Managing the information that drives the enterprise

STORAGE

Vol. 8 No. 7 October 2009

GUIDE TO

Snapshot Technologies

Snapshots can grab images at any point in time, easing backup burdens and dramatically shortening recovery times. But one size doesn't fit all. See what type of snapshot technology is best for your environment. **p. 11**

ALSO INSIDE

TechTarget
The Technology Media ROI Experts


3PAR
Serving Information


i365
A Seagate Company


syncsort
PERFORMANCE DRIVEN


Quantum


ZENITH
Infotech Ltd

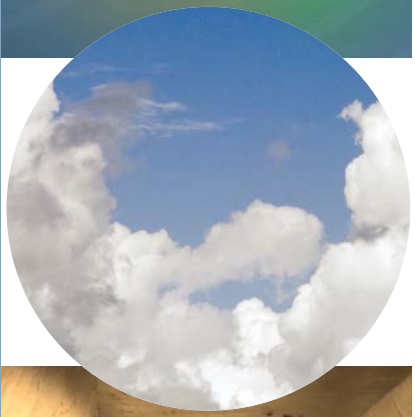

Throwing Caution to the Clouds

Backup in a Snap: A Guide to Snapshot Technologies

Storage Redux: Purchase Plans Reviving

What's Inside Internal Storage Clouds?

NFS 4.1's pNFS: Big NAS Performance Boost

Growing Need for Email Archiving

Vendor Resources

Cover illustration by ENRICO VARRASSO

Managing the information that drives the enterprise

STORAGE

The Essential Guide to Data Deduplication

data deduplication

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

Exclusive Extra for Storage Subscribers

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

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

Throwing caution to the clouds

Cloud storage is coming and it will undoubtedly bring benefits, but there may be residual effects to consider.

BEFORE ANYONE ACCUSES ME of being a Luddite or some sort of cave dweller, let me say that I'm a great fan of innovation and I love technology. But at the same time, I tend not to jump on a new tech bandwagon until it takes a few spins around the block. To put that in storage terms, while I believe cloud storage will play a part in the future of a lot of storage shops, I also think we should put the brakes on that bandwagon a bit and consider all of the ramifications of the new storage world order it promises.

The main concern I hear from storage managers is that they're reluctant to ship their data off to a cloud provider's site. That's worrisome, for sure, but I think that hurdle can actually be overcome relatively soon. Plenty of shops have been shipping data offsite for years by using third-party disaster recovery services. But using a third-party facility to store a third or fourth copy of your data for emergency use is a lot different than using it for primary storage or even backup.

Is it possible that the savings gained from using cloud storage could be offset or even outweighed by the need to build a better telecom infrastructure?

There are a number of other real issues that have gotten some air time but probably haven't been examined as closely as they should be, like the need for faster and wider communications pipes. Is it possible that the savings gained from using cloud storage could be offset or even outweighed by the need to build a better telecom infrastructure?

More perilous, I think, is the commoditization of these storage-related services. A cloud-based storage service—no matter how much customization it might provide—isn't likely to be as tailored, tweaked and tuned to a company's specific needs as an in-house setup. But the relative low cost of these services will be appealing to management, and the urge to bypass the expertise of the resident IT crew will be tough to resist. The traditional role of IT could change dramatically, turning into what could amount to a clearinghouse for external storage services.

Again, staffing and capex savings might be enticing, but what happens when that dwindling expertise is needed again? It's likely that those special skills and expertise will have also drifted off into the cloud. It can be argued that similar evolutions have happened before and IT has survived intact. In a lot of companies, application development teams have evolved into off-the-shelf package evaluators/customizers, with less and less actual development and coding being done in-house. But storage requires different care and feeding. Given the wide variety of storage solutions available and the way they can be applied to unique environments, I don't think any kind of mass commoditization would be beneficial or realistic.

Which brings up another point: If a company's storage moves offsite and into the cloud, it could affect more than just the profile and configuration of its data centers. There could be a big shift in the design and manufacture of all storage systems. Cloud service providers don't use the same storage systems that corporations do—that was tried back when storage as a service was first unveiled approximately 10 years ago. The economics of carving up a big enterprise-class array for a number of clients just never worked out. Today's cloud vendors tend to cobble together their storage systems into more economical and utility-like configurations. Those setups, which would seem rather unsophisticated compared to enterprise storage systems, might work well for delivering storage from the cloud, but they're not likely to suffice as in-house systems where performance may be paramount.

If a company's storage moves offsite and into the cloud, it could affect more than just the profile and configuration of its data centers.

So as cloud storage grows, will storage vendors shift their efforts to building systems for cloud providers? In a recent op-ed piece in *The New York Times* ("Lost in the Cloud," July 19, 2009), Jonathan Zittrain, a Harvard Law School professor, expressed some apprehensions about cloud computing and how it could affect innovation: "But the most difficult challenge—both to grasp and to solve—of the cloud is its effect on our freedom to innovate." Although most of his discussion focuses on PC users, Zittrain is most concerned with software development falling from the hands of the many into the hands of a few service providers. With software the main differentiator among storage systems, a similar scenario could unfold relating to the development of new storage products.

Similar thoughts were expressed by Wikibon.org's David Floyer in a recent Wikibon newsletter, but more from the perspective of storage developers. "Wikibon predicts that [in the next decade], total spending on external IT services will exceed the declining traditional IT budget in most shops," Floyer wrote. He went on to describe how the attention

of storage and other IT providers could (or should) shift, and added, "The traditional model of selling to IT organizations will be a declining market."

I don't know if Zittrain and Floyer raise issues about cloud storage and providers that will never come to pass, or if they're just a step or two ahead of the rest of us (or me, at least). In the end, it's the storage purchasing patterns of corporations that will go a long way in determining if the pictures they paint will develop as described. ⊙

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 Nov./Dec. issue.*

Are you prepared?

Learn why over 22,000 businesses trust i365 EVault Disaster Recovery Solutions

Backup and Recovery

Bare Metal Restore

Remote Disaster Recovery

Data Replication

Disaster Recovery Assessments

Remote and Branch Offices

DOWNLOAD OUR DR WHITEPAPER

i365
A Seagate Company

www.i365.com | 1.877.901.DATA | concierge@i365.com

COMING IN NOV./DEC.

Hot Technologies for 2010

Each year the editors of *Storage* magazine and the SearchStorage.com websites confer with analysts, consultants and users to determine the newer technologies likely to end up in data centers. We describe the technologies and how they'll improve storage shop efficiencies, list those on the verge of having an impact and grade ourselves on our previous predictions.

Storage Salaries

Storage magazine's annual Storage Salary Survey asks storage professionals across the country about their compensation, benefits and perks. As the only salary survey that focuses exclusively on storage-related professions, our analysis provides insight into the state of storage professions, and provides a look at what skills and experience might translate into higher pay.

Quality Awards IV: NAS Systems

In this fourth edition of the *Storage* magazine/SearchStorage.com Quality Awards for NAS arrays, users rank the leading systems for initial product quality, sales-force competence, product features, product reliability and technical support. See if NetApp can repeat its double win from the last NAS Quality Awards survey.

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

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

Subscriptions:
www.SearchStorage.com

Managing the information that drives the enterprise

STORAGE

The Essential Guide for DR Planning

Developing and Refining a DR Plan

Preparing a disaster recovery (DR) plan can be a complex and resource-consuming project. Use these tips to simplify the process and put together a predictable and comprehensive DR strategy.

Exclusive Extra for Storage Subscribers

Preparing a disaster recovery (DR) plan can be a complex and resource-consuming project. Fortunately, The Essential Guide to Developing and Refining a DR Plan highlights a number of expert tips to simplify the process and put together a predictable and comprehensive DR strategy. Discover how to take the complexity out of DR planning with 24/7 access to this new Essential Guide in an easy-to-access, easy-to-read, easy-to-print online format.

Get Your FREE Copy Today: **www.SearchDisasterRecovery.com/DR_Planning**

Backup in a snap: A GUIDE TO snapshot technologies

Snapshot technologies are commonly used to enhance data backup systems and dramatically shorten RTOs and RPOs. But you need to know how snapshot implementations can vary, and what those differences could mean to your environment.

By Marc Staimer

SNAPSHOT is commonly defined as a copy of a set of files, directories and/or volumes as they were at a particular point in time. As its name suggests, a snapshot is very much like a photograph because it captures an image of a certain set of data at a specific moment or point in time.

Snapshot technology was originally architected to solve several data backup problems, including:

- Backing up data that's too large to complete in the allocated time
- Failing to back up data because it has moved from a directory that hasn't been backed up to one that already has
- Corruption of backed up data that can occur when it's being written to while it's being backed up
- The affect on application performance while a backup is in process

All of these backup problems can be resolved with snapshots. But snapshots shouldn't be considered a backup panacea. There are some issues with snapshots that require workarounds (see "Snapshot snafus," p. 14).

A series of steps are required to initiate a snapshot:

1. It starts with a command that a backup is about to occur.
2. This command tells the system to quiesce the file system and apps running at that point in time.
3. The file system is then flushed so that any pending file transactions are completed.
4. The snapshot is then created.
5. Afterwards, the file system and applications are released to resume normal operations.

Snapshot technology has also moved beyond just data protection. Snapshots are an efficient and nondisruptive way to test application software against real data without endangering live production data.

Snapshot technology has also moved beyond just data protection. Snapshots are an efficient and nondisruptive way to test application software against real data without endangering live production data. They're also ideal for data mining and ediscovery. Snapshots have also evolved into a very effective—even preferred—disaster recovery methodology that protects against malware, human errors and data corruption.

WHERE SNAPSHOT TECHNOLOGY LIVES

The common perception may be that snapshotting is a storage system feature, but that's only one place that the technology may reside. Snapshot technologies are generally available in seven different types of implementations:

1. File systems of servers, desktops and laptops
2. Logical volume managers (LVMs)
3. Network-attached storage (NAS)

4. Storage arrays
5. Storage virtualization appliances
6. Server virtualization hypervisors
7. SQL databases

FILE SYSTEM-BASED SNAPSHOTS

File system-based snapshots are available in Microsoft Corp.'s Windows NTFS via Volume Shadow Copy Services (Shadow Copy in Vista); Novell Storage Services (NSS) on NetWare 4.11 or better; Novell's OES-Linux in SUSE Linux; and the Zettabyte File System (ZFS) on Sun Microsystems Inc.'s Solaris and Apple Mac OS X 10.6 (Snow Leopard).

One of the advantages of file system-based snapshot is that it tends to be "free" because it comes with the file system. It also works well and the latest file systems make it pretty easy to use. On the downside, each file system must be managed separately, which can become onerous as the number of systems proliferates. It also means that if snapshot replication is required, each file system must be set up to replicate its own snapshots. In addition, different file systems will likely vary in the kinds of snapshots they provide; snapshot frequency; the amount of capacity that must be reserved (if capacity must be reserved); as well as snapshot set up, operations and manageability. The complexity increases as more servers and file systems must be managed.

One of the advantages of file system-based snapshot is that it tends to be "free" because it comes with the file system.

LVM SNAPSHOTS

LVM snapshot technology is available with Hewlett-Packard (HP) Co.'s HP-UX Logical Volume Manager, Linux Logical Volume Manager and Linux Enterprise Volume Management System; Microsoft's Logical Disk Manager for Windows 2000 and later; Sun Solaris 10 ZFS; and Symantec Corp.'s Veritas Volume Manager (part of Symantec Veritas Storage Foundation).

Logical volume manager snapshot technology can sometimes run across a number of file systems; for example, Symantec's Veritas Volume Manager can function with most common operating systems. LVMs also usually include storage multi-pathing and storage virtualization features.

When using LVMs, there are typically additional costs per server for license/maintenance fees. You may also confront the same issues of coordination and complicated implementations found with file system-based snapshots.

NAS SNAPSHOTS

Network-attached storage is essentially an optimized or specialized file system running on an appliance or an appliance integrated with storage. Most midrange and enterprise-class NAS systems provide snapshot capabilities, including those with proprietary operating systems and the wide variety of NAS systems that are based on Microsoft Windows Storage Server.

There's a lot to like about NAS-based snapshotting, including a common standard for all of the physical and virtual servers, desktops and laptops that connect to the NAS device. It's also very easy to implement, operate

SNAPSHOT SNAFUS

SNAPSHOT PROBLEMS can occur when structured data is involved, such as databases and applications built around databases like email, enterprise resource planning (ERP) or customer relationship management (CRM). Most snapshot technologies aren't integrated with structured data applications, so when a snapshot is executed the snapshot doesn't wait for the database to be quiesced, the cache to be flushed, writes completed, and index and metadata to be updated. If the snapshot is taken when data is in the cache, or before all of the updates are completed, the snapshot isn't crash consistent—it's corrupted.

This is less of an issue for structured data applications running on a Windows server if the snapshot technology takes advantage of Windows Volume Shadow Copy Services (VSS) through its API. VSS is designed to specifically work with structured data applications, and it does all the heavy lifting of quiescing the database, flushing the cache, and completing the writes and updates before initiating the snap.

Unfortunately, there isn't an equivalent service or API in Linux or Unix operating systems. VMware Inc. has a partial solution via its vCenter storage API. The API will allow a snapshot technology to send a command to vCenter telling it to quiesce the virtual machine and then take a snapshot. At this time, it's not structured data application-aware so the snapshot may not be crash consistent.

But there's an excellent workaround for snapshotting structured data applications in a crash-consistency manner without using Windows VSS. The workaround requires backup software that integrates with the snapshot technology's API so it can leverage structured data application agents from the backup software. The agent quiesces the application, flushes the cache, completes the writes and updates, and then tells the backup software to notify the snapshot technology to perform the snapshot. This workaround is a relatively effective solution.

and manage. NAS-based snapshot technology tends to be integrated with Windows Volume Shadow Copy Services (VSS), as well as with backup servers and their agents. Some NAS vendors have their own agents for non-Windows structured data applications. Other NAS snapshot offerings include data deduplication (EMC Corp., FalconStor Software Inc. and NetApp), and some even offer thin snapshot provisioning that minimizes the amount of storage reserved for snapshots.

But there's a price to pay for the convenience and added features: fairly hefty software licensing and maintenance charges that are often system or capacity based. NAS systems tend to proliferate in most companies and, as they do, the number of touchpoints required for snapshots will also increase, making operations and management more complex.

STORAGE ARRAY-BASED SNAPSHOTS

Storage array-based snapshots are included with most block-storage array's operating systems.

The advantages of using snapshotting that comes with the storage array operating system are similar to those of NAS-based snapshots. They provide a common standard and touchpoint for all of the physical and virtual servers, desktops and laptops connected to the array, and are easy to implement, operate and manage. And, like NAS, many storage arrays integrate their snapshot technology with Windows VSS, as well as with backup servers and their agents. Some vendors even provide their own agents for non-Windows structured data applications.

The drawbacks include hefty license and maintenance fees, lack of integration with non-Windows-based structured data applications and increasing complexity as the number of storage systems increases.

SNAPSHOTS WITH STORAGE VIRTUALIZATION APPLIANCES

Storage virtualization appliances are primarily SAN based with the exception of F5 Network Inc.'s Acopia ARX, which is file (NFS) based. Other examples of virtualization appliances (or storage systems that incorporate virtualization) include Cloverleaf Communication Inc.'s Intelligent Storage Networking System (iSN), DataCore Software Corp.'s SANsymphony and SANmelody, EMC's Celerra Gateway blades, FalconStor's IPStor, Hewlett-Packard's XP series, Hitachi Data Systems' Universal Storage Platform V/VM, IBM's SAN Volume Controller, LSI Corp.'s StoreAge Storage Virtualization Manager (SVM) and NetApp's V-Series storage controllers.

Storage virtualization approaches to snapshots have the same advantages as storage array- and NAS-based snapshots, but offer others as well. They provide a common standard and point of management for multiple storage systems from a single or several vendors, aggregating them into fewer or just one image. This greatly simplifies snapshot management, operations and training.

The negatives related to storage virtualization-based snapshots are a bit different. These devices will add some transaction latency, even those that have split-path architectures, which ultimately affects app response time. It also complicates troubleshooting and has the potential to exacerbate multivendor finger-pointing. And while the additional hardware or software comes with a price, it may be offset by lower software license or maintenance fees for the virtualized storage.

SNAPSHOTS WITH SERVER VIRTUALIZATION HYPERVISORS

The ascendancy of server virtualization has made hypervisor-based snapshot technology progressively more popular. This technology is available with virtualization software such as Citrix Systems Inc.'s XenServer, Microsoft's Hyper-V, Sun's xVM Ops Center, and VMware's ESX and vSphere4.

The advantages of using hypervisor-based snapshots are straightforward. The technology comes bundled with the hypervisor; it provides the same snapshot methodology for all virtual machines (VMs); it's integrated with Microsoft's VSS; and it's easy to implement, use and manage.

What's not to like about this approach? Snapshots must be managed separately for each hypervisor, and when snapshots are used for any OS other than Windows, only the entire VM will be imaged. That means restores are coarse grain and time consuming, and the snapshots aren't structured-data-aware outside of Windows and may produce non-consistent images.

The main advantage of using SQL database snapshot technology is that snapshots of the database, and any apps based on the database, will be crash consistent.

SNAPSHOTS WITH SQL DATABASES

In SQL databases, snapshotting is called "snapshot isolation." Snapshot isolation is required for databases such as Oracle and PostgreSQL to guarantee that all transactions are serializable and appear to be isolated and serially executed. Other SQL databases also support snapshot isolation but don't require it for serialization. In general, the SQL databases backup features take advantage of snapshot isolation to provide crash consistent dumps of tables.

The main advantage of using SQL database snapshot technology is that snapshots of the database, and any applications based on the database, will be crash consistent.

When it comes to data de-duplication, most companies only offer one kind of solution. But with Quantum, you're in control. Our new DXi7500 offers policy-based de-duplication to let you choose the right de-duplication method for each of your backup jobs. We provide data de-duplication that scales from small sites to the enterprise, all based on a common technology so they can be linked by replication. And our de-duplication solutions integrate easily with tape and encryption to give you everything you need for secure backup and retention. It's this dedication to our customers' range of needs that makes us the smart choice for short-term and long-term data protection. After all, it's your data, and you should get to choose how you protect it.

Find out what Quantum can do for you.
For more information please go to www.quantum.com

Quantum
BACKUP. RECOVERY. ARCHIVE.

But there are some significant disadvantages. The snapshot technology is very limited and it only works with that particular database and the apps tied to it. It doesn't work with the file system, any other application on the server, or with other databases or servers. So you'll need other snapshot technologies or data protection, thus complicating operation and management.

DIFFERENT TYPES OF SNAPSHOTS AND HOW THEY WORK

There are six general types of snapshot technologies:

1. Copy-on-write
2. Redirect-on-write
3. Clone or split-mirror
4. Copy-on-write with background copy
5. Incremental
6. Continuous data protection

COPY-ON-WRITE (COW) SNAPSHOT

COW requires storage capacity to be provisioned for snapshots, and then a snapshot of a volume has to be initiated using the reserved capacity. The COW snapshot stores only the metadata about where the original data is located, but doesn't copy the actual data at the initial creation. This makes snapshot creation virtually instantaneous, with little impact on the system taking the snapshot.

The snapshot then tracks the original volume paying attention to changed blocks as writes are performed. As the blocks change, the original data is copied into the reserved storage capacity set aside for the snapshot prior to the original data being overwritten. The original data blocks snapped are copied just once at the first write request. This process ensures snapshot data is consistent with the exact time the snapshot was taken, and it's why the process is called "copy on write."

Read requests to unchanged data are directed to the original volume. Read requests to changed data are directed to the copied blocks in the snapshot. Each snapshot contains metadata describing the data blocks that have changed since the snapshot was first created.

The major advantage of copy-on-write is that it's incredibly space efficient because the reserved snapshot storage only has to be large enough to capture the data that's changed. But the well-known downside to copy-on-write snapshot is that it will reduce performance on the original volume. That's because write requests to the original volume must wait to complete until the original data is "copied out" to the snapshot. One key aspect of copy-on-write is that each snapshot requires a valid original copy of the data.

REDIRECT-ON-WRITE (ROW) SNAPSHOT

Redirect-on-write is comparable to copy-on-write, but it eliminates the double write performance penalty. ROW also provides storage space-

efficient snapshots like copy-on-write. What allows ROW to eliminate the write performance penalty is that the new writes to the original volume are redirected to the storage provisioned for snapshots. ROW redirection of new writes reduces the number of writes from two to one. So instead of writing one copy of the original data to the storage space plus a copy of the changed data required with COW, ROW writes only the changed data.

With redirect-on-write, the original copy contains the point-in-time snapshot data, and it's the changed data that ends up residing on the snapshot storage. There's some complexity when a snapshot is deleted. The deleted snapshot's data must be copied and made consistent back on the original volume. The complexity goes up exponentially as more snapshots are created, which complicates original data access, snapshot data and original volume data tracking, and snapshot deletion data reconciliation. Serious problems can occur when the original data set (upon which the snapshot is dependent) becomes fragmented.

CLONE OR SPLIT-MIRROR SNAPSHOT

A clone or split-mirror snapshot creates an identical copy of the data. The clone or split-mirror can be of a storage volume, file system or a logical unit number (LUN). The good thing about clones is that they're highly available. The bad thing is that because all of the data has to be copied, it can't be done instantaneously. A clone can be made in-

QUICK GUIDE: SNAPSHOT TECHNOLOGIES

	SNAPSHOT TECHNOLOGY					
	Copy-on-write (COW)	Redirect-on-write (ROW)	Clone or split-mirror	COW with background copy	Incremental	Continuous data protection (CDP)
Snapshot is tightly coupled to original data	Yes	Yes	No	Initially Yes, until background copy completes	Depends on how the original snapshot is generated	No
Space efficient	Yes	Yes	No	No	No	Yes, when compared to multiple point-in-time snapshots
Original data system IO and CPU resource overhead	High	Medium	Low	Low	Low	Low
Write overhead on original data copy	High	None	None	High	High	High
Protects against logical data errors by rolling back or synching back to original copy	Yes	Yes	Yes	Yes	Yes	Yes
Protection against physical media failures of original data copy	No	No	Yes	After background copy completes	Depends on the underlying snapshot technology	Yes

stantaneously available by splitting a pre-existing synchronous volume mirror into two. However, when a split-mirror is used as a clone, the original volume has lost a synchronized mirror.

A very significant downside to this snapshot methodology is that each snapshot requires as much storage capacity as the original data. This can be expensive, especially if more than one snapshot clone is required to be kept live at any given time. One other downside is the impact to system performance because of the overhead of writing synchronously to the mirror copy.

COPY-ON-WRITE WITH BACKGROUND COPY SNAPSHOT

Copy-on-write with background copy takes the COW instantaneous snapshot data and uses a background process to copy that data from its original location to the snapshot storage location. This creates a clone or mirror of the original data.

Copy-on-write with background copy attempts to take the best aspects of copy-on-write while minimizing its downsides. It's often described as a hybrid between COW and cloning.

INCREMENTAL SNAPSHOT

An incremental snapshot tracks changes made to the source data and snapshot data when the snapshot is generated. When an incremental snapshot is generated, the original snapshot data is updated or refreshed. There's a time stamp on the original snapshot data and on each subsequent incremental snapshot. The time stamp provides the capability to roll back to any point-in-time snapshot. Incremental snapshots allow you to get faster snapshots after the first one, and you use only nominally more storage space than the original data. This enables more frequent snapshots and longer retention of snapshots.

The downside to incremental snapshots is that they're dependent on the underlying baseline technology used in the first snapshot (copy-on-write, redirect-on-write, clone/split-mirror or copy-on-write with background copy). If cloned, the first snapshot will take a while; if COW, there will be a performance penalty on writes to the original data, etc.

CONTINUOUS DATA PROTECTION (CDP)

CDP was developed to provide zero data loss recovery point objectives (RPOs) and instantaneous recovery time objectives (RTOs). It's similar to synchronous data mirroring except that it eliminates the rolling disaster (a problem in the primary data is automatically a problem with the mirrored data long before human intervention can stop it) and protects against human errors, malware, accidental deletions and data corruption.

Continuous data protection is like incremental snapshots on steroids. It captures and copies any changes to the original data whenever they

occur and time stamps them. It essentially creates an incremental snapshot for every moment in time, providing very fine-grain recoveries. Some CDP implementations are both time and event based (such as an application upgrade). A good way to think of CDP is as a journal of complete storage snapshots.

CDP is an excellent form of data protection for email, databases and applications that are based on databases. The ability to roll back to any point-in-time makes recoveries simple and fast. FalconStor's IPStor is an example of a storage system and/or virtualization appliance that provides CDP.

With more and more data to protect and often less time to do it, snapshots will play a bigger role in data protection and daily storage operations. Although the differences among snapshot technologies may seem subtle, how they operate in your environment could have a significant effect on the level of protection provided and how quickly recoveries can occur. ⊙

Marc Staimer is president of Dragon Slayer Consulting.

Storage redux: Purchase plans reviving

By Rich Castagna

"It ain't over till it's over," said Yogi Berra. While the persistent economic doldrums aren't over yet, storage managers may finally be seeing a brighter future.

STORAGE MANAGERS may finally be seeing a light at the end of the economic tunnel or, as the old joke goes, maybe it's an oncoming train. But after a train wreck of an economy, it looks like storage budgets may be creeping back to something resembling normalcy. Of course, "normalcy" is a relative term, as those budgets are still hovering in negative territory.

ABOUT THE SURVEY

The *Storage* magazine/SearchStorage.com Purchasing Intentions survey is fielded twice a year; this is the seventh year the survey has been conducted. *Storage* magazine subscribers are invited to participate in the survey, which gathers information related to storage managers' purchasing plans for a variety of storage product categories. This edition had 826 qualified respondents across a broad spectrum of industries, with the average company size measured as having revenue of $1.5 billion.

Last spring when the full impact of the recession was starting to sink in, storage managers were feeling the same sting that hobbled virtually every business. In that earlier edition of our Storage Purchasing Intentions survey, data storage managers reported that their 2009 storage budgets were likely to be 1.9% less than 2008's. While that might not seem like a huge dip, it did represent a nearly five point drop from what they reported in the fall of 2008. And it was the first time in seven years that we saw a negative number when we asked storage managers to compare their budgets year over year. In our most recent survey, conducted in early September, the picture is a little brighter. Compared with 2008, storage budgets are a mere 0.4% lower—not back into positive numbers, but certainly more encouraging than six months ago. Overall, 30% of those surveyed reported decreased budgets, 28% said their budgets would remain flat and almost 33% said their budgets increased.

The spring-to-fall shift in budget levels and planned spending is a common occurrence. Managers appear to be more cautious in the spring when estimating what they'll have to spend for the year; by autumn, many of them see their budgets rise a bit. New projects, unexpected expenses and newfound funds all seem to contribute to this annual swing. In sheer dollars, the average storage budget reported was $2.9 million, which is the same figure as last spring. However, some very large budgets tend to skew the average upward, as 57% of respondents reported budgets less than $1 million.

STILL CAUTIOUS AND COPING WITH CAPACITY

The survey results suggest that, rather than slashing and burning their way through their storage budgets, storage managers are more likely to couple a "trim here, trim there" approach with newer technology implementations to stem costs.

Regardless of budget size, funds are allocated in the familiar proportions that we've seen over the years, with disk system expenses gobbling up the biggest chunk (39%) of the overall budget. One might expect that as disk prices have dropped significantly over that period, the percentage that disk represents would also have proportionately diminished. But the "*x* factor" here is the need for new capacity; so as capacity demands soar, users need to buy more disks.

The need for new capacity hasn't let up, although this survey's numbers suggest that storage managers might get a little needed relief through the

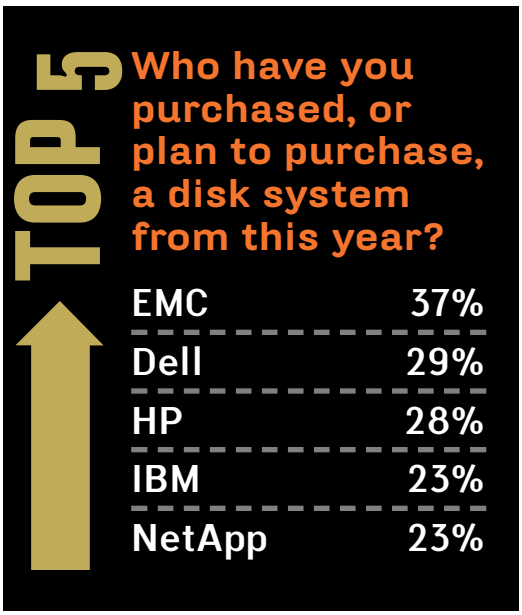

rest of the year. On average, companies will be adding 34 TB of new disk capacity, a fairly hefty figure, but down significantly from last spring's 43 TB (although larger companies are still looking to add an average of 68 TB). It appears that data storage managers have made some adjustments to deal with capacity demands, like employing tools such as thin provisioning, compression and data deduplication to make better use of already installed capacity and perhaps forestall some new purchases. Of course, they may simply have less money to spend, causing some projects to get pushed into next year or beyond.

When storage managers do make disk purchases, they're more likely to try to fill in the capacity of their existing arrays than opt for completely new systems. Thirty-seven percent said their primary disk expenditures would be for new disks for old systems. We've seen this "build out vs. buy new" trend for the past three years; it will be interesting to see if it shifts back to new system purchases as installed gear approaches its end of life or end of lease.

For those planning to buy new storage systems, midrange products will figure into the buying plans of nearly 50% of respondents. Although Fibre Channel (FC) systems are still the top choice, iSCSI continues to make gradual inroads; 12% said they'll buy iSCSI this year, up a couple of points from last spring, while other disk system categories stayed flat or lost a little ground.

iSCSI systems are becoming fixtures in the data centers of all sizes of companies. Among our respondents, the average disk capacity that they have installed is 75 TB, and most of that is on FC SANs (62%) or network-attached storage (NAS) systems (63%). But 35% said they're using iSCSI storage now, up a few points compared with last spring's or last fall's numbers. And that number looks like it will be augmented, as 43% said they plan to or have deployed iSCSI systems this year. That's approximately 3 points higher than last fall and the highest number we've seen to date. The way iSCSI systems are being used also reflects the technology's maturity; 47% said they'll use their iSCSI systems for mission-critical apps, the highest number we've seen to date. For smaller companies, iSCSI has taken on an even greater role, with 61% using it for critical applications (also a new high).

Among the storage vendors respondents purchased from in 2009, EMC Corp. is still top dog, but its lead over Dell Inc. and Hewlett-Packard (HP) Co. has dwindled by a few points; IBM and NetApp Inc. round out the top five. The primary reason for selecting a particular disk vendor is

still the features and functions the product offers (29%) but economic times have had an effect on selection criteria. The second-most-cited selection criteria is that a vendor already supplies other technology to the company (23%); the comfort factor certainly figures in, but it's also likely that a vendor with its foot firmly in the door might cut some special deals for storage purchases. Price, traditionally a lesser selection factor, is up to fourth at 16%, just a point behind tech support.

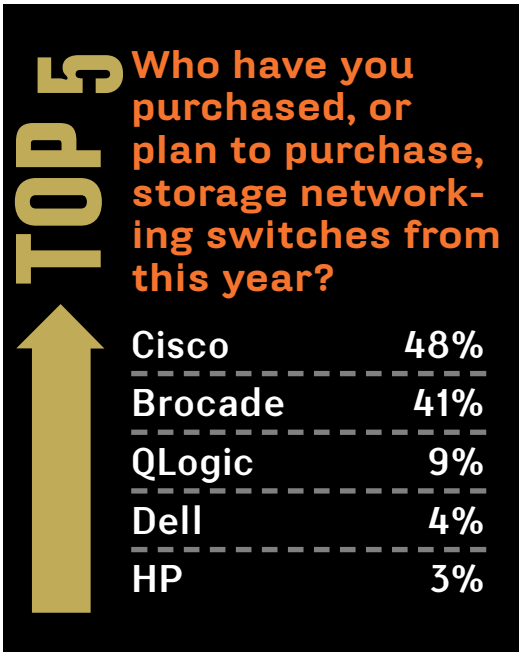

NEW STORAGE TECHNOLOGIES GAIN GROUND

Solid-state storage is the current darling of the storage world, garnering the lion's share of the new tech buzz. For many, however, solid state is still suffering from the "terrible toos"—too new, too untested and too expensive. But our current survey turned up some interesting results: 8% of respondents said they're using solid-state drives (SSDs), with another 3% planning implementations this year. Thirty-five percent are evaluating the technology.

Of those using SSDs, 48% have them in their arrays, 25% are using SSDs as direct-attached storage (DAS) in servers, and approximately 14% report implementations in their arrays and servers. Apparently, the efforts of solid-state drive vendors to counter high price-per-gigabyte claims with more performance-oriented comparisons to hard drives are paying off. And SSD users may also be reaping the benefits of much lower power consumption, although power conservation still isn't top of mind with most data storage managers. Thirty-nine percent of those surveyed said energy efficiency is either the most important criteria or a major factor in choosing an array, essentially the same percentage reported last fall.

Higher-speed storage networking protocols are also nudging SSD and dedupe for the storage spotlight, with 10 Gbps Ethernet (10 GbE) and 8 Gbps FC now available. And while changes to the networking infrastructure are usually painfully slow developments, there's been some noteworthy movement toward these two new protocols.

Currently, the most widely used storage networking protocols are 4 Gbps FC (48%) and 1 Gbps Ethernet (33%). But 8 Gbps FC is used by 14% of respondents, with 13% reporting that they've made the move to 10 GbE.

THE VALUE OF VIRTUALIZATION

Server virtualization is ubiquitous, with 84% of respondents noting that they've virtualized all or some of their servers. And the effects of server

Zenith ARCA

Advanced Recovery and Continuity Appliance

Zenith Infotech Ltd

Compliment or Replacement to Tape Backups

Exchange Recovery

Rapid Failover Server

Bare Metal Restore to Dissimilar Hardware

15 Minute Snapshots

Offsite Replication

It's time for IT departments to have an affordable solution that provides a simple, reliable and flexible way to keep their business running at full speed. That's why Zenith Infotech has created the ARCA.

For IT departments of all sizes, the Zenith ARCA can be a lifesaver if there is a catastrophic server failure. But what about the day-to-day issues that have you scouring disk and tape backups for a deleted file? With the Zenith ARCA, time spent goes from days and hours to minutes and seconds.

Backup • Recover • Restore

See how at www.ZenithARCA.com/OnDemand

Learn more at www.ZenithARCA.com/Learn

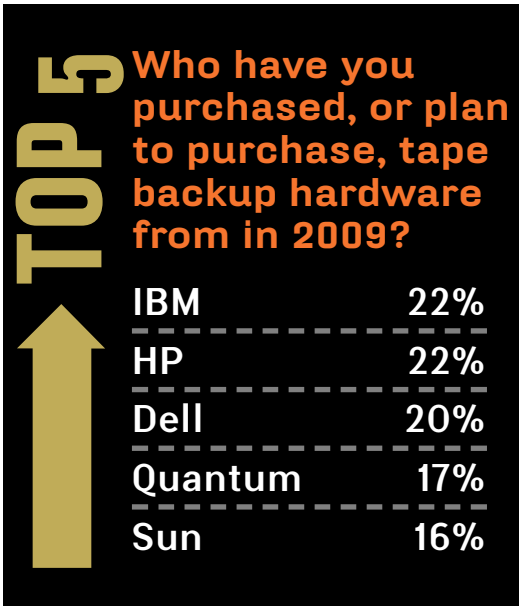

virtualization on storage have become key issues for many of those installations. Forty-nine percent of those surveyed chose Fibre Channel for their virtual server storage, but iSCSI is gaining in popularity. Last fall, only 12% said they hook their virtual servers up to iSCSI storage systems; this time, 18% noted their preference for the technology. In smaller companies, the iSCSI preference is even stronger, with 25% indicating they use it for their virtual servers. (DAS is often dismissed as a viable storage solution for virtual servers, but 11% of those surveyed said DAS is their choice.)

Backup has been the most prominent storage pain point for virtualized server environments. Forty-four percent of our survey takers said they use traditional backup software and methods for their virtual servers, putting agents as required on each virtual machine (VM). VMware Consolidated Backup (VCB), touted as a backup method to save on software licensing costs, is used by only 20% of respondents—a somewhat surprising figure as VCB has been so widely endorsed. VM-specific backup products, like PHD Virtual Technologies' esXpress, Veeam Software's Veeam Backup & Replication and Vizioncore Inc.'s vRanger Pro, have also drawn a lot of attention, but only 6% of those surveyed said they were using products like these in their VM backup operations.

The biggest problems users have run into when backing up virtual servers is backing up too much data, cited by 27% of respondents, and possibly linked to their use of traditional backup methods. Twenty-three percent said the VM backup is just too complicated, while 17% indicated that access to individual files (a frequent issue with VM backups) was their biggest source of frustration.

Although outstripped by server virtualization in terms of implementations, storage virtualization is inexorably gaining favor in storage shops. Among the reasons for the relatively slow uptake for storage virtualization is that it's considerably harder (and more expensive) to do than server virtualization and brings with it the ominous "vendor lock-in" that looms with most virtualization choices.

Still, 31% report virtualizing at least some of their storage, which is up from 26% last fall—a fairly significant jump of 5 percentage points in a short time. Block storage is still the initial candidate of choice for virtualization, with 18% of those surveyed having virtualized all of their block storage and 61% indicating they've virtualized some of their block storage. Approximately 7% of respondents have virtualized all of their file storage, while 44% have virtualized some file storage. These are

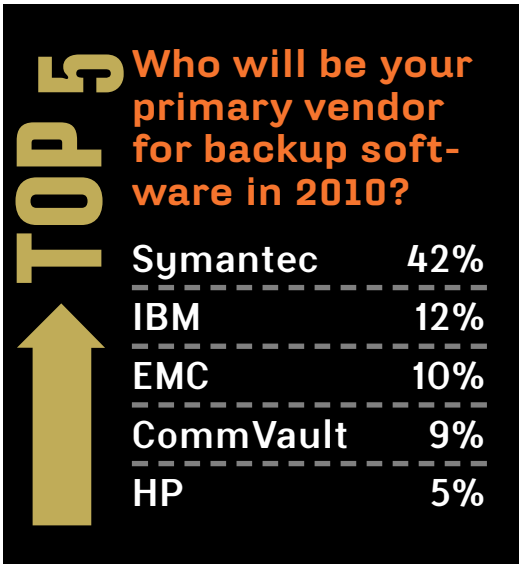

perhaps still modest numbers considering how long virtualization products have been available, but the numbers cited here all represent increases vs. last spring.

MORE EFFICIENT BACKUP OPERATIONS

Space-saving storage technologies, most notably data deduplication and compression, are still key parts of storage managers' efficiency arsenals for coping with increasing capacities and declining budgets.

Each year we see that tape figures less and less prominently in most companies' storage operations, but that trend also illustrates the efforts of storage managers to squeeze their data down to more manageable sizes. Last spring, we saw the lowest number ever (19%) for respondents who said they planned to increase their use of tape and, conversely, the highest percentage yet (29%) for those who plan to decrease their reliance on tape. Those numbers moderated a bit in the current survey, but the trend is clear, although 80% say all or some of their backup data will eventually find its way to tape (down from 86% two years ago).

Even among those respondents who have already bought or plan to buy a tape library in 2009, most are opting for smaller libraries. The average number of slots in the tape libraries they plan to purchase was 109, which is slightly higher than the 101 reported last spring, but still the second lowest number we've seen. But declining slot numbers have a positive side as they point to greater tape library efficiencies, most notably from the use of newer drives. Forty-nine percent of respondents indicated their new libraries would have LTO-4 drives that deliver higher capacity and speed, as well as encryption. Interest in encryption is also starting to translate into actions, with 51% of those surveyed saying they're encrypting at least some backup data, a 7 point gain over last fall. And, once again, encryption ranked second among the newer technologies that respondents plan to implement or evaluate this year (see "Top 10 techs in 2009 implementation/evaluation plans," p. 30).

Of course, disk is de rigueur for backup operations these days, and storage managers will continue to invest in disk-to-disk (D2D) technologies that can improve efficiency. Overall D2D spending appears to be bouncing back from last spring; only 10% said they'll decrease D2D spending (vs. 13% last spring), while 44% plan to increase spending (vs. 34% in the spring). These are positive signs not only for backup vendors, but for beleaguered backup administrators.

Most shops (43%) have integrated disk into their backup systems as a file system-based target or cache; virtual tape libraries (VTLs) are the

BY THE NUMBERS

Change in Storage Budget Compared to Previous Year (%)

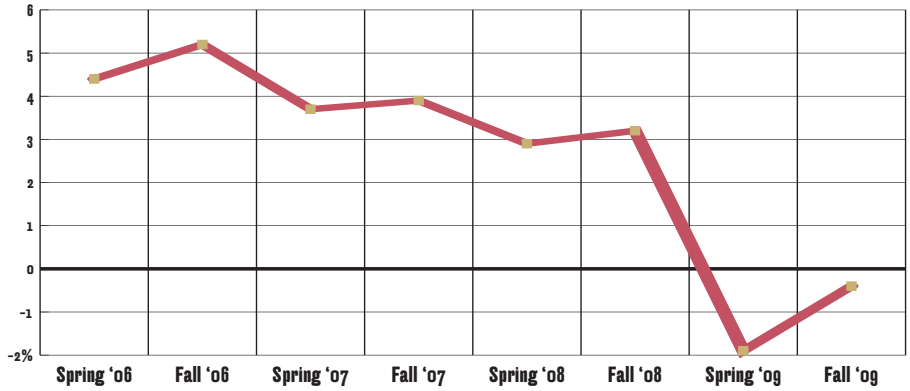

Compared to 2008, storage budgets dropped into negative territory (by -1.9%) in the spring, but have rebounded a bit in this fall's survey, showing only a -0.4% change compared to 2008.

How the Average Storage Budget is Allocated

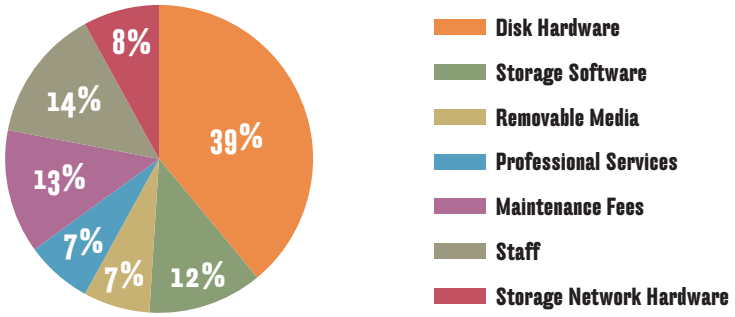

Storage budget allocations have remained virtually unchanged across years of survey data. Because of ever-increasing capacity demands, disk hardware still commands approximately 40% of the budget despite plummeting disk prices.

Average Amount of Disk Capacity to be Added This Year (TB)

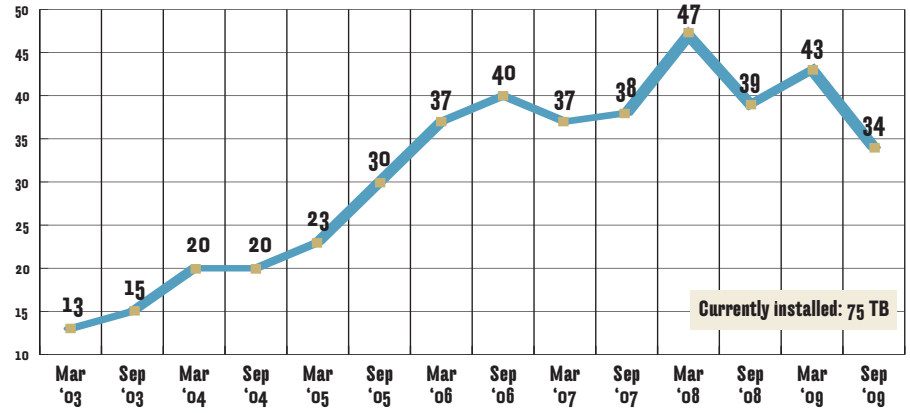

Although the average disk capacity that respondents say they will add this year dipped compared to previous surveys, 34 TB is still a significant amount of new capacity. Larger companies expect to add 68 TB of new disk storage.

BY THE NUMBERS

Solid-State Storage Beginning to Show up in Data Centers

Are you currently using, planning to use or evaluating solid-state storage devices?

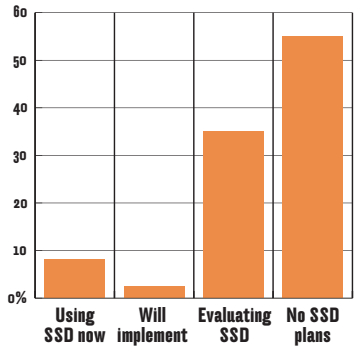

Where have you implemented solid-state storage devices?

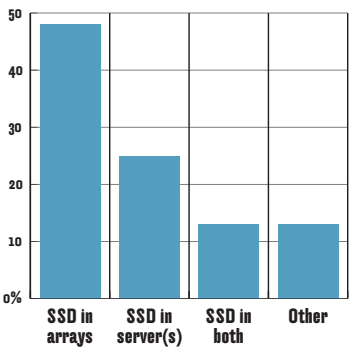

While solid-state storage is still a practical storage solution for a small number of companies, there has been some implementation activity and overall interest is very high.

What's the Main Factor in Your Choice of Primary Storage Vendor?

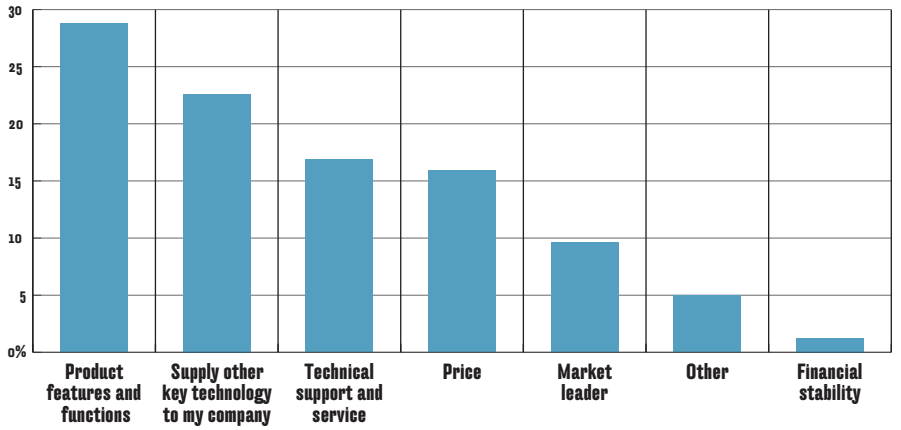

When choosing disk subsystems, the available features and functions have always been the most important factor, but dealing with a familiar vendor that already provides other IT technology has grown in importance, surpassing even price and technical support.

Top 10 Techs in 2009 Implementation/Evaluation Plans

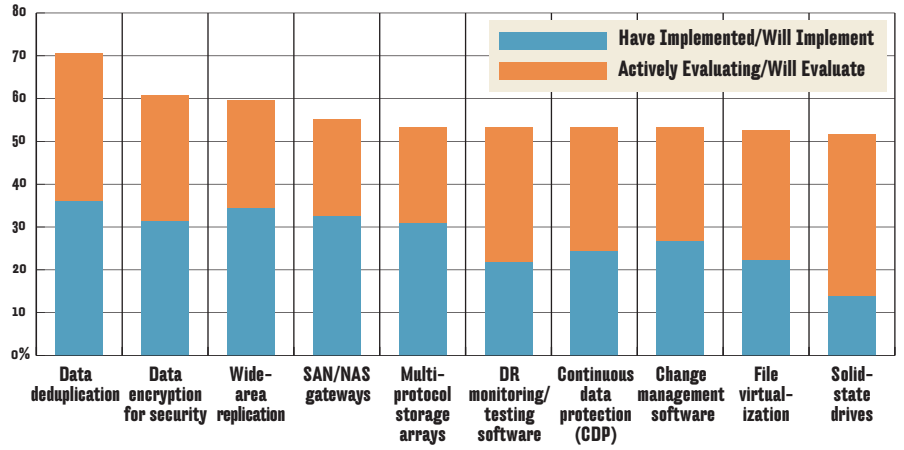

Among newer technologies, data deduplication continues to garner the most attention in terms of implementations and planned evaluations. Encryption, though still not widely used, also figures prominently in many shops' plans.

second most popular approach (26%). In our most recent survey, 35% of respondents said they'll add file system targets to their backup environments, while 24% will opt for VTLs. That figure represents a modest resurgence for that technology.

But the big story is still deduplication. Twenty-one percent of respondents are using dedupe (the highest number we've recorded), and 26% have added it or plan to this year. That, too, is new high for our survey. As far as budgetary considerations go, dedupe is a priority, with 38% planning to increase their dedupe spending (up 8 points from last spring) and only 6% anticipating reduced spending.

Continuous data protection (CDP), another space and time saver, is used by 12% of respondents, just 1 point shy of the highest mark recorded for CDP last spring. Sixteen percent plan to implement CDP this year, also just 1 point lower from the previous high.

The volume around "cloud backup" has been turned up considerably in the past year. While it's still more hype than happening among our respondents, the economy has tweaked interest and purchases. Last spring, we saw a big jump in those looking to trim backup costs by using outsourced or online services for backup when 21% reported using some kind of backup service. In the most recent survey, those numbers have trailed off a bit, with 19% using cloud backup. But use of some specialized backup services—for desktop/laptop files and databases—has actually increased slightly.

BELTS STILL TIGHTENED

For most data storage managers, battling a storage budget that never seems to stretch far enough isn't new, but that annual grind may have helped them weather the most recent economic storms. Storage vendors have also made their contributions with efficiency-oriented products and technologies that seemed to come onto the scene at just the right time. One thing seems certain: Confronted by capacity demands and the need to protect more and more data—with fewer funds to do it all—storage managers have proven their resourcefulness once again. ⊙

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

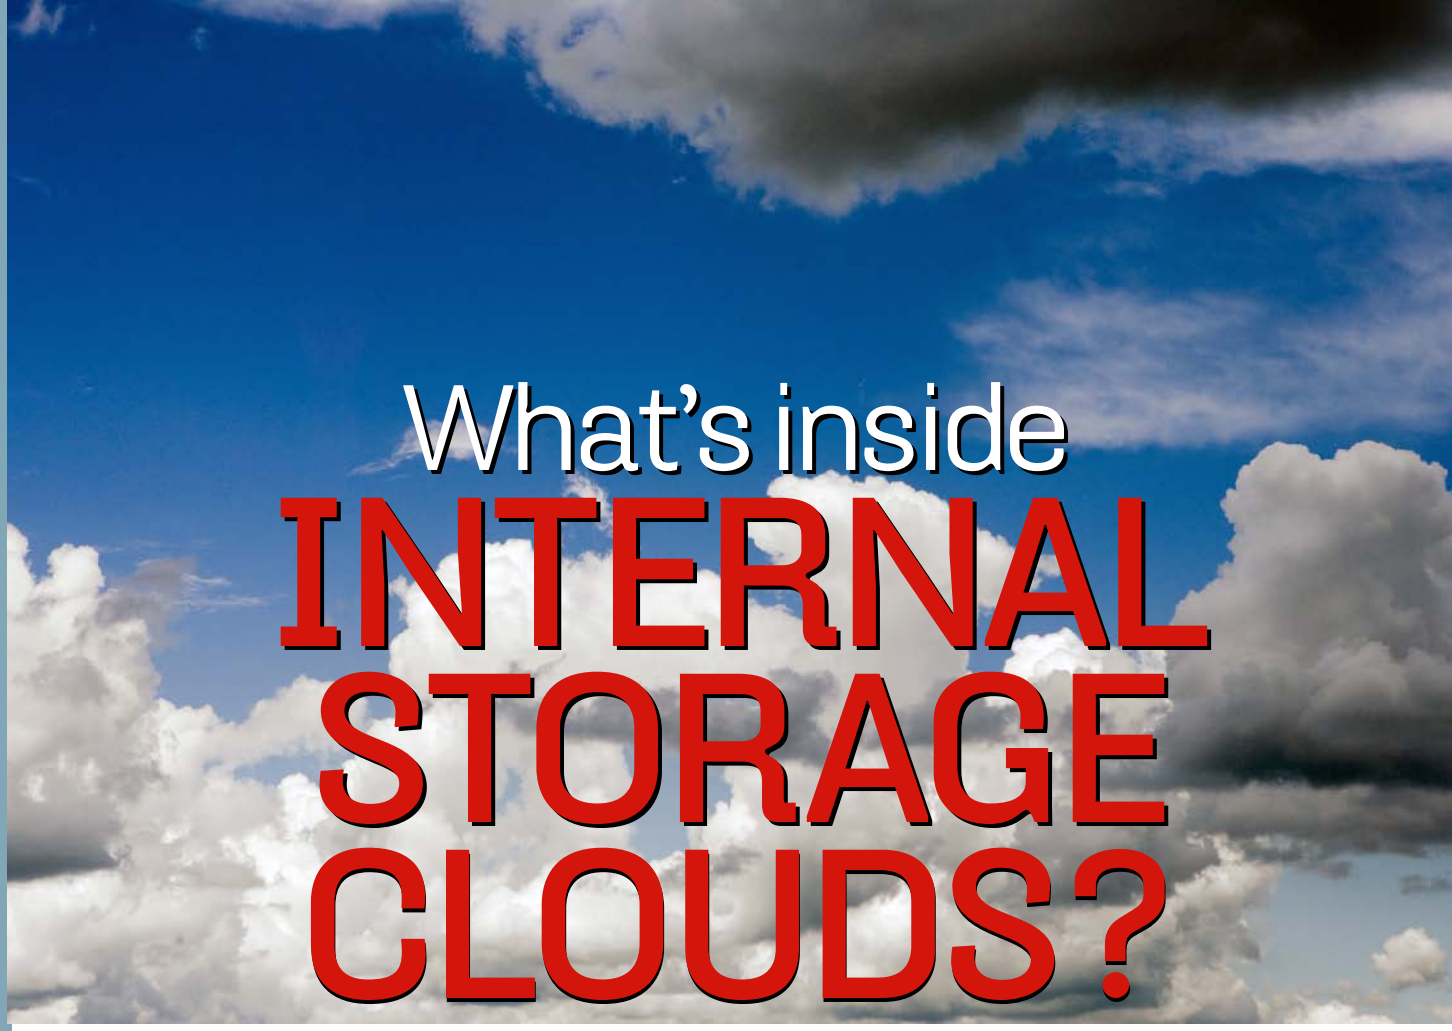

What's inside INTERNAL STORAGE CLOUDS?

Just about every vendor is touting some kind of cloud storage product or service. Here's the lowdown on what constitutes an internal storage cloud.

By Alan Radding

THERE IS NO SUCH THING as a private storage cloud today," declares Stephen Foskett, director of consulting at Nirvanix Inc., a public cloud storage vendor. Maybe so, but that didn't stop the General Services Administration (GSA) from issuing a request for quotes (RFQ) in early August for what appears at first to be a private storage cloud.

But what the GSA considers a private or internal storage cloud may differ considerably from what most enterprises would consider an internal cloud. As noted in the RFQ: "The initial acquisition of these services will be facilitated by GSA through the GSA Cloud Computing Storefront Site—which will enable Government purchasers to buy (using a credit card or other acceptable payment option) Infrastructure as a Service (IaaS) offerings as needed through a common Web Portal, called the Cloud Computing Storefront, which will be managed and maintained by GSA."

Even given that there's no commonly accepted definition for internal storage clouds, the GSA's RFQ seems to describe something completely different. The Feds are asking not for an internal storage cloud or a public storage cloud but for what they label an internal Cloud Computing Storefront, a portal or gateway through which Federal agencies can purchase and access public cloud storage services for their internal use. Even Foskett at Nirvanix, which is preparing a response to the RFQ, seemed puzzled.

The government seems to be on the right track in one regard. However you define internal storage clouds, they promise to reduce storage costs and simplify the storage process. According to the GSA, "Cloud computing has the capability to reduce the cost of IT infrastructure by utilizing commercially available technology that is based on virtualization of servers, databases and applications to allow for capital cost savings ...". The GSA initiative encompasses both storage and compute clouds.

The problem with internal storage clouds isn't that they don't exist, but that there are too many versions of what an internal storage cloud *could* be. "The cloud refers to a layer of abstraction" said Greg Schulz, founder and senior analyst at Stillwater, Minn.-based StorageIO Group. "Almost any storage product can be configured as part of an internal storage cloud. It comes down to your definition. A vendor will define the storage cloud to fit whatever he is selling."

Although there's no widely accepted definition for an internal storage cloud, industry analysts have been identifying the elements needed to create one and explaining how those pieces might be connected. And despite the cloud mystique, "anybody can do this," said John Webster, principal IT advisor at Nashua,

DIFFERENT TYPES OF STORAGE CLOUDS

PUBLIC STORAGE CLOUDS. Services like Amazon's Simple Storage Service (S3) and Nirvanix Inc.'s Storage Delivery Network make massive amounts of file storage available at low cost. Multi-tenancy allows the providers to keep each customer's storage and apps separate and private. Portions of the public storage cloud can be carved out to create what amounts to a private storage cloud.

PRIVATE STORAGE CLOUDS. With a private storage cloud, a company owns or controls the infrastructure and how applications are deployed on it. Private clouds may be deployed in an enterprise data center or at a co-location facility. Private clouds can be built and managed by a company's own IT organization or by a service provider.

INTERNAL STORAGE CLOUDS. This type of storage cloud is similar to a private storage cloud except that it remains inside the organization's firewall. It may be built with the help of consultants or integrators, but it's hosted and maintained by the IT department.

HYBRID STORAGE CLOUDS. A hybrid storage cloud combines attributes of both public and private/internal clouds. It's mainly used to access on-demand, externally provisioned capacity on a temporary basis. The ability to augment a private or internal cloud with capacity from a public cloud can help a company maintain service levels in the face of rapid workload fluctuations or planned workload spikes. Hybrid clouds, however, introduce the complexity of determining how to distribute applications across both a public and private cloud.

SOURCE: Sun Microsystems Inc.

N.H.-based Illuminata Inc. Internal cloud storage isn't brain surgery.

Although internal storage clouds are a rarity today, it's clear what their appeal will be. "This is about performance vs. cost. The internal storage cloud is focused on cost," said Carter George, vice president of products at Ocarina Networks. Conventional storage consisting of sophisticated storage arrays, storage-area networks (SANs), high-performance disk drives, and elaborate backup and recovery, by contrast, focuses on performance and data protection.

But low cost need not be the primary focus, according to Abbott Schindler, an independent storage consultant in Bend, Ore. Cost is top of minds today, Schindler said, because "most start with clouds by thinking about archival storage or data protection so they design it for cheap and slow. There is nothing inherent in the cloud concept, however, that says it cannot be used for transactional data."

ESSENTIAL INTERNAL STORAGE CLOUD COMPONENTS

HERE'S A DO-IT-YOURSELF PARTS LIST FOR BUILDING AN INTERNAL STORAGE CLOUD.

- Global or clustered file (network-attached storage) system (includes virtualization and management capabilities)
- Commodity servers and low-cost storage (SAS DAS, JBOD, PCIe RAID)
- Cost-effective network bandwidth services
- Cloud personality interface depending on application need (NFS, SMTP/POP, HTTP, DICOM, REST, SOAP, XML)
- Object or metadata management layer, including applicable data protection and security tools

SOURCE: Greg Schulz, StorageIO Group

INTERNAL STORAGE CLOUD DEFINED

You could say an internal storage cloud is the same as a public storage cloud—storage delivered as a service over the network—except the components of an internal storage cloud sit behind the firewall. But even that definition isn't completely accurate. A public storage cloud provider, for instance, can reserve a portion of its capacity for the exclusive use of one customer, making it a private storage cloud although it's not internal to the customer (see "Different types of storage clouds," p. 33).

Rather than specifically define the internal storage cloud, industry analysts and consultants prefer to describe its attributes.

For example, the focus clearly is on low cost and easy scalability. "There's a big financial aspect to storage clouds" said Anand Prahlad, CommVault's vice president of product development. "Not only is it

expected to be low cost, but you pay only for what you use." Simply put, internal storage clouds are expected to deliver cheap storage.

And not only cheap storage, but slow as well. Consultants like Schindler, however, don't rule out better storage performance or different service levels as part of the internal storage cloud.

Manageability represents another distinguishing factor. "With an internal storage cloud you want to abstract away the complexity of the storage," said David Allen, chief technology officer (CTO) at i365, A Seagate Company. As a result, the private storage cloud should be easier to manage, enabling a single administrator to handle hundreds of nodes and petabytes of storage. However, the administrator's responsibilities may be limited to a handful of simple tasks.

Finally, how the private storage cloud is accessed can be a key distinguishing factor. HTTP will be the dominant access protocol. "All you want is HTTP or HTTPS connectivity and a Web browser," suggested Ken Satkunam, CTO at SentryBlue in Fargo, N.D.

"A big difference with the internal storage cloud is that it's accessible through an API, not a protocol," Nirvanix's Foskett said. "It will have a programmable API just like a website, maybe use REST over HTTP." Representational State Transfer (REST) is a stateless protocol that includes the state with every communication, the opposite of Fibre Channel (FC). REST provides access to Web services using HTTP; for storage clouds, REST would be used to access storage resources as services.

"A big difference with the internal storage cloud is that it's accessible through an API, not a protocol. It will have a programmable API just like a website, maybe use REST over HTTP."

—STEPHEN FOSKETT, director of consulting, Nirvanix Inc.

In a recent whitepaper, Sun Microsystems Inc. insists on this type of programmability in the storage cloud. "Instead of physically deploying servers, storage, and network resources to support applications, developers specify how the same virtual components are configured and interconnected, including how virtual machine images and application data are stored and retrieved from a storage cloud. They specify how and when components are deployed through an API."

But the industry hasn't standardized on a cloud API, StorageIO Group's Schulz noted, and every cloud provider offers its own. In late July, however, Rackspace Hosting made the API specifications for its public Cloud Servers and Cloud Files open under the Creative Commons 3.0 Attribution license. This might eventually give would-be internal storage cloud builders an open API to get started.

One final characteristic—multi-tenancy—defines the public storage cloud. “Multi-tenancy is an important part of the storage cloud and even the internal storage cloud,” CommVault’s Prahlad said. With internal or private cloud storage, multi-tenancy would let the organization separate departments, projects and workgroups as needed.

So what is an internal storage cloud? The consensus definition appears to be private storage capacity owned or at least controlled by the company, accessible programmatically over an HTTP connection and capable of delivering low cost, highly scalable storage with easily managed multi-tenancy. ParaScale Inc. adds that an internal storage cloud can be small (as few as three to five nodes), and still deliver the economies of cloud storage as well as the ease of management and scaling associated with the cloud.

INTERNAL STORAGE CLOUD OPTIONS

If the internal storage cloud seems familiar, it is. “The storage grid has morphed into the private storage cloud,” consultant Schindler said. Before the storage grid, utility computing packaged computing and storage resources as a metered service. Both concepts are similar, although the technology and architecture is different. “They were all about storage nirvana: accessing the data you want, where and when you want it, and at the cost you want,” Schindler added, and without regard for what the actual storage device was or where it resided on the network.

> "I'm not sure clustered NAS will scale to true storage cloud size."
>
> —ANAND PRAHLAD, vice president of product development, CommVault

The internal storage cloud is also similar to a network-attached storage (NAS) cluster, but with some caveats. “I’m not sure clustered NAS will scale to true storage cloud size,” CommVault’s Prahlad said. Although an internal storage cloud can start small, companies will want it to scale out by adding more devices.

When it comes to internal storage cloud products, the current choices are pretty thin or remarkably wide, depending on how you define the internal storage cloud. For actual products, EMC Corp. offers Atmos, which it describes as an offering for information storage and distribution. With Atmos, EMC stores and replicates a company’s data through its global network depending on the service level you want. It uses business policies, policy-driven automation and metadata to manage a company’s data in this vast storage cloud, and promises operational efficiency, reduced management complexity and cost savings.

AT&T is EMC’s showcase customer for Atmos as a private storage cloud. But AT&T isn’t really using it as a private cloud. Instead, it will offer services involving storage through Atmos to its own customers,

which is more like a public cloud reseller.

Contrary to popular assumptions, there are no giant EMC storage arrays behind Atmos. "That would be way too expensive," Nirvanix's Foskett said. Instead, Atmos' scalable capacity is delivered as JBOD. With Atmos, you get what amounts to a box in your data center with an API and a NAS interface. Or you can use a chunk of the public Atmos storage cloud as a private cloud.

ParaScale offers software specifically for creating and managing an internal storage cloud. Unlike cloud service providers, it sells only the tools that let companies build their own storage clouds. Its software runs on standard x86-based Linux servers and aggregates the direct-attached disks on multiple servers into petabyte-scale file storage in a single namespace.

Beyond Atmos and ParaScale, commercial internal storage cloud products are pretty scarce. "After those, anyone that talks about private cloud isn't really a cloud," Foskett said. Rather, they probably offer storage products that incorporate virtualization at some level, which they're presenting as a cloud. "Often, they're offering their usual product and just sticking the 'cloud' term on it," he added. Similarly, almost any NAS cluster can be presented to look like an internal storage cloud.

BUILDING INTERNAL STORAGE CLOUDS

"DIY is a big thing with internal clouds," consultant Schindler said. Do it yourself is popular because, as Illuminata's Webster noted, it simply isn't that difficult to assemble a private storage cloud (see "Essential internal storage cloud components," p. 34).

There are many ways to design and build a private storage cloud. The simplest may be to "start with a NAS cluster, preferably with a global file system, and put on a cloud [Web] front end," i365's Allen said.

> "DIY is a big thing with internal clouds."
>
> —ABBOTT SCHINDLER, independent storage consultant

The actual storage behind the internal private cloud varies. You probably won't have a storage array as part of the private cloud. "Most will use commodity servers and fill the disk slots with low-cost drives," Nirvanix's Foskett said.

A variation: "Use racks or blade cabinets filled with Linux server blades and disk," SentryBlue's Satkunam said, adding that "the ability to use locally attached disk makes it much less expensive than SAN storage."

The key to building a scalable internal storage cloud "is to start with a lot of little boxes and scale out by adding more boxes," CommVault's Prahlad said. You get data protection through redundancy by replicating the data among the many nodes. To get quality of service, different nodes can have different service performance attributes.

The glue that ties it all together is "a global file system presenting a single name space," Prahlad noted. This may also entail a virtualization and metadata layer.

Management of the internal storage cloud should be simple. "You have to look at websites like Amazon and Facebook for your model. You want whole file storage over HTTP," Ocarina Networks' George explained. For simplicity, limit your file management options to create, read, update, delete and move/copy.

The internal storage cloud doesn't replace an organization's tier 1 storage. Production data continues to run on the high-performance FC SAN or primary iSCSI SAN where it's backed up and protected. Instead, the internal cloud would be used for all the file-based data eating up primary disk space and complicating backup and recovery strategies, as well as for email, archival, media and compliance data. That data is still active, widely used and changed; it needs to be stored and shared but without the expense, performance and service levels associated with tier 1 production storage.

The latest Wave study (January 2009 to May 2009) from New York City-based TheInfoPro asked about interest in clouds in general. "The interest level was light, maybe 12% to 15%," reported Robert Stevenson, TheInfoPro's managing director of storage research. "Most [respondents] had no plans for the cloud." Large companies apparently aren't clamoring for internal storage clouds or cloud computing at this point.

They may, however, already be mimicking internal storage clouds but not realize it as they pop virtualized servers with attached disk onto the network. It's a small step from that to an actual internal storage cloud. ⊙

Alan Radding is a frequent contributor to *Storage* magazine.

NFS 4.1's pNFS: Big NAS performance boost

With pNFS enforcing a standardized approach to parallel file delivery, users will see a NAS performance boost in NFS 4.1, no matter whose NAS storage they have deployed.

ATE LAST YEAR, as Network File System (NFS) 4.1 moved from Last Call to Request for Comment status at the Internet Engineering Task Force, there was a lot of press about how Parallel NFS (pNFS), included in the protocol, would create a quantum leap in network-attached storage (NAS) performance for bandwidth-intensive applications like those used in high-performance computing (HPC) applications. NFS 4.1 represents a major step in the journey to enable NFS to better serve the throughput needs driven by ever-increasing file sizes and demanding HPC environments. While the pNFS buzz has died down, vendors are developing solutions that incorporate the standard and pNFS-supporting products will be on the market in approximately six months. The 600-plus-page Request for Comment is in the editing process and full ratification should happen soon.

NFS 4.1 represents a major step in the journey to enable NFS to better serve the throughput needs driven by ever-increasing file sizes and demanding HPC environments.

THE NETWORK FILE SYSTEM PROTOCOL

The NFS protocol enables users to remotely access and share directories and files stored on a central file server or NAS array as if the data were stored on a local disk drive. By deploying dedicated NAS devices, you gain centralized management of storage resources, easier file sharing and collaboration, better data protection and disaster recovery (DR) planning, storage optimization, space savings via quotas and so on. With the Milford, Mass.-based Enterprise Strategy Group (ESG) projecting file-based data making up more than 70% of total storage capacity by 2012, these benefits are becoming more important to storage administrators.

One of the key challenges with NFS is that performance is gated by the bandwidth of the NAS head or processor node that controls, or "owns," the directory and file being accessed. NFS 4.0 limits file ownership to a single node (there are ways to get around single-node ownership, but not without tradeoffs). When a file is requested by a client, all data delivered to the client must be routed through the NAS head. In the meantime, the

NAS head is also handling NFS tasks such as locking, permissions and file metadata management. One person accessing a large file can bring the performance of the NAS head to its knees, leaving other users with file shares accessed via that head waiting for their files. This issue is exacerbated in HPC environments, which have experienced a shift to parallel processing where multiple processors accessing shared data can easily overwhelm a single NAS head.

A number of vendors have introduced parallel file serving technology to meet this demand, but adoption of these products has been limited thanks to their proprietary natures and the need to add special clients into the mix. Widespread adoption of parallel file services, if it is indeed going to take off, requires a standard approach. This is where pNFS comes in.

pNFS

Parallel NFS takes NAS performance up a level. Files can be broken up and striped across NAS heads and, leveraging multiple data paths and processors, delivered in parallel to the requestor to provide a major performance boost. pNFS also introduces the ability to bypass NAS heads for file delivery altogether.

One of the keys to providing parallel data delivery is the addition of an out-of-band metadata server. The metadata server contains a map, referred to in the NFS 4.1 standard as the "Layout," detailing how and where data is stored. The metadata server also handles file semantics and permissions. When a file request is made by a client, that request is routed to the metadata server first. The metadata server returns information to the client about where the file "lives" on the associated file servers, and then the client can get the information directly. If the file is striped across multiple processor nodes, all of the processor nodes can be leveraged to fill the request, providing a boost in both bandwidth and processing power.

pNFS takes the equation one step further than just parallel data delivery over an IP network by introducing support for direct block data access and Object-based Storage Devices, essentially bypassing NAS heads entirely in the delivery of file data. When file access is requested by an authorized client in block data mode, the actual block layout of the file is returned to the requesting client rather than a file layout. The client can then go directly to the storage devices themselves, rather than NAS heads, to get data leveraging the SCSI protocol. In HPC-type environments, where the clients are often servers in the data center, this means they can be connected directly to block storage devices via fast pipes like 10 GbE or InfiniBand, and access files (as block data) via multiple parallel paths—a huge performance boost vs. accessing shared files over NFS and a single NAS head. A request for an object would follow a similar path.

With pNFS enforcing a standardized approach to parallel file delivery, users will see a NAS performance boost in NFS 4.1, no matter whose NAS storage they have deployed. ⊙

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

Growing need for email archiving

EMAIL, once considered a convenience to augment other forms of communication, is now clearly a mission-critical application in most companies. As such, it requires the special handling and attention—especially for regulatory compliance and legal preparedness—that email archiving provides. In this month's survey, 53% of respondents said they use email archiving, up from 45% last year. Just getting a handle on burgeoning email stores contributed to that swing into the archiving camp, with 30% citing better email management as their primary reason for archiving. In addition, legal readiness (28%) and compliance (26%) loomed large. Nearly half (47%) of those surveyed have opted for third-party archiving apps, but 37% still rely on the email system's capabilities. Outside services are increasingly being considered, with 8% preferring that archiving route, which is twice as many as last year. Of those who have yet to take the email archiving plunge, 52% said it's on their to-do list. With allowable mailbox limits growing—65% allow mailboxes of 500 MB or greater vs. 49% last year—it might be a good idea to move email archiving to the top of your priority list.

—Rich Castagna

Do you archive your company's email?

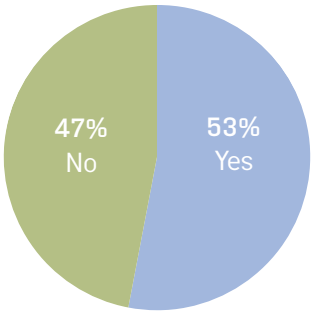

How long do emails remain in users' mailboxes before they're archived?

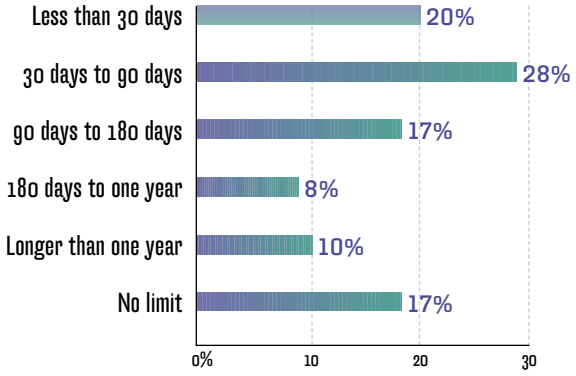

What are your top three email archiving challenges?*

64%	Managing the volume of archived emails
41%	Setting up archiving policies
40%	Setting up retention policies
36%	Recovering archived emails
29%	Addressing the performance impact of the archiving application
28%	Searching for archived email
23%	Meeting compliance requests
14%	Certifying the destruction of old archive data

**Respondents could make three selections*

63

Percent of respondents who archive all email rather than doing it selectively.

"Lots of hidden costs to doing it properly; difficult to explain or justify to a non-IT person."

—Survey respondent

Check out the following resources from our sponsors:

3PAR

3PAR Utility Storage Pays Off Big for Travel Giant Priceline.com

CareerBuilder Positioned for Global Expansion Thanks to 3PAR Utility Storage

i365, A Seagate Company, page 8

Smart Move: Spend Less While Improving Data Protection

12 Best Practices for Data Backup and Recovery

Quantum Corp., page 17

City Where Fantasy Thrives Solves Reality of Backup Challenges with Quantum DXi7500 Express

Case Study: Cutting Backup Times from Eight Hours to Twenty Minutes at U.S. Army Command

Zenith Infotech Ltd., page 26

On Demand ARCA Demos

Learn More About the Zenith ARCA