

STORAGE

MANAGING THE INFORMATION THAT DRIVES THE ENTERPRISE



Software-defined storage—does hardware matter anymore?

Software-defined storage promises low-cost, scalable storage, but you need to read the fine print.

SNAPSHOT 1

For block file flexibility, think unified array

EDITOR'S NOTE / CASTAGNA

Capacity, performance overshadow storage reliability

CLOUD ARCHIVING

Cheap, elastic storage: Good for cloud archiving

STORAGE REVOLUTION / TOIGO

NAS arrays morphing into cloud appliances

SNAPSHOT 2

Unified arrays aren't just for user shares

HOT SPOTS / SINCLAIR

Flash ups its endurance ante

QUALITY AWARDS: NAS

NetApp and Synology snap up NAS users' kudos

READ-WRITE / WHALEN

The advantages of data protection clouds



The day you get back, challenges will be solutions

➤ **Storage Decisions** CONFERENCES AND SEMINARS



Storage Decisions events are guaranteed to deliver practical advice and technical guidance that you can put to use as soon as you return to the office. Featuring the world's most renowned storage experts, our events offer tips and vendor-neutral advice for tackling current issues. Our innovative vendor showcases help you avoid falling behind by providing exclusive insight into market trends and emerging technologies. Plus, only you and your true peers in IT can qualify for FREE admission, ensuring an unmatched experience at an unbeatable value.

Visit www.StorageDecisions.com to register for an event near you.



Home

Castagna: Capacity, performance overshadow storage reliability

Toigo: NAS arrays morphing into cloud appliances

Software-defined storage—does hardware matter anymore?

Snapshot 1: For block-file flexibility, think unified arrays

The lowdown on cloud archiving

Snapshot 2: Key apps for unified arrays

NetApp and Synology win NAS users' kudos

Sinclair: Flash endurance keeps getting better

Whalen: Data protection clouds vs. mega clouds

About us

Vendors offer speed, space but overlook reliability

Storage buyers can get the capacity and performance they need, but they may not get the reliability they want.

WHEN YOU BUY a data storage system, what do you expect? Let me put that another way: Your company just shelled out hundreds of thousands (millions?) of dollars for a new array—what now? That, of course, depends on how you measure your needs and how well that new big chunk of iron meets those needs—or seems to.

A few recent surveys—a couple sponsored by data storage vendors so their results are taken with the requisite grain of salt—suggest that storage buyers not only have pretty diverse purchasing motivations, but also equally diverse criteria for determining how well a system meets their needs. Call it “survey convergence” if you like, but comparing bits from this survey and that—along with our own research—starts to tell a pretty interesting story about [data storage purchases](#).

SEPARATING SERIOUS FROM SELF-SERVING SURVEYS

The survey from Conduvix, a data storage vendor that sells an I/O hop-up kit called [V-locity](#), endeavors to demonstrate that storage buyers are prowling the aisles with their shopping carts because their current storage systems are performance-challenged. Self-serving? Of course, but no less interesting when you get past some of the sillier questions and look at some of the more interesting details.

For example, this survey says that 51% of respondents bought a new SAN to help improve performance due to the stresses that virtual environments put on storage. If a performance lift was all that was needed, buying a whole new rig is a fairly radical reaction. Some of the less surprising options were related to [solid-state storage](#), such as adding flash to existing [data storage systems](#) (27%), putting flash into servers (17%) and purchasing PCIe-based flash cards, which presumably would also go into servers.

The second survey that caught my eye was fielded by [Western Digital](#) and queried CIOs and other “IT decision makers.” There was some fluffy stuff in WD’s survey, too—like citing the 85% of respondents who said they “recognize the positive impact data can have on their bottom line”—I guess the remaining 15% just said, “Data, shmata...” And given their line of business, it doesn’t

Home

Castagna: Capacity, performance overshadow storage reliability

Toigo: NAS arrays morphing into cloud appliances

Software-defined storage—does hardware matter anymore?

Snapshot 1: For block-file flexibility, think unified arrays

The lowdown on cloud archiving

Snapshot 2: Key apps for unified arrays

NetApp and Synology win NAS users' kudos

Sinclair: Flash endurance keeps getting better

Whalen: Data protection clouds vs. mega clouds

About us

shock us that WD reports that 55% of their survey audience aren't yet storing all the data they need for that positive impact to happen. Other numbers from the survey support the capacity angle for motivating data storage buyers—89% of respondents think all data has value.

For me, the most interesting data point in the WD survey is that 39% of their respondents said reliability is the most important consideration when buying storage systems, compared to 25% who said cost. In the dozen-plus years that we've fielded our own [Storage Purchasing Intentions survey](#), we've seen that cost is diminishing as a key factor when buying an array.

Money's always important, but it's not the most important thing on storage managers' minds when they evaluate data storage vendors' trinkets.

WD's capacity angle holds up vs. our home-grown data, too. This month's [Snapshot bears that out](#)—70% said capacity was the main motivator for their considering to purchase a new unified array. I'm sure that number is a bit high because it specifically relates to those combo unified systems that mix block and file in one box, but even if you chisel off a few points, the capacity consideration was still miles ahead of the next most popular reason to buy—performance (46%).

Hmm... so even though those two surveys come from data storage vendors with pretty obvious objectives in mind, they do seem bolster the notion that storage buyers are looking for capacity and performance—maybe not in that order, but both are definitely in the picture.

SERVICE, RELIABILITY MATTER

WD's reliability data point is less compelling. We've conducted our [Quality Awards surveys](#) for 10 years now, querying real users about the service and reliability of their enterprise and midrange arrays, and their NAS boxes.

All of the Quality Awards surveys include a section that specifically addresses reliability. In all three of the storage systems categories that we survey, users' ratings for reliability have gone down, down, down over the years. In our most recent survey for enterprise arrays, for example, our users gave those systems the lowest reliability ratings we've seen in eight years. In the midrange category, you have to go back six years to find a lower set of reliability scores than what was recorded in the most recent survey. The latest NAS survey findings were even grimmer: the lowest reliability scores ever—and that's in 10 years of getting user feedback.

It looks like data storage vendors have the answers to the capacity and performance issues. With big disks growing to 8 TB and 10 TB and beyond, there should be ample capacity to handle big data, archives and any other capacity-gobbling app. Performance, too, appears to be addressed with solid-state gizmos shooting data down the express lanes in [hybrid and all-flash arrays](#) or in servers.

On the other hand, data storage vendors might want to spend a little more time and effort on the seemingly irksome reliability issue. ■

RICH CASTAGNA is TechTarget's VP of Editorial.



Look out, NAS—here come the cloud appliances!

NAS may seem old-fashioned, but it's still a better bet than proprietary storage.

NOT LONG AGO, I received an email message from Promise Technology regarding their “first consumer product,” which they unveiled at the Consumer Electronics Show in January. Called Apollo, the product was an “appliance” that stored “up to 4 TB of data assets” including “treasured family photos and videos that proliferate like wildfire on family member phones, tablets and laptops, and at increasingly high resolutions.” The product sounded a bit to me like a network-attached storage array but it was, according to the vendor, “a personal cloud appliance”—I suppose, because it's accessed via an app.

I don't want to pick on Promise, whose products I use and generally find to be of good quality, but this “cloud storage” product concept has become increasingly pervasive in the market these days, whether aimed at home or

business users. In the process, the [concept of a traditional NAS](#)—a network-attached file server connected to its own storage array—is being usurped. In a way, the marketing around the Promise product provides a bit of a microcosm for a broader discussion of the architectural pluses and minuses of “cloud appliances” aimed at small, medium and large businesses.

THE PROPRIETARY PROBLEM

In the case of Apollo, as with other cloud appliances, we have a problem of a proprietary nature. Promise generally makes a pretty good kit, and one that is pretty open. In this case, the notion of proprietary is introduced as a function of operational compatibility. Apollo will be sold “exclusively through Apple stores.” That already hints at a proprietary product. I'm not a big fan of Apple anything, though we have some iMacs, MacBook and even an iPhone or iPad participating on the sprawling home network in my house. While a product like Apollo might get some play, I suppose, I doubt that the non-Apple devices would find it very useful, especially if the app required for access is only iOS- or OSX-compatible. If it's like other Apple storage products, there may also be an issue of [file system compatibility](#) with the many Windows and Linux systems and devices deployed in the network.

[NAS vendors](#) have worked for many years to ensure that

Home

Castagna: Capacity, performance overshadow storage reliability

Toigo: NAS arrays morphing into cloud appliances

Software-defined storage—does hardware matter anymore?

Snapshot 1: For block-file flexibility, think unified arrays

The lowdown on cloud archiving

Snapshot 2: Key apps for unified arrays

NetApp and Synology win NAS users' kudos

Sinclair: Flash endurance keeps getting better

Whalen: Data protection clouds vs. mega clouds

About us

Home

Castagna: Capacity, performance overshadow storage reliability

Toigo: NAS arrays morphing into cloud appliances

Software-defined storage—does hardware matter anymore?

Snapshot 1: For block-file flexibility, think unified arrays

The lowdown on cloud archiving

Snapshot 2: Key apps for unified arrays

NetApp and Synology win NAS users' kudos

Sinclair: Flash endurance keeps getting better

Whalen: Data protection clouds vs. mega clouds

About us

their products could be accessed via “open” network file system protocols so that data could be stored and retrieved by any device. No proprietary apps or protocols were needed to get my files down from an Isilon, NetApp or Promise NAS array as long as I knew the IP address of the box. Sacrificing this universality on the altar of simplicity or ease of setup is concerning.

The key point here, however, is that proprietary cloud appliances are by nature exclusionary and limited in their support of the broadest number of use cases. Buying something that is aimed at a specific OS, hypervisor or file system—even if it “simplifies” deployment and use in the short term—is creating a technology stovepipe in your environment. That’s tantamount to shooting yourself in the foot. While the product may plug and play today, it may pose some serious and expensive challenges in the future.

THE CAPACITY CONUNDRUM

I also have a problem with the capacity of the basic Apollo product. These days, a home data storage repository with 4 TB capacity isn’t all that grand. (Indeed, individual disk drives that exceed that capacity are available!) Apollo’s 4 TB size seems to reflect current analyses that peg Internet data consumption rates within typical households at about 3 GB per week (12 GB per month, 144 GB annually). Basically, product designers are making the rather modest assumption that something less than 144 GB of data per year will be added to local storage in the home, making 4 TB more than adequate for keeping all the pictures and videos safe.

The real story is much different. Internet or mobile data use is in no way a gauge of how much data is created and stored locally in a home—or in a business. In my house, 8 TB hard disks fill up very quickly with games, movies, music, artwork, writing projects, and so on.

THERE IS ALSO A NEED, IN THE BUSINESS WORLD AT LEAST, TO THINK ABOUT CAPACITY BURN RATES IN TERMS OF OPERATIONAL ECONOMICS.

In a business, the capacity burn rate is even greater—upwards of 40% per year according to IDC, or 10 to 20 times that rate if your servers are virtualized. This helps to explain why the NAS folks have gone to such pains to build new scalability techniques into their products—whether [capacity growth](#) is achieved by adding more storage media or by linking to back-end [Linear Tape File System tape](#) or by attaching to a public cloud via a network gateway.

There is also a need, in the business world at least, to think about capacity burn rates in terms of operational economics. It makes little sense to use expensive hard drives or flash storage in a NetApp rig to store files that are hardly ever accessed. Typically, smart administrators want to place data where its re-reference and revision characteristics require it to be placed. “Hot” data that changes a lot and is accessed frequently needs to go on SSDs or fast

Home

Castagna: Capacity, performance overshadow storage reliability

Toigo: NAS arrays morphing into cloud appliances

Software-defined storage—does hardware matter anymore?

Snapshot 1: For block-file flexibility, think unified arrays

The lowdown on cloud archiving

Snapshot 2: Key apps for unified arrays

NetApp and Synology win NAS users' kudos

Sinclair: Flash endurance keeps getting better

Whalen: Data protection clouds vs. mega clouds

About us

disk, while “colder” data needs to be stored on capacity media, preferably tape. And if you are just putting data to sleep, think public cloud.

That’s the philosophy of products like Crossroads Systems’ StrongBox, or that product used in conjunction with the dTernity Media Service from Fujifilm, which puts that extra layer of mass, low-cost, tape-based cloud storage into the mix. Such a “bottomless” network-accessible storage repository begins to appeal at about 30 TB of data, according to Crossroads spokespersons, which doesn’t make it an alternative to Apollo for home use, but in a small, medium or larger office, the economics of NAS-to-tape-to-cloud make a lot of sense.

SOME VENDORS OFFER WORKABLE ALTERNATIVES

Another alternative is [object storage](#), like the one advanced by Caringo and a few other vendors. Cluster a number of smaller, inexpensive, “cloud appliances” (actually hyper-converged server/storage nodes) and present them to users via either a standard NAS access method or a “cloud” interface (RESTful commands via a browser or app). Objects can be segregated by policy, protected in a manner appropriate to their importance and use, and accessed via familiar and simple interfaces as though they were all on the same big disk drive. An object storage model is already seen as the future of network storage and

is likely to become part of any “[cloud appliance](#)” pitch in the future—even from Promise, which has done some significant work on object storage in recent years.

Bringing this to a speedy conclusion, my recommended “best practice” for home-based network-accessible storage is to buy and deploy whatever you want and can afford, remembering that losing your pictures and movies will not bring about the apocalypse in any case.

However, if your business depends on the data you’re storing, be a bit more circumspect before you ditch NAS for simple “cloud appliances.” First, think about compatibility and look for storage that will work with all of your data. Second, look at the scalability and economics of the product. For the money, capacity and resiliency, [mixing HDD/SSD](#), tape and cloud makes a lot of sense—even if it is more challenging to access with a smartphone. Finally, stay tuned for real cloud storage—not just storage in an off-site data center on the Internet, not just storage accessed via a smartphone or tablet app, but an object-based repository featuring the clustering and virtualization of physical storage infrastructure and a robust data management and access component. That’s a real cloud appliance. ■

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

Everything you need to know about software-defined storage

There are several distinct variations of software-defined storage—here are the pros and cons of each.

BY MARC STAIMER

BARANOZDEMIR/ISTOCK

HOME

NOTHING IN THE storage world elicits more divergent opinions than the term “software-defined storage.” With no universally accepted definition, software-defined storage (SDS) is vendor-specific. Vendors shape the SDS definition to match their storage offerings. The result is that every storage vendor appears to offer SDS.

The closest the software-defined storage market has come to a SDS consensus is more [marketecture](#) than architecture.

Software-defined storage separates the data storage hardware from the software that manages it. The storage software is itself hardware-independent. The storage control plane is usually, but not always, separated from the data plane.

That broad definition enables just about every variation of storage currently available. So it’s up to the software-defined storage market consumer to determine which ones work best for them.

DRIVING FORCES BEHIND THE SDS TREND

All storage systems have always been software-defined. What’s changed is that the software has become portable.

[Storage system software](#) historically was tied to the

Home

Castagna: Capacity, performance overshadow storage reliability

Toigo: NAS arrays morphing into cloud appliances

Software-defined storage—does hardware matter anymore?

Snapshot 1: For block-file flexibility, think unified arrays

The lowdown on cloud archiving

Snapshot 2: Key apps for unified arrays

NetApp and Synology win NAS users' kudos

Sinclair: Flash endurance keeps getting better

Whalen: Data protection clouds vs. mega clouds

About us

hardware it managed. When the hardware ran out of capacity or performance, it had to be replaced and the software licensing was repurchased along with the data storage hardware.

What made matters significantly worse was that storage system architectures created isolated silos. Unique infrastructures made everything from storage provisioning, data protection, disaster recovery, tech refresh, data migration, power and cooling more and more untenable. Compound that with the ongoing trend of rapid data growth and the need to store ever-increasing amounts of data, and the available architectures made storage systems management too complicated, difficult, expensive and ultimately unmaintainable.

Several technological factors contributed to the software-defined storage market phenomenon as well. The first is the direct result of the continuous x86 compute architecture performance improvements. The x86 performance improvements and availability of cores for specific storage functions has led to x86 architectural standardization for storage systems.

An additional technological factor aiding SDS is the general acceptance of x86 virtualization of servers, desktops, applications and networking ([SDN](#)). That has helped condition IT into accepting separation of the data image from the hardware upon which it resides.

The popularity of cloud technologies has also had a major effect on driving the software-defined storage market. The cloud data centers needed a new and much lower-cost storage architecture based on industry standards and commodity hardware.

Other technological factors driving the software-defined storage market phenomenon include [server-side flash storage](#) and the software that allows memory and server storage to be transparently shared with other physical server hosts.

All of these technology changes eroded the differentiation between server and storage hardware while expediting storage software portability and flexibility, and, not inconsequentially, also radically reducing storage costs.

SDS CATEGORIES PROS AND CONS

With no working standard SDS definition, a variety of technologies have emerged in the software-defined storage market. For our purposes, the four categories of SDS include:

1. Hypervisor-based SDS
2. Hyper-converged infrastructure (HCI) SDS
3. Storage virtualization SDS
4. Scale-out object and/or file SDS

There are both significant differences and equally significant similarities between these categories, and several products may actually fit into multiple categories. And some products are unique enough to be in their own category such as PernixData or Saratoga Speed.

Since SDS is focused on flexibility, simplicity, scalability with performance and total cost of ownership (TCO), we'll use those criteria to evaluate the pros and cons of each SDS approach.

1. HYPERVISOR-BASED SDS

VMware invented this category with [VMware vSphere Virtual SAN](#). This is the only category that is a specific product. Virtual SAN is architected to be a part of vSphere and operates as a feature of vSphere and works with all vSphere virtual machines and virtual desktops. Virtual SAN runs in the ESXi layer which means it's not a virtual storage appliance and doesn't require a VM to execute.

Flexibility

PRO: Virtual SAN works with both hard disk drives (HDD) and solid-state drives (SSD) including DIMM-based flash drives, PCIe, SAS, SATA and even NVMe. VMware Virtual SAN supports both HDDs and SSDs in a hybrid mode or all SSDs in all-flash mode.

CON: Virtual SAN is a closed-loop SDS in that it only works with VMware vSphere 5.5 or better. Older ESXi implementations, other hypervisors, or physical machines don't work with Virtual SAN. It can't be used by virtual or physical machines that are not part of the vSphere cluster. There is an element of do-it-yourself (DIY) to Virtual SAN. For example, running on inexpensive commoditized hardware is somewhat limited to VMware's hardware compatibility list (HCL). If hardware isn't on the list, it's not supported.

Scalability and performance

PRO: Virtual SAN is highly scalable while delivering high levels of performance. It scales out through vSphere clustering and can support up to 64 vSphere hosts per cluster. Each vSphere host supports approximately 140 TB raw

storage capacity and well north of 8 PB of raw storage capacity per cluster. On the performance side, each Virtual SAN host can supply up to 90,000 IOPS, yielding more than 5 million IOPS per cluster.

CON: Virtual SAN clusters cannot exceed 8.8 PB. If more capacity is required, it is not a good fit. If a VM requires more IOPS than the 90,000 available in their vSphere host, it can get them from other nodes in the cluster, but with at a considerable latency penalty. Inter-cluster storage performance is another issue. Most Virtual SAN clusters use 10 Gbps to 40 Gbps Ethernet and TCP/IP to interconnect the hosts. This architecture essentially replaces a deterministic system bus with a non-deterministic TCP/IP network so latencies between hosts become highly variable. Unless the cluster uses more sophisticated and faster interconnections, its storage performance from one clustered host to another will be highly variable and inconsistent.

Simplicity

PRO: Virtual SAN is simple because it's natively integrated as part of the VMware stack. It feels and acts like all other vSphere features so it's intuitive for a vSphere administrator. Virtual SAN automates storage tasks on a per-VM basis such as provisioning, snapshots/data protection, high availability, stretch clusters, disaster recovery and business continuity. Even data migration to a Virtual SAN can be accomplished relatively simply via [vSphere Storage vMotion](#).

CON: Converting from a siloed storage environment to a pure Virtual SAN requires converting non-VM images to

Home

Castagna: Capacity, performance overshadow storage reliability

Toigo: NAS arrays morphing into cloud appliances

Software-defined storage—does hardware matter anymore?

Snapshot 1: For block-file flexibility, think unified arrays

The lowdown on cloud archiving

Snapshot 2: Key apps for unified arrays

NetApp and Synology win NAS users' kudos

Sinclair: Flash endurance keeps getting better

Whalen: Data protection clouds vs. mega clouds

About us

VMs first. It's a time-consuming process for non-vSphere environments.

TCO

PRO: Compared to legacy storage architectures, its TCO should be less. The saving comes from the difference in the price of drives (HDDs and SSDs) in a storage system compared to the same drives in a server. Those drives are typically three times more expensive in the storage system. Some of the other Virtual SAN cost advantages come from the predictable pay-as-you go scaling, unified storage management, unified data protection, disaster recovery and business continuity; and consolidated storage networking.

CON: Until the most recent release—version 6.2—Virtual SAN lacked deduplication and compression capabilities. This raises costs per usable TB considerably versus SDS products that include data reduction. In addition, making sure data and VMDKs on a specific clustered vSphere host remain available to the rest of the cluster in case that host fails currently requires multi-copy mirroring. Best practices require at least two copies of the original data and many administrators opt for three copies. This practice eliminates the drive price advantages. And because Virtual SAN is a vSphere exclusive option, it has its own license costs that can be substantial.

2. HYPER-CONVERGED INFRASTRUCTURE (HCI) SDS

HCI combines the server, storage, networking and hypervisor, and packages them into clustered nodes. [HCI](#)

[is designed](#) to eliminate do-it-yourself integration headaches, expensive server hardware, the need to over-provision storage, high availability issues, complex storage management and hardware compatibility issues. There are many HCI options from server and other vendors, including: Atlantis, Cisco, Dell, EMC, Fujitsu, Gridstore, Hitachi, HPE, IBM, Lenovo, Maxta, NEC, Newisys, Nutanix, Quanta, Saratoga Speed, Scale Computing, SimpliVity, StarWind, StorMagic and SuperMicro.

Flexibility

PRO: As with VMware Virtual SAN, a VM administrator can control the storage. In fact, several HCI implementations are based on VMware vSphere and Virtual SAN including VMware's [EVO:RAIL reference design](#). Several HCI vendors offer a choice of hypervisors from vSphere, Hyper-V, KVM or XenServer and some have “bare metal” offerings available with Linux using Docker containers or application virtualization (Saratoga Speed) without requiring a hypervisor. Many HCI implementations allow different capacity-sized nodes within the cluster. A few are software-only such as Maxta, StarWind and StorMagic. Maxta partners with most of the main server vendors including Dell, Quanta and SuperMicro.

CON: HCI are closed-loop [SDS systems](#) so their storage only works with the server nodes in the cluster. Any physical or virtual host not in the HCI cluster cannot access the HCI storage. (There are exceptions: Saratoga Speed provides what they call all-flash ultra-converged infrastructure (UCI) which can act as a target for physical or virtual hosts outside the UCI cluster.)

Home

Castagna: Capacity, performance overshadow storage reliability

Toigo: NAS arrays morphing into cloud appliances

Software-defined storage—does hardware matter anymore?

Snapshot 1: For block-file flexibility, think unified arrays

The lowdown on cloud archiving

Snapshot 2: Key apps for unified arrays

NetApp and Synology win NAS users' kudos

Sinclair: Flash endurance keeps getting better

Whalen: Data protection clouds vs. mega clouds

About us

Cluster hardware is limited to what the HCI vendor provides, or hardware certified by the software-only HCIs. As with VMware Virtual SAN, there is vendor lock-in and replacing the vendor requires migrating everything from the old HCI to the new, which can be time-consuming and tedious.

Scalability and performance

PRO: Scaling HCI is as simple as adding a node to the cluster. Scaling storage capacity just requires adding drives (HDD or SSD) up to a node's maximum or adding additional nodes. Each HCI product has its own scalability and performance limitations; however, most scale well into the PBs and add performance linearly with each server node added to the cluster.

CON: HCI cluster capacity is limited by the number of nodes supported in the cluster and the amount of capacity supported per node. If a VM requires more IOPS than the 90,000 available in their vSphere host, it can get IOPS from other nodes, but with a considerable latency penalty. Inter-cluster storage performance is another issue. Most Virtual SAN clusters use 10 Gbps to 40Gbps Ethernet and TCP/IP to interconnect the hosts so latencies between hosts can be highly variable.

Simplicity

PRO: Plug it in, turn it on, configure and you're done. Few systems are simpler. No DIY, and there's just one throat to choke for support.

CON: Converting from a siloed storage environment to an HCI cluster requires first converting both non-VM

images and VMs to the HCI VMs or Docker containers, a time-consuming process.

TCO

PRO: Total cost of ownership (TCO). Similar to VMware Virtual SAN. Many HCI vendors include inline deduplication and/or compression that can reduce total capacity requirements by as much as 83% to 90% depending on the data, which reduces TCO significantly.

CON: HCIs—like Virtual SAN—have an issue in making sure data, VM images, VD images and Docker container images on a specific HCI node remain available to the rest of the cluster should that node fail. Today, that requires multi-copy mirroring. Best practices require minimally two copies of the original data and more commonly three copies which increases total capacity requirements and related costs.

3. STORAGE VIRTUALIZATION SDS

Storage virtualization SDS is the most mature variation of SDS in the software-defined storage market. It's been around since the early 2000s, when it was just called storage virtualization. [Storage virtualization SDS](#) is primarily the entire storage software stack including all storage services, optimized to run in the x86 architecture and convert hosts into powerful full-featured storage controllers. It virtualizes server storage and external storage systems to create a single or multiple virtual storage pools with different capacities, data protection policies and performance characteristics. Storage virtualization

Home

Castagna: Capacity, performance overshadow storage reliability

Toigo: NAS arrays morphing into cloud appliances

Software-defined storage—does hardware matter anymore?

Snapshot 1: For block-file flexibility, think unified arrays

The lowdown on cloud archiving

Snapshot 2: Key apps for unified arrays

NetApp and Synology win NAS users' kudos

Sinclair: Flash endurance keeps getting better

Whalen: Data protection clouds vs. mega clouds

About us

SDS essentially converts x86 servers into storage systems; some products can also run as VMs as a virtual storage appliance (VSA). Storage virtualization SDS is primarily scale-up architecture, but some products can scale-out as well. They're architected to eliminate costly proprietary hardware, take advantage of lower-cost server drives, repurpose older storage systems and simplify data migration. Some of the better-known players/products include: DataCore Software SANSymphony, EMC VIPR, IBM SVC, Microsoft Windows 2012 R2 (and higher), NetApp Ontap Cloud, Nexenta Systems NexentaStor, QuantaStor and Starwind Software.

Flexibility

PRO: It works with most x86 physical hosts or VMs as long as the hardware or hypervisor is certified and supported by the vendor. It converts all storage that sits behind it into the virtual storage pool, enabling repurposing of older storage. The scale-out versions permit physical or VM access to any node. Multi-copy mirroring isn't necessary to protect against a single controller failure, although it's available. Storage virtualization SDS can be provided as software or bundled with server hardware similar to HCI.

CON: Most storage virtualization SDS can only run on the specific commodity hardware certified and supported by the vendor. Products that can run as VSAs require hypervisors certified and supported by the vendor.

Scalability and performance

PRO: Scaling is multi-dimensional as each node in the cluster can scale-up and more nodes can be added to scale

out. Generally, storage virtualization SDS is equivalent to most active-active silo storage systems.

CON: On paper, these systems support tremendous capacity scalability, but the pragmatic approach is a bit different. Storage virtualization SDS capacity is constrained by x86 server limitations. Each server can handle only so much capacity before performance declines below acceptable levels. Storage virtualization SDS scale-out is constrained by clustering because the number of storage controller nodes supported is limited. Performance may also be constrained by the same limitations.

Simplicity

PRO: When bundled with hardware, storage virtualization SDS is a very simple storage system. It leverages commodity off-the-shelf hardware, has better scalability and in some cases provides both block (SAN) and file (NAS). But in the end, it's still a siloed storage system in an inexpensive container.

CON: Storage virtualization SDS is primarily DIY system integration which requires testing, QA and efforts to make sure the software works correctly with the hardware. Implementation may require professional services or a systems integrator.

TCO

PRO: The biggest cost savings in storage virtualization SDS comes from commodity hardware and server-based drives. Another cost saving comes from inline data reduction technologies. Compared to equivalent storage systems, most storage virtualization SDS will yield a much more

Home

Castagna: Capacity, performance overshadow storage reliability

Toigo: NAS arrays morphing into cloud appliances

Software-defined storage—does hardware matter anymore?

Snapshot 1: For block-file flexibility, think unified arrays

The lowdown on cloud archiving

Snapshot 2: Key apps for unified arrays

NetApp and Synology win NAS users' kudos

Sinclair: Flash endurance keeps getting better

Whalen: Data protection clouds vs. mega clouds

About us

favorable total cost of ownership.

CON: Licensing can be a bit pricey depending on the vendor. And not all storage virtualization SDS products provide inline deduplication and/or compression. These issues can have a deleterious impact on TCO.

4. SCALE-OUT OBJECT AND/OR FILE SDS

Recently we've seen the introduction of scale-out object SDS. Object storage manages data as objects which contain the data, metadata and a unique identifier. There are quite a few object storage players along with two open-source variations in OpenStack Swift and Ceph (distributed and supported by RedHat).

Scale-out file SDS is a highly scalable NAS often with special characteristics such as object storage resilience or unique metadata (Qumulo). Some scale-out file SDS products actually sit atop object storage (Exablox) and others are essentially clustered scale-out implementations of IBM's General Parallel File System (Spectrum Storage).

Flexibility

PRO: Both scale-out [SDS architectures](#) are designed from the ground up for x86 servers. Some can be implemented as software on hardware certified by the vendor while others are bundled with server hardware. They are not designed to be VSAs and typically are intended for secondary or non-mission critical applications.

Many scale-out object or file SDS products can act as HDFS storage for Hadoop implementations. That can significantly lower the cost of HDFS storage by reducing

the number of mirrored copies required and allowing re-purposing NFS or SMB data.

CON: Whether delivered as software or bundled with hardware, the hardware must be certified and supported by the vendors.

Scalability and performance

PRO: Scaling is multi-dimensional: each node can be scaled individually and generally the cluster itself can add nodes for capacity or performance. Performance for both will never approach that of high-performance block storage

CON: Scale-out file SDS generally doesn't scale as high as scale-out object storage, but object will have somewhat higher latencies. Object storage has significant additional latencies from the metadata and data resiliency functions. Both types are best suited for secondary applications where high performance is not a requirement.

Simplicity

PRO: When bundled with hardware, scale-out object or file storage is very simple to setup, configure, and manage. Implementing it as software requires DIY systems integration. Both types leverage commodity hardware, have exceptional scalability and—in the case of scale-out object storage—unmatched data resilience and longevity via erasure coding.

CON: When scale-out file or object storage SDS is purchased as software, it's a DIY project, so special skills, professional services or a systems integrator may be required.

In addition, when these types of SDS are used for secondary applications such as archive, the data must be

Home

Castagna: Capacity, performance overshadow storage reliability

Toigo: NAS arrays morphing into cloud appliances

Software-defined storage—does hardware matter anymore?

Snapshot 1: For block-file flexibility, think unified arrays

The lowdown on cloud archiving

Snapshot 2: Key apps for unified arrays

NetApp and Synology win NAS users' kudos

Sinclair: Flash endurance keeps getting better

Whalen: Data protection clouds vs. mega clouds

About us

moved from its current location; some vendors have products that can do this, most rely on third-party software.

TCO

PRO: Both types are designed to be low-cost and offer very few add-on functions, usually licensed on an annual subscription basis. Scale-out object storage with erasure codes can lower the overall cost per GB as it requires less overhead than traditional RAID and replication data protection.

CON: Data reduction—deduplication or compression—currently isn't available with scale-out object SDS and rarely with scale-out file SDS. That adds to the TCO.

SOFTWARE-DEFINED STORAGE MARKET

BOTTOM LINE

SDS is a broad marketing term with a variety of software-defined storage market implementations, each with its own pluses and minuses.

Selecting the right SDS for the job requires accurate understanding of the application, storage capacity and performance requirements, the organization's skill set, and what the software-defined storage market is capable of handling. ■

MARC STAIMER is the founder, senior analyst and CDS with Dragon Slayer Consulting in Beaverton, Oregon.

Snapshot 1

Unified arrays add flexible capacity to data center infrastructure

Home

Castagna: Capacity, performance overshadow storage reliability

Toigo: NAS arrays morphing into cloud appliances

Software-defined storage—does hardware matter anymore?

Snapshot 1: For block-file flexibility, think unified arrays

The lowdown on cloud archiving

Snapshot 2: Key apps for unified arrays


NetApp and Synology win NAS users' kudos

Sinclair: Flash endurance keeps getting better

Whalen: Data protection clouds vs. mega clouds


About us

➤ Top five reasons for planned unified array purchase



* UP TO THREE SELECTIONS ALLOWED

➤ What type of data access will you use with your new unified array?



43
Percentage who need
16 Gbps FC for block

69
Percentage who
want 10Gig Ethernet
for block



Deep dive into cloud archiving

Archive is often touted as a “killer app” for cloud storage but there are some issues to consider before sending your data into the ether.

BY CHRIS EVANS



KMLMTZ66/FOTOLIA

ONE OF THE best features of public cloud storage has been the ability to easily store large volumes of data without having to deal with the headache of managing the infrastructure that supports it. With the kind of exponential data growth that continues in most organizations, the business of managing storage can be a chaotic chore, coping with growth and infrastructure refreshes. Cloud storage is potentially a lifesaver for IT departments that want to get off the treadmill of managing infrastructure so they can focus on the data. But is cloud archiving really a unique service and could the effort expended to use a cloud archive be more than the benefits achieved?

WHY YOU SHOULD ARCHIVE

Before looking at the technical aspects of cloud-based archiving, it's worth discussing why there's a need to archive data in the first place. The most obvious reason is cost; production systems (databases, files and unstructured data) consist of large amounts of inactive and rarely accessed information, all of which is sitting on expensive [primary storage](#). Companies are looking to retain data “forever”—or at least for a very long time—on the assumption that there is some future value to be obtained from the content. In some instances, regulatory constraints require data to be kept for long periods of time (decades) or in the case

Home

Castagna: Capacity, performance overshadow storage reliability

Toigo: NAS arrays morphing into cloud appliances

Software-defined storage—does hardware matter anymore?

Snapshot 1: For block-file flexibility, think unified arrays

The lowdown on cloud archiving

Snapshot 2: Key apps for unified arrays

NetApp and Synology win NAS users' kudos

Sinclair: Flash endurance keeps getting better

Whalen: Data protection clouds vs. mega clouds

About us

of medical record-keeping, the lifetime of the patient as a minimum. Moving data out of the production environment might also have a significant impact on other production costs, including savings on database licenses, smaller virtual machines or physical hosts where licenses are based on data volume.

There are also operational aspects to storing large volumes of data in primary systems; the bigger the system, the longer the backup/restore or recovery process, and the bigger the backups will be. There's no benefit to continually backing up data that never changes when it can be archived off primary systems and stored and protected in a different way. The performance of production systems can also be affected, too. There's much more overhead in accessing and storing data in a database with 100 million rows, compared to one with only five hundred thousand, for example.

WHY YOU SHOULD CONSIDER CLOUD ARCHIVE

So the need for an archive can be clearly seen, but why choose the cloud as the target destination? There are a number of operational benefits that are inherent in [using cloud services](#) that make it an attractive destination for archive data. These include:

■ **Elasticity.** The cloud storage provider takes the responsibility and associated headaches of making sure that the archive capacity is always available to meet demand. The customer simply consumes the resource on the assumption that there is an infinite capacity available. There's no

need to think about data center space, power, cooling or other physical aspects.

■ **Abstraction.** The customer doesn't need to know or care how the data is being stored in the cloud, only that the service is being delivered by the cloud provider at some agreed-upon service level. This means the data could be on disk, tape, optical or any combination. The cloud vendor takes the responsibility for managing and refreshing the storage media and associated infrastructure over time, as technology ages and needs replacing.

■ **Durability.** Primary storage resilience is measured in terms of availability or how much uptime the system delivers. We typically see figures quoted of five, six or now seven 9s, meaning 99.999% uptime or better. In the archive world, measurement is based on durability with a lower level of availability as the data is assumed to be accessed less frequently, but needs to be there in 5, 10, 20 or 50 years' time. [Amazon Web Services' S3 offering](#), for example, touts durability levels of 99.999999999% or "eleven 9s."

■ **Cost.** The cost of cloud storage is predictable and based on access profile and volume of data stored (more on this later), making rebilling and accounting easier.

So cloud archive makes sense, the question is, how do



The cost structure for cloud-based archive can be very different to on-premises and is typically based on the volume of data stored, plus a charge for recalling and accessing the data in the future.

Home

Castagna: Capacity, performance overshadow storage reliability

Toigo: NAS arrays morphing into cloud appliances

Software-defined storage—does hardware matter anymore?

Snapshot 1: For block-file flexibility, think unified arrays

The lowdown on cloud archiving

Snapshot 2: Key apps for unified arrays

NetApp and Synology win NAS users' kudos

Sinclair: Flash endurance keeps getting better

Whalen: Data protection clouds vs. mega clouds

About us

IT operations teams get the data in and out in a way that meets operational requirements?

CLOUD ARCHIVING CONSIDERATIONS

Probably the most obvious concern about a cloud archive is that of security. How will my data be secured both in-flight across the network and at rest once it reaches the service provider's data center? The in-flight issue is easily resolved, as data in and out of cloud archives is managed through secure HTTPS protocols (SSL). So data transferred across the public network will be safe in flight.

Most providers now also offer the ability to [encrypt data](#) stored within their clouds. As an extra level of security, customers can provide their own encryption keys to be used by the provider to encrypt data on the customer's behalf. Alternatively, data can be encrypted before sending it to the cloud. The choice of encryption option is dictated by the risk profile of the customer; provider-based encryption may be sufficient, whereas compliance rules or flat-out paranoia may dictate using personal encryption keys. In that instance, the customer must maintain the keys for future data retrieval, which can be a significant effort if data is intended to be stored for many years.

The second issue to consider is that of performance, or how quickly data can be stored and retrieved from the cloud. Depending on the type of connectivity in place, the [latency](#) or round-trip time to write data into the cloud could be as high as 20 to 30 milliseconds. This level of response time is fine for sequential transfers but not so great for "random" access. In reality, most archive processes will

not have a problem with latency issues as they work on storing and retrieving large volumes of data but updates to metadata, if cloud-based, could be a problem.

Two other issues affect the performance of accessing data. First, the providers themselves may place restrictions on access. Amazon Web Services' Glacier, for example, provides a lower-cost alternative to S3 (Simple Storage Service) but provides access through a staging process that takes 3 to 5 hours to retrieve the data that is then available for up to 24 hours (after which it needs to be retrieved again). There are also transfer costs to access data above a free 1 GB limit with Glacier, which we will discuss later. Not all vendors have performance restrictions on data access; [Google Cloud Storage Nearline](#), for example, offers response times (or access to first byte) of around three seconds for long-term archive data. There's clearly a trade-off in choosing the right price of service versus the performance the service offers.

Accessibility and data format is another area of concern when using cloud archives. Archive platforms are typically object stores accessible over Web-based protocols. On-premises data, however, may be in the form of structured data (like a database), semi-structured data (like emails) or unstructured data-like files. Each of these data formats will be associated with metadata that's used to describe the content. So how does this data get transformed into a generic object format? One answer is to use



Cloud-based archiving removes headaches associated with planning and maintaining large archives, such as regular hardware and data format refreshes.

Home

Castagna: Capacity, performance overshadow storage reliability

Toigo: NAS arrays morphing into cloud appliances

Software-defined storage—does hardware matter anymore?

Snapshot 1: For block-file flexibility, think unified arrays

The lowdown on cloud archiving

Snapshot 2: Key apps for unified arrays

NetApp and Synology win NAS users' kudos

Sinclair: Flash endurance keeps getting better

Whalen: Data protection clouds vs. mega clouds

About us

products that act as either gateways or archive platforms to provide the bridge between the local and the archive format. Examples include the AWS Storage Gateway, [EMC CloudBoost](#), Microsoft Azure's [StorSimple](#), Nasuni [Cloud NAS](#) and NetApp's [AltaVault](#). Most of these products are conduits to cloud storage and don't directly integrate with a particular application. However, they do offer a more consumable protocol for archive data and the ability to cache some content locally and reduce the impact of always reverting to cloud storage to access data. Application integration work may still be required, but that could be the case whether the data is on or offsite.

Finally, you should consider cost. Most on-premises archiving systems are typically based on the cost of the infrastructure itself whereas cloud-based archiving will be driven by the volume of stored data and access profiles. As more data is stored and recalled from the archive, so the monthly costs increase. IT organizations need to be ready to deal with re-billing the cost back to their end users (where appropriate) and this will mean creating policies on data retention and retrieval as well as partitioning archive data into logical repositories (like vaults) that can be reported individually. The cost of using cloud storage could become a real issue where large volumes of data are placed with a single provider. This is because moving data between archives (and providers) could be cost-prohibitive even though it may be desirable for redundancy and risk reduction purposes.

One opportunity to reduce costs is to look at implementing data reduction technologies, such as deduplication and compression. These can be implemented in the

application before the data is archived, or deployed within a cloud gateway. One such product is [StorReduce](#) from a startup company of the same name. The StorReduce appliance sits in the public cloud and accepts data in S3 format, writing data back out to AWS S3 in deduplicated format. The company claims up to 95% savings on data stored, which can result in significant cost reductions with large archives.

CLOUD ARCHIVING: BOON OR BANE?

Archiving can work as a cloud-only application, with consideration to the points raised here: security, performance, accessibility and cost. In deciding whether cloud is the right place to put archive data, the issues of flexibility in billing and management have to be weighed against the requirement to implement on-premises systems that translate data into object compatible formats.

One final point to consider is how archive data will be accessed in the future. Having the data already in the cloud provides the ability to run cloud-based analytics against the archive. Accessing data within the cloud from cloud-based applications running as virtual instances typically doesn't attract any additional access costs and so cloud might prove the stepping stone to actually doing something useful with all of that cold data. ■

CHRIS EVANS is an independent consultant with [Langton Blue](#).



Cloud archiving provides the capability to store large volumes of data as objects, so some kind of protocol or format conversion (with metadata) is required to exploit cloud storage effectively.

Castagna: Capacity, performance overshadow storage reliability

Toigo: NAS arrays morphing into cloud appliances

Software-defined storage—does hardware matter anymore?

Snapshot 1: For block-file flexibility, think unified arrays

The lowdown on cloud archiving

Snapshot 2: Key apps for unified arrays

NetApp and Synology win NAS users' kudos

Sinclair: Flash endurance keeps getting better

Whalen: Data protection clouds vs. mega clouds

About us

Snapshot 2

Unified arrays destined for databases, user shares, VMs

➔ What apps will you run on your new unified array?



Databases
71%



**File data/
user
shares**
33%



**Virtual
server
support**
31%



**Web/
application
serving**
27%



**Virtual
desktops**
20%



**File
streaming**
15%



Big data
14%



**Science/
engineering**
10%

* MULTIPLE SELECTIONS PERMITTED

➔ Media mix: Anticipated types of storage devices in planned unified array

53%

Solid-state storage (flash)

49%

High performance hard disks

32%

High-capacity hard drives

25%

Mid-level performance hard disk

* MULTIPLE SELECTIONS PERMITTED

➔ Average TB capacity of new unified array

181



NetApp and Synology sew up top NAS honors

In our tenth Quality Awards NAS survey, a storage pioneer and a relative newcomer come out on top.

BY RICH CASTAGNA

NAS STANDS FOR network attached storage—but it’s also the first three letters of “nasty.” And judging by the tough evaluation standards of our survey respondents, NAS system vendors better sit up and take notice—or run for cover.

The respondents to our 10th [Quality Awards survey](#) for NAS systems, continuing a trend we saw developing last year, were tough critics of their network-attached storage vendors and systems in both the midrange and enterprise categories.

Across the board, scores were down from last year, when we saw some of the lowest marks ever doled out by respondents. It’s likely more a “honeymoon is over” situation than deep-seated disappointment, as the spiraling growth of file data is taxing all types and sizes of NAS boxes in companies big and small. So what we’re seeing is probably more about heightened expectations of network-attached storage system vendors than broad disappointment with the product category.


Against that backdrop, we had a couple of very tight races, with NetApp prevailing among enterprise NAS systems and a newcomer—Synology—rising to the top of the midrange class.

Our 505 valid survey responses yielded 616 product evaluations, with six enterprise and seven midrange product lines comprising our finalists in the survey for NAS systems.

Overall Rankings: And the winners are ...

Enterprise: With an overall average score of 6.10, NetApp outdistanced a cluster of rivals who finished second, third and fourth separated by only .04 points. For NetApp, this is its second [enterprise NAS](#) win, its first coming back in the third edition of this survey. And while NetApp's margin over second-place EMC (6.01) was relatively narrow, the NAS storage systems pioneer actually won quite handily, earning top marks in four of the five rating categories and coming in second in the fifth. EMC led in the final category and enjoyed an incredibly consistent cruise across all categories with only a .08 spread from its lowest to highest scores. DataDirect Networks (DDN) was just a hair behind EMC with an overall 5.99 for its best finish to date, with the same margin separating it from fourth-place IBM (5.97).


Enterprise NAS: Overall rankings



NAS NOTE: The group's average overall score was 5.93—just barely above last year's lowest ever at 5.92.

Midrange: The enterprise race was close, but the midrange competition was almost too close to call. The top three vendors—Synology, HPE and NetApp—were separated by .03 points. Synology's .01 point lead over second-place HPE was as close to a dead heat as you can get. Those three leaders divvied up rating category honors with Synology winning two, HPE snagging one and NetApp coming out on top in the remaining two (tied with EMC in one). Hands-down, this was the closest competition ever among midrange NAS storage systems, but they didn't stand out relative to past surveys. In fact, the group's average overall score was the lowest we've seen to date across the 10 sets of survey results. The group's average category scores were either the lowest or second lowest ever.

Midrange NAS: Overall rankings



NAS NOTE: Before this survey, the narrowest margin between overall first- and second-place midrange finishers was .11 points.

Castagna: Capacity, performance overshadow storage reliability

Toigo: NAS arrays morphing into cloud appliances

Software-defined storage—does hardware matter anymore?

Snapshot 1: For block-file flexibility, think unified arrays

The lowdown on cloud archiving

Snapshot 2: Key apps for unified arrays

NetApp and Synology win NAS users' kudos

Sinclair: Flash endurance keeps getting better


Whalen: Data protection clouds vs. mega clouds

About us

Sales-Force Competence: Caveat Emptor

Enterprise: NetApp scored highest for only one of the rating statements in the sales category, but it related to an important bottom-line aspect of a customer relationship: NetApp's 6.06 easily outdistanced the field for "My sales rep is easy to negotiate with." NetApp came in second or third on four of the five other statements. DDN led for three statements to place third in this category; it garnered a 6.33 for having a knowledgeable sales support team. EMC accounted for the highest marks on the remaining two statements, squeaking by IBM (6.06) with 6.07 for "My sales rep is flexible" and earning a 6.09 for sales reps who are knowledgeable about customers' industries. Fourth-place IBM fared best for flexible reps (6.06) and for "My sales rep understands my business" where its 6.03 was only .01 behind statement leader DDN.


Enterprise NAS: Sales-Force Competence



NAS NOTE: NetApp was the lone vendor to score above 6.0—the first time that has happened on a NAS storage systems survey.

Midrange. Looking at the midrange NAS results for sales competence, you can easily conclude that HPE has the customer thing down pat. Of the six statements in this category, HPE broke the tape first for five of them on its way to an average score of 6.18. HPE's strongest showings were for having a knowledgeable sales support team (6.67) and for reps who understand customers' businesses (6.24). The lone statement it didn't win—"My sales rep keeps my interests foremost"—second-place NetApp snagged with a 5.95. NetApp had four 6.0-or-better marks and two just shy of 6.0. EMC's performance was also relatively consistent with five scores near or above 6.0, and only stumbling a bit on the "My sales rep is easy to negotiate with" statement as did several others in the group.

Midrange NAS: Sales-Force Competence



NAS NOTE: The group's average for the sales category—5.68—tied last year's lowest-ever score.

Castagna: Capacity, performance overshadow storage reliability

Toigo: NAS arrays morphing into cloud appliances

Software-defined storage—does hardware matter anymore?

Snapshot 1: For block-file flexibility, think unified arrays

The lowdown on cloud archiving

Snapshot 2: Key apps for unified arrays

NetApp and Synology win NAS users' kudos

Sinclair: Flash endurance keeps getting better


Whalen: Data protection clouds vs. mega clouds

About us

Initial Product Quality: Rating the out-of-the-box experience

Enterprise: All of the enterprise NAS products fared well for initial quality, resulting in an exceptionally tight race, which NetApp won by a .03 margin over DDN and IBM. The others were also packed in tightly, yielding a .18 point spread spanning all six vendors—by far, the smallest spread we've ever seen from top to bottom. NetApp finished first without coming out on top on any rating statement. DDN's 6.44 was tops for the bellwether statement, "This product delivers good value for the money"; it also did well for ease of use (6.32). IBM prevailed on two statements, highlighted by a category-high 6.58 for products that install without defects. EMC and HPE grabbed one apiece: EMC for products requiring little vendor intervention (6.07) and HPE for getting products up and running easily (6.17).


Enterprise NAS: Initial Product Quality



NAS NOTE: The enterprise group's overall average for quality of 6.05 improved on last year's lowest ever.

Midrange: Synology customers are clearly pleased with the quality of the company's products, as the SMB-focused vendor notched a category average of 6.57 by racking up the highest marks for each statement (with one tie). That solid mark was underscored by a pair of 6.77s for installing without defects and products that get up and running quickly. HPE's 6.21 score was good for second place as it tied Synology for user satisfaction with the level of professional services that products require (6.33). HPE was the only vendor in addition to Synology to earn 6.0 or higher marks on all six category statements. Dell (6.04) rounded out the top three, a showing anchored by its best score of 6.30 earned for products that can get up and running quickly.

Midrange NAS: Initial Product Quality



NAS NOTE: As a group, initial quality was midrange systems best category, with a category average of 5.99.

Castagna: Capacity, performance overshadow storage reliability

Toigo: NAS arrays morphing into cloud appliances

Software-defined storage—does hardware matter anymore?

Snapshot 1: For block-file flexibility, think unified arrays

The lowdown on cloud archiving

Snapshot 2: Key apps for unified arrays

NetApp and Synology win NAS users' kudos

Sinclair: Flash endurance keeps getting better


Whalen: Data protection clouds vs. mega clouds

About us

Product Features: What's under the hood counts

Enterprise: NetApp was the only vendor to rack up 6.0-plus scores on all the rating statements for features and, in doing so, netted the highest marks for three of the seven statements. NetApp's best performances came for "Overall, this product's features meet my needs" (6.40) and a 6.37 for mirroring features. As usual, NetApp also did well for snapshotting features, garnering a 6.10, bested only by DDN's 6.15. In addition to that statement win, second-place DDN led the group for two others, with a 6.25 for replication and a 6.16 for interoperability. EMC leveraged a solid 6.33 for scalable capacity that won that statement and helped EMC nip IBM 5.99 to 5.94 to finish third. Although there were plenty of good results, as a group, the overall rating was only .01 point higher than last year, which was the lowest group score for features ever.


Enterprise NAS: Product Features



NAS NOTE: NetApp was the only vendor to merit a score better than 6.0 for management features (6.15).

Midrange: Synology matched its initial quality win with a second category victory in the features category to outpace NetApp (6.21) and HPE (6.04). Synology and NetApp were the only vendors to put up 6.0-plus scores for all features rating statements, and they split statement wins with Synology for five and NetApp copping the last two. Synology's strength was highlighted by sterling scores for overall feature satisfaction (6.65), management features (6.46) and snapshotting (6.38). NetApp earned ratings from 6.12 to 6.28, with statement wins for interoperability with other vendors' gear (6.22) and replication features (6.19). Most of HPE's grades were between 5.80 and 5.90, but it fared well for overall satisfaction with its feature sets (6.48) and management capabilities (6.38).

Midrange NAS: Product Features




NAS NOTE: Midrange NAS systems typically do well here, but this group's overall average was lower than all but one previous survey.

Reliability: Meeting expectations

Enterprise: NetApp (6.13) won this critical category, although it split individual statement honors with DDN and IBM. IBM and EMC had identical 5.99 scores to tie for second, followed by Dell (5.86). NetApp's best result was a group-best 6.27 for products that meet users' service level agreements; its second leading score—5.90—was picked up on the tough "This product requires very few unplanned patches" statement. IBM snapped up the highest statement mark of 6.32 for products that experience very little downtime; co-second-place finisher EMC didn't score highest for any single statement but was strongest for meeting service levels. Finishing farther back in the pack, DDN outranked the rest for two statements: 6.17 for non-disruptive patches and 6.16 for providing good upgrade guidance.


Enterprise NAS: Reliability



NAS NOTE: Overall, the group's lowest averages were for requiring few unplanned patches and non-disruptive patches.

Midrange: NetApp (6.17) dominated the reliability category without winning any statements, but scoring above 6.0 on all of them. NetApp flexed its muscles with a 6.32 for products that have very little downtime, and a pair of 6.24s for meeting service levels and requiring few unplanned patches. Second-place HPE (6.10) bucked the group's trend and fared well on the patches statements, with a 6.50 for its patch guidance and 6.45 for non-disruptive patching. Synology stumbled a bit on the non-disruptive patching statement with a 5.32, but still rode into third place with a group high 6.63 for "This product experiences very little downtime" as well high marks on the service level (6.40) and unplanned patches (6.25) statements.

Midrange NAS: Reliability



NAS NOTE: Like its enterprise siblings, the midrange group had its worst overall marks on two of the patching statements.

Castagna: Capacity, performance overshadow storage reliability

Toigo: NAS arrays morphing into cloud appliances

Software-defined storage—does hardware matter anymore?

Snapshot 1: For block-file flexibility, think unified arrays

The lowdown on cloud archiving

Snapshot 2: Key apps for unified arrays

NetApp and Synology win NAS users' kudos

Sinclair: Flash endurance keeps getting better


Whalen: Data protection clouds vs. mega clouds

About us

Technical Support: NAS safety nets

Enterprise: Four enterprise NAS system vendors finished with only a tenth of a point separating them, but EMC eked out the win to outrun NetApp, DDN and IBM—all clustered within .05 points of each other. Statements honors were well-distributed, but EMC led for three, including 6.18 for taking ownership of problems, 6.08 for resolving problems in a timely manner and 5.87 for problems that rarely require escalation (tied with IBM). NetApp tied IBM for the lead for supplying support as contractually specified (6.23) and led for having knowledgeable third-party partners (6.05). Third-place DDN earned the highest single statement mark, with a 6.50 for “Support personnel are knowledgeable” and stood out in the group with a 6.00 for its user training. HPE copped the last statement with a 6.22 for its documentation and support materials.


Enterprise NAS: Technical Support



NAS NOTE: Tech support is becoming an Achilles heel for enterprise NAS—the two lowest group averages in the last two surveys.

Midrange: Scores were low but the competition was high as both NetApp and EMC picked up 5.96 marks. HPE won a couple of statements in rolling up a 5.88 score, while Synology, nipping at HPE’s heels, beat the bunch on three statements on the way to a 5.86 rating. NetApp’s 6.27 for delivering support per contract was tops; it tied Synology for having knowledgeable third-party partners (6.09) and also tied EMC on taking ownership of problems (6.00). Co-winner EMC added a first for its training (5.97). HPE led for its documentation (6.38) and also prevailed with a 6.29 for its knowledgeable support personnel. Although it finished just out of the top three, DDN had the highest tallies for issues that rarely required escalation (6.19) and timely problem resolution (5.95).

Midrange NAS: Technical Support



NAS NOTE: Techies, not teachers? As a group, the midrange NAS storage system vendors scored their lowest for customer training.

Would you buy this product again?

As we've done on past Quality Awards surveys, we capped off our specific category rating statements with the more general statement, "All things considered, I would buy this product again." As often as not, the results of that final query seem to run counter to our users' responses in the category rating sections. But some things are clear: We know that if users have a uniformly terrible experience with a product or vendor, they're not likely to say they'd repeat the purchase, and we also know that familiarity can be comforting—as a result, so-so category scores might still yield good "buy again" marks.


Against that context, in the enterprise group, 92% of

HPE's users are willing to take the plunge again. That was despite HPE's less than lustrous overall performance. First- and second-place finishers NetApp and EMC also have loyal users with 89% and 86%, respectively, registering buy-again votes—a less-surprising result for sure.


Hewlett Packard Enterprise also led the midrange NAS group with a repeat-purchase score of 95%—somewhat less surprising as HPE came in second overall. EMC (93%) snuck past NetApp (88%) and a Synology-Hitachi tie (86%). ■

RICH CASTAGNA is TechTarget's VP of Editorial.

Enterprise NAS: Would you buy this product again?



Midrange NAS: Would you buy this product again?



About the Quality Awards

The *Storage* magazine/SearchStorage 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 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 to 8.00 scale, where 8.00 is the most favorable score. This is the 10th edition of the Quality Awards for NAS systems; there were 505 valid responses to the survey providing 616 sets of ratings for vendors' products/product lines.

Products in the survey: These products were included in the tenth Quality Awards for NAS survey. The number of responses for each finalist is shown in parentheses.

Enterprise NAS systems

- DataDirect Networks NAS Scaler/GRIDScaler/EXAScaler (26)
- Dell Compellent FS8600 (NAS)/EqualLogic FS7500/FS7600 (NAS) (42)
- EMC VNX 5000/7000/8000 NAS or Isilon X-Series (87)
- Hewlett Packard Enterprise (HPE) StoreEasy 3000/5000 or StorageWorks X9000 Storage Systems (37)
- Hitachi (HDS) NAS 4080/4100 or Essential NAS Platform 1000 Series or HNAS Platform 3000 Series*
- IBM N6000 or N7000 or Scale Out Network Attached Storage (SONAS) or Storwize V7000 Unified (31)
- NetApp FAS6000 or FAS8000 (both with NAS interface) (64)
- Oracle ZFS ZS3-4 or ZFS ZS4-4*

Midrange NAS systems

- Dell PowerVault NX Series (57)
- EMC VNXe 3000 Series NAS or Isilon S-Series (40)
- Exablox OneBlox*
- Hewlett Packard Enterprise (HPE) StoreEasy 1000 or StoreAll 8000/9000 X300/X500/X1000/X3000 Network (21)
- Hitachi (HDS) NAS 4040/4060 or HNAS AMS2000/1000/500/200 (21)
- IBM N3000 or N5000 (29)
- NetApp FAS2000 or FAS3000 (both with NAS interface) (69)
- Nexsan NST2000/NST4000/NST6000 (with NAS interface)*
- QNAP SS-ECxxxxx series, TS-ECx80 Series or TVS-ECx80 Series*
- Oracle ZFS ZS3-2*
- Synology RackStation RS2414 or RS3614 Series (28)

*Too few responses to qualify



Flash delivers reliability and performance

It's common knowledge that solid-state is fast—but its reliability should not be overlooked.

CONVERSATIONS ABOUT solid-state storage have focused primarily on performance—understandably so. As IT demands grow, so too do the demands for performance. As a result, interest in solid-state storage has grown regardless of the form—all-flash arrays, hybrid arrays, or even solid-state [server-side caching products](#). While much has been said about the obvious performance benefits, we are starting to see evidence that there may also be benefits to data reliability.

Enterprise Strategy Group (ESG) recently polled 373 storage decision-makers on a wide variety of enterprise storage technology topics. Respondents familiar with solid-state storage technology were asked to select the most important factor that led them to consider solid-state

storage. It wasn't surprising that the top response was improved performance, but the second most popular answer—improved solid-state drive (SSD) reliability—was less anticipated.

It may seem logical that fewer moving parts translates into higher reliability, but a quick glance at spec sheets to compare the mean time between failure or annualized failure rate statistics for SSDs and hard disk drives ([HDDs](#)) shows little difference.

CHEAPER PRICES + LOWER ENDURANCE = HIGHER SSD ADOPTION?

The fact that organizations are moving to solid-state storage for higher SSD reliability may be even more significant when you consider that the industry is currently in the middle of a massive undertaking to drive down the price of solid-state storage by reducing one aspect of SSD reliability, the endurance level. As solid-state deployments shift from higher-endurance SLC flash to MLC or TLC, the relative cost of storing data is reduced by increasing the capacity density an individual cell can store. As the level of data density increases, the amount of program-erase (P/E) cycles, or wear, an individual cell can endure subsequently decreases.

With solid-state storage, each write operation generates a [P/E cycle](#), and each P/E cycle damages the media to a

Home

Castagna: Capacity, performance overshadow storage reliability

Toigo: NAS arrays morphing into cloud appliances

Software-defined storage—does hardware matter anymore?

Snapshot 1: For block-file flexibility, think unified arrays

The lowdown on cloud archiving

Snapshot 2: Key apps for unified arrays

NetApp and Synology win NAS users' kudos

Sinclair: Flash endurance keeps getting better

Whalen: Data protection clouds vs. mega clouds

About us

Home

Castagna: Capacity, performance overshadow storage reliability

Toigo: NAS arrays morphing into cloud appliances

Software-defined storage—does hardware matter anymore?

Snapshot 1: For block-file flexibility, think unified arrays

The lowdown on cloud archiving

Snapshot 2: Key apps for unified arrays

NetApp and Synology win NAS users' kudos

Sinclair: Flash endurance keeps getting better

Whalen: Data protection clouds vs. mega clouds

About us

small extent. Different levels of solid-state technology, such as SLC or MLC, support a different number of P/E cycles over the life of the drive. That's the endurance level. SLC-based solid-state may, for example, support endurance levels as high as 100,000 write cycles. With TLC solid-state, the endurance levels may be as low as 1,000 write cycles per cell.

While endurance levels can be used to project how much data can be written to solid-state storage over the life of the media, it's also important to note that the calculation is not as simple as multiplying the endurance level by the capacity. Other software factors come into play. Behind the scenes activities, such as [wear leveling and garbage collection](#), alter the number of P/E cycles endured by each write operation. These impacts vary by storage manufacturer. Even if those activities increase the number of P/E cycles per write by a factor of five to one, a 1 TB SSD with an endurance of 1,000 write cycles could be able to support 200 TB of writes over its lifetime. While this represents a significant decrease from the higher endurance levels supported by SLC, the real-world impact may not be as significant.

The industry is still learning how many writes different applications require over a typical three- to five-year [storage lifecycle](#). The strong reliability perception of solid-state storage in the wake of decreasing endurance levels may simply be the result of initial endurance levels offered by SLC being over-engineered in the first place so that they exceeded the requirements of a typical application environment.

Another more plausible explanation may be that per-

ceived SSD reliability is more dependent on reducing the risk to data, rather than how often a particular technology component fails. If an SSD or an HDD fails, the data must be recreated in a separate location. [Solid-state storage performance](#) dramatically accelerates this rebuild process. So even in cases where solid-state storage and HDDs experience the same failure rates, the amount of time spent in a degraded state will likely be dramatically reduced with solid-state storage.

Additionally, the manner in which a device fails can also affect data unavailability or loss. Spinning hard drives are susceptible to a variety of unpredictable mechanical issues, such as scratches or head crashes, plus failures due to excessive heat or dust. Solid-state storage isn't susceptible to mechanical failures and can survive at higher temperatures. Common solid-state failures are often software-related, such as firmware issues, or from exceeding the endurance levels discussed earlier. But those types of failures are often more predictable and trigger remedial actions before a failure actually occurs.

Solid-state storage offers numerous capabilities that help reduce the risks of storing data, and organizations familiar with solid-state are catching on. As solid-state endurance levels and prices continue to decline, the resulting [increased adoption rate](#) will help the industry become better educated on the SSD reliability trade-offs for different application types. For now, however, many concerns about endurance may be unwarranted. ■

SCOTT SINCLAIR is a storage analyst with Enterprise Strategy Group in Austin, Texas.



Data protection clouds offer more focused services

AWS, Azure and Google can all beat the competition on price—but may come up short for data protection services.

CLOUD STORAGE IS becoming more popular, with total capacity growing by double-digit percentages every year. With costs as low as \$0.01 GB/month from vendors such as Amazon, Google and Microsoft, it's affordable and readily accessible.

But the big players aren't the only game in town if your primary needs are data protection and disaster recovery. Backup is still the prevailing use of cloud storage, and many traditional [backup appliance vendors](#) now provide their own cloud storage services. These smaller data protection vendors are surprisingly competitive with the hyperscale providers when total cost is considered, but they tend to focus on service and added-value functionality, instead of competing on cost alone.

With the [hyperscale](#) cloud storage vendors, in addition to the ongoing cost of storage, you'll also pay network egress fees as high as \$0.23/GB and data retrieval fees which vary widely, but tend to average about \$0.01/GB. These costs escalate as you store and use more data. So, it's important to understand your data usage patterns when evaluating whether or not it's beneficial to go with one of the major players.

Some data protection vendors may not actually [own and operate their own cloud](#)—some have chosen to partner instead of setting up their own data centers—but they all compete by offering a service model that's a natural extension of their basic business. Prices vary among vendors, but tend to run in the range of \$0.10 to \$0.15/GB/month. Typically there aren't any egress or data retrieval fees, so your bill is usually more predictable than one from a hyperscale vendor.

VENDOR CLOUD ADVANTAGES

A key advantage of using a data protection vendor's cloud is that it was built specifically for [data protection and disaster recovery](#). With enormous economies of scale, the big cloud vendors can offer very competitive prices, but their clouds are also designed for general-purpose uses. However, the big cloud vendors can offer multiple data center locations around the world, which some data

Home

Castagna: Capacity, performance overshadow storage reliability

Toigo: NAS arrays morphing into cloud appliances

Software-defined storage—does hardware matter anymore?

Snapshot 1: For block-file flexibility, think unified arrays

The lowdown on cloud archiving

Snapshot 2: Key apps for unified arrays

NetApp and Synology win NAS users' kudos

Sinclair: Flash endurance keeps getting better

Whalen: Data protection clouds vs. mega clouds

About us

Home

Castagna: Capacity, performance overshadow storage reliability

Toigo: NAS arrays morphing into cloud appliances

Software-defined storage—does hardware matter anymore?

Snapshot 1: For block-file flexibility, think unified arrays

The lowdown on cloud archiving

Snapshot 2: Key apps for unified arrays

NetApp and Synology win NAS users' kudos

Sinclair: Flash endurance keeps getting better

Whalen: Data protection clouds vs. mega clouds

About us

protection cloud vendors are unable to do.

Data protection vendors generally offer two levels of service. The first is basic “archiving” which is offsite storage capacity for the long-term retention of backup copies. To distinguish this offering from the more generic storage offered by the hyperscale providers, the archiving providers usually offer a retention service that allows customers to easily manage daily, weekly, monthly and yearly backups. And many data protection vendors offer a [physical cloud seeding option](#) to avoid having to move the initial data set over the wire.

The next service level addresses DR requirements. First off, replication copies local [backups to the cloud](#) where they can be accessed in the event of a failure. When a DR event occurs, a customer simply restores the appropriate backup data from the cloud and starts running again. Many data protection vendors also allow spinning up virtual machines in the cloud so operations can continue remotely while the local site is rebuilt.

When a user needs to download a lot of backup data to get their site back up and running, RTOs can stretch into days or weeks. To get around this problem, vendors of purpose-built backup appliances such as Datto and Unitrends give customers the option of having a replacement

appliance with the customer's data pre-loaded and shipped overnight to the downed site for quicker recovery.

TEST YOUR DR PLAN FREQUENTLY

Knowing that a [DR plan](#) will work when you need it most is critical, and you can only ensure that if you test the environment frequently. Data protection cloud vendors generally allow users to manually spin up and test VMs in the cloud. The better providers streamline the process by providing an automated method for testing VMs.

Cloud storage is a big market and there's room for many players with different business models. When [choosing a cloud provider](#), keep in mind that there are specific tools for specific jobs. hyperscale cloud vendors unquestionably provide the lowest cost for a stored bit, but when viewed holistically in terms of total cost, functionality specific to data protection and DR and white-glove service, especially during a disaster, some data protection vendors offer services that are far superior to those from the hyperscale cloud vendors. ■

JIM WHALEN is a senior analyst and consultant at Taneja Group with a focus on data protection.

Home

Castagna: Capacity, performance overshadow storage reliability

Toigo: NAS arrays morphing into cloud appliances

Software-defined storage—does hardware matter anymore?

Snapshot 1: For block-file flexibility, think unified arrays

The lowdown on cloud archiving

Snapshot 2: Key apps for unified arrays

NetApp and Synology win NAS users' kudos

Sinclair: Flash endurance keeps getting better

Whalen: Data protection clouds vs. mega clouds

About us



TechTarget Storage Media Group

STORAGE MAGAZINE

VP EDITORIAL Rich Castagna

EXECUTIVE EDITOR James Miller

SENIOR MANAGING EDITOR Ed Hannan

CONTRIBUTING EDITORS James Damoulakis, Steve Duplessie, Jacob Gsoedl

DIRECTOR OF ONLINE DESIGN Linda Koury

SEARCHSTORAGE.COM

SEARCHCLOUDSTORAGE.COM

SEARCHVIRTUALSTORAGE.COM

SENIOR NEWS DIRECTOR Dave Raffo

SENIOR NEWS WRITER Sonia R. Lelii

SENIOR WRITER Carol Sliwa

STAFF WRITER Garry Kranz

SITE EDITOR Sarah Wilson

ASSISTANT SITE EDITOR Erin Sullivan

SEARCHDATABACKUP.COM

SEARCHDISASTERRECOVERY.COM

SEARCHSMBSTORAGE.COM

SEARCHSOLIDSTATESTORAGE.COM

EXECUTIVE EDITOR James Miller

SENIOR MANAGING EDITOR Ed Hannan

STAFF WRITER Garry Kranz

SITE EDITOR Paul Crocetti

STORAGE DECISIONS TECHTARGET CONFERENCES

EDITORIAL EXPERT COMMUNITY COORDINATOR Kaitlin Herbert

SUBSCRIPTIONS

www.SearchStorage.com

STORAGE MAGAZINE

275 Grove Street, Newton, MA 02466

editor@storagemagazine.com

TECHTARGET INC.

275 Grove Street, Newton, MA 02466

www.techtarget.com

©2016 TechTarget Inc. No part of this publication may be transmitted or reproduced in any form or by any means without written permission from the publisher. TechTarget reprints are available through [The YGS Group](#).

About TechTarget: TechTarget publishes media for information technology professionals. More than 100 focused websites enable quick access to a deep store of news, advice and analysis about the technologies, products and processes crucial to your job. Our live and virtual events give you direct access to independent expert commentary and advice. At IT Knowledge Exchange, our social community, you can get advice and share solutions with peers and experts.

COVER IMAGE AND PAGE 8: BARANOZDEMIR/ISTOCK



Stay connected! Follow [@SearchStorageTT](#) today.