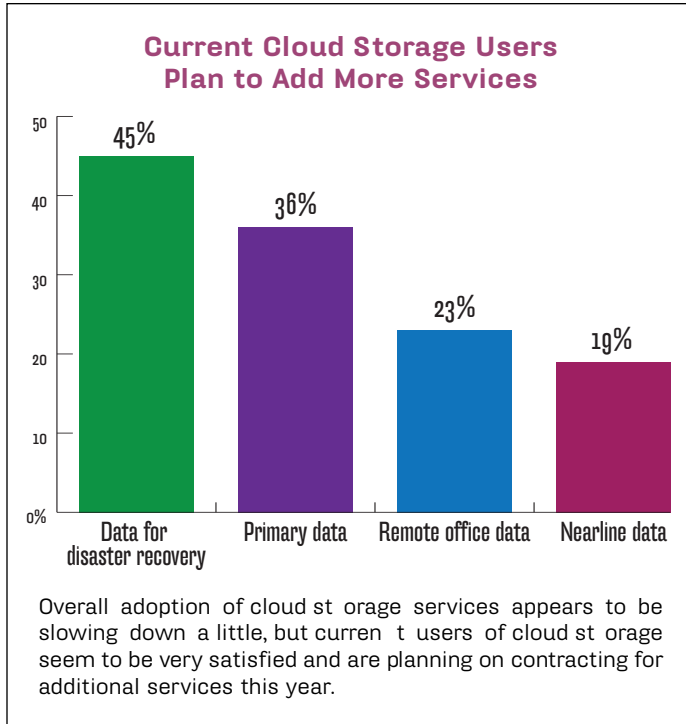
Storage budgets on the rise

Spotlight on unified storage

Use of cloud is evolving

Email archiving

Sponsor resources



with 32%. But the next most important factor—surpassing tech support and dealing with a familiar vendor—was price, as indicated by 22% of those surveyed. That may not be a particularly high number, but it’s the highest we’ve seen price rated as a factor.


CLOUD GETS LESS CLEAR

Six months ago we saw some pretty impressive numbers for cloud storage adoption that, frankly, surprised us a bit. This time, it appears the early enthusiasm for cloud storage of primary data may be wearing off a little.

Sixteen percent of respondents said they use a cloud storage service for non-backup purposes, which is considerably lower than last fall but still an improvement of two percentage points over last spring’s tally. The current numbers are actually quite good, just not as eye-popping as those from last fall. And the dip in usage may be attributable to pilot programs that were in place at the end of last year that were one-off projects or that have not evolved into production implementation.

Still, approximately 46% of survey takers said they’ll start using at least

one primary or nearline data cloud storage service in 2011, a figure that’s also a little off from the rosier 52% recorded last fall. But cloud storage service providers should take heart from the success they’ve apparently had among current users. Those users seem eager to add to their cloud storage portfolios, with 45% expecting to [add cloud storage for disaster recovery \(DR\)](#) in 2011 and 36%



expecting to add it for primary data.

Overall, considering both non-users and current cloud users, one-third plan to evaluate each of these cloud-related technologies or services:

- Private storage cloud products
- Hybrid storage arrays (integrated local storage and cloud storage)
- Cloud-based file sharing and synchronization
- Cloud-based archiving services


SOLID-STATE STORAGE SHOWING UP IN MORE SHOPS

Cloud storage might be experiencing some growing pains, but solid-state storage appears well on its way to becoming a data center mainstay. Slightly more than 16% are using solid-state storage now, representing a year-over-year gain of nearly six percentage points. Another 10.5% said they'll implement solid-state this year (vs. last year's 5.9%). And one-third will evaluate the technology in 2011, leaving only 40% without any specific [solid-state storage plans](#).

On average, current solid-state users have 6.8 TB of the stuff installed, which is a pretty impressive figure when you consider the cost of solid-state. Those who said they'll be adding solid-state storage this year will be upping their installed capacity by an average of 5.9 TB.

Most of that solid-state storage (75%) is finding its way into arrays; that option has emerged as the preferred implementation choice. Thirty percent—

much of which likely represents solid-state in PCIe form factors—is installed in servers, and another 27% of solid-state storage is used in laptops and desktop PCs.



STRIVING TOWARD EFFICIENCY

“Efficiency” has become the byword of many storage shops over the last few years, not because it has a catchy marketing ring, but because a shifting economy has permanently altered the data storage landscape. Storage managers are eager to pursue technologies that can help them make better

A barrage of buzzwords

New way to protect files

Automated storage tiering

Cloud-based DR

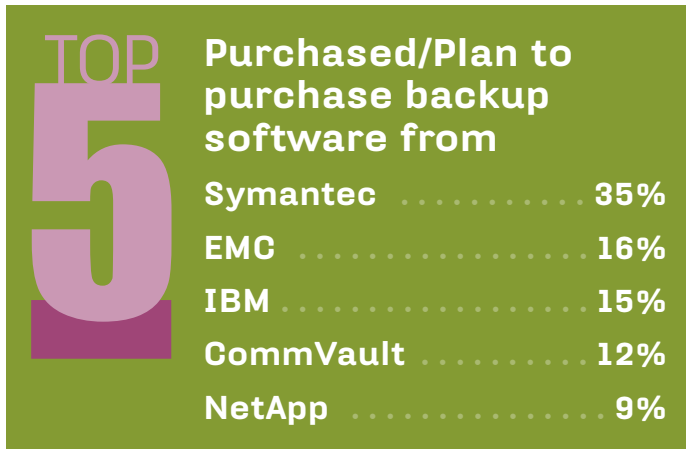
Storage budgets on the rise

Spotlight on unified storage

Use of cloud is evolving

Email archiving

Sponsor resources




use of their installed systems by ensuring that data resides on the appropriate gear, by using available capacity effectively and by removing data that's no longer accessed.

By pooling available storage resources, storage virtualization can help achieve some of these efficiencies. But despite improvements in the technology and its implementation alternatives, adoption of storage

virtualization has been relatively slow. Now, however, our survey reveals that 34% of respondents have virtualized at least some of their storage. That may be a relatively modest number that hasn't budged an awful lot over the past year and a half, but if the 39% of survey takers who said they plan to acquire storage virtualization technology this year follow through on those plans, adoption rates should improve.

On the other end of the virtualization spectrum, there are still issues related to administering storage for virtualized server environments. Fibre Channel storage (47%) is still the top choice for virtual server storage, with iSCSI making some modest gains but still far behind at 20%.



A barrage of buzzwords

New way to protect files

Automated storage tiering

Cloud-based DR

Storage budgets on the rise

Spotlight on unified storage

Use of cloud is evolving

Email archiving

Sponsor resources

There are still kinks to work out: Nearly two-thirds said they're using more storage with virtualized servers than they did before. And while only one-third said virtualizing servers has made storage management a tougher job, 50% indicated they'll be shopping for management tools in 2011 to better manage their storage for virtual servers.

Among other efficiency technologies, data reduction for primary storage is getting a lot of attention: 37% have already implemented it or will this year, and another 37% plan to evaluate data reduction products. Those numbers place the technology just behind deduplication for backup, which once again tops our list of "hot technologies."

Automated tiering software, which can help preserve high-cost disk real estate, has been or will be implemented by 23% of those surveyed, with 37% expecting to evaluate it. It's also interesting to note that 10 Gbps Ethernet products were third on the list (46% implementing and 27% evaluating); the effects of high-speed Ethernet will likely ripple through the data center, affecting both data and storage networks, along with the viability (and cost-effectiveness) of iSCSI storage systems.

LIGHT AT THE END OF THE TUNNEL?

The improving storage budget statistics over the last two years are pretty solid proof that IT shops are emerging from their economic doldrums. And that's good news as storage systems—and storage staffs—continue to be stressed by capacity and performance issues. Typically, year-over-year budget change numbers are somewhat lower on our spring surveys than on the autumn editions. Maybe this spring's good news will be even better by fall. ☉

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

A barrage of buzzwords

New way to protect files

Automated storage tiering

Cloud-based DR

Storage budgets on the rise

Spotlight on unified storage

Use of cloud is evolving

Email archiving

Sponsor resources

MirrorCloud Knocks Out the Competition!



Announcing a Business Continuity Solution in a Weight Class By Itself

Introducing MirrorCloud, the next-generation business continuity solution from Zenith Infotech! While the competition was playing catch-up we were creating the next breakthrough in business continuity.



MirrorCloud is a feature-rich, add-on to the robust SmartStyle Computing platform which continuously mirrors data from Windows-based servers and desktops to the scalable SmartStyle Cloud Servers. It is expandable up to 100 terabytes (TB) and offers redundancy over a thousand times more reliable than RAID 6!

Score a 'technical' knock out with your customers today!

Learn More!

Get the free mobile app at <http://gettag.mobi>



The Next-Generation in Business Continuity:

- Continuous Data Mirroring
- Virtual Failover
- Live Bare-Metal Restore
- Granular Exchange Recovery
- Off-Site Storage
- Colo-Level Virtualization



Infotech Ltd



Virtualization, cloud shine spotlight on unified storage

Unified storage adoption is starting to ramp up as data storage pros see the need for simplifying storage so it can be powered, cooled and managed in one pool.

A barrage of buzzwords

New way to protect files

Automated storage tiering

Cloud-based DR

Storage budgets on the rise

Spotlight on unified storage

Use of cloud is evolving

Email archiving

Sponsor resources

ALTHOUGH IT SEEMS like we've been hearing about unified storage forever, it's still relatively new—and that means we're fairly early in the adoption cycle. But it's clear that unified, or multiprotocol, storage has a pretty attractive value proposition. In a unified storage environment, data storage becomes a shared resource pool, available to store either block or file data that can be configured to meet application needs as they arise. So it comes as no surprise that there's significant user interest in deploying unified storage platforms. In a recent survey of 306 IT professionals with storage planning or decision making responsibilities, Enterprise Strategy Group (ESG) found that 70% of those surveyed have either deployed or are planning to deploy unified storage: 23% have deployed the technology, while 47% are still in the planning phase.

WHY UNIFIED STORAGE

Our figure of one out of every four surveyed IT users deploying unified storage is significant in that data storage users are notoriously conservative when it comes to adopting new technologies, and for good reason. The adage "If it ain't broke, don't fix it" is alive and well in storage infrastructure teams.

If a storage array fails and data is inaccessible or lost, it could cost a firm millions of dollars and the storage administrator could lose their job. Users have been dealing with having separate systems for block and file data, and are used to it. They'll continue their current, stovepiped approach until they're sufficiently comfortable the technology has matured and there's no risk in adoption, or their corporate budgets demand a more affordable, flexible and efficient solution. Our research indicates it may be a matter of both.

Unified storage can increase operational efficiency by providing a single shared pool of storage that can be used where and when needed, eliminating the need to deploy, power, cool, and manage separate block and file systems. This simple reduction in the number of systems to deploy can go a long way in reducing operational costs, never mind the flexibility afforded to the business

from having a system that can be deployed in whatever capacity needed (without having to pay the price of having guessed wrong when doing their capacity planning exercise).

Virtualized environments present an even greater challenge. Using standards-based commodity physical servers, new virtual servers and applications can be deployed in a fraction of the time it used to take in a physical world, and the virtual machines could need either file or block storage to support apps. A fluid virtual server environment creates a requirement for a fluid, responsive storage environment. Yet storage continues to be fragmented and specialized. Unified storage goes a long way in alleviating these issues.

USAGE TRENDS

ESG research finds a clear correlation between the number of systems under management and unified storage adoption. A whopping 80% of those with 26 to 100 discrete storage systems, and 83% of those with 100 or more systems, have either deployed or plan to deploy unified storage—and those with 100 or more systems are leading the early adopter category, with 32% having already deployed unified storage. This corresponds

to ESG's spending data results that users continue their drive to reduce their overall cost of doing business, especially on the operational cost front.

It follows that we would see a strong correlation between unified storage adoption and satisfaction with utilization rates, as unified storage eliminates specialized block or file stovepipes, and that's what our research shows. Eighty-nine percent of early adopters are mostly or completely satisfied with their utilization rates vs. 77% of those currently not using unified storage. We see the biggest differential with those reporting they're completely satisfied, with nearly a third of early adopters falling in this category, two-and-a-half times the number of non-adopters that are completely satisfied. Significantly, not a single unified storage adopter responded they were "not at all satisfied."

Eighty-nine percent of early adopters are mostly or completely satisfied with their utilization rates vs. 77% of those currently not using unified storage.

UNIFIED STORAGE DEPLOYMENT ALTERNATIVES

Today, users have multiple approaches to deploying unified storage; they can deploy a unified storage system, which is an integrated system that supports both block and file data, or they can deploy a file gateway that

A barrage of buzzwords

New way to protect files

Automated storage tiering

Cloud-based DR

Storage budgets on the rise

Spotlight on unified storage

Use of cloud is evolving

Email archiving

Sponsor resources

attaches via a storage-area network (SAN) to block storage shared with other applications. Our research indicates there isn't a strong preference for either approach, with 30% of respondents using or planning to use a unified system, 32% a gateway and 35% planning to use both approaches.

There are certainly business cases that can be made for both. Gateways allow users to redeploy existing block storage investments to support file data by adding a "file personality" to the front end. But the downside is that the SAN-attached block storage and the gateway are truly two distinct components that need to be managed. Unified systems don't carry the attraction of allowing users to tap into existing SAN assets, but they do reduce the number of systems under management. ESG expects to see the continued trend of users taking both approaches to unify their data storage environments because users must deal with properly allocating existing investments in concert with adding new systems.

THE BOTTOM LINE

While specific implementation strategies may still be undetermined, ESG's research clearly finds unified storage will become more common. It's attractive in terms of both IT and financial efficiency—a winning combination by any standard. ESG's findings reveal a clear desire for improved system efficiency as IT groups look to optimize their current storage infrastructure investments in light of continuing data growth and the ongoing tough macro-economic climate.

In addition to covering up past IT sins such as poor capacity utilization, unified storage can help IT organizations accelerate infrastructure consolidation and resource optimization, which are crucial components to future visions of dynamic, highly virtualized or private cloud computing environments. Indeed, as "cloud" becomes a more common model for the consumption of IT resources, there's another explicit value for the standardization that unified storage can deliver. ☺

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

A barrage of buzzwords

New way to protect files

Automated storage tiering

Cloud-based DR

Storage budgets on the rise

Spotlight on unified storage

Use of cloud is evolving

Email archiving

Sponsor resources

Memorizing RAID level definitions and knowing which level does what can be:

- A Confusing
- B Hard to Remember
- C Useful
- D All of the above

So how much do you think you know about RAID?

Find Out For Yourself and Test Your Knowledge with Our Exclusive RAID Quiz!

And don't forget to bookmark this page for future RAID-level reference.

Test your knowledge at [SearchSMBStorage.com/RAID Quiz](https://SearchSMBStorage.com/RAID_Quiz)



SearchStorage.co.UK

The Web's best storage-specific information resource for IT professionals in the UK



Where is the cloud storage market headed?

Break down the cloud storage services market and you'll find players both big and small jockeying for position in key segments.

THE CLOUD STORAGE market is just beginning to hit its stride. For the past few years, cloud storage was largely the province of developers, who have used it as a cost-effective, pay-as-you-go resource to park data for particular projects. But now we're beginning to see the cloud being embraced by traditional IT teams for a whole new set of storage applications. Based on conversations with vendors and users, we believe 2011 will be a cross-over year with mid-sized and enterprise IT stepping up to drive the cloud storage agenda and, increasingly, the adoption of cloud storage technologies.

This shift from development to production is one of the macro trends shaping the market for cloud storage products, profiled in Taneja Group's "Emerging Market Forecast for Cloud Storage." Based on our research, the cloud storage products market is currently a \$4 billion space that will grow to almost \$10 billion by 2014. The cloud will sharply influence the characteristics of next-generation data storage technologies, including how and where they get deployed.

In looking at where the cloud storage market is headed, we find it useful to divide the market into two broad areas: primary storage technologies behind the cloud; and technologies that enable users, systems and applications to connect to the cloud. Much of the first wave of competitive activity falls into the latter bucket, so let's focus on that first.

Based on our research, the cloud storage products market is currently a \$4 billion space that will grow to almost \$10 billion by 2014.

CLLOUD-CONNECTING TECHNOLOGIES

We see three major technology categories that enable connections to the cloud:

- **General-purpose gateways.** As public and private clouds become more

A barrage of buzzwords

New way to protect files

Automated storage tiering

Cloud-based DR

Storage budgets on the rise

Spotlight on unified storage

Use of cloud is evolving

Email archiving

Sponsor resources

pervasive, users will need faster and more cost-effective access to their cloud-based storage. Improved access will come in several forms, including general-purpose gateways, which are devices that connect users to content and primary I/O storage. Vendors such as Cirtas, Nasuni and TwinStrata, have already introduced such products. While small today, this segment promises to grow well in excess of 100% per year through 2014.

- **Cloud-based backup.** A second category of access solutions will enable cloud-based backup, which lets users connect backup data to cloud repositories across the wire. Established suppliers such as CommVault, Riverbed Technology (with its Whitewater product) and Symantec are already offering solutions. This segment will grow rapidly, though not quite at the two-times-per-year clip of general-purpose gateways.

- **Data movement and access.** Buoyed by the continuing growth of virtual machines, applications and storage repositories, and the need to overcome the constraints of long distances and increased latency, data movement and access products will play a big role in allowing users to efficiently move large chunks of information and interact with cloud-resident content. Cisco Systems, Juniper Networks and Riverbed (Steelhead products) will be among the primary participants here. Riverbed, in particular, could emerge as a breakout leader in this market segment. They've always been about accessing distributed stuff; now they're also connecting distributed stuff in new ways.

Improved access will come in several forms, including general-purpose gateways, which are devices that connect users to content and primary I/O storage.

REINVENTING PRIMARY STORAGE FOR CLOUD INFRASTRUCTURES

Primary storage behind the cloud represents a market that will undergo significant change as traditional storage players—including industry behemoths—adapt their technologies and offerings to the new storage model. We divide primary storage in the cloud into two major buckets: content and I/O.

Content will need to be stored, accessed and distributed differently than primary I/O storage. File technologies that have met demands for content in traditional infrastructures typically don't have the scalability and accessibility required to service content needs in the cloud. Instead, content in the cloud

A barrage of buzzwords

New way to protect files

Automated storage tiering

Cloud-based DR

Storage budgets on the rise

Spotlight on unified storage

Use of cloud is evolving

Email archiving

Sponsor resources

will largely be supported by object technologies, which will enable content and archival storage to thrive in highly scalable, multi-tenant, web-accessible repositories. This market will be driven primarily by service providers in the near term, but will eventually find uptake in private clouds within enterprise walls. We expect smaller players such as DataDirect Networks (with Web Object Scaler), Nirvanix (hNode) and Mezeo (Cloud Storage Platform) to join major vendors like EMC, Hewlett-Packard (HP) and NetApp as platform providers for cloud-based content storage. The growth will be solid, but not as spectacular as what we'll see in most of the cloud-connecting markets profiled above.

That brings us to the largest cloud storage opportunity of all: the market for primary I/O behind cloud infrastructures. Already more than \$2 billion in size, this market is being served principally by a subset of next-generation Fibre Channel technologies, although unified storage products are also playing a role. We believe primary I/O storage will experience a renaissance in the cloud, driven in large part by intelligent block technology. Intelligent block will rapidly displace legacy systems as the storage behind both private and public cloud infrastructures, and will largely differentiate winners from losers among storage system vendors. We believe that Dell (EqualLogic), HP/3PAR and NetApp will all prosper as providers of primary I/O storage behind the cloud. HP's 3PAR platform, in particular, is a system to watch. 3PAR has long targeted this space as a utility storage innovator across service providers and enterprises, and has some unique business programs currently under the Cloud Agile banner.

CLOUD: THE NEW BATTLEGROUND

While it's too early to definitively pick winners and losers, we're confident the rapidly growing cloud market will significantly shuffle positions on the data storage vendor leader board. The winners in this battle will find success by executing the right business model on top of the right platforms that enable scale-out and utility storage. ☉

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

A barrage of buzzwords

New way to protect files

Automated storage tiering

Cloud-based DR

Storage budgets on the rise

Spotlight on unified storage

Use of cloud is evolving

Email archiving

Sponsor resources

Capacity and the law drive email archiving

More than half of the respondents to our most recent survey are archiving their companies' email, approximately the same number as when we asked about it nearly two years ago. But nearly 33% of non-archivers say tucking away old email is on their to-do list. The toughest part of email archiving, say 64%, is managing the volume of archived mail. More than 75% are trying to cope with mail volume by imposing mailbox quotas, with an average upper limit of a little more than 1.5 GB, about 20% less generous than in our 2009 survey. Forty-six percent let users do their own archiving using .PST files, while 10% use Exchange Server 2010's "better-than-PST" archive mailbox feature. In any event, email admins are letting messages mellow longer before archiving: only 37% say mail is archived after 90 days vs. 47% in 2009. Why bother with email archiving at all? Twenty-eight percent claim it's for legal issues, while 26% use it for capacity management. —Rich Castagna

A barrage of buzzwords

New way to protect files

Automated storage tiering

Cloud-based DR

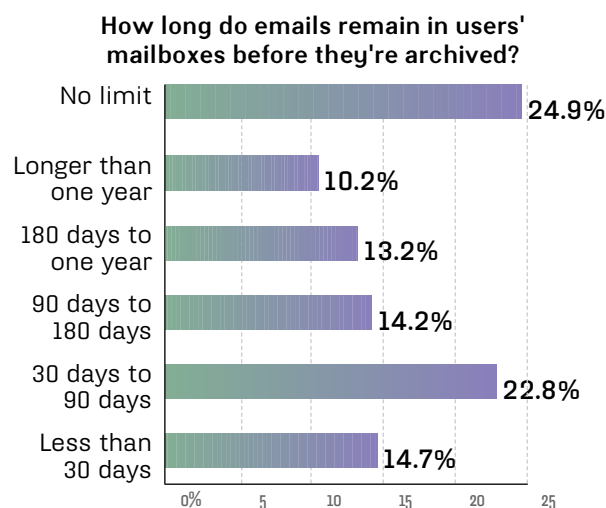
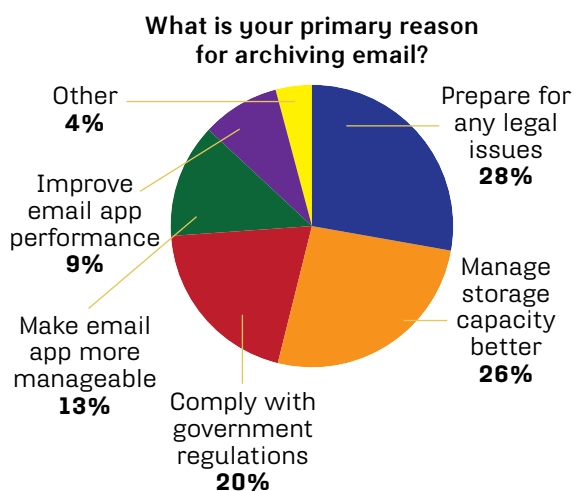
Storage budgets on the rise

Spotlight on unified storage

Use of cloud is evolving

Email archiving

Sponsor resources



What are your top email archiving challenges?*

- 64% Managing the volume of archived emails
- 38% Searching for archived email
- 36% Recovering archived emails
- 29% Setting up archiving policies
- 27% Setting up retention policies
- 20% Meeting compliance requests
- 20% Addressing the performance impact of the email application
- 19% Certifying the destruction of old archive data

*Respondents selected their top three choices

46%

Allow users to continue to maintain their own .PST files.

"We're focusing on our legal and regulatory requirements, and giving direction on what to keep and what to pitch in order to reduce the volume of stored information."

—Survey respondent

TechTarget Storage Media Group

STORAGE

Vice President of Editorial Mark Schlack

Editorial Director Rich Castagna

Senior Managing Editor Kim Hefner

Executive Editor Ellen O'Brien

Creative Director Maureen Joyce

Contributing Editors

Tony Asaro, James Damoulakis,
Steve Duplessie, Jacob Gsoedl, W. Curtis Preston



Executive Editor Ellen O'Brien

Senior News Director Dave Raffo

Senior News Writer Sonia Lelii

Features Writer Carol Sliwa

Senior Managing Editor Kim Hefner

Associate Site Editor Megan Kellett

Editorial Assistant Allison Ehrhart



Senior Site Editor Andrew Burton

Managing Editor Heather Darcy

Features Writer Todd Erickson



Senior Site Editor Sue Troy



Senior Site Editor Sue Troy

UK Bureau Chief Antony Adshead

Storage Decisions

TechTarget Conferences


Director of Editorial Events Lindsay Jeanloz

Editorial Events Associate Jacquelyn Hinds

Storage magazine Subscriptions:
www.SearchStorage.com

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

STORAGE



COMING IN JUNE

Data reduction in primary storage

Data reduction in primary storage might be the hottest topic in storage systems today. While some systems vendors can already boast data reduction features, others are scurrying to add those capabilities, even scooping up the startups that provided much of the innovation in primary storage data reduction. We'll cover the players and methods, and offer implementation suggestions.

Storage management apps for virtual storage environments

Virtualized servers have created numerous problems for data storage managers. Some storage management products have adapted to this new environment, providing the ability to track virtual servers, the apps they host and the storage they use.

Quality Awards VI: Backup and recovery software

In this sixth round of the *Storage* magazine/SearchStorage.com Quality Awards, we survey storage managers to gauge their satisfaction with the backup apps they're using. CommVault has dominated this category, winning four of the five previous surveys.

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

A barrage of
buzzwords

New way to
protect files

Automated
storage tiering

Cloud-based DR

Storage budgets
on the rise

Spotlight on
unified storage

Use of cloud
is evolving

Email archiving

Sponsor
resources



[See ad page 4](#)


- **Storage Management: Be more responsive to business needs by strategically managing data storage volumes and performance**
- **Smarter Storage Management: Fine-tuning storage infrastructure can help you do more without spending more**



Compellent

[See ad page 28](#)

- **Webcast: Learn how to manage data differently with Fluid Data storage. Register today!**
- **White Paper: iSCSI vs. Fibre Channel SANs: Three Reasons Not to Choose Sides**



[See ad page 24](#)

- **Reduce your data storage footprint and tame the information explosion**
- **Leverage the IBM Tivoli advantages in storage management**
- **Virtualize Storage with IBM for an Enhanced Infrastructure**



[See ad page 8](#)

- [Live Demos: Automated PC Backup Solutions and More](#)
- [Tour: Offsite Tape Vaulting](#)



[See ad page 11](#)

- [ESG Report - Scale-out 2.0 Storage](#)
- [Isilon Scale-out Storage Platform Overview](#)



[See ad page 20](#)

- [Quantum DXi Validation Report](#)
- [Deduplication for Dummies](#)



Infotech Ltd

See ad page 37

- **Backup To The Compute Cloud**
- **Checklist: Key factors in planning a virtual desktop infrastructure**
- **The first step toward a virtual desktop infrastructure: The assessment**