Managing the information that drives the enterprise

STORAGE

Vol. 8 No. 9 January 2010

vSphere AND Storage

Virtual servers are a sys admin's dream, but they can be a nightmare for storage managers. But VMware says it has worked out many of the kinks with its new vSphere. p. 10

ALSO INSIDE

TechTarget
The Technology Media ROI Experts


3PAR
Serving Information

Cover illustration by ENRICO VARRASSO

Thin

Start Thin. Get Thin. Stay Thin.

Why stop at thin provisioning?

3PAR revolutionized data storage with the introduction of 3PAR Thin Provisioning.

Today, 3PAR's thin technologies combine hardware and software innovations that take thin provisioning to the next level by enabling your organization to not only *start thin*, but to *get thin* and *stay thin*.

What are you waiting for? Take it to the next level today:
3PAR.com/thin

3PAR
Serving Information

Get Agile. Get Efficient. Get 3PAR.

Storage should get more interesting in 2010

There was plenty of buzz about solid state and VMware in 2009, but here's hoping that the real storage innovations will come in 2010.

IT COULD BE one of those "out with the old, in with the new" things, but with the calendar turning a page and our economy begin to shrug off the effects of a numbing recession, I think it's a good time to look at the state of the data storage industry.

We've seen aggressive sales promotions like the "cash for clunkers" discounts on new storage arrays that a few vendors have been pushing. And I'm sure there are plenty of storage managers who are reveling in the moment where it looks like the vendors' backs are up against the wall for a change. But it's more instructive to look at how product lines have been shaped or reshaped over the last year or so.

Our Purchasing Intentions surveys have shown that over the last few years storage buyers' criteria for selecting particular disk system vendors over others has shifted. Once upon a time it was strictly a features game, with storage managers finding a way to come up with the extra bucks to buy the products with the most impressive and useful feature sets. But the key reasons for sticking with a vendor—or switching to a new one—are beginning to shift in favor of criteria like tech support, familiarity with a vendor and, yes, price.

While vendors struggle to differentiate themselves from the competition, the products themselves are looking more and more like commodities.

That shift seems to have put storage vendors in something of a bind. While they struggle to differentiate themselves from the competition, the products themselves are looking more and more like commodities; highly sophisticated commodities perhaps, but commodities nonetheless. Storage buyers expect a comprehensive set of features (for them it's a given) and are looking beyond what's built into the hardware or software.

For a lot of vendors, this has caused something of a race to the middle. They sometimes endeavor more to match their competition than to beat

The way we back up remote offices
Ford has a better storage idea
Quality Awards: Top NAS products
SAS puts squeeze on FC
vSphere and storage

it. In this "Be like Mike" environment, we're seeing less effort exerted to be a leader in a particular technology, and a lot more energy spent trying to avoid being the vendor that lacks something on their checklist of features or options.

A couple of examples stand out. With solid state being touted as the future of storage (led by EMC Corp.), storage vendors have been scurrying to get solid-state drives (SSDs) on their product lists. But, for the most part, all they've really done is sign up with a solid-state OEM that can provide flash storage in a form factor and with a compatible protocol so they can simply plug it into their existing arrays. Are there benefits? Sure, but they could be even greater if the flash drives were integrated in a way that took advantage of their unique technology. But there hasn't been much of that, and what we've seen so far is vendors just playing catch-up to try to keep the playing field level.

Another example: VMware Inc. is a juggernaut that's not only putting the squeeze on its direct competition but turning the rest of IT—especially storage—on its ear. Server virtualization puts significant new demands on storage systems, but the response from the storage industry has been fairly tepid. You'd be hard-pressed to find a storage vendor not boasting that it's "VMware certified," but it's quickly become a fairly hollow claim. I don't know what a vendor has to do to get VMware's stamp of approval, but I bet it's roughly equivalent to having a pulse and probably doesn't involve much more than being able to present storage capacity to a virtual server in some reasonable manner. It looks like the storage gang has been content to wait for VMware to smooth out the storage wrinkles, which it has apparently done to a great degree with vSphere (see "Everything you need to know about vSphere and data storage," p. 10).

It looks like the storage gang has been content to wait for VMware to smooth out the storage wrinkles, which it has apparently done to a great degree with vSphere.

There are, of course, some vendors that are still innovating despite the overall tenor of the storage market. Compellent Technologies has recently received some much-deserved attention for its Data Progression software, which looks like the biggest development in years in the otherwise moribund information lifecycle management (ILM) market. Kudos should also go to Symantec Corp. for its OpenStorage (OST) API, which—when hardware vendors add support for it—will provide more capable and comprehensive management of disk-based backups from within Veritas NetBackup.

Those are just a couple of examples, and I'm sure I'm slighting a few

The way we back up remote offices
Ford has a better storage idea
Quality Awards: Top NAS products
SAS puts squeeze on FC
vSphere and storage

other companies. But vendors that continue to pursue innovation seem fewer these days.

Maybe I'm being a little too hard on the data storage industry; maybe I expect too much during a time when nearly all vendors have seen sales dip, and it isn't a sin to just try to maintain the status quo until things start looking up again. But rushing to the middle of the pack with cookie-cutter feature sets doesn't make a lot of sense either. In storage, as in much of IT, it takes a long time for things to change. Years pass before storage managers swap out arrays or switches or backup apps. Things move slowly. Slow enough for vendors to take the time to develop some truly unique applications and features. ⊙

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 February 2010 issue.*

Ensure you're dealt a winning hand.

CIO Decisions Ezine provides CIOs and business strategists with expert and peer insight "from the trenches" on pressing issues decision-makers face every day, including:

- Project and Portfolio Management
- Virtualization
- Business Process Management
- Outsourcing
- And much more!

Subscribe Today!

CIO DECISIONS
ezine

The way we back up remote offices
Ford has a better storage idea
Quality Awards: Top NAS products
SAS puts squeeze on FC
vSphere and storage

COMING IN FEBRUARY

2009 Products of the Year Awards

The results of the seventh annual *Storage* magazine and SearchStorage.com Storage Products of the Year awards are announced. See who cops the Gold, Silver and Bronze awards in the following five categories: disks and disk subsystems, storage networking equipment, backup and disaster recovery hardware, backup and disaster recovery software and services, and storage management tools.

Virtualization Comes to Storage Networks

While a great deal of attention has been focused on virtualizing servers or data storage, products have appeared in the last few years that can virtualize storage networks.

Virtualized networks can offer unprecedented flexibility, allowing storage administrators to perform configuration changes in software. See how these new approaches to storage networking stack up against the competition.

Take Control of NAS Systems

As every storage manager knows, network-attached storage (NAS) systems tend to proliferate. With file data growing at a staggering rate, managing multiple NAS boxes is becoming a management nightmare. In this story, we describe how to get a handle on burgeoning file stores, including using clustered NAS systems, global file systems, file virtualization and other NAS management tools.

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

Everything you need to know about vSphere and data storage

VMware has been a boon for system administrators, but not so much for storage pros. vSphere addresses many of the storage-related shortcomings of its predecessor.

By Eric Siebert

WHEN VMware Inc. released vSphere in May 2009 it included more than 100 new features and enhancements, many of which addressed data storage. These data storage management enhancements include:

- Thin provisioning enhancements
- iSCSI improvements
- Support for Fibre Channel over Ethernet (FCoE) and jumbo frames

The way we back up remote offices
Ford has a better storage idea
Quality Awards: Top NAS products
SAS puts squeeze on FC
vSphere and storage

- New ability to hot extend virtual disks
- New ability to grow VMFS volumes
- New Pluggable Storage Architecture (PSA)
- New paravirtualized SCSI adapters
- New VMDirectPath for storage I/O devices
- VMware Storage VMotion enhancements
- New vStorage APIs
- New storage views and alarms in vCenter Server

There are many vSphere storage enhancements and some may have a profound effect on your environment, so we'll look at each one in detail.

THIN PROVISIONING ENHANCEMENTS

Thin provisioned disks are virtual disks that start small and grow as data is written to them. With a "thick" disk, all of its space is allocated when it's created; but a thin disk starts out at 1 MB (or up to 8 MB depending on the default block size) and then grows its defined maximum size as data is written to it by the guest OS.

Thin provisioning was available in Virtual Infrastructure 3 (VI3), but a number of changes make it more usable in vSphere:

- Thin disks can be created using the vSphere client at the time a virtual machine (VM) is created; with VI3, the vmkfstools command line utility was used to create them.

Thin provisioned disks are virtual disks that start small and grow as data is written to them.

- Existing thick disks can be converted to thin disks using Storage VMotion while a VM is running; VI3 required powering off the virtual machine.
- The vSphere client lets you see the actual size of thin disks (previously a command line function in VI3).
- New configurable alarms in vCenter Server provide alerts for overallocation and usage percentages.
- A new safety feature automatically suspends VMs with thin disks when free space is critically low.

These improvements make thin disks more manageable and much easier to use. Users often ask if they should use VMware's thin disk if their storage array supports thin provisioning ("thin-on-thin"). Use both if you can, but make sure you carefully monitor the array and VMware to ensure that you have adequate space available. Another concern with thin disks is the impact on performance as they grow and the increased fragmentation that may occur. According to VMware, thin disks have a negligible effect on performance.

The way we back up remote offices
Ford has a better storage idea
Quality Awards: Top NAS products
SAS puts squeeze on FC
vSphere and storage

iSCSI IMPROVEMENTS

iSCSI storage arrays have become a popular storage choice for virtual hosts due to their lower cost and acceptable performance. Using iSCSI software initiators has always resulted in a slight performance penalty vs. hardware initiators with TCP offload engines. For vSphere, VMware rewrote the entire iSCSI software initiator stack to use CPU cycles more efficiently and to improve throughput compared to VI3.

VMware enhanced the VMkernel TCP/IP stack, optimized the cache affinity and improved internal lock efficiency. Other iSCSI improvements include easier provisioning and configuration, and support for the bi-directional Challenge Handshake Authentication Protocol (CHAP), which improves security by requiring both the initiator and target to authenticate each other.

SUPPORT FOR FCoE AND JUMBO FRAMES

vSphere added support for newer storage and networking technologies, including:

- **Fibre Channel over Ethernet.** vSphere now supports FCoE on converged network adapters (CNAs).
- **Jumbo frames.** Conventional Ethernet frames are 1,518 bytes long; jumbo frames are typically 9,000 bytes, which can improve network throughput and CPU efficiency. VMware added jumbo frame support in ESX 3.5 but didn't officially support it for use with data storage protocols. With the vSphere release, they officially support using jumbo frames with software iSCSI and NFS storage devices using 1 Gbps or 10 Gbps network interface cards (NICs).

HOT EXTEND VIRTUAL DISKS

With VMware vSphere you can increase the size of an existing virtual disk (VMDK file) while it's powered on as long as the guest operating system supports it.

Once you increase the size of a virtual disk, the guest OS can then begin using it to create new disk partitions or to extend existing ones. Windows Server 2008, Windows Server 2003 Enterprise and Datacenter editions, and certain Linux distributions support this feature. Previously, you had to power down a VM before increasing its virtual disk size.

GROW VMFS VOLUMES

With vSphere you can increase the size of VMFS volumes without using extents and without disrupting virtual machines. To do this in VI3, you had to join a separate logical unit number (LUN) to the VMFS volume as an extent, which had some disadvantages. vSphere lets you grow the LUN of an existing VMFS volume using your storage-area network (SAN) configuration tools and then expand the VMFS volume so it uses the additional space.

This means you don't have to use extents and can avoid moving VMs to other data stores to destroy existing VMFS volumes to create larger ones.

The way we back up remote offices
Ford has a better storage idea
Quality Awards: Top NAS products
SAS puts squeeze on FC
vSphere and storage

PLUGGABLE STORAGE ARCHITECTURE

VMware has given vSphere a new modular storage architecture that allows third-party vendors to interface with certain storage functionality. The Pluggable Storage Architecture (PSA) allows vendors to create plug-ins for controlling storage I/O functions like multipathing.

There's built-in functionality that allows for fixed or round-robin path selection when multiple paths to a storage device are available. Vendors can expand on this and develop their own plug-in modules for optimal performance through load balancing and more intelligent path selection. To achieve this, the PSA leverages the new capabilities provided by the vStorage APIs for multipathing.

PARAVIRTUALIZED SCSI ADAPTERS

Paravirtualization is a technology available for certain operating systems that use a special driver to communicate directly with the hypervisor. Without paravirtualization the guest OS doesn't know about the virtualization layer and privileged calls are trapped by the hypervisor using binary translation.

Paravirtualization allows for greater throughput and lower CPU utilization for virtual machines, and is useful for disk I/O-intensive applications. Paravirtualized SCSI adapters are separate storage adapters that can be used for non-primary OS partitions and can be enabled by editing a VM's settings and enabling the paravirtualization feature.

This may sound similar to VMDirectPath, but the key difference is that paravirtualized SCSI adapters can be shared by multiple VMs on host servers and don't require dedicating a single adapter to a single virtual machine.

VMDIRECTPATH FOR STORAGE I/O DEVICES

VMDirectPath is similar to paravirtualized SCSI adapters where a VM can directly access host adapters and bypass the virtualization layer for better throughput and reduced CPU utilization. But with VMDirectPath you must dedicate an adapter to a VM and it can't be used by any other virtual machines on that host.

VMDirectPath is available for specific models of both network and storage adapters, however, only the network adapters are currently fully supported in vSphere; storage adapters only have experimental support. Like PVSCSI adapters, VMDirectPath can be used for VMs that have very high storage or network I/O requirements like database servers. VMDirectPath enables virtualization of workloads that you might previously have kept physical. A downside to using VMDirectPath is that you can't use features like VMware VMotion and Distributed Resource Scheduler (DRS).

STORAGE VMOTION ENHANCEMENTS

While VMotion moves a running VM from one host to another leaving the virtual machine location intact, Storage VMotion (SVMotion) keeps

The way we back up remote offices
Ford has a better storage idea
Quality Awards: Top NAS products
SAS puts squeeze on FC
vSphere and storage

the VM on the same host and only changes the VM's storage location. SVMotion was first introduced in ESX Version 3.5, but was only available as a command line utility. In vSphere, it's integrated in the vSphere Client, allowing quick and easy SVMotion moves. In addition, SVMotion now allows thick-to-thin disk conversion (and vice versa).

SVMotion can also be used to re-shrink a thin disk after data has been deleted from it. Typically, you use Storage VMotion to move a VM location to another storage device; however, you can also leave the VM on its current storage device when performing a disk conversion. SVMotion can be invaluable when performing data storage maintenance as running virtual machines can be easily moved to other storage devices.

Some under-the-covers enhancements make the whole migration process much more efficient. Instead of using a snapshot when copying the disk to its new location and then committing it when the operation is complete, SVMotion now uses a new changed block tracking feature to keep track of blocks that changed after the move process started and then copies them after it completes.

NEW vSTORAGE APIs

vStorage APIs are a collection of interfaces that third-party vendors can use to seamlessly interact with data storage in vSphere. There are four categories of vStorage APIs:

Array integration. These APIs are being co-developed with specific storage vendors. When complete, they will allow access to array-oriented capabilities such as array-based snapshots, hardware-offloaded storage device locking, integration between VMware and array-level thin provisioning, storage provisioning, data replication and more. In some cases, greater efficiencies may be realized by allowing the storage array to perform certain operations. VM cloning or template-based deployment can be hardware accelerated by array offloads rather than file-level copy operations at the ESX server. Storage VMotion would be able to leverage the storage array features to copy data more rapidly and with less ESX host impact; and rather than the traditional "SCSI lock" mechanism used by VMware, the array can lock only specific blocks being updated, which dramatically increases the number of VMs that can be deployed on a data store.

Multipathing. These APIs are used by the Pluggable Storage Architecture to allow storage vendors to more intelligently use multipathing for better storage I/O throughput and storage path failover. Storage vendors must certify their multipathing extensions modules with VMware for use with ESX(i). There are several vStorage multipathing APIs: A path selection plug-in (PSP) can extend the path selection algorithms for any given I/O; a storage array-type plug-in (SATP) allows new/changed path discovery and ongoing path state management; and a multipathing plug-in (MPP) can extend the entire path management model of vSphere, including path management and path selection.

The way we back up remote offices
Ford has a better storage idea
Quality Awards: Top NAS products
SAS puts squeeze on FC
vSphere and storage

Site Recovery Manager (SRM). These APIs are used to integrate SRM with array-based replication for block and NAS storage to allow SRM to seamlessly handle both VM and host failover and storage replication failover. They also allow SRM to control the underlying array-based replication that it relies on.

Data protection. These APIs replace VMware Consolidated Backup (VCB) that was introduced in VI3. While they include VCB functionality, they also add new features such as changed block tracking and the ability to directly interact with the contents of virtual disks via the VDDK. These APIs are for backup and data protection app vendors to provide better integration.

vCENTER SERVER'S NEW STORAGE VIEWS AND ALARMS

VMware also improved storage-related reporting and alarms in vCenter Server. The most conspicuous is a new storage view that shows detailed information on storage metrics. Alarms have been expanded to include specific storage-related issues like datastore overcommitment and low disk space.

vCenter Server's storage view is a plug-in that must be installed and enabled. Once enabled, an additional Storage View tab will appear in the right pane after selecting any object in the left pane. The

CBT CHANGED BLOCK TRACKING

Changed block tracking (CBT) is a significant new data storage feature that's especially important for backup, replication and other data protection applications. vSphere's VMkernel can now track which disk blocks of a virtual machine (VM) have changed from a particular time. By tapping the VMware vStorage APIs for data protection, applications can get the information from VMkernel rather than figuring it out on their own. CBT also enables near-real-time continuous data protection (CDP) when replicating VM disk files. CBT can also speed up incremental backups because backup apps can easily find out which changed blocks need to be backed up. Restoring data is much easier, too, with CBT because backup apps will know exactly what blocks need to be put back on the virtual disk for the restore point selected.

CBT is disabled by default because there's a slight performance overhead associated with it. It can be enabled only on the VMs that require it by adding a configuration parameter to the VM; backup applications that support changed block tracking can also enable it on VMs. Once enabled, CBT stores information about changed blocks in a special –ctk.vmdk file that's created in each VM's home directory. To do this, CBT uses changeIDs that are unique identifiers for the state of a virtual disk at a particular point in time. New changeIDs are created anytime a snapshot of a VM is created by a backup application. Using the changeID, a backup application will know which blocks have changed since the last backup. CBT is only supported on VMs that have virtual machine hardware Version 7 (which is new to vSphere), so older VMs will need to have their virtual hardware upgraded to use CBT.

The way we back up remote offices
Ford has a better storage idea
Quality Awards: Top NAS products
SAS puts squeeze on FC
vSphere and storage

vSPHERE vs. HYPER-V AND XENSERVER

VMware's not the only game in town: Microsoft Corp.'s Hyper-V and Citrix Systems Inc.'s XenServer are also popular server virtualization tools. This table compares vSphere's storage features with those of Hyper-V and XenServer.

	VMware vSphere 4.0	Microsoft Hyper-V R2	Citrix XenServer 5.5
Thin provisioning of virtual disks	Yes	Yes	Limited, requires StorageLink and supported SAN
Fibre Channel over Ethernet (FCoE) support	Yes	Supports FCoE but not InfiniBand	Limited, no FCoE or InfiniBand
Jumbo frame support	Yes	Yes	No
Hot extend of virtual disks	Yes	No	No
Dynamically grow of storage volumes	Yes	No, requires third-party SAN tools to grow and shrink LUNs	No
Paravirtualization support for storage adapters	Yes	Yes	Yes, for Linux guests
Direct access to storage devices (VMDirectPath)	Yes	No	No
Live VMware Storage VMotion	Yes	No, Quick Storage Migration is not live	No
APIs for data protection	Yes	No	No
APIs for multipathing	Yes	No	Yes
Monitoring and alerting for thin disks	Yes	No	No

storage view has selectable columns that display information such as the total amount of disk space a VM is using (including snapshots, swap files, etc.), total amount of capacity used by snapshots, total amount of space used by virtual disks and other capacity usage statistics.

This is a great tool to quickly see how much space is being used in your environment for each component and to easily monitor snapshot space usage. There's also a map view to see relationships between virtual machines, hosts and storage components.

In VI3, the only storage alarm was for host or VM disk usage (in KBps). Hundreds of new alarms have been added, with many of them related to storage, such as an alert for a datastore that's close to running out of free space. This is especially important when you have a double threat from both snapshots and thin disks that can grow and use up all the free space. Other storage-related alarms include:

- Datastore disk overallocation percent
- Datastore state to all hosts
- Datastore created/increased/deleted/discovered/expanded
- Degraded storage path redundancy
- Lost storage connectivity

The way we back up remote offices
Ford has a better storage idea
Quality Awards: Top NAS products
SAS puts squeeze on FC
vSphere and storage

vSPHERE DOES STORAGE BETTER

In addition to vSphere's major enhancements for storage operations, there are many smaller improvements not covered here. Taken together, these enhancements provide better performance, improved usability and easier administration. And they may be compelling enough to convince current VMware users to upgrade to vSphere. ⊙

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

SAS challenges Fibre Channel drives

In their 3 Gbps incarnation, SAS drives have proved popular in low-end and nearline data storage systems, but at 6 Gbps, SAS-2 poses a serious threat to Fibre Channel interface drives. *By Jacob Gsoedl*

SINCE ITS INTRODUCTION in 2004, serial-attached SCSI (SAS) has become the prevailing data storage interface in servers and is on a trajectory to become the same for external storage systems. Those systems are still dominated by Fibre Channel (FC) and serial ATA (SATA) disk drives, but next-generation SAS-2 components, available since early 2009, are hastening this transition.

With 6 Gbps transfer rates and enhancements that neutralize some of the shortcomings of SAS 1.1, SAS-2 has overtaken FC as the most advanced and fastest disk interface. Because FC drives with their 4 Gbps transfer rate and 3.5-inch form factor have reached the end of the road, it's only a question of time as to when SAS will emerge as the drive interface of choice for high-end storage systems and servers. Thanks to its

The way we back up remote offices
Ford has a better storage idea
Quality Awards: Top NAS products
SAS puts squeeze on FC
vSphere and storage

prevalence in servers today, SAS—mostly 3 Gbps SAS—already owns approximately 40% of the multiuser disk interface market, trailing SATA's 41.7% share and outpacing the 16% market share of FC-AL/4 Gbps, according to recent analysis by Stamford, Conn.-based Gartner Inc. "SAS will account for 57% and SATA for 43% of the total multiuser HDD [hard disk drive] market by 2013," predicted John Monroe, a Gartner research vice president, who pegged the extinction of FC drives to the year 2013.

THE IMPACT OF SAS ON FC AND SATA DRIVES

Despite some disagreement on timing, most data storage analysts agree that the days of FC as a disk interface are numbered. Fibre Channel disks were never able to extend beyond high-end servers and storage arrays and, as result, costs remained high (e.g., a SAS host bus adapter [HBA] costs approximately one-third the price of a comparable FC HBA). Moreover, users have been coping with disparate disk interfaces—FC for high performance and SATA for large-capacity storage—requiring conversion chips for the two to coexist in a single storage system. Conversely, SAS and SATA can be supported by a single chip, reducing interface costs, complexity and power requirements. "Enterprise-class SAS drives are on par with high-end FC drives, except SAS has a smaller cooling and power footprint, and enables a single back-end interconnect technology for both SAS and SATA," said Tony Palmer, senior engineer and analyst, ESG Lab at the Milford, Mass.-based Enterprise Strategy Group (ESG).

Contrary to the unanimous vision on the fate of FC as a hard disk drive interface, the storage industry is divided on the impact of SAS on the future of SATA as an enterprise disk interface. SAS and SATA have much in common: Both are serial point-to-point interfaces with advanced features like command queuing, support for hot swapping and commensurate transfer rates (Seagate Technology LLC shipped the first 6 Gbps SATA drives with the Barracuda XT 2 TB in October 2009).

But SATA is impeded by severe shortcomings as a result of its pedigree and original purpose. Conceived to replace parallel ATA in desktop computers and never intended for the enterprise space, SATA is handicapped by its lack of dual porting required for redundant configu-

THE EVOLUTION OF SCSI

SCSI STANDARD	MAXIMUM SPEED	BUS WIDTH (BITS)
SCSI-1	5 MBps	8
Fast SCSI	10 MBps	8
Fast Wide SCSI	20 MBps	16
Ultra SCSI	20 MBps	8
Wide Ultra SCSI	40 MBps	16
Ultra2 SCSI	40 MBps	8
Wide Ultra2 SCSI	80 MBps	16
Ultra3 SCSI	160 MBps	16
Ultra320 SCSI	320 MBps	16
3 Gbps SAS	3 Gbps	Serial
6 Gbps SAS	6 Gbps	Serial
12 Gbps SAS (Production deployment expected in 2012)	12 Gbps	Serial

The way we back up remote offices
Ford has a better storage idea
Quality Awards: Top NAS products
SAS puts squeeze on FC
vSphere and storage

rations, its limited ATA command set vs. SAS's rich SCSI command set, signal integrity challenges with 6 Gbps (and beyond) SATA, inferior command queuing as a result of its shallow queue depth, and limited error handling and diagnostic capabilities. To mitigate some of these shortcomings, storage array vendors have used SAS-to-SATA bridges and SATA multiplexers (MUX), which are basically interposers that sit between SATA drives and the SAS backplane to enable higher reliability, dual porting and better diagnostics for SATA drives. "We offer a $30 to $40 active-active SATA MUX adapter option for our VTrak RAID and JBOD subsystems to provide dual porting for SATA drives," said Ray Bahar, vice president of the Americas at Promise Technology Inc.

While the interposer approach has been an acceptable workaround to compensate for SATA's shortcomings, it's questionable if SATA will survive in enterprise storage beyond 6 Gbps. Seagate has been offering so-called nearline SAS drives that compete with SATA drives in the enterprise space. Designed for high capacity and lower cost, nearline SAS drives are identical to SATA drives with the exception of the interface; in other words, they peak at 7,200 rpm and 2 TB capacity. Contrary to smaller capacity, mission-critical tier 1 10K to 15K high-performance SAS drives, nearline SAS drives are intended for tier 2 applications and for more static data. Not burdened by the limitations of the SATA interface, nearline SAS drives perform significantly better than comparable SATA drives. "We see at least a 30% performance improvement with identical drives from the same vendor by simply having a SAS interface," said Howard Shoobe, senior marketing manager at Dell Inc. "Although it's

IDEAL ENVIRONMENTS FOR HARD DISK TYPES

	DESKTOP COMPUTING	ENTERPRISE: HIGH CAPACITY (NEARLINE)	ENTERPRISE: HIGH PERFORMANCE
Interfaces	PATA and SATA	SATA and SAS	SAS and Fibre Channel
Form factor	3.5-inch	2.5-inch and 3.5-inch	2.5-inch and 3.5-inch
Capacity	80 GB to 2 TB	160 GB to 2 TB	73 GB to 600 GB
Rotational speed	5,400 rpm	7,200 rpm	10K and 15K rpm
Sustained data rate	95 MBps to 138 MBps	95 MBps to 150 MBps	122 MBps to 204 MBps
MTBF (hours)	700,000	1,200,000	1,600,000
Usage profile	8x5; on as needed	24x7; always on	24x7; always on
Workload	Low	Medium	High
OEM warranty	1 year	5 years	5 years
Data integrity features	Desktop	Enterprise	Input/Output Error Detection Code (IOEDC); Input/Output Error Correction Code (IOECC)

The way we back up remote offices
Ford has a better storage idea
Quality Awards: Top NAS products
SAS puts squeeze on FC
vSphere and storage

the same disk media, SAS can run faster because of better error handling, better command queuing and dual porting."

In comparison to SATA drives in systems with SAS-to-SATA bridges, nearline SAS drives still fare approximately 5% better, according to Ian Williams, executive director, nearline storage solutions at Seagate. On the downside, nearline SAS drives are substantially more expensive than SATA drives; the cost has slowed adoption as many storage system vendors take a wait-and-see position. "At present, nearline SAS drives are 30% to 35% more expensive than equivalent SATA drives and, as a result, it is more cost effective for us to use a MUX adapter with SATA drives to achieve high availability," Promise Technology's Bahar said.

"At present, nearline SAS drives are 30% to 35% more expensive than equivalent SATA drives and, as a result, it is more cost effective for us to use a MUX adapter with SATA drives to achieve high availability."

—RAY BAHAR, vice president of the Americas, Promise Technology Inc.

Nearline SAS drives are likely to eventually force SATA drives out of higher-end enterprise storage systems where performance and features matter more than cost. But for lower-end arrays where cost is king, the price of nearline SAS drives needs to come down to that of SATA drives, and that won't happen until more drive manufacturers offer SAS drives. Furthermore, "the personal computing market would need to shift away from SATA drives and move to USB or other interface options," said Steve Gardner, director of outbound marketing at LSI Logic. "As long as the PC market maintains SATA drive volumes, there will always be some enterprise users integrating the technology to get the absolute lowest cost."

SAS-2 ENHANCEMENTS

While SAS has garnered a significant share of the disk drive market, the limitations of 3 Gbps SAS (SAS 1.1) kept it mostly in low-end to midrange systems. But the enhancements that SAS-2 (6 Gbps) bring to the table have catapulted SAS into a leadership position with the potential of becoming the single enterprise disk interface for high-performance tier 1, high-capacity tier 2 and tier 0 solid-state drives (SSDs). The key enhancements of SAS-2 are:

6 Gbps transfer rate. Doubling the transfer rate from 3 Gbps to 6 Gbps is the most obvious improvement, eclipsing Fibre Channel's 4 Gbps while pushing up the transfer rate in lockstep with SATA. Backward compatible with SAS 1.1, SAS-2 provides for a smooth transition.

Decision Feedback Equalization and longer cable length. Short cable length and the lack of an optical interconnect option hindered the deployment

The way we back up remote offices
Ford has a better storage idea
Quality Awards: Top NAS products
SAS puts squeeze on FC
vSphere and storage

of SAS drives in the data center. The use of a technique called Decision Feedback Equalization (DFE) to reduce interferences enables SAS-2 to extend cable length to 10 meters for improved rack-to-rack connectivity. It's an improvement over SAS 1.1, but it will take SAS 2.1 for copper cable lengths up to 20 meters and support for optical connections up to 100 meters. The SCSI Trade Association (STA) recently said the official release of SAS 2.1 should happen in early 2010. "We don't expect the new advanced connectivity options to have a significant part in the adoption of SAS for disk drives, although we do believe that the external storage market will benefit from the improved cabling options," Seagate's Williams said.

Standardized expander zoning. While some 3 Gbps SAS expanders have supported proprietary zoning, SAS-2 provides standardized zoning similar to the zoning capabilities in FC switches, enabling enhanced multi-host support and larger configurations. Zoning is the ability to limit access to drives to certain servers and connections to improve security and enable segregation within the storage topology.

Expander self-discovery. SAS-2 greatly improves the efficiency and scalability of SAS discovery. In SAS 1.1, discovery of SAS components was limited by the lack of standardized zoning and because end devices performed the discovery process. Combining standardized zoning and moving the discovery task from end devices to SAS expanders greatly reduces the time to discover large topologies.

"We don't expect the new advanced connectivity options to have a significant part in the adoption of SAS for disk drives, although we do believe that the external storage market will benefit from the improved cabling options."

—IAN WILLIAMS, executive director, nearline storage solutions, Seagate

SAS DESIGN CONSIDERATIONS

Unlike the shared link protocol of parallel SCSI, SAS and SATA are serial point-to-point protocols supporting connectivity between source and destination devices. To go beyond two devices, a SAS expander is required; akin to an FC switch, it provides a switching matrix to interconnect disk drives, controllers and other expanders. Available from LSI Logic and PMC-Sierra Inc. for both 3 Gbps and 6 Gbps SAS, 24- and 36-port SAS expanders are by far the most popular configurations and typically reside within data storage enclosures.

Almost since inception, so-called wide ports have been instrumental to SAS-based storage systems. SAS wide ports aggregate multiple ports—typically four—into a single "fat pipe" that allows multiple simultaneous connections to different destination devices. Wide ports are

The way we back up remote offices
Ford has a better storage idea
Quality Awards: Top NAS products
SAS puts squeeze on FC
vSphere and storage

used to connect expanders to storage controllers and other expanders. With 6 Gbps SAS, a typical four-port wide port provides a single 24 Gbps conduit. In addition to the performance benefit, wide ports are resilient against port failures, allowing them to function if at least one port within the assigned port group remains available. “A typical use of wide ports in a redundant array configuration is to connect one four-port wide port to one controller, another four-port wide port to a second controller and an additional four-port wide port to another expander, which leaves 12 ports for drives if a 24-port expander is used,” said Paul Vogt, senior director of product management at Xyratex Technology Ltd., a supplier of networked storage products for OEMs.

"A typical use of wide ports in a redundant array configuration is to connect one four-port wide port to one controller, another four-port wide port to a second controller and an additional four-port wide port to another expander, which leaves 12 ports for drives if a 24-port expander is used."

—PAUL VOGT, senior director of product management, Xyratex Technology Ltd.

To reap the benefits of 6 Gbps SAS, all SAS components within a storage system need to operate at 6 Gbps. Simply replacing 3 Gbps SAS drives with 6 Gbps drives won’t result in a performance boost unless expanders, controllers and HBAs all operate at 6 Gbps. To foster the adoption of 6 Gbps SAS, SAS-2 has provided for an optional multiplexing feature that allows a 6 Gbps controller to multiplex two 3 Gbps connections into a single 6 Gbps pipe. While the multiplexing features enabled SAS-2 controllers to communicate with 3 Gbps SAS drives at 6 Gbps speed, it has lost its significance with the advent of 6 Gbps SAS drives.

Even though 6 Gbps SAS is 50% faster than 4 Gbps Fibre Channel and twice as fast as SAS 1.1, most data storage systems won’t see that kind of performance boost. A storage system is only as fast as its slowest component, so having the interface speed doubled won’t change the performance limitation of the underlying disk media. “The fastest SAS hard drives will peak at about 200 MBps for pure sequential access and perform at tens of megabytes per second for random access where mechanical head movements are the limiting factor,” said Andy McNeill, distinguished engineer at IBM Corp. Depending on storage system architectures and their particular congestion points, the performance impact of SAS-2 will vary among storage systems. “We are seeing a big performance boost by moving front-end host ports from 3 Gbps to 6 Gbps SAS, but we are seeing little improvement by moving the back-end drives to 6 Gbps, even with 15K drives, and expect even less of an impact with

The way we back up remote offices
Ford has a better storage idea
Quality Awards: Top NAS products
SAS puts squeeze on FC
vSphere and storage

nearline SAS drives because the bottleneck is in the connection to the host not to the drives," explained Scott McClure, senior product manager at Dot Hill Systems Corp.

SAS ADOPTION LANDSCAPE

While SAS-2 components from disk drives and expanders to HBAs and RAID controllers have been available from vendors like Seagate, LSI Logic and PMC-Sierra since 2009, SAS-2 storage systems—with the exception of arrays from small vendors—aren't expected to be widely available until later in 2010. The main reason is a rigorous and time-consuming qualification process of the various SAS-2 pieces in enterprise-grade arrays to ensure that all components, firmware and software function properly. We'll see some entry-level and midsized storage systems transition to SAS-2 in 2010, but high-end systems based on SAS-2 will take much longer. "It will take years for high-end arrays to transition from FC to SAS because it's an extremely risk-averse space," said Greg Schulz, founder and senior analyst at Stillwater, MN-based StorageIO Group.

> "It will take years for high-end arrays to transition from FC to SAS because it's an extremely risk-averse space."
>
> —GREG SCHULZ, founder and senior analyst, StorageIO Group

Here are some data storage vendors' SAS-2 products and plans:

EMC. EMC currently uses 3 Gbps SAS in its Clariion AX4 and Celerra NX4 storage systems with no roadmap for transitioning these or other platforms to 6 Gbps SAS, according to an EMC spokesman.

Dell. Targeted at small- and medium-sized businesses (SMBs) and branch offices, the Dell PowerVault MD3000 modular array family currently has a 3 Gbps SAS back end. "Although no timeline has been set, it'll be transitioned to 6 Gbps SAS in the near future," Dell's Shoobe said.

Hewlett-Packard. With the HP StorageWorks D2000 JBOD connecting to an HP Smart Array controller with 6 Gbps SAS support, HP is one of a few vendors with a shipping end-to-end 6 Gbps storage product targeted for infrastructure applications like Microsoft Exchange and Web 2.0 service providers. The StorageWorks MSA20000 G2 with its current 3 Gbps SAS back end will be available with a 6 Gbps SAS back end in 2010, according to Kyle Fitze, marketing director, HP StorageWorks storage platforms. "As far as the midrange EVA family is concerned, we are looking at SAS as a disk back end, but no availability date has been set," he said.

Hitachi Data Systems (HDS). HDS has been using a 3 Gbps SAS back end in its Adaptable Modular Storage (AMS) 2000 family and will transition it to 6 Gbps SAS in the next technology refresh sometime within the next couple of years. Within the same timeframe, the AMS family and high-end Universal Storage Platform (USP) with its current FC disk back end

The way we back up remote offices
Ford has a better storage idea
Quality Awards: Top NAS products
SAS puts squeeze on FC
vSphere and storage

will share the same SAS-2 back-end arrays, according to Hubert Yoshida, HDS vice president and chief technology officer.

IBM. The IBM System Storage DS3000 series has been shipping with a 3 Gbps SAS back end and will be transitioned to 6 Gbps SAS by mid 2010, according to IBM's McNeill. "As far as IBM's higher-end storage systems are concerned, no timeline for transitioning the disk back end from FC to SAS has been set," he said.

LSI Logic. With SAS-to-SATA bridges, expanders, HBAs, I/O and RAID controllers all available for 6 Gbps SAS, LSI covers the SAS component spectrum, with the exception of disk drives. LSI has been offering arrays with a 3 Gbps SAS back end with its Engenio 1000 series for the OEM channel (e.g., IBM DS3000 and Sun StorageTek 2500 series), which it will transition to 6 Gbps SAS by mid 2010, according to Steve Fingerhut, LSI's senior director of marketing.

NetApp. The FAS2040 is currently the only NetApp controller shipping with a built-in SAS HBA. "We'll add 6 Gbps SAS support to other families as they get refreshed," said Sandra Wu, NetApp's director of solutions marketing. Aiding the transition to 6 Gbps SAS is the DS4243 chassis with Storage Bridge Bay (SBB) support introduced in mid 2009; it allows NetApp to adjust to various disk form factors and interfaces by simply changing SBB containers within the same enclosure.

Xyratex Technology. As a manufacturer of networked storage systems for the OEM channel (e.g., Dell and NetApp), Xyratex has been offering its OneStor family of storage systems with a 3 Gbps SAS back end and plans to offer a 6 Gbps SAS version in the first half of 2010.

SUMMARIZING SAS-2

With SAS-2, Fibre Channel as a disk interface seems doomed to disappear, but the transition will take some time. We'll see the adoption of 6 Gbps SAS in entry-level arrays pick up speed in 2010, with midrange and high-end arrays following at a more cautious pace. Although the FC disk interface is likely to expire, there's still plenty of life left in SATA, which should continue as the disk interface of choice in entry-level arrays for some time. ⊙

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

The way we back up remote offices
Ford has a better storage idea
Quality Awards: Top NAS products
SAS puts squeeze on FC
vSphere and storage

HP and IBM unseat NetApp

Hewlett-Packard and IBM prevail in our latest NAS survey, topping NetApp, which had won all three previous Quality Awards for midrange NAS systems and one for enterprise NAS products.

By Rich Castagna

AFTER DOMINATING the *Storage* magazine/SearchStorage.com Quality Awards for networked-attached storage (NAS) systems, NetApp has been nudged aside by Hewlett-Packard (HP) Co., among midrange vendors, and by IBM Corp., in the enterprise class, in the fourth cycle of our Quality Awards program. Respondents to the service and reliability survey made HP's midrange NAS product line the clear winner by ranking it No. 1 in all five rating categories and giving HP a comfortable margin over second-place finisher Dell Inc. Respondents' comments about the two winners were succinct and enthusiastic; an HP user simply remarked "Great!" and one very satisfied IBM customer raved, "IBM rocks!"

The way we back up remote offices
Ford has a better storage idea
Quality Awards: Top NAS products
SAS puts squeeze on FC
vSphere and storage

For enterprise-class NAS systems, NetApp nearly repeated its Quality Awards III win but was upset by IBM by the narrowest margin we've ever seen in the NAS evaluations—an extremely slim 0.04 points. But NetApp should find some consolation in IBM's close victory as IBM resells NetApp's NAS products. Because the Quality Awards seeks feedback on a range of activities that affect users' experiences—not just the speeds, feeds and features of the product—it's not uncommon for two vendors selling essentially the same product to end up being ranked very differently.

HP's and IBM's wins mark firsts for both companies, as neither one has ever come out on top on any of our NAS Quality Awards surveys or in our other two storage system categories: enterprise arrays and midrange arrays.

HP's across-the-board sweep in the midrange NAS field isn't unprecedented; NetApp has done it twice and narrowly missed doing it a third time by a few hundredths of a point. In edging out NetApp among enterprise NAS vendors, IBM's win was much like a split decision in a prize fight. It finished first in two of the five rating categories, as did NetApp; HP won in the fifth category. But in the two categories where NetApp came out on top, IBM scored well and was a very close second.

IT STARTS WITH SALES

Our survey respondents delivered high praise to HP's midrange NAS sales efforts, giving the company the highest score for all six statements in the sales-force competence category. For the statement "The vendor's sales support team is knowledgeable," HP garnered a 6.52 rating, just barely ahead of Dell's 6.48. A key factor in making a system purchase that meets your firm's requirements and fits into its environment is a sales rep who takes the time to learn about your business and configure a solution appropriately. For the statement "My sales rep understands my business," HP eked out another narrow win with a 6.34 rating vs. IBM's 6.32. And for "My sales rep is knowledgeable about my industry," HP

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. The results are derived from a survey of qualified *Storage*/SearchStorage.com readers who assessed products in five main categories: sales-force competence, product features, initial product quality, product reliability and technical support. Products are rated on a 1.00-8.00 scale, where 8.00 is the most favorable score.

In this, the fourth edition of the NAS Quality Awards, 340 respondents rated 626 products/product lines.

The way we back up remote offices
Ford has a better storage idea
Quality Awards: Top NAS products
SAS puts squeeze on FC
vSphere and storage

scored a similarly close top mark, besting EMC 6.31 to 6.29.

Although Jeff McMorran, director of IT at NMG, a purchasing services company in Newmarket, Ontario, bought his HP products through a reseller, he appreciated HP's participation in the process. "We had one account manager who was excellent," McMorran said. "He'd bring people from HP into it."

IBM's overall top ranking among enterprise NAS competitors can be

PRODUCTS IN THE SURVEY

The following products were included in the Quality Awards for NAS IV survey (the number of ratings responses are in parentheses).

MIDRANGE NAS

- BlueArc Corp. Titan 1100, Mercury 50/100*
- Dell Inc. PowerVault NF100/NF500/NF600, NX300/NX1950/NX3000/NX4 (60)
- EMC Corp. Celerra NX4, NS-120 (53)
- Hewlett-Packard (HP) Co. ProLiant DL Storage Servers, StorageWorks Network Storage System X1000/X3000 (63)
- Hitachi Data Systems Essential NAS Platform 1000 Series (13)
- IBM Corp. N3000/N5000 (38)
- LSI Corp. (formerly ONStor) Cougar 3000 NAS Gateway*
- NetApp FAS200/FAS900, FAS2000 (all with NAS interface) (69)
- Overland Storage Inc. (formerly Adaptec) Snap Server 620/650*
- Panasas Inc. ActiveStor 7 Series/ 8 Series*
- Silicon Graphics Inc. (SGI) InfiniteStorage NAS2000/NAS3000, NEXIS 2000/NEXIS 7000*
- Sun Microsystems Inc. StorEdge 5210/5220, Sun Storage 7110/7210/7310 Unified Storage System (28)

ENTERPRISE NAS

- BlueArc Titan 2000/3000 series*
- DataDirect Networks Inc. GRIDScaler/EXAScaler/xSTREAMScaler*
- EMC Celerra NS-480/NS-960 (58)
- Exanet Inc. EX1500*
- Hewlett-Packard StorageWorks EFS Clustered Gateway, 4400 Scalable NAS File Services, 9100 Extreme Data Storage System (30)
- Hitachi Data Systems Lightning Blade, TagmaStore Blade, AMS1000/500/200, WMS100 with NAS Option, Hitachi NAS Platform 3000 Series (27)
- IBM N6000/N7000 (31)
- Isilon Systems Inc. Isilon IQ*
- LSI (formerly ONStor) Cougar 6000 Series NAS Gateway*
- NetApp FAS3000/FAS3100, FAS6000 (all with NAS interface) (83)
- Panasas ActiveStor 9 Series*
- SGI InfiniteStorage NEXIS 9000*
- Sun StorEdge 5310/5320, Sun Storage 7420 (20)

*Too few responses to be included in final tabulations; at least 13 complete sets of responses are required.

credited to its strong showing in the sales-competence category. Big Blue led the pack on four of the six statements, trailing only Hitachi Data Systems on "My sales rep is knowledgeable about my industry" by the slimmest possible margin (6.27 to 6.26) and NetApp on "The vendor's sales support team is knowledgeable" by slightly more than a tenth of a point. But IBM came out on top for statements typically related to confidence building during the sales process, such as "My sales rep keeps my interests foremost" and "My sales rep understands my business."

For Mike MacNeill, director of architecture and infrastructure at Cross Country Healthcare Inc. in Boca Raton, Fla., a good relationship with a solid reseller helped him navigate the NAS waters when he was negotiating with both Dell/EMC Corp. and IBM for a system. "We had a really good rep who would kind of help us behind the scenes," MacNeill said. They eventually settled on the IBM gear.

PLUG AND PLAY, OR PLUG AND PRAY?

The old saying "You never get a second chance to make a first impression" holds true for network-attached storage purchases, too, although our enterprise NAS respondents seem to be a little more forgiving. Initial installation and power-up experiences often shape overall perceptions about the quality of a product. Apparently, HP's midrange NAS systems made a very good first impression among our survey respondents. HP didn't quite match its dominating performance in the sales-competence category, but it came awfully close for initial product quality, snagging the top ratings for five of the six statements, although for one—"This product was installed without any defects"—Dell matched HP's leading score of 6.80. Dell also kept HP from another sweep with a 6.68 rating for "This product was easy to get up and running" compared with a nearly-as-strong 6.61 for HP.

"[HP's NAS box is] very easy to set up, and it's very easy to maintain. It came with everything it needed."

—JEFF McMORRAN, director of IT, NMG

HP's 6.83 for "This product is easy to use" was the highest rating for a single statement among all products (midrange or enterprise) in the survey. "It's very easy to set up, and it's very easy to maintain," NMG's McMorran said of their HP NAS box. "It came with everything it needed."

Among enterprise NAS systems, NetApp prevailed with a 6.35 in the initial product quality category with top scores for three of the six statements. IBM continued to duke it out with NetApp and finished a close second with a 6.30 rating. Hitachi, which scored the highest for "This product was easy to get up and running" (6.37), wasn't far behind the two leaders with a category rating of 6.26. Respondents' comments bolstered Hitachi's strong showing, with one indicating that expectations had been met: "Works as they claim."

MIDRANGE NAS

- Dell PowerVault NF100/NF500/NF600, NX300/NX1950/NX3000/NX4
- EMC Celerra NX4, NS-120
- Hewlett-Packard (HP) ProLiant DL Storage Servers, StorageWorks Network Storage System X1000/X3000 (63)
- Hitachi Data Systems Essential NAS Platform 1000 Series
- IBM N3000/N5000
- NetApp FAS200/FAS900, FAS2000 (all with NAS interface)
- Sun Microsystems StorEdge 5210/5220, Sun Storage 7110/7210/7310 Unified Storage System

PRODUCT FEATURES

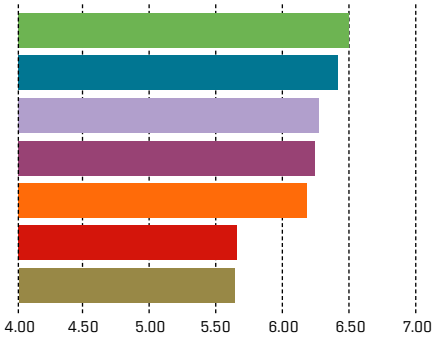

OVERALL RANKINGS

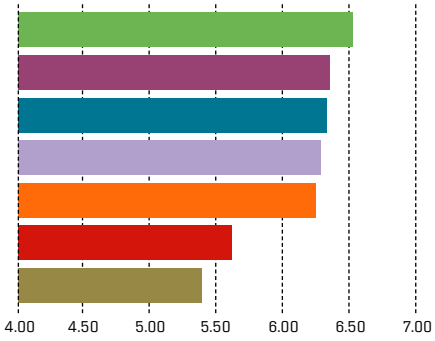

PRODUCT RELIABILITY

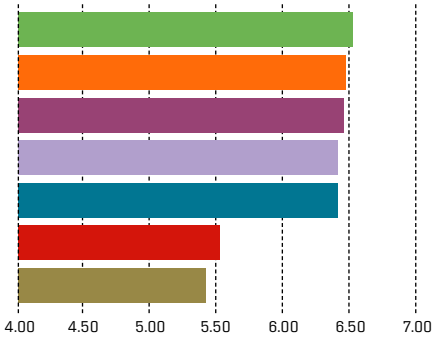

SALES-FORCE COMPETENCE

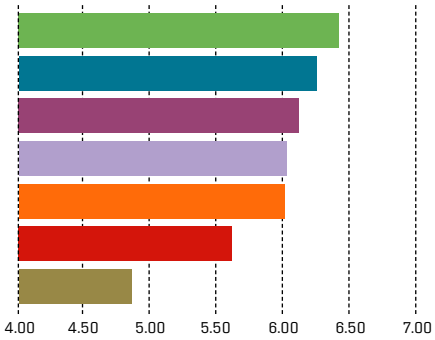

TECHNICAL SUPPORT

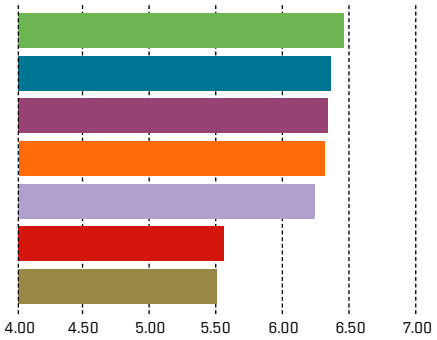

INITIAL PRODUCT QUALITY

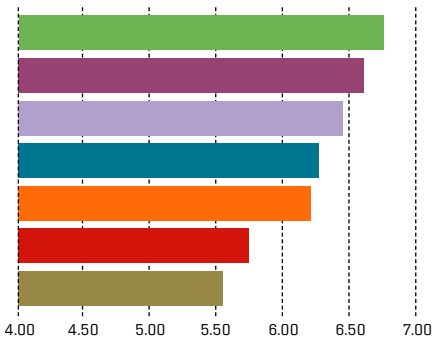

WOULD YOU BUY AGAIN?

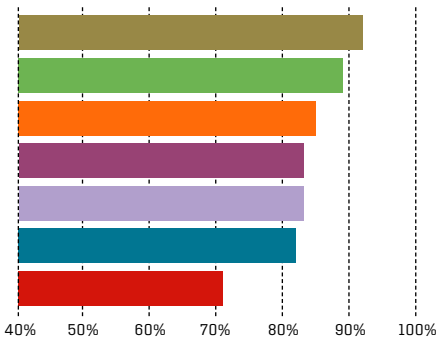

Based on a 1.00-8.00 scoring scale

The way we back up remote offices
Ford has a better storage idea
Quality Awards: Top NAS products
SAS puts squeeze on FC
vSphere and storage

Even though IBM offers "to rack it for you and do the initial setup," Cross Country Healthcare's MacNeill eschewed the help. "The actual setup once you power it on is 15 or 20 minutes," he said. "We did it ourselves this time because it's that easy."

PRODUCT FEATURES

A good sales experience and an uneventful installation are certainly important, but the rubber truly hits the road when it comes to a product's feature set. HP's 6.50 for this category in the midrange group was highlighted by a 6.62 for "This product's capacity scales to meet my needs" and nearly the same score (6.61) for "The product's management features meet my needs." Overall, HP was top dog on four feature statements, with NetApp and IBM taking the other three. NetApp's extensive snapshotting capabilities were acknowledged with the highest rating (6.45) among the midrange group for "This product's snapshot features meet my needs." IBM's two leading scores came on data protection-related statements: "This product's mirroring features meet my needs" (6.48) and "This product's remote replication features meet my needs" (6.56).

A good sales experience and an uneventful installation are certainly important, but the rubber truly hits the road when it comes to a product's feature set.

Continuing their duel in the enterprise group, NetApp's 6.41 for features just bested IBM, which received a 6.39 rating in the product feature category. HP finished a strong third by scoring highest on three of the category's seven statements for a total of 6.34. NetApp prevailed on the strength of a 6.66 for "This product's snapshot features meet my needs" and a 6.59 rating on the general satisfaction statement "Overall, this product's features meet my needs." "Very satisfied with the product," noted one respondent, who, with an eye to an upgrade, added, "Looking forward to the new features from Ontap 8.x."

IBM's good showing in the product features category hinged on top scores for scalability (6.47) and interoperability (6.33). For Cross Country Healthcare's MacNeill, manageability was an important factor in the product decision process, and the IBM/NetApp combination was a winner: "With the IBM/NetApp, we've seen a huge increase in productivity just using the management interface."

RELIABILITY RATINGS

Among midrange products, HP finished on top for the reliability statements with an excellent 6.53, but EMC (6.48) and Dell (6.46) were hard on its heels, and IBM and NetApp weren't far behind as both posted scores of 6.40. In the most closely contested midrange rating category, EMC and IBM each ranked highest on one statement, with EMC picking

ENTERPRISE NAS

- EMC Celerra NS-480/NS-960 (58)
- Hewlett-Packard StorageWorks EFS Clustered Gateway, 4400 Scalable NAS File Services, 9100 Extreme Data Storage System
- Hitachi Data Systems Lightning Blade, TagmaStore Blade, AMS1000/500/200, WMS100 with NAS Option, Hitachi NAS Platform 3000 Series
- IBM N6000/N7000
- NetApp FAS3000/FAS3100, FAS6000 (all with NAS interface)
- Sun StorEdge 5310/5320, Sun Storage 7420

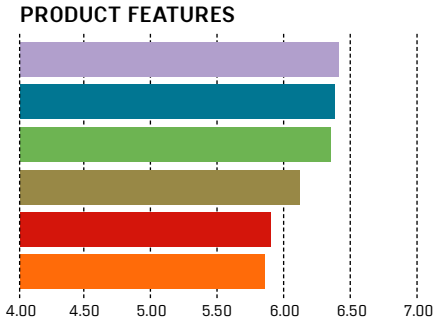

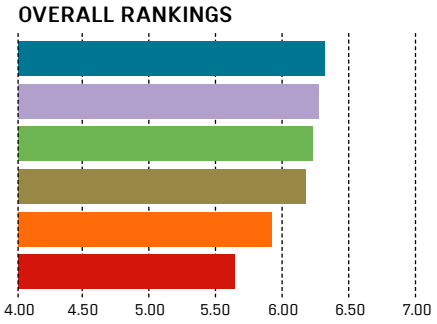

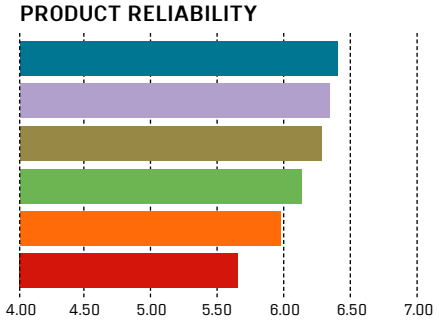

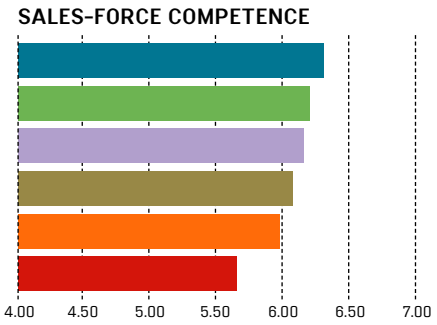

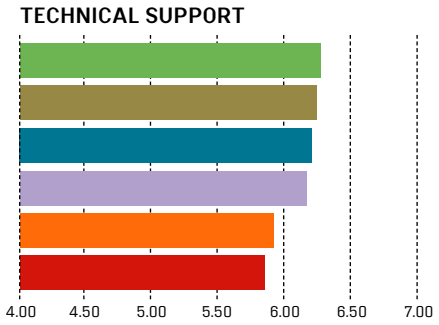

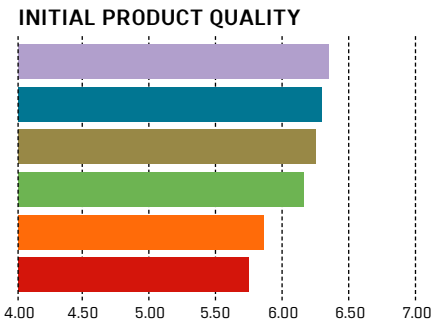

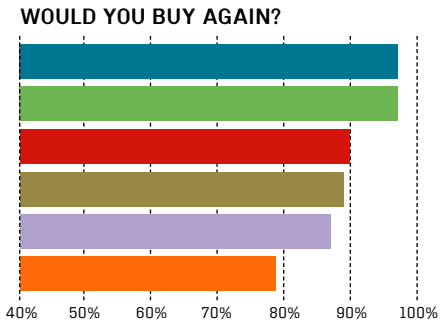

Based on a 1.00-8.00 scoring scale

The way we back up remote offices
Ford has a better storage idea
Quality Awards: Top NAS products
SAS puts squeeze on FC
vSphere and storage

up a 6.70 for "The product meets my service-level requirement" and IBM coming out on top for "Vendor provides comprehensive upgrade guidance." Still, despite the narrow range in total scores for this category, HP led for the other three statements, which relate to limited downtime, unplanned patching and being able to apply patches without disrupting operations.

On the enterprise side, IBM and NetApp once again went toe-to-toe, with each scoring highest for two statements in the reliability rankings and finishing in a dead-heat tie for the fifth category statement. Overall, IBM's 6.40 rating outpaced NetApp (6.34) and third-place finisher Hitachi (6.29). IBM showed its strength with a 6.55 for the key reliability statement "This product experiences very little downtime," while NetApp's best showing was a 6.53 for meeting service-level requirements.

THE MEASURE OF SUPPORT

HP completed its grand slam with a 6.48 rating in the technical support category among midrange products with a relatively comfortable margin over a cluster of contenders, including IBM (6.37), Dell (6.34) and EMC (6.31). Of the eight statements in this evaluation category, HP notched the highest scores on five and was bested by EMC by a hair (6.48 vs. 6.45) for the support statement "Vendor takes ownership of the problem."

HP completed its grand slam with a 6.48 rating in the technical support category among midrange products with a relatively comfortable margin over a cluster of contenders, including IBM (6.37), Dell (6.34) and EMC (6.31).

NetApp scored high for meeting support expectations with a 6.72 for "Vendor supplies support as contractually specified," and Dell stood out for sending smart techs to solve its customers' problems with a 6.53 for "Support personnel are knowledgeable."

Among enterprise candidates, the NetApp-IBM brawl was overshadowed in the support category as HP nosed out Hitachi, 6.28 to 6.25, for the top spot. But only 0.08 points separated the top four, with IBM (6.23) and NetApp (6.20) rounding out a very closely bunched group. HP's strengths were for training and documentation, while Hitachi received top marks for taking ownership of problems, timely problem resolution, knowledgeable personnel and delivering as promised.

REPEAT CUSTOMERS

In every Quality Awards survey we ask respondents whether they would buy the same product again today, having had some experience with it and knowing what they know now. Often, the "buy again" rankings sync up pretty closely with the overall ratings, but just as frequently, they

The way we back up remote offices
Ford has a better storage idea
Quality Awards: Top NAS products
SAS puts squeeze on FC
vSphere and storage

seem to run counter to the general scores. When the latter occurs, it's assumed that users, having worked through their ups and downs with a product, are more comfortable sticking with what's now familiar. That appears to be the case with the latest NAS midrange tallies: Hitachi, lowest in the overall rankings, nonetheless has apparently built customer loyalty, as 92% said they would consider buying the Hitachi product again. For enterprise NAS systems, the "buy again" numbers followed more closely to form, with 97% of both IBM's and HP's users saying they'd do it all over again.

The repeat-buy numbers for enterprise NAS systems were particularly high across the board, indicating that these vendors are ultimately delivering when it comes to customer satisfaction. With so much focus lately on the rapid growth of file-based storage, it's a promising sign that NAS system vendors are doing a good job of meeting users' requirements and expectations. ⊙

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

The way we back up remote offices
Ford has a better storage idea
Quality Awards: Top NAS products
SAS puts squeeze on FC
vSphere and storage

Ford's storage ledger balances capacity decisions

Data storage shops often struggle with anticipating new capacity requirements and ensuring that business needs can be met. Ford Motor Company took a unique approach to the problem and made capacity planning as simple as balancing your checkbook.

By Thomas Woods

A FUNNY THING HAPPENED to the Ford Motor Company enterprise storage team: Storage Sunday maintenance post-mortem review meetings became boring. The meetings were no longer spiced up with horror stories about outages caused by maintenance action failures. Thanks to the enterprise virtual storage program, many high-risk and high-impact maintenance actions had been mitigated by the ability to transparently, and without server interruption, move critical virtual logical unit numbers (LUNs) from physical storage arrays before scheduled maintenance took place.

Aside from the woeful exploits of the Detroit Lions, Sundays would have been perfect. For those attending the weekly enterprise data storage infrastructure capacity forecasting meetings, the stress-reducing benefits of virtual storage were less evident.

The way we back up remote offices
Ford has a better storage idea
Quality Awards: Top NAS products
SAS puts squeeze on FC
vSphere and storage

Ford's enterprise storage architecture progressed in stages, from attaching servers directly to storage arrays, to attaching servers to storage-area networks (SANs) to physical storage arrays, to the current virtualized environment where SAN-attached servers link to virtualized physical storage arrays.

This transition occurred even as IT budgets were under intense scrutiny, with the requirement that new infrastructure purchases be directly traceable to end-user requests. But in the virtual world, the correlation of server to storage array is no longer a direct line. Adding to the complexity was a requirement that overhead and white space (capacity that's allocated but not used) had to be tracked and explained.

This transition occurred even as IT budgets were under intense scrutiny, with the requirement that new infrastructure purchases be directly traceable to end-user requests.

VIRTUALIZED STORAGE RAISES NEW ISSUES

After the introduction of virtual storage into the Ford environment, the data storage capacity forecasting meetings at the company reflected the new reality: All stakeholders had a piece of the action, but the teams lacked the tools, methods and approach to assemble a single consolidated view they could effectively communicate to the business side of the house. Some of the missing pieces were:

- The enterprise storage teams, specifically storage operations and storage engineering, didn't have an effective way to communicate with finance and purchasing.
- Capacity planning couldn't translate forecasts into storage capacity requirements.
- Finance and purchasing didn't fully grasp the multidimensional nature of virtual storage environments.

The frustration in the capacity planning meetings soon exacerbated long-standing trust issues among operations, engineering, sales and finance.

Because the communication problem crossed multiple team boundaries, a cross-functional team was created to address the issues, with representatives from all interested parties—finance, storage operations, business management, capacity planning and storage engineering. The team reached a consensus that a systematic, repeatable and traceable method to track and forecast network-attached storage (NAS), SAN and backup infrastructure was needed. This method would have to facilitate

The way we back up remote offices
Ford has a better storage idea
Quality Awards: Top NAS products
SAS puts squeeze on FC
vSphere and storage

understanding for all storage capacity stakeholders by providing the following capabilities:

- **Infrastructure ordering personnel** would be able to place customer orders without being required to be "storage Gurus."
- **Storage administrators** would be able to use customer orders to create forecasts against the key physical and virtual storage components, while meeting customer requirements and business constraints, in a repeatable fashion.
- **Business managers** would be able to track customer demand to proposed infrastructure projects.
- **Storage managers** would be able to justify non-customer infrastructure capacity increase requirements, such as temporary storage needed for a data migration or to accommodate the organic growth of a specific storage subsystem.
- **Storage engineers** would need to verify that storage infrastructure is deployed consistently and adheres to engineering standards.

A cross-functional team was created to address the issues, with representatives from finance, storage operations, business management, capacity planning and storage engineering.

THE STORAGE MANAGEMENT SOLUTION

The cross-functional team implemented a storage ledger, or storage "checkbook" with multiple "currencies" related to key storage consumables. For an ordinary checkbook, money is the only consumable, but a storage ledger checkbook has many consumables. The team decided to create three separate ledgers to track all consumables related to three major storage activities: NAS, SAN and backup. A single ledger was considered as it would have reduced the work for the groups that order storage, but the ledger would have been more complicated and considerably larger.

HOW A STORAGE LEDGER WORKS

Each ledger consists of two types of spreadsheets: a main ledger sheet and the component aggregation sheet used to calculate the top line (beginning balance) for each storage consumable that needs to be tracked. The main ledger sheet is divided into two sections. The first includes ordering columns where customers of storage services can write post-dated storage capacity checks. The information entered in the ordering columns drives the entries placed in the second set of columns, the storage capacity columns. The storage capacity columns aren't visible to the customers writing checks. The storage capacity columns include information needed by data storage administrators to map requests to a set of storage consumables associated with the appropriate environment and technology.

The way we back up remote offices
Ford has a better storage idea
Quality Awards: Top NAS products
SAS puts squeeze on FC
vSphere and storage

The ledger process consists of the following steps:

Step 1. Update the main ledger sheet top-line balances that are on the capacity side of the main ledger sheet. The top-line balance is calculated by adding all of the component capacities for the specific subsystems. The specific subsystems are tracked in separate tracking sheets and map to top-line ledger balances as the enterprise architecture dictates. The capacity columns are divided not only by technology, but also by other factors such as location or functional environment.

Step 2. Normalize the top-line balance for percent overhead. Many storage systems require a certain amount of overhead to run efficiently and can't be managed at 100% capacity. The percent overhead should be determined so that it will indicate the point at which new infrastructure is needed. The component aggregate sheet allows a business analyst to reference the source data if there's a concern that the percent overheads are too generous.

Step 3. Reconcile the ledger entries. Forecasted storage actions that have been executed or cancelled should be marked as completed and taken off the active forecasting ledger.

Step 4. Analyze the updated ledger. After the ledger has been updated, a storage infrastructure capacity forecasting meeting is held. Using the ledger's graphical output, the team determines what actions, if any, are required to address forecasted activities. The team works together to consider different scenarios, such as the effects of:

- Extending leases vs. buyouts of storage infrastructure
- Placing new or existing data on different storage performance tiers
- Backup and disaster recovery (DR) scenarios
- NAS vs. SAN scenarios

After the analysis and scenario assessments, the team recommends a best course of action to management. Because the ledger process has been initiated, the management review process has been considerably more collaborative because management is now provided with a set of high-quality and transparent options.

Step 5. Open ledger for general use. The ledger, like your checkbook, is to be used throughout the week. Steps 1 through 4 are intended for periodic baselining and reconciliation activities.

Step 6. Begin reconciliation. Close ledger for general use and return to Step 1.

SAN STORAGE LEDGER

The SAN storage ledger ordering columns are as follows:

- Storage performance level
- Type of storage performance
- Type of replication (none, local or remote)
- Environment
- Capacity requested
- Forecast date
- Project name

The way we back up remote offices
Ford has a better storage idea
Quality Awards: Top NAS products
SAS puts squeeze on FC
vSphere and storage

The SAN storage ledger storage capacity columns are determined by:

- Environment (data center, virtual storage array environment, major business unit)
- Disk performance and type (high speed, SATA, etc.)
- Storage types (mirrored data center to data center, instant copy using local snapshots of data)

Besides storage end users, data storage administrators also place entries into the storage ledger. SAN storage admins may create ledger entries to forecast storage infrastructure needs because of lease expirations, to reserve temporary staging capacity needed for storage maintenance, and to document the impact of future storage realignments such as the effect of transferring data between storage tiers or technologies (e.g., SAN to NAS).

NAS and backup administrators may also write checks against the SAN storage environment. A NAS administrator is required to create a SAN ledger entry to track removal of storage from a NAS gateway that's slated to migrate to NAS Fibre-attached storage (FAS). A backup admin might write a check to request more disk pool storage for online disk backups.

Date	Action	Storage Type	Capacity (GB)	Regular Storage	Performance Storage	Economy Storage	Regular Storage Total	Performance Storage Total	Economy Storage Total
2009-12-01				18895	5500	11630	18,895	5,500	11,630
2009-12-15	Storage Host Add	Regular	200	(200.00)			18,695	5,500	11,630
2009-12-15	Storage Host Add	Performance	333		(333.00)		18,695	5,167	11,630
2009-12-15	Host Return Storage	Performance	-20		20.00		18,695	5,187	11,630
2009-12-15	Storage Host Add	Regular	916	(916.00)			17,779	5,187	11,630
2009-12-15	Storage Host Add	Regular	1831	(1831.00)			15,948	5,187	11,630
2009-12-20	Storage Host Add	Economy	37			(37.00)	15,948	5,187	11,593
2009-12-20	Storage Host Add	Performance	262		(262.00)		15,948	4,925	11,593
2009-12-20	Storage Host Add	Performance	146		(146.00)		15,948	4,779	11,593
2009-12-20	Storage Host Add	Performance	131		(131.00)		15,948	4,648	11,593
2009-12-20	Storage Host Add	Performance	262		(262.00)		15,948	4,386	11,593
2009-12-20	Storage Host Add	Performance	262		(262.00)		15,948	4,124	11,593
2009-12-20	Storage Host Add	Economy	37			(37.00)	15,948	4,124	11,556
2009-12-20	Host Return Storage	Economy	-100			100.00	15,948	4,124	11,656
2010-01-10	Storage Host Add	Economy	37			(37.00)	15,948	4,124	11,619
2010-01-10	Storage Host Add	Economy	37			(37.00)	15,948	4,124	11,582
2010-01-10	Host Return Storage	Economy	-193			193.00	15,948	4,124	11,775
2010-01-10	Storage Host Add	Economy	37			(37.00)	15,948	4,124	11,738
2010-01-10	Storage Host Add	Economy	37			(37.00)	15,948	4,124	11,701
2010-01-11	Add Temp Swap	Economy	-10000			10000.00	15,948	4,124	21,701
2010-01-10	Storage Host Add	Economy	109			(109.00)	15,948	4,124	21,592
2010-01-10	Storage Host Add	Regular	342	(342.00)			15,606	4,124	21,592
2010-01-10	Storage Host Add	Regular	512	(512.00)			15,094	4,124	21,592

This simplified example of the ledger process that Ford uses for provisioning storage capacity shows provisioning activity for three general types of storage: economy, regular and performance. The first row shows the beginning "balances" or available capacity. When a group makes a request for capacity and it's allocated, the amount of storage is entered on the ledger with the balance adjusted accordingly. The Action column describes the transaction: Storage Host Add means a user is returning unused capacity to the pool; Add Temp Swap indicates an allocation for temporary capacity for a migration or similar activity; and Planned New Purchased Storage lists any capacity that's added to one of the tier pools. The actual ledger that Ford uses is far more detailed than this sample.

The way we back up remote offices
Ford has a better storage idea
Quality Awards: Top NAS products
SAS puts squeeze on FC
vSphere and storage

NAS STORAGE LEDGER

When the NAS storage ledger process started, the NAS environment at Ford Motor Company was more complicated than it is today. The environment consisted of multiple vendors with both gateway-based NAS (NAS devices connected to external storage-area networks administered by the SAN team) and FAS-based network-attached storage (NAS gateway heads built into OEM vendor-provided Fibre-attached storage arrays).

For gateway-based NAS, the NAS administrators are required to write SAN storage ledger checks to increase or decrease NAS gateway SAN capacity. The old system used tape-based Network Data Management Protocol (NDMP), which meant entries had to be made in the backup ledger. Ford has been migrating to a pure FAS-based NAS environment with a non-tape backup infrastructure supplied by a single vendor. In addition, only a single performance type of NAS storage is now offered. As a result, the NAS ledger has become simpler to reflect the less-complicated environment.

From a customer perspective, the NAS entries are the amount of data needed and the location of primary storage; the capacity columns track the amount of primary and backup mirrored storage requested. When there are a lot of small entries for the NAS ledger, they're bundled together to reflect a forecastable amount. NAS team entries consist of mainly major projects, a bundle of small projects for a specific month and organic growth predictions. NAS file systems that aren't protected by file system size quotas are the main causes of organic growth.

Ford uses a progressive incremental backup approach, meaning that after an initial full backup only changed files are backed up.

BACKUP STORAGE LEDGER

The backup ledger is the most complicated ledger; Ford Motor Company is currently in the process of rolling out this ledger. Backup capacity isn't as straightforward as NAS or SAN capacity. For NAS and SAN capacity, the actual amount of storage is the dominant key performance indicator (KPI). But for backup, the amount of storage is just one of many criteria that must be tracked. Ford uses a progressive incremental backup approach, meaning that after an initial full backup only changed files are backed up. The team couldn't create a forecasting ledger without first understanding the backup system design criteria. The first step in creating a backup forecast strategy is to develop criteria to address design tradeoffs, otherwise the backup systems could inadvertently slip into sub-optimal performance. The team developed the following design criteria:

The way we back up remote offices
Ford has a better storage idea
Quality Awards: Top NAS products
SAS puts squeeze on FC
vSphere and storage

- Daily backup disk pools are large enough to contain a single day's worth of file system incremental backup.
- Weekly backup disk pools are large enough to contain one week's worth of backup data before it's pushed to tape.
- Database size shouldn't exceed 60 GB (to ensure that DR recovery targets can be met).
- Data migration infrastructure must be able to dump disk pools to tape within four hours.
- Server backups will be co-located to dedicated tapes.
- Workstation backups will be co-located to dedicated tape groups.

For the daily incremental disk pool, a 14-day high-water mark of the disk pool average is derived and then averaged with all of the servers in that same environment. For example, assume there are 10 backup servers at a data center, each disk with 500 GB of disk pool for an aggregate disk pool of 5,000 GB. The 14-day average of each of the 10 servers is 350 GB (aggregate of 3,500 GB). The aggregate high-water mark backup capacity shouldn't go above 90%, so additional disk pool capacity should be provided before the 90% threshold is reached. Therefore the capacity is 5,000 GB times 90% minus 3,500 GB or 4,500 GB minus 3,500 GB = 1,000 GB of available aggregate disk pool.

It should be noted that there's a difference between capacity forecasting and day-to-day tuning and monitoring. Capacity forecasting is used primarily to determine future infrastructure capacity requirements; efficiently balancing current infrastructure should be a normal part of a storage admin's tasks.

THE LEDGER'S BOTTOM LINE

Ford Motor Company's data storage forecasting and purchasing meetings are now more effective and collaborative, which helps ensure that the right amount of infrastructure is delivered at the right time. Storage forecasting capacity meetings are now data-driven and leverage dynamic what-if scenarios that can be created instantly. If there are questions on source data, stakeholders can quickly view ledger or spreadsheet data for more information. The ledger process also allows Ford and its IT service providers to work together better to ensure that capacity is tracked and forecasted appropriately. Any organization can implement new technology, but what truly makes a difference is how the organization adapts its actions to maximize the return on the investment.

Tom Woods is currently global ITIL services transition manager at Ford Motor Company. At Ford, Tom has held storage operations, engineering and architecture positions, and has supervised the backup and NAS teams.

Cloud storage: Just another storage tier?

If you dig a little deeper, you'll find cloud storage has been adopted in some sectors as a data archive tier, and has been delivering cost-saving benefits for quite some time.

VERY FEW YEARS a technology comes along that gets hyped to the point that no one is truly sure what it means anymore. In the storage world, storage virtualization and information lifecycle management (ILM) are two of those technologies that received a lot of vendor buildup and held a lot of promise for reducing the cost and complexity of the storage infrastructure but never really played out. The latest and greatest technology to hit the peak of the hype cycle is cloud storage.

Data storage vendors aren't doing themselves any favors with all the hype and cloudy (forgive the pun) long-term vision around cloud storage, but it doesn't appear to be headed down the same road as ILM and storage virtualization. Unlike with ILM and the original vision (2001-2002 timeframe) of storage virtualization, there are some real and immediate benefits that can be derived for adopting a cloud storage tier:

IT managers are paranoid about what could happen to corporate information once it leaves the boundaries of corporate IT, which is a good trait in an IT manager.

- Subscription-based services that roll up as an operating expense.
- Reduced upfront and ongoing storage costs and capacity.
- Near-perfect economics, with the subscriber paying only for the capacity that's used.
- Reduced operating expenses because the service provider can manage the infrastructure more efficiently thanks to multitenancy and scale-out virtualized storage platforms.
- Improved business flexibility, including easy deployment of temporary capacity.

With all these benefits, why haven't users jumped on the cloud bandwagon en masse? Well, there are various concerns about information security in the cloud and some confusion as to what constitutes cloud storage. IT managers are paranoid about what could happen to corporate information once it leaves the boundaries of corporate IT,

The way we back up remote offices
Ford has a better storage idea
Quality Awards: Top NAS products
SAS puts squeeze on FC
vSphere and storage

The way we back up remote offices
Ford has a better storage idea
Quality Awards: Top NAS products
SAS puts squeeze on FC
vSphere and storage

which is a good trait in an IT manager. However, if you dig a little deeper you'll find cloud storage has been adopted in some sectors as a data archive tier, and has been delivering those cost-saving benefits for quite some time.

EFFECTIVE DATA ARCHIVE TIER

Cloud storage makes sense for data that needs to be retained for long periods of time, is shared by multiple users, needs to be easily accessed and requires latency tolerance, which perfectly describes the requirements for digital medical images and healthcare records. Add to those requirements the huge explosion of semistructured data in healthcare—driven by electronic medical records and advances in content capture devices that create ever-denser images—and it becomes clear why cloud storage has become a viable data archive tier for the healthcare industry.

Deploying cloud storage archiving within a private storage cloud mitigates security concerns, as the storage service isn't shared outside the healthcare network.

There are numerous examples of regional healthcare systems sharing a centralized digital image archive to store images such as CT scans and x-rays. These shared archives can be considered private storage clouds in which archive storage services are offered to healthcare network members as a way to cost-effectively store images and medical records accessed over IP networks. Deploying cloud storage archiving within a private storage cloud mitigates security concerns, as the storage service isn't shared outside the healthcare network.

The massive growth of unstructured data isn't just a healthcare problem; Enterprise Strategy Group estimates that unstructured data will make up the vast majority of data in commercial data centers by 2012, which affects companies of all sizes in every industry. With the advent of Web 2.0 applications, increasing amounts of regulatory oversight over how data is stored and for how long, as well as growth in the use of rich media, commercial enterprises will see unstructured data growth rates that exceed anything they've experienced in the past. Most unstructured data is frequently accessed only for the first two weeks after it's created; access then slowly tapers off and stops, yet it needs to be retained for long periods of time for business or regulatory reasons. This data is a good near-term candidate for being moved off the expensive tier 1 or tier 2 system it was created and onto a public cloud-based archive tier.

While there's a huge potential for saving on both CAPEX and OPEX

The way we back up remote offices
Ford has a better storage idea
Quality Awards: Top NAS products
SAS puts squeeze on FC
vSphere and storage

fronts by moving long-tail data into the cloud, IT organizations considering a public cloud archiving option need to go into any cloud storage service provider evaluation with their eyes wide open. There's a high degree of variance in the transparency provided by cloud archive service providers into how data is stored, protected and secured, as well as the user's ability to audit access and authenticity to ensure regulatory compliance.

So the technology is here, but you need to be cautious. The Internet has reached every corner of the world, effectively creating a flat global network with few, if any, barriers to connectivity. The combination of wide-area network (WAN) acceleration and ubiquitous network connectivity allows business to be conducted anywhere. On the platform front, scale-out, commodity-based platforms that provide massive scalability, parallel data transfers and economies of scale—while maintaining ease of use and management—are available. And the application profiles that can withstand latency associated with storing data remotely are better understood. Cloud storage can now be leveraged as part of a data storage tiering model for persistent data. And it will become that, provided vendors get past pushing hype and get down to discussing what can be delivered today. ⊙

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

The way we back up remote offices
Ford has a better storage idea
Quality Awards: Top NAS products
SAS puts squeeze on FC
vSphere and storage

Remote-office backup not getting much easier

THE LAST TIME we asked *Storage* magazine readers about backing up remote offices, 53% of respondents said they did it at each individual remote site. That might not be the best way to protect corporate data, but it's still the most popular, with 50% still not centralizing remote-office backup. To be fair, nearly two-thirds (63%) of those doing site-by-site backups have centralizing the process on their to-do lists. The other half—those who back up to a central site—use a variety of means, the most popular of which is pumping backup data through a WAN optimization device (30%). Data deduplication, which wasn't even a choice a couple of years ago, is now used by 27% of centralizers. But despite all of the recent technical developments in disk-based backup, a lot of remote sites are still backing up to tape; 55% of those surveyed said their remote offices back up directly to tape, while another 16% said their remote site backups go first to disk and then to tape. Only 7% said they use an online backup service for their remote offices, but 17% of those planning to get a better handle on remote backups said they'll look to the cloud for help.

—Rich Castagna

Where are backups for your company's remote sites performed?

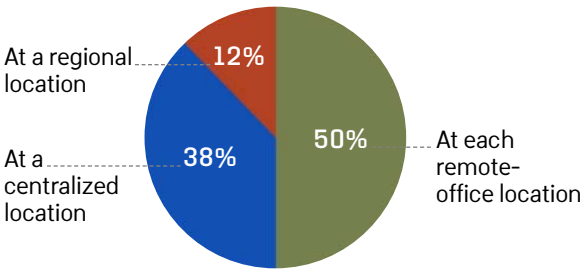

How do you back up data at remote offices?

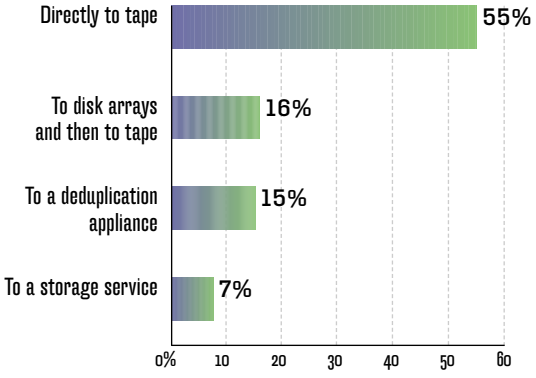

Currently, what is the most common problem you encounter with regard to backing up remote or branch office data?

21%	Sending data across the WAN to corporate results in intolerable packet loss/throughput levels
19%	Remote site backups are done by inexperienced staff and/or aren't performed regularly
18%	Backing up to tape at each site is costly and/or unreliable
11%	Remote site backup tapes (or other media) aren't sent offsite
11%	Added backup data from remote sites makes backup at the main data center difficult
4%	Cloud (online) backup services are inadequate

20%

Of remote backups are performed by non-IT personnel

"The potential exposure of sensitive data due to lost or misplaced tapes is a huge issue with using tape backups for remote sites."

—Survey respondent

Check out the following resources from our sponsor:

3PAR, page 4

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

CareerBuilder Positioned for Global Expansion Thanks to 3PAR Utility Storage