Managing the information that drives the enterprise

STORAGE

OCTOBER 2013
VOL. 12 | NO. 8

GOODBYE, OLD BACKUP APP

Say hello to new backup methods that leverage snapshots, CDP and replication to build a better backup. **PAGE 9**

TechTarget

Who knows where flash shines brightest?

Flash is cool, crazy fast and it's going to remake our data centers; but it can be pretty confusing, too.

RECENTLY, I WAS on a panel at the Flash Memory Summit and Exhibition conference held in Silicon Valley. The topic of the session was why flash storage should be considered an investment rather than just a cost. The panel consisted of three storage vendors and me. I'm not sure if I was there as an antidote to the vendors or for comic relief. If the latter, I still took my job on the panel seriously.

The three amigos were smart guys who really knew storage and had some strong opinions about implementing solid-state storage. So, it didn't take long for them to start going at each other (in a very gentlemanly manner) about whose approach to "flashing up" storage was better. It was a minor turf skirmish based on a completely normal impulse to stand one's ground: Defend motherhood, the flag and the product portfolio—but not necessarily in that order.

Basically, the tiff centered on how much flash technology you need and where you should put it. It's a debate reminiscent of the vendor squabbles from a few years back over where storage virtualization should be implemented. Ultimately, that debate was resolved with the familiar conclusion of "it depends." Like all things in IT, it depended on the problem being solved, the current environment, and available budget and expertise. It's the same with flash today.

The whole flash-in-the-enterprise scene has developed a lot faster than most observers expected, and the technology's potential to be truly transformative can't be understated. But one of the coolest things about the flash revolution is that it's battling traditional storage on all fronts with an array (pun intended) of deployment options and technical implementations. But that's also one of the most confusing things about flash technologies.

Ignoring the end user or consumer-targeted flash stuff such as USB thumb drives, CompactFlash cards, Secure Digital cards and those tiny fingernail-sized MicroSD slivers that add gigs of storage to your phone, there are still a bevy of enterprise choices.

You can get a flash drive that looks like a typical 3.5- or 2.5-inch SATA or SAS hard disk and plugs into one of those bays. Or get your flash on a PCI Express (PCIe) card that slots into a server. You can also get a traditional array that blends some flash storage with hard disk drives. Or throw tradition out the window and opt for an all-flash storage array. But wait, there's more. There are flash appliances that operate as caches in front of hard disk arrays, and now there's even flash storage that plugs into the DIMM slots that are also used for DRAM.

But before you make the form-factor/implementation decision, you need to decide how you're going to use flash for the greatest impact on performance. As persistent storage, it's the fastest stuff around and can easily eliminate bottlenecks caused by latency and deliver mind-boggling IOPS. But it can also be used as a cache, essentially augmenting and super-sizing server memory and releasing I/O-bound apps from the constraints of limited amounts of expensive DRAM.

If you need fast flash storage, any of the previous implementations will do the trick—depending, of course, on your environment. If your flash is going to be cache, the implementation choices aren't reduced by much, as the caching can happen in the server (PCIe and SAS/SATA plug-ins), in an appliance or even at the array.

There's also a fair amount of confusion about the difference between caching and automated tiering, which solid-state has helped bring back into vogue.

And as with any newish technology, it gets even more complicated when you look at specific products and how they leverage flash. For example, some caching apps are hands off when it comes to server DRAM, while others will integrate it and effectively create a tiered cache setup.

There's also a fair amount of confusion about the difference between caching and automated tiering, which solid-state has helped bring back into vogue.

Flash technologies are ushering in a new age of storage, and some of the old paradigms just don't cut it anymore. We didn't have these discussions, debates and contentions when we lived in a hard disk-only world. I don't remember ever hearing a storage vendor extolling the caching cachet of a collection of hard disks. And while you might grapple with the question of where to put flash, that's a pretty basic discussion with hard disks: You can put it here (server) or there (array).

What's a storage manager to do? First, learn as much as you can about the technology and products (shameless plug: Try SearchSolidStateStorage.com first). But you also need to lean on your storage vendors to get accurate, useful and comparable information about their solid-state products. One problem we're having is that vendors are very selective about the metrics used to describe a product's capabilities. Too often, a storage manager gets stuck with an apples-to-oranges comparison when trying to evaluate products. Greater standardization in this area would not only help the IT professionals struggling to evaluate these products, but help the vendors selling them. ■

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

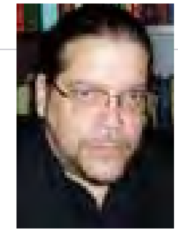

What I did on my virtual summer vacation

A storage hypervisor comes to the rescue as the elements wreak havoc in my test lab.

I**N THE TAMPA BAY** area of Florida, where I live and mostly work, we haven't been in the direct path of a significant hurricane for nearly 90 years. We've picked up a lot of collateral damage from legendary storms that have either churned their way up the Gulf of Mexico or made landfall on the East Coast over the last few decades, but we've missed the direct hit. That record holds true as I write this column.

However, that fact shouldn't be viewed as a boast. A few months ago, a tropical storm system (precursor to a hurricane) came ashore within 70 miles of my office and home bringing high winds and torrential rains that showed me it was time for a new roof. Water found its way into eight different parts of my house, including my on-premises storage lab where a rack of Dell servers (not virtualized) and a variety of storage systems from different vendors (all virtualized using DataCore Software Corp.'s SANsymphony-V) were under the spray. Electricity and water don't mix; damage was done.

The good news was that most of my data was protected via a combination of tape-based and disk-based replication. Key to my data's survival was storage virtualization technology, which lets me consolidate all my disk-based protection policies in one place—in a "storage hypervisor" that is, though the use of that term is likely to cause acid reflux for the VMware fan boys reading this month's column.

But I believe "storage hypervisor" is a correct way of describing the functionality of SANsymphony-V (SS-V). Using the software, which I run on a pair of clustered Windows 2008 R2 servers, I collectivize my storage capacity for ease of allocation in the form of virtual volumes—a process that ignores the brand names on the rigs themselves. It doesn't matter one whit whether my box of Seagate hard disks says EMC, HDS, IBM, NetApp, Promise Technologies, Sun, X-IO or whatever on its bezel plate. Disk is disk.

What might matter to some folks are the "value-add" services that vendors insist on joining at the hip to commodity gear so they can charge more for it. The bulk of these "services"—software functions running on the most commodity storage hardware controller of all, an x86 PC motherboard—can, as VMware likes to say, be abstracted away from their commodity host hardware (read "virtualized"). That's exactly what SS-V does. It's a software-based controller that runs on a server and competes well with appliance-based approaches such as IBM SAN Volume Controller or Hitachi Data Systems' latest variation on its Universal Storage Platform. Those vendors just include the server with the software for those who like to sniff the tin.

Anyway, I call SANsymphony-V a storage hypervisor because it's just that: A software abstraction layer that allows me to pool and allocate/deallocate hardware resources on demand while dynamically associating the right set of value-add services with a virtual volume created to support workload data. If the data from one application is mission-critical, I can allocate continuous data protection (CDP), mirroring and even remote replication services to it by clicking on checkboxes as I create the volume and entering a few specifics (how frequently I want to make a CDP snapshot or where I want the mirrored or replicated data to go). I don't have to worry whether the gear beneath the volume has EMC or IBM or HDS on its faceplate, which is a huge improvement over the all-too-common mirroring lock-in that requires the same vendor's brand name on every participant in the replication scheme. Most importantly, all my data protection policies (except for my tape processes) can be managed from one place and tested on an ad-hoc basis.

In the end, for a fraction of the price of a name-brand array, I had built a very sizable storage rig, which I then virtualized and added to my DataCore pools.

The big problem from my lab's water damage and the subsequent havoc created by a waterspout a few days later that took out our local power substation (a story for another time), was the loss of physical gear. I was running on my redundant cluster components and wanted to get back to my fully redundant and highly available storage infrastructure. But I couldn't afford to replace all the damaged gear in the current economy.

Instead, I decided to put the hypervisor theory to the acid test. From StarTech, I acquired some generic enclosures with eSATA connections and built a server with a bunch of external eSATA ports and SATA III internal drive bays. The bays and enclosures were then fitted with 4 TB SATA drives from Seagate. In the end, for a fraction of the price of a name-brand array, I had built a very sizable storage rig, which I then virtualized and added to my DataCore pools. SS-V uses DRAM and flash in the server to front end all the disk devices, so I receive very respectable performance even out of high-capacity, slow-speed drives.

It took a lot longer and was much more expensive to replace the roof, a task just completed. So, how was your summer? ■

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

MODERN BACKUP

With more and more data to protect, weekly fulls and nightly incrementals may not cut it anymore. It's time to consider alternatives to traditional backup apps and processes.

By Brien M. Posey

PLENTY OF IT shops still perform nightly incremental backups along with weekly full backups, but many organizations are increasingly finding that their data—and the recovery requirements for that data—are breaking the backup models they've relied on for so long. For storage managers addressing inadequate backup operations, this may mean confronting the difficult but critical task of backup modernization.

Backup modernization can be a somewhat painful process; you not only need to choose a backup technology, you need to consider the impact the transition will have on key business processes and requirements.

BACKUP ALTERNATIVES TO CONSIDER

When it comes to modernizing your backups, there are many solutions available ranging from the mundane and utilitarian to the exotic. Even so, there are three main flavors of data protection in use today:

- Continuous data protection (CDP)
- Snapshots
- Image-based backups

CDP technology protects data on a nearly continuous basis. Rather than running a large monolithic backup overnight, CDP products back up data every few minutes, 24 hours a day.

CDP products work by initially replicating data to a disk-based backup on a block-by-block basis. The software then monitors data for changes to the stored blocks or the creation of new blocks. When a block is created or modified, it's backed up. An index tracks versioning information and data deduplication ensures only unique blocks are stored on the backup media.

Snapshots aren't the same thing as backups because they never create a copy of the data. Instead, snapshots provide you with a way to roll a virtual machine (VM), file or application back to an earlier point in time. Snapshots can be based on the use of differencing disks or pointers. Because snapshots aren't actually backups, some backup vendors offer snapshots as a way of augmenting their product's recovery capabilities rather than offering snapshots as a standalone protective mechanism.

Image-based backups represent one of the newer approaches to backup, and are used to back up VMs. The idea behind this type of backup is that the backup process captures a VM as a whole. If a recovery operation is required, then a copy of the VM is usually mounted in a sandbox environment so the data can be extracted. The sandbox mounting capability is also sometimes used to provide native recovery testing capabilities or even virtual lab capabilities. Image-based backups offer tremendous flexibility as long as your protected resources are all on virtual servers.

IMPORTANT BUSINESS CONSIDERATIONS

Regardless of what type of backup technology you choose to implement, there are some critically important considerations to take into account with regard to your organization's business needs. Some of these factors need to be considered before you purchase a new backup system; others need to be taken into account once the new backup process is in place.

Retention requirements. One of the first things you need to think about when choosing a modern backup alternative is your backup retention requirements. In other words, how far back in time do you need to go to retrieve data?

The reason why this is important is because most modern backup solutions are disk based, cloud based or both. Tape-based backups provide a nearly unlimited retention span because you can back up to tape and then keep the tape for as long as you like, which isn't necessarily the case with disk-based backups. Disks have a finite capacity, and that capacity will impact the total amount of historical data you're able to retain within your backups.

> Snapshot rollbacks can cause database corruption unless the snapshot product is specifically designed to work with the application running on your server.

Even if disk capacity weren't an issue, some modern backup applications impose synthetic limits. For instance, some CDP products differentiate between short-term protection (disk) and long-term protection (tape), and place very strict limits on the total number of recovery points that can be stored within short-term protection.

Agent compatibility. If the backup solution you're considering is agent based, then agent compatibility must be a major consideration prior to purchase. Although most of the major backup application players offer agents for the most popular operating systems, you need to verify that agents exist for the operating systems you're running in your own environment.

Another consideration that's sometimes overlooked is compatibility with future operating systems. For example, Windows Server 2012 R2 is soon to be released. Some backup vendors already offer support for the new operating system, but others don't. If you plan to migrate to Windows Server 2012 R2 in the near future, you'll need to ensure that any backup vendor under consideration will support the new operating system.

Application awareness. Application awareness is one of the most important criteria in selecting a backup application. If you're backing up anything other than file data, your backup software must support the applications you're running.

For CDP or image-based backup products, ensuring application awareness usually means verifying that the backup product includes a Microsoft Volume Shadow Copy Service (VSS) writer for the applications running on the servers you're backing up. In the case of snapshot products, however, you'll need to look for granular application rollback capabilities.

Although most snapshot utilities support rolling back the entire server, this can have disastrous consequences for database applications because snapshots don't capture transactions stored in the server's memory at the time the snapshot is taken. As such, snapshot rollbacks can cause database corruption unless the snapshot product is specifically designed to work with the application running on your server.

The initial backup. After you've purchased and implemented a modern backup solution, there are some things you'll need to consider regarding your first backup.

Don't ditch that legacy backup app yet

WHEN IT COMES to replacing outdated technology, it's often a matter of "out with the old, in with the new." For backup infrastructures, however, it would be foolhardy to immediately dispose of your legacy backup system. Most organizations have data retention requirements, so the legacy backup hardware and software needs to remain at least until the last backup that was created with it ages beyond the required backup retention.

When the backups created using your legacy backup system outlive their useful lifespan, you'll need to find a way to securely dispose of the legacy backup media. If the old backup system is tape based, for example, you may be able to demagnetize and recycle the old tapes. There are also services available that will physically destroy outdated backup tapes and other media. ■

Because the new backup system has never been used to back up your data, it will have to start by making full backups of everything. Depending on how much data you have on hand, this initial seeding process can be quite time-consuming.

In addition, your servers may be left vulnerable during the initial seeding process. For example, let's say the initial backup takes three days to complete and a server fails a day and a half into the backup process. You may or may not be able to recover the server depending upon what has been backed up at that point.

One tempting solution to this problem might be to run your legacy backup software in addition to the new backup software. However, this can cause problems. For example, if both products manipulate the archive bit on file data then the backups can become confused as to what they have/have not backed up. Similarly, database applications such as Microsoft Exchange require backups to commit transaction logs to a database. Running two separate backup products can result in missing log files, which can impact your ability to perform a restoration.

A better solution is to implement the

new backup solution gradually, rather than initially configuring it to back up every resource in your entire organization. For instance, you might start by backing up one application or one virtual machine at a time. You can still safely use your legacy backup product to back up everything else. This approach minimizes the amount of data that is left vulnerable to loss during the transition process. It also greatly reduces the strain that the initial seeding process places on your network and storage infrastructure.

Recovery testing. This is an important part of the data protection process in any organization. However, recovery testing becomes even more important if you have recently transitioned to a new backup solution.

Once you begin the transition process, you should perform recovery testing on at least a weekly basis for the first six months. During the first few months the new backup system is in use, you're likely to make adjustments to the backup configuration and these adjustments can sometimes have unexpected consequences. The only way to verify that you're still being protected is by regularly testing your ability to perform recovery operations.

Modernizing a backup infrastructure may not be as simple a process as some backup vendors would lead you to believe. There's a lot of work that goes into choosing backup alternatives and then making sure they're properly implemented and are adequately protecting your data. But once you're past the initial phases of implementation and verifying operations, you'll have a much more flexible and scalable data protection process in place. ■

BRIEN POSEY *is a Microsoft MVP with two decades of IT experience. Before becoming a freelance technical writer, Brien worked as a CIO for a national chain of hospitals and health care facilities. He has also served as a network administrator for some of the nation's largest insurance companies and for the Department of Defense at Fort Knox.*

NETWORKING SERVER-ATTACHED STORAGE

Direct-attached storage — long seen as second class to networked storage — is gaining renewed attention as emerging technologies offer easy ways to pool and share this highly scalable storage resource.

By Jacob N. Gsoedl

IN RECENT YEARS, traditional storage systems have been challenged by alternate approaches to building storage arrays. Those traditional systems typically package storage controllers, disks, interfaces and firmware into a proprietary array whose components and inner workings are esoteric and usually only privy to the array vendor. Because those systems are costly, complex and lack openness, alternatives have emerged that have gained the attention of some storage shops.

A changing computing landscape that increasingly favors mobile and cloud computing has also been an impetus for more open and cost-effective storage platforms to surface.

Among the most promising alternatives to traditional storage array architectures are products that network directly attached server-based storage into a shared storage pool—storage systems that are also known as networked server-based storage. Being able to leverage commodity x86-based servers and their directly attached disks, and combine that with storage software that runs on standard servers to provide a shared storage pool, enables more open storage platforms at a price point far below that of traditional proprietary systems. Networked server-based storage is offered by a growing list of startups and has been the favorite storage architecture of cloud service providers, but it's now finding its way into the product portfolios of established storage vendors.

MAKING THE CASE FOR NETWORKED SERVER-BASED STORAGE

Networked server-based storage is succeeding for several reasons. First and foremost, it's cost-effective. Traditionally, storage arrays have been based on proprietary hardware and controllers optimized for storage processing to cope with the low-latency and high-bandwidth requirements of storage arrays that had to be accessed by a host of systems. Storage vendors have been charging dearly for that approach, establishing a lucrative, high-margin storage hardware business. But as multicore x86-based systems, interfaces and networks became faster and the performance gap compared to proprietary storage hardware narrowed and eventually closed, traditional storage vendors stayed the course, keeping storage hardware proprietary and margins high. It was just a matter of time before networked storage systems that run on standard X86-based servers emerged.

"Networked server-based storage that uses commodity server hardware is less expensive, because you're not paying the vendor markup for [host bus adapters] HBAs, adapters, disks and other components," said Terri McClure, a senior analyst at Milford, Mass.-based Enterprise Strategy Group (ESG).

Networked server-based storage reaps the benefits of virtualization. Decoupling the storage software from the underlying hardware enables the storage stack to run on all types of hardware, including low-cost commodity x86-based systems and their direct-attached storage. By taking advantage of virtualization, networked server-based storage helps increase storage utilization by pooling otherwise underutilized direct-attached server storage into a shared storage resource. Some networked server-based storage products combine storage and virtual machine (VM) processing on the same server hardware to further increase server resource utilization; this means a more cost-efficient data center that reduces data center hardware, and lowers power consumption and space requirements.

Networked server-based storage is also easier to manage. Storage arrays, especially Fibre Channel systems, are complex to deploy and manage, and require the expertise of storage specialists. Networked server-based storage, on the other hand, can be managed by the server team, which takes fewer IT resources and results in lower operational costs. Ease of management is the other main reason for networked server-based storage having done so well with cloud service providers and consumer-facing Web 2.0 companies.

THE MANY FACES OF NETWORKED SERVER-BASED STORAGE

Contemporary networked server-based storage products come in a variety of flavors and from different roots, but they have one thing in common: They all use standard server hardware and aggregate directly attached server storage into a single shared storage pool.

Virtual storage appliances (VSAs) run the storage software within VMs and are distributed as virtual machine images that are deployed to one or multiple physical hosts. They combine the directly attached storage of the hosts they run on into a shared storage pool. The number of hosts supported by currently available VSAs varies by vendor. VMware's vSphere Storage Appliance combines the local storage of up to three hosts into a shared storage resource that's accessed via the Network File System and managed through vCenter. The Hewlett-Packard (HP) StoreVirtual VSA, which is based on the LeftHand OS, eclipses the VMware VSA by pooling the local storage of up to 16 hosts to present it as shared storage that's accessible via iSCSI. NetApp's Data Ontap Edge is a VM that runs Data Ontap, currently supporting only a single server node, but able to seamlessly interact with other NetApp storage.

The current crop of VSAs is well-suited for small environments, such as small- and medium-sized businesses (SMBs) or branch offices of larger companies. Being able to deploy them on existing servers, without the need for additional storage hardware, makes them very cost effective, as well as easy to deploy and support. VSAs clearly have the potential to expand beyond the SMB market and today's limits are mostly imposed by VSA vendors that are either protecting their lucrative high-end storage business—in the case of NetApp and to some extent HP—or trying to avoid upsetting the existing ecosystem that could harm adoption, as in the case of VMware. VMware in particular, with its VSA and the potential of VMware Virtual SAN (vSAN), is walking a fine line so it won't agitate storage vendors that have integrated their storage systems with the various VMware APIs.

Converged storage systems. To maximize server resources, a group of startups has emerged that offers products that combine storage and VM processing on the same server hardware. The idea is to dedicate a certain percentage of the computing resources of the underlying hosts to storage tasks and the remaining resources to virtual machines. Besides high server resource utilization, ease of management and low cost are the key benefits of

Networked server-based storage product category comparison

	VALUE PROPOSITION	LIMITATIONS/CHALLENGES	TARGET MARKET
Virtual Storage Appliances	• Simplicity	• Currently limited to the small- and medium-sized business (SMB) market	• SMBs
Converged Storage Systems	• Maximize server resource utilization • Simplified administration	• Potential of virtual machines interfering with storage services • Challenge to expand beyond the SMB market	• SMBs • Some vendors target enterprises
Distributed File Systems	• File-based storage on commodity server hardware	• Possible support challenges • Competition with traditional NAS vendors	• SMBs • Enterprises • Service providers
Object Storage	• Scalability • Designed and optimized for cloud storage	• Not appropriate for transactional apps • Requires some degree of integration with traditional storage	• Enterprises • Service providers
Windows Server 2012 Clustered Storage Spaces	• Integrated into the OS with high-end storage features	• Limited to Windows Server 2012 and a small number of server nodes	• Windows Server 2012 customers

combining storage and VM processing in a single system that usually consists of multiple server nodes. On the downside, combining storage and VMs in a single system can limit the ability to scale and increases the likelihood of VM processing interfering with storage processing and vice versa.

Currently, a handful of vendors offer converged virtualization and storage products. Scale Computing's HC3 runs the open source Red Hat KVM hypervisor on top of its server-based multinode scale-out network-attached storage (NAS) to target the SMB market. Nutanix runs its storage stack as an additional service parallel to the VMs, virtualizing storage from physical server nodes into a unified pool of scale-out converged storage, targeting both SMBs and enterprises. The SimpliVity OmniCube scale-out NAS, with features like real-time deduplication and compression, can also host VMs and is intended for both SMBs and enterprises. By adding virtual machine processing, converged storage systems literally take the value proposition of networked server-based storage to a new level.

Distributed file systems. Even though the NAS market is dominated by large storage vendors, open source distributed file systems have enabled others to build NAS systems that run on commodity x86-based servers. A case in point is Red Hat Storage Server, which is based on the GlusterFS scale-out NAS file system Red Hat obtained with its acquisition of Gluster in 2011. Another example is the Oracle Zeta File System (ZFS), which was originally developed by Sun and is used in Oracle's ZFS Storage 7000 series; ZFS is open source software that has found its way into products from vendors such as Nexenta, which markets its ZFS product as a software-defined storage offering.

"Nexenta uses ZFS and puts its management software [around it] to sell it to [users to] create NAS storage or NAS gateways," said Greg Schulz, founder and senior analyst at Stillwater, Minn.-based StorageIO. Similarly, the Apache Hadoop Distributed File System runs on commodity servers, and it supports anywhere from a few to thousands of server nodes to create large shared-storage pools.

Networked server-based storage products based on a distributed file system can be a cost-effective storage option for large amounts of file-based content. The challenges of foregoing a more established NAS product for a more novel approach is offset by significant cost savings and product-specific features and behaviors that may not be available on a traditional NAS system.

Object storage. Object storage is in many ways similar to scale-out NAS, but instead of files it manages objects with unique identifiers. Objects support rich metadata far beyond the file-system attributes of NAS. Object storage is usually accessed through HTTP APIs such as REST, and data redundancy is achieved by storing objects multiple times on different nodes.

To meet the requirements of the cloud storage market, object storage is designed to run distributed on multiple server nodes, and it scales by simply adding more server nodes. In other words, object storage is designed as networked server-based storage. Object storage is best suited for storing file-based content for both internal and external applications. It's unfit for transaction-based systems, such as databases and transaction processing applications. While many Web 2.0 companies created their own proprietary object stores, object storage products, such as Caringo CAStor, EMC Atmos, Hitachi Content Platform and NetApp StorageGRID among others, have been available for several years.

Windows Server 2012 Clustered Storage Spaces. With Windows Storage Spaces (WSS), Server Message Block 3.0 enhancements, data deduplication and thin provisioning support, Windows Server 2012 has all the ingredients for building high-end, networked server-based storage systems. By combining WSS with the failover clustering feature, Windows Server 2012 provides for building clustered storage spaces that consist of multiple server nodes. Clustered Storage Spaces combine a small number of servers, typically two or four, with a set of serial-attached SCSI (SAS) JBOD enclosures connected to all servers. Access is unified into a single namespace via Cluster Shared Volumes that can be accessed by all server nodes, regardless of the number of servers, JBOD enclosures and provisioned virtual disks. A January 2013 ESG lab report concludes that the storage efficiency, agility and transparency of new and improved features make the high-performing, cost-effective Windows Server 2012 a no-brainer for small and large businesses alike.

TRENDING TOWARD SOFTWARE-DEFINED INFRASTRUCTURE SERVICES

The benefits of low-cost, high resource utilization and simplified management bode well for networked server-based storage platforms. The rapid growth of unstructured data and a need to scale seamlessly internally and into the cloud are key drivers that should help networked server-based storage solutions expand the market. From a longer term perspective, virtualization and the move toward software-defined infrastructure and data centers will result in an increased decoupling of infrastructure services from the underlying hardware. Networked server-based storage is just a step toward a future where VMs, networking, storage, security and other infrastructure services run on top of shared physical hardware components. ■

JACOB N. GSOEDL *is a freelance writer and a corporate director for business systems. He can be reached at jacobslab@live.com.*

10 TIPS FOR VIRTUAL DESKTOP STORAGE

Boot storms and antivirus scans are just a couple of the virtual desktop events that can bring your storage to its knees. Find out how to configure VDI storage the right way.

By David Davis

STORAGE IS ONE of the hottest topics in virtualization, which is one of the hottest topics in storage. These once distinct silos of data center technology are quickly converging. And it's increasingly easy to see why: They desperately need each other.

A virtual infrastructure can only be robust if it's backed by shared storage that supplies the endless low-latency I/O demanded by virtual machines (VMs). Virtual environments place more demands on a storage infrastructure than ever before in most data centers. Similarly, infrastructure virtualization is more demanding than ever on the physical servers in the environment.

A resilient and responsive storage infrastructure is a key part of any successful desktop consolidation project. While that's true for server virtualization, it's even more of an issue with virtual desktop infrastructure (VDI) because virtual desktops demand so much more I/O than most virtualized servers. Users do things that servers never would, such as Web surfing, viewing videos and running remote display protocols. As a result, users create much more diverse and, in many cases, even more demanding I/O workloads than those created by virtualized servers. In addition, virtual desktops may have agents on them (such as antivirus agents) that can (conceivably) kick off all at the same time and create massive amounts of disk scanning as the agents look for virus signature matches. Finally, when a physical server that runs hundreds of VMs is powered on, those VMs will also power on and create a "boot storm" of storage I/O that can bring shared storage to its knees.

Because storage is so crucial when deploying virtual desktops, here are 10 tips for configuring and managing storage in support of virtual desktops.

1

Use a VDI I/O calculator. Before you decide on a storage system, make an actual purchase or roll out your VDI implementation, you need to do some I/O and capacity calculations. There are free calculation tools available that will help you to estimate how many I/Os per second (IOPS) your storage will require and how much capacity your VDI will consume. By determining these requirements up front, you could avoid losing time and money down the road. Examples of VDI calculators are MyVirtualCloud and Unidesk.

2

Use linked clones. It's common to clone virtual servers from golden images for a particular server operating system (OS). VMware Horizon View takes cloning a step further; not only are all the VMs based on a single image (or sometimes multiple images) but the only additional storage consumption from each new VDI VM created are the changed blocks from the golden image. VMware calls this "linked clones." All VMs are based on a single VM image, and the other VMs are all linked to that image but still retain the power to modify their own guest OS image.

3

Consider user data management. Much of the storage consumed and many of the I/Os created from virtualized desktops are related to user data. And the truth is that user data is a mess at most companies. User profiles are bloated, corrupt and missing information a user needs. Ideally, a VDI infrastructure should consist of a VM that's a linked clone of a "golden VM" with access to packaged virtualized applications, and with user profile data properly managed and layered on top. Offerings like LiquidwareLabs ProfileUnity have a good track record for helping to get user data under control, reduced in size, properly managed, and layered on top of virtual and physical VMs as needed.

4 **Use flash or hybrid storage.** Since VDI requires so many more IOPS than more traditional virtual infrastructures, many IT shops are turning to storage products that use solid-state storage in dedicated flash arrays or hybrid storage systems that augment hard disk drives with flash drives. These systems can provide far more IOPS with fewer disks. Examples of vendors offering flash or hybrid storage solutions include Tegile Systems and Tintri.

5 **Consider virtualization-aware storage.** As noted at the beginning of this article, with virtualization and storage converging, there are new storage products that are commonly referred to as "virtualization-aware storage." While some storage vendors will claim they're virtualization-aware because they support the VMware vStorage APIs for Array Integration (VAAI) and the vStorage APIs for Storage Awareness (VASA), support for VAAI and VASA are minimum requirements for storage used in any new virtual infrastructure—they don't make storage comparable to other storage arrays that actually communicate with VMware vCenter to learn about your VMs. With that kind of knowledge, the array can give you per-VM I/O utilization and other important operational statistics. That's much closer to being a true virtualization-aware storage system. Keep in mind that depending on the storage system you select, there may be various levels of virtualization-aware functionality.

6 **Know your I/O and measure it.** No matter what storage you select, you need to have an idea of its typical I/O utilization. Ideally, if you're using virtualization-aware storage, this should be much easier. However, even if you use a more traditional storage array, there are many virtualization performance analysis tools—such as VMware's vCenter and vCenter Operations Manager, or third-party tools such as Dell's Foglight for Virtualization and the Xangati Management Dashboard (XMD) product suite—that can help you measure your I/O utilization.

7 **Simulate a boot storm.** In addition to the typical bursty I/O demands from virtual desktops, the other major concern is a boot storm. A boot storm is when all the VMs on a single host or all VMs in the VDI are powered on and boot at about the same time. While that probably wouldn't happen on a regular basis, it is a concern that at some point would warrant testing to measure what happens to your storage array and virtual machine performance.

8 **Use agentless antivirus.** With an antivirus storm, all antivirus agents in your VDI wake up at the same time and begin scanning the VMs that will flood the storage system with I/O. You shouldn't use agent-based antivirus applications in VDI environments, as there are much better alternatives that don't have the same effect. An agentless solution, such as a VMware vShield Endpoint-compatible product with your vSphere VDI, will eliminate any chance of an antivirus storm as it will offload antivirus scanning from the end-user virtual desktops. Some examples of vShield Endpoint-compatible antivirus products include Symantec Endpoint Protection, McAfee Management for Optimized Virtual Environments and Trend Micro Deep Security.

9 **Don't back up.** While it should go without saying that backups just aren't needed, it's important enough to emphasize it again. If you use the layered approach described earlier with the linked clone VM OS, virtualized apps and user customizations/data on top, then there's no need to back up your VMs. This is because the OS is a clone and the applications are virtualized and on a share (already backed up). As with physical desktops, the end-user data must still be backed up but that data, in the case of VDI, should be stored on a network share.

10 **Consider host-based caching.** Host-based caching is an emerging product category—so new that many products are in stealth mode or otherwise under the radar. But there will soon be numerous host-based caching offerings available. These products are designed to improve storage I/O performance. They typically use solid-state storage on each host as a cache to make VMs run faster. The additional I/O performance they can squeeze out can make a huge difference in VDI virtual machine response. But not all host-based caching offerings are the same, so do a thorough investigation to ensure that a product will suit your specific needs.

STORAGE CAN MAKE OR BREAK YOUR VDI

The storage you use for deploying virtual desktops is crucial. If you don't have the right storage infrastructure in place to support the I/O demands of the virtual desktop infrastructure, you'll experience slow response times for your applications and may even experience outages.

Follow the tips outline here and your virtual desktop project should get the storage configuration it needs. ■

DAVID DAVIS *is a virtualization expert who blogs at VirtualizationSoftware.com. His video training courses covering vSphere and Hyper-V are available from Pluralsight.com. Follow him on Twitter using @DavidMDavis.*

If you think data protection is just about disk, think again

Disk is great for backups and speedy recoveries, and should play a key role in a DR plan. But tape is still clearly the best choice for data retention.

THE CALENDAR MAY read 2013, but I still hear IT professionals and vendors assert that disk-based storage is always best when it comes to protecting business data. To me, that's more of a "1999" mentality.

Admittedly, from a performance perspective, disk will usually be your best bet. You can compress and deduplicate data more efficiently to disk, and it's certainly a faster data-restoration medium when you're trying to re-establish part of your environment from a backup. In general, disk is superior to tape- or cloud-based storage as the first tier of recovery.

But when it comes to long-term retention—adhering to a seven-, 10- or 25-year data preservation policy—using disk alone is nearly always impractical.

TAPE'S TARNISHED PAST

So why has tape been reduced to lingering in the shadow of disk as a viable long-term data retention tier? In part, it may be because data deduplication became incredibly popular and dedupe is a disk-centric process.

But it was mainly because every five to seven years, tape vendors revised their tape formats and form factors. Over my career, I've used 4mm DAT, 8mm, DLT and Linear Tape-Open (LTO), just to name a few. While vendors tried to incorporate backward compatibility with the new formats, it didn't stop IT folks from viewing tape as a rather old-fashioned, hard-to-manage medium with reliability issues.

Imagine finding out that the data you just stored is now on an obsolete tape format, and regulations mandate keeping the information for another 20 years. You have two options: hang onto the old tapes and the compatible tape libraries required to read them for two more decades, or transfer all the data to the newer tape format. The second option is not only a major undertaking, but it might make it harder to prove that the original data is intact and unaltered after the migration.

LTO SETS TAPE STANDARDS

But there's good news. The biggest "anti-tape" arguments are more or less invalid now due to the following:

- Most tape vendors have standardized on LTO, an open-format cartridge. LTO tape libraries not only work with the medium's current iteration, LTO-6, they read data on LTO-5, -4 and -3 tapes. It's a consistent retention medium, so organizations no longer need to maintain old tape libraries or migrate data to keep it readable.

- Believe it or not, today's tape cartridges with LTO have a lower mean time between failures than individual hard drives. The claim that today's tapes are failure-prone compared with disk is simply outdated and untrue. Granted, most storage operations don't write data to a single spindle; they write it to an array engineered for fault tolerance across spindles. But the point is that data written to tape is secure from both a reliability and encryption perspective.

Tape has re-emerged from disk's shadow, serving as yet another example of the pattern of constant change in IT. Tape is now longer lived, more reliable and faster than it was in the old days.

Still, that doesn't mean tape has overtaken disk as the go-to medium for disaster recovery and short-term backup. Disk remains the best medium to recover from using a modern disk-to-disk backup solution or snapshot mechanism.

Retention is a different story, where no single offering fits every use case. After all, the phrase "long-term retention" is relative. For three years of retention, disk might be fine. For 10 years, tape is likely the better choice. (Note: The cloud is another good option but any savvy cloud provider holding your data for 10-plus years is using a large-scale, highly economical tape farm to do it.)

SOME PRACTICAL ADVICE FOR IT

Most busy IT teams are trying to achieve multiple data protection and retention goals in parallel. To support all those simultaneous efforts, it's a good idea for the team to determine which combination(s) of disk, tape and software will work best for their individual business.

You should start by pinpointing exactly what and how you need to recover. For example, if you have to recover data within seconds, you should implement snapshots. If there's a need to recover data across long distances, replication should be incorporated. When there's a need to restore from a range of previous versions of data, backup is a good choice. And if you need to recover data generated 10 years ago, you need an archive … a *tape* archive.

Current research from Enterprise Strategy Group shows that (regrettably) organizations tend to use tape more for backup than for archiving. However, "best practice" data protection and archiving requires an architecture that balances the performance benefits of disk for backup/fast restore with the cost and reliability benefits of tape for long-term retention. You want a strategic, diverse set of media and mechanisms protecting your data.

That doesn't mean deploying disjointed offerings from multiple vendors. A few providers now offer tape, disk, snapshot and replication products that have been cohesively converged to support common data protection goals for big data centers and small remote offices alike. It's even possible (a challenge, but possible) to corral the management of it all under one umbrella.

If you're trying to define your organization's data protection strategy, the worst thing you can do is base the definition on the capabilities/limitations of the systems and software you have on the floor. Instead, think strategically and objectively about how you need to recover data and accomplish the other data protection activities you need to deal with. Then choose the right technologies—which may include disk *and* tape—to help you get there. ■

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.*

Good-bye LUNs, you served us well

In a virtual server world, the concept of LUNs and the amount of attention they require just won't cut it anymore.

THE ERA OF LUNs and volumes, as we have known them for decades in the data storage industry, is quietly coming to an end. And if you ask me, it's for all the right reasons, even if storage administrators may feel threatened by the change.

In the world of physical servers, our standard practice of ganging up a set of disk drives into a RAID set to create a LUN has served us very well for decades. This LUN was created recognizing the type of application it was to serve and it was associated with all the appropriate storage services (replication, compression, snapshot and so on) the application warranted, based on its importance. That's all well and good. But then we started servicing several applications from the same LUN, and unless we overdid it or the applications were erratic, the LUN was able to serve multiple applications. If an application was important enough it got its own LUN and associated services, even if the utilization of either capacity or performance was sometimes less than ideal.

But then server virtualization came along and all hell broke loose. One or a few LUNs serving a multitude of virtual machines (VMs), maybe even several hosts, each with tens of VMs or more representing a variety of applications/workloads, simply couldn't cut it. The infamous I/O blender effect is now well understood. The precisely tuned LUN of yesteryear was wrestling with totally random I/O coming in an unpredictable onslaught from a large number of VMs. The storage controllers were overwhelmed and application performance plummeted.

The storage industry responded with a variety of solutions to this problem. The traditional storage array vendors added solid-state drives to deliver more I/Os. They also virtualized the arrays internally so all disk drives would participate in serving all VMs, rather than only a few drives getting hit constantly. And they added features to more tightly integrate their systems with VMware or Hyper-V using application programming interfaces such as vStorage APIs for Array Integration, vStorage APIs for Storage Awareness or Hyper-V's Offloaded Data Transfer. We saw a new crop of vendors come to market with storage built from the ground up for the virtual world. That list includes Gridstore, Nimble Storage, Nutanix, Scale Computing, SimpliVity, Tintri and others. Hewlett-Packard responded with its StoreVirtual VSA, presented either as a VM to be used with any storage or as a complete appliance offering. Yet other newcomers developed software that worked with existing storage arrays but delivered an order of magnitude improvement in performance, latency and capacity reduction. The classic example was Virsto Software, which was acquired by VMware.

Those latter two product categories have taken a very different path to delivering storage services to applications. They are 100% VM-centric. They completely get rid of the LUN-centricity that storage shops have been addicted to since the beginning of the SCSI era. Via policy, one assigns the level of importance to a virtual machine. That policy establishes the type and amount of storage the VM would receive, where that data would be placed, how many times it would be replicated and what type of data protection it would receive. In many such offerings, the VM can even be assigned a quality of service (QoS) to ensure it has priority to receive the appropriate resources if there's contention. These offerings are monitored on a virtual machine basis and deliver performance data, capacity utilization and other relevant information

Storage administrators can now be relieved of all the mundane chores of delivering storage to applications and instead focus on strategic matters.

on a per-VM basis. You don't back up a LUN, you back up a specific VM. You don't clone a LUN, you clone a VM. You don't manage storage from a storage console, you manage it from VMware vCenter or Microsoft System Center.

The server administrator becomes the person in charge of provisioning storage at a VM level, not the storage admin. There are still some management tasks that require the storage administrator's deep understanding of how storage works under the covers, but day-to-day storage management now shifts to the server administrator.

Activities such as creating new LUNs, tuning LUNs for changing conditions of applications, choosing the right RAID groups, and replicating volumes synchronously or asynchronously are no longer needed. A lot of those functions have been automated and are triggered (you guessed it) at the VM level. Performance metrics, including those we've historically associated with storage, are now viewed from VMware vCenter or Microsoft System Center VMM at a VM-level granularity. Working at a LUN level is simply not meaningful anymore.

Of course, there are differences in the way the various products mentioned here have implemented these functions. Products from Nutanix, Scale Computing and SimpliVity are more than storage; they're infrastructure in a box. Nutanix and SimpliVity only support VMware; Scale Computing only supports KVM today. Regardless, the storage is 100% VM-centric and you couldn't find a reference to a LUN if your life depended on it. Tintri is storage for VMware only today, while Nimble Storage supports all major hypervisors. Gridstore implements a virtual storage controller as a miniport driver for Hyper-V. It uses the virtual controller's awareness of each application's assigned priority and I/O needs to enable sophisticated QoS and auto-tuning functionality. But despite these differences, these products are conceptually bound by their VM-centricity.

To be sure, the world has not fully transformed yet. Legacy storage players are burning the midnight oil to make their array capabilities work more effectively with VMware, Hyper-V and XenServer. They're also surfacing storage management information at the server virtualization admin's console. But there is no doubt that the era of cumbersome LUN management is over. LUNs may not be dead under the covers, but they're certainly dead as the main weapon wielded by a storage administrator. LUN creation and selection is much more easily done by the storage system, working in conjunction with the application and the hypervisor, in a totally automated fashion. Doing things at a VM level, by definition, provides application awareness—knowledge that can be used for a vast number of purposes, such as load balancing, auto-alignment and delivering the right QoS.

Storage administrators can now be relieved of all the mundane chores of delivering storage to applications and instead focus on strategic matters, such as infrastructure planning, analyzing information to improve efficiency and the like. Of course, data protection, archiving and disaster recovery still need to be dealt with, as do all things related to cloud storage. It might take the industry another three years to get there but the handwriting is on the wall. A smart storage administrator is already preparing for the day when life will be LUN-free. LUN technology has served us well, but it's time to put LUNs to rest. ■

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

File storage growing, but mostly under control

ASK 10 STORAGE managers what their biggest headache is and most are likely to say "capacity." Even before being drubbed into a big data daze, we knew run-away files/unstructured data was a threat to our data centers (and our sanity). NAS systems are installed at 74% of the companies that responded to this month's survey, with an average of nine NAS systems on the floor—11% have more than 30 of the file-chomping critters installed. Overall, approximately 54% of installed NAS capacity is currently used, occupied by an average of 118 TB of data, representing a 35% rate of growth compared to last year. The biggest problem with handling all that file storage is backup (cited by 46%), followed by managing multiple systems (29%). Storage pros employ a number of different approaches to manage file storage: 34% have installed clustered NAS systems, 29% use solid-state devices to pep up performance and 13% are on the cutting edge with object-based storage systems. *—Rich Castagna*

WHAT ARE THE BIGGEST PROBLEMS YOU HAVE REGARDING YOUR FILE STORAGE SITUATION?*

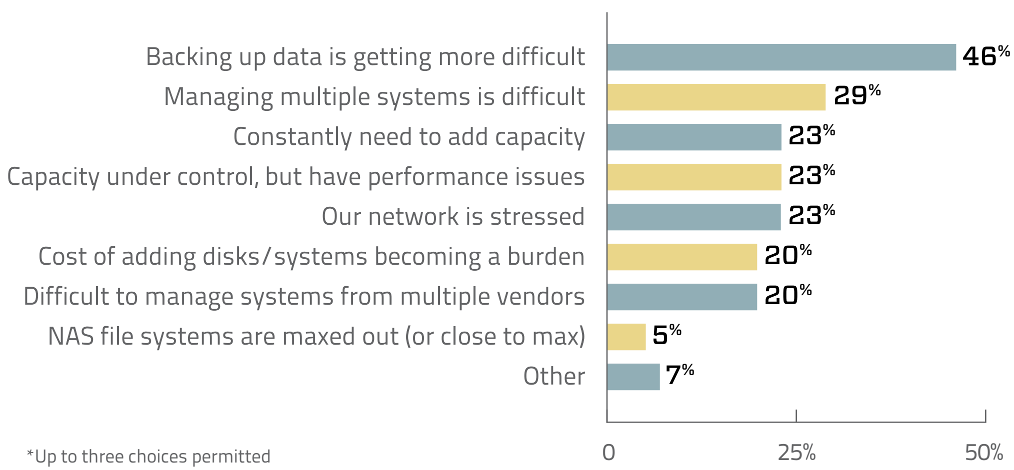

*Up to three choices permitted

COMPLETE THIS SENTENCE: "FILE/UNSTRUCTURED DATA STORAGE AT MY COMPANY IS …

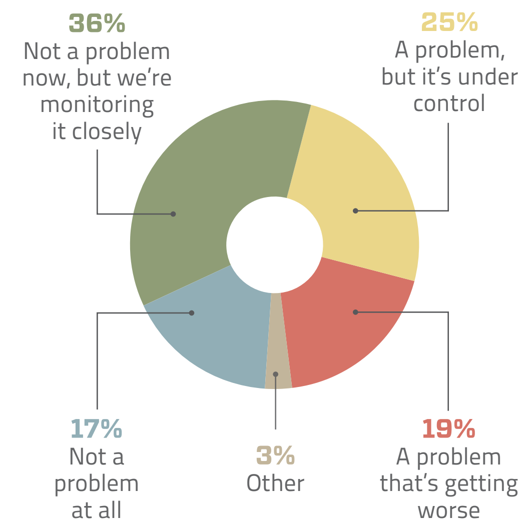

23

PERCENT OF RESPONDENTS WHO REPORT HAVING UNSTRUCTURED FILE GROWTH OF 50% OR HIGHER YEAR-OVER-YEAR

TechTarget Storage Media Group

STORAGE MAGAZINE

EDITORIAL DIRECTOR Rich Castagna
SENIOR MANAGING EDITOR Kim Hefner
EXECUTIVE EDITOR Ellen O'Brien
CONTRIBUTING EDITORS James Damoulakis, Steve Duplessie, Jacob Gsoedl

SEARCHSTORAGE.COM

EXECUTIVE EDITOR Ellen O'Brien
SENIOR NEWS DIRECTOR Dave Raffo
SENIOR NEWS WRITER Sonia R. Lelii
SENIOR WRITER Carol Sliwa
SENIOR MANAGING EDITOR Kim Hefner
ASSOCIATE SITE EDITOR Sarah Wilson

SEARCHCLOUDSTORAGE.COM
SEARCHVIRTUALSTORAGE.COM

EXECUTIVE EDITOR Ellen O'Brien
SENIOR MANAGING EDITOR Kim Hefner
ASSOCIATE SITE EDITOR Sarah Wilson

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

SENIOR SITE EDITOR Andrew Burton
MANAGING EDITOR Ed Hannan
FEATURES WRITER Todd Erickson

STORAGE DECISIONS TechTarget Conferences

EDITORIAL EVENTS MANAGER Jacquelyn Hinds

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

COVER IMAGE AND PAGE 9: MARK AIRS/IMAGEZOO/GETTY IMAGES