IT OWNER'S MANUAL

IT WORLD

Optimizing your virtual storage

The storage world is changing. Here's a look at how different technologies can meet your specific needs, and things to watch for as you move into a more virtualized environment.

Sponsored by:

TINTRI

Zero Management Storage

Is your virtual infrastructure squeezing your storage capacity?

By David Strom

AS DATA CENTERS become more virtualized, many IT managers are finding that their storage arrays can't keep up with the demand. In the past, storage was simple: You would transition from using direct-attached drives to building storage area network (SAN) arrays, and that was the only decision involved.

But as data centers migrate their physical servers to virtual machines (VMs), the world of SANs seems quaint now. Using VMs can greatly increase storage requirements by several orders of magnitude, and specialized VM storage repositories are needed to keep things under control and ensure additional storage capacity is used productively. It isn't just that VMs can gobble up gigabytes quickly as you clone new instances of your servers or migrate your physical servers to the virtual world. Now there are different kinds of virtual storage to match the particular needs of different virtualized applications and services. Storage requirements change when you add virtual desktops or virtualize your network connections.

Here's a look at the changing storage world, how your needs can impact storage technologies, and things to watch for as you move into a more virtualized environment.

Top network and data center initiatives

PROJECTS THAT ARE ON THE RADAR OR ARE ACTIVELY BEING RESEARCHED OR PILOTED INCLUDE:

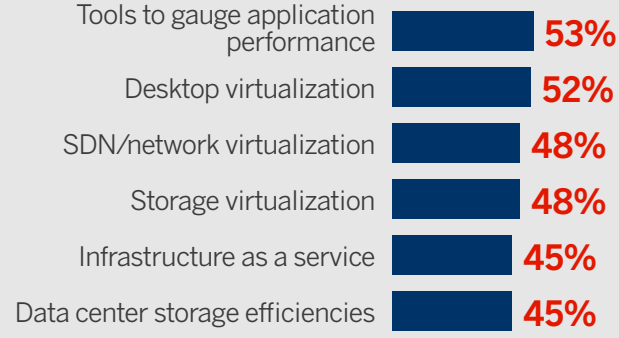

Source: 2014 State of the Network Study, Network World, 282 respondents

Virtualized storage options

SANs have a lot of benefits: You can deliver large amounts of storage to where it is needed on a network without having to rebuild a physical server's hard drives. You can centralize your storage to make backups and recovery easier, and have a single place to track and manage all of your storage needs. SANs work well in a world where most of the servers are running on physically separate machines.

But as the number of virtual servers escalates, managing your storage needs becomes much more complicated, and the plain-Jane SAN isn't always the best solution.

There are several different kinds of storage products that go beyond the traditional SAN. One alternative is a plug-in flash card that boosts performance and cuts down on overall network latencies. Like a graphics card for increasing your computer's processing power, a **plug-in flash card** is simple to install, but it must have the right kinds of drivers to work efficiently with the overall storage array. It is

Cover image by Thinkstock

also a costly and disruptive addition, since you have to bring down your SAN to install it. Further, this option is not as flexible to manage as others.

Another option is **all-flash solid-state arrays**. These are very costly, require custom hardware and are hard to scale. But they are fast. All-flash technology makes sense when data doesn't change often and network latency matters, such as for real-time online database access.

However, all-flash isn't always the best solution, even though its cost has come down. Some vendors have begun to combine flash with traditional rotating media, automatically moving data between the two types of storage as it ages or becomes more of an archive. Such solutions combine some of the advantages of flash with the lower cost of magnetic media.

Finally, another choice is to use software-defined storage to serve up blocks of storage for NoSQL or big data applications. Vendors design an entire rack of equipment that works together to provide a more coherent and converged solution that bridges storage, computing power and integrated networking. The advantage of these arrays is that the vendor tests all the different components to ensure that they work together and deliver optimum performance. In addition, they can be installed and deployed quickly and offer some advanced virtualization features.

Prevent empty space

No matter which type of storage array you use, one of the problems with virtualized storage is that you can end up provisioning a lot of empty and unused space on your physical hard drives that never gets

Do you know your storage options?

Storage architecture options with and without flash exist in a dizzying array of flavors. Make sure you're educated on the pros and cons of your architectural options before making a choice on what is best for your application environment.

	Pros	Cons
Plug-in flash card	✓ Delivers IOPS boost ✓ Low latency to local apps	✗ Flash local to server ✗ Uses host resources ✗ No built-in resiliency ✗ High cost per gigabyte of capacity ✗ Not good for random data access
All-flash solid-state array	✓ High IOPS per dollar ✓ Hundreds of thousands of IOPS	✗ High cost per gigabyte of capacity ✗ Requires custom hardware ✗ Doesn't scale out for continued growth ✗ Hard for most apps to make use of all the performance all of the time
Hybrid flash and magnetic array	✓ Great price/performance (Cost per IOPS and cost per gigabyte) ✓ Flash always available to all hosts ✓ Tiering of data between flash and disk	✗ Limited scale ✗ Controller and network bottlenecks ✗ Must force-fit flash into traditional array form factor with SATA or SAS
Software-defined scale-out objective Storage	✓ Scale-out enables storage capacity to be grown easily from TBs to PBs ✓ No fork lift upgrades required to increase capacity ✓ Hardware flexibility, leveraging desired ratios of flash and disk ✓ Intelligent object/workload based on tiering with direct application integration available ✓ Entire system from TBs to PBs is addressed as a single namespace	✗ Older object storage solutions designed for capacity, not performance, so check that the architecture was designed for flash ✗ Some object storage implementations are not backwards compatible with existing apps, so verify support for physical and virtual apps ✗ Traditional scale-out systems incur latency as the number of nodes increases, so check for how bottlenecks are overcome in the architecture

Source: Coho Data

Should virtualized storage use solid-state flash drives?

The notion of using solid-state drives to speed up storage operations isn't new; such technology has been a part of desktop computers for more than a decade. For many years, storage vendors have been offering some form of tiered storage array where data can be moved from rotating media to SSDs. What is new is the proliferation of SSD solutions for network-attached storage devices and specialty storage appliances that understand virtualized servers and application requirements. And as the cost premium for SSDs over traditional rotating media decreases, the idea of using flash for storage becomes more compelling.

There are more than a dozen vendors that offer some kind of SSD storage or combine solid-state electronics with traditional rotating media, for both dedicated storage and SANs that support virtual servers. These include Pure Storage, Tintri, Fusion-io, Violin, Tegile, Whiptail, Drobo and Nimble, along with traditional storage vendors such as Hewlett-Packard, Dell, IBM and Hitachi. These vendors all offer different tradeoffs between cost and performance.

All of the hybrid arrays (those that have both solid-state and magnetic storage mechanisms) have some kind of management software that moves frequently used data to the SSD storage area. But depending on your own situation, this may or may not be satisfactory because of the amount of data that needs to be on the SSDs or the amount of time that operator intervention is required to manage the process. Some of the newer products have better or more automated tools to handle drive management and data migration, and some are more closely tuned to particular virtualized networks. Before you purchase an SSD-based array, understand what its role will be in simplifying your overall storage technology portfolio, and examine how it can help reduce your overall application latency and improve performance and response times. And if you are considering using SSDs in backup or recovery processes, test to see how short your backup window can be as a result or how the drives can improve your disaster recovery processes.

One downside of flash is limited read/write cycles, meaning that eventually the memory will wear out and no longer be able to store any data.

Finally, while still in beta, VMware's Virtual SAN allows you to pool storage across vSphere hosts, create virtual storage policies and set up simple two-click storage provisioning. It can also use SSDs as a cache.

touched by any of your VMs. This is because when you provision your SAN, you generally don't know exactly how much storage you'll need at first and so you tend to err on the high side, with volumes large enough to meet your needs for the life of the server. Then the same thing happens when you create the individual VMs on each virtual disk partition.

A typical scenario might start out with a terabyte of storage allocated to a virtual server, and by the time you get to your actual applications, you are using less than 100GB. This is because you have to take bits and pieces off for various management and overhead reasons — first to partition a usable storage repository from the array, then to divide it into logical volumes, then to provision it for the various VMs that are going to commit to the virtual drive. With desktop computers, we don't think about the entire overhead because typical PCs have such big drives that we don't ever fill them up and there is a lot of empty space. However, in the virtual world, that empty space carries a cost, which means we have to be more efficient about how it is split up. Many of the storage vendors use clever tricks that enable the drives to expand beyond the initial capacity that the VM thinks they have — but in a way that is flexible and doesn't require a lot of operator intervention.

Virtualization optimization techniques

Because virtual storage has many different dimensions, the ideal storage product should be configured to meet particular needs with a number of features. These include dynamic thin provisioning, automatic deduplication, cloning and VM snapshots, pooling of disk volumes, load balancing across multiple volumes to the same server and use of storage hypervisors.

For example, the kind of storage

available for desktop or application virtualization will be focused on deduplication, while virtual servers might be configured for making quick snapshots and image clones of their storage repositories. Here's a look at each of these features in more detail.

Being thin is in. To avoid the problem of having lots of empty space in the modern data center, you want to be able to thinly provision volumes. Thin provisioning means that commonly used operating system and application files can share one physical repository instead of being stored in multiple virtual hard drives, so you save on the copious space you need for these files. But you should be able to take things a step further and provision your virtual drive to have room to grow without adjusting the size of your virtual disks. This is where dynamic thin provisioning comes into play. Dynamic provisioning means that you can adjust the size of a storage volume while a VM is running and using the storage. Most of the current storage products offer this feature.

Another way to avoid duplicate data is to use a storage product that supports **deduplication**. This technology is especially useful when backup jobs are done on magnetic storage, since most of the data stored in backups is duplicates of operating system and commonly used documents. And the savings can be huge. For example, California's Emergency Management Agency in Sacramento was able to cut its storage needs by 50:1, which is a typical savings using deduplication. "We have lots of geospatial data and saved tons of space when we did the dedupes on the backups," says Lovell Hopper, an IT manager at the agency.

Top technology opportunities for productivity initiatives

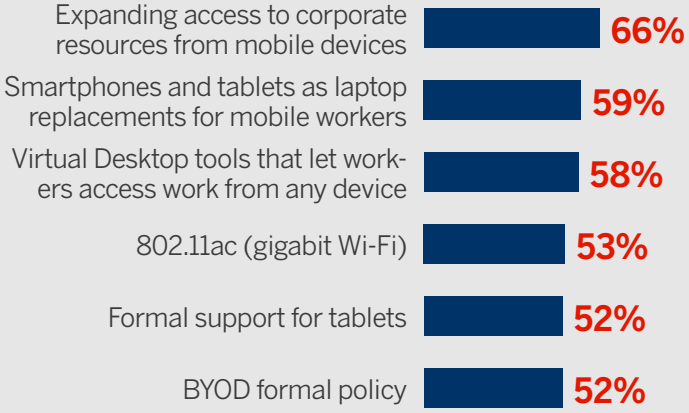

Source: 2014 State of the Network Study, Network World, 282 respondents

Scaling up your environment. As mentioned earlier, one issue for VM-aware storage is that space gets consumed quickly as new VMs are added or new application workloads are virtualized. This means that a storage solution that worked for a smaller number of VMs can easily get tapped out if the storage repository can't keep up. As an example, Beyond Nines, a provider of SaaS-based fundraising and nonprofit management hosting services, had customers who had SQL queries that were taking hours to complete on a Dell server that was running only 10 VMs. After looking into the problem, it found a bottleneck in its storage configuration and decided it needed something more robust.

Now the company is able to run more than 150 high-performance customer environments on a single Tintri VMstore T540, according to Glen Kendell, president of hosting and operations at Beyond Nines. "SQL queries that used to take six hours now take four minutes," he says. This is because Tintri storage can directly handle virtualization abstractions such as VMs and virtual disks, along with providing a hybrid flash and magnetic array. The T540 unit contains a total of 13.5TB of storage, with a supplemental 1.8TB of flash capacity.

Another part of easily scaling your environment is using one of the storage hypervisors to more easily migrate storage from one place to another.

Understanding latency issues. When it comes to measuring application performance across our local enterprise networks, we think we know what network latency is and how to calculate it. But when we move these applications off-premises and onto private clouds, a lot of subtleties can impact latency in ways that we might not immediately realize. Part of the problem lies squarely at the feet of the current crop of developers, who take low-latency, high-bandwidth networks

Mixing different hypervisors

Many enterprises employ more than one vendor's hypervisor to deliver their VMs, which can be an issue. In the past, SANs had their own specialized commands to connect their storage repositories to various applications and file servers, set up particular volumes, and keep track of the storage consumed on each array. As storage requirements changed, or as new hypervisors were added, you generally didn't have much flexibility in moving unassigned capacity to where it was needed or from one hypervisor to another. Enter the role of the storage hypervisor or storage resource manager. These tools, including products from DataCore Software and IBM's SmartCloud Virtual Storage Center, can manage repositories that can be used across multiple VM hypervisors. A new entrant in this market is from Convirture, which can manage multiple VM hypervisors, including VMware, Amazon Web Services and OpenStack, and various storage types, including NFS, iSCSI, Fibre Channel and ATA over Ethernet.

These products have powerful features, such as the ability to automatically migrate data from expensive flash to less costly disk drives or the cloud, depending on usage. They also provide for near-real-time backup, often called continuous backup, of key server databases. "The kicker comes from being able to non-disruptively move those slices around. In the case of a storage hypervisor, you can move a slice (or virtual volume) from tier to tier, from vendor to vendor, and now, from site to site, all while the applications are online and accessing the data," noted Ron Riffe, program director, software defined environments at IBM at a panel presentation at the Storage Networking World conference last year.

Bryan Peterson is the associate director of technical services for the Utah Education Network in Salt Lake City. The statewide network connects all public K-12 and higher education campuses, and provides central IT services for various education courseware and back-office applications. The organization purchased DataCore Software's SANsymphony to manage its virtualized storage network. "This allows us to be array-agnostic, since we manage everything with DataCore," Peterson says. "We would like to stick with one array vendor, but we are very price-sensitive being a state agency, so this makes it a lot easier. Plus, we can attach all different kinds of storage to our network and manage it centrally."

for granted in building their internal applications. However, with cloud applications, endpoints aren't fixed: People can connect to the cloud from all over the world, and these connections can be unpredictable. Also, cloud applications are inherently more distributed than ones that reside on a few servers inside your typical data center, introducing yet another degree of variation and network quality. And with virtualization, a third degree of uncertainty is added to the mix.

Many cloud providers are cagey when it comes to providing details about their infrastructure and how it's connected to the Internet. Amazon, for example, doesn't even tell you where its data centers are located, using general terms such as "US-East" to describe where its sites are, let alone specify what gear it uses in them.

Such factors are what online gaming vendor Digital Chocolate faced when it began hosting and delivering its games on a public cloud. To cut back on latency problems, it decided to build a private cloud using an ordinary SAN. However, this didn't scale or perform well. The company decided to move to Tintri storage servers and now has better performance and visibility into application resource usage and latency stats across its applications.

Last but not least, consider these three questions as you look to optimize your virtual storage:

- How dynamic is your data? Make sure your storage architecture supports scaling your capacity under the right circumstances.
- Can your storage management solution work with automating and provisioning your applications?
- How seasonal is your storage demand, and how will your storage solution handle the peaks and valleys?

How VDI changes storage requirements

ONE CHALLENGE OF implementing virtual desktop infrastructure is that it consumes storage differently than a typical desktop or server PC. A desktop operating system naturally assumes that its storage will be dedicated, under its exclusive use, and have complete control over the entire computing stack from the lowest level to the applications. That isn't the case with VDI. Files are being accessed by a number of different VMs concurrently, and boot storms can occur when dozens or hundreds of VMs want to boot up their virtual operating systems within a very narrow time window. This causes havoc with a traditional storage system that is expecting this one-to-one relationship with the desktop OS. The net result is that VDI workloads have more frequent and random disk writes to storage and can be more write-intensive in their disk operations than traditional desktops. This means that you must design storage for VDI implementations that takes these demands into consideration.

For those considering a VDI deployment, VMware has published a white paper that shows you how to adjust your expectations and ensure that the virtual replacement can outperform your traditional desktop PC. It also provides new formulas for calculating storage performance in a VDI environment and includes an appendix with a sample worksheet that takes you through how to calculate storage requirements for a typical 500-user installation.

VDI can also multiply your storage needs if you end up using a large number of persistent desktop VM images, or ones that need to be stored separately and have their user states preserved. A better solution is to employ as few golden master disk images as possible. This reduces overall desktop support costs, since the IT staff knows what end users are starting from. The more images there are,

"We can update our lab configurations in mid-semester, or even anytime, without the complexity and the associated downtime that we used to have."

–JASON STRICKLAND, IT MANAGER AT SOUTHEASTERN COMMUNITY COLLEGE IN WHITEVILLE, N.C.

Thinkstock

the more updates that have to be applied and maintained across the board. Having a small number of master images means that everyone starts with, for example, the same basic Windows 7 or 8 image.

Jason Strickland, IT manager at Southeastern Community College in Whiteville, N.C., originally had to deal with 24 different Windows images across the campus. Now, with the college's present VDI deployment, Strickland and his staff have to patch only one Windows golden image, and they were able to toss out 5-year-old Windows Vista PCs that didn't support their current device drivers. In reducing all these variations, the team has improved its support and responsiveness to users. "Our students were pleased at how quickly they could get online to their desktops with our VDI implementation," Strickland says. But there also have been benefits elsewhere. "We can update our lab configurations in mid-semester, or even anytime, without the complexity and the associated downtime that we used to have."

There are different mechanisms to reduce overall VDI storage needs. Those include using a VDI-aware intelligent storage array from vendors such as Tintri, deploying virtualization add-on products from VMware, Unidesk and others that make linked clones or layered storage possible, or using storage hypervisors to eliminate redundant disk images. For example, in a test of a 500-node VDI installation, Enterprise Strategy Group saw vast performance gains using Tintri's storage array and deployed 500 virtual desktops within 30 minutes. According to ESG's report, "Individual desktops were provisioned with 40GB virtual disks, and at the end of testing, each desktop had consumed approximately 3.26GB of capacity."

Whatever solution you choose, the goal is the same: Avoid having multiple copies of the same desktop image tying up your storage repository and more quickly delivering the bits on boot time to boost virtual desktop performance.

David Strom writes and speaks about networking and communications topics. He can be reached through his web site or @dstrom.