Managing the information that drives the enterprise

STORAGE

More Choices for Virtual Server Backups

Backing up virtual servers used to be a painful and uncertain process, but new tools and better backup apps make it easier than ever. **PAGE 9**

APPLIANCES LINK YOUR STORAGE TO THE CLOUD

DELL OUTDUELS NETAPP FOR TOP MIDRANGE ARRAYS

CASTAGNA:
DATA, TELL US ALL ABOUT YOURSELF

TOIGO:
NO CURE FOR SUMMER-TIME STORAGE BLUES

McCLURE:
WHO'S MAKING THOSE IT BUYING DECISIONS?

BOLES:
SOFTWARE-DEFINED STORAGE LOOKS LIKE STORAGE VIRTUALIZATION

SNAPSHOT:
STORAGE SETUP FOR VMs GETTING EASIER

TechTarget

The day you get back, challenges will be solutions

Data, define thyself

Protecting burgeoning data stores and securing mobile data is an uphill battle, and it's one that we're bound to lose using conventional methods.

THERE'S A LOT of cool stuff going on with backup these days, with a lot of interesting technical innovation making backup easier and more flexible than ever. Yet despite these new developments in data protection, I can't help but wonder if the efforts are somehow missing the mark.

Key data protection methods like deduplication and snapshot management continue to develop, while other techs—like continuous data protection (CDP)—are getting dusted off, spiffed up and finding new niches. The idea of backup-less backups is starting to take hold, but most companies are still looking for ways to streamline or simply enable more traditional approaches to data protection, like daily incremental backups plus weekly fulls.

We need to figure out ways to squeeze a ton of backup into a 10- to 12-hour window or even eschew windows altogether in favor of CDP methods that do backup in dribs and drabs instead of a torrent of data at day's end.

So coming up with new ways to protect ever-growing data stores is a good thing, right? Well, yes and no.

Most data management revolves around the age of the data, if it happens at all. New stuff is backed up, older stuff gets archived and so on. It's a process based on a single dimension of the data and it's purely reactionary. Granted, it's worked for years and still works at plenty of companies, but a changing IT environment is applying new pressures to old tools. It used to be that a storage manager's job was to give data a place to live and ensure it was protected by tucking away copies in case the originals got lost or damaged. But with smartphones, tablets and all manner of bring-your-own-device (BYOD) things invading

the enterprise, it's not enough to house and protect data—you have to make sure no one lets it wander away.

So while the tools are barely adequate to maintain the status quo, they're definitely going to fall short in the future; maybe they're already starting to break down at your company.

I think that's because our typical approach to data management is quickly becoming obsolete. Most of what IT has done and continues to do assumes it has control over the data, and that the amount of data is controllable. Neither case is true anymore.

Dealing with data using software or hardware tools is essentially a process that's external to the data. The intelligence in the process doesn't come from the data, but from the tools used to move, copy and otherwise deal with data. And today, tools aren't all that smart. So the data itself needs to be smarter; it needs to know what it's supposed to do, how long it should do it, and where it can and can't go.

It's time for metadata on steroids—mega-metadata or whatever you want to call it—that will add enough intelligence to data so that we humans using other relatively crude tools don't have to worry about it.

This mega-metadata would be self-describing and able to trigger autonomic actions, such as delete, move/don't move and copy. The goal would be to create data that is smart enough to take care of itself after a little initial guidance, and would know what would keep it safe and

what would get it into trouble.

We're already part of the way there, with technologies such as the expanded metadata enabled by object storages, and the application of Active Directory or LDAP to data management. On the BYOD front, being able to remotely wipe a phone or tablet is an effective data se-

> **If all the necessary disposition information was packed in with the data, the data itself would know what to do.**

curity response although it requires too much human intervention.

If all the necessary disposition information was packed in with the data, the data itself would know what to do. Of course, it would need cooperative operating systems, file systems and applications.

For example, a file might be tagged to remain "live" until a certain date. It would also be able to let a backup or replication application know how many copies to make, be able to tell an iPhone user that it couldn't be copied or sent to a sync-and-share cloud, and it would let the archiver know when it was time to retire or simply shuffle off this mortal coil.

My mega-metadata fantasy continues with an "ultimate dashboard" that displays daily alerts describing what files are scheduled for self-destruction, which ones will be moved near-line or off-line, and which files have been moved to different user devices. It would be a kind of "Where's Waldo?" for data, but you wouldn't have to search to find all the instances of Waldo.

A lot of those things happen now, but you have to pull the levers and push the buttons on a myriad of tools. And just keeping track of everything is a full-time job.

There are probably a bunch of products out there that I'm just not aware of that do some of the things I've described. But for this vision to become reality, applications have to get smarter, and OSes and file systems have to become more aware.

If data is the most important thing—and I can't imagine anyone would dispute that—then we have to focus more on the data itself if we're ever going to protect and secure it in an information-crazy, ever-mobile world. ∎

RICH CASTAGNA *is editorial director of TechTarget's Storage Media Group.*

There ain't no cure for the summertime storage blues

LTFS is heating up, clouds are still rather cool and the forecast may be changing for solid-state.

HAD THE EXTRAORDINARY pleasure of dining with real IT people in several cities over the past few weeks as TechTarget's "The New Rules of Backup and Data Protection" seminar winds its way around the country. It seems that summer's extreme weather always increases interest in disaster recovery and data protection.

At these casual meetings, there are always the expected inquiries about the latest "shiny new thing" in data storage technologies—flash storage—but I'm sensing less enthusiasm about adopting the technology on anything more than a one or two PCI Express card basis than I might have expected. Some analysts might say flash in server has peaked, but that isn't what I'm seeing. People seem desperate to do anything to speed up the doggedly slow performance of their server hypervisor and virtual machine complex even if I/O isn't the problem.

I also had the opportunity to interview Erik Eyberg, who came to IBM with the Texas Memory Systems acquisition and serves as technology savant and evangelist at Big Blue. I recorded the interview (and stuck a two-part video on my blog for anyone who's interested) because I was amazed at the man's candor. No, he admitted, flash memory probably would do nothing to speed up VMware or any other application that was processor or network bound (conditions you can readily check with performance monitors available in most OSes). Flash only works its magic when I/O binding is the issue, signaled by overly long queues in storage performance monitors.

Another highlight of the Eyberg interview was his desire to switch the narrative about flash from one of speeds and feeds (you know, how the technology even makes grits cook faster), to one of latency. Latency is a bit

nuanced in his usage, referring to how long it takes for a transaction to complete. From a business standpoint, this is all that matters and not the 18 million IOPS achievable with an all-flash array under test conditions. I quite agree.

I've also found in my conversations with the plain folk of IT that the blush just doesn't quite seem to be reaching the rose of cloud storage. I receive considerable plaudits for my skepticism on this subject with only some occasional criticism. In the latter category, only one fellow, a vendor, has actually gotten "up in my grill" about my questioning the solvency of public clouds. He said that "real analysts like IDC and Gartner," not to mention "real technology leaders" like Joe Tucci of EMC and John Chambers of Cisco Systems, would "laugh me out of the room" for my doubts about the future of clouds.

Less caustic was a cloud evangelist who listened closely to my questions, nodded approvingly at my concerns about security, service-level agreements and so on, and then responded with the observation that "for clouds to succeed, the culture of IT and business must change." While I appreciated her tone (no shouting or growling), her view struck me as kind of out there: If we could change the culture so that everyone felt comfortable with clouds, then we could work out all the technical hurdles that still remain. That struck me as sort of like saying that if we could get everyone to drive a car, we could work out the problems with all those dials, gauges, brakes and even the safety gear in due course.

One topic that I did see gaining a lot of interest was the Linear Tape File System (LTFS) and how it can be used to build huge-capacity storage for infrequently accessed files. Surprisingly, my seminar presentation was the first that most attendees had even heard of the data storage technology. I was able to bring them up to date quickly, not only with the basic concept, but with the

> **You can use LTFS as part of a file-system migration and archive strategy without needing to leave behind clumsy pointers or stubs. It could become a low-cost way to store historical data you access later for big data analysis.**

Enterprise Edition software that debuted last month at IBM's Edge conference in Las Vegas. The key talking point around the latest version is its support for IBM's General Parallel File System, which allows the system of files stored to LTFS tape to be included in a hugely scalable, clustered and cross-storage file namespace. In short, you can use LTFS as part of a file-system migration and archive strategy without needing to leave behind clumsy pointers or stubs. Moreover, LTFS tape may well become

a low-cost way to store historical data you may wish to access later for big data analysis.

To most of the folks I talked to, however, the greatest perceived value of LTFS was its potential use to store a ton of files that are never accessed, clutter up the storage junk drawer and that everyone is afraid to throw away. No one I spoke to trusts clouds enough to put their old files there, owing mainly to cost, accessibility and privacy concerns. LTFS tape seems like a great way to leverage existing investments in LTO tape libraries to slow the rate of disk growth and its associated cost. Despite the green shoots of economic improvement, no one seems to want to rush out and buy an all-flash array or cloud service, despite what analysts and vendor/innovators are saying. ∎

JON WILLIAM TOIGO is a 30-year IT veteran, CEO and managing principal of Toigo Partners International, and chairman of the Data Management Institute.

THE STATE OF VIRTUAL SERVER BACKUP

Backing up virtual servers was once a kludgy and network-choking process, but backup applications have evolved to handle the special needs of virtualized servers. Here's what you should look for in a backup app.

By George Crump

SERVER VIRTUALIZATION IS without question one of the most significant technologies introduced into the data center in the last five years. It has changed almost every aspect of how architectures are designed, including networks, storage and the servers themselves. Data protection is one of the key operations that has been most impacted by the shift to a virtual environment. Gaps in data protection for virtualized infrastructures led to the origin of startup vendors that focused solely on providing virtual machine-specific backup and recovery solutions.

THE IMPACT OF VIRTUALIZATION ON BACKUPS

Prior to virtualization, applications ran on dedicated servers with access to all the resources (storage, memory, CPU, network) available to that server. When a backup was triggered for that application it could, for the most part, use all the available resources to complete

the task at hand to copy data from the server to a backup destination.

Virtualization changed things. Resources are now shared across multiple virtual machines (VMs) each running an application of their own. If the backup process doesn't adjust to this new reality, then all the VMs could start sending all their data at the same time—all from a single server. That could lead to a potential server crash as the hypervisor runs out of memory resources, or at least produce mediocre performance as it runs out of CPU and networking resources.

EARLY ATTEMPTS TO FIX VM BACKUP

In the "early" days of VM backup, most data centers protected VMs as if they were standalone servers, and administrators would balance backup schedules so that only one or two VMs were backed up at the same time. This meant IT managers could continue to use their legacy backup applications. But as virtualization continued to grow and VM density increased, the scheduling balancing act became untenable and alternatives were needed.

THE VIRTUALIZATION BACKUP ADVANTAGE

Despite the negative impact of virtualization on data protection performance, it did bring its own set of advantages. A "server" was now encapsulated as a single large

file instead of thousands or potentially millions of small files. And that file is accessible by multiple servers via the virtualized cluster put in place to enable features like live migration of VMs between hosts and automated resource balancing.

That added up to fairly easy access by an alternative server to back up the "file" (the server). In addition, most hypervisors had snapshot capabilities built into their clustered file systems so they could be snapshotted and protected by the alternative server without impacting the primary host server's resources and performance. Essentially, the capability for off-host backup was born.

This led to the rise of companies like Nakivo Inc., PHD Virtual Technologies, Veeam Software and Vizioncore Inc. (bought by Quest and then acquired by Dell). They leveraged the above capabilities and expanded them to include granular recoveries from virtual server systems.

In the early days of virtual server backups there were a limited number of ways in which the backup software could interface with the hypervisor to perform the task at hand. As a result, when hypervisors were altered or upgraded, compatibility issues with the backup application sometimes arose. While this was an acceptable risk for smaller backup vendors, larger enterprise software vendors were more conservative in providing VM-specific backup capabilities. With traditional backup apps lumbering along, startups were able to capture an early lead in

VMware data protection.

Today, hypervisor vendors are providing API sets that backup software companies can leverage as part of their code bases. In theory, at least, this means their backup applications should work despite revisions to the hypervisor code, as the amount of backup application code rewriting should be minimized.

VM BACKUP TODAY

With the availability of an API set, most vendors, whether legacy or VM-specific, can provide off-host VM backup, something that should now be considered a basic requirement for VM data protection. But there are specific features beyond off-host backup that IT planners should consider.

● **Agent versus agent-free backup.** An agent is software that's installed in the VM to assist in the backup process. Even though the above-mentioned APIs allow for off-host backups, some vendors still rely on agents installed in VMs. The agents may be used to help with application-awareness (allowing for granular backup and recovery of databases or email stores) and, in some cases, may accelerate raw backup performance.

While agent-free backup offerings don't install code into the virtual machine, granular recovery of application data is still available; however, the backed up VM image may have to be mounted as a separate VM and have that data copied out of it. Some agent-free backup products have developed "helper" applications that will allow for scanning, searching and extracting granular data components from well-known data types such as Microsoft Exchange, SQL Server and Oracle without having to mount the VM image.

● **Changed block backup.** Hypervisor APIs have increasingly added capabilities such as VMware's Changed Block Tracking (CBT) that allows the backup software to understand which parts of a virtual machine's image file have changed since the last backup. This is a key feature that allows backups to occur more frequently since the amount of data transferred is minimized and should result in reduced data loss in the event of VM corruption.

● **Enhanced restores.** Restores have also been significantly improved in virtualized environments. First, instead of having to restore the entire VM image by tapping into the hypervisors' API, most off-host backups can now recover a single file or set of files when a recovery is needed. Some vendors also leverage CBT to provide changed block restores. For example, if a large database is corrupted, a changed block recovery would just restore the parts of the database that changed since the last backup was made.

Restores can be further enhanced in products that

allow for the execution of the VM directly from the recovery device, often called "in-place recovery." With an in-place recovery scenario, no data needs to be transferred across the network and the VM and its data can be returned to operation in a matter of minutes. For many organizations, this capability combined with hourly CBT backups can eliminate the need for separate business continuity software.

Some vendors are extending this capability to the cloud. Where the "in-place" part of the recovery actually occurs is in the remote data center. In those architectures, data is typically backed up locally, then replicated to the cloud and placed in position for recovery in the event of a site disaster. This not only solves the local protection and availability issues, but also provides DR readiness.

There is a tradeoff between in-place recovery and changed block restores. With in-place recovery there will come a time when the VM needs to be moved back to its primary storage destination. Also, it's unlikely the backup device has the performance and redundancy of the primary storage device, something that's especially true with the cloud recovery model described earlier. CBT recovery, on the other hand, incurs downtime up front, but eliminates the more prolonged downtime required to move the whole VM into place. Ideally, IT planners should look for a product that offers both methods.

● **Tape support.** It may seem surprising that tape support is working its way into VM-specific applications that were originally disk-only, but tape is inexpensive, portable and ideal for long-term storage of VMs. It's a perfect addition to the rapid backup and restore capabilities of disk because it allows the disk investment to stay small and be used for the most immediate of recoveries.

The changing role of the backup disk

THANKS TO FEATURES like Changed Block Tracking, cloud-based recovery and in-place recovery, the design of disk backup devices needs to evolve. In the past, data transferred to the disk backup appliance was bandwidth-focused (large files, a lot of data all at once); now, it's much more random in nature (small block changes transferred throughout the day). In addition, since virtual machines (VMs) may now be executed directly from the device, the performance of the disk backup appliance matters. We may soon see disk backup appliances that have some solid-state storage installed for the execution of VMs. ∎

Tape support should be given strong consideration, even in disk-only environments. The long-term storage capacity savings plus the ability to "overnight" a VM on tape can pay big dividends.

● **Physical server support.** A major differentiator among backup applications is their ability to back up physical servers. Many of the new VM-specific backup apps are VM only. While many data centers are striving for 100% server virtualization, the majority of them aren't even close. That means that if a VM-specific application is selected, you must be prepared to deal with at least two separate backup and recovery processes.

Most legacy enterprise solutions support physical and virtual server protection, but tend to be behind in some of the VM-specific features described earlier. You may need to choose between the luxury of a single data protection product or running two products to capture best-of-breed functionality. In general, the choice comes down to how much mission-critical data resides on physical systems.

VM BACKUP BOTTOM LINE

The state of virtual server backups has improved significantly over the past few years, thanks in large part to vendors like VMware that established a robust API set that allowed for innovation and integration. The capabilities aren't only improving backup, they're helping to eliminate the need for separate business continuity and disaster recovery applications. ∎

GEORGE CRUMP *is president of Storage Switzerland, an IT analyst firm focused on storage and virtualization.*

APPLIANCES LINK ON-PREMISES STORAGE TO THE CLOUD

Cloud-integrated storage appliances allow hybrid storage configurations that seamlessly link data center storage with cost-effective, scalable cloud storage.

THE WORD "HYBRID" has a variety of definitions when it's used to refer to cloud storage or computing. For our purposes, we'll define hybrid cloud storage as storage that transparently and effectively integrates on-premises storage and in-the-cloud storage to create a greater overall value. That means hybrid cloud storage must deliver increased value in one or more of these dimensions: cost reduction, scalability, manageability, performance, data protection, business continuity, degree of automation and security. And it must do so by integrating transparently with—and without altering—the existing storage on-premises infrastructure. Lastly, it must not require any changes to the applications.

While the benefits of cloud storage are well understood and appealing, adoption has been somewhat slower than expected because cloud storage stores and accesses data based on the Representational State Transfer (REST) protocol. Some early gateways translated iSCSI and CIFS

By Arun Taneja

protocols to REST protocols to make it possible to load and extract data to/from the cloud. It was a step in the right direction, but not enough to significantly boost the cloud storage market.

That changed when cloud-integrated storage (CIS) appliances appeared a few years ago. These appliances enable hybrid cloud storage, allowing true integration of on-premises and in-the-cloud storage with 100% transparency to existing storage environments. There are four use cases that illustrate what's possible with these state-of-the-art systems:

- Backup
- Disaster recovery
- Archiving
- Tier-two primary storage (both NAS and SAN)

ANATOMY OF A CLOUD-INTEGRATED STORAGE APPLIANCE

A CIS appliance presents an iSCSI, CIFS or NFS interface to on-premises hosts, which effectively removes the main issue associated with cloud storage: the use of the REST protocol. A typical appliance has a suite of technologies, including compute, caching, tiering, deduplication, compression, encryption, thin provisioning, WAN optimization, replication, data protection, protocol conversion, snapshots and cloning. Most appliances come with solid-state and/or hard disk drives built in. The back end of

the appliance connects to the cloud via the Internet and speaks REST protocol.

When a write request arrives, it's written to cache and acknowledged; the data is then deduped at the block level, or it's combined with other contiguous blocks and converted into a chunk. Some appliances also apply compression algorithms. The chunks are then stored on local drives but also optimized for transmission across the WAN to the cloud. The appliance does the protocol conversion and only sends the chunks that are unique and not already stored in the cloud.

Snapshots can be taken periodically and, depending on the appliance, may be immediately transferred to the cloud, with or without retaining a local copy. Metadata maps containing pointers that describe the makeup of chunks are stored in the cloud (and in the appliance in some cases) along with the chunks. Some appliances only use their local storage for caching while others also use it for tiering. Cloud service support varies, but AT&T Synaptic, Amazon Web Services (AWS), Google, IBM SoftLayer, Microsoft Azure, Nirvanix and Rackspace are broadly supported.

Users don't need to know any details about how the data is stored in the cloud, how it's protected or how it's managed. The interaction with the appliance is exactly as if it were an iSCSI volume or a CIFS or NFS share. CIS products may include a physical or virtual appliance, with both often available.

USE CASE 1: BACKUP AND RESTORE

A cloud-integrated storage appliance can fundamentally transform backup and restore, and eliminate most headaches associated with traditional data protection. One implementation works in conjunction with the backup software currently in use, including all the major backup apps. Alternatively, a CIS appliance can eliminate the need for traditional backup software completely, which can significantly reduce costs and simplify the overall data protection environment.

● **CIS appliance with backup software.** In this scenario the CIS appliance is seen by the media server as an iSCSI (or CIFS) target. All traditional backup management happens as usual, but the backup streams are deduped, compressed, WAN optimized and encrypted by the appliance before being sent to the cloud. Such an appliance can be very cost effective versus buying a backup appliance such as EMC/Data Domain, HP StoreOnce or IBM ProtecTIER. In addition to getting full data protection in the cloud, customers also get geographically remote storage for their data. That represents substantial savings over buying another set of backup appliances, locating them at a remote site and replicating to them. And since only unique chunks are sent to the cloud, storage costs are further minimized. Depending on the cloud-integrated storage product, a local copy of the latest backup may reside on the appliance, allowing for fast recoveries. Backup

management is still performed using the backup software, and data must be recovered as usual before it can be made available on primary storage. Riverbed Whitewater, StorSimple (Microsoft) CiS and TwinStrata CloudArray are examples of appliances built for this use case.

● **CIS appliance without backup software.** This functionality is enabled when the CIS appliance is used as primary storage (more details later). The CIS appliance may be used to improve data protection, with periodically scheduled snapshots integrated with cloud storage. A local copy of the snapshot may be kept in the appliance for faster recoveries. Hundreds of snapshots can be stored in the cloud since only changed data is kept, and the concept of full and incremental backups disappears. In this case, the recovery happens via the CIS user interface since no backup software is involved. Both file-level and volume-level recoveries are possible; and if the application server is running virtual machines (VMs), the recovery can be at the VM, volume or file level. Because the data is never reformatted by backup software, snapshots can be mounted immediately without having to recover first.

USE CASE 2: DISASTER RECOVERY

Probably the most novel use case for a CIS appliance is disaster recovery (DR). A CIS appliance has the potential to deliver a cost-effective, testable "on-demand" DR offering

for all applications. By uploading cloud-integrated snapshots into the cloud, when the local site is disabled a new CIS appliance can be fired up at another site and connected to the cloud-based snapshot. In the most sophisticated appliances, the application can be started without performing a recovery, and no data needs to be down-

> **The burden of keeping the data safe and with integrity shifts to the cloud provider, who performs technology refreshes and migrates data to newer technologies when necessary while maintaining 100% data integrity.**

loaded into the appliance. The recovery time objective (RTO) is determined by the amount of time it takes to download the metadata map that describes the contents of the snapshot—a trivial process compared to downloading the entire volume of data. Some performance may be sacrificed, but the application will be up and running with a very short RTO.

A huge benefit is that the company doesn't require a second site for DR purposes. The StorSimple CiS appliance from Microsoft and TwinStrata's [CloudArray](#) are probably the best examples for this use case. Both of these appliances are also available as VMs so users can choose

to fire up the application(s) in the cloud, assuming compute capability is offered by the cloud vendor, such as Microsoft Azure and AWS Elastic Compute Cloud (EC2).

This use case is also ideal for protecting data at remote offices. A CIS appliance installed at each remote office can transfer snapshots to the cloud; a disaster at a site can be handled by recovering at the data center or at one of the other remote sites. This flexibility enables simple, inexpensive DR for remote offices.

USE CASE 3: ARCHIVING

The archiving use case is similar to backup/restore, except that cold data is literally separated from active data and moved from [primary storage to the cloud](#). This relieves pressure on primary storage, improves application performance and delays new purchases, but all archived data is still available online when needed. Archive data may be stored for very long periods of time and must not lose its integrity. The burden of keeping the data safe and with integrity shifts to the cloud provider. The provider performs technology refreshes and migrates data to newer technologies when necessary—all while maintaining 100% data integrity and without involving the customer.

When a CIS appliance is used as primary storage, the implications for archiving are even greater. Unlike traditional on-premises primary storage where cold data

occupies expensive real estate, a cloud-integrated storage appliance only keeps active data in the appliance, with cold data getting moved to the cloud. All data is always available online, and IT doesn't have to worry about periodic housecleaning of primary storage.

Applications can be run against this archived data in the cloud for best performance. But not all the bells and whistles, such as the ability to place legal holds against specific data, perform audits for compliance and so on are currently available. Still, cloud-based archiving solutions can be very cost effective for most companies not in highly regulated industries.

USE CASE 4: PRIMARY STORAGE

The use of cloud storage as primary storage is the most challenging implementation because of latency issues. Most latency-sensitive applications can't deal with the tens of milliseconds (or more) of delay typically associated with accessing cloud storage. However, with a well-architected CIS appliance it's possible to enable excellent performance for all but the most latency-sensitive critical applications. For example, Microsoft's StorSimple appliance offers excellent performance for its SQL Server, Exchange and SharePoint applications. These applications require low latency, which is achieved by a judicious use

Three ways hybrid cloud storage can improve disaster recovery

- Recovery point objective (RPO) and recovery time objective (RTO) matter in disaster recovery (DR). RPO is controlled by the frequency of snapshots taken; most cloud-integrated storage (CIS) appliances support snapshotting. However, not all CIS appliances deliver the same RTO. For the fastest RTO, look for products where applications can be started without requiring all the data to be recovered first.

- The ability to start higher-priority applications first, without having to treat all application recovery alike, is important. Look for offerings that allow this.

- In a DR scenario, if an application is restarted at another site, make sure you start protecting it right away with new cloud snapshots. That way there won't be a lapse in protection. ∎

of the appliance's storage for caching. This scenario provides scalable, on-demand storage for applications, effortless provisioning and a greatly reduced on-premises storage footprint. The bulk of the storage is delivered online from the cloud, yet it appears and behaves exactly like local storage. Data is also protected in the cloud by the cloud provider, further relieving IT of some data management tasks.

To be sure, using cloud storage for primary storage can't replace on-site tier-one storage, at least at this stage of hybrid cloud storage development. But tier-two storage generally makes up the largest portion of storage in most companies, and that tier is certainly a candidate to be moved to a hybrid cloud storage architecture. Not all CIS appliances do a good job at delivering reasonable performance for primary storage, however, so you must weigh your options carefully.

PRODUCT SAMPLER: HYBRID CLOUD STORAGE APPLIANCES

For iSCSI block storage StorSimple is worthy of consideration, and for file Nasuni has a solid offering. TwinStrata's CloudArray is primarily intended for use for backup, archiving and DR, but it can also be considered for primary storage use. Amazon offers a gateway product, but it's mostly designed to move on-premises storage to AWS Simple Storage Service. Nirvanix's CloudNAS is an

appliance to move files back and forth between on-premises NAS and the Nirvanix cloud. Nirvanix also offers customers the ability to actually create a Nirvanix cloud on the premises using Nirvanix technology, and to integrate it with Nirvanix's public cloud for what might be considered a homogeneous hybrid cloud. It uses its hNode appliance as the interface between the two. A number of other variations will become available as the concept of hybrid cloud storage gains traction.

HYBRID CLOUD IS READY FOR PRIME TIME

Hybrid cloud storage has finally matured to a point where organizations of all sizes can consider it for their data storage environments. Security, once a major issue with cloud storage, has been dealt with by most cloud-integrated storage appliances, since all data is transferred to and stored in the cloud in an encrypted fashion. In addition, all the usual authentication methods used with on-premises storage are now available with these appliances. The benefits of hybrid cloud storage are so overwhelming for the majority of organizations that it should definitely be evaluated when upgrading or augmenting a storage infrastructure. ∎

ARUN TANEJA *is founder and president at Taneja Group, an analyst and consulting group focused on storage and storage-centric server technologies.*

DELL NIPS NETAPP FOR TOP MIDRANGE ARRAY HONORS

Dell beat NetApp in a seesaw battle to lead another strong field in the eighth Quality Awards for midrange arrays.

OVER THE COURSE of eight years of annual *Storage* magazine/SearchStorage.com Quality Awards, users have registered their satisfaction with their midrange storage array vendors by consistently rating them highly in all evaluation categories. This year is no exception. Dell Inc. managed to edge out NetApp Inc. in a very tight competition, and all seven finalists once again racked up solid scores.

Midrange arrays are in the "sweet spot" of networked storage. They boast enough performance and capacity to handle all but the most demanding of enterprise applications, but they also offer a reasonable entry point with room to grow for smaller and midsized companies. In our annual Storage Purchasing Surveys, midrange storage arrays have figured prominently in purchasing plans, accounting for nearly half of all storage system spending. Over the years, the midrange segment has grown and the products have added high-end features previously only available on the most expensive enterprise arrays.

By Rich Castagna

OVERALL RANKINGS

Dell and NetApp waged a seesaw battle through the five rating categories in our survey. Ultimately, Dell prevailed with an overall 6.80 rating based on a consistent set of scores that was good enough to place first or second in each category. But NetApp was hardly a slouch, matching Dell's feat of all first and second finishes en route to an overall score of 6.75.

NetApp took top honors in the sales-force competence, product features and product reliability rating categories; Dell led the group for initial product quality and technical support. Dell managed to overcome NetApp's three-to-two category advantage by winning its two categories by fairly substantial margins.

But there were solid performances from the other finalists, too. EMC Corp. (6.56), Hitachi Data Systems (6.46) and Hewlett-Packard (HP) Co. (6.42) all demonstrated similar consistency on their paths to overall scores that were high enough to have been winners in past surveys.

Of the previous seven Quality Awards, Compellent came out on top on two, EqualLogic nabbed one and Dell scored highest on another—so Dell's win this year isn't so surprising considering its shrewd acquisitions over the past few years.

MIDRANGE ARRAYS: OVERALL RANKINGS
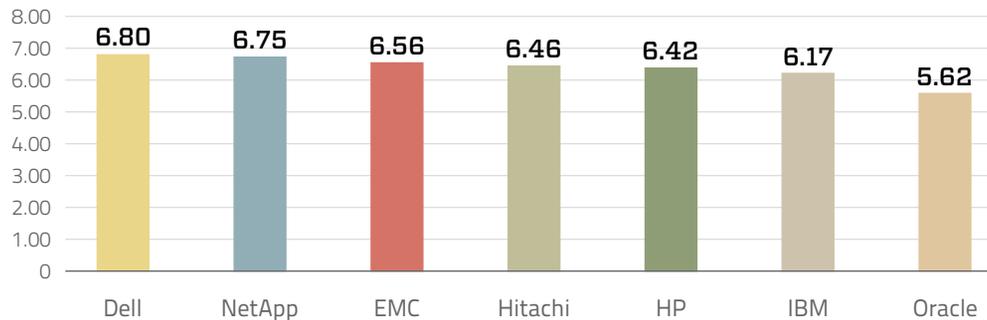
BY THE NUMBERS

Dell had the highest single category score, a 6.99 for initial product quality.

Dell "lost" three rating categories to NetApp, but by a total of only 0.09 points.

The average overall score for all seven finalists was 6.40, which was the third highest ever recorded for midrange arrays.

KEY STAT

0.05

Dell's overall winning margin over NetApp; four of the last seven midrange Quality Awards were decided by 0.05 points or fewer.

SALES-FORCE COMPETENCE

NetApp netted the highest score in the sales-force competence rating category, edging out Dell by 0.01. Dell had the highest scores on three of the six statements in the category, NetApp was tops on two statements and EMC had the best rating on the last. But the NetApp victory was the result of consistency, as it ranked second on three of the four statements that it didn't manage to win.
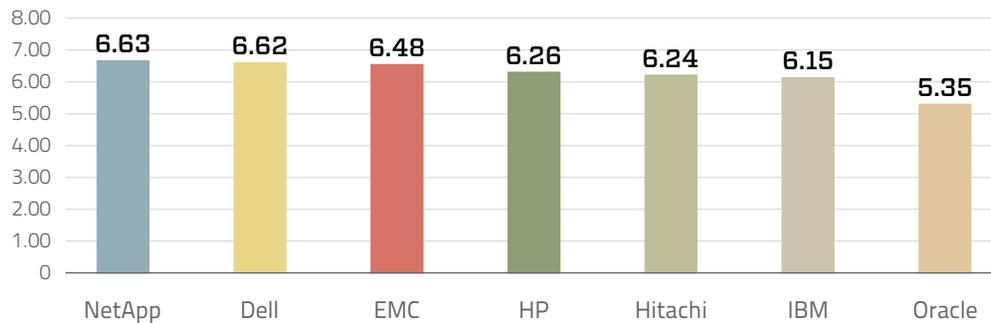
Dell had the highest single statement score, a 6.97 for "The vendor's sales support team is knowledgeable." Midrange vendors should pat themselves on the back for providing such effective sales support, as there were high scores across the board for the statement, including HP's second-place 6.80 rating, followed by NetApp (6.78) and EMC (6.75).

NetApp's highest category score was earned for "My sales rep is flexible" (6.71); its other winning statement was for having reps that understand customers' businesses (6.62).

EMC nosed out NetApp (6.74 to 6.72) for having sales reps knowledgeable about their customers' industries.

Dell's other two statement-leading scores were for having reps who are easy to negotiate with (6.64) and for the statement "My sales rep keeps my interests foremost" (6.52).

MIDRANGE ARRAYS: SALES-FORCE COMPETENCE

Bar chart values — NetApp: 6.63, Dell: 6.62, EMC: 6.48, HP: 6.26, Hitachi: 6.24, IBM: 6.15, Oracle: 5.35

BY THE NUMBERS
NetApp's sales-force competence scores were very consistent, ranging from a high of 6.78 (knowledgeable sales support team) to 6.42 (reps who keep customer interests foremost).

EMC won "My sales rep is knowledgeable about my industry" with a 6.74, but its highest category score was for the knowledgeable support team statement (6.75).

All the vendors can boast of having smart sales support teams, as that statement earned the highest average across all vendors (6.51).

KEY STAT
6.11
The lowest across-the-board average earned on "My sales rep keeps my interests foremost."

INITIAL PRODUCT QUALITY

Easy implementation has always been a hallmark of midrange arrays, and a key reason they appeal to large and smaller businesses alike. The initial product quality rating category assesses how quickly and effectively a new midrange system can be put into service. Mid-range storage array vendors have proven their implementation mettle, with this category receiving the highest average marks in seven of eight Quality Awards surveys to date.

The current set of finalists continues that legacy, led by Dell's impressive 6.99 category score. Dell rated highest on five of the category's six statements, with 7.00-plus scores for ease of use (7.04), products that require little vendor intervention (7.04), and being easy to get up and running (7.02). Dell just missed 7.00 scores for the other three statements, with ratings ranging from 6.91 to 6.98 all adding up to a remarkably consistent performance.

Second-place NetApp was hardly a slouch, with a leading 7.03 for "This product was installed without any defects," a 6.86 for products that require an acceptable level of professional services and a 6.82 for requiring little vendor intervention.

EMC wasn't off the pace by much, netting a solid 6.67 for third place, with Hitachi (6.49) and HP (6.46) neck-and-neck and just behind.

MIDRANGE ARRAYS: INITIAL PRODUCT QUALITY

Dell	NetApp	EMC	Hitachi	HP	IBM	Oracle
6.99	6.79	6.67	6.49	6.46	6.12	5.80
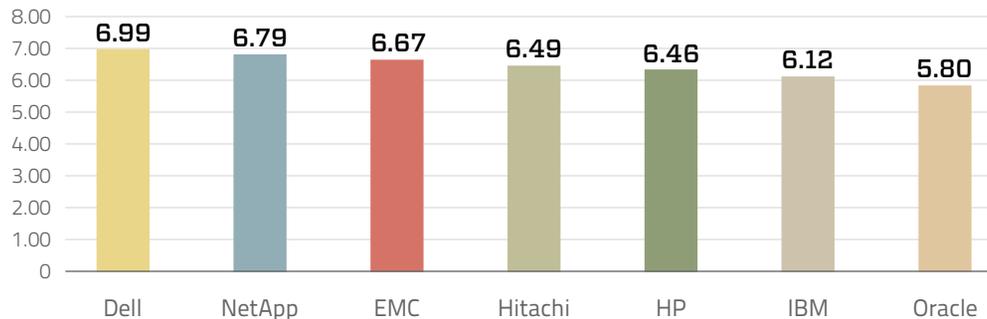
BY THE NUMBERS

The highest group average for any statement was 6.67 for "This product was installed without any defects."

Dell racked up an excellent 6.97 for the key statement "This product delivers good value for the money." But the statement earned the lowest group average (6.35) in this category.

KEY STAT
6.99

Dell's rating for initial product quality was the highest category rating in this edition of the Quality Awards.

PRODUCT FEATURES

A satisfying sales experience and an easy implementation are great ways to start a rela-tionship with a midrange array, but eventually you'll settle into daily operations and the product's features will take center stage. Of course, if a product rates highly in all the other Quality Awards categories, coming up a bit short feature-wise might be overlooked. But in our surveys, midrange systems have traditionally shown strength for their feature sets.

NetApp led the field for features ratings, garnering a 6.71 that put it just ahead of Dell's 6.67 and HP's 6.61. The top statement scores were shared by more vendors in this category than any other, with NetApp leading on three statements, Dell on two, and HP and Hitachi with one each.

NetApp's win came on the strength of its group-leading scores of 6.82 for capacity scaling, 6.76 for data mirroring and 6.74 for its overall feature set. Dell's ratings included three iden-tical 6.71 scores, but its two winners were a 6.84 for snapshot features and 6.76 for man-agement features.

HP had the highest rating for remote replication capabilities, while Hitachi won for "This product is interoperable with other vendors' products" (6.61).

MIDRANGE ARRAYS: PRODUCT FEATURES

6.71	6.67	6.61	6.48	6.38	6.09	5.91
NetApp	Dell	HP	Hitachi	EMC	IBM	Oracle
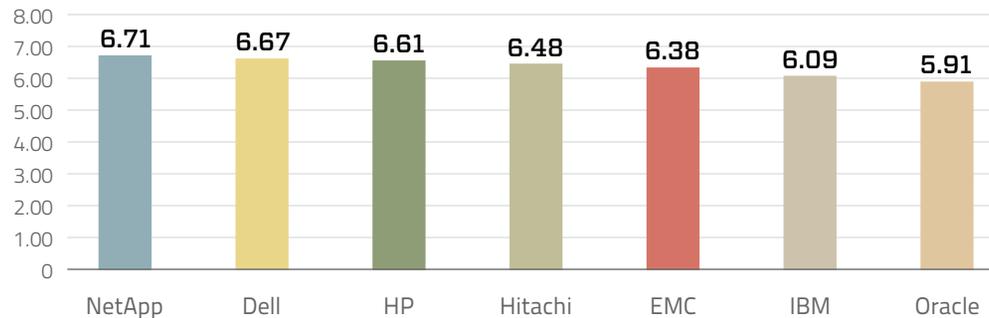
BY THE NUMBERS

NetApp had the best or second-best rating for six of the seven statements in this category.

Midrange systems are known for their flexibility, which was reflected by solid scores by all vendors for scalability.

All product lines also had ratings of 6.00 or better for "This product's mirroring features meet my needs."

KEY STAT

6.87

The highest single statement score in the features category, which was earned by HP for remote replication.

PRODUCT RELIABILITY

NetApp registered its third category win for product reliability, once again nudging out Dell by a small margin (0.04) and with the rest of the field also earning outstanding marks. Measuring performance over time, the reliability category represents the key factors that are likely to make the difference between satisfied and disappointed users.

NetApp earned its endurance chops by putting up the top scores for three of the five category statements and coming in second on the final two by slim margins. Once again, Dell scooped up the statements that NetApp didn't win to roll up a 6.84 category mark that fell a bit short of NetApp's 6.88.

Third-place EMC received a very good category score of 6.76, built on excellent grades for experiencing very little downtime (6.93) and meeting service-level requirements (6.90).

Hitachi also fared well for reliability. Its mark for the category was 6.63, highlighted by strong scores for meeting service-level requirements (6.77) and needing few unplanned patches (6.68).

The overall average of the finalist vendors in this category had the third highest rating for product reliability we've ever seen.

MIDRANGE ARRAYS: PRODUCT RELIABILITY

	NetApp	Dell	EMC	Hitachi	HP	IBM	Oracle
Score	6.88	6.84	6.76	6.63	6.42	6.22	5.80
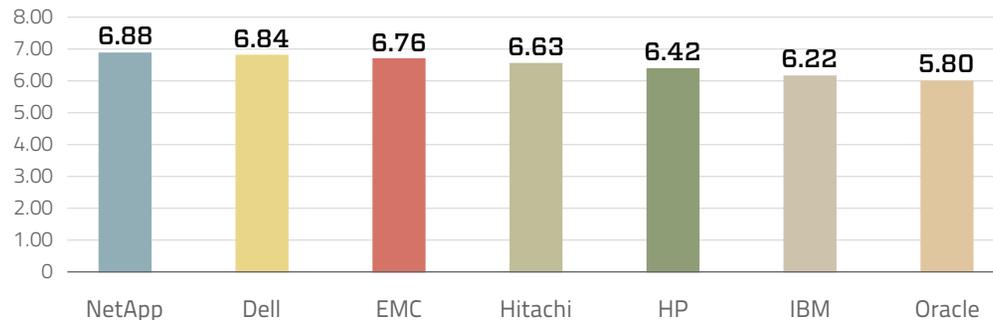
BY THE NUMBERS

Dell posted the highest statement mark (7.06) for "This product experiences very little downtime …"

… while NetApp neared that mark with a 7.04 for meeting service-level requirements.

NetApp just missed a couple more 7.00-plus scores on the downtime statement (6.96) and requiring few unplanned patches (6.94).

KEY STAT

6.62

As a group, the highest statement scores were a pair of 6.62s for meeting service-level requirements and experiencing little downtime.

TECHNICAL SUPPORT

Tech support can be a great equalizer—poor support can make users regret purchasing products that tout high performance and extensive feature lists, or make a less-than-spectacular product seem greater than the sum of its parts.
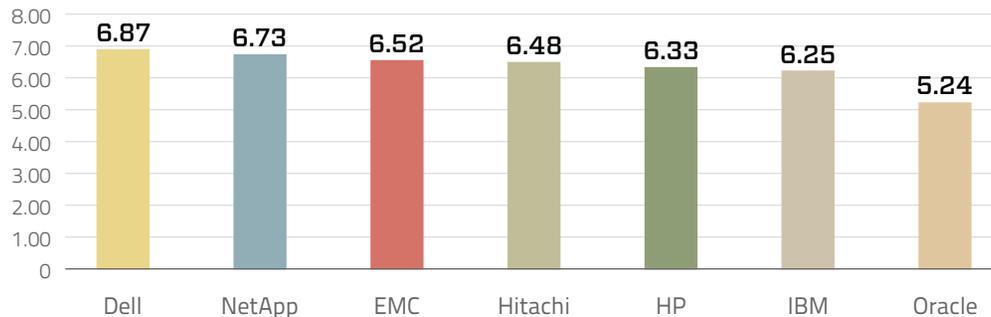
Once again, Dell and NetApp dominated the category—particularly Dell, which rode three 7.00-plus scores on the way to snagging six of the eight category statements. NetApp wasn't far off that pace, gaining the highest marks on the remaining statements and finishing second on all the others. EMC had another sturdy showing, finishing third in a category for the fourth time.

The 7.13 posted by Dell for "Vendor supplies support as contractually specified" was the highest statement rating across all categories; Dell also had a 7.03 for having a knowledgeable support staff and a 7.00 for resolving problems in a timely manner.

NetApp led the education statements by receiving a 6.85 for "Vendor's documentation/support materials are adequate" and a 6.73 for providing adequate training.

EMC's strengths were for delivering support as promised (6.90) and having knowledgeable support personnel (6.65).

MIDRANGE ARRAYS: TECHNICAL SUPPORT

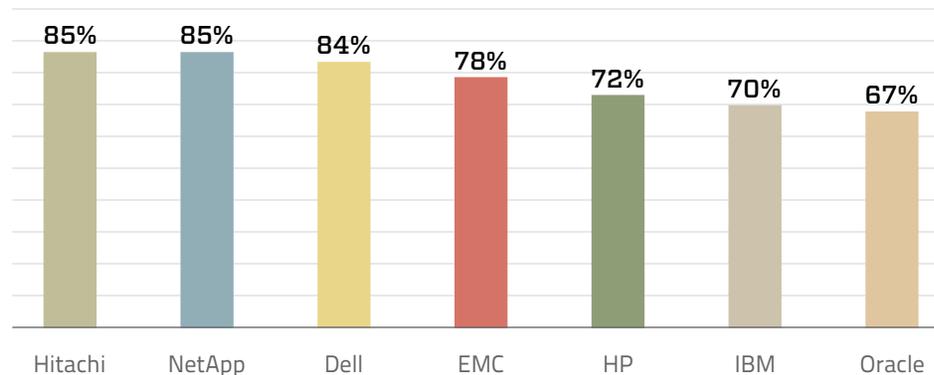Chart showing Technical Support scores: Dell 6.87, NetApp 6.73, EMC 6.52, Hitachi 6.48, HP 6.33, IBM 6.25, Oracle 5.24

WOULD YOU BUY THIS PRODUCT AGAIN?

In each Quality Awards survey we close out our product evaluation sections by asking users if they'd purchase the product again considering their real-life experiences with it. Often, the responses seem at odds with the category ratings, but this time our respondents' inclinations to repeat their purchases track fairly closely with the other ratings those products earned.

The top four finishers were the leaders for the buy-again question, but in a somewhat shuffled order, with Hitachi, NetApp and Dell in a virtual three-way tie followed by EMC.

MIDRANGE ARRAYS: WOULD YOU BUY THIS PRODUCT AGAIN?

Hitachi	NetApp	Dell	EMC	HP	IBM	Oracle
85%	85%	84%	78%	72%	70%	67%

RICH CASTAGNA *is editorial director of TechTarget's Storage Media Group.*

About the Quality Awards

The *Storage* magazine/SearchStorage.com Quality Awards are designed to identify and recognize products that have proven their quality and reliability in actual use. Results are derived from a survey of qualified readers who assess products in five main categories: sales-force competence, initial product quality, product features, product reliability and technical support. Our methodology incorporates statistically valid polling that eliminates market share as a factor. Indeed, our objective is to identify the most reliable products on the market regardless of vendor name, reputation or size. Products are rated on a scale of 1.00 to 8.00, where 8.00 is the best score. A total of 243 respondents provided 454 midrange storage array evaluations.

Products in the survey: The following vendors and midrange array model lines were included in the Quality Awards survey. The number of responses for each finalist is included in parentheses after the product name(s).

Celeros SmartSAN*

DataDirect Networks S2A Series*

Dell EqualLogic PS Series, Dell Compellent Storage Center or Dell CX Series (68)

Dot Hill AssuredSAN 3000/4000/5000 Series*

EMC VNX Series or Clariion CX Series (94)

Fujitsu Eternus DX400 Series*

Hewlett-Packard 3PAR StoreServ 7000, E/F/T Series, StoreVirtual 4000, EVA Series or P4000/P6000 Series (47)

Hitachi Data Systems Unified Storage Series, Unified Storage VM or Adaptable Modular Storage (AMS) Series (39)

IBM Storwize V7000, DS3950, DS5000 or DS6000 (43)

IceWEB 3000/6000/7000 Series*

Infortrend ESVA E60/F60/F70 Series*

NEC D4 Series/M-Series*

NetApp FAS2000/FAS3000 series (75)

Nexsan (now part of Imation) E Series, NST5000 Series or iSeries*

Oracle Sun ZFS Storage Appliances, Pillar Data Axiom or Sun Storage 25xx-M2 Array (21)

Overland Storage SnapSAN S3000 or SnapSAN S5000*

SGI InfiniteStorage 5000 Series*

Starboard Storage AC2000/AC4000 Series Storage Systems*

X-IO (formerly Xiotech) ISE or Hyper ISE Series*

Too few responses to qualify as a finalist

HOT SPOTS | TERRI McCLURE

Who is making today's IT purchasing decisions?

Users are making more and more rogue IT purchases.
Here's why it's happening and how IT can mitigate further risks.

ARE YOU AN information technology (IT) manager who feels as if you're losing control of crucial business application purchasing decisions? If you are, you're not alone. Recent ESG data finds users are making more and more rogue IT purchases, so it's time for you to get in front of the curve.

IT is losing its grip on technology buying decisions when it comes to business applications. The world has changed significantly from the command and control world where IT makes all the decisions related to what tools and applications can be used for business processes and data. We've all been hearing anecdotal data from software vendors, especially those in the online file-sharing and collaboration space, who report that knowledge workers have become increasingly involved in business application purchases. And we've strongly suspected that a shift was underway. But ESG research has now confirmed this trend with hard data. Indeed, a survey of 509 North American corporate knowledge workers found that not only are these non-IT personnel gaining more influence in this process, but, in many cases, are the sole decision makers in software purchases.

One of the most surprising outcomes of the research is that knowledge workers not only influence and decide on software purchases, they often do so without involving IT. This trend is popularly being called "shadow IT." Nearly half of the knowledge workers we surveyed either accessed or downloaded applications that weren't provided by IT. It would be easy to assume these are programs for personal use, such as Facebook, but that's not the case. Almost one-third said they accessed or downloaded software

for business use, while an additional 44% said they did so for both personal and business use.

While it's clear that knowledge workers are becoming more involved in official IT purchases, it's also the case that they're often doing so without IT involvement.

WHAT'S BEHIND THE MOVE?

A number of factors are conspiring to empower knowledge workers to become their own IT buyers. One factor is the consumerization of IT. As more consumers regularly buy Web-based and mobile applications, they become used to the idea of buying their own software. And as knowledge workers get used to the ability to download or access the software they want or need without IT involvement, they're more likely to continue doing so.

Another reason knowledge workers are buying more software on their own is simply because it's fast and easy to do. Cloud deployment turns everything on its head and changes the purchase model from one that favors IT to one that favors the non-IT buyer. In the past, purchasing a business application required it to be installed on-premises, leveraging company servers or desktops. The only entity even capable of doing that effectively was IT, whose concerns about security and support had typically led to the lockdown of desktops or laptops. But this approach is time-consuming and much more difficult to enforce with a Web-based or mobile application, especially

when the device is owned by the knowledge worker. Cloud-based applications put IT responsibilities in the hands of vendors, make it possible for knowledge workers to buy applications themselves, and the time to get applications up and running is minimal. In addition, the subscription model of many cloud applications makes it easier to purchase software from department or line-of-business operating expenses instead of having to incur a large capital expense as part of a big IT project, often putting these purchases under the radar of deep financial inspection.

TIME FOR IT TO CATCH UP

On the surface, this shift may seem to be a good one for companies. Knowledge workers no longer have to deal with project backlogs in IT, so they can order applications immediately to answer a timely need. Subscription pricing allows expenses to flex with headcount or requirements, eliminating the need to incur a large capital expense that may never be recouped and aligning expenses with the health of the business. And IT support, long a sore spot with knowledge workers, is also assumed by the independent software vendor, removing a major headache for IT and end users alike.

However, knowledge workers lack the collective training and knowledge of IT departments when it comes to data governance, regulatory compliance, data protection,

and security. They're also not skilled in assessing the potential complications of cloud-based and mobile applications arising out of legal issues such as legal holds due to litigation, IP leakage and implementation of defensible deletion policies.

To mitigate the risk that may come with non-IT-approved software purchases, IT has to approach the software purchase process differently. Instead of trying to control the process, IT should help guide and assist.

Recognizing that coercive attempts to restrict knowledge worker purchases will only result in more inventive ways to circumvent the rules, IT should position itself as a partner with specific expertise that the knowledge worker buyer can leverage. This will help to protect the business and the knowledge worker from a ruinous mistake. ∎

TERRI McCLURE *is a senior storage analyst at Enterprise Strategy Group, Milford, Mass.*

Software-defined storage might not be so radical after all

Software-defined storage isn't as new an idea as it may seem, as storage virtualization vendors have been working toward it for years.

REALIZING THAT THIS column will be published shortly after the 2013 VMworld conference, I'm quite confident that the trendiest marketing buzz among us is now—and will be post-VMworld— "software defined."

The idea is that virtualized services based on software will create an infrastructure that can be dynamically defined and able to adapt to changing business needs faster than infrastructures based on physical systems that must be managed independently. It'll be possible to provision optimal logical services irrespective of the makeup or even the location of the physical infrastructure. Software-defined everything is a hot topic, so it's no surprise that vendors have been touting software-defined storage (SDS).

IS SOFTWARE-DEFINED STORAGE FOR REAL?

The concept of software-defined storage strikes many data storage veterans as a pretty radical one. We've long been able to time, slice, dice and virtualize CPUs, memory and even Ethernet networks so that they can be shared and flexibly used by many different applications.

But it's a stretch to overlay that same model on storage. In contrast to a CPU or DRAM that processes and stores a transient digital bit that's on its way after a few microseconds, storage is intrinsically physical. Digital storage permanently parks a bit on disk as we collect bunches of bits over time. Essentially, there's a bucket of "data" that accumulates in our environment that must be managed over time.

So, how is it possible to software-define something that is so inherently physical?

SOFTWARE-DEFINED STORAGE
NOT SUCH A NEW IDEA

In reality, we've been chasing the possibility of software defining storage for years, and we're finally approaching a point where it might be practical. Moreover, the route we've been following to get to where we are now makes the idea of SDS seem a little less scary, while shedding some light on the real potential of software-defined storage.

The industry started its SDS quest when the first storage virtualization pioneers rolled out products. Those innovators were trying to make storage more malleable in the face of ongoing data growth amid constant environmental change. Arguably, storage virtualization had some rough spots for a number of years. Many vendors didn't seem to get the basic recipe right, but a number of them prevailed and are still running strong today, most notably Hitachi Data Systems with its Universal Storage Platform line, IBM's SAN Volume Controller and NetApp with its V-Series arrays. All three are molding heterogeneous storage virtualization offerings into tools that can work more closely with a virtual infrastructure. On the software-only side, DataCore and FalconStor had early and successful storage virtualization entries.

But storage virtualization still faces a couple of obstacles when it comes to creating "software-defined" storage. The biggest one is that storage virtualization is still pretty physical. Virtualizing storage might make heterogeneous collections of storage more dynamic and capable, but storage remains connected to a specific physical point—an appliance or controller—in the fabric. For many users, this is no longer OK; they need storage that can match the newfound mobility and fluidity of the rest of their infrastructure.

STORAGE VIRTUALIZATION HAS MADE ITS MARK

Fortunately, virtualization has had an influence on nearly all storage system architectures and, when combined with another trend in storage system architectures, it is bringing us closer to the possibility of software-defined storage.

First, the impact of storage virtualization has changed how storage systems of all types handle physical controllers and disks, even when they're just inside of a single array. This homogeneous, in-array storage virtualization has allowed storage vendors to make much better use of devices inside the array, making them less tied to the underlying physical controllers and disks.

Second, over the past couple of years, storage systems have moved increasingly toward a software-centric architecture, dispensing with requirements for specialized hardware and running entirely on standard x86 hardware. While there's still specialized hardware at the high end for systems that are built to operate at extreme scale and performance, a majority of midrange storage systems

run on standard x86 hardware.

Those two evolutions in storage systems appear poised to usher in the age of virtualization and further the quest for software-defined storage. A number of vendors offer their storage systems as virtual machines that run within the virtual infrastructure. The storage system no longer depends on any particular type of disk and the system runs on standard x86 hardware; this makes virtualization of the entire storage system an easy step for the storage vendor. Today, most of these implementations are packaged as virtual storage appliances (VSAs). Among the vendors offering VSAs are FalconStor, Hewlett-Packard, NetApp, Nexenta Systems, StorMagic and VMware. The idea is that these VSAs can be provisioned on top of a larger pool of physical storage, often direct-attached storage, but it can also be a SAN or network-attached storage. A VSA makes it easy to carve up the storage space, reclaim any stranded storage capacity and may deliver enhanced storage functionality that's more easily managed in the virtual infrastructure.

STORAGE IS STILL PHYSICAL, BUT MORE FLEXIBLE

VSAs don't make storage any less physical, but they offer several important benefits.

- **Storage can become more mobile.** While storage might still be tied to physical bits in a virtual storage appliance, it can be moved around (often without disruption), which could help put an end to disruptive hardware changes and data migrations.

- **VSA storage can be more adaptable than physical storage systems.** Capacity expansion can look just like expanding the capacity of any VM, rather than the process that expanding physical storage requires. Moreover, if the VSA can scale out, adding more capacity can be accomplished as easily as deploying another VSA.

- **Users gain the ability to deploy advanced storage capabilities** anywhere a workload needs it, whether on the premises or in a remote cloud.

Recent hands-on testing in the Taneja Group's labs has demonstrated that VSAs aren't just the toys or small storage products they were initially perceived to be. VSAs can compete with their hardware brethren, and they make pretty efficient use of virtual infrastructure resources. While they may not be the epitome of software-defined storage that's highly orchestrated and programmatically operated, they come pretty close. More importantly, VSAs are here now and are a practical enablement of software-defined storage that can add agility to a data center infrastructure. ∎

JEFF BOLES *is a senior analyst at Taneja Group.*

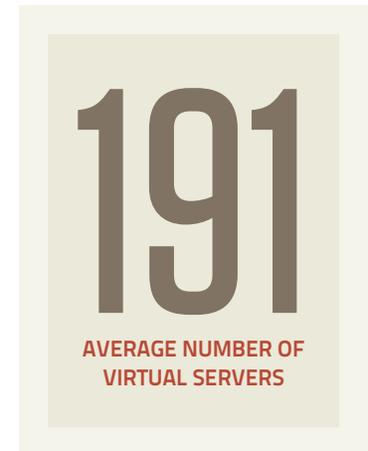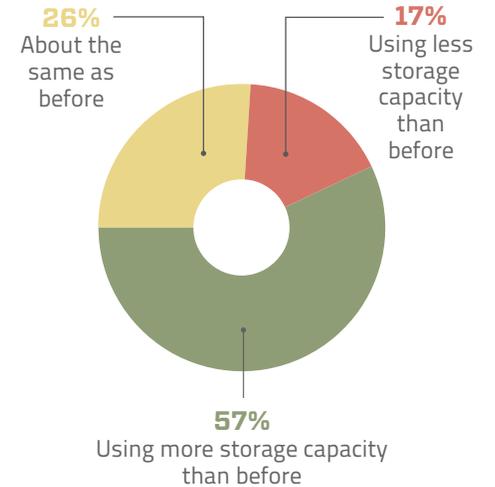## Storage for virtual servers getting easier, but it's still a big job

A WHOPPING 93% of companies have virtualized at least some of their servers, and most have had to make corresponding adjustments to their storage environments. With an average of 68% of their servers virtualized, IT managers still rely primarily on block storage to support the load: Fibre Channel SAN is the main virtual machine (VM) storage resource (56%), while 43% use iSCSI SANs; 38% use NAS and 16% use multiprotocol arrays. Sixty-four percent have at least one array dedicated to supporting VMs, and 69% have specifically purchased a new storage system for their VMs. And it's a good thing they have those new arrays, as 57% are using more capacity with VMs than they did before virtualization. How hard/easy is it to manage all this stuff? It's a pretty even split: 37% say it's a little/lot harder than before, 28% claim it's about the same and 34% find it a little or a lot easier. —*Rich Castagna*

SINCE VIRTUALIZING SOME/ALL OF YOUR SERVERS, HOW WOULD YOU DESCRIBE YOUR COMPANY'S STORAGE CAPACITY USAGE?

26%
About the same as before

17%
Using less storage capacity than before

57%
Using more storage capacity than before

WHAT TYPE OF STORAGE DO YOU USE TO SUPPORT YOUR VIRTUAL SERVERS?*

Fibre Channel SAN	56%
iSCSI SAN	43%
NAS	38%
Directly attached	26%
Multiprotocol (Unified)	16%
Other	1%
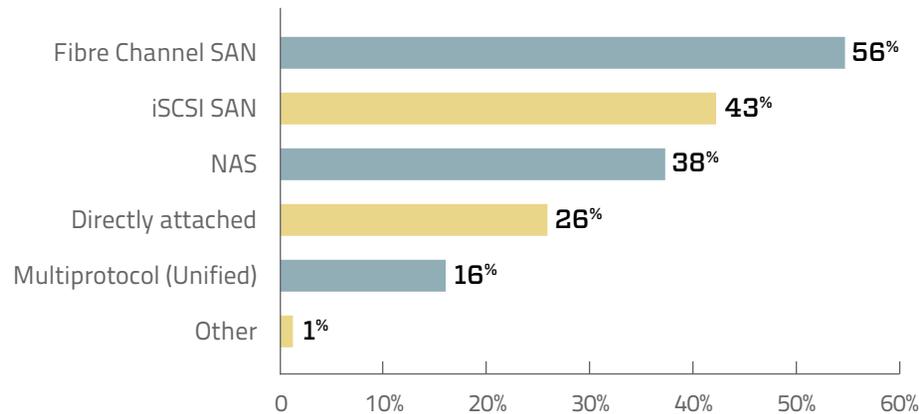
*Multiple selections allowed

191

AVERAGE NUMBER OF VIRTUAL SERVERS

TechTarget Storage Media Group

HOME

DATA, TELL US ALL
ABOUT YOURSELF

NO CURE FOR SUMMER-
TIME STORAGE BLUES

MORE CHOICES
FOR VIRTUAL SERVER
BACKUPS

APPLIANCES LINK
YOUR STORAGE
TO THE CLOUD

DELL OUTDUELS NETAPP
FOR TOP MIDRANGE
ARRAYS

WHO'S MAKING THOSE
IT BUYING DECISIONS?

SOFTWARE-DEFINED
STORAGE LOOKS LIKE
STORAGE VIRTUALIZATION

STORAGE SETUP FOR
VMs GETTING EASIER