

STORAGE

MANAGING THE INFORMATION THAT DRIVES THE ENTERPRISE

Hot data storage technologies for 2017

These storage technologies are ready to step up and play a major role in your data center.



SNAPSHOT 1

Most enterprises confident about (and test) DR plan

EDITOR'S NOTE / CASTAGNA

If you love your data, set it free

CLOUD DR

Step-by-step guide to better cloud DR planning

STORAGE REVOLUTION / TOIGO

Time to take a holistic view of storage

SNAPSHOT 2

Mission-critical apps foundation of DR planning

HOT SPOTS / BUFFINGTON

How to align data protection and compliance

COMING SOON IN STORAGE

Stay tuned for "Products of the Year"

READ-WRITE / KATO AND BYRNE

Hyper-convergence meets requirements of private cloud





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: If you love your data, set it free

Toigo: Time to take a holistic view of storage

Hot data storage technologies for 2017

Snapshot 1: Most enterprises confident in (and test) their DR plan

Step-by-step guide to better cloud DR planning

Snapshot 2: Mission-critical apps foundation of DR planning

Buffington: How to align data protection and compliance

Kato/Byrne: Hyper-convergence meets private cloud needs

About us

If you love your data, set it free

We're drowning in data and don't have the tools we need to save ourselves.

COMPUTER USERS HAVE a hard time saying goodbye. Not to each other, or to you or me, but to the data they create, receive or just stumble upon. They like to keep everything because letting go is just so (sniff!) wrenching. Whether it's a Word doc; spreadsheet; or a .gif, .jif or .dif file, bidding adieu to data is too much to bear for many users.

But while your users spare themselves tearful ta-tas, it's your storage shop that has to shoulder the burden of the massive amounts of data that keep on growing.

A recent publication from Veritas—aptly titled “The Data Hoarding Report”—cited a recent survey sponsored by the vendor that reveals 62% of office professionals admit that they are (gasp!) [data hoarders](#). The IT pros ostensibly responsible for managing that data are even guiltier, however, as 81% are self-confessed data accumulators as well.

The report went on to note the issue with data storage

capacity isn't just where to put all this stuff—it's more of a case of why do we have all this stuff, because 86% of the data companies store is “[Redundant, Obsolete \[or\] Trivial](#).” That conclusion also offers up one of the best—and most accurate—acronyms we've seen in a long time: ROT.

The report notes that storing and managing all that ROT ends up costing businesses billions and billions of dollars. Whether or not you buy the report's nearly \$1 trillion cost estimate, you've got to figure ROT, whatever the cost, is a [black hole for storage budgets](#). And while data storage capacity has gotten pretty cheap—and is bound to get even more affordable—it still costs something.

Data storage capacity, meanwhile, is really only one part of the equation. In addition to providing a home for [voluminous data stores](#), you have to back it up and then create some copies of those backups to tuck away here and there for safekeeping. You'll also have to manage all of that to ensure copies are made properly and everything is up to date and in sync. For many companies, the amount of data requiring this kind of care has simply overrun the IT staff's ability to handle it all within a day that stubbornly still lasts only 24 hours.

DATA HOARDERS 'R' US

We're up to our ears in ROT because of knee-jerk reactions to compliance and an enthusiastic embrace of big

Home

Castagna: If you love your data, set it free

Toigo: Time to take a holistic view of storage

Hot data storage technologies for 2017

Snapshot 1: Most enterprises confident in (and test) their DR plan

Step-by-step guide to better cloud DR planning

Snapshot 2: Mission-critical apps foundation of DR planning

Buffington: How to align data protection and compliance

Kato/Byrne: Hyper-convergence meets private cloud needs

About us

data. When the likes of the [Health Insurance Portability and Accountability Act](#), [Sarbanes-Oxley](#), and a whole new generation of Securities and Exchange Commission and financial rules first raised compliance consciousness, there were essentially two camps.

One said “delete everything” to avoid uncomfortable situations like smoking gun email trails that put your company in litigation hot water. The other said that approach was shortsighted, so instead we should keep everything—you know, just in case.

While the latter camp was gaining converts, and we amassed more and more data, along came the allure of big data analysis. Big data is turning all of us into data hoarders. The whole [premise of big data analysis](#) is that even seemingly useless data, when matched up in some creative way with other bits of data, can yield useful information that will help your company squash the competition and achieve world dominance.

Unfortunately, a lot of that ROT will always be ROT, regardless of how cleverly it’s twisted and turned and manipulated. But merely the possibility of unearthing hidden gems is enough to inspire most companies to keep every bit and byte they create.

TIME FOR AN INTERVENTION

Now, what if we had a way to separate the ROT from truly useful data? It would be nice if vendors would look up from feverishly writing [SAN and NAS orders](#) long enough to see that most storage shops desperately need help

managing information. Seems like I go on this rant every year, but there are so few developments in this space that the issue just doesn’t go away.

A handful of companies are really trying to address the out-of-control data storage capacity issue. Data Gravity, still really a startup, is practically a poster child for the better data management movement. As is Actifio, which pioneered the concept of actually *not* creating a zillion copies of data, but rather just spin off one or two duplicates at a time and then create others as needed.

With its capability of storing detailed user-generated metadata, object storage still promises to be one of the key building blocks of a truly manageable, data-centric storage environment. Caringo’s FileFly app takes a big step in that direction by allowing storage managers to create policies embedded in metadata that describe the disposition of each piece of data. I know there are other vendors doing good things in this area, but there aren’t enough of them, and most of the big iron storage vendors are still basically ignoring the issue.

Users need help determining what to keep and what to ditch. The Veritas report listed the top reasons why users are data hoarders: 47% are afraid they’d delete something they might need later, and 43% simply don’t know what to delete or keep.

Storage pros can help users make those decisions, but they need the tools to turn those data storage capacity decisions into constructive action. ■

[RICH CASTAGNA](#) is TechTarget’s VP of Editorial.



Home

Castagna: If you love your data, set it free

Toigo: Time to take a holistic view of storage

Hot data storage technologies for 2017

Snapshot 1: Most enterprises confident in (and test) their DR plan

Step-by-step guide to better cloud DR planning

Snapshot 2: Mission-critical apps foundation of DR planning

Buffington: How to align data protection and compliance

Kato/Byrne: Hyper-convergence meets private cloud needs

About us

Time to take a holistic view of storage

Pay less attention to storage component stories and more to system-focused narratives.

THE COGNITIVE DISSONANCE engendered by the disparity in how various trade shows and conferences defined storage's place in the enterprise in 2016 is telling. Hypervisor computing and server-centric shows tended to cater to the app/dev crowd, with lots of woo about high-performance ([NVMe](#) and [NVMe over Fabric](#)) and agile (hyper-converged) storage. By contrast, shows catering to large-scale "enterprise-strength" data centers and "industrial farmers" of the cloud tended to skip over storage component stories in favor of a more data management system-focused narrative.

In the very first general session of IBM Edge 2016 in Las Vegas in September, for example, one of the very first statements to emanate from the show ringmaster was

this simple declaration: "Performance is not the result of any single component but of the successful and efficient integration of a complete system."

It was nice to hear such a principle echoed to all the app/dev'ers in the 5,500-person assembly. Because, frankly, I've grown tired of the limited perspective offered by those hawking component-level storage wares rather than a data management system.

COMPONENT-LEVEL OVERKILL

To illustrate my point, here's a summary of a few of pitches that crossed my desk over the past year:

1. 3D NAND flash (the first generation, at least) is now shipping, with second-generation chips not scheduled until late 2017, at the earliest. Truth is, while [3D NAND flash](#) is more capacious than 2D, it's also significantly more expensive to manufacture—a potential gating factor on uptake.

2. Nonvolatile memory express (NVMe) flash is beginning to show up in the market, but remains an answer to [problems most of us don't have](#). While theoretically accelerating the performance of workloads where bus congestion is a problem, by parallelizing I/O transports to individual flash chips, NVMe does nothing to parallelize the unloading of sequential I/O from x86 chips to the bus where

Home

Castagna: If you love your data, set it free

Toigo: Time to take a holistic view of storage

Hot data storage technologies for 2017

Snapshot 1: Most enterprises confident in (and test) their DR plan

Step-by-step guide to better cloud DR planning

Snapshot 2: Mission-critical apps foundation of DR planning

Buffington: How to align data protection and compliance

Kato/Byrne: Hyper-convergence meets private cloud needs

About us

latency develops in the first place. So the logjam in the I/O path slowing virtual machines (VMs) and other apps hosted in multicore chips remains unresolved by NVMe, although the technology may find use in large in-memory databases in the future.

3. Shingled magnetic recording (SMR) disk is available on many—primarily archival—disk targets. In truth, [SMR](#) is pretty much a bust. Without a lot of workarounds, users have to rewrite an entire piece of high-capacity media if a file or object is updated. Without such gyrations or a near-perfect knowledge about what data will remain unchanged for the next half decade or so, you're better off using high-capacity helium drives or tape.

4. Seagate continues to promise large capacity hard drives with good random access performance leveraging [heat-assisted magnetic recording \(HAMR\)](#), but the dates of arrival keep slipping. Who knows if bit-patterned media or acoustically assisted magnetic recording will ever deliver the goods?

SUCCESS STORIES

Unlike the component-level technologies listed above, tape continues to gain traction as a data management system, especially among cloud service providers who see it as the only way to store the mass quantities of data—measurable in thousands of exabytes—expected to come their way by 2020. Tape's on track to deliver between 140 and 220 TB of capacity (uncompressed) per cartridge by

the next decade, and is the only hope for cloudies looking to forestall the [Zettabyte Apocalypse](#).

I'd also like to tip my hat to Nutanix for demonstrating with its initial public offering (IPO)—one of a very few this year—that appliance-ized “server + hypervisor + software-defined storage stack + storage hardware” has a market among virtualization administrators and hypervisor vendors who know very little about storage. The IPO has caused many of the third-party software-defined storage (SDS) vendors allied with server makers like Cisco UCS, Dell, Fujitsu, Hewlett Packard Enterprise, Huawei or Lenovo to pine for their own soup-to-nuts appliance bearing their own logo art on the bezel.

I would caution that the SDS stack is very much in flux, though. Which could mean current hypervisor-centric stacks may end up becoming [Jurassic infrastructure](#) in a relatively short period of time.

INTEGRATION: THE ADULT IN THE ROOM

Now, if you went to the enterprise data center shows, you were likely treated to a more grown-up treatment of technology—one in which the whole data management system is greater than the sum of its parts. Less important than the latest flash performance metrics is how you design a balanced data management system that levels the differences between the speeds and feeds of discrete components to deliver the greatest overall efficiency.

For example, IBM is once again talking about “FLAPE” and “FLOUD,” both of which involve the use of high-speed storage to capture data and lower-speed storage for

Home

Castagna: If you love your data, set it free

Toigo: Time to take a holistic view of storage

Hot data storage technologies for 2017

Snapshot 1: Most enterprises confident in (and test) their DR plan

Step-by-step guide to better cloud DR planning

Snapshot 2: Mission-critical apps foundation of DR planning

Buffington: How to align data protection and compliance

Kato/Byrne: Hyper-convergence meets private cloud needs

About us

archiving. These are essentially “triangulated” architectures, much like hybrid cloud, intended to meet the goals of application developers for speed and agility, with no latency, and the ops crowd, which tends to prize resiliency, continuity and security in addition to performance.

In this vein, the storage tech I’m watching closest nowadays is cognitive data management (CDM). Essentially, CDM is the integration via cognitive computing of (1) an internet of things (IoT) approach to storage resource management, (2) an analytics approach to storage service management and allocation and (3) a fresh take on data lifecycle management.

The CDM platform sits over the storage infrastructure to direct data where it needs to go to receive the appropriate services (per policy linked to file/object metadata) in the most efficient manner. All data migrations (tiering) and copies (backups) are handled by the cognitive computing platform so as not to create latency in applications or VMs. “Universal translator” functionality enables many types of file systems and object systems to coexist and share data, with a global namespace providing all locational information for every stored bit across the infrastructure, even in the cloud.

I like what IBM has in mind for this technology, almost as much as I like what StrongBox Data Solutions is already doing with it under the moniker “StrongLINK.” IBM is

taking an IoT-meets-Watson approach to storage resource management in order to place data where it should go on tiered on-premises and cloud infrastructure storage components sporting the IBM moniker. StrongLINK,

IT IS CDM ... THAT HOLDS THE GREATEST POTENTIAL FOR ADDRESSING PRACTICAL STORAGE CHALLENGES THAT WILL COME TO A HEAD BY THE NEXT DECADE.

the CDM from Strongbox Data Solutions, takes a hardware-agnostic approach that will likely deploy more effortlessly in a heterogeneous shop.

In the final analysis, it is CDM—which goes well beyond the current boundaries of software-defined storage—that holds the greatest potential for addressing practical storage challenges that will come to a head by the next decade. ■

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



Hot data storage technologies for 2017

What storage tech will (or won't) set data centers ablaze in 2017? Check out our annual hot list.

BY RODNEY BROWN, RICH CASTAGNA, PAUL CROCKETTI,
GARRY KRANZ, SONIA LELII, JAMES ALAN MILLER,
DAVE RAFFO, CAROL SLIWA

IT'S OUR FAVORITE season. For the 14TH year running, we get to compile the data storage technology trends we believe will have the largest impact on the world of storage in the coming year. Welcome to Hot Techs 2017!

As in past years, there's nothing bleeding-edge or impractical here, only newer storage tech that's been proven practical. Hence, while our list of storage technology trends represents the best and brightest the storage industry has to offer, it only includes technologies you can buy and deploy today.

Climb aboard, fasten your seatbelts and get ready to discover our take on what technologies will have the most profound effect on storage shops in 2017.



CLOUD-TO-CLOUD BACKUP

Just because data's in the cloud doesn't mean it's adequately protected. Cloud-to-cloud backup fills that role by allowing enterprises to copy data stored on one cloud service to another cloud. Among the storage technology trends poised to have a big year in 2017, cloud-to-cloud backup vendors continue to add capabilities as user interest in the services grows.

Storage expert Brien Posey thinks cloud-to-cloud backup will likely become the norm by 2018. Why it's becoming popular is twofold. "First, backup technology is finally

Home

Castagna: If you love your data, set it free

Toigo: Time to take a holistic view of storage

Hot data storage technologies for 2017

Snapshot 1: Most enterprises confident in (and test) their DR plan

Step-by-step guide to better cloud DR planning

Snapshot 2: Mission-critical apps foundation of DR planning

Buffington: How to align data protection and compliance

Kato/Byrne: Hyper-convergence meets private cloud needs

About us

starting to catch up to the public cloud, making it more practical to do cloud-to-cloud backups,” Posey wrote in an email. “Second, and this is the big one, is the economic factor.”

For organizations moving data to the public cloud because it’s cheaper, backing up to another cloud provider makes economic sense and offers the advantages of off-site backup.

Posey still sees a role for local backups and storage, but thinks local storage requirements will likely decrease over the next few years.

“We may see cloud backup moving to be the de-facto standard, with snapshots retained on-prem for user error type restores,” storage expert and consultant Chris Evans wrote in an email. “Backup software vendors need [to] and have started to adapt. The biggest losers could be backup appliances in this instance.”

Specifically, with private-to-public cloud backups, there are tools to back up and restore applications into the cloud, saving money and improving operations, Evans added. And, he noted, software-as-a-service (SaaS) applications and data also need backing up, which is most easily done through the cloud.

As an internet-based software delivery model, SaaS has become a major option for businesses looking to remove the overhead of providing a range of IT services (e.g., email, collaboration, CRM and so on) themselves. So, as SaaS grows and serious work moves to the cloud, more organizations are recognizing the value of cloud-to-cloud backup, said storage consultant Jim O’Reilly.

“Rather than returning data to the in-house data center,

economics and operational efficiency suggest that a formal backup of cloud data into another cloud namespace is the best mechanism for totally protecting data in the cloud, whether from SaaS efforts or from owned applications,” O’Reilly explained. “Increasing comfort with the cloud as a site for running serious apps and the increasing use of SaaS will make this a must-have approach for larger IT operations through 2017.”

“BACKUP SOFTWARE VENDORS NEED [TO] AND HAVE STARTED TO ADAPT. THE BIGGEST LOSERS COULD BE BACKUP APPLIANCES IN THIS INSTANCE.” –CHRIS EVANS






Major players in SaaS-oriented cloud-to-cloud backup include Asigra Cloud Backup, Barracuda Cloud-to-Cloud Backup, Datto Backupify and Dell EMC Spanning. Barracuda, among those that enhanced their cloud-to-cloud backup platforms in 2016, reduced how long it takes to [complete incremental backups](#) of hosted versions of Microsoft applications.

Cloud-to-cloud backup is essential for protecting data created by SaaS applications. Notably, in May, a Salesforce outage prevented customers from accessing their data for several hours. Relatedly, although SaaS vendors perform their own backups, it’s only for their protection. If, for

(Continued on page 11)

Grading last year's hot storage tech picks

"Prediction is difficult, especially about the future," Yogi Berra is said to have noted once. So while we earned an additional A for our predictions over 2015, we had to give ourselves a couple of C's for picks that turned out not to be quite ready for primetime. That's still a pretty good batting average, though, one that we're eager to build upon with this year's lineup of hot storage technology trends for 2017.

GRADE	TECHNOLOGY	WHAT HAPPENED?
	Copy Data Management	Copy data management is practically the poster child of hot storage products. A year ago, there were just a few vendors with products along-side category pioneer Actifio. Now, every backup app rev includes CDM features.
	Software-Defined Storage Appliances	We're either geniuses or we're riding the wave of vendor hype, where just about every product is a "software-defined storage appliance." Either way, SDS created plenty of buzz and a whole lot of interest this past year.
	Object Storage	It might not be flying off the shelves just yet, but object storage has established itself as a solid alternative to scale-out NAS. If you didn't buy object in 2016, you will in 2017.
	Erasure Coding	We were ready to witness last rites over RAID with erasure coding nudging it out of the data protection business. As they say, reports of RAID's death have been exaggerated.
	Next-Generation Storage Networking	It's the end of 2016, and everybody's storage network is screaming along at 32 Gbps if it's FC, while Ethernet networks are jetting data around at 25 Gbps, 50 Gbps or even 100 Gbps—right? Well, not really. We missed on this one—maybe 2018?

Home

Castagna: If you love your data, set it free

Toigo: Time to take a holistic view of storage

Hot data storage technologies for 2017

Snapshot 1: Most enterprises confident in (and test) their DR plan

Step-by-step guide to better cloud DR planning

Snapshot 2: Mission-critical apps foundation of DR planning

Buffington: How to align data protection and compliance

Kato/Byrne: Hyper-convergence meets private cloud needs

About us

(Continued from page 9)

example, one of your SaaS users accidentally deletes something, without your own backup copy, you would have to pay Salesforce to restore the data—and the cost of that starts at \$10,000.

When considering a cloud-to-cloud or SaaS backup plan, Evans advised administrators to apply the same standards to backup and restore processes as they would to on-premises deployments. You can do this by, for example, testing how well a cloud-to-cloud backup service meets recovery time and recovery point objectives.



CONTAINERS

Propelled by technology advances that allow application microservices to directly consume persistent storage, container virtualization made significant inroads [in enterprise storage](#) in 2016. We fully expect open source containerization to remain one of the hot storage technology trends in 2017 as well, as container technology has advanced in key areas such as data protection, persistent storage consumption and portability.

Development and testing remain the dominant container uses, but experts say storage admins are learning to selectively manage Docker instances also. “The biggest change, in one word, is persistence,” said Henry Baltazar, a research director for storage at IT advisory firm 451 Research. “Containers have moved beyond ephemeral storage to persistent storage to hold data and protect your applications.”

Although open source Docker rival CoreOS released

the first commercial version of its Rocket runtime for Linux container this year, Docker remains the container kingpin—claiming more than 5 million software downloads and 650,000 registered users. Application teams use Docker to rapidly develop, ship and spawn applications inside containers. The closer those “Dockerized” applications move to real-time deployment, the greater the need to manage and provision Docker storage in stateless containers.

The Docker runtime engine was originally geared for Linux-based storage, but with Windows Server 2016 (released in September), Microsoft now allows admins to manage Docker virtualization on Windows servers. Microsoft also added its own container runtime to its latest server OS, enabling Microsoft shops to launch Windows-based containers on Windows Server 2016 server hardware or inside Hyper-V virtual machines.

While hardware virtualization is here to stay, containerization extends the concept further by virtualizing the operating system itself, thereby allowing workloads to share the underlying code and dependent libraries. Enterprises with highly virtualized storage could therefore deploy hundreds or perhaps thousands of containers on a single node, all running as lightweight instances.

Enabling [persistent storage in containers](#) is a top priority for storage vendors, said Greg Schulz, founder and senior advisor at IT infrastructure firm Server StorageIO. “Support for stateless and stateful containers will be the norm within 18 months. Containers for Docker, Linux and Windows will become a more viable unit of compute, joining physical bare metal along with other software-defined

Home

Castagna: If you love your data, set it free

Toigo: Time to take a holistic view of storage

Hot data storage technologies for 2017

Snapshot 1: Most enterprises confident in (and test) their DR plan

Step-by-step guide to better cloud DR planning

Snapshot 2: Mission-critical apps foundation of DR planning

Buffington: How to align data protection and compliance

Kato/Byrne: Hyper-convergence meets private cloud needs

About us

virtual machines and cloud machine instances.”

Legacy storage vendors Dell EMC, Hewlett Packard Enterprise (HPE), Hitachi Data Services, IBM and Net-App are differentiating their storage arrays to deploy and

manage Docker environments at a large scale. [Portworx](#), [Rancher Labs](#) and [StorageOS](#) are among container software startups tackling data management and secure migration of container data between server nodes.

Not quite hot ... yet

The five storage technology trends listed here, while heating up, simply weren't hot enough yet to make this year's Hot Tech lineup: Maybe next year ... or perhaps the year after that.

THE TECH	THE VERDICT ...
NVMe over Fabrics	NVMe suddenly became the biggest flash news around the middle of 2016—and products soon followed. But NVMe over Fabrics, the connective tissue that will best serve the new interface, will take longer to arrive. We're thinking it's more of a 2018 thing.
Flat Backup	Flat backup, backup-less backups, no-backup backup—whatever you call it, it's pretty cool because it sidesteps traditional backup processes and protects data way faster. But backup habits change very, very slowly.
Storage-Class Memory	Bridging persistent solid-state storage and volatile server memory, storage-class memory will occupy all those empty DIMM sockets and blur the line between dynamic RAM and storage. But not for a while.
Disaggregated Storage	Hey, didn't we separate storage from servers when shared storage was invented? No telling what vendors are cooking up with this new marketing lingo. And while the concept may be cool, it certainly isn't hot yet.
Predictive Storage Analytics	Hey, vendors—it's 2016, time to stop treating data like just a bunch of zeroes and ones! A couple of years ago, when we started to write about smarter storage with built-in analytics, there were just a couple of products. There are still just a couple of products.

Home

Castagna: If you love your data, set it free

Toigo: Time to take a holistic view of storage

Hot data storage technologies for 2017

Snapshot 1: Most enterprises confident in (and test) their DR plan

Step-by-step guide to better cloud DR planning

Snapshot 2: Mission-critical apps foundation of DR planning

Buffington: How to align data protection and compliance

Kato/Byrne: Hyper-convergence meets private cloud needs

About us

Red Hat added Red Hat [Gluster software-defined storage](#) as a persistent storage back end for Linux-based application containers. Even virtualization giant VMware has joined the fray. VMware Integrated Containers permit customers to run containers in vSphere.

As with any new technology, the hype cycle for containers has outpaced actual deployment. That means enterprises need to move cautiously as they connect the dots between containers and storage management, Baltazar said.

“Enterprises aren’t going to start running a bunch of Oracle database apps in a container, but there are areas where containers are important,” Baltazar noted. “Mobile apps and analytics are ideal for containers. You get really powerful resource allocation and the ability to do provisioning very rapidly.”



HIGH-CAPACITY FLASH

Samsung introduced a [15 TB 2.5-inch SAS solid-state drive](#) (15.36 TB actual capacity) in 2015. That drive, which started shipping last spring, is currently the largest capacity enterprise SSD available and is now beginning to show up in all-flash arrays from HPE and NetApp. Not to be outdone, Seagate unveiled a 60 TB SAS SSD drive in a 3.5 inch form factor at this year’s Flash Memory Summit and is now partnering with HPE to move it into mass production.

Increasing drive capacity to previously unimagined heights is the latest of the hot storage technology trends in the flash industry. If that sounds familiar, it should.

Just like the early days of hard disk drives, SSD vendors are now competing on high-density capacity levels, or who can cram the most and highest-density flash into a standard-sized drive.

“PRETTY SOON, [SSDs] WILL BE CHEAPER THAN SPINNING DISK. BY 2020, SPINNING DISK PROBABLY WILL BE DEAD.”

—RUSS FELLOWS, EVALUATOR GROUP

Samsung bases its large drives on its 512 GB V-NAND chip. The vendor stacks 512 V-NAND chips in 16 layers to forge a TB package, 32 of which [combine in 32 TB SSD](#). Samsung pointed out its 32 TB will enable greater density than Seagate’s 60 TB SSD because 24 2.5-inch drives can fit into the same space as 12 3.5-inch SSDs. Both Samsung’s 32 TB and Seagate’s 60 TB SSDs will ship sometime in 2017. So it looks like Seagate will be number one in density for a while, at least until Samsung packs its higher-density flash technology into a 60 TB drive of its own.

According to Russ Fellows, senior partner and analyst at Evaluator Group, it will eventually get to the point [where spinning disk becomes secondary to SSDs](#). “I think when density starts going [up] so fast and the dollars per gig go down, it’s going to be even cheaper than fast SATA drives,” he said. “Pretty soon, [SSDs] will be cheaper than spinning disk. By 2020, spinning disk probably will be dead.”

For Mark Bregman, senior vice president and CTO

Home

Castagna: If you love your data, set it free

Toigo: Time to take a holistic view of storage

Hot data storage technologies for 2017

Snapshot 1: Most enterprises confident in (and test) their DR plan

Step-by-step guide to better cloud DR planning

Snapshot 2: Mission-critical apps foundation of DR planning

Buffington: How to align data protection and compliance

Kato/Byrne: Hyper-convergence meets private cloud needs

About us

at NetApp, high-capacity SSDs offer huge reductions in space and savings in power and cooling.

“NetApp’s all-flash arrays that use high-capacity SSDs can now address customer use cases where such significant relief has historically been impractical,” Bregman wrote in a post on the NetApp Community blog. “From a space efficiency standpoint, you can’t beat the new high-capacity all-flash arrays, which give you up to 321.3 TB of raw storage in a single 2U form factor. That means a single 2U system using 15.3 TB drives can provide more than 1 petabyte of effective capacity.”

“To achieve the same [capacity] with even the highest density SFF hard disk drives,” he continued, “would require 52U of rack space and 18 times as much power.”



Solid-state drives have been the largest market for nonvolatile memory express (NVMe) specified storage to date. Latency-lowering, performance-boosting NVMe technology is one of the storage technology trends now starting to heat up in enterprise storage systems, however.

Shipments of server models as well as hybrid and all-flash storage arrays that leverage NVMe technology should more than double in 2017 due to more affordable price points and an expanding ecosystem, according to Jeff Janukowicz, a research vice president at IDC.

“NVMe adoption is still in its infancy,” Janukowicz wrote in an email. “However, we are at an inflection point, and we are beginning to see many more models that are starting to become available to the broader market.”

NVMe is an alternative to the age-old SCSI for transferring data between hosts and peripheral storage. [SCSI became a standard](#) in 1986 when HDDs and tape were the data center’s main storage media. The industry designed NVMe to support faster storage technology such as PCI Express (PCIe) SSDs. The NVMe specification, released in 2011, provides a streamlined register interface and command set to reduce the I/O stack’s CPU overhead.

Eric Burgener, a research director at IDC, singled out real-time, big-data analytics as one type of application workload that would need the level of performance that NVMe can deliver. Vendors targeting those high-performance workloads with NVMe-based storage products include Dell EMC (with [DSSD all-flash storage systems](#)), E8 Storage and Mangstor, he said.

Burgener predicted that NVMe in storage systems, also known as “rack-scale flash,” would remain a relatively small but growing market over the next several years. He said the array market would grow faster once off-the-shelf NVMe devices support enterprise capabilities such as [hot plug](#) and [dual port](#).

Also [on the horizon is NVMe over Fabrics](#) (NVMe-oF), enabling the use of alternate transports to PCIe to extend the distance over which NVMe hosts and NVMe storage devices can connect. NVM Express Inc., a nonprofit organization of more than 100 vendors, finalized the NVMe-oF specification in June 2016.

The long-term growth potential for NVMe is significant. Market research firm G2M Inc. forecasted the NVMe market would hit \$57 billion by 2020, with a 95% compounded annual growth rate.

Home

Castagna: If you love your data, set it free

Toigo: Time to take a holistic view of storage

Hot data storage technologies for 2017

Snapshot 1: Most enterprises confident in (and test) their DR plan

Step-by-step guide to better cloud DR planning

Snapshot 2: Mission-critical apps foundation of DR planning

Buffington: How to align data protection and compliance

Kato/Byrne: Hyper-convergence meets private cloud needs

About us

G2M also predicted that 60% of enterprise storage appliances and more than 50% of enterprise servers would have NVMe bays by the end of the decade. And it projected that, by 2020, nearly 40% of all-flash arrays would be NVMe-based with shipments of NVMe-based SSDs growing to 25 million units.



SOFTWARE-DEFINED STORAGE

There are many loosely defined terms used in the technology world in general, and for storage in particular. The least clear storage term may very well be one of the terms most bandied about nowadays, software-defined storage (SDS).

WhatIs.com [defined SDS](#) as “an approach to data storage in which the programming that controls storage-related tasks is decoupled from the physical storage hardware.” It’s a definition that’s flexible enough in interpretation to cover all sorts of technologies, however. The thing is those technologies actually share some features—mainly a focus on storage services rather than hardware and the use of policy-based management to allow for increased efficiency and reduced complexity when it comes to managing storage.

One of the main areas of confusion comes from how SDS is most often used in relation to virtualized environments. That isn’t a requirement, though. Thankfully, the market is slowly coming around to an agreement on exactly what SDS is. In order to be called software-defined storage, a product has to allow users to allocate and share storage resources across any workload, even if the storage

isn’t virtualized.

In [the March 2016 issue of Storage magazine](#), storage analyst Marc Staimer of Dragon Slayer Consulting established four basic categories of SDS: hypervisor-based, hyper-converged infrastructure (HCI), storage virtualization, and scale-out object or file.

VMware practically owns the hypervisor-based category with vSphere Virtual SAN. As the second-oldest SDS category with products on the market, it is well-established among storage technology trends. However, it is also restricted to hardware that VMware has determined compatible.

The bulk of the SDS market resides with HCI SDS, through products offered by giants like Cisco, Dell EMC and IBM, as well as startups like Nutanix and SimpliVity. The positive aspect of HCI SDS is how everything you need for your storage infrastructure is included and designed to work together. The negative is that means only resources in the HCI get to take advantage of the benefits of SDS.

Storage virtualization SDS is the grandfather of all varieties of software-defined storage, and DataCore, Dell EMC, IBM and even Microsoft offer products in this category. But just because it is the oldest doesn’t mean there are no younger players (including NetApp, Nexenta Systems and StarWind Software) in the game.

Newest to market is scale-out object or file SDS. Even in this area, in addition to newer companies like Scality, you have giants competing, like IBM’s Spectrum Storage and Red Hat with its OpenStack and Ceph-based products.

The continuing drop in price for commodity hardware

Home

Castagna: If you love your data, set it free

Toigo: Time to take a holistic view of storage

Hot data storage technologies for 2017

Snapshot 1: Most enterprises confident in (and test) their DR plan

Step-by-step guide to better cloud DR planning

Snapshot 2: Mission-critical apps foundation of DR planning

Buffington: How to align data protection and compliance

Kato/Byrne: Hyper-convergence meets private cloud needs

About us

is spurring greater adoption of SDS. With SDS, there's less of a need for specific hardware to get high levels of performance, particularly when a performance-focused type of software-defined storage is used in the enterprise. In addition, because object storage has become less of a cloud-only platform for storing data, that particular version of SDS has quickly gained ground in the data center.

In [a recent interview](#), CEO and founder of SDS startup Hedvig Inc., Avinash Lakshman, explained why he thinks scale-out SDS is a hot technology that will continue to grow rapidly.

“The ROI is pretty simple because hardware costs are going nowhere but down. People like Amazon, Google and all these large internet-scale companies are obviously going that route. It's forced the enterprise to take a look at them and ask the question, ‘If they can do a lot more with a lot less, why can't we?’”



32-GIG FC

Most of today's storage technology trends work against Fibre Channel (FC). Hot new architectures like hyper-convergence and the cloud use Ethernet with little need for FC. Ethernet also dominates storage for file-based unstructured data, which is growing much faster than the block-based structured data that often requires Fibre Channel SANs. There are barely a handful of FC networking companies left, and they all now support Ethernet as well.

On the flipside, there's flash. FC vendors and fans are counting on the rapid emergence of all-flash storage SANs

to keep FC relevant, especially while the protocol transitions from 16 Gbps equipment to [32 Gbps switching and adapters](#). That transformation will likely make big inroads in late 2017, providing a new wave for FC to ride.

ADOPTION IS EXPECTED TO PICK UP WHEN STORAGE ARRAY VENDORS SUPPORT 32-GIG, WHICH SHOULD HAPPEN IN A MEANINGFUL WAY NEXT YEAR.

“Storage performance bottlenecks are moving out of arrays and into the storage network, so Fibre Channel will remain the data center storage protocol of choice for the next decade,” Gartner research director Valdis Filks and research vice president Stanley Zaffos wrote in a recent report called “The Future of Storage Protocols.”

Solid state storage—today mostly flash with NVMe arriving and [3D XPoint on the horizon](#)—provides greater throughput and lower latency than hard disk media. Storage networks need more bandwidth if storage media is to reach those performance peaks, though.

Filks and Zaffos added that 16 Gbps—and even 40-gigabit Ethernet—will be too slow to keep up with the next generation of solid-state storage. So they recommend moving to 32-gig FC and 100-gig Ethernet within five years.

Last year, we saw early 32-gig FC products hit the

Home

Castagna: If you love your data, set it free

Toigo: Time to take a holistic view of storage

Hot data storage technologies for 2017

Snapshot 1: Most enterprises confident in (and test) their DR plan

Step-by-step guide to better cloud DR planning

Snapshot 2: Mission-critical apps foundation of DR planning

Buffington: How to align data protection and compliance

Kato/Byrne: Hyper-convergence meets private cloud needs

About us

market, including switches from Brocade and Cisco and adapters from Broadcom and QLogic. Adoption is expected to pick up when storage array vendors support 32-gig, which should happen in a meaningful way next year.

Broadcom further consolidated the FC networking industry with a \$5.9-billion acquisition of Brocade that is expected to close in early 2017.

Broadcom CEO Hock Tan said his company will continue to invest in FC. “If you believe in all-flash, you have to believe in Fibre Channel,” he said on a conference call for the Brocade acquisition. “Even today, iSCSI and Ethernet does not offer that. We expect this market to remain relatively stable as it supports private data centers with a large installed base of Fibre Channel SANs that are constantly upgraded.”

Flash and the move from 16 Gbps to 32 Gbps are expected to drive many of those upgrades.

Gartner recommended waiting a year after general

availability before moving to the latest protocol. That would provide time for prices to come down, early kinks to be worked out, and allow for full compatibility of storage and servers with switching. Early 32-gig FC buyers can still use their 16-gig and 8-gig storage and servers with the new switching products, however. They also can check out [roadmaps to 64-gig and 128-gig FC equipment](#).

Adarsh Viswanathan, senior manager of product management for Cisco’s storage group, expects 32-gig FC to take off late in 2017, after storage array vendors fully embrace it. “A lot of the big customers we talk to are in production with all-flash arrays and attaching mission-critical workloads on flash through Fibre Channel. Flash vendors can fill the 32-gig pipe. We expect it will get big traction in the second half of 2017.” ■

RODNEY BROWN, RICH CASTAGNA, PAUL CROCETTI, GARRY KRANZ, SONIA LELII, JAMES ALAN MILLER, DAVE RAFFO and CAROL SLIWA are all members of TechTarget’s Storage Media Group.

Snapshot 1

Most enterprises confident about (and routinely test) DR plan

Home

Castagna: If you love your data, set it free

Toigo: Time to take a holistic view of storage

Hot data storage technologies for 2017

Snapshot 1: Most enterprises confident in (and test) their DR plan

Step-by-step guide to better cloud DR planning

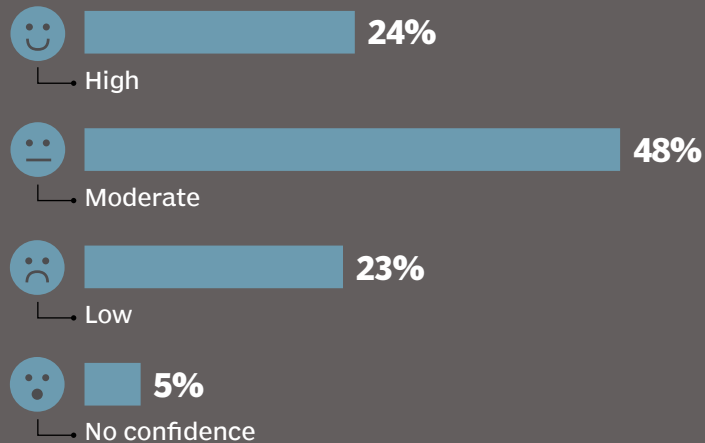
Snapshot 2: Mission-critical apps foundation of DR planning

Buffington: How to align data protection and compliance


Kato/Byrne: Hyper-convergence meets private cloud needs

About us

How confident are you in your organization's disaster recovery plan?



How many times per year do you test your disaster recovery plan?



60
Percentage of respondents with DR plans in place for multiple physical locations

A guide to a better cloud DR plan

DRaaS has quickly become an ideal way for enterprises to tap the cloud to solve their disaster recovery challenges.

BY GEORGE CRUMP



ALASHI-ISTOCK

HOME

DISASTER RECOVERY REMAINS an ongoing challenge for storage professionals. Because, unlike many other areas of IT, circumstances surrounding DR have become even more complex and difficult to get a handle on in recent years.

First, the definition of a “disaster” has expanded to include almost any service interruption. Second, the causes of a disaster are many. There are the natural disasters that we all hear about (earthquakes, floods, hurricanes) and there are man-made disasters like cyberattacks, ransomware, accidental user error or corporate sabotage. Third, the amount of data we have to consider in case of a disaster has grown exponentially. And, fourth, user expectations for a rapid recovery have increased tremendously. Today, users expect little to no interruption in their key applications.

Meanwhile, as a backdrop to these difficulties, most organizations’ [DR budgets](#) remain stagnant. As a result, many enterprises—from the smallest to the largest—are looking to their cloud disaster recovery plan to eliminate these problems while reducing overall spending on DR.

BACKUP, ARCHIVE AND DR IN THE CLOUD

Understanding the cloud is a service delivery model (see “[What is the cloud?](#)”) helps identify which [data protection processes](#) and cases it’s good and not so good for. Most

Home

Castagna: If you love your data, set it free

Toigo: Time to take a holistic view of storage

Hot data storage technologies for 2017

Snapshot 1: Most enterprises confident in (and test) their DR plan

Step-by-step guide to better cloud DR planning

Snapshot 2: Mission-critical apps foundation of DR planning

Buffington: How to align data protection and compliance

Kato/Byrne: Hyper-convergence meets private cloud needs

About us

data protection processes have three basic components: backup, archive and recovery.

■ **Backup** is the copying of data with the goal of having a version of that data available for recovery in the case of a failure or loss. Most organizations keep backups for a finite number of years (three to seven is typical), so [leveraging the cloud for backup](#) means renting the capacity to store each day's backup for at least that length of time. Depending on the amount of data you need to store, this may or may not be as cost-effective as doing it yourself.

■ **Archive** is the long-term retention of data. This should be the place where a final copy of data is stored—ideally, limited to only two copies of each piece of data in two separate locations and, preferably, two formats. The stronger an organization's [archive strategy](#), the smaller the backup footprint will be in terms of the amount of data protected and the length of time it takes to recover that data. Most organizations need to keep archive data for at least seven years, potentially much longer. Renting storage for decades isn't going to be cost-effective for most organizations, so the cloud may not be the right destination for archive data.

■ **Recovery**, specifically disaster recovery, almost always requires the most recent copy of data. Any recoveries outside of that "most recent" time frame should come from an archive, if at all possible. The cloud is a great location for DR, as only the most recent copy of data has to be stored there, and most—if not all—providers allow you to use

What is the cloud?

FIRST, LET'S BE clear, the cloud isn't some magical place that solves all IT problems. It's a data center that has to struggle with the same laws of physics as private data centers. The advantage is that a cloud provider is an entire organization focused solely on delivering IT services with—for the most part—top-notch data center designs.

They also have the advantage of economies of scale, meaning the sheer quantity of IT products purchased gives cloud providers a significant cost advantage. Combine these economies with most heavy investments in process and automation, and you've got operational costs typically much lower on a per-unit basis.

Remember, the cloud is primarily an IT delivery model. Instead of buying IT products and services upfront, enterprises rent them on a periodic basis. The rental nature of IT services means that the cloud is very attractive for temporary use and less attractive for more permanent cases. ■

their compute resources in the event of disaster. This saves on the cost of maintaining your own server and storage hardware at a remote DR site, all without adding too much to backup-disk capacity costs. Consequently, cloud-based DR, or [disaster recovery as a service](#) (DRaaS), has quickly

Home

Castagna: If you love your data, set it free

Toigo: Time to take a holistic view of storage

Hot data storage technologies for 2017

Snapshot 1: Most enterprises confident in (and test) their DR plan

Step-by-step guide to better cloud DR planning

Snapshot 2: Mission-critical apps foundation of DR planning

Buffington: How to align data protection and compliance

Kato/Byrne: Hyper-convergence meets private cloud needs

About us

become an ideal way to solve DR challenges.

Let's take a look at the steps you should take and the things you should consider to develop the most effective cloud disaster recovery plan for your organization.

STEP 1: GET DATA TO THE CLOUD

The first step in building a cloud disaster recovery plan is to get your data to your provider, of course, but also to remember to optimize costs by controlling how much data is stored there—preferably just the latest copy. Also, while cloud data protection options are numerous, they can be typically broken down into two types: products that back up data and those that [replicate data](#). The difference between them is in how data is stored.

Most cloud backup products store data in a propriety backup format. In the event of disaster, data has to be extracted and moved into a format enterprises can access by a virtual machine. Most [cloud backup products](#) leverage an on-premises appliance that captures all backups first and then copies changed data to a cloud location. Nearly all replication products replicate data to the cloud, but—by contrast—store it in a native file-system format that's immediately accessible in the event of a primary-site failure. Customers can even choose to have this data stored on high-performance cloud storage to make “return to operations” even faster.

With backup or replication, there's the challenge of initial data seeding. It can take hours, days or even weeks to transmit the data that will create the cloud baseline that the backup or replication software will need to compare

against. To speed up the process, some cloud vendors ship a ruggedized high-capacity NAS to the customer. The foundational data set is copied onto it and then the NAS is shipped back. Ideally, more [cloud providers](#) would use tape, which is easier to ship and more cost-effective for seeding, however. Once data is in the cloud, daily updates to that data generally happen easily, while technologies like compression, deduplication and changed block replication significantly thin the amount of data that has to move across the network.

There is a third method to move data to the cloud—run production data in the cloud itself. This either involves using an on-premises cache so that local applications don't suffer latency delays or shifting the entire workload to the cloud itself. While the operational gains of placing both primary storage and secondary storage in the cloud may outweigh the costs, customers have to be comfortable with moving all their data to the cloud. And, although this method also eliminates the ongoing transfer of data to the cloud, you have to make sure your cloud vendor provides acceptable resiliency for the data it's storing.

STEP 2: DECLARING A DISASTER IN THE CLOUD

At least once in every IT professional's career, a disaster will occur. In fact, it is more likely today than ever since the definition of “disaster” has expanded from a data center wipe to an important application becoming unavailable. If a disaster is isolated, meaning that it only impacts one workload, then failing over both compute and data to the cloud provider is unnecessary and, frankly, unwanted.

Home

Castagna: If you love your data, set it free

Toigo: Time to take a holistic view of storage

Hot data storage technologies for 2017

Snapshot 1: Most enterprises confident in (and test) their DR plan

Step-by-step guide to better cloud DR planning

Snapshot 2: Mission-critical apps foundation of DR planning

Buffington: How to align data protection and compliance

Kato/Byrne: Hyper-convergence meets private cloud needs

About us

Typically, an organization is better off recovering on premises. Some cloud backup products can leverage their appliance in the on-premises data center to host both the fallen application's data and, in some cases, be the compute needed to drive a virtual machine version of that application.

IN MOST CASES, YOU CAN RESTORE DATA FROM THE CLOUD, ALBEIT SLOWLY, WHILE THE PROVIDER CONTINUES TO HOST YOUR APPLICATION.

Replication-based DRaaS products should be set up to have a local storage target in addition to the cloud. That way, the organization can recover locally or in the cloud.

For more extensive disasters, where your data center becomes unavailable, true failover to the cloud is required. Here, the first step after disaster is to start all the services in the cloud the high-priority applications require to operate, such as DNS and directory services, then all the servers that make up that application. Last, adjust the networking configurations so that users logging in can seamlessly access the now cloud-hosted application.

Obviously, testing the DR process before it's actually needed is critical to make certain all these steps work, especially the networking changes. It's also important to factor end-user conditions into the changes. They likely won't be at a central office and may be logging in from

a virtual private network connection at a coffee shop or home.

The process is similar for enterprises that put all their data in the cloud but keep compute on premises. The new network routing issues are the same, but since the data is already in the cloud, you merely need to start the applications alongside of that data to get up and running.

STEP 3: THE RETURN

At some point, you will want to leave the cloud and return to normal on-premises operations. The cloud exit is one of the more difficult aspects of leveraging a cloud disaster recovery plan, as all the techniques IT uses to facilitate the daily transfer of data to the cloud won't work here since there is no baseline available for comparison.

In most cases, you can restore data from the cloud, albeit slowly, while the provider continues to host your application. Once data's done transferring, perform a quick data sync and then switch operations back to your enterprise's primary data center. Depending on the amount of data to be transferred and available bandwidth, this transfer could take days or even weeks—all while your organization is paying a surcharge on the compute it's using in the cloud.

DRaaS providers should offer the ability to mass ship data directly to a customer's new data center. As mentioned before, this can be accomplished through ruggedized NAS or tape. This would allow the baseline copy to be created much faster while affected applications still run in the provider's cloud.

The ordeal of return illustrates another advantage of the

Home

Castagna: If you love your data, set it free

Toigo: Time to take a holistic view of storage

Hot data storage technologies for 2017

Snapshot 1: Most enterprises confident in (and test) their DR plan

Step-by-step guide to better cloud DR planning

Snapshot 2: Mission-critical apps foundation of DR planning

Buffington: How to align data protection and compliance

Kato/Byrne: Hyper-convergence meets private cloud needs

About us

“already in the cloud” case. Since no data has to be moved, you can start applications in the new data center—with data cached locally—as systems begin to access data in the cloud. Only active data has to be copied on premises.

CONCLUSION

In terms of storage, the cloud offers the greatest value and cost-effectiveness for disaster recovery. That’s because DR storage is lower in capacity than other data protection

products and seldom read, and the cloud provides access to compute that can leverage data for a quick restart of mission-critical applications in event of disaster. A cloud disaster recovery plan also eliminates much of the cost of an organization’s DR strategy, as you only have to pay for compute resources when testing your DR plan or when an actual disaster occurs. ■

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

Snapshot 2

Quick recovery of mission-critical apps is the foundation of disaster recovery planning

Home

Castagna: If you love your data, set it free

Toigo: Time to take a holistic view of storage

Hot data storage technologies for 2017

Snapshot 1: Most enterprises confident in (and test) their DR plan

Step-by-step guide to better cloud DR planning

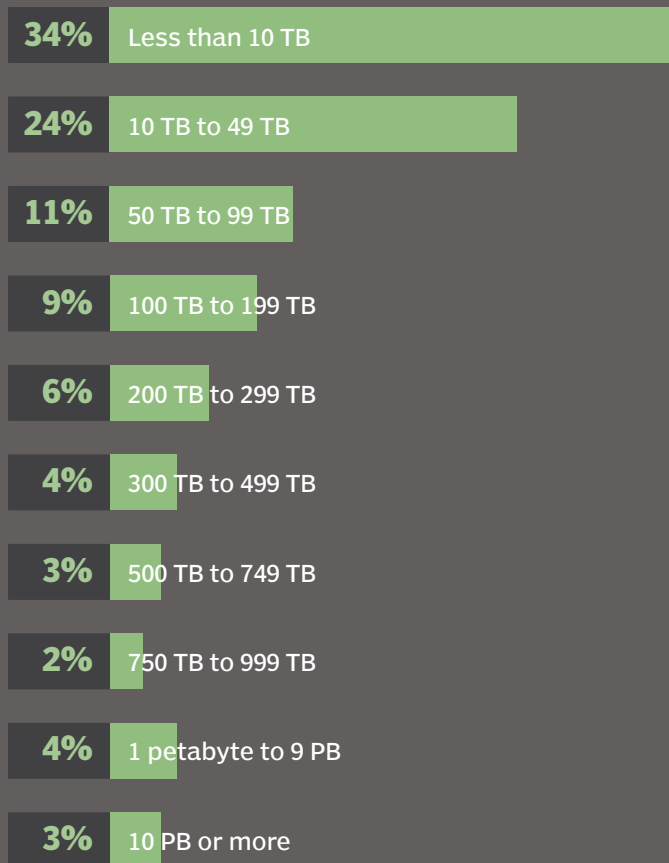
Snapshot 2: Mission-critical apps foundation of DR planning

Buffington: How to align data protection and compliance


Kato/Byrne: Hyper-convergence meets private cloud needs

About us

→ How many terabytes of data will you need to recover quickly in the event of a disaster?



→ How much of your company's overall workload is protected by your disaster recovery plan?



Percentage of respondents using cloud DR services as part of their disaster recovery plan

34



Home

Castagna: If you love your data, set it free

Toigo: Time to take a holistic view of storage

Hot data storage technologies for 2017

Snapshot 1: Most enterprises confident in (and test) their DR plan

Step-by-step guide to better cloud DR planning

Snapshot 2: Mission-critical apps foundation of DR planning

Buffington: How to align data protection and compliance

Kato/Byrne: Hyper-convergence meets private cloud needs

About us

How to align data protection and compliance

While many compliance mandates discuss data protection, not all are as prescriptive as you'd think.

COMPLIANCE MANDATES DICTATE how you should conduct operations or, at the very least, define how someone else believes you should. That “someone else” could be senior execs, a board of directors, the legal department, or external agencies or industry groups.

With so many potential stakeholders, many businesses end up having to follow multiple unaligned regulations. A publicly traded pharmaceutical maker in the U.S., for example, may be simultaneously subject to retention rules handed down by the Securities and Exchange Commission, the Food and Drug Administration, the Department of Agriculture and the Drug Enforcement Agency.

Also, some regulations include language related to data protection methods, but primarily focus on the outcome,

not the method. In other words, these regulations won't instruct you to use specific software products or [storage media](#); they just tell you what your retention durations and destruction standards must be.

Whatever a regulation says or doesn't say regarding the specifics of data protection requirements, if you [rely on backup administrators](#) to ensure adherence, you've already failed. Typical backup admins don't have enough visibility into every task related to data protection requirements and internal policy to ensure compliance. For example, a vAdmin might oversee [protection of certain virtual machines](#). Or perhaps a database administrator may be the person who ensures a particular mission-critical database properly replicates off-site. Backup admins sometimes aren't even aware of every regulation that needs to be followed. That's not to say backup admins aren't a crucial part of compliance. But other equally important participants factor into the equation:

- **Application and workload owners** (e.g., database administrators). They manage how applications should be maintained for uptime by following [best practices for data management](#).
- **Infrastructure managers** (e.g., vAdmins or IT operations staff). They oversee how servers are provisioned and maintained.

Home

Castagna: If you love your data, set it free

Toigo: Time to take a holistic view of storage

Hot data storage technologies for 2017

Snapshot 1: Most enterprises confident in (and test) their DR plan

Step-by-step guide to better cloud DR planning

Snapshot 2: Mission-critical apps foundation of DR planning

Buffington: How to align data protection and compliance

Kato/Byrne: Hyper-convergence meets private cloud needs

About us

■ **Business unit leaders and non-IT employees.** They supervise various teams who rely on the IT infrastructure and create or use regulation-affected data.

■ **Compliance teams.** They comprise HR staff, internal legal staff, and external auditors and attorneys.

Each group plays an indispensable role in achieving compliance because, in most organizations, one group's perspective is usually neither wholly understood nor a primary focus of the others. Some people understand the regulations. Others understand the data. Some know the platforms. Others know how data and platforms alike must be protected and recoverable to ensure availability.

THE KEY IS COMMUNICATION

The single most important part of compliance is *communication*.

The process starts with compliance teams becoming educated on applicable internal or external regulations and desired outcomes. Particularly when outside mandates are in play, consider developing a comprehensive table—a requirements “superset”—listing each mandate. Revalidate the list semiannually to ensure it remains accurate.

Some regulations may conflict with others, but that is a rare case when it comes to a [single data set](#). (For example, a federal law may conflict a bit with a state law. In that case, the organization's legal team might need to determine which regulation supersedes the other or otherwise establish how to proceed in a manner that achieves

“regulatory harmony” defensibly.)

After identifying what adherence outcomes must be achieved—for example, ascertaining that a certain record type should be preserved for seven years instead of six—communication between the compliance team and other groups within the organization should continue. Plan to provide some education regarding what needs to be done and why it's important (i.e., conduct training sessions and provide participants with referenceable handouts that describe relevant regulations and offer business-level context regarding why those rules exist), then trust the other teams to develop viable ways to achieve compliance—for example, by updating the records management procedures that are specific to their business units.

ADDITIONAL REGULATION-RELATED BEST PRACTICES

Now that you've collectively developed plans to align your data protection requirements and [compliance efforts](#), two more actions come into play:

1. Inspect what you expect. There are two ways to test for compliance: “looking to pass” and “looking to improve.” Regarding the former, if you only strive for green checkmarks, you may find them, sure, but you could still fail on a larger scale. Instead, seek out those red X's, albeit in a safe and collaborative manner. By doing that, you'll have a better chance of passing your audits and achieving what you need to *when you really need to*.

2. Seek out outside expertise. The difference between

Home

Castagna: If you love your data, set it free

Toigo: Time to take a holistic view of storage

Hot data storage technologies for 2017

Snapshot 1: Most enterprises confident in (and test) their DR plan

Step-by-step guide to better cloud DR planning

Snapshot 2: Mission-critical apps foundation of DR planning

Buffington: How to align data protection and compliance

Kato/Byrne: Hyper-convergence meets private cloud needs

About us

knowledge and wisdom is experience. Your teams may have the composite knowledge to align IT and business processes with applicable mandates, but outside experts with proven compliance experience will help you effectively assemble all the pieces of the puzzle when it comes to data protection requirements and compliance. For example, they can share their own interpretations of how a given mandate has influenced IT and business processes at companies elsewhere.

To sum up, don't rely solely on a backup admin to maintain an entire organization's compliant status. Emphasize

communication and follow through with training. Don't despair if you encounter a red X [during an internal audit](#); it simply means you have the chance to fix that issue before it becomes an expensive problem. And take advantage of the real-world knowledge of outside experts. They can help you avoid pitfalls you didn't even know existed.

In this way, you'll confidently implement and ultimately succeed in ensuring compliance for your organization. ■

JASON BUFFINGTON is a senior analyst at Enterprise Strategy Group. He focuses primarily on data protection, as well as Windows Server infrastructure, management and virtualization. He blogs at [CentralizedBackup.com](#) and tweets as [@Jbuff](#).

READ / WRITE
JEFF KATO & JEFF BYRNE



Hyper-converged meets requirements of private cloud

Infrastructure choice and integration are fundamental to taking advantage of all the cloud has to offer your organization.

ENTERPRISES LOOKING TO benefit from the cloud are often reluctant to deploy business-critical apps and data in the public cloud due to concerns about availability, security and performance. Most IT managers consider a private cloud platform a more comfortable choice, given the superior visibility into and control over IT infrastructure and peace of mind that comes from housing critical assets on the inside.

Application owners are often skeptical about whether a private cloud platform will really provide the increases in business agility promised by vendors, however. In a similar vein, they're also wary about whether, and over

what timeframe, they'll realize the ROI required to make deploying a fully functional and expensive private cloud platform worthwhile. Meanwhile, most companies aren't willing or able to build their own private cloud infrastructure due to a lack of skilled resources and the perceived risk involved. So they turn to vendors. Unfortunately, until recently, most vendor offerings provided some but not all the pieces and capabilities required to deploy a fully functional private cloud platform.

For example, basic [open source software stacks](#) deliver a private cloud framework that generally includes virtualization, compute, storage and networking components, along with security (identity management and so on), management and orchestration functionality. These layers are loosely integrated at best, however, which means the heavy lifting of integrating and testing components to make them work together is left to the customer (or third-party consultant). Similarly, most vendor-specific products have taken a mix-and-match approach, enabling customers to choose from among different modules or capabilities—necessitating integration on the back end.

Consequently, enterprises that want to avoid the large investment of time and money required to build or integrate private cloud stacks are now looking to adopt preintegrated products based on infrastructure platforms designed to support cloud-enabled apps and data. And, as our recent research reveals, these organizations [prefer](#)

Home

Castagna: If you love your data, set it free

Toigo: Time to take a holistic view of storage

Hot data storage technologies for 2017

Snapshot 1: Most enterprises confident in (and test) their DR plan

Step-by-step guide to better cloud DR planning

Snapshot 2: Mission-critical apps foundation of DR planning

Buffington: How to align data protection and compliance

Kato/Byrne: Hyper-convergence meets private cloud needs

About us

Home

Castagna: If you love your data, set it free

Toigo: Time to take a holistic view of storage

Hot data storage technologies for 2017

Snapshot 1: Most enterprises confident in (and test) their DR plan

Step-by-step guide to better cloud DR planning

Snapshot 2: Mission-critical apps foundation of DR planning

Buffington: How to align data protection and compliance

Kato/Byrne: Hyper-convergence meets private cloud needs

About us

[converged and hyper-converged infrastructures](#) (HCIs) to traditional three-tier architectures to host their private cloud environments.

HYPER-CONVERGENCE WELL-SUITED TO PRIVATE CLOUDS

As companies transition away from the inherent complexities of traditional IT architectures, we are starting to see [hyper-converged infrastructure](#) and—to some extent—converged-infrastructure providers fully integrating private cloud platforms into what was previously only a foundational infrastructure for server virtualization. This development is a natural maturation along the journey toward data center simplification.

“Born-in-the-cloud” application workloads—those purpose-built for cloud delivery—are leading the way for private clouds that play nice with public clouds—private clouds of this type will be predominately based on HCI architectures. [Mission-critical legacy applications](#) will continue to lag in private cloud adoption, though, as these tend to function perfectly well as they sit nicely on top of traditional or converged infrastructures. Over time, we expect almost all workloads will transition to a private cloud platform. But for this article we keep our focus on the applications best suited for HCI architecture today.

At the center of all HCI architectures are [software-defined \(SDS\) storage services](#) that enable flexible and modular scalability, which endears this type of infrastructure to cloud service providers. This ability to add just what you need when you need it is critical to a cost-effective

approach to building a private cloud platform as well. Moreover, placing data closer to where the application resides usually enables a higher quality of service.

The battle for the data services (SDS) layer of the software stack—the stack that enables workload mobility between private and public clouds—is just beginning. Will the winners be full private cloud providers such as Microsoft’s Azure Stack or VMware’s Cloud Foundation? They recognized the need early on to invest heavily in SDS technology that’s tightly integrated into their server virtualization technology. Or will the winner be based on a version of OpenStack or a combination of the myriad of other startups and legacy SDS providers too numerous to list in this article?

WHAT CUSTOMERS WANT

The graph ([Most important IaaS private cloud attributes](#)) lists the attributes our research tells us are most important to would-be infrastructure-as-a-service private cloud customers.

Once you get past self-service capabilities at the top, the majority of the remaining attributes place a heavy burden on the SDS technology underpinning the private cloud platform. That’s why it is crucial the data services layer of the private cloud software stacks provide the following critical capabilities:

- It must assist in the ability to move workloads between [private and public clouds](#)—ideally between disparate

(Continued on page 31)

Home

Castagna: If you love your data, set it free

Toigo: Time to take a holistic view of storage

Hot data storage technologies for 2017

Snapshot 1: Most enterprises confident in (and test) their DR plan

Step-by-step guide to better cloud DR planning

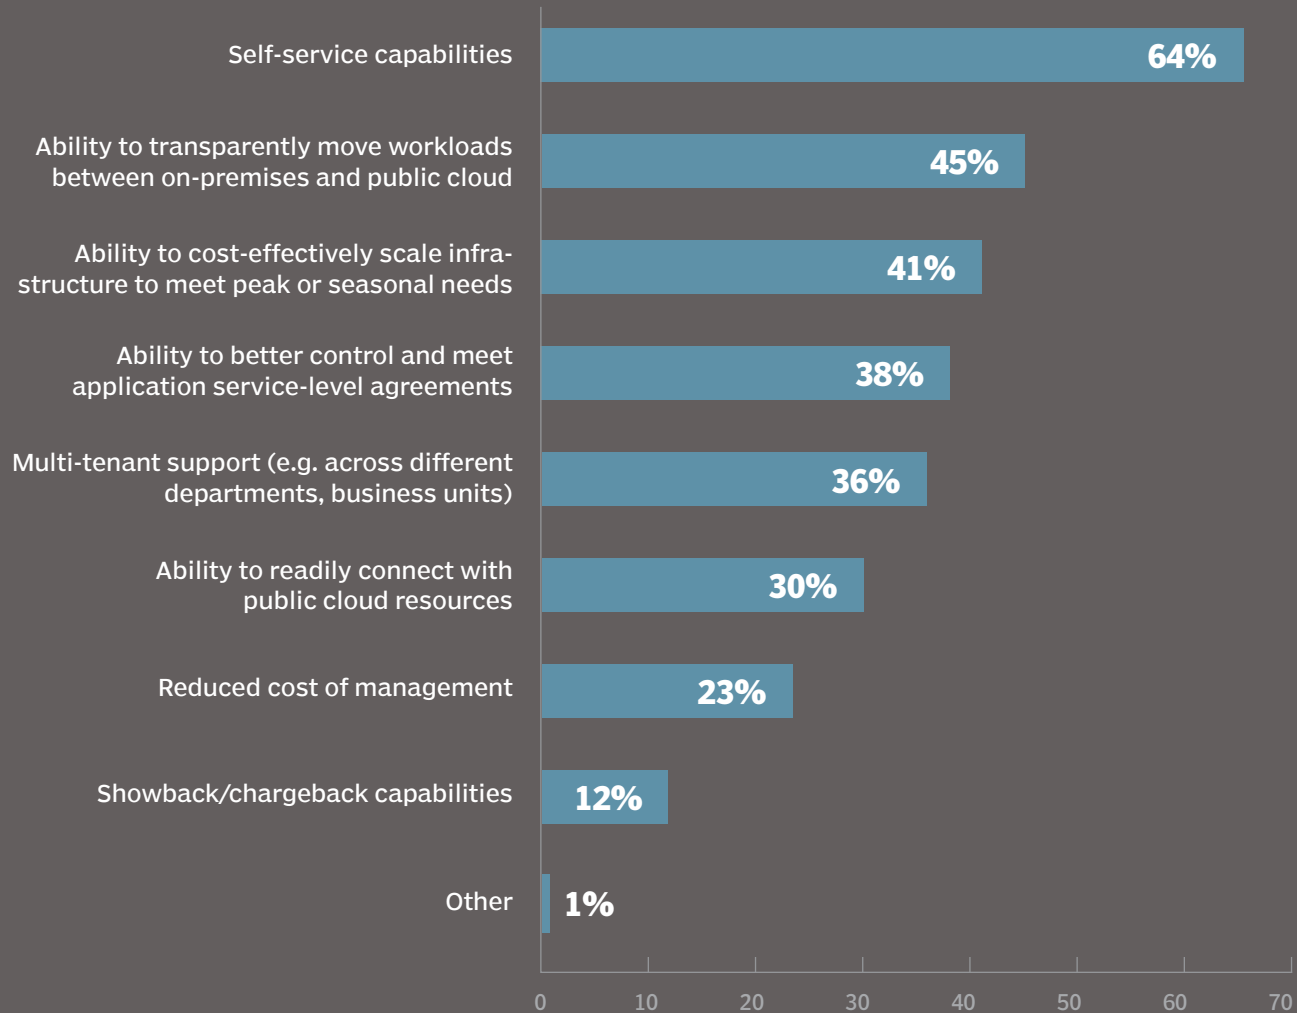
Snapshot 2: Mission-critical apps foundation of DR planning

Buffington: How to align data protection and compliance

Kato/Byrne: Hyper-convergence meets private cloud needs

About us

Most important IaaS private cloud attributes



SOURCE: TANEJA GROUP

Home

Castagna: If you love your data, set it free

Toigo: Time to take a holistic view of storage

Hot data storage technologies for 2017

Snapshot 1: Most enterprises confident in (and test) their DR plan

Step-by-step guide to better cloud DR planning

Snapshot 2: Mission-critical apps foundation of DR planning

Buffington: How to align data protection and compliance

Kato/Byrne: Hyper-convergence meets private cloud needs

About us

(Continued from page 29)

public and private cloud providers such as between Amazon Web Services (AWS), Azure and VMware.

- It must provide a high degree of quality of service across the broadest range of workloads, by enabling the transition of a broader set of workloads into a private cloud environment.
- And it should reduce management complexity through seamless integration within private and public cloud orchestration frameworks.

We already see that some vendors recognize this [need for cross-platform support](#).

Startup Hedvig is building its Universal Data Plane for cross-platform support, including containers, while Nutanix is investing in its App Mobility Fabric, for example.

Hewlett Packard Enterprise, in the meantime, supports all the leading private and public cloud vendors (AWS, Azure, OpenStack and VMware) with HPE Operation Bridge, and has a long-term strategy of data mobility using its Composable Data Fabric based on StoreVirtual VSA. Will AWS, Microsoft and VMware recognize the need for cross-platform support, or will they leave it to others like the vendors listed here to fulfill that promise?

Our customers tell us [vendor lock-in is the number one reason](#) they're concerned about buying technology that goes into building a private cloud platform from a single company. While it will be years before the final winners shake out, we believe those that seamlessly allow workloads to move across platforms—not including their own—will be the ones to come out on top. ■

JEFF KATO and **JEFF BYRNE** are senior analysts and consultants at Taneja Group.

Home

Castagna: If you love your data, set it free

Toigo: Time to take a holistic view of storage

Hot data storage technologies for 2017

Snapshot 1: Most enterprises confident in (and test) their DR plan

Step-by-step guide to better cloud DR planning

Snapshot 2: Mission-critical apps foundation of DR planning

Buffington: How to align data protection and compliance

Kato/Byrne: Hyper-convergence meets private cloud needs

About us



TechTarget Storage Media Group

STORAGE MAGAZINE

VP EDITORIAL **Rich Castagna**
EXECUTIVE EDITOR **James Alan Miller**
SENIOR MANAGING EDITOR **Ed Hannan**
CONTRIBUTING EDITORS **James Damoulakis, Steve Duplessie, Jacob Gsoedl**
DIRECTOR OF ONLINE DESIGN **Linda Koury**
ASSOCIATE MANAGING EDITOR, E-PRODUCTS **Nick Arena**

SEARCHSTORAGE.COM
SEARCHCLOUDSTORAGE.COM
SEARCHCONVERGEDINFRASTRUCTURE.COM

SENIOR NEWS DIRECTOR **Dave Raffo**
SENIOR NEWS WRITER **Sonia R. Lelii**
SENIOR WRITER **Carol Sliwa**
STAFF WRITER **Garry Kranz**
SENIOR SITE EDITOR **Rodney Brown**
SENIOR SITE EDITOR **Maggie Jones**
ASSISTANT SITE EDITOR **Erin Sullivan**

SEARCHDATABACKUP.COM
SEARCHDISASTERRECOVERY.COM
SEARCHSMBSTORAGE.COM
SEARCHSOLIDSTATESTORAGE.COM

EXECUTIVE EDITOR **James Alan 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: BLABLO101/ISTOCK



Stay connected! Follow [@SearchStorageTT](https://twitter.com/SearchStorageTT) today.