Managing the information that drives the enterprise

STORAGE

WHAT TO DO WHEN YOUR NETWORK IS A STORAGE BOTTLENECK

QUALITY AWARDS FOR ENTERPRISE ARRAYS: FIRST-TIMER FUJITSU'S BIG WIN

CASTAGNA:
STORAGE VENDORS PEER INTO THEIR CRYSTAL BALLS

TOIGO:
THE TRUTH BEHIND SOFTWARE-DEFINED STORAGE

McCLURE:
CONSIDERING SYNC AND SHARE? THINK HYBRID

BOLES:
VIRTUAL SERVER STORAGE GETS SMARTER

SNAPSHOT:
CONFIDENCE IS HIGH FOR DISASTER RECOVERY SUCCESS

GET A HANDLE ON CLOUD STORAGE
Cloud storage offers an excellent complement to on-premises storage infrastructure, but you'll need some new tools to manage data stored in the ether.

TechTarget

The day you get back, challenges will be solutions

Storage Decisions CONFERENCES AND SEMINARS

TechTarget

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.
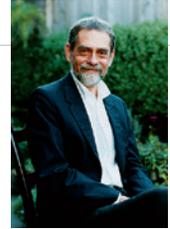
What we can learn from storage vendors' predictions

It seems as if every vendor has a mystic on staff who can see into the future. Here's what they think is in store for you.

WAS PRETTY PLEASED with myself after writing my New Year's [five predictions column](#), but before I could pat myself on the back a second time, I noticed my email inbox was full of messages with subject lines touting five, 10 or even 15 predictions for 2014. I'm not so naïve as to think I had somehow cornered the prediction market, but it looks like everybody is buffing up their crystal ball and spewing out prognostications about the future of data storage like there's no tomorrow.

Who are all these storage seers purporting to have some kind of sixth sense that lets them see into the [data center future](#)? As I sorted through the emails I realized they were from vendors, every last one of them. And all their predictions just coincidentally suggested that in the future—the very near future, they hope—everyone will

have a desperate need for their products. I didn't come across a single prediction that said, "Next year, smart storage managers will realize our system does less and costs more, and the next time you see our marketing team they'll be selling tacos from a truck."

My favorites are the ones that have survey data to back them up. They ask a question like, "Are you absolutely, indisputably, you-bet-your-job, 100% certain that your disaster recovery plan will recover every last scrap of data your company owns in, say, just a few seconds?" Of course, 75% of the survey respondents will say no (the other 25% are delusional), which will lead to headlines telling us that three-quarters of the world's data is at risk. But fear not, the surveying vendor's disaster recovery product is ready to swoop in and rescue the world.

So, here are some of my favorites from this year's

batch of vendor predictions. And to protect those who innocently prognosticate, I won't name the vendors responsible for these gems.

• **40 GigE will go mainstream.** Wow! They're really sticking their neck out with this one, huh? Well, I'm going to go them one better and predict that Ethernet will get even faster than 40 Gig and then *that* will go mainstream.

• **Organizations will seek vendors that can combine technologies like various flash drive types with automated tiering to get all-flash performance at economics equal to disk prices today.** Hmm ... I wonder if that came from a storage vendor that doesn't have an all-flash array in its portfolio?

• **Reducing the total cost of ownership will become table stakes to play in the storage market.** Now that prediction is a real relief. For all these years storage managers have just ignored the total cost of ownership of the gear they buy. But now, finally, people will care about how much money they spend and how they spend it.

• **As stored data volumes continue to see exponential growth, organizations will recognize the value of and opportunity for monetizing data.** Won't they have to? Since total cost of ownership is important now, wouldn't storage managers want to do something with all that data

they're storing? Maybe they could just call it *big data* and then nobody will ask any more questions.

• **In 2014, companies will move beyond the "I have a cloud strategy" to "I have a multi-cloud strategy."** And some of them will move even further to "I have data I can't access in multiple cloud companies that have gone belly up."

• **Online data repositories will grow to sizes that were once inconceivable.** Wasn't everything inconceivable at one time? I mean when we were chatting on princess phones with rotary dials, weren't smartphones with more computing power than NASA's man-on-the-moon mission just a wee bit inconceivable? Additional points off on this one for being way too obvious.

• **In 2014 we'll see organizations finally implement software-defined architectures to achieve continued flexibility and control.** That's easy to say because *software-defined* is so hard to define. So maybe if something works better we can say it's software-defined and if it doesn't we can just blame the hardware.

• **The false start of software-defined everything.** Wait a second, that other company said we're going to achieve flexibility and control but this one is talking about a "false start." Is this one of those glass half-full, glass half-empty situations, and isn't it annoying when people

inject reality into their predictions? The company went on to say that "the shift toward software-defined solutions represents an evolution, not a revolution," which I suppose means that someday someone will actually be able to define software-defined (but don't hold your breath).

● **A shift will begin toward an increase in adoption of high-capacity, enterprise tape drive technology.** Really? And the earth will reverse the direction it's spinning in and the sun will rise in the West. We're going to keep an eye on this one.

This last prediction is mine: As usual, storage managers will have to cut through the clutter of predictions—including mine—and any other nonsense that stands between them and a good purchasing decision. ■

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

HOME

STORAGE VENDORS
PEER INTO THEIR
CRYSTAL BALLS

THE TRUTH BEHIND
SOFTWARE-DEFINED
STORAGE

GET A HANDLE ON
CLOUD STORAGE

WHAT TO DO WHEN YOUR
NETWORK IS A STORAGE
BOTTLENECK

QUALITY AWARDS FOR
ENTERPRISE ARRAYS:
FUJITSU STEALS A WIN

CONSIDERING SYNC AND
SHARE? THINK HYBRID

VIRTUAL SERVER
STORAGE GETS SMARTER

CONFIDENCE IS HIGH
FOR DISASTER
RECOVERY SUCCESS

STORAGE REVOLUTION | JON TOIGO

Storage as Software

Maybe there really is something behind all this "software-defined storage" talk—but maybe it doesn't mean what I think it means.

FEW COLUMNS HERE elicited as much pushback from industry insiders as my [tongue-in-cheek rant about software-defined storage](). In case you missed it, I argued, first, that storage infrastructure should always be defined by the application software that will use it and, second, that *software-defined* was just the latest meaningless market speak from an industry that changes up the rhetoric every six months in an attempt to sound fresh. While many techs agreed with me, marketing folks took umbrage. I dismissed their whining … er, criticisms … until I had occasion to chat with a chief technologist at a storage virtualization software company who gave me a somewhat different view. It changed my thinking, so I figured it would be good to share it here.

There's a theory gaining traction in some circles that the current [flirtations with software-defined infrastructure](), clouds and server virtualization aren't just knee-jerk reactions to economic pressures to cut IT costs, but a trend that has been in the works for a lot longer. Over the last few decades, we've fielded a lot of technology in a haphazard way, with little attention to its proper fit with business application requirements.

Case in point: IBM says we have way too much Tier-1 storage because everyone who's fielding a new app wants it to shine, so they host its data on low-capacity, high-performance storage whether that's necessary or not.

Moreover, we've allowed vendors to sell us what they *want* to sell us rather than what we *need*. This has, in turn, made [management of heterogeneous infrastructures]() a nearly impossible task. Without management, there's only inefficiency and the need for rapid expansion of capacity

until we're oversubscribed and underutilized.

That in turn enables vendors to sell ideas like unified storage (purchasing all storage from a single vendor, aka lock-in) and to make persuasive cases for "value-add" software such as compression, deduplication and so on that contribute little strategic value but a lot of revenue for the vendor.

Eventually, the hardware itself becomes commodity—all parts made in Taiwan, Thailand or Singapore—and the value-add software becomes the mainstay of storage vendor revenues. Try reselling a NetApp filer on eBay; the software isn't transferable with the kit.

If software is the only value-creating element of the kit, then separating storage software from hardware makes sense. Storage is now a software function or should be; nobody makes money from selling hardware anymore.

If you're still tracking, this thinking leads inevitably to the conclusion that software-defined storage has a different meaning than what I critiqued in past columns. Storage is a set of services that need to be intelligently allocated to application data based on business rules and application accessibility requirements. Capacity is a service, as is performance, as is protection, as is retention. Software is increasingly used to carve up storage to create pools offering different combinations of service elements that are appropriate for different kinds of data. To the extent that the carving up is inhibited by hardware, we still operate in an old world of hardware vendor-defined storage. Conversely, to the extent that we've gone beyond array bezel logos and abstracted services away from proprietary kit, we're closer to the realm of software-defined storage.

This is important to understand if you want to build a dynamic data center going forward, one that can turn on a dime to provide the right kinds of services to business

> ## The *Starbuckification* idea puts non-technical users in charge of allocating their own services, which in my experience, is not a good idea.

processes in a fast-paced 24/7 world. At least, that's the story my friend the chief technologist is telling.

This isn't to say that the storage application has become simpler, or that it's analogous to an automatic coffee machine that enables the user to push a button for the kind of storage desired; another button for a large, medium or small serving; one for the proper amount of performance; and yet another button to sweeten with the right amount of data protection.

While that kind of *Starbuckification* of storage is possible today with storage virtualization software products like DataCore's SANsymphony-V, unlike the automatic

coffee machine, its use is hardly drool-proof. You need to actually know something about technology to [pro-vision the right kind of storage](#), and you need to know something about storage to allocate it intelligently. The *Starbuckification* idea, espoused by many cloud service providers, puts non-technical users in charge of allocating their own services, which in my experience, is not a good idea.

Software is and remains a tool. Our storage needs to be managed and allocated by intelligent humans, with software-based controls and policies serving as a more efficient extension of our ability to translate business needs into automation support. That's another way of saying that things are as they always have been.

Maybe what we really need are smarter humans to better use the much-improved storage application. ■

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

MANAGING CLOUD STORAGE

Whether your company is using public, hybrid or a private storage cloud, you need tools to manage, monitor and keep track of the stored data.

THOSE OF US in IT must have a thing about stovepipes. How else can you explain our propensity to create them? That might be a bit tongue-in-cheek, but we do spend an inordinate amount of time deploying technology and then figuring out how to integrate it with everything else we have in the data center. Cloud computing is the most recent example. We deploy applications in third-party data centers to gain the benefits of rapid deployment and lower unit cost only to learn that its management is opaque and full of hidden "gotchas."

When it comes to storage, Storage as a Service, Infrastructure as a Service and Platform as a Service providers may offer some level of insight regarding such things as capacity usage (especially to bill for it), uptime service-level agreements (SLAs) and often not much else.

Even with more advanced reporting capabilities, the provider's environment is still an entirely separate entity from the rest of the organization's IT estate. Any effort

By Phil Goodwin

at cost optimization, deployment efficiency or visibility must be undertaken on an almost entirely manual basis. Having to manage multiple application deployment models further inhibits a key IT goal: agility.

While the deployment-stovepipe-integration progression seems like a vicious cycle, it's the order of technology maturity. Fortunately, some of this maturity is seeping into the cloud environment. Most IT organizations have at least dabbled in cloud computing, making the hybrid cloud model the dominant one demanding attention.

A NEW STORAGE MANAGEMENT MODEL

There will always be a need for low-level storage management that focuses on physical attributes, but correlating private (in-house) storage deployments with cloud deployments requires a higher-level view of the different estates. In public cloud environments, IT managers have little say over specific configurations, other than the providers' SLAs. What this requires, then, is the ability to control deployments through business-level integration points such as authentication, change management and audit/compliance management. Specific requirements will vary depending on whether the deployment is primarily private cloud, hybrid cloud or backup/recovery-specific. The solution will also depend on whether the deployment is conducted in-house by the IT department or provided by the cloud provider.

BMC Software Inc. has addressed the hybrid cloud

use case with two different, but complementary, offerings: BMC Cloud Lifecycle Management (CLM) and BMC Cloud Operations Management. BMC CLM is designed to perform application workload management, rather than just systems and storage management. The key to CLM is the policy engine that drives decision support. Policy engine parameters include such things as performance requirements, security needs, capacity, physical location and lifecycle stage (i.e., development vs. production). The result should be suggestions regarding an "informed choice" about the appropriate platform for optimal deployment.

To make CLM function in a hybrid environment, BMC has implemented API connections to third-party clouds. Supported environments include Amazon Web Services, CenturyLink Technology Solutions (formerly Savvis) and Microsoft Azure. With these APIs, organizations can have a consolidated view of their entire cloud computing estate. It will help to understand how data moves from point to point and how the workload is performing.

BMC Cloud Operations Management provides a lower-level view of the cloud environment. This view includes storage, as well as the other infrastructure and application stacks. It can assist in root-cause analysis, performance analysis, capacity planning and forecasting. Without a tool to look at the IT environment holistically, IT managers would need multiple tools and manual

correlation of events and trends. From a process perspective, BMC tools help pull cloud operations into IT Infrastructure Library (ITIL) and IT Service Management (ITSM) compliance to effectively manage change control, patch management and auditing across the estate, whether the infrastructure is privately or publicly hosted.

KEEPING TABS ON CONTENT SHARING

Content sharing is another popular use case for hybrid cloud storage environments. Organizations that leverage content-sharing providers need to ensure their provider meets enterprise-level data management requirements. Box is an example of a company targeting the enterprise content-sharing market. Box frames its offering around four aspects of access:

- **Users:** profiles, access patterns and security
- **Devices:** thin devices, bring your device and mobile devices
- **Applications:** update control and retention
- **Intelligence:** reporting

The fundamentals of file sharing are inherent to the Box environment. This includes integration with Active Directory, plus single sign-on and two-factor authentication. Security is enhanced based on "behavior-based" security to detect suspicious activity. When this activity is detected, an additional verification step is required.

Activity reporting helps storage administrators to manage the environment effectively.

Box also looks at the data management ecosystem to leverage functionality from other organizations. Integration examples include data-loss prevention products, such as CipherCloud for Box, Code Green Networks' Cloud Content Control and Proofpoint Data Loss Prevention. Mobile device management integration comes from the likes of AirWatch, Good Technology, IBM's Fiberlink MaaS360 and others. This saves an IT organization from having to integrate and manage the various aspects of content collaboration.

TRADITIONAL STORAGE VENDORS OFFER CLOUD CONTROLS

While BMC and Box are examples of more broad-based cloud management, EMC Corp. is leveraging its experience in storage management to address cloud storage specifically. EMC's value proposition is to provide a "cloud-like experience," meaning reduced complexity and faster provisioning regardless of cloud type. The company's ViPR software is designed to give a single view to all cloud storage, whether it's public or private. Moreover, this management view extends to both EMC or non-EMC arrays in the storage pool. The pool is managed as a single entity for such things as provisioning and capacity planning, while the underlying storage retains its native capabilities. Third-party interfaces currently supported

include Amazon S3, OpenStack Swift and other REST-based APIs; others are planned in the future.

Like EMC, Symantec Corp. has a long heritage as a data center storage management vendor. It's working to translate those capabilities into cloud environments through Veritas Operations Manager, the user interface for Symantec Storage Foundation. This transformation is beginning in the private cloud with an emphasis on multi-tenancy, whereby business units are treated like tenants. This includes role-based access that provides greater self-service with a custom application-owner dashboard and reporting capabilities. Even so, it remains a tool focused

Cloud storage management landscape

CLOUD ENVIRONMENT TO MANAGE	BASIC SOLUTION REQUIREMENTS
Hybrid cloud	▪ High-level view of cloud estate, beyond just storage ▪ Provides service management view across disparate infrastructures ▪ May include integration with authentication, security and broader management offerings
Cloud storage	▪ Storage-specific view of storage deployment ▪ A lower-level view than hybrid cloud; gives more traditional insight to storage management issues, but across public-private clouds
Backup and recovery	▪ Purpose-specific cloud implementation to simplify backup and recovery of data ▪ Designed to be an extension of the IT group

on the needs of the storage administrator. Other features include SmartTier, SmartIO and flexible storage sharing that benefits the user, but are certainly "under the hood" from a user perspective. Symantec expects to extend these capabilities to the hybrid cloud environment in the future.

BACKUP CLOUD MANAGEMENT

Among backup and recovery cloud providers, Mozy, a unit of EMC, is a purpose-built, small enterprise and remote office/branch office backup service in the cloud. Mozy provides users with a dashboard to understand the environment. This includes tracking individual machines (i.e., identifying those that haven't been backed up in a certain time period), groups, quota management, performance management and daily reports. Access can be controlled through Active Directory or any Lightweight Directory Access Protocol (LDAP) device. Encryption is available with standard Mozy or custom key management. In late 2013, the firm released APIs to allow access to its environment for third-party reporting products. Presently, this doesn't include special plug-ins for specific products, but it can be customized for any product or to give alerts to remote network monitoring dashboards.

Asigra Inc. is another cloud backup provider, but with a different market angle from Mozy and similar services. Asigra Cloud Backup is installed by providers, such as Amazon and IBM, to give backup functionality to their cloud subscribers. It can be installed in private, public or hybrid environments. Backups can flow from private to public cloud repositories, from public to private cloud repositories, and even public to public cloud repositories. A provided dashboard identifies what data is being backed up, what's being restored and why it's being restored. Reasons for restores may include hardware failure, software malfunction or user error. Regardless, it gives IT managers the ability to remediate problem areas as they're identified.

Data remains a key asset for most organizations. As the cloud storage marketplace matures, organizations need to understand how that data is being managed and not simply trust that the cloud provider has everything under control. While most cloud providers are trustworthy, a measure of verification can avoid unpleasant surprises. Cloud storage management tools give users the control they need to feel comfortable with the solution. ∎

PHIL GOODWIN *is a storage consultant and freelance writer.*

THE STATE OF STORAGE NETWORKING

Storage networks are struggling to keep up with virtualized server environments and speedy flash storage. If your network is becoming a bottleneck, an upgrade may be in order.

THE STORAGE NETWORK is a frequently neglected component of a virtualization initiative or clustered database rollout. Greater priority is typically given to the servers and actual storage system for the project, instead of the network that connects them. That's typically true because in the past the storage network was more than capable of delivering adequate performance. But with the advent of high-density server virtualization and flash-based near-zero latency storage, the storage network is becoming a bottleneck and—like it or not—IT planners need to review their storage network upgrade options.

THE THREE LAYERS OF A DATA CENTER

A data center can be divided into three layers:

- The compute layer that runs the applications
- The storage layer that stores the data created by the compute layer

By George Crump

- The networking layer that connects the compute and storage layers

The compute layer has become significantly more powerful as virtualization has enabled a single server to support dozens of virtual servers, and the storage layer has become significantly more responsive thanks to flash-based storage. These are revolutionary changes that have had a dramatic impact on the data center.

The network layer in most data centers has been evolving as well, an evolution motivated more by obsolescence than by a need for greater performance or capabilities. These upgrades have typically happened gradually as the cost of next-generation storage network components such as adapters and switches continues to drop. Essentially, the storage network is typically on a slow, rolling upgrade as new servers are added to the data center.

THE NETWORK PERFORMANCE CHASM

As a result of these revolutionary performance changes in the compute and storage layers, the storage networking layer has become a performance chasm. The compute layer has the ability to generate tremendous and random storage I/O demand and, thanks to flash, the storage layer can respond to those demands accordingly. But the networking layer no longer has the latency of spinning hard disk storage to hide behind and, as a result, has become the bottleneck.

IT planners are suddenly faced with the need to upgrade their entire storage network at once, but the options can be overwhelming with alternatives such as Gen 5 (16 Gbps) Fibre Channel (FC), FC over Ethernet, 10 Gigabit Ethernet (10 GbE), 40 GbE, InfiniBand and server-side networking.

Let's compare the various infrastructure options that each of these network types provides.

IP NETWORKS (10 GbE AND BEYOND)

Internet Protocol (IP)-based storage networks are generating a lot of interest for storage networking because there's a belief that they're less expensive and easier to maintain. After all, most data centers already use an IP-based network for user-to-server communication and server-to-server communication.

So the inclination is to use the expertise and physical assets invested in those communications to allow servers to communicate with storage. Doing so should lower the price of the storage infrastructure and its operating costs, since it wouldn't require special skills.

Most existing IP storage architectures today use a trunked or bonded set of 1 GbE connections for performance and redundancy. The move to 10 GbE is appealing because those multiple 1 GbE connections can be eliminated. But 10 GbE networks, just like the 1 GbE networks that preceded them, have some issues that network designers need to be aware of. One issue with IP-based

storage networks is that most IP architectures still use the spanning tree protocol (STP), which means that only one connection can be active at a time.

The weakness of STP is that it was designed in the days of Ethernet hubs before switches were commonly available. In general, STP ensures there's only one active path between two points in the network. The goal was to create a loop-free infrastructure.

In the modern network, there are always redundant paths, but with STP they're blocked or turned off. When an active path fails, the network has to re-converge on a new path. In a large network, re-convergence can take a few seconds. As a result, not only is potential bandwidth wasted, but there's a "hesitation" in the entire network when an inactive path needs to be activated.

Having inactive paths in 1 GbE wasn't a big concern because there wasn't much bandwidth being wasted and each link was very inexpensive. The move to 10 GbE makes this a more glaring concern since the amount of wasted bandwidth is potentially 10x greater and the cost

Speeds and feeds: Storage networking protocols

PROTOCOL	CURRENT VERSION/SPEED	NEXT VERSION/SPEED/WHEN
Fibre Channel	▪ Gen 5/16 Gbps	▪ Gen 6 ▪ 32 Gbps ▪ 2015
Ethernet	▪ 10 Gigabit Ethernet (GbE)	▪ 40 GbE ▪ 40 Gbps ▪ Available in limited products
InfiniBand	▪ QDR 40 Gbps	▪ FDR IB 56 Gbps

of each connection is significantly higher.

IP networking vendors are trying to create fabric-based IP networks that don't have these blocking issues based largely on Transparent Interconnection of Lots of Links (Trill). But each implementation is unique and interoperability between providers seems sketchy. In addition, the cost of the hardware for these implementations is significantly higher than their standard Ethernet counterparts.

IP networks have the overhead of translating between SCSI and IP, and the general overhead of TCP. While these issues can be overcome by using special network interface cards that offload the IP stack and help with IP to SCSI conversion, installing these cards increases the cost and complexity of a seemingly simple IP network.

16 Gbps GEN 5 FIBRE CHANNEL

Most storage professionals are very familiar with FC storage networks. Fibre Channel networks have a reputation for being more expensive and complicated than Ethernet networks. But FC is a fabric-based storage topology, so all the links are active. That means none of the 16 Gbps bandwidth of Gen 5 FC is wasted on inactive links. All links are active and there's no reset period if a link fails.

Fibre Channel is also a lossless network, which means data always reaches the destination, and there are no retries as there are with IP-based networks. Finally, no

translation needs to occur, as FC transparently handles SCSI communication packets.

Fibre Channel may represent the "old guard" in storage networking, but it's the Gen 5 architecture that delivers almost all its theoretical potential, and it has specific capabilities to prioritize I/O bandwidth at the granularity of a virtual machine (VM).

SERVER-SIDE NETWORKING

Server-side networks often leverage software-defined storage to aggregate internal storage across the server hosts that make up the compute layer of a virtual infrastructure. That aggregated storage is then shared as a single pool. These forms of storage networking actually collapse the data center layers described above, but don't eliminate them. They all still reside on the same physical devices, the hosts that make up the virtual infrastructure.

Server-side networking relies, for the most part, on IP-based protocols, although some use InfiniBand. They should help with some of the IP issues described earlier because the overhead of dealing with IP communication is addressed by the CPU power within the physical hosts. Also, STP networking issues should be minimized since the connections between servers are private, requiring fewer connections.

Some of these solutions will segment data across the aggregated pool of storage. Doing so provides the data

protection and data sharing needed in clustered environments. But it also introduces network latency, since there's still a network connecting the servers.

A few of these offerings mitigate this problem by allowing a virtual machine's data to be stored on the host the VM resides on and then distributed across the pool. All reads come from the local server, eliminating latency. Writes go to both the local pool and the aggregated pool, enabling data safety and sharing. Essentially, these

systems give up some capacity efficiency for additional performance.

EXTREME NETWORKING: INFINIBAND, 40 GbE AND GEN 6 FC

There are several more extreme or esoteric storage networking choices available today, including InfiniBand, 40 GbE and Gen 6 Fibre Channel. The latter two are essentially upgrades to the current IP and FC standards,

Pros and cons of storage networking alternatives

PROTOCOL	PROs	CONs
Fibre Channel	▪ Still most widely used ▪ Purest use of bandwidth ▪ Standardized way to prioritize traffic	▪ Potentially most expensive ▪ Requires a specific expertise in Fibre Channel
Ethernet (IP iSCSI/NAS)	▪ Compatible with current messaging network ▪ Potential to reuse expertise	▪ Overhead reduces effective bandwidth ▪ Becomes complex as it scales
InfiniBand	▪ Fabric-based architecture ▪ Fast, 40 Gbps with limited overhead	▪ Very limited deployments ▪ Limited tools

providing greater bandwidth. InfiniBand, on the other hand, is a flat network similar to FC but it runs at 40 Gbps.

> **IT planners should consider which of the available architectures will allow them to extract maximum ROI from investments in the compute and storage layers.**

InfiniBand is commonly found within storage systems as an interconnect of clustered nodes in environments that demand extremely high-performance I/O, such as high-frequency trading. It's also a popular choice as a server interconnect for mirroring flash devices between servers in high-availability configurations.

While InfiniBand can be made to work in many storage network configurations, especially server-side networks, the chances of it becoming widely adopted seem slim. It will remain a viable special use-case network.

NET NETWORK

There's no question that many data centers have reached the tipping point where without at least a gradual upgrade to the next-generation storage network architecture, the potential return on investment (ROI) in the compute and storage layers will be limited. Without the raw bandwidth and other capabilities these networks provide, virtualization reaches maximum density and flash storage can't deliver maximum performance.

The question is which network should be chosen for the next-generation storage infrastructure. For the most part, that decision will depend on what's in place today. Most organizations will stick with the topology and protocol they already own, and simply move to its next generation. But IT planners should consider which of the available architectures will allow them to extract maximum ROI from investments in the compute and storage layers. There should be little hesitancy about moving to a new architecture if it can justify itself with greater density and greater performance. ∎

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

FUJITSU TAKES THE QUALITY AWARD FOR ENTERPRISE ARRAYS

This time—our ninth annual Quality Awards survey for enterprise arrays—Fujitsu didn't just make the cut, it took the cake by earning top honors.

By Rich Castagna

THE TERM *enterprise array* may be broadly defined these days, but common wisdom says these data center main-stays share at least three prominent traits: a lot of I/O horsepower, ample capacity and a price tag that usually has more digits than a storage manager cares to count.

But even measured against those high expectations, enterprise array vendors don't disappoint their customers, as underscored by our ninth annual Quality Awards user satisfaction survey for enterprise arrays. All the usual suspects were represented and fared quite well, but this year a surprise winner emerged: Fujitsu, featuring its high-end Eternus systems.

Following Fujitsu are a bevy of more familiar players, including EMC Corp., IBM, Hitachi Data Systems, Hewlett-Packard (HP) Co. and NetApp Inc., in yet another typically tight race for rankings. Over the past few years we've seen some fairly high scores for these big iron products, and this year's survey yields similar results.

Overall rankings

Fujitsu's margin of victory over runner-up EMC (0.06) was the smallest we've seen since NetApp nipped EMC by the same margin in the fifth edition of this survey. (EMC and Net-App tied for first on the fourth survey.) Fujitsu's overall score of 6.57 was accrued by snag-ging the top scores in three of the five rating categories, with one of those wins earned by a mere 0.01 point. IBM and Hitachi prevailed in the remaining two categories, while EMC nipped at all the winners' heels with a consistently competitive set of scores and an overall 6.51 that was just off Fujitsu's mark.

With only 0.07 points separating them, IBM (6.48), Hitachi (6.45) and HP (6.41) were bunched tightly in spots three through five. NetApp, a winner/co-winner of four previous enterprise array survey overall crowns, rounded out the field with a 6.14.

As another indicator of the wall-to-wall quality these products represent, they all earned a score of 6.00 or higher across all rating categories. It's just the second time that has happened with enterprise storage arrays, and a good gauge of customer satisfaction.

For this survey, we had 348 responses that included a total of 602 product evaluations.

ENTERPRISE ARRAYS: OVERALL RANKINGS
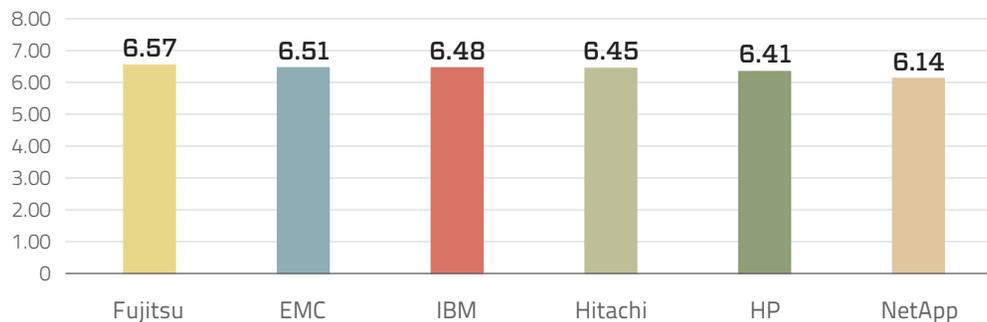
BY THE NUMBERS

6.71
Fujitsu's highest category score (product features), but it placed second to Hitachi's 6.77 (product reliability).

3, 3, 3
Third highest ever group average scores for sales-force competence, initial product quality and product features.

6.43
The overall average score for all products in the survey; excellent, but still just the fourth highest ever.

KEY STAT
This is the first time that Fujitsu gained enough responses and product evaluations to qualify as a finalist among enterprise arrays.
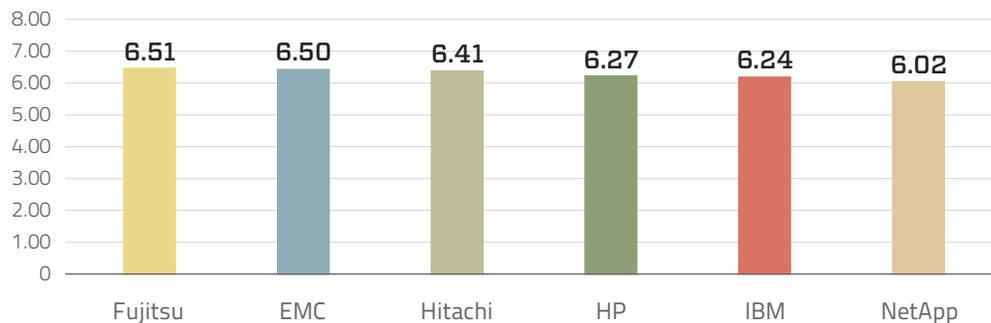
Sales-force competence

Call it the foot in the door, the first impression or whatever catchphrase best sums up the initial vendor-user experience. How well a vendor's sales force and representatives understand a potential customer's environment and business needs is likely to be one of the key make-it-or-break-it factors for whether that enterprise-class rig ever makes it into the data center.

Fujitsu edged out EMC in the sales-force competence category in a real squeaker, winning by the smallest possible margin, 6.51 to 6.50. Fujitsu prevailed on three of the six rating statements in the category, while EMC scored highest on two; the two vendors tied on a third. For all three statements that Fujitsu won, EMC placed second, and Fujitsu had the second-highest rating for one of the statements on which EMC ranked best—further demonstrating just how closely contested this category was.

Fujitsu's winning statement scores were for having flexible sales reps, for reps that keep customers' interest foremost and for reps who are easy to negotiate with. EMC's wins came for the statements, "My sales rep is knowledgeable about my industry" and "My sales rep understands my business."

ENTERPRISE ARRAYS: SALES-FORCE COMPETENCE

Bar chart titled "ENTERPRISE ARRAYS: SALES-FORCE COMPETENCE" with y-axis from 0 to 8.00. Values: Fujitsu 6.51, EMC 6.50, Hitachi 6.41, HP 6.27, IBM 6.24, NetApp 6.02.

BY THE NUMBERS

6.75

The highest statement score in the sales-force competence category, earned by EMC and Fujitsu for having knowledgeable sales support teams.

6.41

Hitachi's score was good for third place in this category, highlighted by an excellent 6.71 for a knowledgeable sales support team.

6+

For two statements— reps knowledgeable about customers' industries and knowledgeable support teams—all vendors had scores higher than 6.00.

KEY STAT

A slim margin of 0.02 separated this group's sales-force scores from the best ever mark.

Initial product quality

Users of IBM's enterprise storage ranked their out-of-the-box experience the best—but only by a small margin, with IBM edging out Fujitsu 6.57 to 6.54. But that difference was even greater than the gap between Fujitsu and third-place HP, which racked up an initial product quality category average of 6.53. IBM outscored the competition on three of the six category statements, while Fujitsu, HP and Hitachi each snagged one.

IBM's strong points were ease of use (6.66), delivering good value (6.55) and having products that require little vendor intervention (6.45). Fujitsu scored 6.71 for "This product was easy to get up and running," and Hitachi's statement leader was 6.75 for "I am satisfied with the level of professional services this product requires." But the highest single statement score was HP's 6.76 for products that install without defects.

All the products scored 6.00 or higher for all rating statements in the initial product quality category, with the third best cumulative average we've seen to date. With so much emphasis these days on the immediate productivity that pre-tested, pre-configured converged systems can offer, it seems vendors are rising to the challenge with enterprise arrays that also install and get up to speed quickly.

ENTERPRISE ARRAYS: INITIAL PRODUCT QUALITY
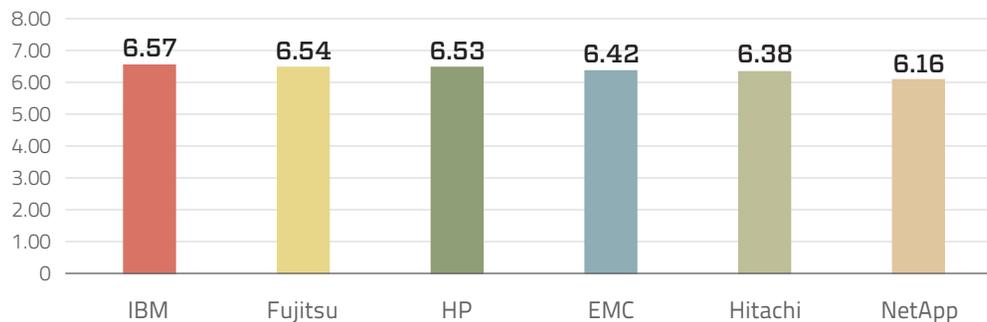
BY THE NUMBERS

6.60
The highest statement average for all products for "This product was installed without any defects."

6.57
IBM's winning category score is the third highest ever for initial product quality, tying Hitachi's mark from last year.

6.39
The group's average for the bellwether statement, "This product delivers good value for the money," which has been exceeded only twice before.

KEY STAT
Model of consistency: The difference between IBM's best and lowest scores for initial product quality was 0.21.
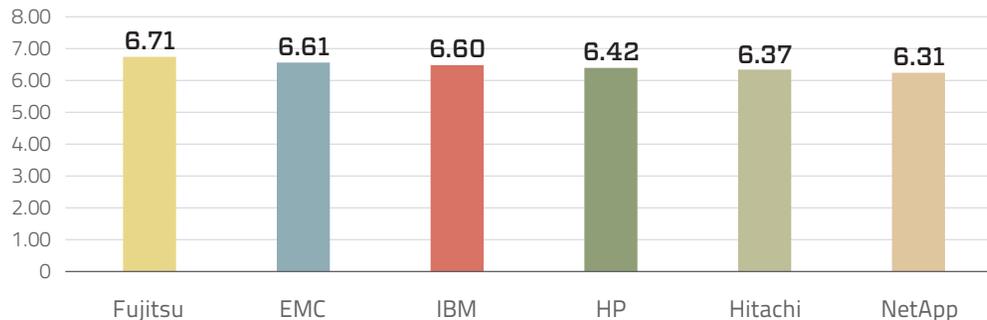
Product features

When a product meets quality and reliability expectations, even a relatively modest product feature list will be acceptable. But we know that high-end storage systems typically don't lack for features, although our survey focuses on core capabilities. As a group, these products certainly seem well-equipped to meet enterprise requirements, as respondents gave them their second highest category rating (6.50). The leader of the pack was Fujitsu once again, collecting a solid 6.71 to outdistance EMC and IBM who nearly tied at 6.61 and 6.60, respectively.

Fujitsu scored highest on four statements, doing so with the four highest scores for all products in the group, highlighted by a 6.91 for interoperability, a pair of 6.83s for replication and capacity scaling, and a 6.79 for management features that meet users' needs. EMC was the winner of mirroring features (6.78), while IBM captured two statements: the comprehensive "Overall, this product's features meet my needs" statement (6.78) and snapshot features (6.53).

HP, Hitachi and NetApp also did well in the product features category; HP scored best on the overall features statement, Hitachi's top mark was for capacity scaling and NetApp logged a 6.52 for snapshots.

ENTERPRISE ARRAYS: PRODUCT FEATURES

|	Fujitsu	EMC	IBM	HP	Hitachi	NetApp
Score	6.71	6.61	6.60	6.42	6.37	6.31

Product reliability

A satisfying sales process, great out-of-the-box experience and cool features won't add up to much if the array doesn't stay up and running. The five rating statements in the product reliability category measure how well products meet those expectations. There are apparently very few sleepless nights for Hitachi users, who rated their products so highly for reliability that Hitachi prevailed on four of those statements and came in second on the fifth. EMC took the fifth statement honors (a 6.85 for meeting service levels) and was second to only Hitachi on three of the others (IBM was second for providing comprehensive upgrade guidance).

Hitachi flirted with a 7.00 score but settled for an impressive 6.98 for "This product experiences very little downtime." Hitachi's other leading grades were for upgrade guidance (6.76), non-disruptive patching (6.69) and requiring few unplanned patches (6.60).

IBM (third place) and Fujitsu (fourth) also fared well in the product reliability category. IBM achieved its best mark for meeting service-level requirements (6.74) and Fujitsu's high grade (6.71) was picked up for the same statement.

ENTERPRISE ARRAYS: PRODUCT RELIABILITY

Vendor	Score
Hitachi	6.77
EMC	6.70
IBM	6.59
Fujitsu	6.51
HP	6.44
NetApp	6.12
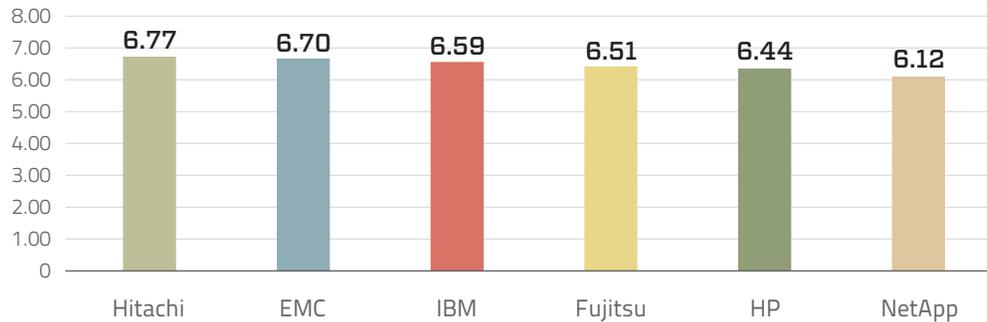
BY THE NUMBERS

6.65

The group's two highest averages for all products were a couple of 6.65s for meeting service-level requirements and products that have very little downtime.

6.52

The group's overall average for reliability was its highest in the survey—that has happened in eight of the nine enterprise array surveys fielded to date.

KEY STAT

This is the sixth time Hitachi has scored highest for product reliability on the Quality Awards for enterprise arrays.
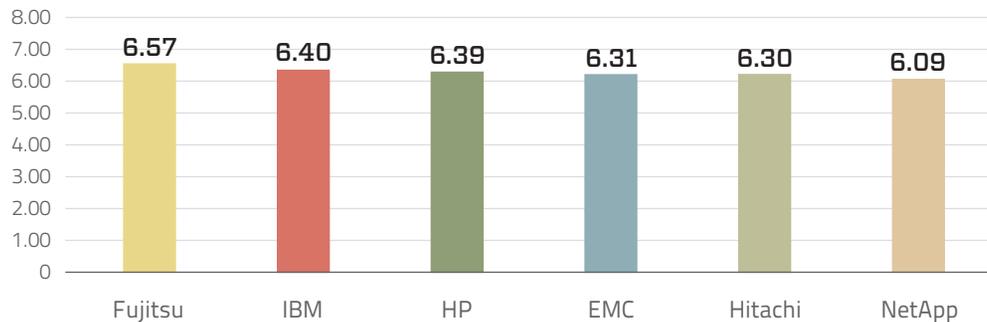
Technical support

We've said it many times before and it's still true: Good technical support can help a user overlook a multitude of shortcomings in a storage product. No matter how dependable the product, timely tech support is always welcome.

Fujitsu was back on top with the tech support category, scoring a 6.57 to give it a fair lead over a tightly bunched pack led by IBM and with HP, EMC and Hitachi all doing some heel nipping. Fujitsu scored highest on six of the eight statements in this category, with its best marks for knowledgeable support personnel (6.79), support issues that rarely require escalation (6.71) and having knowledgeable third-party partners (6.61).

IBM topped the group for documentation and support materials with a 6.52, while EMC had the other winning statement score for supplying support as contractually specified (6.66). Third-place HP's best showing was a 6.58 for knowledgeable support personnel, and fifth-place Hitachi's best rating came for "Vendor supplies support as contractually specified" (6.58).

ENTERPRISE ARRAYS: TECHNICAL SUPPORT

QUALITY AWARDS

Would you buy this product again?

Go figure; Fujitsu is the overall winner for user ratings in five categories, but trails most of the rest of the field on the "add-on" question, "All things considered, I would buy a [name your product] again." Of course, in this case, the difference from top to bottom is a mere 8 percentage points, and any vendor would probably be delighted with an 80-something per-cent satisfaction rate. Still, it's hard to explain the "buy again" numbers on many of these surveys, with familiarity seeming to be a leading factor. With that said, our hats are off to EMC, IBM and HP for finishing one-two-three on this question. ■

ENTERPRISE ARRAYS: WOULD YOU BUY THIS PRODUCT AGAIN?

89%	87%	85%	84%	81%	81%
EMC	IBM	HP	Hitachi	NetApp	Fujitsu

[Rich Castagna](#) is editorial director of TechTarget's Storage Media Group.
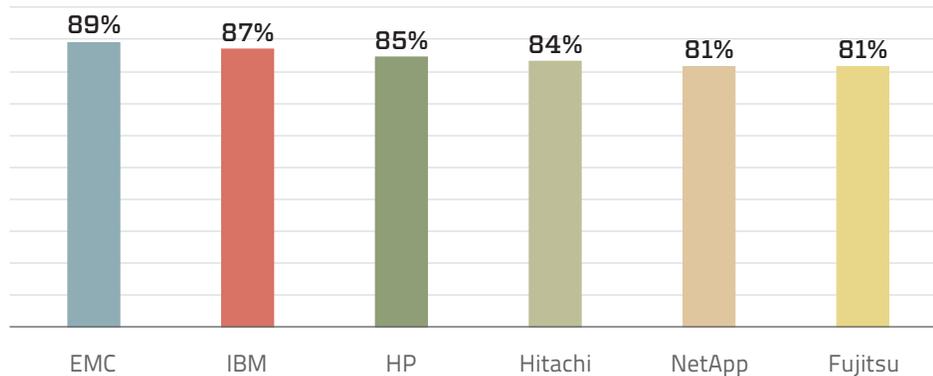
ABOUT THE AWARDS

The *Storage* magazine/Search-Storage.com Quality Awards are designed to identify and recog-nize products that have proven their quality and reliability in actual use. The results are derived from a survey of qualified readers who assess products in five main cat-egories: sales-force competence, initial product quality, product fea-tures, product reliability and tech-nical support. Our methodology incorporates statistically valid poll-ing that eliminates market share as a factor. Our objective is to iden-tify the most reliable products on the market regardless of vendor name, reputation or size. Products were rated on a scale of 1.00 to 8.00, where 8.00 is the best score. A total of 348 respondents provided 602 product evaluations.

Products in the survey: The follow-ing products were included in the ninth Quality Awards for enterprise storage arrays survey; the number of responses for finalists is shown in parentheses.

EMC Corp. Symmetrix VMAX/ VMAXe or DMX-3/DMX-4 (147)

Fujitsu Eternus DX8400/DX8700 or DX200 S3/DX500 S3/DX600 S3 (27)

Hewlett-Packard (HP) Co. XP Series, HP StorageWorks P9000 Series, HP 3PAR StoreServe 7000/ HP 3PAR P10000 Storage Sys-tems (149)

Hitachi Data Systems USP/USP V/VSP Series (63)

IBM DS8000 Series or XIV Storage System (100)

NEC Corp. D8 Series*

NetApp Inc. FAS6000 Series or V6000 Series (108)

SGI 15000/16000/17000 Series*

Too few responses to qualify as a finalist

Hybrid model offers more secure file sync and share

The first popular collaboration and file-sharing services were strictly cloud-based affairs, but companies may be better off with hybrid or on-site implementations.

ARE COMPANIES REALLY comfortable putting their data in the cloud? It's a simple question, but the answer is quite complex. Obviously, some companies are [storing corporate data in the cloud](). Back in June 2012, Amazon said it had 1 trillion objects stored on its S3 storage cloud, and less than a year later it hit 2 trillion. Essentially, it took S3 six years to get to 1 trillion, and less than a year to double that. That's a lot of data in the cloud. Some of it is consumer data, but some is certainly business data. So just what data is going to the cloud?

SECURITY IS THE BIGGEST CONCERN

Of course, not all corporate data is equal. Some data, like patient healthcare information, is regulated and must have multiple copies kept geographically apart or audit logs showing who accessed it. Some data is sensitive, and may be considered a core competitive advantage, and it would be catastrophic to the business if it was leaked. Those types of data should be treated differently than the latest marketing brochure or product specs which, if leaked or obtained by a competitor, wouldn't cause any harm. Depending on the type and sensitivity of the data you're storing, there are different security concerns, and [security is the bottom line]() when it comes to putting corporate data in the cloud.

Back in 2012—ages ago in cloud years—there were big concerns about data security, specifically fears that hackers or even cloud provider employees could access and steal data. In 2013, those concerns persisted, but they were compounded by the [NSA Prism scandal](). So now

with even more security threats on more fronts, are companies still willing to store corporate data in the cloud?

COMPANIES ARE SELECTIVE ABOUT CLOUD STORAGE

In a late 2013 survey of 334 North American IT professionals in companies ranging from small (as few as 20 employees) to very large (more than 20,000 employees),

> **Many of these services now offer IT insight into who's doing the accessing, sharing and downloading; the ability to limit with whom files are shared; and the ability to expire shared links and files.**

Enterprise Strategy Group (ESG) found that many are willing to store their files in the cloud. But most IT pros want a solution that offers the flexibility to choose where specific data sets are stored: on the premises or in a public cloud. This survey looked specifically at corporate deployment preferences for online file sharing and collaboration (e.g., Dropbox-type platforms) and asked IT pros if they preferred data stored on-premises or in the cloud.

Digging a little deeper, we found that enterprise

adoption of online file-sharing (OFS) solutions continues to climb. The initial OFS charge was led by knowledge workers leveraging the consumer versions of these services for work, but now corporate IT groups are starting to embrace these offerings. In addition to enabling users to synchronize, share and access documents from a variety of devices, many of these services now offer IT insight into who's doing the accessing, sharing and downloading; the ability to limit with whom files are shared; and the ability to expire shared links and files. These services give IT real control over corporate file data while still allowing it to be accessed using a variety of endpoint devices.

HYBRID FILE SYNC AND SHARE IS APPEALING

The market for OFS solutions has evolved significantly over the past 24 months. A couple of years ago the only choice was a SaaS-based offering where the application and data live in the service provider's data center. So naturally, current OFS customers overwhelmingly use the public cloud deployment model. But our research uncovered an interesting trend. Among the early cloud-based OFS adopters, there's actually significant demand for hybrid deployments. In fact, a large majority of these respondents are extremely (69%) or somewhat (28%) interested in a hybrid model that would allow them to retain all or at least some data on-premises. Similarly, potential OFS adopters were more likely to leverage some type of hybrid cloud offering than a public cloud or completely

on-premises solution for future corporate deployments.

The preference to store some or all data on the premises is primarily driven by the desire to maintain control over corporate data and to leverage existing storage investments. In fact, more than half of both current public cloud OFS users and potential hybrid or on-premises OFS adopters indicated they wanted flexibility and control over where data is stored. Additionally, many organizations, particularly large enterprises, have made significant investments building out their storage infrastructures and would benefit from leveraging those existing resources.

SECURITY AND COMPLIANCE STILL TOP OF MIND

Security and regulatory compliance are also key concerns among current and potential OFS users and are also reasons for considering services with on-premises storage capabilities. To gain more insight into objections to the purely public cloud-based OFS model, ESG looked at types of data (if any) that current and potential OFS users prohibit from being stored with cloud storage services. Responses ranged from regulated data to intellectual property, but the biggest takeaway is that at least 90% of respondent organizations impose restrictions on at least

one type of potentially sensitive information. This suggests that the vast majority of organizations have some reservations about how service providers handle sensitive company data. In light of those reservations, it's not surprising that customer interest in hybrid and on-premises deployment models has burgeoned.

So, are companies willing to store corporate information in the cloud? The answer is still yes—but not all their company data. There are still concerns about security and control, and an interest in leveraging existing infrastructure investments. And more than one-third of respondents have concerns about third parties being able to access their cloud-based data. In the long-term, cloud-based offerings will likely be a dominant deployment model, at least for OFS. But near-term, enterprise IT organizations want the comfort of having the ability to choose where data lives, including keeping some in-house. The good news is that many OFS vendors have heard the message loud and clear. Vendors such as Accellion, Citrix (ShareFile), EMC (Syncplicity) and Egnyte—among others—all offer this type of flexibility. ∎

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

Storage for virtual servers getting smarter

New storage products that are designed from the ground up to serve virtual servers can save money and time.

O**VER THE PAST** few months, Taneja Group Labs has done a considerable amount of work with products I'll call next-generation virtual infrastructure storage, including Hewlett-Packard's StoreVirtual VSA products, Tintri's VMstore and VMware's Virtual SAN (VSAN). Some of them are described as software-defined storage, while others are more akin to virtualization-specific storage, but they all represent an emerging class of storage that appears poised to fundamentally alter virtual infrastructure administration.

Each product has proven intriguingly useful in making virtualization storage a bit more agile and much more automated, even though each approaches the issue from a different angle. StoreVirtual VSA is deeply injected into the virtual infrastructure and has a uniquely versatile approach to scalability: basically, you can run it anywhere and keep adding more to scale or adapt. In a recent report, we recognized that Tintri's virtual-infrastructure-dedicated VMstore is so deeply integrated with virtual machine (VM) and VM-specific tasks that it makes the act of managing virtual storage almost disappear—to the degree that we called it a 60x reduction in the time administrators spend on storage management tasks in the virtual infrastructure. Meanwhile, VMware's VSAN is wrapped in a logical organizational and configuration layer that highly automates ongoing management, what VMware calls Storage Policy-Based Management (SPBM).

CUTTING STORAGE COSTS IN VIRTUAL ENVIRONMENTS

These are just a few of the products coming to market, but

the makings of a sea change are clearly evident in these new approaches to virtualization storage, which merits serious evaluation. Why? Because virtualization-specific storage may have a bigger impact on storage costs than anything we've seen, and storage costs are a really big deal in the data center scheme of things.

Based on our testing and conversations with vendors, Taneja Group uses a rule of thumb that midrange storage still costs the typical organization $3 to $5 per gigabyte (GB) to purchase, and then $5 to $8 per GB *annually* to operate. Over the life of a typical storage array, operating that storage will cost many times more than the original purchase price, with the vast majority of that cost the result of the time and effort spent managing the system. Since there are still non-virtual systems in almost every data center, virtualization-specific storage may not replace all of a data center's storage but it will replace enough to radically reduce storage costs.

NEW VIRTUAL STORAGE SYSTEMS BRING FUNDAMENTAL CHANGE

Let's examine four ways these new storage approaches can change storage management.

1. Speedy configuration and reconfiguration. Integrating storage with the virtual infrastructure makes it more agile and more easily adaptable to changing virtual workload requirements. Such storage systems can typically be deployed and then drastically altered (for performance, protection or availability) from within the virtual infrastructure, and without the need to reconfigure a separate storage layer.

2. More automation. If storage is integrated closely enough with the virtual workload, it's also often highly automated. This can practically eliminate configuration and provisioning and, in effect, "pre-organize" the storage infrastructure. This can make it easier to deploy VMs or copies of VMs with the right level of performance, availability and protection, and then monitor and confirm how storage is provisioned. This is at the heart of VMware's VSAN SPBM, but similar capabilities are springing up elsewhere. This is also why Fibre Channel storage vendors have long anticipated VMware's long-promised Virtual Volume (vVol) functionality that's beginning to look a bit like vaporware. According to VMware, vVol will change a storage system's focus from LUNs or volumes to its own VMDK file architecture.

3. Easier management. We're also seeing how this concept of automation and storage policies can be extremely powerful in the continued management of the storage infrastructure, and in ways that may not be obvious. For example, consider how an application needs to move from test/development into production today. With virtualization-specific storage, simply changing a policy could alter

the performance and availability configuration for 10, 20 or even 100 different VMs all at once; no cloning, migration or anything else required.

4. Improved insight. Ultimately, when storage becomes deeply integrated and VM-specific it usually enhances visibility, which can fundamentally alter the day-to-day operation of storage. If integrated enough, a virtualization-specific offering may be able to collect highly granular data—from VM latency to host latency to storage latency, for example—and, in turn, provide better visibility into troubleshooting and tuning application performance.

BY ANY OTHER NAME ...

Call it what you want—software-defined storage, virtualization-specific storage or VM-centric storage—but the real story is that these capabilities are obtainable today and make up the capabilities and benefits of some

seriously innovative virtualization storage products.

Still, there are some caveats with virtualization-specific storage. The biggest is that early products look poised to change responsibilities in the data center, with fewer administrators involved in managing storage. Historically, IT hasn't exhibited a great deal of good governance in managing storage. As more tasks become consolidated in the hands of fewer administrators, it's worth thinking about how we maintain responsible use and control of storage. The upside is that a reduction in operational costs might allow us to rethink the operation of storage and, for the first time, invest in much better storage operations. But getting there will take some foresight and planning. Now is the time to be thinking about how you might "do" storage better in your company if you could radically reduce the time and effort spent on its upkeep. ∎
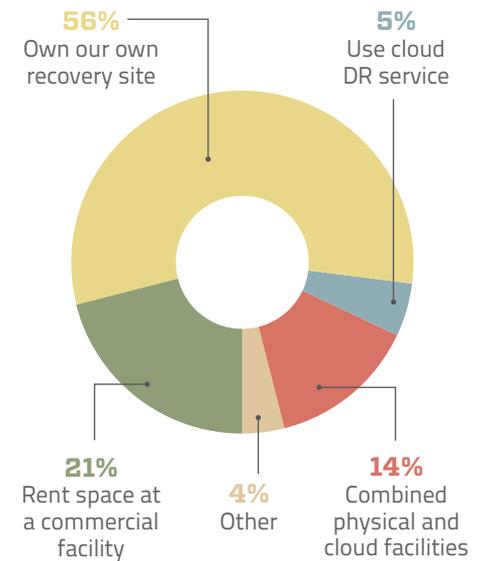
JEFF BOLES *is a senior analyst at Taneja Group.*

Most companies have high expectations for their DR plans

ONCE CONSIDERED AN insurance policy only well-heeled companies could afford, disaster recovery (DR) is now standard fare in most data centers. In our recent survey, 77% of respondent organizations said they have a DR test plan in place, just a few percentage points higher than two years ago. Most (61%) of those plans cover mission-critical applications, but 36% say their DR strategy protects all their company's apps. A DR plan's recovery point objective defines how much data can be lost without problems; 26% of those surveyed had zero tolerance for any data loss, but 30% could forfeit a day's worth. A recovery time objective is how long a recovery will take; 78% said 24 hours or less, but 14% expect lickety-split recoveries of an hour or less. Those numbers are demanding, but with server virtualization and cloud services, recoveries are faster and more comprehensive: 83% are using server virtualization at either or both primary and recovery sites, and 19% are using cloud recovery exclusively or paired with a physical site. And what about the laggards who lack a DR plan? Forty percent say they're working on it now, and 19% "haven't gotten around to it yet." —*Rich Castagna*

DOES YOUR COMPANY OWN A RECOVERY SITE, RENT RECOVERY SPACE OR USE A CLOUD SERVICE FOR DR?

56%
Own our own
recovery site

5%
Use cloud
DR service

21%
Rent space at
a commercial
facility

4%
Other

14%
Combined
physical and
cloud facilities

HOW QUICKLY DO YOU EXPECT YOUR COMPANY WILL BE ABLE TO RECOVER ITS MISSION-CRITICAL APPLICATIONS AFTER A DISASTER?

Average
expected
data recovery
time:
24.6 hours

14%	41%	23%	18%	2%	2%
Within an hour	2 hours to 8 hours	8 hours to 24 hours	1 day to 5 days	More than a week	Other

51

PERCENT OF COMPANIES THAT TEST THEIR DISASTER RECOVERY PLANS AT LEAST TWICE A YEAR

TechTarget Storage Media Group